High Peak and Average Power Mid-Infrared Laser for High Harmonic Generation of Soft X-rays

by

Susannah Rachel Wang

B.S., Massachusetts Institute of Technology, 2009

M.S., University of Colorado, 2012

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Physics

2016

This thesis entitled:
High Peak and Average Power Mid-Infrared Laser for High Harmonic Generation of Soft X-rays
written by Susannah Rachel Wang
has been approved for the Department of Physics

Prof. Henry Kapteyn

Prof. Margaret Murnane

Date

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

Wang, Susannah Rachel (Ph.D., Physics)

High Peak and Average Power Mid-Infrared Laser for High Harmonic Generation of Soft X-rays

Thesis directed by Prof. Henry Kapteyn

This thesis describes the development of a new mid-infrared laser designed to drive high harmonic generation of keV-energy soft x-rays with high flux. The mid-infrared wavelength regime (3 to 5 μm wavelength) is required to generate high harmonics with photon energies reaching 1 keV in the form of isolated attosecond bursts. This light would provide simultaneous few-nanometer spatial resolution and attosecond time resolution that could shed light on physical processes which occur at these length and time scales. Such processes have both scientific and technological importance.

The laser system described in this thesis reaches pulse energies up to 1.25 mJ at 1 kHz repetition rate, 3.1 μm wavelength, and with enough bandwidth to support 60 fs transform-limited pulses. Also, we demonstrate preliminary pulse compression to below 500 fs. This laser is therefore the first table-top mid-infrared laser with enough peak intensity and average power to generate harmonics with sufficient flux to be useful for application experiments. This laser uses Optical Parametric Chirped Pulse Amplification (OPCPA) to convert near-infrared light to the 3 μm wavelength regime, combining fiber lasers, cryogenically cooled solid state lasers, diode lasers, and optical parametric amplification in a unique architecture. In this thesis, we describe the current design of this laser system, the considerations that influenced its design, and its potential for scaling to higher pulse energies and repetition rates in the future.

Dedication

This thesis is dedicated to my loving parents who provided me with an excellent education, to my wonderful husband, Tim, to our little one on the way, and to God, the Father of lights.

Acknowledgements

I would like to thank the many people who have contributed to the work described in this thesis. First, my advisors, Henry Kapteyn and Margaret Murnane have provided valuable input, direction, and resources over the course of my time at CU. I am grateful for the opportunity that they provided to work on a cutting-edge laser development project in their leading research group. Second, I would like to acknowledge Michael Gerrity who has worked with me on this project from the time I started until his graduation and has continued to provide important advice and input on design decisions. Sterling Backus has played an important role in this project, providing a wealth of knowledge on laser design and engineering as well as help in making important design changes. I also received help and advice from other current and former Kapteyn-Murnane Laboratories employees including Matthew Kirchner and Christopher Wood. Finally, I would like to acknowledge Seth Cousin who has been instrumental in developing a vision for the next steps with this laser system and has played an important role in many aspects of this project since his arrival as a postdoctoral fellow in January.

Contents

Chapter

Tables

Table

Figures

Figure

Chapter 1

Introduction

1.1 High Harmonic Generation for Probing Nanoscale Systems on Ultrafast Timescales

As humans, our ability to observe and understand patterns in nature largely hinges on our ability to see the world around us. The sliver of nature we are able to see has increased with inventions such as the telescope, microscope, and, more recently, the laser, revealing both order and new mysteries. Laser science has enabled many new types of quantitative imaging and metrology techniques that let us "see" processes at widely varying time and length scales. Coherent ultraviolet and x-ray light from high harmonic generation, in particular, provides imaging resolution at nanometer length scales in the form of femtosecond-scale ultrashort pulses. These light sources therefore provide a window into the nano- and atomic-scale world in both the spatial and temporal domains.

In this thesis, I will describe the development of a next-generation light source that is designed to enable even better spatial and temporal resolution than is currently available from any table-top laser source. The laser described in this thesis is a mid-infrared, ultrafast laser system, designed to drive the production of keV-energy, coherent soft x-rays through high harmonic generation. This work represents progress toward a high-flux soft x-ray source that is accessible to individual research groups on a day-to-day basis (in contrast to large user facilities such as synchrotrons or free electron lasers). The goal is to provide a new tool to enable progress in understanding physical phenomena at the edges of our current picture of the natural world.

1.2 Mid-infrared Lasers for Strong-Field Physics

The laser science described in this thesis is part of the newest wave in an ongoing effort to develop new laser sources capable of driving high harmonic generation and strong-field physical phenomena in general. As research in strong field physics has progressed, the need for mid-infrared lasers that can simultaneously provide high peak intensity pulses and high average power has increased. These lasers are needed to push the maximum photon energy from high harmonic generation to the keV-level while maintaining the flux needed for experiments.

Historically, mode-locked titanium-doped sapphire lasers (Ti:sapphire) have been the workhorses of laboratory-scale strong field physics experiments due to their ability to simultaneously support few-optical-cycle, ultrashort pulses and handle high average power. Since the development of mode-locked Ti:sapphire lasers in the 1990's, this science has matured to the point of yielding terrawatt-level peak intensity pulses at kilohertz repetition rates. [6, 7] These high peak intensity pulses allow atomic ionization which is required for high harmonic generation. Also, the high repetition rates enable high flux of harmonics. Ti:sapphire-driven high harmonic generation has enabled measurements that have shed light on many areas of scientific and technological interest including ultrafast magnetism, nano-scale heat transfer, and 22-nm resolution coherent diffractive imaging. [8, 9, 10]

As this field continues to progress, the need for better temporal and spatial resolution as well as more x-ray flux has emerged. These requirements correspond to more demands on the driving lasers including the need for mid-infrared wavelengths. It has long been understood that the maximum photon energy that a single atom can produce in the high harmonic generation process scales proportionally to the driving laser's intensity and to the square of its wavelength. [11] Additionally, intensity scaling at a fixed wavelength turns out to be limited due to plasma dispersion. For these reasons Ti:sapphire laser's 800 nm light can only produce phase-matched harmonics with energies up to around 150 eV and long-wavelength driving lasers are a natural direction to take in pursuit of higher energy photons. [12, 13, 14]

Creating a mid-infrared laser capable of generating high peak and average power is not trivial due to the lack of suitable laser materials in this wavelength range. Recently, however, laser science has progressed to the point where this goal is becoming achievable. In Chapter 3 of this thesis, we will describe progress in mid-infrared laser technologies such as solid-state lasers, fiber lasers, semiconductor lasers, and optical parametric amplifiers which has brought each to higher power levels. In the work described in this thesis, we have built on these advances to demonstrate the first mid-infrared laser with enough pulse energy to drive strong-field processes and enough flux to be useful for experiments.

The work presented in this thesis has explored and combined many different laser technologies including fiber lasers, cryogenically cooled solid-state lasers, and optical parametric amplification. This work has not only resulted in the highest pulse energies of any 3–5 μm wavelength laser at a kilohertz repetition rate, but has also explored new laser architectures that will aid the continued development of future mid-infrared lasers for strong-field physics. In the same way that Ti:sapphire laser science has matured over the past two decades, we believe that mid-infrared lasers will continue to develop to become robust, compact, and reliable sources with increasingly impressive capabilities.

1.3 Contents

This thesis will begin with an introduction to High Harmonic Generation (HHG) as it is the key physical process enabling a table-top coherent x-ray source. A simple model of the HHG process will be described as well as a strategy for extending the cutoff photon energy from the ultraviolet into the soft x-ray spectral regime. Specifically, we will motivate the use of mid-infrared lasers as drivers for the soft x-ray HHG process and describe their performance requirements.

In the next section, we will motivate the use of Optical Parametric Chirped Pulse Amplification (OPCPA) for generating high average and peak power mid-infrared light and describe the theory behind this technique as well as a comparison with light amplification in traditional laser gain media. We will give an overview of the OPCPA system design employed in this work and the motivations for our design choices for each major component. Finally, we will describe the current

best performance of this system and discuss its physical and practical limitations as well as future directions for improvement.

Chapter 2

High Harmonic Generation

High Harmonic Generation (HHG) is a process that combines many relatively low-energy photons to create one high-energy photon, allowing visible or infrared laser light to be converted into a temporally and spatially coherent ultraviolet or x-ray beam. This phenomenon is a cousin to familiar nonlinear optical processes like second harmonic generation. Second harmonic generation occurs when atoms respond anharmonically to an intense driving laser field so that pairs of photons combine, creating light with twice the frequency of the original field. High harmonic generation carries this concept to the extreme, driving atoms in the strong field regime so that tens, hundreds, or even thousands of photons combine to produce an ultraviolet or x-ray photon.

The goal of this chapter is to give an introduction to HHG, describing its scaling up to keV photon energies and the need for mid-infrared driving lasers. We will first give a simple description of the single-atom physics underlying HHG, emphasizing the aspects of the process that are wavelength-dependent and influence scaling to higher energies. Next, we will consider the macroscopic description of HHG and introduce the concept of phase matching. Finally, we will outline the requirements for a mid-infrared laser designed to make harmonics up to keV photon energies.

2.1 Single-Atom Picture

In this section, we will give an introduction to the interaction between an intense laser field and an atom which can give rise to high-order harmonics. We will start with a semiclassical

description of this phenomenon, giving some intuition for the process and predicting the highest photon energy that can be generated. Secondly, we will introduce a more accurate quantum picture that gives information about the probability for any particular laser-atom interaction to yield a high-order harmonic photon. This probability affects the total flux we can expect to get from the HHG process and is strongly dependent on the driving laser wavelength.

2.1.1 Three Step Model

The semiclassical description of HHG is commonly broken down into three steps: ionization of an electron from an atom by the electric field of the driving laser, propagation of the freed electron in the continuum, and recombination of the electron with its parent ion, resulting in the emission of an high-energy photon. [11, 15] Simply known as the "Three Step Model," this process is illustrated in Figure 2.1.

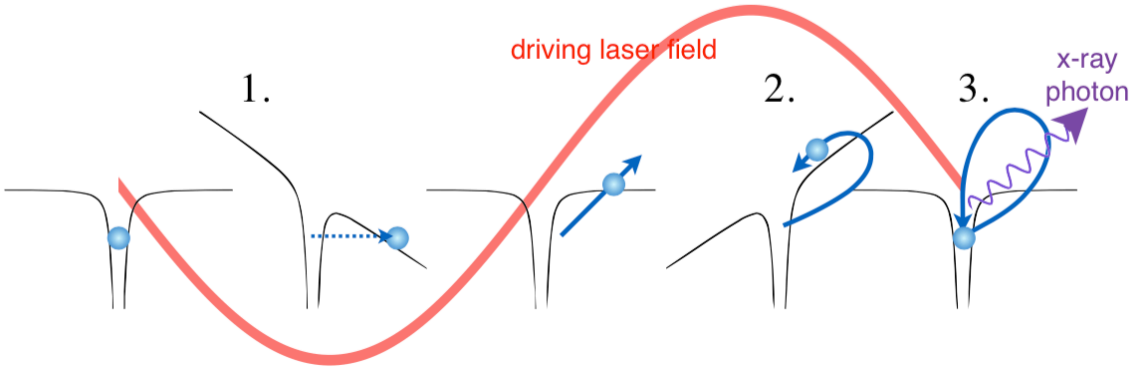

Figure 2.1: Illustration of the semiclassical three step model. The red line represents the field of the driving laser. It distorts the coulomb potential of an atom differently depending on the instantaneous field direction. The interaction can be broken down into three steps: 1. Ionization. 2. Acceleration of the free electron in the continuum. 3. Recombination with the parent ion and emission of a high-energy photon.

The ionization process occurs because of the interaction between the electric field of the driving laser and the Coulomb potential of the atom. A linearly polarized driving laser will induce a force on the valence electron proportional to its magnitude and direction: $\vec{F} = e\vec{E_L}$. Usually (in the non-relativistic regime) we can neglect the Lorentz force due to the magnetic component of the

light field, and the total force experienced by the electron is that of a Coulomb potential that has been tilted to one side. With enough tilt, the electron may tunnel ionize or spill out over the top of the suppressed Coulomb barrier.

Once ionized, the electron's trajectory is controlled by the electric field. For a linearly polarized driving field, the semiclassical trajectory follows Equation 2.1.

$$m\ddot{x} = eE_L(t)\cos(\omega t) \tag{2.1}$$

Here m and e are the mass and charge of an electron respectively, $\ddot{x}$ is the second derivative of the electron's position with respect to time, E_L is amplitude of the driving laser's electric field and ω is its frequency. The velocity and position as a function of time are therefore:

$$\dot{x}(t) = \frac{eE_L}{m\omega}\left(\sin\omega t - \sin\omega t_o\right) \tag{2.2}$$

$$x(t) = -\frac{eE_L}{m\omega^2}\left(\cos\omega t - \cos\omega t_o\right) - \frac{eE_L}{m\omega}(t - t_o)\sin\omega t_o \tag{2.3}$$

Since the electric field of the driving laser is oscillatory, the freed electron will first travel away from its parent ion, then change direction and head back as shown in Figure 2.2, (b). By the time it reaches its parent ion, the electron will have gained some kinetic energy from the electric field and, upon recombination, this energy will be released as a photon. The high harmonic photon energy can therefore be determined by adding the ionization potential of the atom, $U_{ionization}$, to the kinetic energy at the point of return, K:

$$\hbar\omega_{HHG} = K + U_{ionization} \tag{2.4}$$

At any particular ionization time, t_o, K can be determined by finding the the velocity at the time of return, t_r, since $K = \frac{1}{2}m\dot{x}^2$. The return time can be found using Equation 2.3. Numerically solving for t_r and $\dot{x}(t_r)$, we can plot kinetic energy versus the ionization time as in Figure 2.2 a). We find that the maximum kinetic energy is about $3.17U_p$, where U_p is the ponderomotive

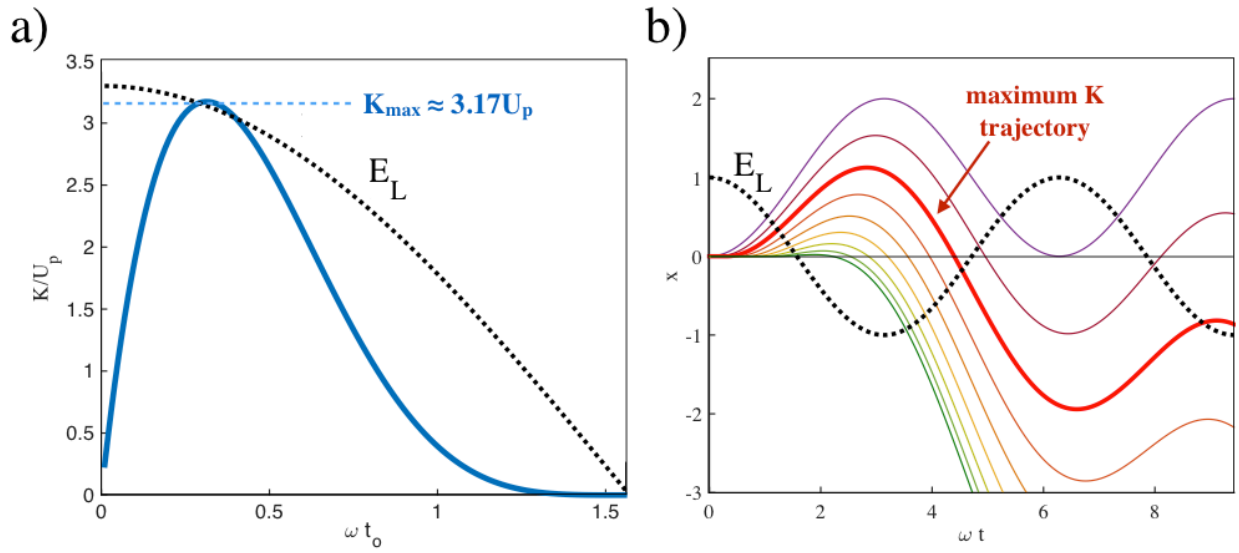

Figure 2.2: The kinetic energy of the electron upon return to its parent ion varies with ionization time with respect to the driving laser field oscillation. a) Numerical calculation of the ratio of kinetic energy, K, to ponderomotive potential, U_p, (solid, blue line) varying ionization time, t_o. $t_o = 0$ corresponds to the point where the electric field of the driving laser (E_L, dashed, black line) is at its maximum value. b) Semiclassical trajectories of electrons with various ionization times (solid lines) and the electric field of the driving laser (dashed, black line). The bold red line illustrates the trajectory of the electron which returns with maximum kinetic energy.

potential. [11, 15] In terms of the laser's field strength and frequency and the electron's mass and charge, the ponderomotive potential can be written as in Equation 2.6.

$$U_p = \frac{e^2 E_L^2}{4m\omega^2} \tag{2.5}$$

$$= \frac{e^2}{8\pi^2 c^3 \epsilon_o m} I_L \lambda^2 \tag{2.6}$$

Here I_L is the driving laser intensity. Equation 2.4 can therefore be rewritten for the maximum photon energy and in terms of the driving laser parameters:

$$\hbar\omega_{max} = 3.17 \frac{e^2}{8\pi^2 c^3 \epsilon_o m} I_L \lambda^2 + U_{ionization} \tag{2.7}$$

The important point here is that the maximum harmonic photon energy (also known as the cutoff energy) is proportional to the driving laser intensity and the square of its wavelength:

$$\hbar\omega_{max} \propto I_L \lambda^2 \tag{2.8}$$

Equation 2.8 therefore suggests two methods for scaling up the energy of our harmonics toward the x-ray regime: 1) turning up the driving laser intensity and 2) using a laser with a longer wavelength. As it turns out, neither intensity nor wavelength scaling is trivial, both facing fundamental and or technical challenges which will be described in more detail in what follows. We will also see, however, that wavelength scaling does present a tenable path toward soft x-ray HHG, motivating the development of mid-infrared driving lasers.

2.1.2 Quantum Picture

The semiclassical three step model is helpful for gaining intuition, but a more detailed understanding of HHG requires a quantum mechanical description. For our purposes, the most important quantum correction to the semiclassical picture is wavepacket diffusion while the electron is in the continuum. Whereas the semiclassical picture treats the free electron like a particle with well-defined position, quantum mechanics dictates that it will spread significantly as it propagates.

This spreading lowers the probability that the electron will recombine with its parent ion and emit a photon.

Intuitively, we expect the wavepacket to spread more the longer it remains in the continuum. From our semiclassical model, we also expect that the amount of time the electron spends in the continuum before returning to its original position increases with the wavelength of the driving laser. This suggests that the single-atom photon yield will decrease as the driving laser wavelength increases. Fully quantum analytic and numerical models of this process show that single-atom HHG yield scales as λ_L^{-n} where n is about 5.5 ± 0.5, assuming constant intensity. [5, 16, 17].

At first look, this strong decrease in photon yield with increasing driving laser wavelength is discouraging for generating a bright soft x-ray source with a long-wavelength driving laser. Fortunately, however, the single-atom picture is not the whole story since we need many atoms coherently emitting x-rays to create a bright source. This macroscopic description of HHG uncovers scaling which is favorable for long-wavelength-driven HHG and compensates for the seemingly unfriendly $\lambda_L^{-5.5}$ scaling.

2.2 Macroscopic Picture

So far we have only considered the physics of HHG from a single atom, yielding a single photon. However, since we are interested in making a bright x-ray beam, we will need many atoms to contribute. What is more, their contributions need to sum coherently and constructively. In what follows, we will discuss the physics of many-atom HHG and the conditions required for bright, coherent x-ray emission. Finally, we will also describe the implications of this macroscopic picture of HHG to the design of our mid-infrared driving laser.

2.2.1 Phase Matching

In order for many atoms to emit high harmonic photons which interfere constructively, the relative phases of the harmonics must be equal. Equivalently, the phase velocities of the driving laser and the emitted photons must be equal as they propagate. This condition is known as phase

matching.

In terms of the wavevector, k, the phase mismatch can be written as:

$$\Delta k = qk_L - k_{x-ray} \tag{2.9}$$

Here q is the harmonic order, k_L is the wavevector of driving laser and k_{x-ray} is the wavevector of the high-order harmonic photon. Perfect phase matching, and thus efficient buildup of harmonic intensity, occurs when Δk is equal to zero.

In traditional nonlinear frequency conversion processes (such as second harmonic generation) phase matching is typically achieved with the aid of birefringence in certain solid materials. This approach will not work for x-rays, however, since many solid materials are highly absorptive in this region of the spectrum. Inert gases, on the other hand, are only weakly absorbing. These gases are isotropic, however, meaning that we cannot rely on birefringence to achieve phase matching.

Fortunately, gases provide other "knobs" for tuning the phase mismatch. Specifically, the phase velocity of the driving laser is affected by the density of the gas, the density of ionized atoms and electrons, and by focusing conditions or waveguide geometry. On the other hand, the phase velocity of the x-rays remains constant (slightly above the speed of light in vacuum) as these parameters are tuned. [18, 12]

To get the highest conversion efficiency from the driving laser to high harmonics, it is generally desirable to use a gas-filled waveguide to maintain the laser intensity over an extended distance, increasing the length of the interaction region. In this geometry, the phase mismatch is given by Equation 2.10. [12, 3]

$$\Delta k \approx q\left[\frac{u_{11}^2\lambda_L}{4\pi a^2} + P\eta N_a r_e \lambda_L - P(1-\eta)\frac{2\pi\Delta\delta + n_2}{\lambda_L}\right] \tag{2.10}$$

Here, the first term is the dispersion due to propagation in the waveguide, the second term is the contribution from neutral atoms, and the third term is the effect of free electrons. q is the harmonic order, u_{11} is a factor representing propagation in the 11 mode of the waveguide, a is the

radius of the interior of the waveguide, P is the gas pressure, η is the ionization fraction, N_a is the number density of atoms per atmosphere, r_e is the classical electron radius, $\Delta\delta$ is the difference in the refractive indices of the gas at the wavelengths of the driving laser and the harmonic, and n_2 is the nonlinear index of refraction per atmosphere at the driving laser wavelength.

Equation 2.10 tells us that the key to phase matching HHG in a gas is balancing the dispersion due to neutral atoms with that of the free electrons and the waveguide. Experimentally, the gas pressure can be tuned to change the relative magnitudes of the three terms in Equation 2.10 and phase matching can be achieved if the ionization levels are moderate. [3, 12, 13, 14]

2.2.2 Phase Matching Cutoff

The next question is how far we can push this process in terms of harmonic flux and photon energy. Naively, with perfect phase matching we might expect that the longer the interaction length, the more harmonic photons we will get. In other words, the longer the waveguide the better. However, in practice, there is an optimal waveguide length determined by the point at which re-abosorption of the harmonics by the gas would make a longer interaction regions counterproductive. [14]

A second way to increase harmonic flux is to push the driving laser to higher repetition rates so that more bursts of harmonics are generated in a given time interval. The limitation here has to do with design constraints of the driving laser such as the amount of available pump power or thermal loading of the laser crystal or other components.

In terms of reaching higher photon energies, Equation 2.8 provides two paths forward: 1) higher driving laser intensity and 2) longer wavelength. The first is the simpler of the two because it does not require a new laser material or frequency conversion. Intensity can be scaled up by increasing the amount of energy per pulse, decreasing each pulse's temporal duration, and/or decreasing the mode area in the interaction region. Furthermore, Ti:sapphire lasers have been made to reach extremely high pulse energies beyond 600 Joules while maintaining pulse durations short enough to reach petawatt-level peak intensities. [19, 20] From the perspective of driving laser

requirements, intensity scaling is therefore possible to a high level beyond what is required to ionize any given gas.

Critical Ionization Unfortunately, intensity scaling with Ti:sapphire lasers meets a roadblock with the physics of phase matching. As intensity increases, so does the ionization fraction (η in Equation 2.10). Once a certain critical ionization level is reached (which depends on the gas species), the plasma dispersion increases the driving laser's phase velocity to the point where the neutral atoms' dispersion can no longer compensate for it. A Ti:sapphire-driven HHG process is therefore limited to photon energies below about 150 eV. [12, 13, 14] An expression for the critical ionization fraction is given by Equation 2.11. [21, 3]

$$\eta_{cr}(\lambda_L) = \left[\left(\frac{\lambda_L^2 N_a r_e}{2\pi\Delta\delta(\lambda_L)}\right) + 1\right]^{-1} \tag{2.11}$$

It is important to note that η_{cr} depends both on the gas species and on the laser wavelength. The harder the gas is to ionize, the more intensity can be used before reaching the critical ionization fraction. The smaller nobel gases (helium and neon) have higher ionization potentials and therefore allow phase matching at higher photon energies. Helium is usually the gas of choice when the goal is to push to high energies because it lacks the inner shell absorption edges that can decrease HHG photon yield in neon and argon. [3] Also, η_{cr} scales approximately as λ_L^{-2}, meaning that longer wavelength driving lasers reach critical ionization at a lower maximum intensity.

Quasi Phase Matching Several methods have been developed to overcome this maximum photon energy limit while sticking with the Ti:sapphire laser. One of these is called quasi phase matching. With this technique, the phase slip, Δk, is periodically corrected for rather than eliminated, allowing coherent buildup of harmonics even in situations where ionization levels do not allow perfect phase matching. Quasi phase matching of HHG has been demonstated by modulating the diameter of the waveguide to periodically change the waveguide contribution to Δk [22] and with a counter propagating pulse train which supresses harmonic generation in regions of the waveguide that would otherwise cause destructive interference. [23, 24, 21]

Quasi phase matching techniques have allowed the cutoff energy to be extended above 280 eV. [22] However, further extension toward the keV level remains challenging due to nonuniformities of the phase matching conditions along the length of the waveguide and due to the technical difficulties in creating optical pulse trains or waveguides with sufficient tunability to match these variations.

Wavelength Scaling While the ultimate limits of quasi phase matched HHG have yet to be found, wavelength scaling has proven to be a simpler alternative. Though the single atom yield has a highly unfavorable scaling for long wavelengths ($\lambda_L^{-5.5}$), a convergence of other favorable factors in the phase matching process save the day. One of these factors is that the phase matching pressure increases with wavelength, scaling as λ_L^2. Also, as the photon energies of the harmonics increase, the absorption in helium is lower, allowing longer absorption-limited interaction regions between the driving laser and the gas. Taking these factors into account and remembering η_{cr}'s dependence on wavelength, one can show that the maximum photon energy that can be phase matched ($\hbar\omega_{pm}$) scales as λ_L^n where n is between 1.5 and 1.7. Though this is a slightly slower increase than the λ_L^2 dependence of Equation 2.8, it is still highly favorable. Furthermore, the maximum HHG photon flux is expected to even exceed that achievable using near-infrared driving lasers. [3, 1, 4, 25]

$$\hbar\omega_{pm} \propto \lambda_L^{1.5-1.7} \tag{2.12}$$

This favorable wavelength scaling has motivated the development of Ti:sapphire-based Optical Parametric Amplifiers (OPAs) delivering ultrashort pulses with central wavelengths near 1.3 μm and 2.0 μm, as well as Optical Parametric Chirped Pulse Amplifiers (OPCPAs) delivering ultrashort 3.9 μm pulses. These lasers have enabled the verification of the wavelength scaling of Equation 2.12, extending the phase matching cutoff in helium from 150 eV ($\lambda_L = 0.8$ μm) to 1.6 keV ($\lambda_L = 3.9$ μm). [1, 2, 3, 4, 5, 25] Figure 2.3 plots these experimental data for helium, neon, and argon, showing that we require a driving laser with wavelength of at least 3 μm to reach

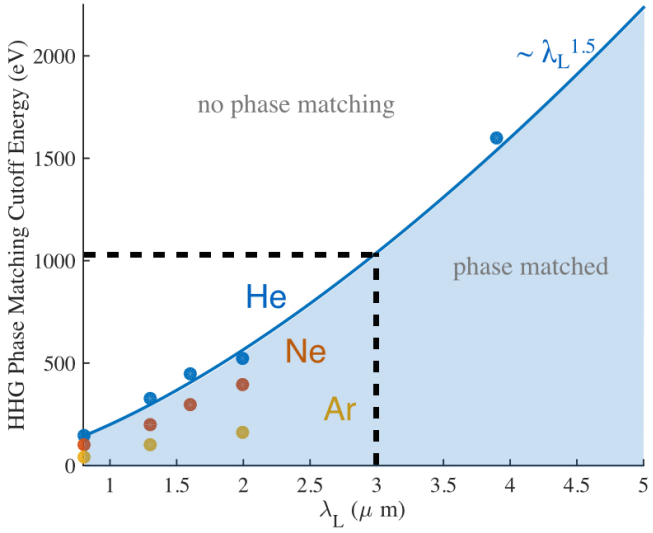

Figure 2.3: Scaling of the phase matching cutoff of HHG with driving laser wavelength. Based on experimental data for helium (blue), neon (orange), and argon (yellow) from Popmintchev et al. [1, 3, 5], Chen et al. [4], and Takahashi et al. [2].

phase-matched, 1 keV harmonics.

2.2.3 Supercontinuum and Isolated Attosecond Pulses

As the HHG driving laser wavelength is increased, the attainable harmonic order can increase into the thousands. For example, Popmintchev et al. used a 3.9 μm laser to reach the 5001^{st} order. [5] In this regime the harmonic peaks blur together into a supercontinuum and the bandwidth of the generated spectrum broadens, spanning the Vacuum Ultraviolet (VUV) to 1.6 keV. These effects are illustrated by the experimental data shown in Figure 2.4.

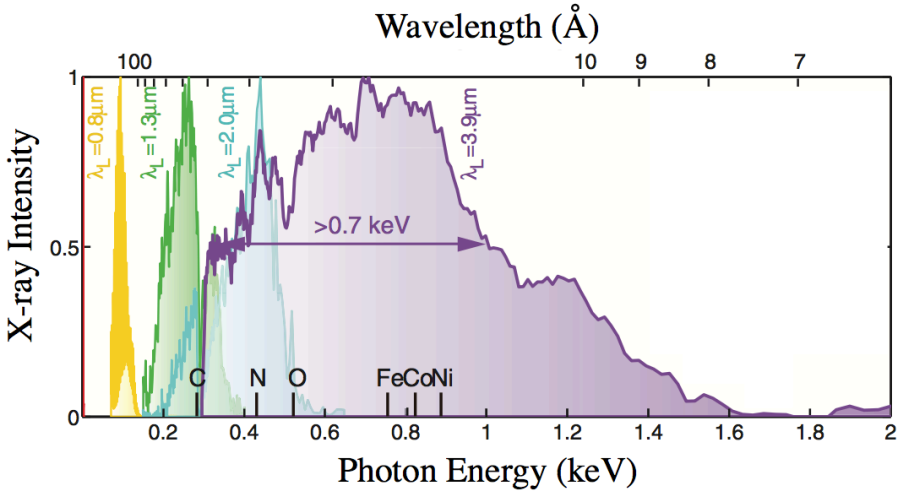

Figure 2.4: Experimental data showing the HHG spectra obtained for 0.8, 1.3, 2.0, and 3.9 μm driving laser wavelengths. Figure adapted from [5].

A broad, smooth spectrum such as the one shown in Figure 2.4 for $\lambda_L = 3.9\ \mu m$ is useful for experiments that require elemental contrast since various elements have absorption edges in the spectral range between 200 and 1000 eV. Some of these elements (namely carbon, nitrogen, and oxygen) have immense importance for imaging biological samples having absorption edges in the "Water Window" between 200 and 600 eV. Other notable elements include iron, cobalt, and nickel which have absorption L-edges between 750 eV and 1 keV. The L-edges are especially interesting for studying magnetic samples because the absorption lines are split by spin-orbit coupling in the presence of a magnetic field. Absorption measurements at these edges can therefore yield

information about the electronic configuration and spin state of materials. [26, 8] The ultrafast spin dynamics of magnetic materials are both not yet fully understood and technologically relevant for next-generation magnetic storage devices.

From Heisenberg's uncertainty principle relating spectral bandwidth to temporal pulse duration, we know that the larger the bandwidth, the shorter the pulse it can support, assuming flat spectral phase. This supercontinuum spectrum therefore also has the exciting possibility of supporting isolated attosecond bursts of harmonics. In contrast to this supercontinuum, Ti:sapphire-driven harmonics typically emerge as a series of peaks in the spectral domain, corresponding to a pulse train in the time domain. Though these harmonics are useful for pump-probe experiments which require picosecond or ≈100 femtosecond time resolution, isolated attosecond pulses are desirable in that they allow sub-femtosecond time resolution. A light source with pulses at this time scale will give experimenters access to a new regime of ultrafast dynamics where there may be new physics and/or keys to elucidate existing theories.

Theoretical work presented in [27] suggests that, with mid-infrared driving lasers, isolated attosecond pulses are not only possible, but arise naturally and automatically under phase matched conditions. This occurs because dispersion distorts the driving laser pulse as it propagates so that phase matching is destroyed except over small temporal region. As the wavelength is increased to the mid-infrared, this phase matching window decreases to one half cycle of the field, allowing only one harmonic burst to emerge.

This theoretical work was also experimentally verified by Chen et al. using Fourier transform interferometry to directly measure the autocorrelation of the x-ray pulse train emerging from the HHG process. Using driving laser wavelengths of 0.8, 1.3, and 2.0 μm under phase matched and non-phase matched conditions they showed: 1) for a given driving wavelength, the number of x-ray bursts decreases in the phase matched case compared to the non-phase matched case and, 2) that the number of phase matched bursts decreases to one as the driving laser wavelength is increased to 2.0 μm. [27]

Though a complete understanding of the physics involved with this process has yet to be

developed, these results tell us something interesting about the optimal driving laser pulse duration. Intuitively, one might expect that the shortest possible pulse duration is always the best as it gives the highest peak intensity for a given energy. However, for a few-cycle pulse, the group velocity walk-off between the fundamental and harmonic fields can be significant compared to the pulse envelope. With the high gas pressures required to phase match mid-infrared-driven HHG, this walk-off can narrow phase matching window to the point of suppressing harmonic emission. In order to preserve the phase matching window, it is better to have a relatively slowly-varying pulse envelope. In practice, this turns out to be between eight and ten cycles ($80 - 100$ femtoseconds) for a 3 μm laser. [28]

2.2.4 Requirements for a Mid-IR Driving Laser

In light of these results, we can put together a list of specifications for a driving laser that will enable high harmonic generation of a supercontinuum spanning the keV region of the spectrum. These are as follows:

1. Mid-infrared Wavelength We know from the work summarized in Figure 2.3 that the a mid-infrared driving laser wavelength allows HHG of fully phase-matched, keV photons. Specifically, the wavelength should be at least 3 μm. In the work described in this thesis, we therefore chose to aim for a laser wavelength in the range between 3 and 5 μm.

It is also worth noting that a complementary technique using ultraviolet driving lasers has recently been shown to generate nearly phase-matched harmonics in the soft x-ray regime, up to several hundred eV. [29] Rather than a supercontinuum, this technique produces a comb-like spectrum of bright, narrow harmonic peaks corresponding to a pulse train in the temporal regime. UV-driven harmonics are well suited for applications requiring narrow spectral line widths, such as Coherent Diffractive Imaging (CDI). [10] On the other hand, the broad spectrum and isolated pulses created by mid-infrared-driven HHG are better suited for applications requiring the highest levels of time resolution or absorption measurements covering multiple elements. Though UV-

driven HHG has great promise and possibly can be pushed to keV photon energies, both avenues are worth pursuing and, in this thesis, we will concentrate on mid-infrared HHG.

2. 8-10 Cycle Pulses As discussed in the previous section, the theory and experimental confirmation described in [28] and [27] show that our driving laser must produce pulses 8-10 optical cycles in duration. For a wavelength of 3 μm, this corresponds to 80-100 fs pulse durations. Our laser must therefore also supply sufficient spectral bandwidth to support these pulse durations and have flat spectral phase. For example, a transform limited gaussian pulse with central wavelength of 3 μm requires 170 nm full width half maximum bandwidth to support an 80 fs pulse. A pulse with a 5 μm central wavelength requires 460 nm to support the same pulse duration. Furthermore, since it is often difficult to eliminate all deviations from a flat spectral phase, we will require our mid-infrared laser to provide more than the minimum bandwidth. So, for a 3 μm laser at 80 fs, we will aim for about 200 nm bandwidth or more.

3. High Intensity to Drive Harmonics in Helium Helium is the most promising medium for generating keV harmonics since it has relatively low absorption in this region of the spectrum and has a high ionization potential. We can get a lower estimate of the intensity required to reach a particular photon energy using the single-atom cutoff rule, Equation 2.7. Typically, this value is in the range of 10^{14}-10^{15} W/cm^2. In terms of laser performance requirements, the obtainable intensity depends on the pulse duration, the amount of energy per pulse, and the mode area of the interaction region. When small interaction regions are required due to limited energy, the demand is placed on the spatial mode quality of the laser beam since this affects the beam's focusability. Table 2.1 gives an idea of the mode size needed for a 1 mJ, 3 μm pulse to drive keV harmonics. We estimate that our driving laser will need at least 1 mJ per pulse.

4. kHz-level Repetition Rate Finally, in order to improve the usability of this x-ray source, we need to maximize the x-ray flux it produces. The most straight forward way to accomplish this is to increase the number of x-ray bursts created per unit time. Previous work by Popmintchev et al. demonstrated keV-level HHG using a 4 μm laser operating at 10 Hz. [5] Though this laser supplied an impressive 8 mJ energy per pulse, the x-ray flux was still too low for

pulse duration (fs)	d (μm)	I_L (W/cm^2)	$\hbar\omega_{max}$ (keV)
100	50	5.1×10^{14}	1.47
100	75	2.3×10^{14}	0.67
100	100	1.3×10^{14}	0.39
80	50	6.4×10^{14}	1.84
80	75	2.8×10^{14}	0.83
80	100	1.9×10^{14}	0.48

Table 2.1: This table assumes a 3 μm, 1 mJ driving laser with temporally square pulses. Given a pulse duration and effective mode diameter, d, the resulting intensity, I_L, and expected single-atom photon energy cutoff, $\hbar\omega_{max}$, are shown.

most applications. Low x-ray flux can necessitate long integration times: enough counts must be collected for a well-resolved measurement to emerge from the background noise. Often the required integration time becomes longer than the time over which the laser source can remain stable in energy, pulse duration, and pointing. We aim to improve the x-ray flux compared to this previous work by a factor of 100, and therefore will be aiming for a 1-100 kHz driving laser.

Chapter 3

High Power Mid-infrared Lasers

3.1 Mid-Infrared Laser Materials

Typically, high harmonic generation is driven by a Ti:sapphire laser amplifier system because these lasers support femtosecond-scale pulses with energy up to tens of millijoules at kilohertz repetition rates. These energetic and short pulses possess the high peak intensities necessary for high harmonic generation. However, the emission bandwidth of Ti:sapphire is limited to the near-infrared whereas we require a mid-infrared driving laser for keV HHG. [30, 5] For this application, we need a mid-infrared laser material that shares Ti:sapphire's ability to support short, high-energy pulses. It must have a high emission bandwidth as well as thermal, thermo-optical, and mechanical properties suited for power scaling to several watts at kilohertz repetition rates.

In fact, there are a number of laser materials with emission spectra in the infrared [31], but the challenge lies in that our application is demanding, requiring average and peak power levels that took a decade of development to reach in Ti:sapphire amplifier systems. [30, 32] In what follows, we will discuss the advantages and limitations of various mid-infrared lasers including crystalline solid-state lasers, fiber lasers, semiconductor lasers, and Optical Parametric Amplifiers (OPAs). In doing so, we will justify our choice of optical parametric amplification as a promising path forward toward high power, high energy, and broad bandwidth.

Having justified our choice of OPA, we will briefly introduce the theory of second-order nonlinear frequency conversion. We will describe it first in general terms and then concentrate on the process of interest: optical parametric amplification using difference frequency mixing. In

doing so, we will point out aspects of the OPA process that require attention and influence design choices. These include power gain, phase matching, and bandwidth. Quasi phase matching will also be introduced as an advantageous method for achieving parametric gain in cases where true phase matching is not possible.

Finally, we will introduce Optical Parametric Chirped Pulse Amplification (OPCPA). This is the technique employed in the work described in this thesis for scaling our optical parametric amplifier to multi-millijoule pulse energies.

3.1.1 Mid-infrared Solid State Lasers

Mid-infrared solid state lasers have a long history, dating from 1960 (the same year as the invention of the laser) when Sorkin and Stevenson demonstrated a U^{3+}:CaF_2 laser operating at 2.6 μm. [33] Since then, many mid-infrared solid-state lasers have been demonstrated based on transition metal ions (Co, Ni, Cr, and Fe), lanthanide ions (Tm, Ho, Er, Dy, Pr, Tb, and Nd), and F-centers (also known as color-centers). [31]

Many of these laser materials can achieve broad bandwidths due to strong coupling between absorbed and emitted photons and phonons in the host lattice. [31] In fact, F-center lasers have achieved pulse durations below 50 fs. [34] However, this photon-phonon coupling also causes quenching due to non-radiative processes which destroy the population inversion. For F-center lasers and some transition metal ion-based lasers, this means cryogenic cooling is needed for lasing in the mid-infrared. [31]

Of these lasers, Cr^{2+}:ZnSe has been one of the most successful in terms of stability, simplicity, bandwidth, and power scaling. [35, 36] Its emission spectrum spans 2-3.4 μm, which is even broader than that of Ti:sapphire and could theoretically support few-cycle pulses. Also, it can be passively mode locked, diode-pumped, and is suited for thin-disk operation to mitigate thermal loading issues.

Cr:ZnSe and its cousin, Cr:ZnS, also have important limitations, however. One is that its power scaling may be limited to several watts in the mode locked or amplifier regime due to quenching from multi-phonon absorption. [35] Also, the power and pulse energy of mode locked

Cr:ZnSe is limited at this point. Active mode locking has resulted in 2-4 ps duration pulses, but with only 400 mW in average power with 2.5 μm central wavelength. [37] Strategies for passive mode locking using saturable absorber mirrors, carbon nanotubes, and graphene have yielded better results. Tolstik et al. report reaching pulse durations as low as 41 fs with 2.3 nJ pulse energy using Cr:ZnS [38], and Moulton et al. report pulse energies up to 0.3 mJ at 1 kHz and with 300 fs pulse duration in Cr:ZnSe. [39] In spite of these advances, however, the limited peak intensity provided by these lasers still precludes driving keV HHG.

Furthermore, the thermal, thermo-optical and mechanical properties of ZnSe are less promising than sapphire as a host material for a high-power laser. First, the thermal conductivity is only 67% that of sapphire, making efficient heat extraction more difficult. Second, the thermo-optic coefficient is about six times higher, exacerbating the problem of thermal lensing. Finally, the Knoop hardness is seventeen times lower than sapphire, making ZnSe susceptible to fracture due to a thermal gradient from the pump laser or from mechanical stresses from mounting to a heat sink. [31] This last issue makes it challenging to mount this crystal to a cold head for cryogenic cooling, a technique that has been instrumental in mitigating thermal lensing in Ti:sapphire amplifiers so that they can reach millijoule energies at kilohertz repetition rates. [40] For these reasons, Cr:ZnSe is not an ideal candidate for a driving laser for keV harmonics, at least at this stage in the technology's maturity.

3.1.2 Fiber Lasers

Fiber lasers doped with rare-earth cations offer another path toward a mid-infrared driving laser. Fiber lasers are promising in that they have proved successful for power scaling in the near-infrared, delivering over a kilowatt of average power at 1 μm. [41] They excel for high average powers because their long, thin gain material allows for efficient heat extraction. Additionally, they have the advantages of robustness, compact architecture, and, in the near-infrared regime, the laser transition has close to 100% efficiency and low quantum defect. [42]

Moving into the mid-infrared, the performance of fiber lasers is unfortunately more limited.

They can be made to work in the mid-infrared at up to 4 μm, but with exponentially decreasing output power as the wavelength increases from 1 to 4 μm. [42] In contrast to the kilowatts available at 1 μm, only 30 W of continuous-wave output power has been reached at 3 μm, [43] 1.5 W of continuous-wave power at 3.5 μm, [44, 45] and only 11 mW at 4 μm (even with the aid of cryogenic cooling). [46] This limited performance in the mid-infrared is partially caused by the fluoride or germanate glasses that must be used for λ>2 μm. These glasses have lower optical damage thresholds than their silicate counterparts, limiting the intensity that can be confined in the small fiber core areas. This is a particularly stringent constraint for pulsed lasers. Additionally, mid-infrared fiber lasers suffer from lower quantum efficiency which can lead to thermal loading in spite of the fibers' high surface-to-volume ratio. [42] These lasers are not currently a feasible direct path to reach 3-5 μm laser pulses with millijoule-level energies and kilohertz repetition rates.

3.1.3 Semiconductor Lasers

Another class of laser that can reach the mid-infrared is the semiconductor laser. These lasers have the advantage that they can directly transform electrical power into optical power, making them some of the most compact and robust laser sources available. In the mid-infrared regime, traditional (non-cascade) diode lasers have been demonstrated but suffer from low power output and the need for cryogenic cooling. [31]

Quantum cascade lasers are more promising, having greatly improved the performance of semiconductor lasers beyond 2 μm and even reaching out to 25 μm. [47] They offer immense design flexibility in terms of both wavelength and bandwidth which enables even supercontinuum spectra and mode locked operation with few picosecond pulses. [48, 49] Additionally, the output power from commercial quantum cascade lasers have reached 20 W in continuous-wave operation. [50]

In spite of these advantages, quantum cascade lasers have yet to be scaled to the high-pulse-energy, sub-picosecond regime. One barrier to this goal is their typically low wall plug efficiency which causes them to generate large amounts of heat in relatively small active areas. [47] Though quantum cascade lasers are rapidly progressing and finding applications in fields such as mid-

infrared spectroscopy and military countermeasures, they are not yet practical as driving lasers for keV harmonics in themselves. In the future, however, these may become attractive pump lasers for solid-state laser materials such as Cr: or Fe:ZnS, pushing this technology to a higher level than previously possible.

3.1.4 Nonlinear Frequency Mixing

Nonlinear frequency mixing provides an alternative approach to generating high-energy and high-power mid-infrared pulses. Frequency mixing can occur in certain materials which respond to a high-intensity driving light field nonlinearly, reradiating light of a different color. Under the right conditions, this nonlinear response can combine two incident frequencies, producing the difference frequency. This process is known as Difference Frequency Generation (DFG).

The usefulness of this process lies in that it allows us to leverage laser technologies that allow power scaling in the near-infrared and later convert that light into the mid-infrared using a nonlinear crystal. Unlike the laser gain process, difference frequency generation is lossless so that no heat is dumped into the crystal. This eliminates thermal loading as a barrier to power scaling, leaving only the crystal's damage threshold and maximum aperture as the primary limitations. Additionally, the gain can be quite high and support a broad range of frequencies, allowing for simple, single-pass amplification stages. For these reasons, we chose difference frequency generation as the most promising route toward our goal of a kilohertz, multi-millijoule, ultrafast mid-infrared laser source.

3.2 Optical Parametric Amplification

As previously mentioned, we have chosen nonlinear difference frequency generation as our route toward a mid-infrared driving laser. For this type of scheme, we first need a material that can support the difference frequency generation process. Then, in analogy to a laser amplifier, we will illuminate that material with intense pump light and a seed beam. In the amplification process, the pump will be depleted and two colors will be amplified: the seed beam and what is known as the "idler." An idler photon has energy corresponding to the difference in energy between a pump

and a seed photon. This process is known as Optical Parametric Amplification (OPA).

OPA differs from laser amplification in that all the energy that is taken from the pump beam is converted into either seed or idler and, in an ideal case, none of it is lost to heat or fluorescence in the nonlinear medium. In other words, the pump photons are split into two, lower-energy photons in a process that conserves energy. The energy of one of the photons is determined by the color of the seed light and the other photon gets the energy that is left over. The efficiency of this process is determined by the nonlinear coefficient of the chosen gain medium and the extent to which the phase velocities of the pump and seed beams match over the whole length of the medium. The latter condition is known as phase matching. In practice, some nonlinear gain media exhibit significant absorption at the pump, signal, and/or idler wavelengths, leading to thermal loading. This can limit parametric gain if the material has poor thermal conductivity, so care must be taken to choose a nonlinear material with low absorption at the wavelengths of interest.

The following two sections will be an introduction to OPA and the equations that govern it. We will see which characteristics of the nonlinear medium are important for amplification, motivating the design choices outlined in the next chapter. Additionally, this description of OPA elucidates requirements for the pump and seed lasers which are relevant to our design.

3.2.1 Nonlinear Frequency Conversion

In this section we will give a general description of the process by which light of one color can be converted to another frequency. This type of frequency conversion can be understood by looking at how high-intensity light propagates in a material. In the low intensity case, light will typically propagate through a transparent medium without a change in frequency. The weak electric field of the light will drive electrons gently and they will respond with harmonic oscillation, emitting light with the same frequency as the incident field. However, in the regime of a high-intensity driving field, the response of the electrons is generally nonlinear rather than harmonic. That is, the emitted light can have higher and/or lower frequency components. The exact spectrum of the radiated light is strongly dependent on the structure of the material and the relative direction of

the light's propagation with respect to it.

To understand this effect more concretely, we will start with a very general expression describing the interaction of light with matter:

$$\vec{D} = \epsilon_o \vec{E} + \vec{P} \tag{3.1}$$

Here $\vec{D}$ is the electric displacement field written in terms of the electric field, $\vec{E}$ and the polarization, $\vec{P}$. The first term in this equation, $\epsilon_o \vec{E}$ represents what the electric field would be in vacuum. The second term, $\vec{P}$, represents the effect of material.

At low intensities, the polarization is approximately linear and the field does not change its frequency content, though its group velocity may be slowed compared to the speed of light in vacuum and it may experience gain or loss due to absorption or emission in the material. In general, however, the polarization can also have a nonlinear part, $\vec{P}_{NL}$.

$$\vec{P} = \epsilon_o \chi^{(1)} \vec{E} + \vec{P}_{NL} \tag{3.2}$$

Here, $\chi^{(1)}$ is the linear susceptibility which is related to the material's index of refraction as well as its gain and absorption coefficients. The nonlinear polarization, $\vec{P}_{NL}$, can be similarly written in terms of higher-order susceptibilities:

$$P_{NL} = \epsilon_o \{\chi^{(2)} E^2 + \chi^{(3)} E^3 + ...\} \tag{3.3}$$

where $\chi^{(2)}$ is the second-order susceptibility and $\chi^{(3)}$ is the third-order susceptibility. These higher-order susceptibilities are tensorial quantities whose elements relate different components of the polarization vector to components of the electric field.

The relationship between the electric field and polarization is easiest to see if we write each as a sum of monochromatic plane waves.

$$\vec{E}(\vec{r}, t) = \sum_{\omega_a \geq 0} \sum_{l=1}^{3} \hat{e}_l \frac{1}{2} \{E_l(\omega_a) \mathrm{e}^{i(\omega_a t - \vec{k_a} \cdot \vec{r})} + c.c.\} \tag{3.4}$$

$$\vec{P}(\vec{r}, t) = \sum_{\omega_b \geq 0} \sum_{l=1}^{3} \hat{p}_l \frac{1}{2} \{P_l(\omega_b) \mathrm{e}^{i(\omega_a t - \vec{k_a} \cdot \vec{r})} + c.c.\} \tag{3.5}$$

Here, $E_l(\omega_a)$ is the complex amplitude of the electric field at frequency, ω_a, with polarization corresponding to the unit vector, e_l. Similarly, $P_l(\omega_b)$ is the complex amplitude of the polarization at ω_b with direction p_l. Also, $\vec{r}$ is the position vector and $\vec{k_a}$ is the wave vector corresponding to frequency ω_a in the direction of propagation of the plane wave.

Now we can relate the components of the electric field to the components of the polarization using the susceptibility tensors. It turns out that different orders of the nonlinear susceptibility correspond to distinct physical phenomena, so it is often convenient to write them separately. [51] The second-order polarization for $\omega_3 = \omega_1 - \omega_2$ is:

$$P_l^{(2)}(\vec{r}, \omega_3) = \epsilon_o \sum_{mn} \chi_{lmn}^{(2)}(\omega_3 : \omega_1, -\omega_2) E_m(\omega_1) E_n^*(\omega_2) \mathrm{e}^{-i(\vec{k}_1 - \vec{k}_2)\cdot\vec{r}} \tag{3.6}$$

Physically, this corresponds to the case where two colors (ω_1 and ω_2) are incident on a nonlinear medium in which the difference frequency, ω_3, is generated. A similar equation can be written for $\omega_3 = \omega_1 + \omega_2$ (known as sum frequency generation) or the case where $\omega_1 = \omega_2$ and $\omega_3 = 2\omega_1$ (known as second harmonic generation). Typically, the nonlinear medium is a crystal and its lattice structure determines the relative magnitudes of the $\chi^{(2)}$ tensor elements. For crystals with a high degree of symmetry, many elements of $\chi^{(2)}$ will have zero magnitude. Equation 3.6 is commonly written in terms of these tensor elements, otherwise known as nonlinear coefficients, d_{lmn}:

$$P_l(\omega_3) = \epsilon_o \sum_{mn} d_{lmn}(\omega_3 : \omega_1, -\omega_2) E_m(\omega_1) E_n^*(\omega_2) \tag{3.7}$$

In practice, crystals are carefully chosen and cut so as to maximize frequency conversion along the propagation direction. Thus, it is convenient to define the effective nonlinear coefficient, d_{eff}. It is a scalar quantity which depends on the frequencies of interest and the crystal orientation. With this definition, Equation 3.7 can be rewritten as:

$$P(\omega_3) = \epsilon_o d_{eff} E(\omega_1) E_n^*(\omega_2) \tag{3.8}$$

Now, to get an idea of how an electric field changes as it propagates in a nonlinear material, we need to employ the wave equation derived from Maxwell's equations:

$$\nabla^2 \vec{E} - \mu_o \sigma \frac{\partial \vec{E}}{\partial t} - \mu_o \epsilon_o \frac{\partial^2 \vec{E}}{\partial t^2} = \mu_o \frac{\partial^2 \vec{P}}{\partial t^2} \tag{3.9}$$

Here, μ_o is the permeability of free space and σ represents surface currents. For clarity and simplicity, we will assume steady-state conditions, the absence of surface currents, and that the envelope of the electric field is slowly varying compared to its frequency. We will let $\vec{E} = \hat{e}E(z,t)\mathrm{e}^{i(\omega t - kz)}$ and $\vec{P}_{NL}(\omega, z, t) = \hat{p}P_{NL}(z,t)\mathrm{e}^{i(\omega t - k_p z)}$. In this case, we obtain:

$$\frac{\partial E}{\partial z} = \frac{-i\omega\mu_o c}{2n}(\hat{e}\cdot\hat{p})P_{NL}\mathrm{e}^{i(k-k_p)z} \tag{3.10}$$

where c is the speed of light, $n = \sqrt{\epsilon/\epsilon_o}$, and $\hat{e}$ and $\hat{p}$ are the unit vectors describing the directions of $\vec{E}$ and $\vec{P}$ respectively.

Substituting Equation 3.8 into Equation 3.10, we get a set of coupled wave equations describing the frequency conversion process for a second-order nonlinearity.

3.2.2 Optical Parametric Amplification

In the previous section, we showed that we can write down an equation representing the transformation of light from one frequency to another as it propagates through a material exhibiting nonlinear polarization. We will now continue this description in terms of the nonlinear process that is most important to this thesis: Difference Frequency Generation (DFG). In this process, a "pump" beam is combined with a "signal" beam in a nonlinear crystal and a third, lower frequency beam called the "idler" is generated, as illustrated in Figure 3.1. This process allows two relatively high frequencies to be down-converted so that mid-infrared light can be generated starting with only near-infrared lasers. A device which uses DFG to amplify the signal and/or idler is called an Optical Parametric Amplifier (OPA). In what follows, we will examine the equations governing DFG and optical parametric amplification in particular, emphasizing aspects of this process which influence OPA design.

In order to apply Equation 3.10 to DFG, we will consider an electric field containing three

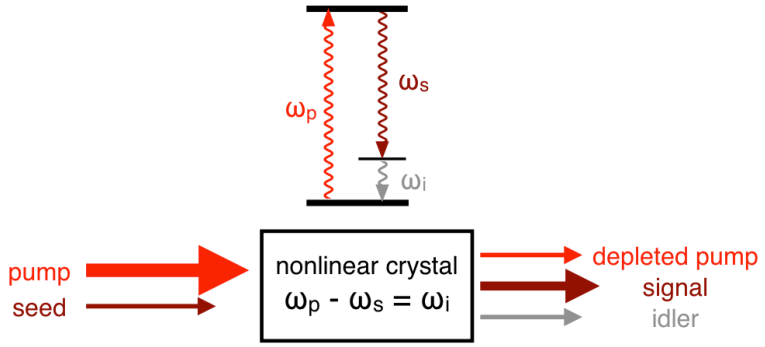

Figure 3.1: Optical parametric amplification employs difference frequency generation in a nonlinear crystal. In the crystal, incident pump photons are split into two lower-energy photons: the signal and idler. The signal energy is determined by the incident seed light and the color of the idler corresponds to the difference between the frequencies of the pump and seed.

colors, ω_p, ω_s, and ω_i corresponding to the pump, signal, and idler respectively. These colors have the relationship $\omega_p - \omega_s = \omega_i$ and $\omega_p > \omega_s > \omega_i$. Equation 3.8 allows us to write the nonlinear polarization in terms of the three colors. Substituting Equation 3.8 into Equation 3.10 results in three coupled differential equations known are the Manley-Rowe Relations [52]:

$$\frac{dE_p^*}{dz} = i\kappa_p E_i^* E_s^* \mathrm{e}^{-i\Delta kz} \tag{3.11}$$

$$\frac{dE_s}{dz} = -i\kappa_s E_p E_i^* \mathrm{e}^{-i\Delta kz} \tag{3.12}$$

$$\frac{dE_i}{dz} = -i\kappa_i E_p E_s^* \mathrm{e}^{-i\Delta kz} \tag{3.13}$$

Here E_p, E_s, and E_i are the electric fields corresponding to the pump, signal, and idler frequencies respectively, $\kappa_j = \omega_j d_{eff}/n_j c$, Δkz is the phase mismatch, and $\Delta k = k_p - k_s - k_i$.

In the regime of a strong, undeleted pump we can look for exponential solutions to this system of equations of the form of $\mathrm{e}^{\pm\Gamma z}$. Writing these solutions in terms of the pump, signal, and idler intensities (I_p, I_s, and I_i), we find: [52, 53]:

$$I_s = I_{so}\left[1 + \frac{\gamma^2}{\Gamma^2}\sinh^2(\Gamma z)\right] \tag{3.14}$$

$$I_i = \frac{\omega_i}{\omega_s} I_{so} \frac{\gamma^2}{\Gamma^2}\sinh^2(\Gamma z) \tag{3.15}$$

where I_{so} is the initial intensity of the signal, and γ and Γ are defined as follows:

$$\gamma = \sqrt{\frac{2\kappa_i\kappa_s I_p}{n_p c\epsilon_o}} \tag{3.16}$$

$$= \sqrt{\frac{8\pi^2 d_{eff}^2 I_p}{n_i n_s n_p \lambda_i \lambda_s \epsilon_o c}} \tag{3.17}$$

$$\Gamma = \sqrt{\gamma^2 - \left(\frac{\Delta k}{2}\right)^2} \tag{3.18}$$

In the case where Δk is zero, Γ and γ are equal. In this case, the signal and idler grow with propagation distance as shown in Figure 3.2. The growth is exponential when the gain is high (i.e., $\Gamma z >> 1$), and in this limit the signal and idler intensities can be written as follows:

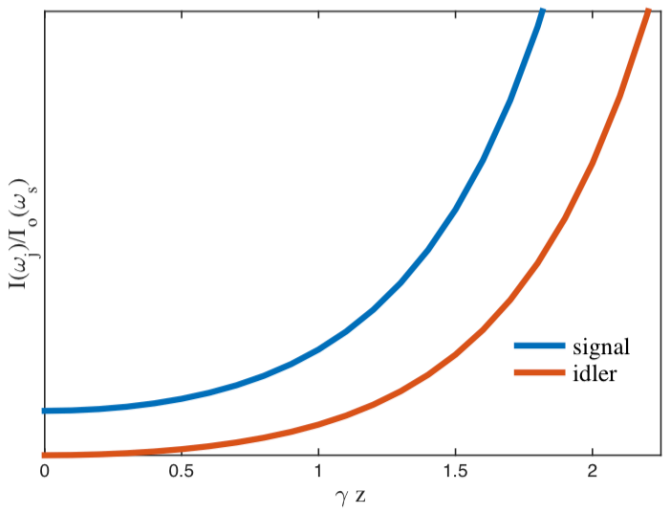

Figure 3.2: The intensities of the signal and idler beams increase as the beam propagates through a second-order nonlinear medium with perfect phase matching and in the undepleted pump approximation.

$$I_s(z) = \frac{1}{4} I_{so} \mathrm{e}^{2\Gamma z} \tag{3.19}$$

$$I_i(z) = \frac{\omega_i}{4\omega_s} \frac{\gamma^2}{\Gamma^2} e^{2\Gamma z} \tag{3.20}$$

$$\tag{3.21}$$

It is useful to define the parametric gain, G, in the high gain limit:

$$G(z) = \frac{I_s}{I_{so}} = \frac{1}{4} \mathrm{e}^{2\Gamma z} \tag{3.22}$$

Evidently, the growth rate of the signal and idler depends on Γ which in turn is proportional to the magnitude of the electric field of the pump and to d_{eff}. This implies that we can increase the parametric gain of the signal and idler by: 1) increasing the propagation distance, i.e., increasing the crystal length, 2) finding a crystal with a high value of d_{eff}, or 3) increasing the pump intensity.

Phase Matching It is also worth noting that the gain is affected by the Δk term in Γ (Equation 3.18). When Δk is zero, the gain is at its maximum – a condition known as perfect phase matching. If Δk is not zero, the phase velocities of the pump, signal, and idler will differ, causing destructive interference rather than exponential growth. What is more, this condition is not readily met in most cases, but has to be obtained through experimental tricks. The normal dispersion of some materials is enough to cause the combined phases of the colors to slip by π radians after less than one millimeter of propagation. Sapphire, for example, has a coherence length (the length required for Δk to reach π radians) of about 20 μm for idler, signal, and pump wavelengths of 3.1, 1.5, and 1.0 μm. [54] Figure 3.3 shows the effect of increasing phase mismatch on parametric gain, leading to reduced or negligible gain. This plot also shows how the coherence length, ℓ_c, shrinks as Δk grows.

One common method for achieving phase matching is to take advantage of a crystal's birefringence. A birefringent crystal will exhibit anisotropies in its index of refraction arising from its lattice structure. An incident beam may therefore experience a higher or lower index depending on its direction of propagation and polarization. [55] By choosing the orientation of a crystal with

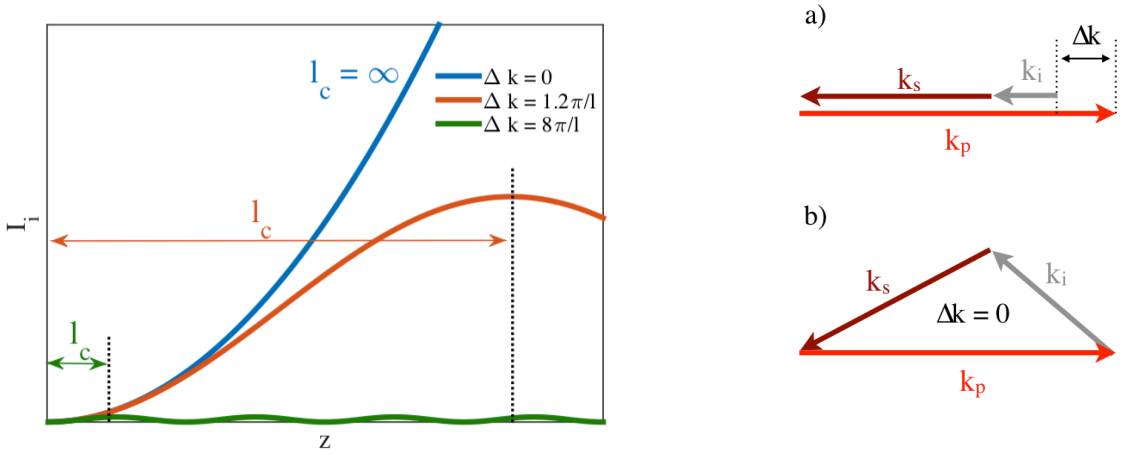

Figure 3.3: Left: The plot on the left shows the effect of phase mismatch on the growth of the idler intensity as it propagates through a nonlinear medium. The blue curve represents the intensity of the idler with perfect phase matching, the orange curve shows the effect of a small phase mismatch, and the green curve shows that a larger phase mismatch can extinguish gain almost entirely.
Right: The diagram on the right shows a) a case where the k-vector of the pump does not cancel that of the signal and idler, resulting a phase mismatch and b) a case where perfect phase matching has been achieved.

strong birefringence, the polarization of the incident beams, and the angles of the beams with respect to one another, it may be possible to make Δk vanish. [51]

Amplification Bandwidth So far, we have considered the simplified case of single, well-defined frequencies for the pump, signal, and idler, but it is important to note that, in practice, these beams have finite spectral bandwidths and that this wavelength spread will affect phase matching. This is an especially important consideration for the work described in this thesis because we are interested in short pulses and short pulses require broad bandwidths. A ten-cycle, gaussian pulse with a central wavelength of 3 μm, for example, requires at least 130 nm full-width-half-maximum bandwidth. It is not necessarily the case that the entire range of frequencies in such a pulse will be phase matched. In fact, since $\vec{k}$ is defined in terms of the wavelength and index of refraction ($k = 2\pi n/\lambda$), we expect that the relationship $\Delta k = k_p - k_s - k_i = 0$ will generally not hold as λ_p, λ_s, and λ_i vary from their corresponding central wavelengths.

Fortunately, some crystals can accommodate frequency mixing over wide spectral ranges due to the specific ways their indicies of refraction vary with wavelength. Quantitatively, the acceptance bandwidth of a crystal can be ascertained by expanding Δk about the central frequency of the signal, ω_s, holding the pump frequency constant. For simplicity, we will keep only the first order term in the expansion, given in Equation 3.23. [56]

$$\frac{\partial \Delta k}{\partial \omega_s}\Delta\omega_s = \left(\frac{\partial k_s}{\partial \omega_s} - \frac{\partial k_i}{\partial \omega_s}\right)\Delta\omega_s \tag{3.23}$$

Here, the partial derivatives of the signal and idler wave vectors with respect to ω_s are recognizable are the group velocities of the signal and idler: v_{gs} and v_{gi}.

The parametric acceptance bandwidth is defined as the wavelength range over which the absolute value of the phase is less than or equal to π. The phase mismatch accumulated over the length of a crystal, ℓ, is Δkℓ corresponding to a drop in power conversion efficiency to about 40 percent. Putting this together, we can write the parametric bandwidth in terms of wavelength as in Equation 3.24. [56]

$$\Delta\lambda_P = \frac{\lambda^2}{c\ell}\left|\left(\frac{1}{v_{gs}} - \frac{1}{v_{gi}}\right)^{-1}\right| \tag{3.24}$$

It is important to note that the parametric bandwidth is inversely proportional to the crystal length whereas the conversion efficiency increases with length. This means that when deciding how to cut a crystal, there is a trade-off between gain and amplified bandwidth. If very broad bandwidth is needed, conversion efficiency must be compromised and vice versa.

Amplification bandwidth can also be defined in terms of gain, which is of practical importance when designing an OPA. This definition gives the maximum width of the spectrum that we can expect to amplify in a given OPA configuration. This gain bandwidth is defined by how far we can deviate from the signal's central signal frequency before the gain is cut in half. Using Equation 3.22 and the high-gain approximation, the gain bandwidth is [53]:

$$\Delta\lambda_G = \left|\left(\frac{1}{v_{gs}} - \frac{1}{v_{gi}}\right)^{-1}\right| \frac{2\sqrt{\ln 2}\lambda^2}{\pi c}\sqrt{\frac{\gamma}{\ell}} \tag{3.25}$$

As we saw with the parametric bandwidth definition, the gain bandwidth decreases as the length of the crystal increases, resulting in a trade-off between bandwidth and gain as a function of crystal length. The good news is that the gain bandwidth is also proportional to the square root of γ. γ, as defined by Equation 3.18, is proportional to the pump intensity and the square of d_{eff}. This means we can increase the amplified bandwidth and the gain simultaneously if we can increase the pump intensity and/or find a crystal with higher d_{eff}. Therefore, it is doubly advantageous to find a crystal that has both a high d_{eff} and a high damage threshold.

Gain Saturation Another important consideration for designing an OPA is that when the number of seed photons is significant compared to the number of pump photons, the undepleted pump approximation may fail. In this case, the gain saturates. That is, the signal and idler intensities no longer increase exponentially with propagation, but the growth rate slows as the pump photons are used up. Eventually the gain will reach zero and back-conversion will begin. [57]

When trying to push the amount of amplified signal as high as possible, making full use of

available pump photons, it is helpful to anticipate where saturation will occur so the OPA can be driven to the point where it just reaches saturation. This allows maximum gain for a set number of pump photons and minimizes the transfer of any amplitude jitter on the seed to the amplified pulse. How far a beam can travel in a crystal before gain saturation occurs depends on all the crystal parameters which influence gain (d_{eff}, indices of refraction, length, etc.) and on the intensity and total energy in the pump. It can be modeled using a split-step propagation algorithm. [57] Alternatively, the point of gain saturation can be experimentally determined by increasing the pump intensity until the amplified beam starts to exhibit concentric rings. These appear since the center of the seed beam has the highest intensity and therefore will deplete the pump faster than the wings of the mode. Thus, after saturation is reached in the center of the beam the wings will continue to be amplified, resulting in ring structure on the mode.

If the amplified pulse is temporally dispersed, saturation will also affect the amplified spectrum. For a smooth spectrum with spectral intensity that peaks at the center wavelength, the center of the spectrum will saturate before the wings. This causes the spectrum to be reshaped, with a dip in the center and shoulders to either side. This can be useful for spectral broadening, but may also add temporal structure to the compressed pulse. [57]

Quasi Phase Matching For any given application, it is necessary to find a nonlinear crystal that can be phase matched for the process of interest, has a high nonlinear coefficient, sufficient acceptance bandwidth, an appropriate transparency range, and desirable mechanical properties such as hardness and high damage threshold. Needless to say, it can be challenging to find a material that fulfills all these requirements. For this reason, a technique known as Quasi Phase Matching (QPM) has been developed which can loosen the phase matching constraint.

Rather than eliminate Δk, quasi phase matching allows for parametric gain in the presence of a phase mismatch by periodically compensating for it. For example, quasi phase matching can be realized in a crystal by periodically reversing its polarization axis. The thickness of each polarization region is designed to be equal to one coherence length for the nonlinear process of interest. This ensures that just as the phases have slipped by π radians the polarization, and therefore also the

phase slip, reverse direction. As a result, the absolute value of the phase slip never exceeds π and the signal and idler continue to grow over the length of the crystal. [58, 51] Figure 3.4 illustrates this process.

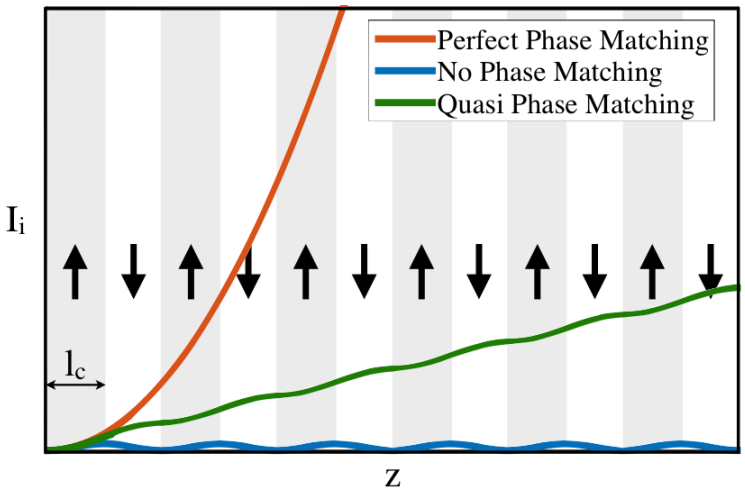

Figure 3.4: Growth of the idler intensity versus propagation distance, comparing the cases of perfect phase matching, quasi phase matching in a periodically poled material, and no phase matching.

It is evident from Figure 3.4 that quasi phase matching is not as efficient per unit length as perfect phase matching, but its advantages are manifold. In addition to allowing parametric amplification in crystals with birefringence that is too weak to achieve perfect phase matching, periodically poled crystals have a large degree of design flexibility. This results in freedom to tailor the gain spectrum for a particular application and the possibility of achieving extremely broad acceptance bandwidths. [59, 60] Though the fabrication of these crystals can be challenging, improvements in crystal growth techniques have allowed periodically poled crystals to become widely available for the practical implementation of quasi phase matching over the last two decades. [61]

Gain Bandwidth of QPM Devices Quasi phase matching can also offer advantages in terms of bandwidth. For a periodically poled crystal, the expression of the gain bandwidth is the same as Equation 3.25 except that Δk must be replaced by the total phase mismatch, κ, which is the

sum of Δk and the contribution from the poled grating itself. When the crystal has poling period, Λ, its wave vector is $K_g = 2\pi/\Lambda$ and, the total phase mismatch is therefore $\kappa = \Delta k - K_g$. [62]

For a perfectly phase matched crystal, Equation 3.25 tells us that the bandwidth is inversely proportional to the square root of the crystal length, and proportional to d_{eff} and the square root of the pump intensity. Though the equation is the same, QPM can be useful for increasing bandwidth when it allows gain along a crystal axis with a much higher d_{eff} than would otherwise be possible. Periodically Poled Lithium Niobate (PPLN), for example, allows mid-infrared difference frequency generation with $d_{eff} = 14.9$ pm/V whereas bulk lithium niobate phase matches this process with only $d_{eff} = 4.0$ pm/V. [63]

Periodically poled crystals have the additional advantage that the periodicity can be designed to broaden the bandwidth. One way to do this is by increasing or decreasing the period of the grating along the length of the crystal. In such a "chirped" crystal, the wavelength for which κ is exactly zero changes through the length of the crystal and the effective gain bandwidth is therefore broadened. [62] Even more sophisticated grating patterns can be used to further broaden the gain bandwidth and/or to control its shape. The poling period, for example, can be apodized on each side of the crystal to ensure smooth nonlinear coupling and therefore a smooth amplification spectrum and temporally clean pulses. [64] Another technique employs adiabatic chirping to achieve a supercontinuum gain bandwidth that spans 3000 nm centered at $\lambda = 3.5\ \mu m$. [65]

3.3 Optical Parametric Chirped Pulse Amplification

This next section addresses the problem of scaling the power and energy of an OPA to the level required for keV high harmonic generation. The challenge is apparent in that our goal is to create such an intense electric field that we can ionize helium and impart large amounts of energy to the freed electrons. That is, around $10^{15}\ W/cm^2$. [5] The challenge is to deliver these light pulses, designed to efficiently ionize the matter they interact with, to the helium-filled waveguide still intact. We need the light to do its work on the helium without damaging optics in its path or accumulating nonlinear phase through interactions with solid material or even air. Such interactions

can distort the beam's spatial profile and add dispersion, lowering the pulse's peak intensity.

There are several ways to control peak intensity without changing the energy per pulse. The simplest of these is to increase the beam diameter, spreading out the energy in space and consequently reducing the intensity. We will make use of this by focusing into the helium-filled waveguide and maintaining a larger beam diameter elsewhere. However, this alone is not enough. If the beam diameter required to reach $10^{15}\ W/cm^2$ were $100\mu m$, then its diameter would need to be expanded by a factor of 1000 to bring the intensity down to the relatively safe level of $10^9\ W/cm^2$. In this example, that would correspond to an unwieldy 10 cm beam. In practice, maximum beam diameter is limited by considerations such as thermal transfer in the the pump laser and by the expense, lower quality, and availability of large-aperture optics.

Another way to address this problem is to stretch the beam in time with a technique called Chirped Pulse Amplification (CPA), allowing the beam diameter to be kept to a manageable size. The stretching, otherwise known as "chirping," is accomplished by imparting a spectral phase to the pulse. After amplification, the stretch can be undone by imparting phase of the opposite sign. The phase can be added using dispersive material, diffraction gratings, Bragg reflectors, or anything else that will delay the reddest frequencies of pulse with respect to the bluest ones in an appropriate way.

CPA can be used to stretch a pulse to many times its original duration and has been repeatedly demonstrated as an effective tool for creating high-intensity pulses. This technique was introduced by Strickland and Mourou in 1985 [66] and has enabled lasers to reach terawatt and even petawatt peak intensities. [19, 67] It is now a well established and routinely used technique.

The effect of spectral phase on pulse duration can be understood through the Fourier transform relationship between the temporal and spectral representations of a laser pulse. Say, for example, we start with a transform-limited pulse with a gaussian envelope and send it through a dispersive material. Dispersive materials, by definition, delay certain spectral frequency components more than others. Mathematically, this is equivalent to taking the Fourier transform of the pulse and multiplying by a phase factor, $e^{i\phi(\omega)}$. The inverse Fourier transform then represents the

output pulse.

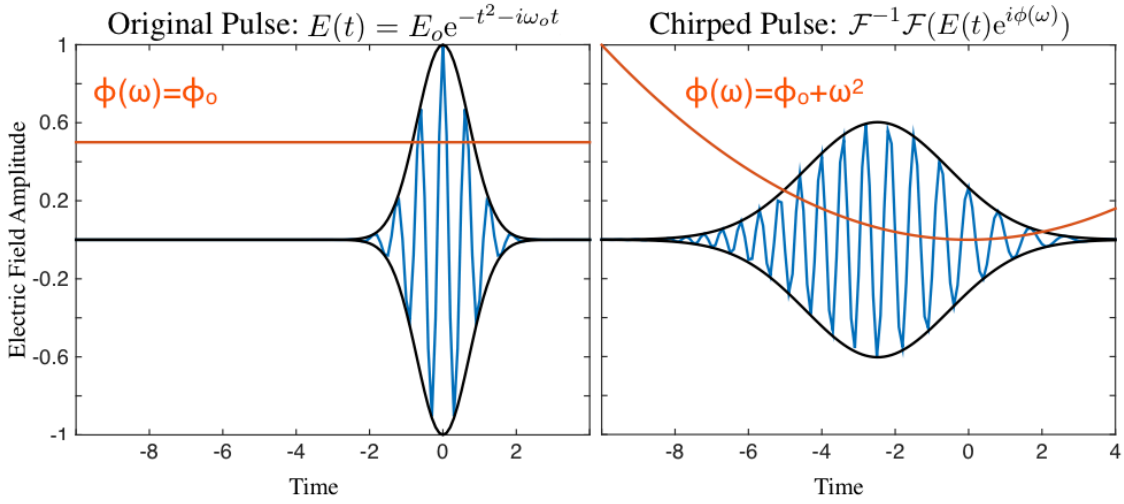

Figure 3.5: The electric field of a pulse before (left) and after (right) the application of a quadratic spectral phase. A quadratic spectral phase imparts a linear chirp and temporally spreads the pulse.

$$E_{in}(t) = E_o e^{-t^2 - i\omega_o t} \tag{3.26}$$

$$E_{out}(t) = \mathcal{F}^{-1}(e^{i\phi(\omega)}\mathcal{F}(E_{in}(t))) \tag{3.27}$$

Figure 3.5 illustrates the effect of a quadratic spectral phase on the pulse. The delay of a particular frequency component with respect to a reference frequency, otherwise known as the group delay, is equal to the frequency derivative of the phase: $T_g = \partial\phi(\omega)/\partial\omega$. Therefore, for a positive quadratic phase, the blue side of the spectrum will be delayed more than the red side and the delay will linearly increase with frequency, effectively stretching the pulse in time. A pulse stretched in this way is said to have a linear chirp. In Figure 3.5, this effect is manifested in that the electric field oscillations on the leading edge of the pulse are slower than those on the trailing edge.

The amount of stretch needed in order to prevent damage can be estimated using the fact that damage threshold scales like the square root of pulse duration for nanosecond to 100 picosec-

ond pulses. For pulses shorter than 100 ps, damage occurs even more readily because damage mechanisms such as ablation begin to dominate. [68] Additionally, the amount of nonlinear phase distortion that the beam will experience can be quantified by integrating the nonlinear index and intensity over the beam path, a quantity known as the B-integral.

CPA can be straight-forwardly applied to Optical Parametric Amplification by replacing the laser gain medium with the nonlinear crystal in which difference frequency mixing will take place. The adaptation of CPA to OPA is known as Optical Parametric Chirped Pulse Amplification (OPCPA). Figure 3.6 illustrates this adaptation: in both cases, the input beams are stretched, then amplified, and finally recompressed. The primary difference is that a CPA chain has only one color seeding it and uses a laser gain material for amplification whereas an OPCPA system is seeded by two colors and amplification occurs in a nonlinear material.

An additional complication unique to OPCPA and relevant to the design of stretchers and compressors for such a system is the phase relationship between the signal and idler beams. This relationship can be derived from Equation 3.13 by separating out the argument of the complex fields on both sides of the equation. We obtain:

$$e^{i\phi_i} = -ie^{i\phi_p - i\phi_s} \tag{3.28}$$

Here ϕ_i, ϕ_p, and ϕ_s are the phases of the idler, pump, and signal respectively. From Equation 3.28 we can readily see that $\phi_p - \phi_s - \phi_i = \pi/2$ and therefore:

$$\phi_i = \phi_p - \phi_s - \pi/2 \tag{3.29}$$

For the difference frequency mixing process we also know:

$$\omega_i = \omega_p - \omega_s \tag{3.30}$$

For a narrow-band pump with a flat phase (which is a good approximation for the work described in this thesis), Equation 3.30 implies that the spectra of the signal and idler will be

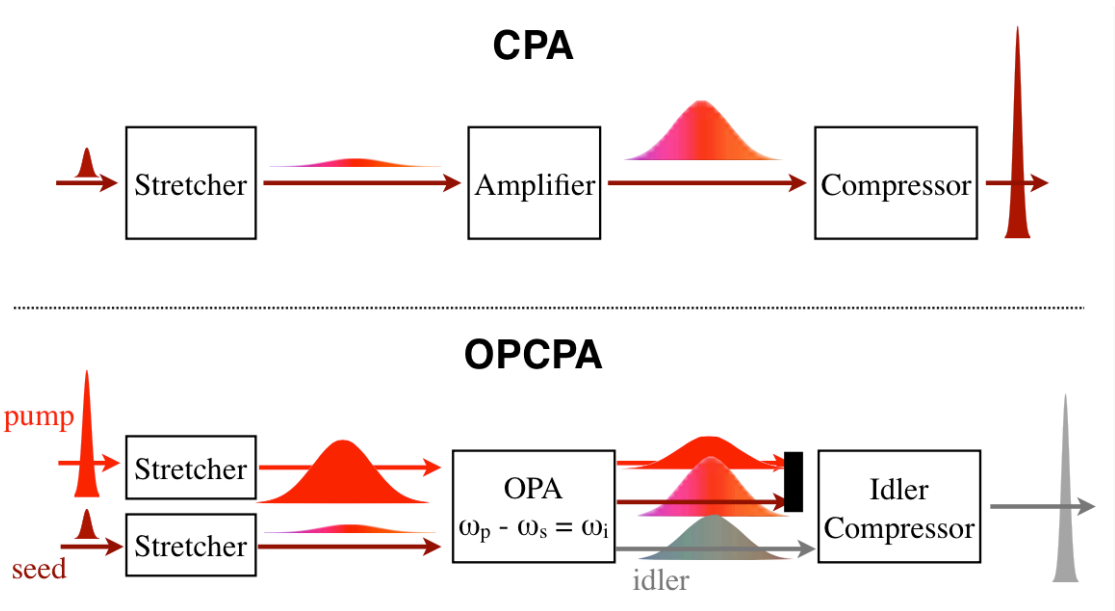

Figure 3.6: Illustration of chirped pulse amplification (top) and optical parametric chirped pulse amplification (bottom).

mirror images of one another when plotted versus frequency. Equation 3.29 then tells us that the phases will not only also be mirrored, but also inverted with a constant offset. With a mirrored and inverted phase, the odd terms in the Taylor expansion of the signal's and idler's phases will have the same sign and even terms will have opposite signs.

This phase relationship is important if you want to stretch the signal, but compress the idler. This is often a necessary arrangement due to the fact that the idler is usually not present before amplification. However, this is challenging because it requires imparting dispersion on the signal that can be taken out by the compressor. This is not hard to do if the stretch is a purely linear chirp (quadratic phase) because there are several techniques for adding both positive and negative linear chirp. The difficulty comes in with the higher-order phase terms: if we impart a positive quadratic and cubic phase (and there is almost always some cubic phase in real situations) then we need to compress by adding negative quadratic phase and positive cubic phase in the correct proportions. This issue can put some serious design constraints on system design and will be discussed in more depth in the next chapter.

Chapter 4

OPCPA Sytem Design

4.1 Overview

The OPCPA system at the center of this thesis is designed to create and amplify an ultrashort, 3 μm wavelength beam for HHG using parametric amplification of chirped pulses. As illustrated in Figure 4.1, it consists of a high-energy pump and a seed, pulse stretchers for each, a series of parametric amplification stages where the pump and signal are combined to create and amplify the mid-IR idler, and a compressor for the idler.

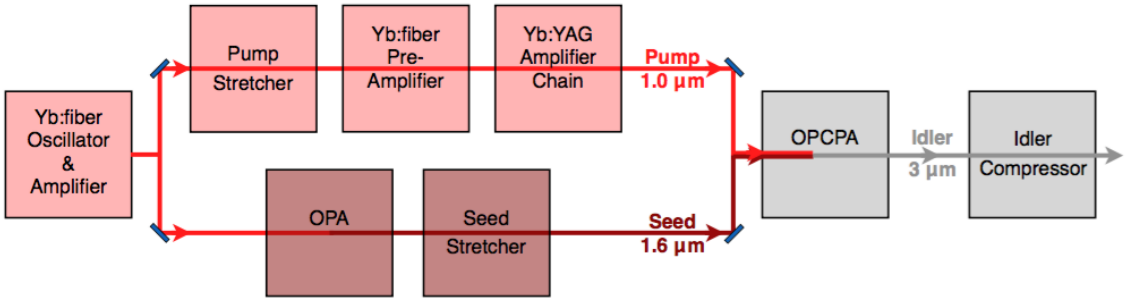

Figure 4.1: Schematic Diagram of the OPCPA system, starting with 1.0 μm light (drawn in red) from the Yb:fiber laser on the left which is split into a 1.0 μm pump arm and 1.6 μm seed arm (drawn in maroon). The pump and seed meet again in the OPCPA itself where 3 μm light is generated (grey).

In what follows, we will describe the design of this laser system, breaking it up into its main components: the pump laser, 1.6 μm seed source, parametric amplification stages, and compressor. We will discuss our reasoning for these design choices, the new problems that we encountered and

addressed along the way, and remaining challenges in improving each part of the system.

4.2 Pump Laser Architecture

4.2.1 Past Work

So far, the development of mid-infrared OPCPA technology primarily fits into two categories: 1) low pulse energy systems with high repetition rates and few-cycle pulses, and 2) multi-mJ pulse energy, low-repetition rate systems. In the first category, 50-160 kHz systems operating with center wavelengths between 3.1 and 3.4 μm have reached pulse energies up to 22 μJ and pulse durations as short as 42 fs (less than 4 optical cycles). [69, 70, 71] Among OPCPA lasers operating above 3 μm wavelength, these have demonstrated the highest average powers at up to 3.2 W.

In spite of their attractive average power capabilities and short pulses, these laser systems are not ideal for driving keV x-ray harmonics due to their relatively low pulse energies. The intensities required to reach keV harmonics with tens of microjoules would require a tight focusing geometry, precluding the achievement of absorption-limited harmonic generation. Additionally, recent work by Hernandez-Garcia et al. indicates that group velocity walk-off between the driving laser and the harmonics severely limits the phase-matching window when using few cycle pulses. Instead, eight to ten cycle pulses are preferable for phase matching mid-IR-driven harmonics. [28]

In the second category, Andriukaitis et al. developed an OPCPA laser that produces 8 mJ pulses with center wavelength of 3.9 μm at a repetition rate of 20 Hz and compressed to 83 fs. [72] Enabled by its impressive pulse energy, this laser was used to demonstrate the first phase matched high harmonic generation out to 1.6 keV ($\lambda = 0.8$ nm). [5] The use of this harmonic spectrum for applications is limited, however, due to the low repetition rate which yields correspondingly low x-ray flux.

In order to produce an x-ray harmonic source suitable for applications, we need to aim for a regime between these two categories. Specifically, our goal is a mid-infrared driving laser with millijoule-level pulse energy that also maintains high average power by running at kilohertz

repetition rates. Compared to the work by Mayer et al., we aim to boost the pulse energy by a factor of 45 and maintain Watt-level average power, but will put a lesser requirement on pulse duration, aiming for roughly 100 fs rather than 40 fs pulse durations. [71] On the other hand, compared to the work described by Andriukaitis et al., this is a decrease in pulse energy by a factor of 8, but an increase in average power by a factor of 60. [72]

In this regime balancing high average and peak power, Deng et al. have developed an OPCPA laser with 1.2 mJ pulse energy, 3 kHz repetition rate, and 2 μm center wavelength. [73] Recently, Lai et al. reported an OPCPA laser with 2.6 mJ pulse energy at 1 kHz repetition rate and 2 μm center wavelength. [74] These lasers are the closest to our goal, but the center wavelength falls below that needed to drive keV harmonics. To our knowledge, there are no such ultrashort pulse, millijoule-energy, and kilohertz repetition rate laser systems operating with center wavelength in the 3 – 5 μm regime. This is precisely the gap we aim to fill with the work described in this thesis.

The OPCPA system described in this thesis represents the first (to our knowledge) femtosecond, millijoule-level laser source with wavelength above 3 μm and kilohertz-level repetition rates. Specifically, this laser operates with 3.1 μm central wavelength and 1.3 mJ pulse energy at 1 kHz. Its is unique in its efficient conversion of pump light to signal and idler while maintaining good idler mode quality and bandwidth, reaching up to 35 percent power conversion efficiency in the final amplification stage. This efficiency is, in part, enabled by the excellent beam quality of the OPCPA pump, a home-built cryogenically cooled Yb:YAG amplifier chain. Also, we use magnesium oxide-doped Periodically Poled Lithium Niobate (MgO:PPLN) as the nonlinear medium, allowing high gain and collinear amplification geometry. A tightly focusable idler mode is important for efficiently coupling the beam into a capillary for HHG.

Another important feature of this laser system is its stable and tunable front end. The first element of this laser system is a commercially-available Yb:fiber oscillator and amplifier with turn-key operation. It enables passive stabilization of the OPCPA pump and seed pulse timing because it supplies enough pulse energy to both generate a 1.6 μm beam and to seed the Yb:YAG amplifier chain which delivers the OPCPA pump. The repetition rates of the pump and seed pulses

are therefore inherently locked. Furthermore, the 1.6 μm beam is generated with an OPA which provides clean, quiet, easily tunable, and relatively broadband seed pulses for the OPCPA.

4.2.2 Requirements for the OPCPA Pump Laser

One of the largest obstacles in reaching this high average and peak power OPCPA regime is building a suitable pump laser. High peak power requires mitigation of unwanted nonlinear effects and care to prevent optical damage of laser optics. High average power, on the other hand, goes hand in hand with increased waste heat, requiring efficient thermal management.

In terms of peak power required, in order to reach the millijoule pulse energy regime at 3 μm, we will need about 16 mJ in the pump assuming about 25 percent overall conversion from the pump to the signal and idler in three stages. We will therefore aim for the 30 mJ range in order to allow for potential losses and room for pushing the idler energy beyond 1 mJ in the future. At these energies the pulses need to be temporally stretched to the 100 picosecond range to avoid most unwanted nonlinear interactions with transmissive optics in the beam path. Also, the stretch will allow us to use beam sizes which easily fit within the aperture of one-inch diameter optics without causing damage.

In terms of average power, we will require a pulse repetition frequency of 1 kHz at least (i.e., 30 W at 30 mJ). We have designed the system to be scalable to higher repetition rates, gaining average power with some cost to pulse energy due to thermal loading.

Finally, we need the pump laser to exhibit excellent mode quality. For the parametric conversion in three cascaded OPCPA stages to be efficient and result in an idler beam that can be tightly focused into a capillary, the pump mode must also be round and smooth, free from strong astigmatism or high-order mode structures.

4.2.3 Gain Material

We chose to use cryogenically cooled Ytterbium-doped Yttrium Aluminum Garnet (Yb:YAG) as the gain material for the OPCPA pump laser. This material has several advantages – mechanical,

thermal, and optical – which make it an excellent choice for high average and peak power operation.

Yb-doped materials have promise for power scaling because they exhibit a low quantum defect. The quantum defect, q, is defined as the difference in the energies of a pump photon ($h\nu_P$) and a laser photon ($h\nu_L$):

$$q = h\nu_P - h\nu_L \tag{4.1}$$

This sets an upper limit to the power efficiency of the laser since the energy difference between the absorbed and emitted photons is dumped as heat in the laser material. The fraction of absorbed pump power that becomes heat, χ_{QD}, is related to the quantum defect according to [75]:

$$\chi_{QD} = \frac{q}{h\nu_P} \tag{4.2}$$

$$= 1 - \frac{\lambda_P}{\lambda_L} \tag{4.3}$$

Where λ_P and λ_L are the pump and laser photon wavelengths, respectively. For Yb:YAG, λ_L is 1030 nm and the most convenient pump wavelength is 940 nm. The fraction of the absorbed pump power that becomes heat is therefore 0.087. In comparison, the Ti:sapphire laser is typically pumped by 532 nm light and lases at about 800 nm, so χ_{QD} is 0.34.

At room temperature, Yb:YAG's small quantum defect has the consequence that thermal excitations can destroy the population inversion of the lower laser energy level. Room temperature Yb:YAG is therefore a quasi-three-level laser suffering from the need for significant pump density before it reaches transparency, high threshold pump power, and low overall efficiency. [76] At cryogenic temperatures, however, the Boltzmann energy, k_BT is very small compared to the energy gap needed to destroy the population inversion and Yb:YAG functions as a four-level laser. At the temperature of liquid nitrogen, for example, k_BT is about 7 meV whereas the energy difference between the ground state and the lower laser level is 76 meV.

Cryogenic cooling has several other positive effects on Yb:YAG. One of these is increased efficiency. Yb:YAG's spectral absorption and emission lines narrow and brighten at low tempera-

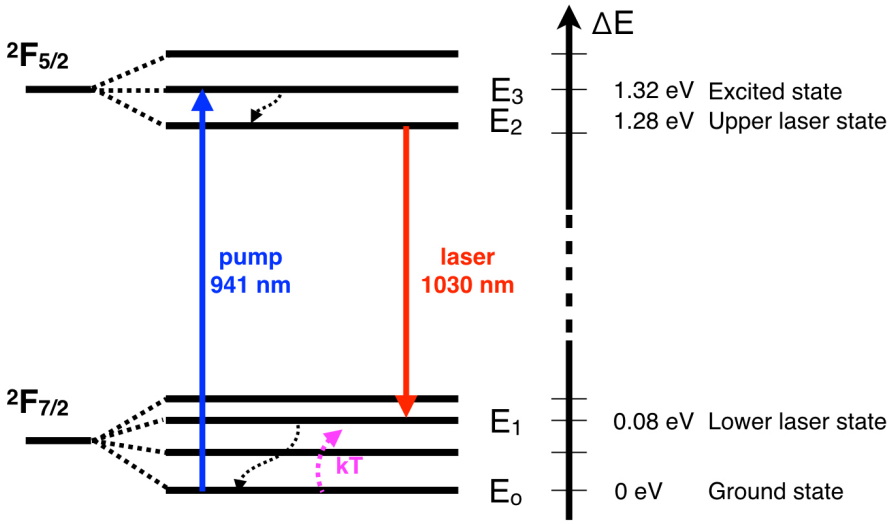

Figure 4.2: Energy level diagram showing the pump and laser transitions. The dashed arrows represent nonradiative energy transfer. [77]

ture: its emission cross section at 1030 nm increases by about a factor of four at 77 K compared to room temperature and its absorption cross section at 941 nm increases by about a factor of two. In terms of laser operation, this reduces the saturation fluence and improves overall efficiency. [76, 78]

Second, at cryogenic temperatures the thermal conductivity is improved by nearly a factor of four from 12 W/mK to 46 W/mK. [75] Not only is quantum defect heating small, but Yb:YAG is also excellent at facilitating the transfer of dumped heat out of the crystal to a cold reservoir.

Third, Yb:YAG's ability to maintain good beam quality at high power is improved. The Coefficient of Thermal Expansion (CTE) is decreased from 6 ppm/K at 300 K to 2 ppm/K at 100 K. As heat is dumped into the crystal by the pump laser, the crystal will expand, bulging at its faces and pressing against its mount. This causes beam distortions and stress that, in certain cases, could lead to severe birefringence or fracture. The dramatically reduced CTE mitigates these issues. Also, the thermooptic coefficient, dn/dT, is reduced from 8 ppm/K at 300 K to 1 ppm/K at 100 K. [75] The thermooptic coefficient is a measure of how strongly the index of refraction depends on temperature gradients. A material exhibiting a non-zero thermooptic coefficient will take on an index profile mirroring the temperature profile caused by an incident pump beam, acting like a lens. This effect is known as thermal lensing and can be difficult to compensate, leading to beam distortions. This factor of eight decrease in dn/dT therefore increases the attainable average power in Yb:YAG.

Another important feature of Yb:YAG is that it can be pumped by commercially available 940 nm diode lasers. The crystal is not naturally birefringent so the pump can be unpolarized. Also, it has a long (1 ms) upper state lifetime, allowing for continuous wave pumping. Finally, although the absorption spectrum narrows at cryogenic temperatures, the absorption line is still broad enough to facilitate diode pumping without wavelength stabilization. Together, these properties allow us to use readily available, compact, and robust diode pump lasers which are much less expensive than the pulsed green lasers commonly used to pump Ti:sapphire.

In summary, we chose Yb:YAG as the gain material for our OPCPA pump laser because its mechanical robustness, thermal properties, and minimal optical distortions all facilitate power

scaling. Furthermore, the ability to use diode lasers makes the system far simpler and less expensive than some other common solid-state laser systems.

4.2.4 Oscillator Front End

The front end of this laser system is a commercially available Y-Fi laser from Kapteyn-Murnane Laboratories: a Yb:fiber-based oscillator and amplifier. It both seeds the Yb:YAG amplifier chain with 1.0 μm light and is used to pump an OPA that generates the seed for the OPCPA signal at 1.5-1.7 μm. Since both the OPCPA seed and pump originate from the same source, our laser system's timing is passively locked. Also, the Yb:fiber oscillator and amplifier are capable of running continuously and stably, allowing the OPA and Yb:YAG amplifiers to function consistently over long periods of time.

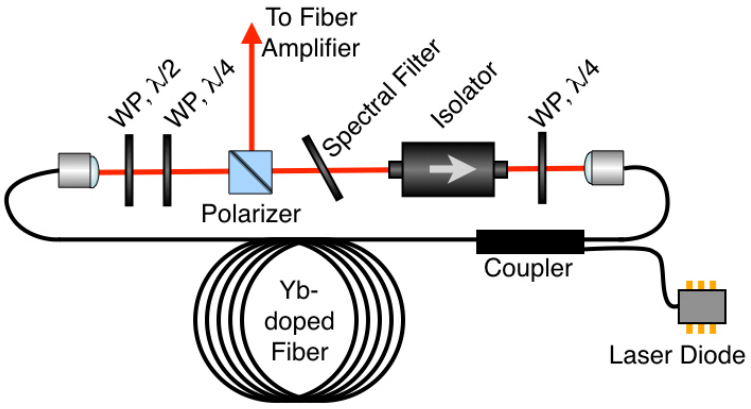

Figure 4.3: Illustration of an ANDi fiber oscillator.

Fiber oscillators are attractive for the front end of a large laser system because they are generally more stable, compact, and relatively insensitive to misalignment compared to their solid-sate counterparts. Furthermore, the oscillator at the heart of this laser system has the advantage of being an All-Normal-Dispersion (ANDi) Yb:fiber laser, as illustrated in Figure 4.3. This type of fiber oscillator relies on the combination of polarization rotation, spectral filtering, and self-phase modulation to produce a pulse train. In contrast to fiber lasers that rely on dispersion compensation to create pulses, the ANDi laser runs in a modelocked state with highly chirped pulses. This chirp

(i.e., temporal stretch) reduces peak intensity and therefore allows high pulse energies in the small fiber core that would otherwise cause damage. [79, 80]

Our ANDi laser operates with about 400 mW output at 60 MHz repetition rate (7 nJ per pulse). It is then pulse-picked down to 5 MHz using an acousto-optic modulator before being further amplified to 10.2 W in a Large Mode Area (LMA) ytterbium-doped gain fiber. The resulting 2 μJ pulses are then compressed down to 130 fs using a transmission grating pair compressor. The 5 MHz repetition rate was chosen as the highest repetition rate possible while maintaining this high pulse energy and short pulse duration.

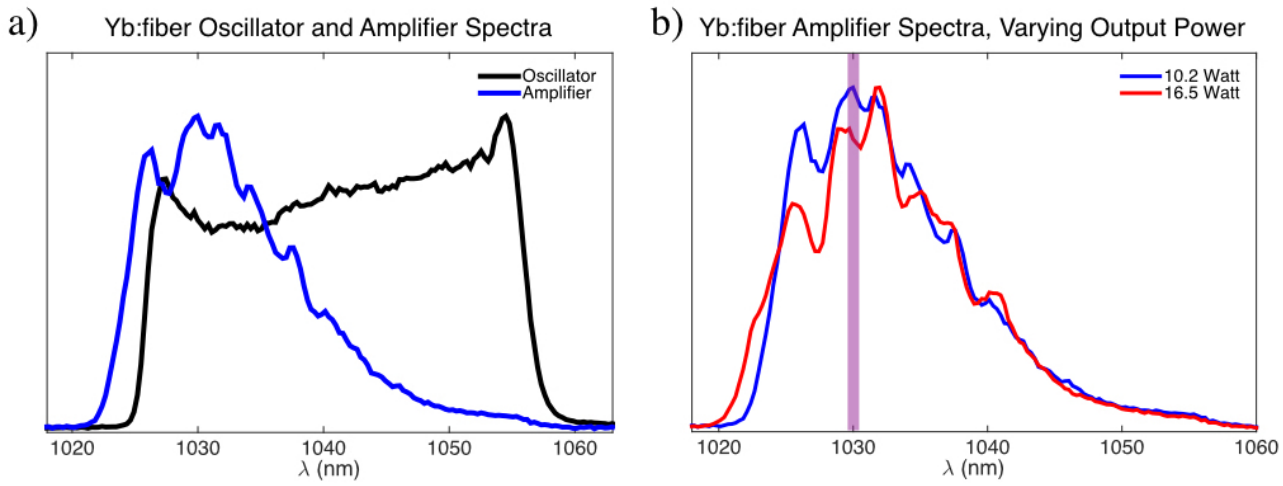

Figure 4.4: a) Yb:fiber oscillator and amplifier spectra under optimal operating conditions for the OPCPA. b) Amplifier spectra optimized for clean pulses (10.2 W, blue) and for power at the expense of compressibility (16.5 W, red). The spectral region used to seed the Yb:YAG amplifiers is highlighted in purple.

The spectra of the oscillator and amplifier are shown in Figure 4.4 a). The oscillator exhibits the ANDi laser's characteristic square spectrum with "ears" and has about 30 nm of spectral bandwidth. The LMA amplifier shifts the spectrum slightly to the blue and reshapes it.

Depending on the amplifier gain and the oscillator pulse that seeds it, this reshaping may result in strong amplitude modulations in both the spectral and temporal shapes of the amplifier output. Severe modulations will limit the pulse compressibility and may result in unacceptable

levels of noise in the OPA and Yb:YAG amplifier's output. In our laser, these modulations were minimized by choosing the pump power for the amplifier such that it gives the smoothest spectrum. This results in 10.2 W output. The power can be increased to 16.5 W, but at the expense of pulse compressibility. The difference in the amplified spectra at the optimal operating power versus full power are shown in Figure 4.4 b).

Alternative Oscillators Explored: Yb:Fiber Oscillator Only In an earlier design of this laser system we used a home-built Yb:fiber ANDi oscillator as the front end of this laser system rather than the commercial Y-Fi described above. [81] Once the Y-Fi became available, we decided it was a better option due to its higher pulse energy, however the Yb:fiber oscillator alone has some advantages that are worth mentioning.

First, the oscillator output gave us 1.0 W at 63 MHz (16 nJ per pulse). The main advantage is that higher repetition rate (the Y-Fi operates at 5 MHz) makes it easier to match the delay of the OPCPA seed with that of the pump. A maximum of only 16 ns delay is needed for a 63 MHz pulse train whereas up to 200 ns delay is needed at 5 MHz. With the higher repetition rate, shorter sections of fiber delay are needed, reducing potential for losses, time-dependent polarization rotation, and dispersion in the delay fiber. These are not insurmountable problems with a 5 MHz pulse train, but the higher repetition rate makes them trivial.

Since it does not have enough pulse energy to both saturate an OPA and to seed the Yb:YAG amplifiers, we used this Yb:fiber oscillator to drive an Optical Parametric Oscillator (OPO) instead. The OPO makes efficient use of the oscillator's output by circulating the light in an actively length-controlled cavity so that the spontaneously emitted signal from a DFG crystal builds up, forming a coherent beam with wavelength in the 1.4-1.7 μm range to seed the OPCPA. That the OPO depends on active cavity length control and a feedback loop makes its output spectrum more likely to drift than that of an OPA. Also, its alignment is much more sensitive to small changes, compromising the overall stability of the OPCPA system. This stability problem was the main motivation for moving to a high pulse energy oscillator/amplifier combination capable of pumping an OPA.

Alternative Oscillators Explored: Yb:YAG Oscillator In an even earlier approach to

the front-end design of this laser system, we explored the possibility of using a passively modelocked, solid state Yb:YAG oscillator. We built such an oscillator based on cryogenically cooled Yb:YAG, demonstrating its high efficiency, high power, and passive mode locking. [82]

This oscillator achieved continuous-wave-modelocking with the combination of a saturable absorber mirror and an intracavity frequency doubling crystal which acts as a negative saturable absorber. Compared to a fiber laser, this oscillator can be designed with a large mode size and with less unwanted nonlinear dispersion so that it is capable of much higher power at a similar repetition rate: 12 W (200 nJ) rather than 1 W (16 nJ). Cryogenically cooled Yb:YAG also has the advantage of high pump-to-laser conversion efficiency – we reached 75 percent power conversion in continuous wave operation.

In spite of its attractive power output, this Yb:YAG oscillator lacks the robustness of a fiber laser that is key for the front end of a complex laser system. It was susceptible to falling into Q-switched modelocked operation and to pulse amplitude bifurcation. Additionally, alignment for modelocking was difficult and sometimes time-consuming. Therefore, fiber lasers offer a much more robust and compact alternative. The addition of the LMA fiber amplifier to the ANDi oscillator makes it by far the best of these choices for our design, bringing the pulse energy a factor of ten above that of the solid state oscillator.

4.2.5 Stretcher

When seeded with a an unmodified pulse from the Yb:fiber laser, the cryogenic Yb:YAG amplifiers will produce pulses on the order of ten picoseconds in duration. At the desired pulse energy of 30 mJ and assuming a 6 mm beam diameter, a 10 ps pulse has an average pulse intensity of $10^{10}\ W/cm^2$ – about twice the damage threshold of most optical components. In order to avoid the need for even larger beam diameters on every optic (which are more susceptible to distortions, decrease gain in the laser crystals, and require larger, more expensive optics), we need a significantly longer pulse. Specifically, we will aim for intensity just below $5 \times 10^9\ W/cm^2$ for a 1.5 mm diameter beam, corresponding to a final pulse duration of about 100 ps.

There are many ways to stretch pulses, but we chose to use a Chirped Volume Bragg Grating (CVBG) stretcher as the most effective and practical option. In this section, we will describe the stretcher options considered and the characteristics of our laser that ultimately led to the choice of the CVBG as our best available option.

Possibility 1: Grating Stretcher A grating stretcher was the first option we explored due our familiarity with these stretchers in Ti:sapphire CPA systems. Another attractive aspect of this scheme is that it is designed around diffraction gratings which are easily available, inexpensive, and highly efficient in this spectral region. The Achilles heel of this design, however, lies in the dramatic difference in bandwidth of cryogenically cooled Yb:YAG compared to Ti:sapphire. The former is a fraction of one nanometer whereas Ti:sapphire crystals support hundreds of nanometers.

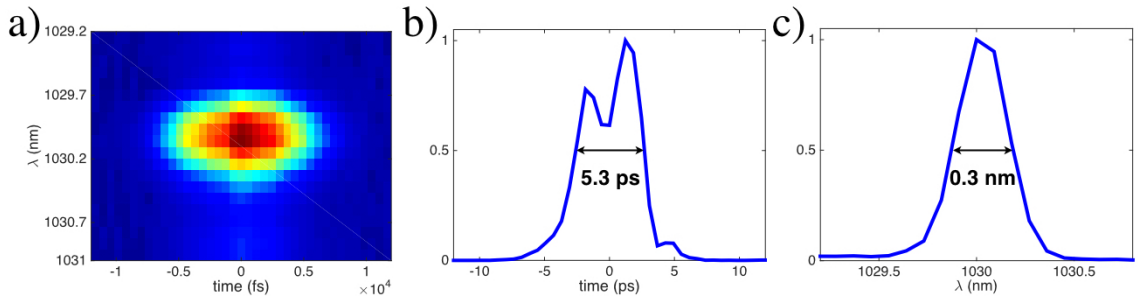

Figure 4.5: a) FROG scan of the output of a cryogenically cooled Yb:YAG laser amplifier with an unstretched seed. b) Retrieved temporal profile, showing 5.3 ps FWHM pulse duration. c) Spectrum, showing 0.3 nm FWHM bandwidth.

To measure the pulses emerging from the first stage Yb:YAG amplifier without a stretcher, we performed a Frequency Resolved Optical Gating (FROG) measurement. Figure 4.5 shows the FROG scan and the resulting pulse duration and spectrum information. Here we see that the spectral bandwidth is 0.3 nm. For a grating stretcher, this means that we need both high groove density gratings to give some angular spread to this narrow spectral region and we need to let that light propagate a long way between gratings so it will result in the necessary temporal stretch factor.

More concretely, we can consider a pair of high-groove-density gratings with 1740 grooves per

millimeter and 64 degree angle of incidence which are set up in a two-grating, double pass stretcher configuration (as illustrated in Figure 4.6). Starting from 10 ps, we need a grating separation of about 125 cm to reach 100 ps. Although not outside the realm of possibility, this large grating separation either requires a large footprint or additional fold mirrors. In either case, the design is unwieldy and sensitive to misalignment even with these exceptionally dispersive diffraction gratings.

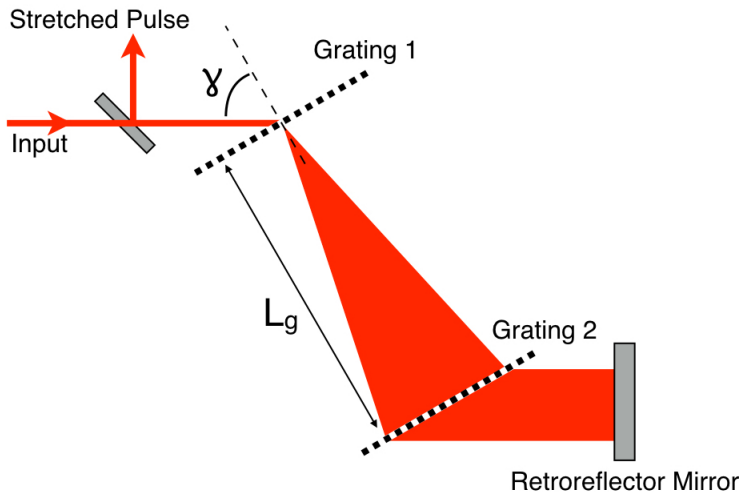

Figure 4.6: Illustration of a diffraction grating stretcher with transmission gratings. L_g is the perpendicular distance between the gratings and γ is the angle of incidence with respect to the grating normal.

Option 2: Fiber Stretcher Another possibility is to use the dispersion experienced in a pulse traveling through an optical fiber to stretch the pulse. A major advantage of this design is that even large lengths of fiber can be coiled to make a compact stretcher that is completely insensitive to misalignment except where the light is coupled into the fiber. Also, the beam emerging from a single mode fiber stretcher will be a perfect gaussian mode due to the guiding.

Unfortunately, this design also has some major drawbacks. One is that it limits the amount of energy that can be in the stretched pulse because an unstretched pulse must propagate in the fiber. Due to the small mode diameters in single mode fiber (typically about 5 μm), even tens of milliwatts at 5 MHz can cause nonlinear distortions to the spectrum and add spectral phase due to self phase modulation. Also, to stretch a 0.3 nm bandwidth that starts at 10 ps to 100 ps requires

a long section of fiber – about six kilometers, to be specific. At this length, the beam is extremely sensitive to birefringence in the fiber. In a very long section of fiber, polarization rotations can arise from imperceptibly small shifts in the fiber's curvature, leading to an unstable polarization at the output. Fibers are also not perfectly lossless, and at the several-kilometer length scale, the loss can be significant. Corning's HI980 single mode fiber, for example, specifies loss < 2.5 dB/km, so at 6 km, this could be up to 15 dB loss (a factor of 30 decrease).

In summary, in spite of the elegant simplicity of the fiber stretcher idea and the advantages of compactness and excellent mode quality, this design compromises pulse energy and stability to a high degree and is therefore not the best choice for our design.

Solution: Volume Bragg Grating Stretcher Our final stretcher design combines the compactness of the fiber stretcher with the power handling capability of the diffraction grating stretcher while also delivering sufficient dispersion to reach 100 ps. This design is based on an optical component called a Chirped Volume Bragg Grating (CVBG).

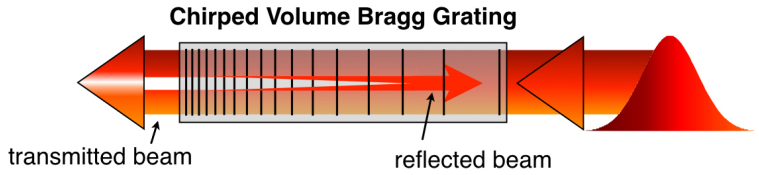

Figure 4.7: Illustration of a Chirped Volume Bragg Grating (CVBG). A narrow spectral component of the incident light is reflected and emerges temporally stretched.

An illustration of the CVBG is shown in Figure 4.7. It is a 34 mm-long, patterned block of glass that is transparent to most of the incident light from the fiber amplifier. Its key features are precisely spaced, internal reflective surfaces which result in the Bragg reflection of different spectral components of the incident light at different propagation distances within the material. Our volume Bragg grating is designed to reflect a 0.66 nm band around 1029.4 nm with dispersion of 520 picoseconds per nanometer of bandwidth. The long-wavelength side of this spectral range is reflected at the front of the CVBG and the short-wavelength side, at the back so that the pulse emerges positively chirped with 340 ps pulse duration. The spectral bandwidth of the pulse will be

further narrowed in the Yb:YAG amplifier, resulting in about 120 ps pulse duration.

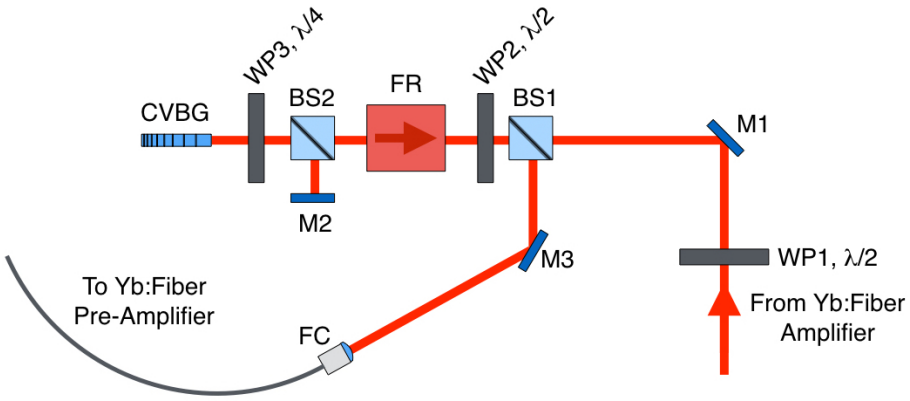

Figure 4.8: Stretcher for the 1 μm wavelength light that seeds the OPCPA pump laser. The design is based on a double reflection from a Volume Bragg Grating. The beam is shown in red. Waveplates are marked WP, mirrors, M, polarization beamsplitters, BS, Faraday rotator, FR, chirped volume Bragg grating, CVBG, and fiber collimator, FC.

We set up our stretcher using two reflections from the CVBG and coupling the output into a single mode fiber which will deliver it to the amplifier. The design is illustrated in Figure 4.8. Note that since the pulse is stretched to hundreds of picoseconds before being coupled into the fiber, we do not have to worry about limiting the pulse energy into the fiber to avoid self phase modulation. We have some of the advantages of a fiber stretcher (clean output mode and robustness against misalignment) without severely limiting the seed energy delivered to the Yb:YAG amplifier chain. The CVBG is a perfect tool for the narrow bandwidth we are working with and the large stretch factor we need, resulting in a compact and robust 1 μm stretcher.

To measure the pulse duration resulting from this stretcher in combination with the bandwidth-narrowing effect of the cryogenically cooled Yb:YAG amplifiers, we performed an autocorrelation measurement on the output of the amplifier chain. We used a home-built autocorrelator consisting of two beam paths, one with variable delay, which were overlapped at an angle in a BBO crystal. When the pulses are overlapped in time and space, they produce a second harmonic signal through type-I phase matching in the BBO. We recorded this signal using a photodiode. The strength of the signal as a function of delay corresponds to a convolution of the pump pulse shape with itself.

The results of this measurement for a single pass and a double pass through the CVBG are shown in Figure 4.9.

From the autocorrelation, we can infer that the amplified pulse duration using a single pass through the CVBG stretcher is about 40 ps. Also, two passes through the CVBG result in a pulse duration of about 120 ps. We therefore chose the two-pass configuration for our design.

4.2.6 Regenerative Amplifier

From the CVBG stretcher, the 1 μm light is amplified in four cryogenically cooled Yb:YAG amplifiers. The first of these is a regenerative amplifier designed for the highest gain of the three stages. It must cut down the repetition rate from 5 MHz to the desired operating level (between 1 and 100 kHz) to mitigate heating issues associated with high average power. It must also give good pulse contrast and preserve excellent mode quality. Regenerative amplifiers are well suited to these requirements because they allow the pulses to take many trips through the gain material, employ a Pockels cell to control pulse timing and repetition rate of the amplification, and their design is based on a resonant cavity which will enforce the cavity mode.

A diagram of the regenerative amplifier in this laser system is shown in Figure 4.10. For operation at 1 kHz, 5 W of continuous-wave, 940 nm, unpolarized pump light is supplied by a fiber-coupled laser diode. The pump light is focused onto the crystal with a diameter of 250 μm. The crystal is a 5 mm long, 8 mm diameter cylinder that is mounted to a 60 Watt capacity cold head in a vacuum chamber. While operating, the crystal mount temperature is typically about -220°C.

The timing of this amplifier is controlled by the Pockels cell labeled PC1 in Figure 4.10. It uses fast, high voltage switching to rotate the polarization of light passing through an internal crystal by a quarter wave. In conjunction with the surrounding polarization optics (BS1, BS2, WP4, WP5, and FR) this forms a gate that either closes the cavity so that whatever pulse(s) are inside stay there and no new ones get in, or opens it so that pulses are free to come and go. When the Pockels cell is off, the gate is open, and when on, it is closed. The timing is adjusted so that

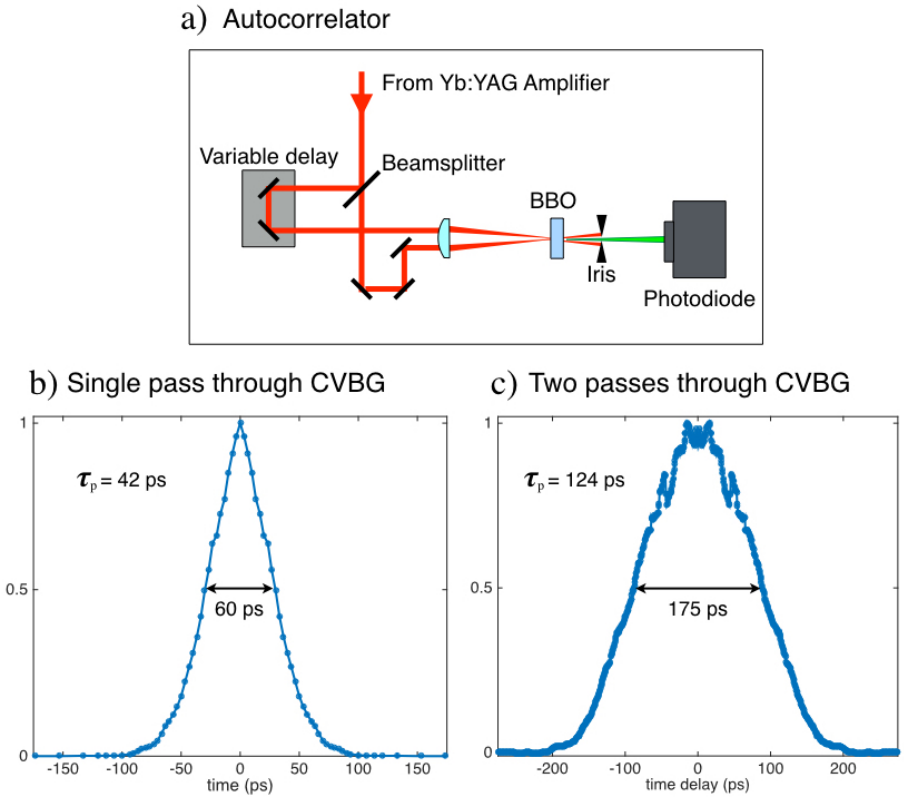

Figure 4.9: a) Diagram of the autocorrelator. The 1 μm beam is shown in red and its second harmonic is shown in green. b) Autocorrelation measurement of the output of the Yb:YAG amplifier chain when the pulse is stretched using one pass through the CVBG. Assuming a gaussian temporal beam profile, the width of the pulse corresponding to this autocorrelation is 42 ps. a) Autocorrelation of the Yb:YAG amplifier output when stretched using two passes through the CVBG. Assuming a gaussian temporal pulse shape, the pulse duration is 124 ps.

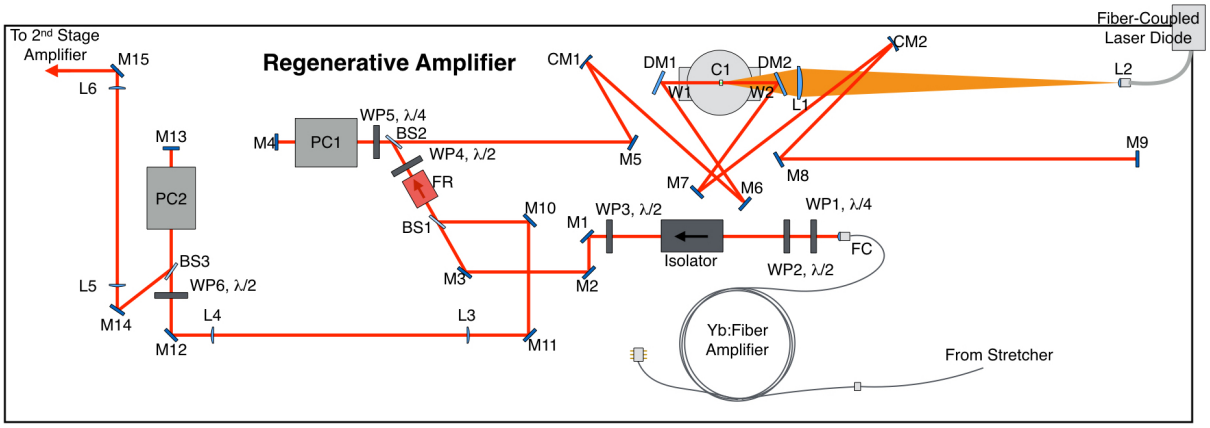

Figure 4.10: Drawing of the first stage of the cryogenically cooled Yb:YAG amplifier chain: a regenerative amplifier. The 1 μm beam is shown in red and the 940 nm pump beam, in orange. Waveplates are marked WP, mirrors, M, dichroic mirrors, DM, curved mirrors, CM, lenses, L, windows, W, crystal, C, polarization beamsplitters, BS, Faraday rotator, FR, fiber collimator, FC, and Pockels cells, PC.

one pulse will be trapped in the cavity at a time and it will pass through the crystal enough times to give it the desired gain. One pulse will be let in every period of the desired repetition rate. So, for example, at 1 kHz repetition rate the Pockels cell will cycle on and off once every millisecond with the duration of the "on" period determining the number of round trips that pulse takes in the cavity before being let out.

The maximum energy that can be imparted to these pulses is primarily limited by two constraints: the damage threshold of the optical components and by an instability known as bifurcation. To prevent optical damage, we keep the energy per pulse at or below 0.5 mJ. This level is determined by the smallest beam size on an optical component in the cavity which is found in the laser's gain medium. Increasing the pump spot size would raise the damage-limited maximum output energy of this laser, but would also lower the gain. Instead, we chose to work with the 0.5 mJ maximum. Once this pulse energy is reached, we can know not to increase the length of the "on" period of the Pockels cell any further.

The bifurcation instability is less straightforward. Bifurcation occurs when consecutive seed pulses extract different amounts of energy from the crystal and is common in continuously pumped regenerative amplifiers run at high repetition rates. Also, Yb:YAG has a relatively long upper state lifetime of about 1 ms. [83] This can make the laser more prone to bifurcation because if the first of two seed pulses does not extract all of the available energy in the crystal, the population inversion will be even higher for the second pulse and it will be amplified more than the first. The output of a laser experiencing bifurcation instability will often appear to be either alternating between two output energies or exhibit chaotic pulse energy output. [84, 85]

We have three knobs that can help suppress bifurcation: the seed energy, the pump power, and the repetition rate set by the Pockels cell. For a given pump power, higher seed energies can help to suppress bifurcation because a pulse that starts with more energy can extract more. [85] For this reason, we added a Yb:fiber preamplifier to boost the seed energy from roughly 1 nJ emerging from the CVBG stretcher to 200 nJ. At 1 kHz, this amount of seed suppresses birefringence as the pump power is turned up to the level required for 0.5 mJ laser output. Lower repetition rates

also tend to be more stable. [84] Above 10 kHz, bifurcation begins to be problematic at the 0.5 mJ level. For the work presented in this thesis, we chose to maximize pulse energy, beginning the development of this laser at 1 kHz where bifurcation is not an issue with the seed energy we have and with 0.5 mJ output pulses.

The output of the amplifier cavity consists of a train of amplified pulses at the repetition rate chosen with the Pockels cell, superimposed on a 5 MHz un-amplified pulse train. We use a second Pockels cell to reject the un-amplified pulses. This is labeled PC2 in Figure 4.10.

The regenerative amplifier's output beam is finally expanded and recollimated on its way to the second stage Yb:YAG amplifier. The beam profile is shown in Figure 4.11. It displays some astigmatism from the off-axis bounces from the two curved mirrors in the regenerative amplifier cavity (CM1 and CM2 in Figure 4.10). The ratio of the Full Width at Half Maximum (FWHM) of the beam shape in the horizontal direction to the FWHM in the vertical direction is 1.3. Other than this slight departure from roundness, the mode closely resembles a gaussian beam, without additional structure.

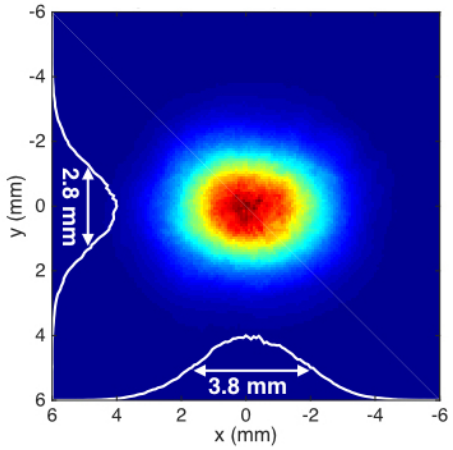

Figure 4.11: Spatial beam profile of the collimated regenerative amplifier output mode when this amplifier is operating at full power at 1 kHz repetition rate. The white curves are lineouts of the beam at its peak in the x- and y-directions.

In summary, the first stage of the Yb:YAG amplifier chain is a regenerative amplifier running at 1 kHz repetition rate, 0.5 mJ pulse energy, and with a nearly Gaussian mode. The total gain of

the regenerative amplifier is a factor of 2500 (34 dB) and that factor increases to 5×10^5 (57 dB) when one includes the fiber pre-amplifier.

4.2.7 Multi-Pass Amplifiers

The output of the regenerative amplifier is fed into three further stages of amplification. All three are collinear multi-pass amplifiers, each designed to handle progressively more pulse energy than the previous one at the expense of gain. The collinear geometry in each helps to maintain good mode quality. To achieve collinear operation, we use waveplates, Faraday rotators, and polarizing beamsplitters to separate the amplified beam from the seed after the desired number of passes.

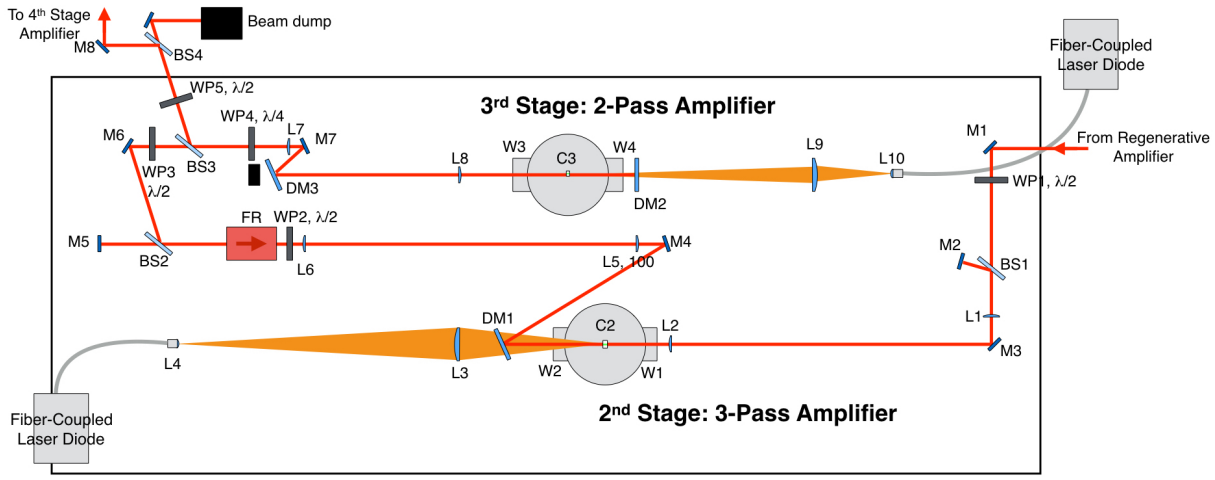

Figure 4.12: Drawing of the second and third stage Yb:YAG amplifiers. The input from the regenerative amplifier is shown on the far right and the output of the third stage is at the top left. The 1 μm beam is shown in red and the 940 nm pump beam, in orange. Waveplates are marked WP, mirrors, M, dichroic mirrors, DM, lenses, L, windows, W, crystals, C, polarization beamsplitters, BS, and Faraday rotator, FR.

Figure 4.12 illustrates the layout of the second and third stage amplifiers and Figure 4.13 shows the fourth stage. The second stage amplifier uses three passes through its crystal whereas the third and fourth use two passes. In all cases, the beam is collimated as it passes through the crystal and when it enters and exits each amplification stage. The beam diameter through each crystal is set by the gain required as well as the need to keep the peak intensity below the damage

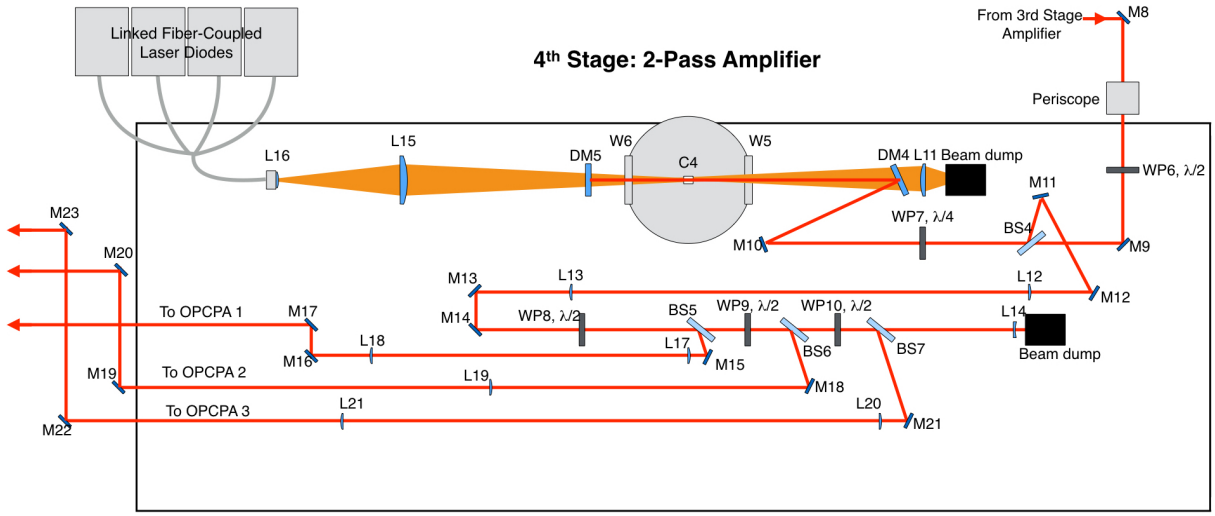

Figure 4.13: Drawing of the fourth (and final) stage Yb:YAG amplifier. The input from the 3^{rd} stage amplifier is shown on the top right. The output is split into three beams in preparation to pump three stages of OPCPA. These beams are shown exiting on the left. The 1 μm beam is depicted in red and the 940 nm pump beam, in orange. Waveplates are marked WP, mirrors, M, dichroic mirrors, DM, lenses, L, windows, W, crystal, C, and polarization beamsplitters, BS.

threshold of the optics.

The number of passes we use in each stage is set to maximize the amount of energy extracted from the crystal while keeping beam distortions minimal. Most of the beam distortions in this laser system come from nonuniform birefringence in the laser crystals induced by pressure from their mounts at cryogenic temperatures. The beam is much more sensitive to these distortions when it is larger, allowing many passes in the regenerative amplifier where the beam diameter is only 250 μm, three passes in the second stage where it is just under 1 mm in diameter, and only two passes in the third and fourth stages where the beam diameter is 2.2 mm and 4 mm, respectively.

The pump mode diameters were optimized with the goal of optimizing output pulse energy. We iteratively decreased the pump mode and seed beam sizes to maximize the amplified pulse energy while keeping the intensity near or below the damage threshold. We increased the number of passes until cumulative beam distortions from birefringence became noticeable. The pump power is generally limited by the power-handling capacity of the cold heads – above a certain pump power, the crystal starts to heat up, causing the gain to decrease.

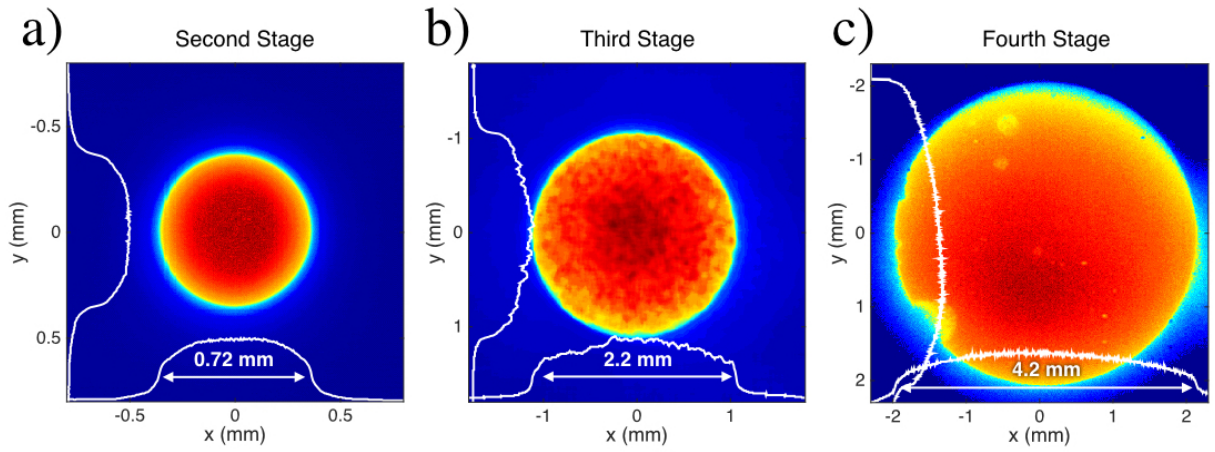

Figure 4.14: Beam profiles of the 940 nm pump at the crystal plane in (a), the second stage, (b), the third stage, and (c), the fourth stage Yb:YAG amplifier. The white curves are lineouts of the beams at their peaks in the x- and y-directions.

The pump profiles as they are on the second, third, and fourth stage crystals are shown in Figure 4.14. In each stage, the pump lasers are fiber-coupled laser diode arrays. The output of

each fiber is imaged onto the corresponding crystal with the desired magnification, resulting in a flat-top pump beam profile. The flat-top pump helps to ensure that the seed is amplified uniformly in the spatial domain.

The second stage is pumped with 24 W using a 0.72 mm diameter spot size. The crystal is mounted to a 90 W capacity heat exchanger. The beam makes three passes through the crystal, using a Faraday rotator, half waveplate, and two polarizing beam splitters to couple the beam out after three passes. The final energy is 5 mJ, limited by the damage threshold of the optics. To minimize risk of damage while maintaining the designed gain, the beam has a 1 mm diameter as it passes through the crystal but is up-collimated with a telescope to about 4 mm diameter elsewhere.

The third stage uses two passes through its crystal which is pumped by 80 W in a 2.2 mm-diameter spot and mounted to a 90 W capacity heat exchanger. As in the second stage, the third stage's seed beam diameter is decreased to the desired size (2.2 mm to match the pump spot) as it passes through the crystal and up-collimated elsewhere. The beam makes two passes through the crystal, coupled out using a quarter waveplate and polarizing beamsplitter. Each time the beam passes through the crystal it experiences a thermal lens when the pump is at full power. We compensate for the thermal lens by adding divergence to the beam as it passes through the crystal. The telescope comprising lenses L7 and L8 in Figure 4.12 is set up to add enough divergence so that the output of the amplifier emerges roughly collimated under full-power operating conditions. The final output energy is about 14 mJ, limited by thermal loading of the crystal.

Like the third, the fourth stage is also a two-pass amplifier. It is illustrated in Figure 4.13. This stage uses a 4.2 mm-diameter pump beam, and therefore can handle pulse energies up to 63 mJ in a matching 4 mm-diameter beam. It is pumped by 187 W. Like the third stage, this amplifier uses a quarter waveplate and a polarizing beam splitter to couple the beam out of the amplifier after two passes. It is designed to have minimal transmissive optics for simplicity and to mitigate the effects of self focusing in reshaping the beam. Therefore, we do not use a telescope to compensate for the thermal lens until after the output of the laser. The thermal lens is significant, but the amplified beam remains large enough to avoid damage at the output. The specifics of the

thermal lens will be discussed further in the following section. The final output energy is typically 27 mJ, and up to 36 mJ. As in the third stage, this is limited by thermal loading of the crystal. Increasing the number of passes to four, we have demonstrated amplified pulse energy up to 50 mJ, but this comes at the expense of mode quality: cumulative losses due to birefringence in the crystal add significant distortions and structure to the mode.

Amplifier Stage	P_{out} (W)	P_{pump} (W, 1 kHz)	Gain	Pump Mode Diameter (mm)	I_{max} (W/cm^2)
1	0.5	5	2500	0.25	1×10^{10}
2	5	24	10	0.72	6.4×10^{9}
3	14	80	2.8	2.2	3.7×10^{9}
4	27 (up to 36)	187	1.9	4.2	2.9×10^{9}

Table 4.1: Summary of the parameters of the four cryogenically cooled Yb:YAG amplifier stages. P_{out} is the output power of the indicated stage, P_{pump} is the pump power, "Gain" is the ratio of the output power to the input power, "Pump Mode Diameter" is the diameter of the pump spot, and I_{max} is the maximum peak intensity reached by the amplified beam in each stage.

Table 4.1 gives a summary of the main parameters of each amplification stage and Figure 4.15 shows the beam profiles at the output of each stage. Note that the second stage mode is almost perfectly round, getting rid of the slight distortion from the curved mirrors in the regenerative amplifier cavity. This illustrates the effectiveness of the collinear geometry in maintaining excellent mode quality. The third stage's output is also quite good with a nearly gaussian beam and horizontal-to-vertical beam diameter ratio of 1.14. The fourth stage mode is also smooth, though it has slightly more energy in the wings of the pulse and worse ellipticity (horizontal-to-vertical width ratio of 1.29). The ellipticity is likely from a nonuniform thermal lens – the heat is extracted from the top of the crystal's mount, leading to a nonuniform thermal gradient across the crystal. In spite of this ellipticity, however, the fourth stage output is still focusable to a smooth, round spot for pumping the OPCPA and supplies more than sufficient pulse energy.

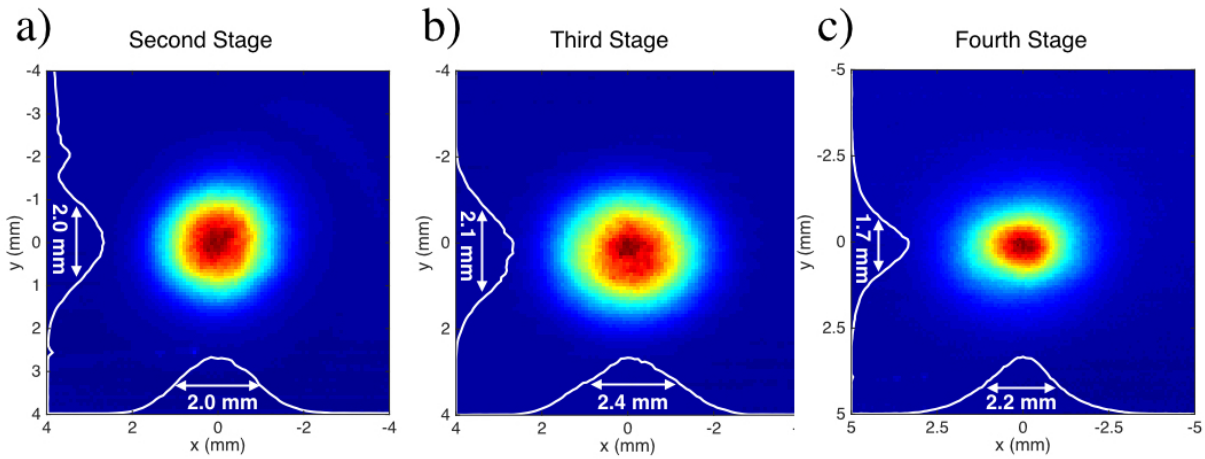

Figure 4.15: Far-field spatial beam profile of the output of (a), the second, (b), the third, and (c), the fourth stage Yb:YAG amplifier. The white curves are lineouts of the beams at their peaks in the x- and y-directions.

4.3 Pump Laser Design Considerations

In this section, we will describe some of the design challenges of building a pump laser with both high pulse energy and kilohertz repetition rate. The pump laser for an OPCPA system in this regime is the most difficult component because it must routinely and stably handle the most heat and pulse energy of any part of the system. It also must be simple enough to enable consistent day-to-day and all-day operation so the OPCPA idler will also be sufficiently stable for experiments – especially ones requiring long integration times. In this section, we will describe the constraints imposed by the need for good mode quality, avoiding optical damage, thermal lensing, thermal loading, crystal birefringence, nonlinear effects, and other instabilities. We will discuss how we addressed each and mention ways that this pump laser can be further optimized in the future.

4.3.1 Optical Damage

One of the most important design considerations for a stable, high pulse energy laser is avoiding optical damage. Since the optics in this laser are all very close to non-absorbing at 1.03 μm, the damage we need to worry about has to do with peak power rather than average power. Peak power depends on the pulse energy, the temporal duration of the pulse, and mode area:

$$I_{peak} = \frac{E_{pulse}}{\tau_{pulse} A} \tag{4.4}$$

Here E_{pulse} is the pulse energy, τ_{pulse} is the pulse duration, and A is the mode area assuming a beam that is both temporally and spatially flat-top.

To lower the optical intensity and avoid damage, the beam area or pulse duration must be increased. Since the stretcher sets the pulse duration at about 120 ps, the main variable we have to play with is the mode area. There are two constraints that prevent us from making the mode area arbitrarily large. The first is a practical one: the beam should ideally remain small enough to fit on standard-sized optics without clipping. Smaller optics are both less expensive than larger ones

and typically introduce fewer distortions due to non-uniformities on their surfaces. The second constraint is a tradeoff between mode area and the efficiency with which energy can be extracted from the laser crystal per pass.

The amount of energy a seed pulse will be able to extract from a given laser crystal is characterized by the saturation energy, E_{sat}, of the material. The saturation energy is defined as the amount of seed energy required to reduce the gain to 1/e of its starting value. It is related to an intrinsic property of the material, the saturation fluence (F_{sat}), multiplied by the mode area of the pulse [86]:

$$F_{sat}(\nu) = \frac{h\nu}{\sigma_{em}(\nu) + \sigma_{abs}(\nu)} \tag{4.5}$$

$$E_{sat}(\nu) = AF_{sat}(\nu) \tag{4.6}$$

Where ν is the laser frequency, σ_{em} is the emission cross section, σ_{abs} is the absorption cross section, and A is the mode area. For cryogenically cooled Yb:YAG at ν = c/1030 nm, the emission cross section is approximately $11 \times 10^{-20}\ cm^2$ and reabsorption of the laser light is small, so the saturation fluence is approximately 1.8 J/cm^2. [76, 83] The saturation energy is therefore 1.8 J/cm^2 times the area of the beam:

$$E_{sat,\ Yb:YAG} = A \times (1.8\ J/cm^2) \tag{4.7}$$

Based on experience with the optics in our laser system, we also know that damage typically occurs at intensities near $10^{10}\ W/cm^2$. We therefore use $I_{max} = 5 \times 10^9\ W/cm^2$ as the maximum safe intensity, incorporating a safety factor of two. The damage-limited maximum pulse energy for a 100 ps-long pulse is therefore:

$$E_{max,\ Yb:YAG} = A \times (0.5\ J/cm^2) \tag{4.8}$$

Comparing Equations 4.7 and 4.8, we see that, with a 100 ps pulse, we can never quite reach the saturation energy of cryogenically cooled Yb:YAG without first running into the damage threshold. This means that we should maximize gain by making the beam as small as possible before reaching 5×10^9 W/cm^2 and we should use multiple passes through the crystal to get the best energy extraction possible.

4.3.2 Mode Quality

In the previous section, we noted that efficient energy extraction requires multiple passes through each Yb:YAG crystal. However, we need to balance this consideration with the need for a focusable, spatially smooth beam to pump the OPCPA. The need for excellent mode quality was a driving factor in our choice to use collinear rod geometries in all four stages of our Yb:YAG amplifier. This need also limits the number of passes used in each multi-pass amplification stage. Alternative amplifier geometries were explored, each with advantages. In this section, we will describe a "bow tie" and a thin-disk geometry, explaining our ultimate design choice.

Bow Tie Amplifier A bow tie amplifier design differs from a collinear amplifier in that passes through the laser crystal are angularly separated so that the output can be picked off with a mirror rather than relying on polarization-manipulating optics to couple out the amplified beam. One advantage of the bow tie design over a collinear geometry is that it is not inherently sensitive to polarization and therefore can avoid beam distortions due to birefringence in the crystal. The intensity profile of the output beam will be undistorted by even strong birefringence in the crystal as long as the beam does not pass through a polarizer. Using a collinear geometry, the number of passes in the third and fourth stages are limited to two due to birefringence, so extraction could theoretically be improved by the bow-tie geometry.

However, in spite of its insensitivity to birefringence, the amplified beam in a bow tie arrangement gets distorted when it passes through the gain medium at an angle. In practice, we found that these distortions compromised beam quality in just a few passes. This geometry therefore loses its advantage over the collinear geometry. In addition, it requires more challenging alignment

and more beam path length which can decrease the robustness and ease of operation of the laser. For these reasons, we chose not to use the bow tie geometry in any amplification stage.

Thin-Disk Amplifier Another option is to use a wide, thin crystal instead of the relatively thick rods that we currently employ. The main advantage of using a thin-disk crystal rather than a rod is that one side of the crystal can be given a reflective coating and mounted directly to a heat sink. The resulting crystal mount has much better thermal contact to the heat sink than our rod crystals, mitigating thermal loading. [87] Indeed, this geometry has been used successfully to build amplifiers with extremely high average power – up to 10 kW in continuous wave operation. [88] Thin-disk lasers have also been demonstrated to exhibit high gain and final output power. Metzger et al., for example, report amplification up to 25 mJ at 3 kHz in a regenerative amplifier that also provides excellent mode quality (M^2 <1.1). [89] More recently, Klingebiel et al. reported a thin disk laser operating with 220 mJ at 1 kHz. [90]

The drawbacks of the thin-disk laser scheme are mostly practical. First, the fact that the laser crystal is thin means that the gain per pass is low and many round trips are required. Metzger et al., for example, used 150 round trips in the cavity and therefore had to work in an intermittent stable regime to avoid bifurcation. [89] Though this scheme was successful, we preferred a rod-type laser with lower gain, but a more easily identifiable regime of stable operation and simpler alignment. Second, a thin-disk laser requires a different crystal mounting scheme than the rod geometry. We preferred the rod geometry so that we could capitalize on our previous experience in cryogenically cooling laser rods. For these reasons, we chose to use rod amplifiers are the most straight-forward path toward a simple and robust OPCPA pump laser. However, the recent result by Klingebiel et al. shows potential for a single stage thin disk laser that could replace all four stages of our current amplifier and surpass its performance in terms of pulse energy. This makes a strong case for pursuing this design in the future.

4.3.3 Thermal Lensing

As mentioned in previous sections, thermal loading limits the amount of pump power we are able to use in the third and fourth stage Yb:YAG amplifiers. In this section, we will concentrate on one consequence of thermal loading: thermal lensing.

When a laser crystal is subjected to a thermal load from the pump laser, there is a temperature gradient across the face of the crystal. In many materials, including Yb:YAG, this temperature gradient also creates a gradient in the index of refraction which looks like a lens to an incident laser beam. This effect is known as thermal lensing. To better understand this effect, we can model the temperature distribution a simple situation similar to what we have in our Yb:YAG amplifier. Specifically, we will consider a cylindrical, rod-shaped laser crystal with 12 mm diameter and 15 mm long, pumped by a 4 mm diameter flat-top pump beam incident from one end of the rod. This is illustrated in Figure 4.16

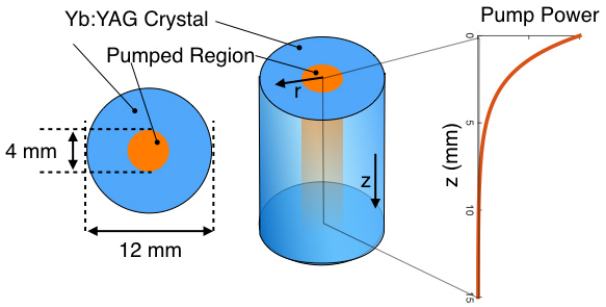

Figure 4.16: Cylindrical laser rod, end pumped by a flat-top pump beam.

The flat-top beam is realistic for a fiber-coupled diode laser pump. Its power distribution in the laser crystal will decay exponentially along the length of the crystal with a rate of decay that depends on the doping density of the crystal. For the sake of comparison, we will take the doping density to be that of the fourth stage laser crystal, 3.2 percent. The normalized power distribution in the laser crystal is therefore given by [91]:

$$p(r,z) = \frac{\alpha e^{-\alpha z}}{\pi r_p^2 \left(1 - e^{-\alpha \ell}\right)} \Theta(r_p^2 - r^2) \tag{4.9}$$

Here α is the absorption coefficient at the pump wavelength, r_p is the radius of the pumped region of the crystal, ℓ is the length of the crystal, and Θ is the Heaviside step function. Also, z is the direction along the length of the crystal and r is the radial coordinate. The fraction of the absorbed pump power that becomes heat is χ_{QD} which is given in terms of the quantum defect in Equation 4.3. Given this heat source, we can find out the steady state temperature distribution in the crystal using the heat conduction equation for isotropic media [92]:

$$\nabla^2 T = -\frac{Q_o}{\kappa} \tag{4.10}$$

Where Q_o is the heat density in the material and κ is the material's thermal conductivity. For a given z-position, if Q_o is constant and nonzero, the solution to Equation 4.22 is quadratic in the radius, r. If Q_o is equal to zero the temperature distribution decays logarithmically. Therefore, we expect that the temperature distribution in our model laser rod will have a parabolic character inside the pumped region and a decaying logarithmic character outside. Assuming that the thermal conductivity is constant with temperature, this temperature distribution can be calculated analytically and is given by Equation 4.11 [91]:

$$T(r,z) = T_o + \frac{\chi_{QD} P_{abs}}{4\pi\kappa} \frac{\alpha e^{-\alpha z}}{1-e^{-\alpha\ell}} \left\{ \left[\left(1 - \frac{r^2}{r_p^2}\right) + \ln\left(\frac{r_c^2}{r_p^2}\right) \right] \Theta(r_p^2 - r^2) + \ln\left(\frac{r_c^2}{r^2}\right) \Theta(r^2 - r_p^2) \right\} \tag{4.11}$$

Here, P_{abs} is the total absorbed pump power, T_o is the temperature at the edge of the crystal, and r_c is the radius of the crystal. Figure 4.17 shows plots of the pump power distributions (Equation 4.9) and of this temperature distribution in the crystal (Equation 4.11). For this calculation, we used parameters matching the fourth stage Yb:YAG amplifier, given in Table 4.2.

This temperature distribution creates a thermal lens through the thermootic effect and through thermal expansion. The thermooptic effect is a change in index of refraction in response to a temperature gradient. It is characterized by the thermooptic coefficient, β:

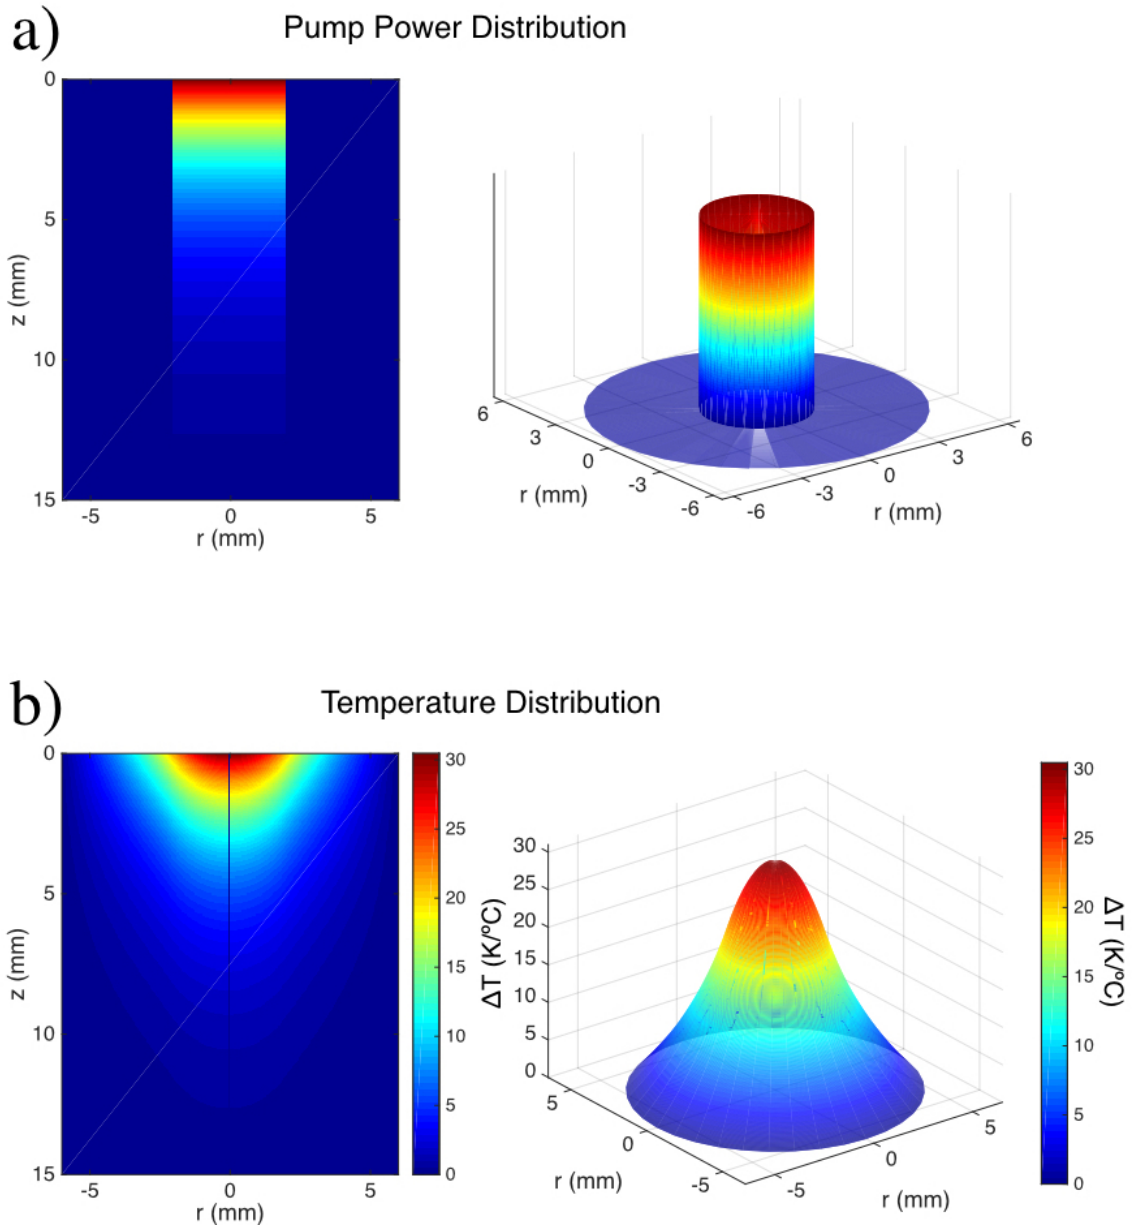

Figure 4.17: a) Pump power distribution in a 12 mm diameter, 15 mm long crystal with a 4 mm diameter flat-top pump. The graph on the left shows the pump power distribution in a slice through the center of the crystal and the graph to its right shows the pump distribution on the pumped face of the crystal, sometimes called the top hat distribution. b) Temperature distributions in the same crystal. The temperature is given relative to the temperature at the edge of the crystal. The graph on the left shows ΔT in a slice through the center of the crystal, and the one on the far right shows ΔT on the pumped surface of the crystal.

Parameter	Value	Reference
r_c, crystal radius	6 mm	
ℓ, crystal length	15 mm	
r_p, pump radius	2 mm	
P_{abs}, absorbed power	187 W	
κ, thermal conductivity at 80 K	45 W/mK	[93]
χ_{QD}, fractional thermal loading	0.087	[75]
α, absorption coefficient	328.8 m^{-1}	[75]
β, dn/dT, thermooptic coefficient at 80 K	0.6 ppm/K	[93]
ξ, coefficient of thermal expansion at 80 K	1.3 ppm/K	[93]
n_o, unperturbed index or refraction of YAG	1.83	

Table 4.2: Parameters of the crystal used to model the temperature distribution. The crystal dimensions, doping density, and pump power match that of the fourth stage amplifier.

$$\beta = \frac{dn}{dT} \tag{4.12}$$

Thermal expansion contributes to the thermal lens by leading to bulging in the material where it is hottest and by changing the index of refraction due to to stress and strain from this expansion. The contribution of stress and strain is typically small enough to ignore to a good approximation. [76] Putting this all together, we can write the index of refraction as follows:

$$n(r) = n_o + \Delta n_\beta(r) + \Delta n_{r,\phi}(r) \tag{4.13}$$

$$\approx n_o + \Delta n_\beta(r) \tag{4.14}$$

Here, n_o is the unperturbed index of refraction, $\Delta n_\beta(r)$ is the contribution due to the thermooptic effect, and $\Delta n_{r,\phi}(r)$ is the contribution due to stresses and strains on the crystal. The thermal lens itself comes from the resulting radially-dependent change in optical path length. Equation 4.15 gives the corresponding phase accumulated by incident light at the laser wavelength that passes through a thermal lens [76]:

$$\Delta\phi(r) = \frac{2\pi}{\lambda_L}\left[\Delta n + (n_o - 1)\xi\Delta T(r)\right] \tag{4.15}$$

Here, ξ is the coefficient of thermal expansion and $\Delta T(r)$ is the relative temperature distribution. The first term is the contribution of the thermooptic effect and the second represents material bulging due to thermal expansion. We can get an idea of what phase distortion we expect in our laser by considering the simple case of a uniformly distributed heat load. This is not a bad approximation for the region of the crystal filled by the pump beam since we use a flat-top distribution. Also, this is the region that the seed beam will pass through, and therefore the region where the thermal lens strength is relevant.

In this scenario, we can assume a constant heat density, Q_o, which is related to the total absorbed pump power by dividing by the pumped volume: $Q_o = P_{abs}/(\ell\pi r_p^2)$. From this, we can solve Equation 4.22 to find [94]:

$$T(r) = T(r=0) - \frac{Q_o r^2}{4\kappa} \tag{4.16}$$

$$\Delta T = T(r) - T(r_p) = \frac{Q_o}{4\kappa}\left(r_p^2 - r^2\right) \tag{4.17}$$

Also, the change in index of refraction due to the thermooptic effect can be expressed as [76]:

$$\Delta n_\beta(r) = \beta\left(\frac{Q_o r_p}{2h} + \frac{Q_o r_p^2}{4\kappa}\right) - \frac{Q_o}{4\kappa}\beta r^2 \tag{4.18}$$

Where h is the coefficient of convective heat transfer over the interface at r_p. In our case, since the radius of the pump beam, r_p, is smaller than the radius of the crystal, there is no interface material and we will neglect the first term in the parenthesis of Equation 4.18. This gives $\Delta\phi$ a simple, quadratic form:

$$\Delta\phi(r) = \frac{2\pi}{\lambda_L}\frac{Q_o}{4\kappa}\left[\beta + \xi(n_o - 1)\right]\left(r_p^2 - r^2\right) \tag{4.19}$$

Now we can determine the expected focal length of the thermal lens by noting that a lens also gives a quadratic phase of the form:

$$\Delta\phi_{lens} = -\frac{\pi}{\lambda_L f} r^2 \tag{4.20}$$

Where f is the focal length. Using equations 4.19 and 4.20 to solve for the focal length, we find:

$$f = \frac{2\kappa}{Q_o} \frac{1}{\beta + \xi(n_o - 1)} \tag{4.21}$$

The thermal lens is plotted in Figure 4.18 c) as a function of pump power. As intuitively expected, the focal length is infinite when the incident pump power is zero. At the full power used in the fourth stage amplifier (187 W) it decreases to f = 620 mm. As a check, we also calculated the focal length based on the 3D model, Equation 4.11. This was done by summing ΔT along the z-direction, dividing it by the number of z-steps and using that to calculate the corresponding value of $\Delta\phi$. The center of $\Delta\phi$ within the pump radius, r_p, could then be fitted with a parabola to calculate the focal length. Figure 4.18, c), shows that the two methods of calculating the focal length agree precisely.

To get an idea of what thermal lens we see in practice, we used a camera to measure the size of the laser beam passing once through the fourth stage laser crystal while varying the pump power between 0 and 215 W. Figure 4.19 shows the result. The same procedure was also performed for the third stage laser crystal. As expected for a laser exhibiting a thermal lens, the beam clearly shrinks as the pump power is increased. However, the result deviates from what we expect in that the divergence does not change significantly as the pump power is increased. Also, the strength of the lens is smaller than expected from theoretical calculations. Specifically, using the change in the measured beam diameter at a known distance from the crystal, we estimated the strength of the thermal lens at 187 W to be 6.3 m. This is roughly a factor of ten larger than what we calculated for the same power.

There are a few possible reasons for the discrepancy between theoretical and observed thermal lensing behavior. The first is that our model assumes constant thermal conductivity, thermooptic

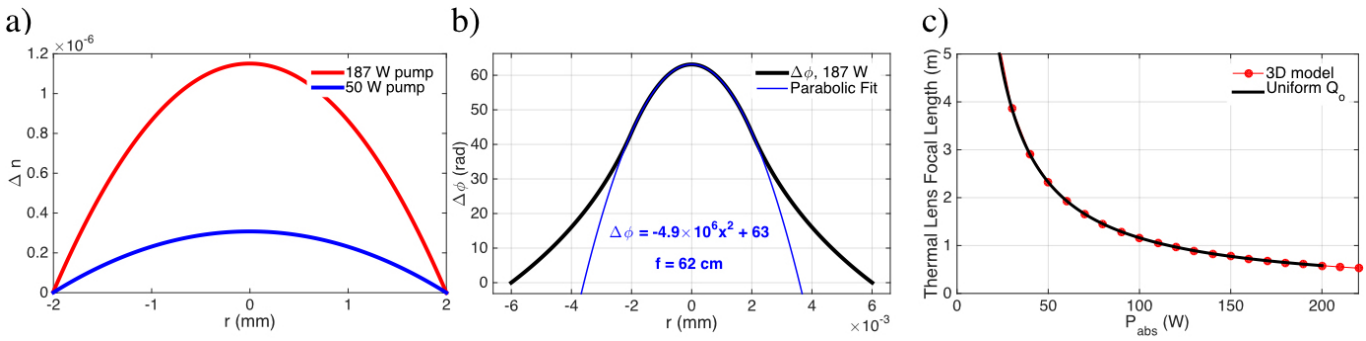

Figure 4.18: a) This plot shows the variation in index of refraction with radius in the pumped region of a cryogenically cooled Yb:YAG crystal due to thermal loading. The effect depends strongly on the amount of power dumped into the crystal. The red line shows Δn for 187 W of absorbed pump power, and the blue line shows Δn for 50 W absorbed. b) The second plot shows the relative phase change across the crystal due to the thermooptic effect and thermal expansion. This curve was calculated from the 3D temperature distribution given in Equation 4.11. c) The third plot shows the variation of the thermal lens' focal length with absorbed pump power. The black line is based on Equation 4.21 for a uniform heat density, Q_o and the red circles are calculated by fitting the phase profile derived from the 3D temperature distribution. The agreement is excellent.

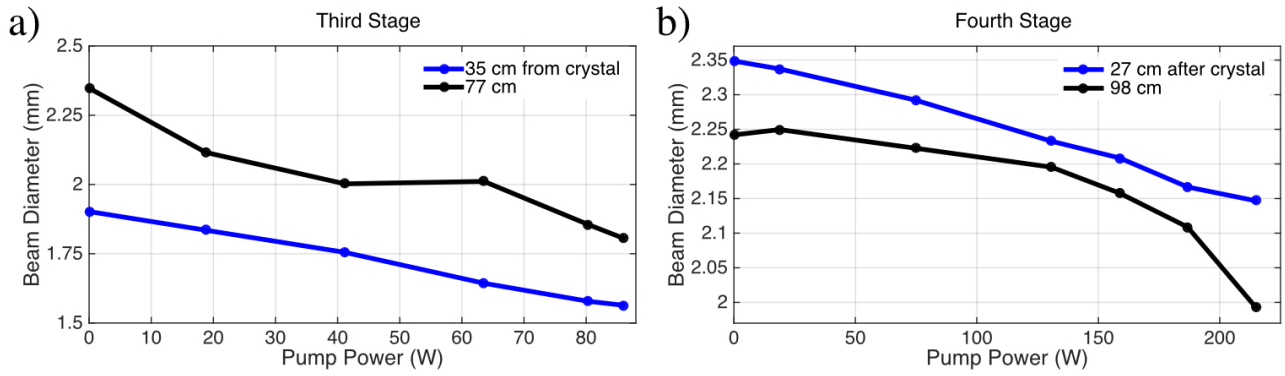

Figure 4.19: Graphs showing measurements of change in FWHM beam diameter as pump power is increased in the third (a) and fourth (b) stages. In both cases, the beam passes through the laser crystal once and is measured at two distances from the laser crystal.

coefficient, and thermal expansion coefficient over the whole surface of the laser crystal at an average temperature of 80 K. Though these assumptions greatly simplify the calculations and give some picture of what we should expect, they are not entirely valid. All three of these parameters (κ, β, and ξ) are temperature-dependent and strongly vary at temperatures below 100 K. [93, 75] A more realistic picture would show decreased thermal conductivity and increased thermooptic and thermal expansion coefficients in the center of the crystal. This would lead to an increased thermal gradient across the face of the crystal since heat would be more difficult to remove from the middle. Also, we would expect more phase distortion in the center. The cumulative effect is a more severe thermal lens than our current calculation shows. Since our experiment instead shows a much smaller thermal lens, a different effect must be dominant.

Another possible source of error is the variation in average temperature of the crystal with pump power. Our model assumes a constant average temperature of 80 K, but we know that the average temperature of the crystal rises as the pump power is increased. Exactly how this temperature varies depends on the specifics of heat transfer from the crystal to its mount, the interface with the mount, and the interface between the mount and the cold head. We do not have a way to directly measure the temperature of the laser crystal, but we can monitor the temperature of its mount, giving a lower limit to the temperature of the crystal. When the pump power is zero, we observe the mount's temperature to be -215°C (58 K) and when it is pumped by 187 W the temperature rises to -185°C (88 K). Therefore we expect that the crystal's temperature to depart significantly from 80 K for many values of P_{abs}. At 187 W pump, however, the temperature is most likely close to or above 80 K, meaning that our estimated thermal lens focal length should be close or too large compared to the actual value. Since our estimate is instead too small, we can again conclude that inaccuracy in our assumption of constant average temperature is not the main cause of this model's failure.

Finally, we must also note that thermal lensing is not the only change our beam would have experienced as it passed through the crystal. That is, the probe beam used here was the 1030 nm seed beam which sees significant gain from the pump in addition to a thermal lens. Gain can

reshape the beam, obfuscating the thermal lensing effect. This could explain both the lack of a clear change in beam divergence as the pump power was increased and the unexpectedly long focal length.

The problem of separating the effect of thermal lensing from gain could be addressed by using a probe beam with wavelength that does not overlap with the gain spectrum of the laser crystal. Unfortunately, this is technically difficult with our setup because none of the optics in the beam path are coated for such a wavelength. We made an attempt to measure the thermal lensing using a Helium-Neon laser, but could not extract useful information due to poor beam quality arising from interference of back reflections. In the future, a better measurement could be performed if the vacuum chamber's windows and dichroic mirrors were replaced with optics coated for 632 nm as well as 1030 and 940 nm.

In summary, we can theoretically estimate the thermal lens that will be experienced by the laser beam in each stage of amplification based on the power load, but this estimate is not a good predictor of the beam's behavior in terms of changes in size or divergence. For our laser in its current configuration, thermal lensing must be accounted for to avoid damage due to shrinking beams, but is a relatively small effect. It can be compensated for using telescopes that change the beam diameter, but even these are not absolutely necessary. In fact we do without such a telescope in the fourth stage to minimize the number of transmissive optics. We designed our laser amplifier based on direct observation of the beam's size and divergence changes rather than on a model of its behavior. If an accurate, predictive model is desired, however, it must account for the gain shaping in the crystal as well as thermal lensing.

4.3.4 Gain Reduction due to Thermal Loading

Thermal lensing is not the only effect of temperature on laser performance – we also see limitations on pump power due to thermal loading. More specifically, our third and fourth stage amplifiers are currently limited to a critical pump power, P_c, above which the amplified power begins to drop rather than rise as the pump power is increased.

To assess the effect of temperature on the output power of the Yb:YAG amplifiers, we measured the output power of the second stage amplifier with a fixed pump power, fixed seed power, and varied temperature. We began with the crystal at room temperature and the cooling system turned off. After turning on the cooling system, we monitored the output power and the temperature of the crystal's mount using a thermocouple. Figure 4.20 shows the results of this measurement.

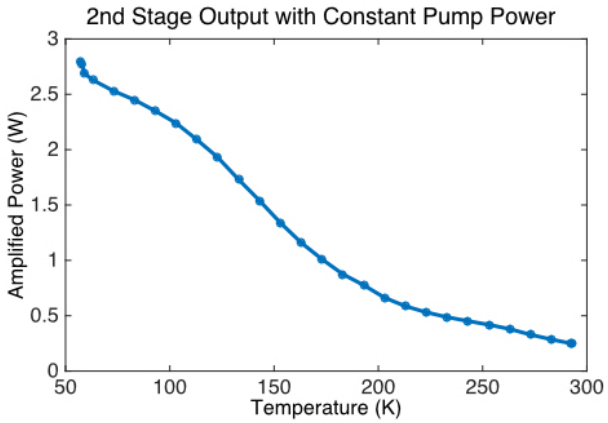

Figure 4.20: Measurement demonstrating the effect of crystal temperature on laser output power. The pump power on the crystal was chosen to be well below the thermal loading threshold of the operational cooling system. The temperature was varied from room temperature to just above 50 K. The second stage Yb:YAG amplifier's output power was measured as a function of temperature as the crystal cooled down. The temperature reported is that of the crystal's mount.

We observed a dramatic rise in output power over the whole temperature range from 300 to 50 K. This is due to Yb:YAG's increased transparency to the laser wavelength at lower temperatures as well as decreased thermal excitations into the lower laser level. One interesting and important observation is that the laser's performance continues to improve dramatically between 100 and 50 K with a sharp upswing in power as the temperature approaches 50 K. This indicates that the cooler the crystal, the better the gain we can expect for a given pump power. Also, as we increase the pump power to improve gain, we can only go so far before the crystal warms up enough to cancel out the advantage of higher pump power. This explains our observation of a critical pump power.

When running the third and fourth stage amplifiers under normal conditions, the relationship

between the pump power and the amplified power is nearly linear until the thermal loading point is reached and the amplified power first levels off and then decreases. Figure 4.21 a) shows a typical amplified power vs. pump power curve for the third stage amplifier. Here, we stopped increasing the pump power when the amplified power leveled off. Figure 4.21 b) shows a case where the pump power was increased beyond the critical power, causing the amplified power to fall.

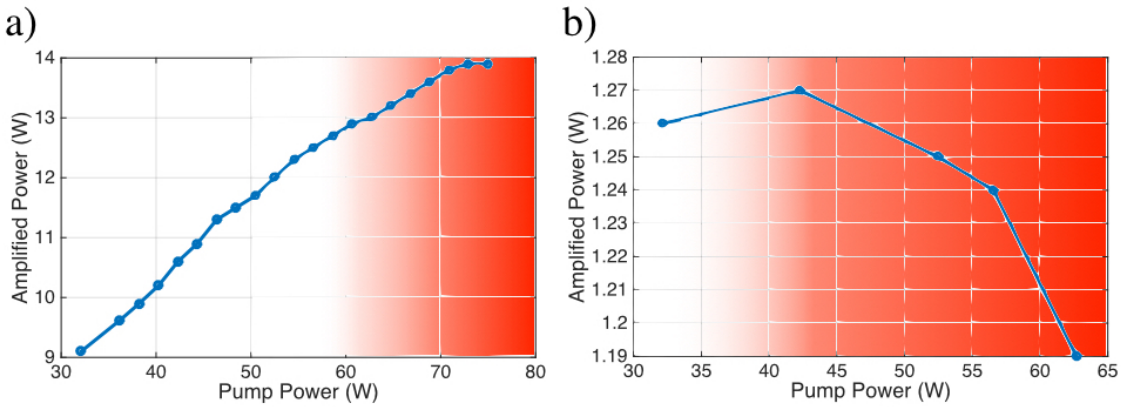

Figure 4.21: These two plots show the limitation on pump power due to thermal loading. Both data sets were taken using the third stage Yb:YAG amplifier. The red color qualitatively indicates the presence of thermal loading. Graph a) shows the linear rise in amplified power with pump power followed by flattened gain. Graph b) zooms in on the region of the curve at and beyond the thermal loading point, showing a decrease in amplified power as pump power is increased further. Graph b) shows the thermal loading point occurring at lower pump power than graph a) because the crystal mounting techniques were different. The mount used in a) has a higher overall thermal conductivity than the mount used in b).

To make the critical power as high as possible, we need to optimize the heat transfer away from the laser crystal. One approach is to find cooling systems which can reach lower temperatures and handle higher heat loads. This can be accomplished to some extent – we use heat exchangers with 90 W capacity at 77 K for the second and third stages and a 200 W-capacity heat exchanger for the fourth stage. At this temperature, and with our need for relatively low vibration on the optical table, it is both very expensive and difficult to find heat exchanges with higher capacity and similar or lower temperature than these.

Given that we are constrained to these heat exchangers, we can still optimize the heat flow

from the crystal to the cold head. The interface here is the crystal mount. Variables under our control are the material and geometry of the mount as well as the type of interfaces between the crystal and the mount and between the mount and cold head.

The equation governing heat conduction is:

$$\vec{q} = -\kappa \vec{\nabla} T \tag{4.22}$$

Where κ is the thermal conductivity of the material, $\vec{q}$ is the heat flux density, and $\vec{\nabla} T$ is the thermal gradient. Integrating both sides of Equation 4.22 over a surface, S, we find:

$$\frac{\partial Q}{\partial t} = -\kappa \int\int_S \vec{\nabla} T \cdot \mathrm{d}\vec{S} \tag{4.23}$$

Where $\partial Q/\partial t$ is the heat transfer per unit time. To get an intuitive picture of the implications of this heat transfer equation, we will now consider a uniform material of cross sectional area, A, and a temperature gradient, ΔT, across its length, L. This is not so different from the geometry of a crystal mount which, in essence, is a block of material with a heated crystal at one end and a heat sink at the other. A more realistic picture would show heat radiating radially from the crystal rather than uniformly distributed at one end of the material, but this simple picture still gives some insight. For the simple case described above, Equation 4.23 becomes:

$$\frac{\Delta Q}{\Delta t} = -\kappa A \frac{\Delta T}{L} \tag{4.24}$$

This immediately gives us a few principles for optimal mount design: 1) Heat flow is highest for materials with high thermal conductivity, so we should choose the best thermal conductor we can. 2) Larger cross sectional area of the mount improves heat flow – like a larger electrical wire facilitates current flow or larger pipe, water flow. We should therefore make the crystal mount as wide as possible. 3) Heat flow is inversely proportional to the length of the material, L, so the distance between the crystal and the cold head should be as short as possible.

In the process of optimizing the mount for the third stage, we tried three different configurations, illustrated in Figure 4.22. These consisted of two materials, Material 1 having higher thermal conductivity, and Material 2 being ideal for eliminating stress on the crystal that can lead to beam distortions. We tested each mount design to find the approximate critical power for thermal loading, using the same pump spot size in each test. The first mount was made completely of Material 2. This is the simplest zero-stress mount, but has the worst overall conductivity. We found that this mount has $P_c = 32$ W, which, with the current gain, is too little to get up to the 15 W-level of output that we were aiming for.

This motivated us to test the second mount (Figure 4.22 b)) which is like the first except that the section of Material 2 below the crystal was replaced with a solid disk of Material 1. This reduces the path length that the heat must take through the lower conductivity material. This design change increased the critical pump power to about 38 W. Though improved, this was still not good enough for the desired output power, motivating the use of an all-Material 1 design shown in Figure 4.22 c). The critical pump power of this last mount was 60 W, nearly double that of the solid Material 2 mount.

4.3.5 Stress-Induced Birefringence

As alluded to in the previous section, stress-induced beam distortions are another challenge that must be addressed in designing this high-power cryogenically cooled Yb:YAG amplifier system. The phenomenon to blame for these beam distortions is stress-induced birefringence, the strength of which is characterized by the stress-optic coefficient, C.

A material which exhibits stress-induced birefringence (also known as photoelasticity) will become birefringent or experience a change in its birefringent properties under the influence of stress. YAG is particularly susceptible to this effect because it is an optically isotropic crystal, having no natural birefringence. In the presence of stress, Yb:YAG's stress-optic coefficient is related to optical path retardance, R, according to Brewster's Law:

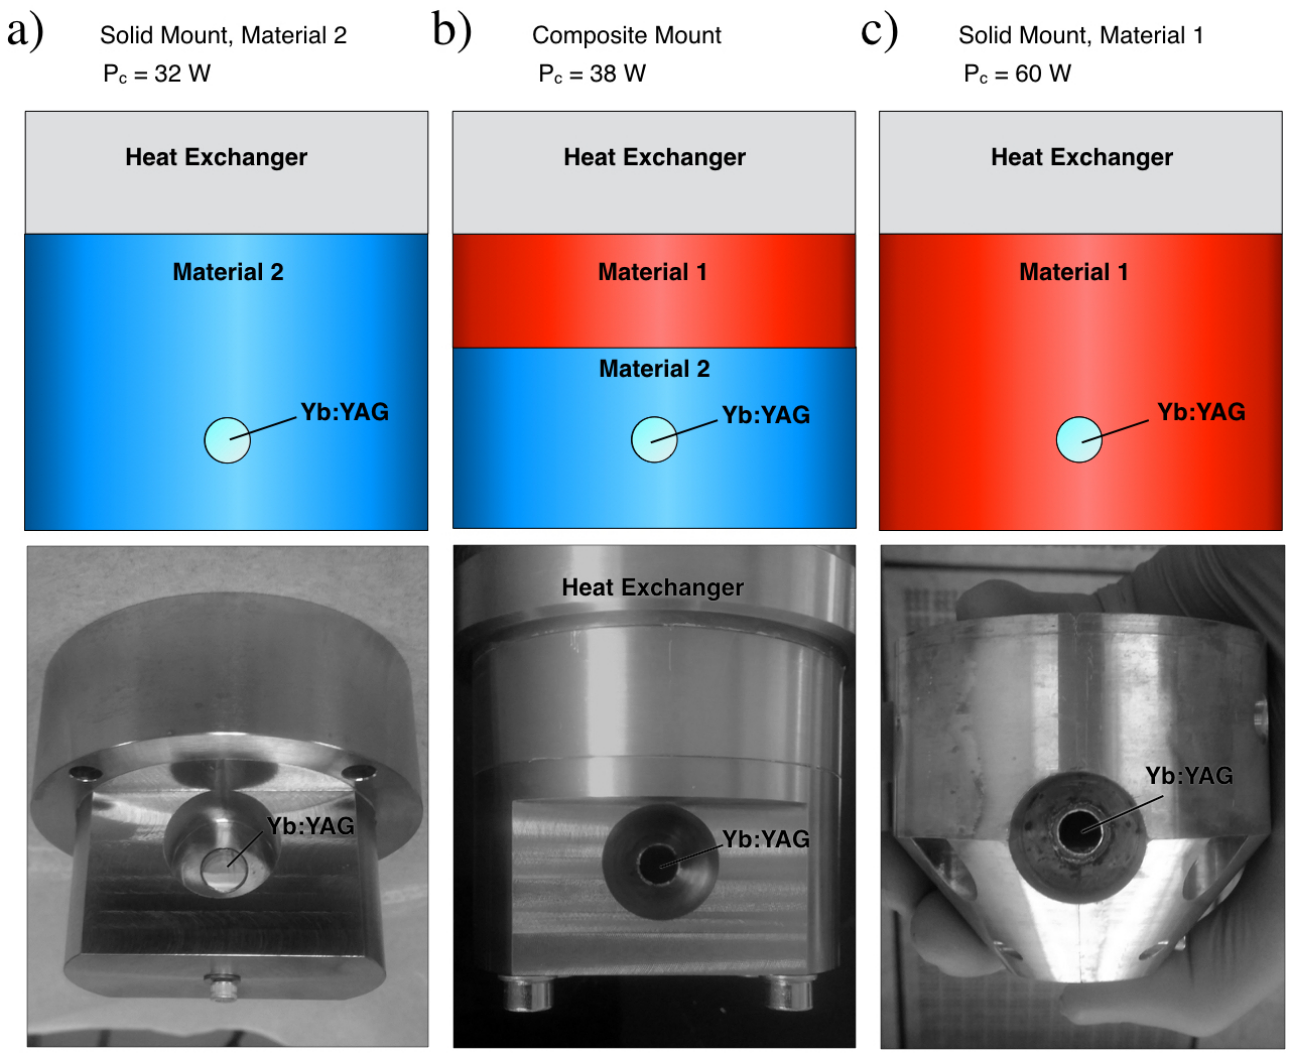

Figure 4.22: Drawings and photographs of three different crystal mount designs for the third stage Yb:YAG amplifier. All three have the same distance from the crystal to the cold head (heat exchanger). a) The first mount is made entirely out of Material 2, having lower thermal conductivity than Material 1, but eliminating stress on the crystal. b) The second mount is made of Material 2 in the area that surrounds the crystal, but sits on an block made of Material 1 to improve thermal conductivity. c) The third mount is entirely made of Material 1.

$$R = Ct(\sigma_1 - \sigma_2) \tag{4.25}$$

Here, t is the material thickness and σ_1 and σ_2 are the principal stresses in orthogonal directions. [55]

In our first attempt to cryogenically cool Yb:YAG for use with few-mm diameter laser beams, we used a mount made entirely out of Material 1. This material has a larger coefficient of thermal expansion than Yb:YAG, causing the mount to squeeze the crystal as it cools to its operating temperature around 70 K. Using this mount, we noticed that we were not able to efficiently couple the beam out of the amplifier cavity and saw distortions in the beam profile after the output polarizer. These effects are indicative of nonuniform depolarization which comes from stress-induced birefringence of the laser crystal.

In the regenerative amplifier and second stage, it was possible to correct for the polarization rotation using waveplates in the cavity. In these two amplifiers, the beam is small enough that it sees a uniform depolarization. In the third and fourth stages, however, the birefringence patterns were generally not uniform over the diameter of the beam, requiring a decrease in the amount of stress on the crystal or at least more uniform stress that would result in uniform depolarization over the region seen by the beam. Figure 4.23, a), shows an example of severe birefringence in the third stage crystal.

In order to continue using the mount made of Material 1, we first tried an alternative mounting technique to make the stress more uniform over the center of the crystal. This technique was partially successful in the third stage. Though not always repeatable, we found that it could be successful in making the stress birefringence pattern more uniform. Figure 4.23, b), shows a case in which the birefringence was sufficiently uniform to allow laser operation without distorting the beam. Unfortunately, this mounting technique was not good enough for the fourth stage amplifier where the crystal is larger and the amplified beam has twice the diameter of the beam amplified in the third stage. With these dimensions, the fourth stage is much more sensitive to stress.

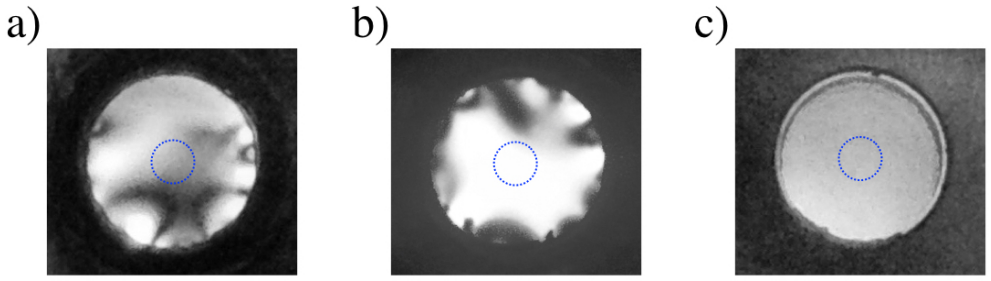

Figure 4.23: These images show stress-induced birefringence in the third stage laser crystal. The images were obtained by shining white light from an LED flashlight through the crystal with crossed polarizers on either side. The light fringes indicate regions where the incident light's polarization is rotated by π radians (or an odd multiple) and the dark regions indicate places where the polarization is not rotated or is rotated by a multiple of 2π radians. The blue circles on the images indicate the size of the laser beam. a) In the first image, the crystal was mounted using a mount made of Material 1. The birefringence was not uniform over the beam diameter. b) The second image shows a crystal mounted in the same way as the first, but which exhibits a large region of uniform birefringence, sufficient for laser operation. c) The third image shows a crystal that was mounted in a mount made of Material 2. It exhibits no signs of stress-induced birefringence.

To test the fourth stage mount for birefringence, we sent a linearly polarized laser beam through the mounted crystal in its vacuum chamber with another polarizer at the end of the line. Starting from room temperature, we used a power meter to measure the amount of light that maintained its original polarization state and another to measure the amount that had been projected onto the perpendicular polarization state. Figure 4.24 shows the setup used for this measurement as well as a sample of the results obtained. In a typical test of a mount made of Material 1, we observed a significant fraction of the beam's power being rotated. The rotation was also generally not uniform over the diameter of the beam.

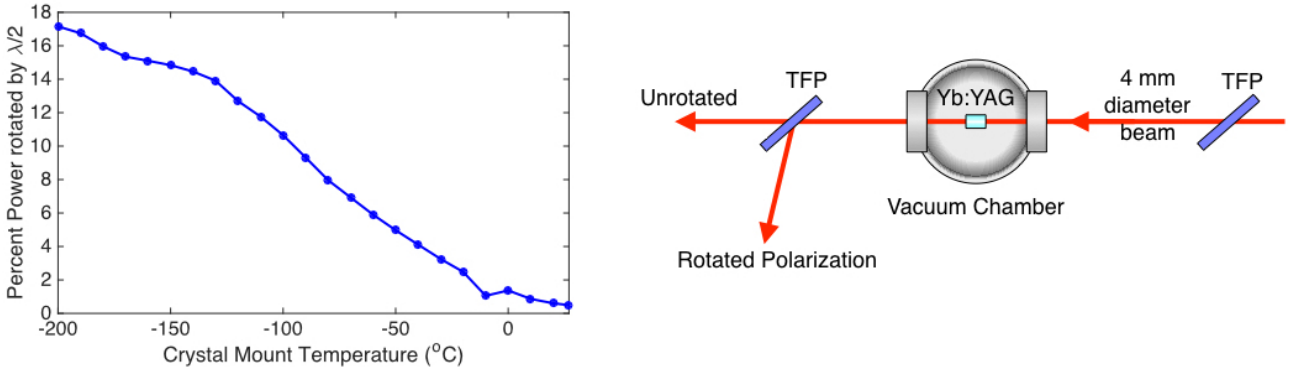

Figure 4.24: Plot of the percentage of light passing through the mounted fourth stage Yb:YAG crystal that was rotated from the p-polarization state to the s- state. The polarization rotation arises from stress-induced birefringence. This measurement took place while the crystal was cooled from room temperature down to -200°C. The birefringence worsens as the mount shrinks around the crystal at lower temperatures. The illustration on the right shows the setup used for this measurement.

In certain cases, we found that stress-birefringence occurred even at room temperature. To verify that the birefringence was caused by stress, we measured the amount of depolarization while heating the crystal, approaching the temperature at which the crystal was mounted. The incident beam nearly filled the crystal and was linearly polarized. It passed through a polarizing beam splitter after the crystal so we could measure the fraction of the light that had been rotated by a half wave. Figure 4.25 shows the results of this measurement. At room temperature, roughly a third of the incident light was depolarized and the depolarization was not uniform over the crystal

surface. As the crystal was heated, the depolarization dropped close to zero as it would have been when the mounting took place. This clearly demonstrates that the depolarization is stress-induced and arises from a mismatch between the thermal expansion of the crystal and that of the surrounding material.

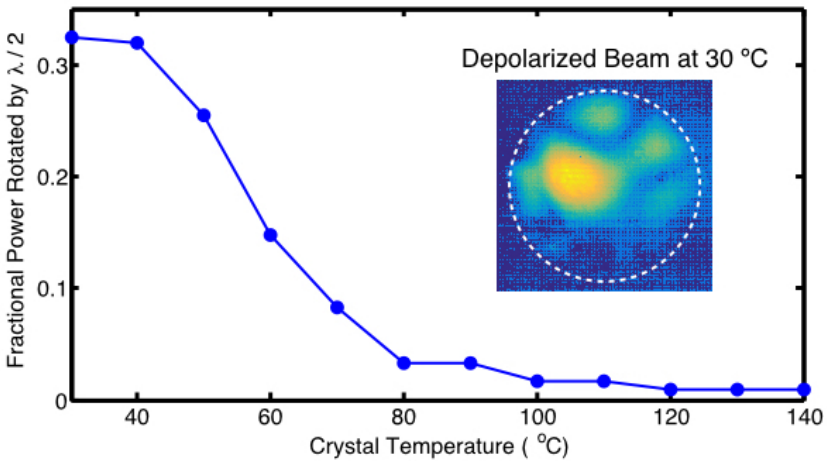

Figure 4.25: The plot shows the fraction of light incident on a mounted Yb:YAG crystal that was rotated by a half wave. The inset shows an image of the depolarized beam.

This clearly problematic level of stress-induced birefringence in the fourth stage motivated us to consider different mounting materials including ones that would match the coefficient of thermal expansion of Yb:YAG such as Material 2. To help, we created a simple model of the stresses in this mount, treating the crystal, a surrounding cylinder of material (or sleeve), and outer mounting material as a series of three concentric, thick-walled cylinders. The equations for the azimuthal and radial stresses in a thick-walled cylinder are:

$$\sigma_\theta = \frac{p_i r_i^2 - p_e r_e^2}{r_e^2 - r_i^2} - \frac{r_i^2 r_e^2 (p_e - p_i)}{r^2 (r_e^2 - r_i^2)} \tag{4.26}$$

$$\sigma_r = \frac{p_i r_i^2 - p_e r_e^2}{r_e^2 - r_i^2} + \frac{r_i^2 r_e^2 (p_e - p_i)}{r^2 (r_e^2 - ri^2)} \tag{4.27}$$

Where p_i is the pressure on the inner wall of the cylinder, p_e is the pressure on the external wall, and $r_{i(e)}$ is the internal (external) radius of the cylinder. We set up our problem with the

crystal in the center, having a zero inner radius, the sleeve surrounding it, and a cylinder outside of that made of Material 1. An illustration of this setup is shown in Figure 4.26.

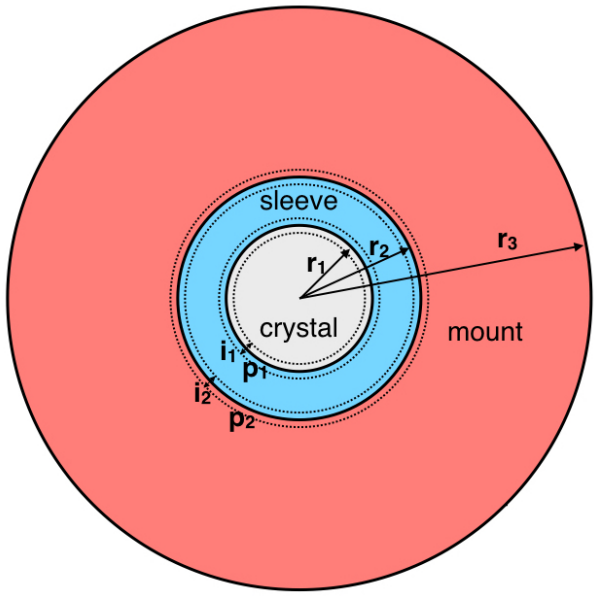

Figure 4.26: Calculations of stress on a laser crystal using a thick-walled cylinder stress model. The model was based on concentric cylinders starting with the Yb:YAG crystal in the middle, a sleeve of variable thickness and material surrounding the crystal, and an outer layer of Material 1 with a fixed radius.

Equations 4.27 and 4.28 can be written out for each interface using unknown pressures, p_1 and p_2 at the inner and outer interfaces, respectively. The goal is to find out the effect of varying the material and thickness of the sleeve on the pressure on the crystal's surface, p_1.

First, we also know that a given pressure will cause a deflection, u, of each material. The magnitude of the deflection will depend on the malleability of the material under pressure:

$$u = \frac{r}{E}(\sigma_\theta - \nu\sigma_r) \tag{4.28}$$

Here, ν is the Poisson ratio of the material, E is its Young's modulus, and r is the radius at which the pressure is exerted.

Secondly, we can use the Coefficient of Thermal Expansion (CTE), α, to find out how much each material would naturally expand or shrink if it were cooled while not under any additional

pressure. The final radius after a change due to thermal expansion (or shrinking) is:

$$R = R_o(1 - \alpha\Delta T) \tag{4.29}$$

Where R_o is the initial radius and ΔT is the temperature change. For our model of concentric cylinders, the outer material may naturally want to shrink to a smaller interface radius than does the material underneath. The difference in natural radii is called interference. Since this conflict in natural shrinking is exactly the effect that causes the pressure we are interested in, we can use it to solve for p_1 and p_2. Specifically, we have two equations for the interference:

$$\begin{aligned} i_1 &= R_{inner} - R_{outer} && (4.30)\\ &= R_o\Delta T(\alpha_{outer} - \alpha_{inner}) && (4.31)\\ i_1 &= u_{outer} - u_{inner} && (4.32) \end{aligned}$$

Where the subscripts *inner* and *outer* refer to the inner and outer materials. This is enough to solve for p_1 and p_2.

Assuming that the outermost material is Material 1 and that the innermost material is YAG, we can gain insight into the best strategy for a low-stress mount with maximum thermal conductivity. Specifically, this model yields helpful information in two regimes: 1) where the sleeve thickness is small compared to the surrounding material, and 2) where the sleeve material approaches the size of the mount. In the first regime, where the sleeve is relatively thin, we find that a soft, deformable material would yield the best results, regardless of whether its CTE matches that of Yb:YAG. In this case, the sleeve acts as a cushion. That is, the deformability of the sleeve prevents the inward squeezing of Material 1 from being transferred to the surface of the crystal.

In the second regime, the CTE of the sleeve material is more important than its hardness or elasticity. Even a very malleable material with a CTE larger than Yb:YAG can exert significant pressure on the crystal when it is sufficiently thick. On the other hand, when the mounting material surrounding the crystal has a CTE matching or close to that of Yb:YAG, there will be very little

or no pressure at the interface because the mount and crystal will shrink by the same amount as the assembly is cooled to cryogenic temperatures.

With the results of this simulation, we decided first to try a Material 1 mount having a thick sleeve made of indium, which is very malleable. This would have the advantage of the good thermal conductivity of Material 1 and the cushioning effect of indium. This design was an initial success, resulting in no visible birefringence as the mount was cooled to cryogenic temperatures. Also, with the relatively high thermal conductivity of Material 1, we were able to use this mounting scheme to reach 51 mJ output at 1 kHz from the fourth stage amplifier – our record for this amplifier system.

However effective in the short term, the above-mentioned mounting scheme turned out not to be very durable. Though indium is soft and compressible without transferring that pressure to the crystal, it is not good at springing back to its original state. The mount undergoes daily temperature cycling which, over time, deforms the sleeve and eventually causes gaps to open up that compromise conduction of heat away from the crystal. This mounting scheme could be effective in a facility that allows continuous operation of the cryogenic cooling system, but is not practical for our situation.

Our more permanent solution was to use a solid mount made of Material 2 with no sleeve. This design was successful in effectively eliminating stress birefringence down to -215°C. Figure 4.23, c), shows the image of a crystal mounted in this way between two crossed polarizers. There is no visible fringing from stress. This design is also robust, allowing temperature cycling without losing good thermal contact between the crystal and mount. We did not see any degraded thermal performance of this mount over roughly a one-year period of operation. Even at that point we needed to re-solder the crystal due to optical damage to the crystal, not because of a degraded interface.

The main drawback of the solid Material 2 mount is that its thermal conductivity is poor compared to Material 1. This turned out to be too much of a hit in thermal conductivity for the third stage amplifier which requires a gain of roughly 3× in two passes and has a 90 W capacity heat exchanger. This stage also has a sufficiently small beam (2 mm diameter) to allow operation

with a Material 1 mount, provided the birefringent pattern is sufficiently uniform. The fourth stage, however, has more than double the capacity in its heat exchanger (200 W) and could be operated with 187 W pump power and 36 mJ output even with the lower-conductivity mount. Material 2 was the better option for the fourth stage because it ensured preservation of good mode quality with a tolerable decrease in gain.

In summary, we found that mounts made out of material with coefficient of thermal expansion close to that of the laser crystal were the most effective and reliable method for removing stress-induced birefringence at cryogenic temperatures. The main drawback in this design is a loss of thermal conductivity compared to our original mount design. Where this conductivity is needed and some birefringence can be tolerated, a mount made of Material 1 is a better choice.

4.3.6 Low-frequency Amplitude Modulation Instability

In this section we will describe an instability in the Yb:YAG laser output that manifested as slow fluctuations in the amplitude. These fluctuations were visible to the unaided eye on a fluorescent detector card and looked like the beam was blinking. We observed a modulation depth as much as ten percent of the average pulse's amplitude. This is enough to translate into a temporally modulated signal and idler pulse trains after the OPCPA amplification stages. Fortunately, this instability can be corrected for and eliminated.

Figure 4.27 shows a typical "blinking" waveform consisting of modulations on the amplitude of pulses emerging from the regenerative amplifier. The amplitude of these modulations relative to the average is not significantly changed by the following amplification stages. When the amplifier chain was exhibiting this instability the slow (0.5-10Hz) modulations such as the ones depicted in Figure 4.27 would typically occur over 1-3 minute time intervals, separated by 5-15 minutes during the frequency would increase above 60 Hz and/or the modulation would temporarily stop. These spurts of unstable operation could be separated by hours of apparently stable pulses, making the instability difficult to track down.

Fortunately, we were able to find that the first appearance of this blinking pattern occurs at

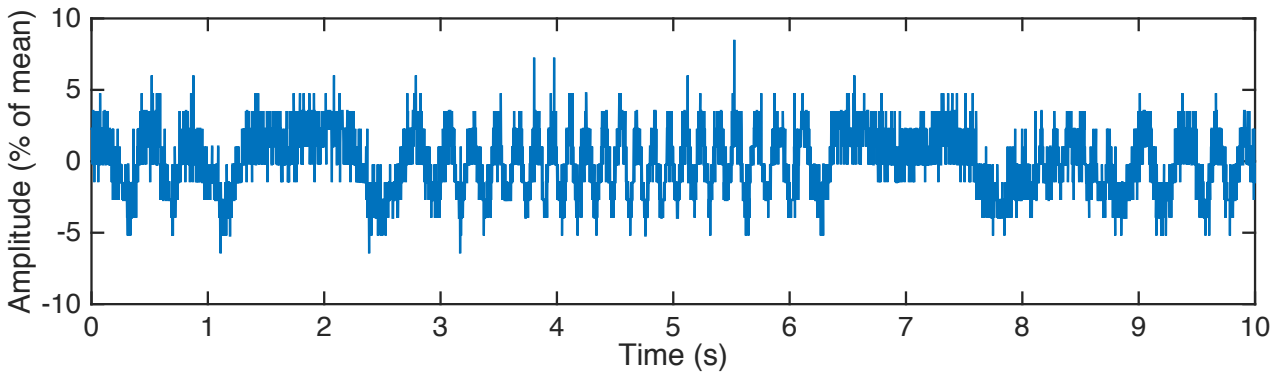

Figure 4.27: An oscilloscope trace showing the peak values of pulses as a percentage of the average amplitude. Over a period of ten seconds, this plot shows amplitude modulations with frequency varying between roughly 0.5 and 7 Hz.

the output of the volume Bragg grating stretcher, though it was not as pronounced as that seen at the output of the regenerative amplifier. This suggests that this instability is a result of spectral shifting in the Yb:fiber amplifier spectrum. The volume Bragg grating acts as a spectral filter, translating a spectral modulation into an amplitude modulation. The regenerative amplifier would exacerbate the problem by effectively applying an even narrower spectral filter. This hypothesis was confirmed using a narrow bandpass filter placed at the output of the fiber amplifier. After the filter, these amplitude modulations were evident whereas in the unfiltered pulse train, they were not.

Fortunately, the character of the Yb:fiber amplifier spectrum and FROG trace can be used to predict whether the amplitude modulation will occur and can be changed to eliminate the modulation altogether. Figure 4.28 shows examples of fiber amplifier spectra which illustrate the difference between fiber amplifier outputs that are likely to result in this blinking instability and those which are not. Spectrum a) in Figure 4.28 corresponds to a modulated pulse train. It exhibits two red flags that alert us to the possibility of an unstable output: 1) The 0.25 nm spectral region centered at 1030 nm is on the edge of the spectrum so there is relatively little light in the band needed to seed the cryogenically cooled Yb:YAG amplifiers. 2) The spectrum us strongly modulated. This feature is also reflected in time-domain modulations, manifested as side lobes on a FROG scan.

The danger of the first red flag, that 1030 nm is on the edge of the amplified spectrum, is twofold. First, if the spectrum shifts further so that the power at 1030 nm drops, the amount of light seeding the regenerative amplifier will drop also. This could result in bifurcation of the regenerative amplifier output. The second danger is that the slope of the spectrum at 1030 nm is high, meaning that a small shift in the spectrum will cause a relatively large change in the amount of power in that spectral band.

Generally the blinking occurs without bifurcation. In that case, the regenerative amplifier has enough seed energy, but still may exhibit low-frequency power modulations. The regenerative amplifier's output power is sensitive the amount of seed energy it receives up to a point where it saturates. Figure 4.29 shows this saturation effect as the fiber pre-amplifier's power is turned up.

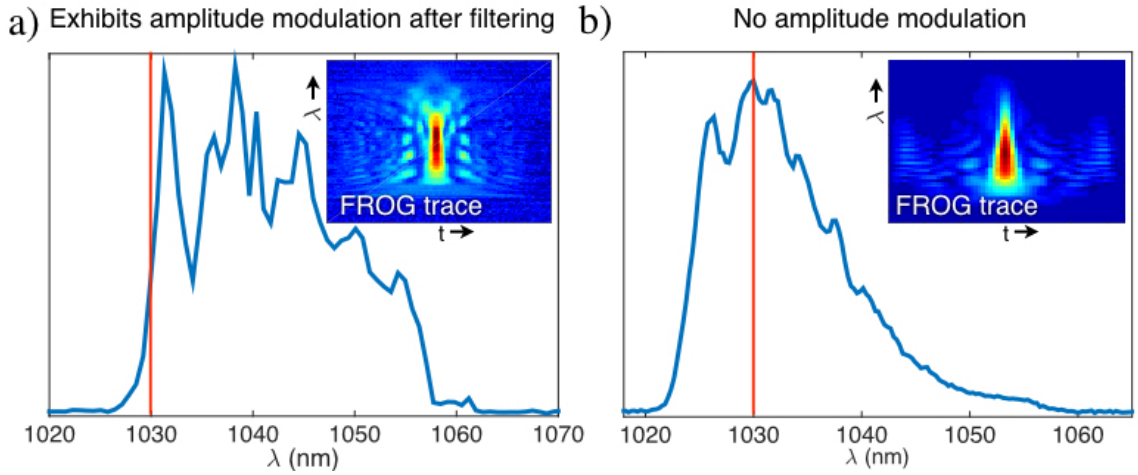

Figure 4.28: Two Yb:fiber amplifier spectra. The red bars indicate the spectral region that seeds the regenerative amplifier. The insets on the upper right of both graphs show FROG traces of the pulses. a) This spectrum leads to blinking behavior when spectrally filtered. It is highly modulated and has relatively little light at 1030 nm. The FROG trace shows significant side lobes in the temporal domain. b) This spectrum results in stable operation of the Yb:YAG amplifiers. It is relatively smooth and peaks near 1030 nm. The FROG trace is correspondingly clean, showing very little structure around the main pulse.

We can see from this plot that the regenerative amplifier's output power varies strongly when the seed from the fiber amplifier is less than about 500 mW (100 nJ). The fiber amplifier's power is sensitive both to the current of its pump laser and to the amount of light that is coupled into it after the CVBG stretcher. Therefore, if the spectral input of the CVBG varies in time and brings the regenerative amplifier's seed to the 500 mW power level in some of these modulations, we expect to see the modulations translated onto the output of the regenerative amplifier as well. The effect will be amplified by the spectral narrowing in the regenerative amplifier – the output of the fiber amplifier covers a 0.6 nm bandwidth whereas the regenerative amplifier only supports 0.3 nm.

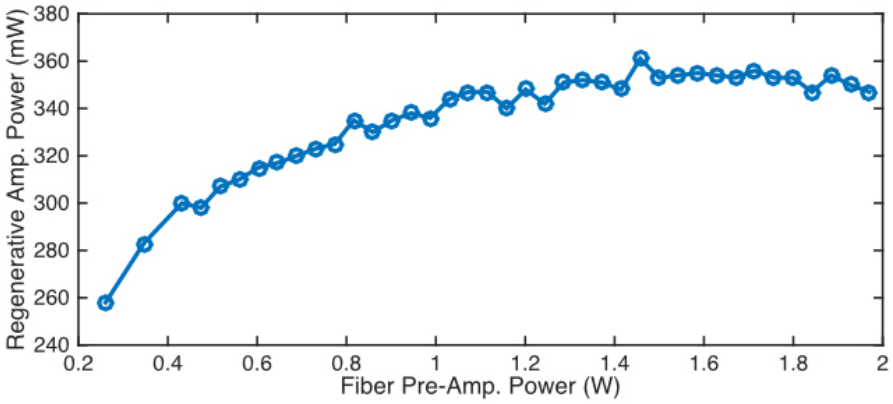

Figure 4.29: Plot showing the output power of the regenerative amplifier (1 kHz) as a function of the seed power received from the Yb:fiber pre-amplifier (5 MHz). The plot shows that the output power of the regenerative amplifier depends strongly on the input power at low seed energies, but saturates as the seed energy is turned up further.

In attempt to eliminate the blinking, the regenerative amplifier's seed band can be shifted away from the edge of the fiber amplifier's spectrum by finding a new oscillator modelocking regime that yields a red-shifted spectrum. This adjustment is a relatively straightforward matter of tuning the temperature of the fiber laser's pump diodes and/or adjusting the orientation of the inter-cavity waveplates. It turns out that this type of adjustment is not necessarily enough in itself, however. To eliminate blinking we must also eliminate any strong modulations in the Y-Fi amplifier's spectrum. Even if the 1030 nm band is centered, a strongly modulated spectrum usually still has a steep slope at that point, making the seed energy sensitive to small spectral shifts.

A truly stable fiber amplifier state will exhibit a relatively smooth spectrum like the one shown in Figure 4.28 b), and the 1030 nm section will be toward the middle. Smooth spectra typically also correspond to short, clean compressed pulses from the Yb:fiber amplifier since spectral modulations typically translate to temporal modulations through the Fourier transform. The spectrum shown in Figure 4.28 that resulted in blinking corresponded to a 160 fs pulse, whereas the other spectrum corresponds to a 130 fs pulse and was free of this amplitude modulation. Temperature stabilization of the fiber oscillator and amplifier also helps prevent spectral shifts.

We found that with this type of smooth, stable Y-Fi amplifier spectrum the Yb:YAG amplifier chain could run indefinitely without blinking. We also added a temperature-regulated water-cooled plate to the Y-Fi to help keep it stable. This setup was effective enough that we were able to stably run the Yb:YAG amplifier chain and pump the OPCPA even when the lab's cooling system failed and the room experienced a 15°F temperature increase.

4.3.7 Nonlinear Mode Reshaping

Another effect that influences the design of the Yb:YAG amplifier chain is self-focusing. This is a nonlinear optical phenomenon in which a high-intensity laser beam can induce an intensity-dependent change in refractive index of whatever material it passes through. It occurs whenever the power of the laser pulse is greater than a critical power that depends on the material: [51]

$$P_{crit,sf} = \frac{\pi(0.61)^2\lambda_o^2}{8n_o n_2} \tag{4.33}$$

Here λ_o is the center wavelength of the laser, n_o is the linear index of refraction of the material, and n_2 is its nonlinear index of refraction. If the power of the pulse passing through a given material is greater than the critical power (regardless of its mode size), it will experience a lens-like gradient in index of refraction. If the lensing is strong enough, the beam will collapse on itself after a distance, z_{sf}: [51]

$$z_{sf} = w_o\sqrt{\frac{n_o}{2n_2 I}} \tag{4.34}$$

$$= \frac{2n_o w_o^2}{\lambda_o}\frac{1}{\sqrt{P/P_{crit,sf}-1}} \tag{4.35}$$

Here w_o is the initial radius of a collimated laser beam, I is its intensity, and P is its power. This equation uses the paraxial approximation, assuming that the distance over which the beam collapses is large compared to the initial beam size. From here, we can get an idea of what we expect to see in our Yb:YAG amplifiers. The laser beam passes through vacuum cell windows, lenses, polarizers, the Yb:YAG crystals, waveplates, Faraday rotators, and, of course, air. These materials and their indices of refraction are listed in Table 4.3.

Material	n_o	n_2 (m^2/W)	$P_{crit,sf}$ (W) at 1030 nm	z_{sf}/w_o for I = 5×10^9 W/cm^2
YAG	1.835	7.2×10^{-20}	1.2×10^6	500
Fused Silica	1.458	2.7×10^{-20}	3.9×10^6	730
NBK-7 Glass	1.5071	3.4×10^{-20}	3.0×10^6	660
Sapphire, Al_2O_3	1.768 (o)	2.5×10^{-20}	3.5×10^6	840
Terbium Gallium Garnet (TGG)	1.95	1.7×10^{-19}	4.6×10^5	340
Air	1.0	4×10^{-23}	3.9×10^9	16000

Table 4.3: Linear and nonlinear indices of refraction for materials that the laser beam passes through in the Yb:YAG amplifier system, their corresponding critical power for self focusing, and the ratio of the self-focusing collapse distance, z_{sf}, to the initial beam radius, w_o, at a laser intensity of 5×10^9 W/cm^2. [95]

The full power per pulse in the second, third, and fourth stage amplifiers is about 50 MW, 150 MW, and 300 MW, respectively at 100 ps pulse duration. Since these powers are significantly above the critical power for all the transmissive materials in our laser except air, we may see some self-focusing.

We first noticed this effect in the fourth stage amplifier. We expected the pulse duration to be about 120 ps, but it had dropped to about 40 ps, increasing the power per pulse by more than a factor of two. Under these conditions, we observed that, with the fourth stage pump power fixed

at 187 W, the beam profile of the laser output changed significantly depending on the amount of energy we used to seed the fourth stage. At low seed power, the output beam was nearly gaussian, but became increasingly narrow in the center as the seed power was increased. This led to optical damage because the beam diameter shank down to about 1 mm, raising the intensity.

To check whether this mode reshaping was primarily a consequence of self focusing rather than thermal lensing or gain reshaping, we repeated the test with the pump laser off. We found that the beam shrinking and reshaping persisted. We noticed a particularly marked change in beam shape due to propagation through a Faraday rotator made out of terbium gallium garnet, which has an especially high nonlinear index of refraction and a correspondingly low critical power.

To confirm that this shrinking was a result of the beam's nonlinear interaction with material, we measured its change through a 20.4 cm-long fused silica block. With this much material, even a beam with only a few hundred milliwatts of power was reshaped. When the power through the block was turned up to 3 W, the beam reshaping became severe. Figure 4.30 shows images of the beam without the glass and with it at low power and at 3 W. The character of the beam reshaping is the same as we observed in the fourth stage amplifier: a tightly focused center developed while a low-intensity ring remained at the original size.

The shape change of this beam is consistent with the expected effect of self-focusing: an intensity-dependent refractive index change should have more effect on the center of the beam than on the edges. Also, this is distinct in character from beam reshaping that would occur from thermal lensing in a crystal pumped by a flat-top beam. Thermal lensing would result in an index-gradient equivalent to a lens with uniform focal length over the whole diameter of the beam. Gain reshaping should also have a different character from what we observe: gain should result in either uniform amplification over the beam profile or preferentially amplified edges if the center of the seed beam saturates the gain.

There are two main strategies for mitigating the effects of self-focusing: 1) Change the layout of the laser so that the beam interacts with less material and/or with material that has a higher critical power or 2) Change the laser beam so it will not interact with transmissive materials as

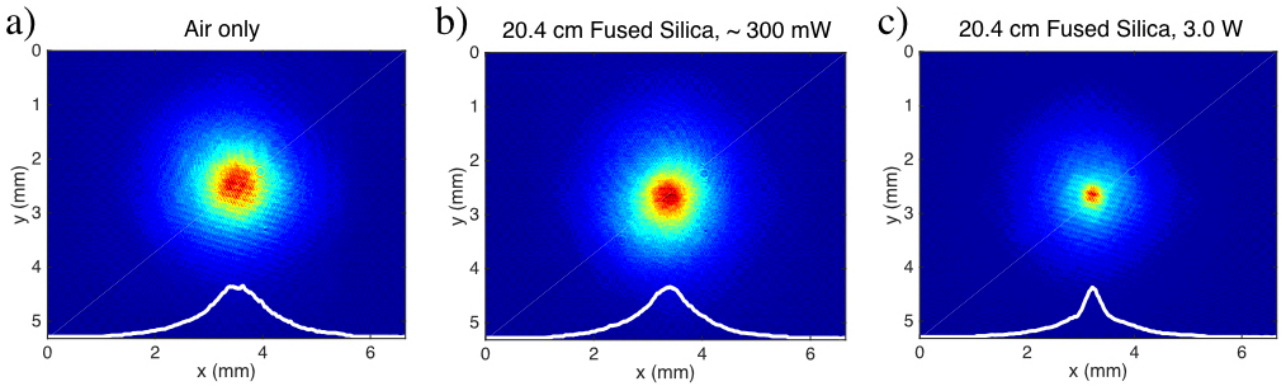

Figure 4.30: Beam images captured by a camera at a fixed position. a) The beam hits the camera after passing through air only (no self-focusing is expected due to the air at this power level). When only passing through air, the beam profile does not change when the power is increased. b) A 20.4 cm-long fused silica block is inserted into the beam path before the camera. About 300 mW of power is incident on the block at 1 kHz and 40 ps pulse duration, causing a small amount of beam reshaping. c) The fused silica block is still present and the beam's power is turned up to 3 W, resulting in severe beam reshaping due to self focusing. The white lines in each plot show a cross section of the beam profile through its maximum in the x-direction.

strongly.

We first tried to mitigate self-focusing by removing excess transmissive material from the amplifier beam path. We noticed evidence of self-focusing in the second and third stage amplifiers, so we began there, removing one pentagonal fused silica polarizing beam splitter and replacing three other beamsplitters with thin-film polarizers. Overall, this effectively removed 10.9 cm of fused silica from the beam path of the second and third stage amplifiers, accounting for multiple passes through certain optics. At a 40 ps pulse duration, this significantly reduced the self-focusing effect. The fourth stage amplifier was also simplified as much as possible, removing two telescopes and an optical isolator consisting of a terbium gallium garnet Faraday rotator and two polarizing beam splitters.

We also changed the location of the fourth stage quarter waveplate so that the light passing through the fourth stage crystal and vacuum chamber windows would have circular rather than linear polarization. The idea is that the nonlinear interaction between circularly polarized light and material is reduced by one third compared to the interaction with linearly polarized light. [96] The 15 mm-long crystal and two 9.5 mm-thick fused silica windows constitute the bulk of the transmissive material in the fourth stage, where the power is highest, so this improvement could be important.

Figure 4.31 shows the results of both self-focusing mitigation efforts: the effect of removing material from the second and third stage amplifiers and the effect of circular rather than linear polarization passing through the fourth stage Yb:YAG crystal and vacuum chamber at 6 W power. Both were effective, leading us to test the newly designed amplifier with the pump laser on, revealing whether the beam profile would still exhibit severe self-focusing at higher power.

Indeed, in spite of the notable improvements described above, as the seed power was turned up under full-power operating conditions in the the fourth stage Yb:YAG amplifier, we still observed significant beam reshaping and shrinking characteristic of self-focusing. This is shown in Figure 4.32 a). In light of this result, we increased the pulse duration to 120 ps from 40 ps using a double pass through the volume Bragg grating stretcher.

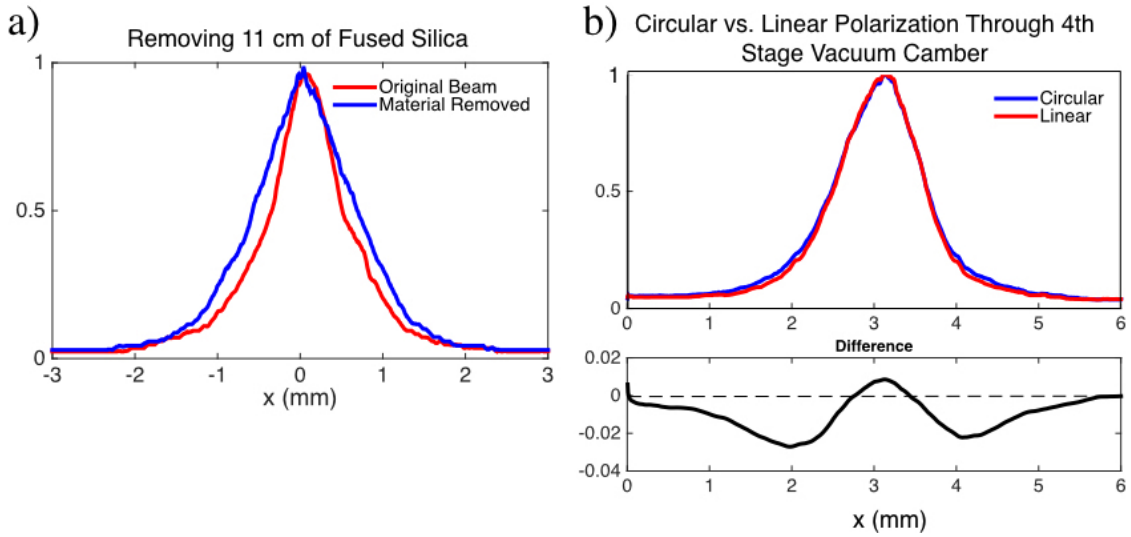

Figure 4.31: a) Beam profiles at the output of the third stage amplifier. The original full-power beam is plotted in red. The blue profile shows the beam at the same power after 10.9 cm of fused silica was removed from the beam path. It is significantly broader, especially toward the center, indicating that self focusing has been reduced. b) Beam profiles at the output of the fourth stage amplifier at 6 W power. In both cases the beam has passed through the crystal and vacuum chamber twice. In one case (plotted in red) the beam was linearly polarized when passing through and circularly polarized in the other (plotted in blue). The plot at the bottom right shows the difference between the linear and circular beam profiles, indicating that the linear polarized case is narrower in the center.

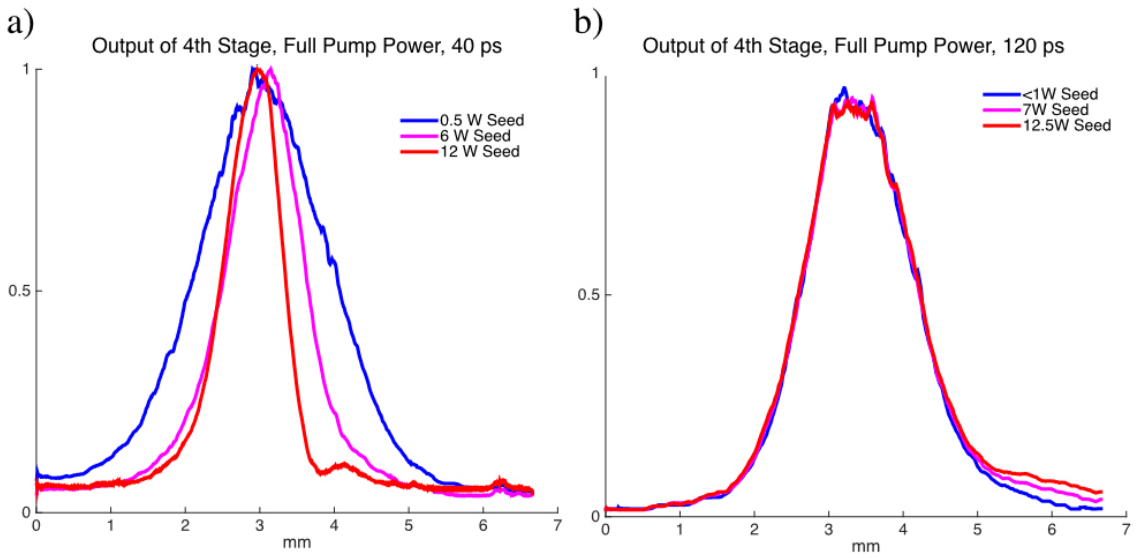

Figure 4.32: Both plots show beam profiles at the output of the fourth stage Yb:YAG amplifier with the pump laser on at full power (187 W). a) The plot on the left shows the change in beam profile due to self-focusing as the seed power is increased to 12 W. The pulse duration in this case is 40 ps. b) The plot on the right shows the same test using 120 ps pulses. In this case, the beam profile does not change noticeably as the seed energy is increased to 12.5 W. The slight increase at the right-most tail of the plot is due to scattered light hitting the camera.

Again running the fourth stage amplifier at full power but with the longer pulse duration, we found that self-focusing was negligible. That is, the beam profile remained unchanged as the seed power was increased, as shown in Figure 4.32 b). At 120 ps, the power per pulse is still higher than the critical powers for the laser's transmissive materials (except air), so the self-focusing effect can still exist, but is small enough not to make a noticeable difference in the beam diameter or shape.

Another useful way to see why the pulse duration increase helped remove self-focusing is to consider the total accumulated nonlinear phase of the beam as it propagates through the amplifier. This quantity is called the B-integral and is defined as:

$$B = \frac{2\pi}{\lambda} \int n_2 I(z) \mathrm{d}z \tag{4.36}$$

Where λ is the laser wavelength, n_2 is the nonlinear index or refraction, and $I(z)$ is the intensity as a function of the beam's propagation direction. When the B-integral is greater than about 3, self focusing may occur if the critical power is reached. At 120 ps pulse duration, the B-integral in the second, third, and fourth stages amplifiers are 2.7, 0.8, and 0.7, respectively. At 40 ps, these B-integrals become 8.0, 2.5, and 2.1, respectively (without taking into account beam shrinking which will drive the number even higher in the third and fourth stages). In the case of the shorter pulse duration, self-focusing is likely to start in the second stage and will therefore become increasingly severe in the final two stages. At 120 ps, however, the second stage only approaches the danger zone for self-focusing in terms of the B-integral and is therefore not likely to cause severe beam shrinking at that point. The combined B-integral through each of the next two stages is less than 2, so these alone are not likely to cause self-focusing.

In summary, self-focusing is a potential problem for our Yb:YAG amplifiers, but can be mitigated by minimizing the amount of material in the beam path. Also, for our current design, it is effectively eliminated by increasing the pulse duration to 120 ps.

4.3.8 Repetition Rate Scalability

In this section, we will look at the extent to which this Yb:YAG laser system can be scaled to higher repetition rates. Increasing the laser's repetition rate is advantageous because it will ultimately result in higher x-ray flux from high harmonic generation. A perfectly scalable laser system would allow for the pulse energy at a low repetition rate to be maintained as the repetition rate is increased. In a realistic laser system, however, there is usually a trade-off between average power and pulse energy. That is, higher pulse energies are easier to reach at lower repetition rates and, at high repetition rates where the average power is higher, the pulse energy is compromised.

The Yb:YAG pump laser in this thesis is typically run at 1 kHz repetition rate where it gives about 30 mJ per pulse. However, the Pockels cell in the regenerative amplifier which controls the repetition rate is capable of running at repetition rates as high as 100 kHz. Furthermore, due to the availability and relatively low price of 940 nm diode lasers, we have ample unused pump power which could help in scaling up the average power output. In the regenerative amplifier, we are currently pumping with 5 W, but have 50 W available, are using 24 W in the second stage but have 250 W available, are using about 80 W in the third stage but have 250 W available, and are using 187 W in the fourth stage but have 1000 W available.

In practice, there is indeed room to scale up the repetition rate, especially in the regenerative amplifier and second stage which are not limited by thermal loading. To test this, we compared the laser performance at repetition rates between 1 and 10 kHz. In increasing the repetition rate of the regenerative amplifier, one must be careful not to run into a regime of bifurcation instability. Bifurcation can typically be avoided by changing the number of passes through the cavity in tandem with the pump power. At higher repetition rates, more passes can be used so the pump power does not have to be increased as much. For this test, we were able to maintain 0.2 mJ per pulse from the regenerative amplifier at 1, 2.5, 5, and 10 kHz.

The second stage amplifier also is scalable to 10 kHz without sacrificing pulse energy. In this test, we maintained 3.6 mJ per pulse between 1 and 10 kHz. Even at 10 kHz, this stage of the

laser system is damage threshold limited rather than thermally limited, so the pump power can be increased without compromising pulse energy. At 10 kHz, the laser was pumped with 88 W yielding 36 W output power compared to 20 W pump at 1 kHz, yielding 3.6 W output.

The third stage Yb:YAG amplifier is a different story, however. Even at 1 kHz, this stage is limited by thermal loading, meaning that increased pump power actually decreases the laser's output due to warming of the laser crystal. In this case, as the repetition rate was scaled up, the third stage pump power could only be increased to the extent allowed by improved energy extraction from the laser crystal by the amplified beam. Specifically, as the repetition rate was scaled from 1 to 10 kHz, the pump power was increased from 80 W to 149 W, limited by thermal loading at each step. The pump power increase was not enough to maintain the pulse energy, causing it to decrease from 15 mJ at 1 kHz to 6.5 mJ at 10 kHz.

With this loss in pulse energy, we have the consolation of significant gain in average power. As the repetition rate is scaled from 1 kHz to 10 kHz, the average power increases from 15 W at 1 kHz to 65 W at 10 kHz. Notice that this is an increase in pump-to-output conversion efficiency from 19% to 44%. The power increase is therefore due to dramatically increased extraction of the stored energy in the laser crystal. Cryogenically cooled Yb:YAG has a relatively long upper laser state lifetime of about 1 ms and high saturation fluence of 1.8 J/cm^2. The corresponding saturation energy is above the typical damage threshold of the laser optics. This means that, at 1 kHz, one pulse making two passes through the crystal will attempt to extract all the energy the pump can provide in 1 ms, but it fails due to the fact that its energy is well below the saturation energy. At 10 kHz, on the other hand, ten pulses will pass through the gain medium in the same amount of time, extracting much more energy before it is lost. The output of the third stage amplifier as a function of repetition rate is shown in Figure 4.33, illustrating the simultaneous decrease in pulse energy and increase in average power.

We expect the same behavior in the fourth stage which, like the third, is limited by thermal loading: a decrease in pulse energy and increase in average power as repetition rate is scaled up. This result indicates that repetition rate scaling is feasible to the extent that a decrease in pulse

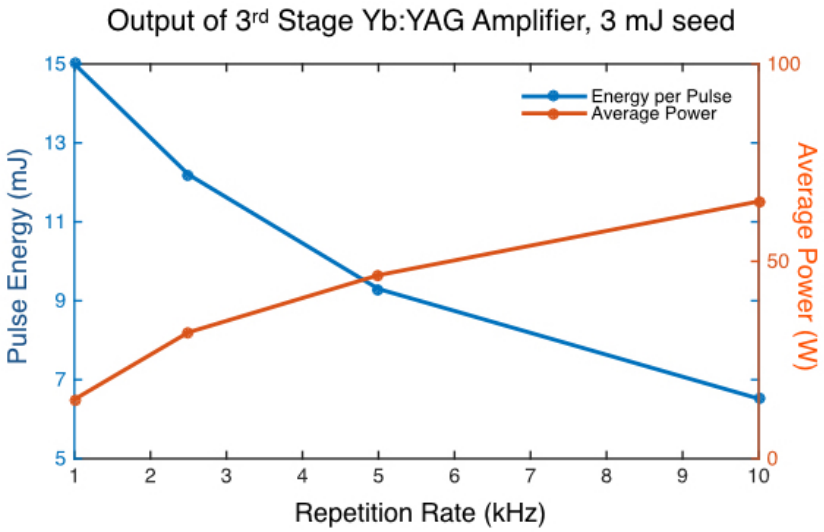

Figure 4.33: Output pulse energy and average power of the third stage Yb:YAG amplifier with fixed seed pulse energy of 3 mJ and pump power fixed at about 115 W.

energy can be tolerated. To mitigate the decrease in pulse energy, the beam size and the pump spot size in the third and fourth stage amplifiers could also be decreased. This would result in higher gain per pass at a given pump power (and therefore more efficient extraction) and would allow the pulse energy to be brought up to the point where it nears the damage threshold at the new beam size. This will be somewhere between the pulse energy it started with when the repetition rate was increased and that reachable at the lower repetition rate for which the laser was originally designed.

Another way to scale up the repetition rate with a smaller compromise in pulse energy is to improve the heat extraction from the third and fourth stage laser crystals. As discussed in a previous section, this must be done while balancing the need for low stress-induced birefringence.

Finally, it is possible to scale this laser system up beyond 10 kHz. For applications other than high harmonic generation for which millijoule-level energies are not needed in mid-infrared pulses, it can be advantageous to run at higher repetition rates. For example, our collaborators, Luther et al., built a 100 kHz Yb:YAG laser for pumping a 3 μm OPCPA system for 2D IR spectroscopy. [97] They used a Yb:YAG laser architecture similar to our regenerative amplifier and second stage amplifier to reach a 70 W output power with 140 W pump power. This corresponds to 0.7 mJ per pulse at 1 μm. The advantage of the higher repetition rate is shown by the fact that at 100 kHz they were able to reach more average power in two stages than we reached in three at 10 kHz.

In summary, this 100 kHz laser illustrates the flexibility of the Yb:YAG amplifier system described in this thesis, showing its utility for applications such as spectroscopy. For pumping an OPCPA intended for keV high harmonic generation, however, it is advantageous to keep the repetition rate below 10 kHz. In the future, smaller beam sizes in the third and fourth stages as well as improved heat extraction from the laser crystals will help maintain the pulse energy at higher repetition rates.

4.4 Seed

In this section, we will describe the arm of this laser that generates the 1.6 μm seed for the OPCPA. Beginning with 1.0 μm light from the Yb:fiber oscillator and amplifier, we used an Optical

Parametric Amplifier (OPA) based on Difference Frequency Generation (DFG) in MgO:PPLN to generate a broadband 1.6 μm beam. Once generated, this light is temporally stretched and delayed so it arrives at the first OPCPA stage at the same time as the output of the Yb:YAG amplifier chain.

4.4.1 OPA Architecture

The OPA consists of two inputs: a 1.0 μm wavelength pump beam and a broadband seed with spectral content at 1.6 μm. These undergo nonlinear difference frequency mixing in an appropriate nonlinear medium, depleting the pump, amplifying the seed (now called the signal) and an idler with photon energy corresponding to the difference in energy between the pump and signal photons.

In our design, both the pump and the seed originate from the Y-Fi Yb:fiber oscillator and amplifier. This allows the repetition rate of the OPCPA seed to be passively synchronized to that of the OPCPA pump. Also, the OPA idler will automatically have a consistent Carrier Envelope Phase (CEP). If the OPA idler were chosen to seed the OPCPA rather than the signal, this would result in a CEP-stable OPCPA output.

The OPA setup is illustrated in Figure 4.34. Starting from the right, the 1.0 μm light is split into three arms: one to generate the OPA seed, one to pump the OPA, and finally one to seed the Yb:YAG amplifiers. The OPA seed is generated by focusing the 1.0 μm light into a 3 mm-thick piece of sapphire, generating white light that spans the 1.6 μm spectral region. This seed is then recombined with the pump in a Magnesium-Oxide-doped Periodically Poled Lithium Niobate (MgO:PPLN) crystal. The seed and pump are overlapped in time and space so that phase matched parametric amplification occurs in the crystal. The resulting 1.6 μm signal beam is then separated from the depleted pump and the idler to be routed into a stretcher.

The stretcher uses a polarizing beamsplitter and a quarter waveplate to separate the stretched light from the input. The stretching takes place in a pair of transmissive diffraction gratings. The beam passes through the pair of gratings and is retroreflected back through them a second time before being ejected from the stretcher and coupled into a delay fiber. The delay fiber delivers the

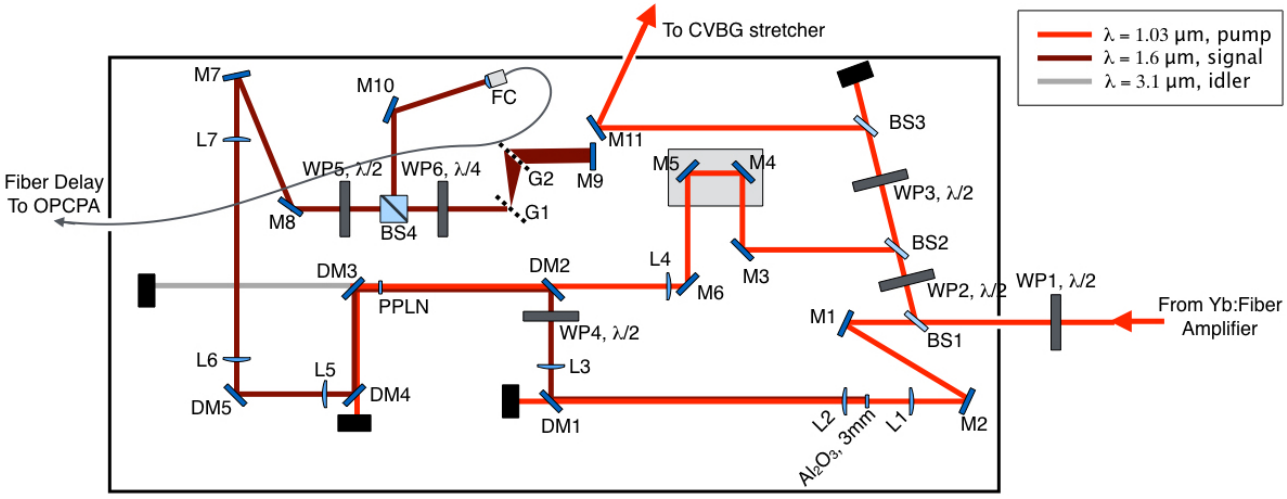

Figure 4.34: Diagram of the OPA and 1.6 μm stretcher. The input from the Y-Fi amplifier is shown on the far right and the stretched signal exits the OPA via a delay fiber (exiting on the far left). In this diagram, the 1.0 μm beams are drawn in red, 1.6 μm in maroon, and 3.1 μm in grey. The abbreviation WP indicates a waveplate, BS, a beamsplitter, M, a mirror, L, a lens, Al_2O_3, a piece of sapphire, DM, a dichroic mirror, PPLN, the nonlinear crystal, G, a diffraction grating, and FC a fiber collimator.

stretched 1.6 μm beam to the OPCPA.

4.4.2 White Light Generation

In this section we will describe the OPA seed generation process in more detail. The requirements for the OPA seed are 1) that it have at least as much bandwidth as the desired OPA signal, 2) it provide light in the spectral region around 1.6 μm, and 3) it be stable and free of amplitude jitter.

We need the seed to be broadband because the OPCPA pump is spectrally narrow with about 0.3 nm bandwidth. The seed therefore must supply enough bandwidth for the OPCPA idler to be compressible below 100 fs. This corresponds to at least 140 nm of bandwidth at 3.0 μm and 35 nm of bandwidth at 1.6 μm. We will therefore aim for at least 40 nm bandwidth in our OPA signal.

In order to create a broadband signal, the seed must be broadband as well. We accomplish this by generating a supercontinuum in bulk sapphire. Supercontinuum generation (also called white light generation) is a result of several nonlinear processes occurring simultaneously when an intense light pulse interacts with a material (sapphire, in this case). More specifically, supercontinuum generation is caused by optical shock occurring at the trailing side of the pump pulse due to self-steepening and space-time focusing. The specifics of the process are also influenced by higher-order nonlinearities. [98] This process is aptly named, yielding a broad, continuous spectrum.

We use a 2.3 W beam of 130 fs pulses at 5 MHz repetition rate (460 nJ per pulse) to drive the white light generation. This light is focused onto a 3 mm-long slab of sapphire with a focal spot diameter of approximately 25 μm. The resulting white light has some energy in the visible spectrum and extends beyond 1600 nm on the red side. The infrared section of the white light continuum spectrum is shown in Figure 4.35.

There are several advantages of this method of generating white light. The first is the simplicity of the setup, requiring only a fast focusing lens and a piece of bulk sapphire. As long as no aberrations are added by the focusing optic, this process results in an excellent gaussian mode. Also, if the pump light which drives the white light generation consists of short pulses, free

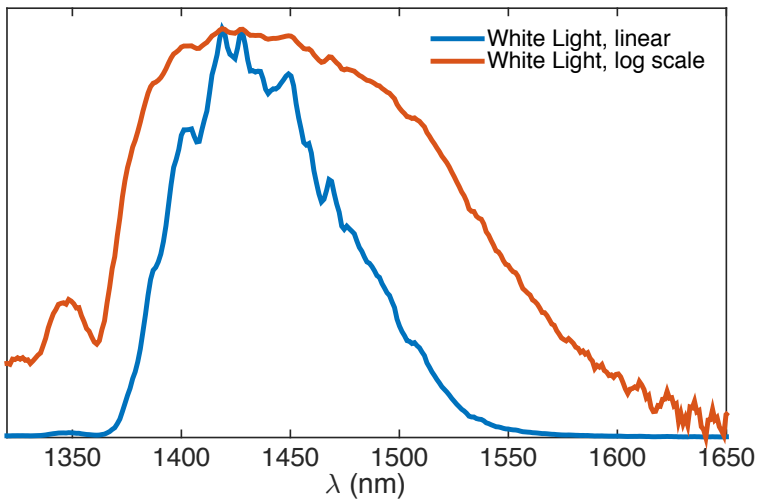

Figure 4.35: White light spectrum generated in bulk sapphire and used to seed the OPA. The spectrum plotted in blue is plotted on a linear scale. The orange plot shows the same spectrum on a log scale, showing that there is some power beyond 1600 nm.

of temporal structure such as a leading or trailing pulse, it will result in a stable, jitter-free white light spectrum. Since the focused light must be intense for white light generation, we sometimes see slow damage of the sapphire causing the white light generation to drift to lower power. This can be easily corrected by slightly moving the sapphire slab so the damaged spot is no longer in the beam path. The sapphire slab is mounted in a rotation mount so that this adjustment is quick and simple. Another advantage of white light generation in bulk sapphire is the fact that the spectrum in the region between 1400 and 1600 nm is continuous and smooth. This will allow the OPA signal's center wavelength to be tuned anywhere within that range. Consequently, the OPCPA idler's center wavelength will also be tunable between 2900 and 3900 nm.

The main drawback of generating white light in bulk material is that it requires a larger amount of driving energy than other methods. The Y-Fi fiber laser provides ample power (10 W at 5 MHz; 2 μJ per pulse), so this constraint is not a hardship for us. However, in a previous design of Y-Fi, we only had 500 nJ per pulse to work with, making our current white light generation method, which uses 460 nJ of driving light, impractical. To get around this limitation, we first tried recycling the driving light left over from white light generation for use in pumping the OPA. The problem with this strategy is that the nonlinear phase accumulated by the 1.0 μm light is significant and causes the OPA signal to have shot-to-shot amplitude jitter as well as longer-term spectral drift. These issues lead us to explore an alternative method of white light generation using a Photonic Crystal Fiber (PCF) rather than bulk material.

Alternative method for white light generation using a PCF Photonic crystal fibers are optical fibers which are internally patterned with microstructures that alter the fiber's guidance properties. The dispersive characteristics of the fiber can be tailored and nonlinear effects can be enhanced by tight mode confinement. The nonlinear interaction between the fiber and the light it guides can be strong over a long length so that even a relatively small amount of energy (nanojoule-level) can result in significant spectral broadening. The spectral broadening is due to various nonlinear effects including self phase modulation and Raman scattering. This can result in dynamics such as soliton formation and fission. The specific nature of the spectrum that emerges

from the fiber depends on many factors including the zero dispersion wavelength of the fiber, its length, and the energy, pulse duration, spectral phase, and polarization state of the incident light. [98, 99]

We used a polarization-maintaining PCF with zero dispersion wavelength at 1040 nm, near our driving laser's wavelength. The polarization-maintaining property of this fiber helps prevent instabilities due to the inherent birefringence of the fiber. [99] We chose to use a fiber roughly 7 cm in length, resulting in a broad white light spectrum with about 150 nm bandwidth, covering the 1550 nm region. This requires about 65 nJ of driving light – much less than the 460 nJ used to drive white light generation in bulk sapphire.

The length of the fiber was chosen as a tradeoff between bandwidth and conversion efficiency. For a given length of PCF, as the pump light is turned up, some light will eventually appear at the output with wavelength above 1300 nm, shifting farther to the red as the power is increased. With a longer fiber, the process is much more efficient. As illustrated by the spectra shown in Figure 4.36, a 125 cm-long fiber can generate light at 1500 nm with only 10 nJ of pump energy. That light emerges as a soliton, recognizable as an isolated spectral peak which carries a significant fraction of the pump energy. With 12 nJ pump, we measured 1.2 nJ at 1550 nm – a ten percent conversion efficiency from the pump wavelength. This is enough energy that it is possible to seed the seed the OPCPA directly, without additional amplification in an OPA. We tried this with some success, but found that the OPCPA power and bandwidth output can be boosted when the seed has more energy.

The main disadvantage in using the 125 cm PCF for white light generation is that the spectral bandwidth that emerges is relatively narrow: typically about 40 nm in width. This bandwidth is near the minimum needed for a 100 fs compressed idler, so we are counting on not having spectral narrowing in the OPCPA and being able to compress the pulses very close to their transform limited pulse duration. More bandwidth in the white light would therefore be beneficial and can be obtained by shortening the fiber. As mentioned earlier, for a gain in bandwidth we need to use more pump power. To reach 1550 nm, a 5 cm fiber requires 80 nJ pump, as compared to 19 nJ

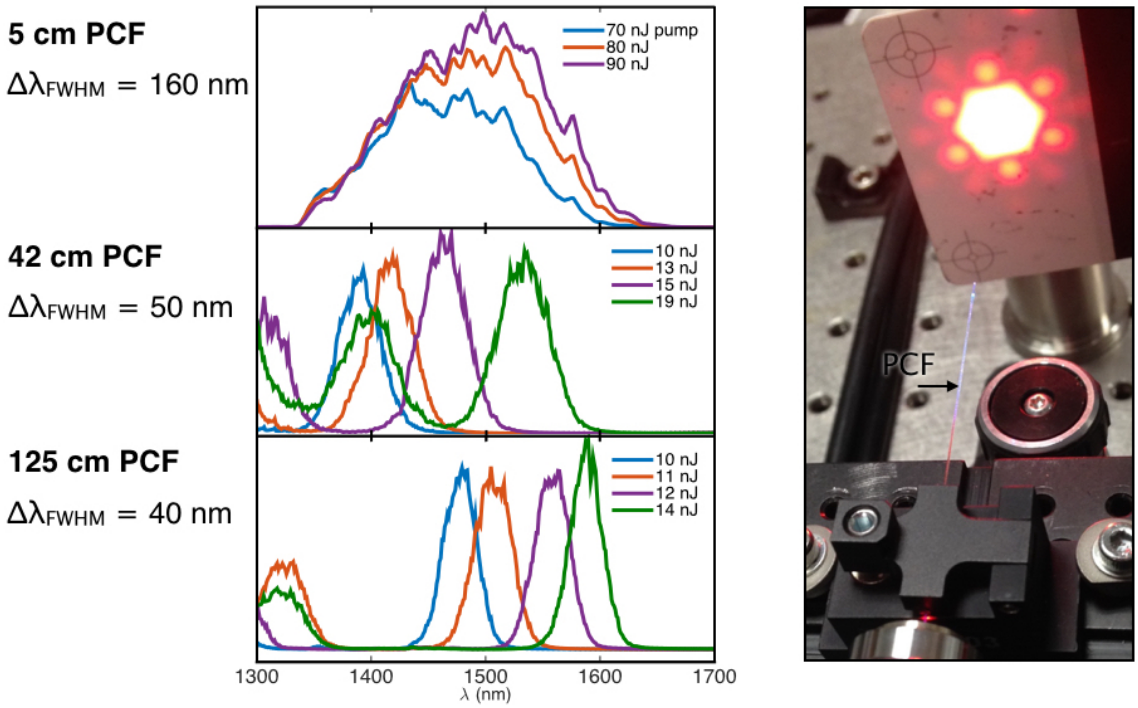

Figure 4.36: Left: Plots showing the spectra emerging from three photonic crystal fibers of different lengths (5, 42, and 125 cm). For a given fiber length, we varied the driving laser pulse energy to observe the change in output spectrum. For each change of pulse energy, the spectrum was optimized by changing the polarization state of the driving light with a half waveplate placed before the entrance of the fiber. Right: Photograph of a PCF and the output on an IR detector card. The PCF generates visible light as well as infrared. The "snowflake" structure on the beam reflects the shape of the structured fiber core.

required for a 42 cm fiber and 12 nJ required for a 125 cm fiber.

It is also worth noting that as the fiber is shortened and/or pump power is increased, multiple solitons will be generated. This decreases the amount of energy in each one, and can lead to amplitude jitter if the solitons overlap spectrally and interfere. With a 5 cm fiber, the light appears as one broad peak. This peak contains substantially less energy than the 1.6 μm output of the 125 cm fiber, but it is enough to amplify in the OPA. Amplitude jitter caused by multiple soliton interference can be an issue, but can be mitigated by changing the input polarization state of the driving light and by making sure the incident pulses are optimally compressed. Any remaining jitter can be eliminated by saturating the OPA process.

This work showed that photonic crystal fibers are a good solution when pulse energy is limited. There are two main drawbacks of this technique, however. The first is that the output of the PCF is extremely sensitive to the stability of the driving light, more so than supercontinuum generation in bulk material. The spectrum and stability can be dramatically changed by instabilities in the driving laser's pointing, polarization, and dispersion characteristics. Pointing changes decrease the coupling efficiency into the fiber and therefore can shift the spectral peak to smaller wavelengths. Also, with a short fiber, the power needed to generate light at 1.6 μm is close to the fiber's damage threshold. If the power of the driving laser fluctuates, this can damage the fiber (usually the ends), requiring repair.

The second drawback of PCF-based white light generation is that it imparts a large amount of complicated spectral phase to the OPCPA seed. To design a functional compressor it is important to be able to predict the spectral phase that will end up on the OPCPA idler, originating from the seed. The supercontinuum generation process in a PCF involves high-order nonlinearities, resulting in a phase that is not easy to predict. We found that the PCF-based white light generation results in group delay dispersion with roughly the same magnitude as that imparted by our stretcher. In contrast, white light from bulk sapphire does not, in our experience, impart significant spectral phase to the OPA signal.

Even with these issues, we used a photonic crystal fiber-based OPA to successfully operate

the 3 μm OPCPA system at the full power reported in this thesis. It is a viable technique for this application, but we preferred the increased stability and predicability of bulk white light generation when more driving laser energy became available.

4.4.3 Difference Frequency Generation

The core of the OPA is in the nonlinear crystal in which Difference Frequency Generation (DFG) occurs. In this section, we will give the parameters of the OPA in its typical operating state.

The nonlinear crystal used for this OPA is a MgO:PPLN crystal. This crystal is patterned with domains of alternating optical axis orientation, allowing quasi phase matched DFG. MgO:PPLN has a relatively high d_{eff} coefficient of 14.9 pm/V, allowing efficient collinear amplification. The center wavelength of the amplification is determined by the poling period and can be fine tuned by changing the crystal's temperature. [63] The bandwidth that can be amplified depends on the length of the crystal. Longer crystals allow a longer interaction region between the pump and signal and therefore the potential for more gain, but the bandwidth of the signal will be larger for a shorter crystal. The MgO:PPLN crystal used in this OPA has an array of poling periods, but we typically use the 30.5 μm period to phase match a signal with center wavelength at about 1550 nm. The crystal's length is 2 mm, allowing amplification of about 72 nm bandwidth. [63]

The seed is amplified in a single pass through the crystal during which it is temporally and spatially overlapped with the pump beam. The pump supplies 3.4 W of power at 5 MHz (680 nJ) and is focused to a 165 μm diameter spot. We use a dichroic mirror to make the pump and the seed beams collinear before they reach the crystal. Typically, we generate about 170 mW of amplified signal with excellent mode quality.

To ensure the signal is free of amplitude jitter and to optimize its power and bandwidth, we make sure to saturate the OPA process. That is, push the pump power to the point where the signal gain levels out, but before it begins to back-covert. Typically, at the point of saturation a faint ring appears around the signal mode. Also, faint visible white light may be generated when

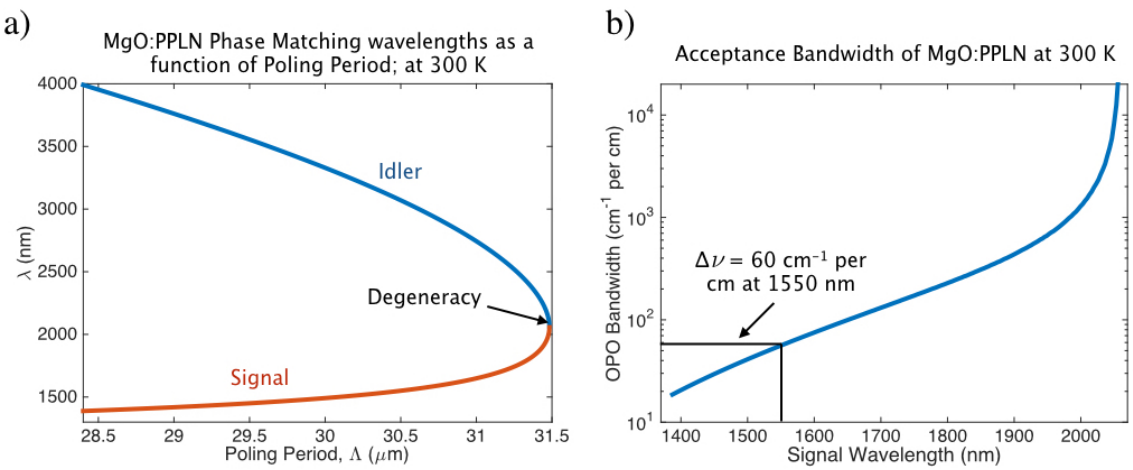

Figure 4.37: a) The plot on the left shows the DFG phase matching curve of MgO:PPLN as a function of poling period, Λ, at 300 K. b) The signal bandwidth in the DFG process in MgO:PPLN is plotted as a function of the signal wavelength. Here, the bandwidth is defined as the inverse of the difference between the signal and idler group velocities: $1/(v_{g,signal} - v_{g,idler})$. It is reported in wavenumber per unit length of the crystal. These data were obtained from [63].

the pump intensity nears the damage threshold of the crystal (around 5×10^9 W/cm^2).

The spectrum of the OPA signal can be tuned between between 1400 and 1600 nm, limited by the bandwidths of the coatings of the dielectric mirrors in the OPA. The tuning is accomplished by adjusting one or more of the following parameters: the poling period of the MgO:PPLN, the crystal's temperature, and the timing between the pump and seed.

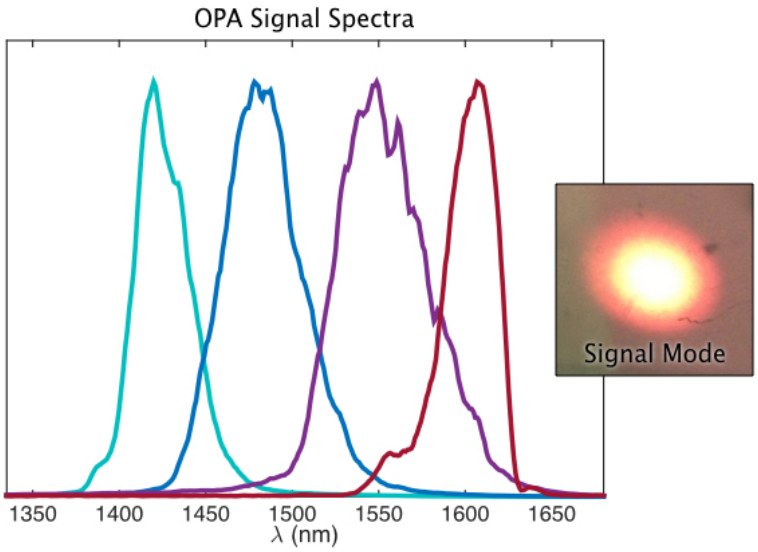

Figure 4.38: This plot shows spectra of the OPA signal with central wavelength tuned between 1420 and 1610 nm. The inset shows the signal mode on a fluorescecent IR detector card.

Examples of OPA signal spectra are shown in Figure 4.38. The spectrum centered at 1550 nm is typically what we would use to seed the OPCPA. This spectrum has FWHM bandwidth of about 60 nm and spans more than 100 nm tail-to-tail. This is more than sufficient bandwidth for an OPCPA idler that must be compressed below 100 fs. If this bandwidth is maintained through the OPCPA, it will result in a 240 nm idler with a transform limited gaussian pulse duration of 59 fs.

4.4.4 Stretcher

The final step in preparing the OPA signal to seed the OPCPA is to temporally stretch it so that it will be closer to the pulse duration of the pump and the resulting idler will be compressible.

There are many possible approaches to this problem. In this section, we will begin by describing the current configuration of the stretcher and discussing its limitations. Next, we will describe other approaches used in the past, their merits, and their shortcomings. Finally, we will discuss two new designs that will address the current stretcher's limitations.

Currently, the 1.6 μm beam amplified in the OPA is stretched using a double pass through a pair of transmissive diffraction gratings, as illustrated in Figure 4.34. Also, a diagram of this type of stretcher is shown in Figure 4.6. This results in a spectral phase given by [100]:

$$\phi(\omega) = \frac{4\omega L_g}{c}\left[1 - \left(\frac{2\pi c}{\omega d} - \sin\gamma\right)\right] \tag{4.37}$$

Here ω is the angular frequency, L_g is the perpendicular distance between the gratings, c is the speed of light in vacuum, d is the grating groove spacing, and γ is the angle of the incident beam with respect to the grating normal. The second derivative of this phase with respect to the angular frequency gives the Group Delay Dispersion (GDD, or second-order dispersion), which corresponds to a linear spreading of the incident light's frequency components as a function of time. The third derivative, or Third-Order Dispersion (TOD), will tell us about the quadratic component of the frequency-distribution. Likewise, the Fourth Order Dispersion (FOD) will tell us about the cubic frequency distribution, and so on. For a grating pair in the present arrangement, the GDD is negative, meaning that the blue side of the stretched pulse will lead. Also, the TOD will be positive, and the FOD, negative. The magnitude of these dispersion components for our grating stretcher are:

$$GDD = \frac{\partial^2\phi}{\partial\omega^2} = -\frac{\lambda^3 L_g}{\pi c^2 d^2}\left\{1 - \left(\frac{\lambda}{d} - \sin\gamma\right)\right\}^{-3/2} \tag{4.38}$$

$$TOD = \frac{\partial^3\phi}{\partial\omega^3} = -\frac{3\lambda}{2\pi c}\frac{\partial^2\phi}{\partial\omega^2}\left\{\frac{1 + \frac{\lambda}{d}\sin\gamma - \sin^2\gamma}{1 - \left(\frac{\lambda}{d} - \sin\gamma\right)^2}\right\} \tag{4.39}$$

$$FOD = \frac{\partial^4 \phi}{\partial \omega^4} = \frac{3\lambda^3}{4\pi^2 c^2 d} \frac{\partial^2 \phi}{\partial \omega^2} \left\{ \frac{\sin\gamma + \frac{2\left(1-\sin^2\gamma+\frac{\lambda}{d}\sin\gamma\right)\left(\frac{\lambda}{d}-\sin\gamma\right)}{1-\left(\frac{\lambda}{d}-\sin\gamma\right)^2}}{1-\left(\frac{\lambda}{d}-\sin\gamma\right)^2} \right\} \tag{4.40}$$

All three derivatives scale linearly with the grating separation, L_g. The ratio between the second and third orders can be tuned somewhat by changing the incident angle, γ. Due to the fact that gratings are generally designed with optimal diffraction efficiency at a particular value of γ, tuning the ratio between the second and third order dispersion involves a tradeoff with the power throughput of the stretcher.

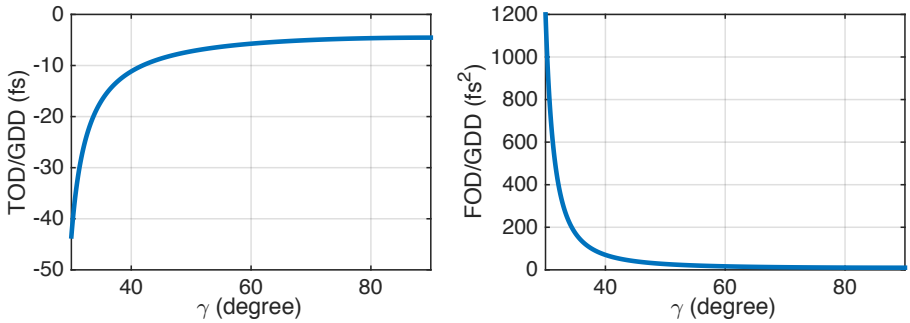

Figure 4.39: Ratios of third- and fourth-order dispersion to group delay dispersion as a function of the incidence angle, γ. These calculations were performed for 1550 nm light with the seed stretcher diffraction gratings. The groove density is 940 grooves/mm and the incidence angle with optimal efficiency is 47.5°.

We use transmission-mode diffraction gratings close to the incidence angle that gives optimal efficiency (47.5°). These gratings are largely polarization insensitive, so we are able to use a polarizer and quarter waveplate combination to extract the stretched pulse without vertical spatial chirp distortions. We typically use a grating separation of about 25 mm, corresponding to a stretched pulse duration of 50 ps. With 170 mW incident, we get 100 mW at the stretcher output. The losses come from the gratings and imperfect polarization rotation over the seed bandwidth by the quarter waveplate.

The advantages of this grating stretcher design are that the amount of stretch is easily tunable by changing the grating separation, that it is relatively efficient, and that the resulting beam is

undistorted, allowing for optimal coupling into the delay fiber. In particular, the tunability of the stretch is helpful for preliminary optimization of the OPCPA's performance because the seed's temporal stretch affects both the bandwidth and amplification efficiency of the OPCPA idler.

On the other hand, this stretcher design has one major drawback: the TOD cannot be compensated by the idler compressor. As discussed in Chapter 3 of this thesis, when seeding with 1.6 μm, the OPCPA idler will inherit the opposite sign of even orders of dispersion and the same sign of odd orders. Therefore, a grating compressor which gives negative GDD will be able to compensate for the GDD of a grating stretcher, but not the TOD. The TOD of both the stretcher and compressor will add, resulting in a limit to the compressibility of the idler. Though tolerable for a first attempt at OPCPA idler compression, this limitation must be addressed in order to compress the 3 μm light below 100 fs.

Stretcher Alternative: Highly Nonlinear Fiber One alternative stretching method is to simply couple the light into a Highly Nonlinear Fiber (HNLF). The light will undergo nonlinear processes such as self-phase modulation and self-steepening which will both stretch it in time and broaden its spectral bandwidth as it propagates through the fiber. The extent of the broadening and stretching the light experiences depends on its initial intensity and the length of fiber.

This method can be useful when seed pulse energy is low (100's of picojoules) and lacking in bandwidth. This was the case before the Y-Fi fiber laser became available to us. At that time, we used a ANDi oscillator to pump and OPO which generated about 20 mW of 1.6 μm light at 61 MHz (330 pJ per pulse). The bandwidth of the OPO output was typically around 30 nm. When coupled into a highly nonlinear fiber, the bandwidth expands to more than 150 nm over the length of a 1 m fiber. The measured spectra before and after the fiber are shown in Figure 4.40 a).

To find out the corresponding temporal stretch, we simulated the pulse's propagation through the fiber using a symmetrized split-step Fourier propagation algorithm including Self Phase Modulation (SPM) and self-steepening nonlinearities. The results of the simulation assuming 20 mW at the input are shown in Figure 4.40. It has reasonably good agreement with the measured spectrum, accurately predicting the spectral broadening and the number of lobes in the spectrum, though the

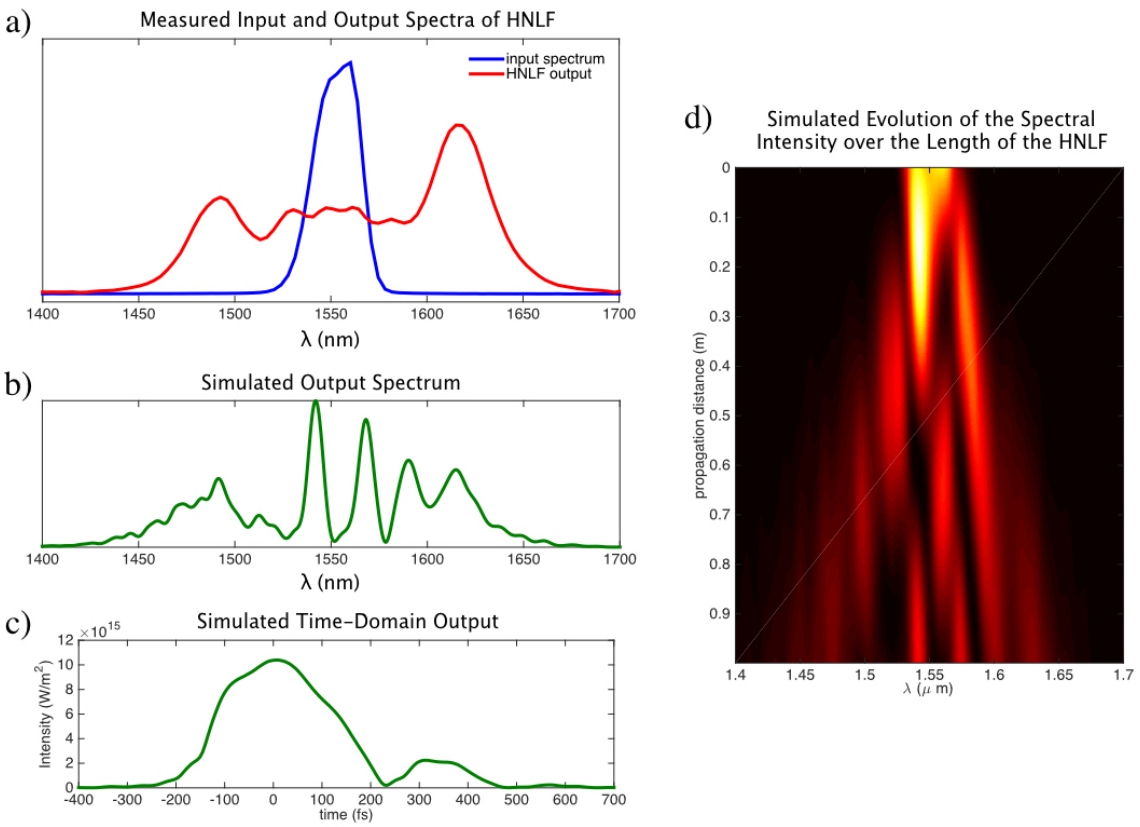

Figure 4.40: Spectral broadening in a highly nonlinear fiber with 20 mW input at 60.8 MHz. a) The measured spectrum at the input and output of the fiber. b) Spectral intensity at the output of the HNLF as predicted using a split-step Fourier propagation simulation. c) Time-domain output according the simulation. d) Simulated spectral intensity as a function of propagation distance through the HNLF.

modulations it shows are more severe than what we measure. It also tells us that the stretched pulse duration is only about 300 fs.

This points to the first major problem with the HNLF stretcher: that its ability to temporally stretch is outmatched by the spectral broadening it causes. A 300 fs pulse will be amplified in the OPCPA, but the efficiency of this amplification will be poor. Only a small fraction of the pump pulse (which is stretched to 120 ps) will overlap in time with the seed, so only a small fraction of the pump pulse has potential to amplify the seed. If the seed is stretched further by lengthening the HNLF, the spectrum will also broaden further, losing power to the spectral regions outside of the OPCPA crystals' phase matching bandwidth.

Additionally, an HNLF stretcher does not readily produce a seed with stable spectral phase. The output of the HNLF is highly sensitive to the input, often leading to a rapidly shifting spectrum at the output.

Finally, the HNLF stretcher is not easily scalable to higher seed pulse energies. At 34 nJ (the current pulse energy from the OPA), self-phase modulation would occur even in a few mm-long section of single mode fiber without enhanced nonlinear properties. We found that pulses which experience strong self-phase modulation typically have shot-to-shot spectral and amplitude jitter. Therefore, though attractive for its ability to expand the available bandwidth of the seed, the HNLF stretcher is not feasible as part of a robust front end for the OPCPA laser in its current design.

Stretcher Alternative: Chirped Fiber Bragg Grating This second alternative stretcher design is similar in concept to the volume Bragg grating used to stretch the 1 μm pulses. Like a CVBG, a Chirped Fiber Bragg Grating (CFBG) uses layered patterns to reflect the input light, imparting a temporal stretch. A CFBG accomplishes this over the length of a fiber rather than in bulk material. With the extended length of a fiber, the CFBG is capable of handling bandwidth between 1440 and 1640 nm and the imparted dispersion can be tailored to a specified value.

We chose the dispersion to give a stretch equivalent to 60 ps for a compressed 70 nm bandwidth input pulse. We also chose the GDD, TOD, and FOD to all be negative and have relative magnitudes matching that of our 3 μm grating compressor design: GDD of -1.13 ps^2, TOD of

-7.47^{-3} ps^3, and FOD of -7.60$\times$10^{-5} ps^4. This solves the problem of compounding TOD that is faced by the grating stretcher design. In contrast, the CFBG stretcher allows for transform-limited compression of the resulting OPCPA idler pulses.

The disadvantages of the CFBG are that it is not flexible, it has only ten percent throughput efficiency, and it is fiber based so that the amount of energy that can be coupled into it is limited to a few nanojoules by the onset of self-phase modulation. The flexibility issue comes from the fact that the dispersion the CFBG imparts is written in the structure of the fiber – once it is designed and built, there is no way to change it. If the dispersion of the compressor and any phase accumulated in the OPCPA stages is well known and constant, however, this inflexibility is not a problem. Fine tuning of the compression can be accomplished by changing the separation and angle of the compressor gratings, compensating for small inaccuracies in the second and third orders of the CFBG's dispersion.

The low efficiency of the CFBG stretcher is a bigger problem, especially when coupled with the input energy limit enforced by self-phase modulation. With 5 nJ incident, only 500 pJ will be left to seed the OPCPA. This is enough to allow OPCPA operation, but results in lower gain and reduced final idler energy compared to operation with tens of nanojoules of seed energy. Since maximum idler energy is necessary for the usefulness of this laser for HHG, this low throughput is a significant problem for the CFBG stretcher design and must be addressed if this will be used in the future.

Stretcher Alternative: Grating Stretcher using OPA Idler Another alternative stretcher design makes use of the OPA idler rather than the signal. The idler power is roughly half that of the signal (85 mW, 17 nJ for a 170 mW signal), so it is possible to use this beam rather than the signal to seed the OPCPA. Seeding with the idler has a few notable advantages. First, the three stages of OPCPA will automatically be more efficient, building on some light at 3 μm that is already there rather than starting from zero. Second, the OPA idler is CEP stable, so the OPCPA idler would inherit this feature from its seed. Third, the idler compression would be more straight-forward, allowing for third order dispersion compensation using a standard positive

dispersion grating stretcher and negative dispersion grating compressor. This type of setup is what is routinely used in CPA lasers. A positive dispersion grating stretcher is similar to a negative-dispersion stretcher except for the inclusion of an additional lens or curved mirror which effectively flips the output so that the red side leads.

The drawbacks of using the idler for the OPCPA seed are technical rather than fundamental. First, a 3 μm beam at the 100 mW level is much harder to "see" than a 1.6 μm beam of the same power. The reason for this is that most detector cards and cameras do not work at 3 μm and/or are not sensitive enough to detect this amount of light. Liquid crystal thermal detection cards are an exception, but these do not yield much information about the beam shape. Therefore, aligning a grating stretcher for an 85 mW, 3 μm beam is difficult. Second, the 3 μm beam cannot be delayed using a standard silica-based single-mode fiber in the same way that our 1.6 μm beam is on its way to the first stage OPCPA. The reason is that silica-based fibers and not transmissive beyond 2.3 μm. We would therefore require a different method of setting the timing between the OPCPA seed and pump so that they overlap in the crystals. Fluoride-based single mode fibers are available and could be used for this purpose, though they would lead to about 25% loss in power over a 4 meter length. These problems are by no means insurmountable, and may be worth addressing in future work.

Stretcher Alternative: Grism Another alternative stretcher design makes use of dispersion due to refraction in prisms in addition to diffraction from gratings. This type of stretcher, made up of prisms and diffraction gratings is called a grism stretcher. At the cost of some complexity of design, the advantage of a grism stretcher is that it allows for the ratio of TOD to GDD to be either positive or negative depending on the arrangement. Therefore, we could build a grism stretcher for the OPA signal with third-order phase that can be cancelled out by an OPCPA idler compressor consisting of a grating pair. [100]

We have not implemented this type of stretcher at present, but have plans to do so in the future. One potential issue with this design is relatively high losses due to multiple Fresnel reflections from the prism surfaces as well as imperfect diffraction efficiency of the gratings. Losses

from grism stretchers are often greater than 70 percent. Even so, because this stretcher design uses bulk optics rather than fiber, the throughput is not fundamentally limited by nonlinear effects like the CFBG or HNLF stretcher designs. This loss, though high, may be tolerable, still allowing optimal OPCPA operation.

Stretcher Alternative: Pulse Shaper Finally, one can construct a pulse-shaping stretcher by including an electronically controllable phase element. The advantage of this design is the ability to add arbitrary phase to the OPCPA seed to optimize the idler pulse compression. We have designed such a stretcher based on a liquid crystal Spatial Light Modulator (SLM).

The arrangement of our SLM-based stretcher is shown in Figure 4.41. This stretcher uses a grating to disperse the incident light before it hits a lens that focuses each spectral component onto a reflective SLM. The phase imparted by the SLM is controlled through a computer interface which varies the voltage applied to the SLM's pixels. Since the SLM is in the Fourier plane of the stretcher, the phase it imparts will be written onto the spectral phase of the output pulse. With the SLM off, this stretcher can be set to either impart negative or positive GDD. The GDD value is determined by the distance between the grating and the lens: negative GDD is obtained when L_g is greater than zero, and positive GDD when it is less than zero. We plan to use negative GDD similar to that currently given by our grating pair stretcher, using the SLM to adjust the higher-orders of spectral phase.

This design has not yet been tested in our OPCPA laser but is promising in that it allows compensation for both the compounded TOD and any other high-order phase that may be accumulated in the OPCPA. It has the advantage of tunability so that pulse stretch can be used as a knob to optimize the OPCPA performance and is able to handle the full power of the OPA signal without nonlinear effects. Also, it is more flexible and easier to align than a grism stretcher.

Summary In summary, the results reported in this thesis reflect the use of a negative-GDD grating pair stretcher. It is used for its tunable GDD and ability to handle the 34 nJ pulse energy of the OPA signal without nonlinear effects. However, it has the major disadvantage of imparting TOD on the stretched light that is compounded in the compressor, preventing compression to the

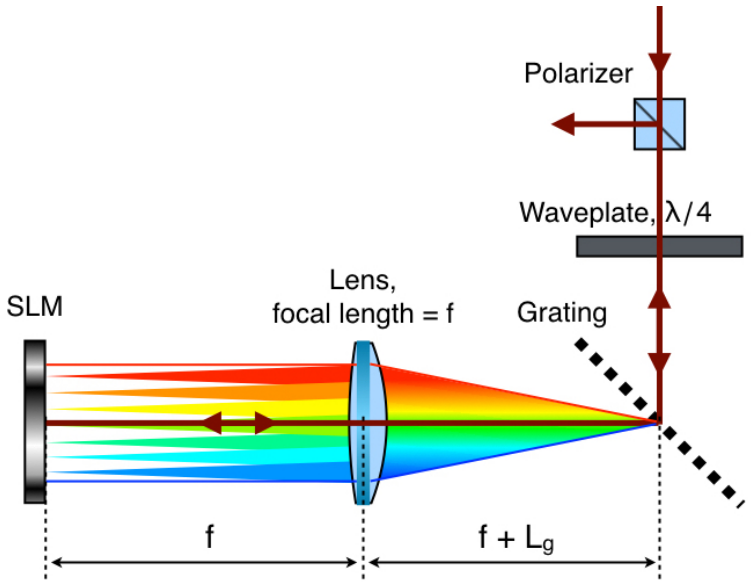

Figure 4.41: Diagram of a pulse stretcher which incorporates a spatial light modulator at its Fourier plane.

transform-limit of the OPCPA idler's bandwidth. In the future, we have two alternative plans that may be able to solve this issue. The first is the implementation of a grism stretcher which can impart the correct sign of TOD. The second is a stretcher that incorporates a spatial light modulator at the Fourier plane of the compressor, allowing active addition of arbitrary phase to the OPCPA's seed.

4.5 Parametric Amplification Stages

In this section, we will give a detailed description of the OPCPA stages themselves. We will first look at the architecture of the three stages and explain the reasoning behind our design choices. Next, we will discuss optimization of the OPCPA output by varying the pulse duration of the seed. Finally, we will review the system's typical performance.

4.5.1 Architecture

A diagram of the OPCPA itself is shown in Figure 4.42. It consists of three amplification stages designed to result in the highest possible idler pulse energy while maintaining the bandwidth needed for compression below 100 fs. Each stage is designed to handle progressively more power. The first stage has the smallest beam diameters so as to optimize gain and the next two stages use progressively larger beams to accommodate higher pump power without damage.

Timing In order for amplification to occur, the pump and seed pulses must be overlapped in time and space in each of the three OPCPA stages. The repetition rates of the pump and seed are passively locked due to their shared source, so timing can be matched by simply adding or subtracting beam path length from the seed or pump. Since the seed has nanojoule-level energy in stretched pulses, this path length control is easily attained using single-mode fiber as a delay line. The ability to make fine timing adjustments in each OPCPA stage is provided by micrometer-controlled translation stages placed before each one.

The single mode delay fiber also has the secondary advantage that it spatially filters the OPCPA seed. This ensures that the seed is delivered to the OPCPA breadboard with a nearly

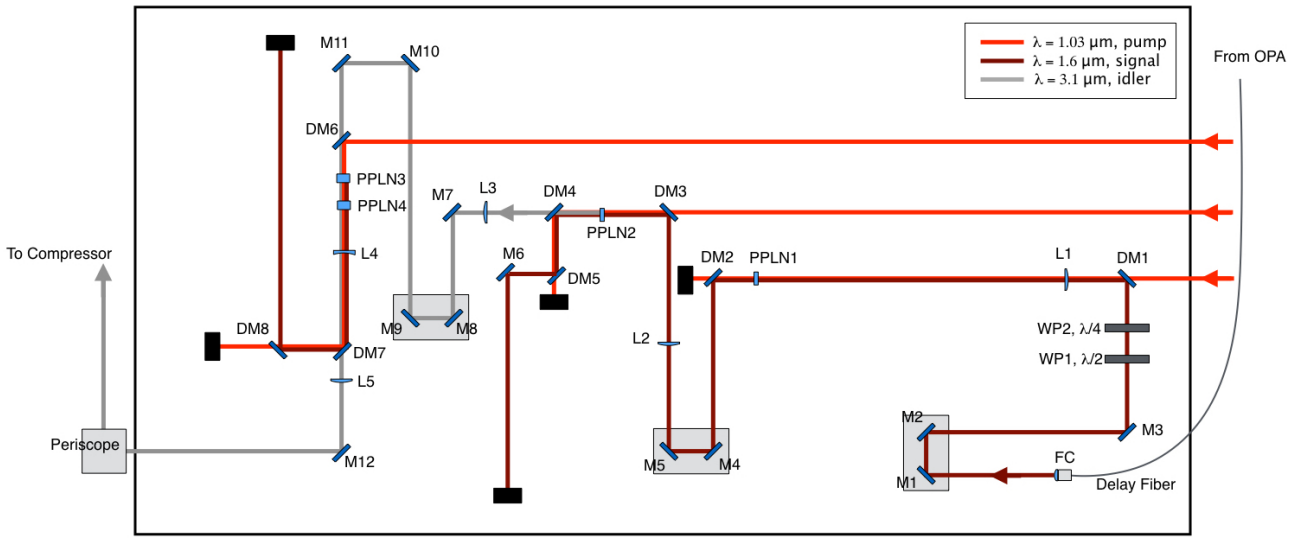

Figure 4.42: Diagram of the three-stage OPCPA. The 1.6 μm seed is delivered to the first stage from the OPA through a delay fiber shown on the right of the diagram. The first stage of amplification takes place in the crystal labeled PPLN1 where the seed is met by the pump. The amplified signal from the first stage seeds the second (PPLN2). The idler from the second stage then seeds the third and final stage. The third stage OPCPA uses two crystals (PPLN3 and PPLN4) to amplify the idler. Finally, the idler is separated from the pump and signal for compression. In this diagram, mirrors are labeled M, fiber collimator, FC, waveplates, WP, dichroic mirrors, DM, lenses, L, and MgO:PPLN crystals, PPLN. The red lines represent the pump beams at 1.0 μm, the maroon lines represent the 1.6 μm seed and signal, and the grey lines represent the 3 μm idler.

perfect TEM00 mode. Additionally, the fiber always delivers the seed to the same spot on the OPCPA breadboard so that the alignment of the OPCPA stages are insensitive to any alignment changes in the OPA, improving the robustness of the design.

Crystals We used MgO:PPLN crystals in each stage of amplification. The first stage crystal is a 1-mm aperture, 3 mm-long, single-period crystal with poling period of 30.1 μm, operated at 80°C. The second stage also used a 3 mm-long, single-period crystal with poling period of 30.1 μm at 80°C. The second stage crystal differs from the first in that it has a 3 mm aperture to accommodate a larger beam size. The third stage uses two 3 mm aperture, 10 mm-long chirped PPLN crystals, both operated at 100°C. In this case, the chirp refers to a varied poling period across the length of the crystal. These crystals both have period varying between 28.1 and 30.4 μm. The varied poling period increases the acceptance bandwidth of these crystals, though it also lowers the gain per crystal length. [61, 62] To compensate for the decreased gain, we increased the interaction length between the pump and seed by using two long crystals. The 3 mm aperture of these crystals is, to our knowledge, the largest commercially available aperture at this time.

Pump In each stage of amplification, the pump power is set so that the intensity reaches about 5×10^9 W/cm^2. In the first and second stages, the size of the pump spot is determined by balancing high gain with the need to keep the intensity below the damage threshold of the MgO:PPLN crystal. A smaller pump spot will yield higher gain per amount of pump power, but a slightly larger spot will allow for more pump power to be used. In the first stage, we found that a spot size of 0.6 mm (FWHM diameter) was optimal for the typical 6 nJ of seed energy received from the OPA. At this intensity, one can see a halo of faint, visible pink light starting to form around the beam after the crystal. This indicates that the intensity is approaching the point where damage could occur. In the second stage, the optimum beam size was found to be about 1.1 mm FWHM diameter. The third stage beam diameter is limited by the aperture of the MgO:PPLN crystal (3 mm). We fill the third stage crystals to the point where the edges of the pump beam are just clipping the crystal aperture, giving a FWHM beam diameter of 1.6 mm. Figure 4.43 shows images of the pump mode at the crystal plane for each stage.

The mode shape of the pump at each stage is also important for amplification. These OPCPA stages have high overall gain, so the pump profile will be written onto the modes of the amplified beams where they overlap with the seed. Also, optimal amplification efficiency requires good mode matching between the pump and seed. From Figure 4.43, we can see that the pump modes we used are all fairly round, though somewhat oblong (a characteristic inherited from the output of the fourth stage Yb:YAG amplifier). The first stage uses a lens to focus the pump and seed at a point behind the crystal. In the second stage, the pump beam is diverging as it interacts with the seed in the crystal. Also, the mode in the second stage is approaching a flat-top image of the beam in the fourth stage Yb:YAG crystal. We found that this results in efficient amplification of the signal and idler which take on the shape of the second stage pump. In the third stage, the beam is nearly gaussian and collimated so that it can maintain mode matching over the length of the two third stage crystals. The amplified idler emerges with a similar gaussian shape.

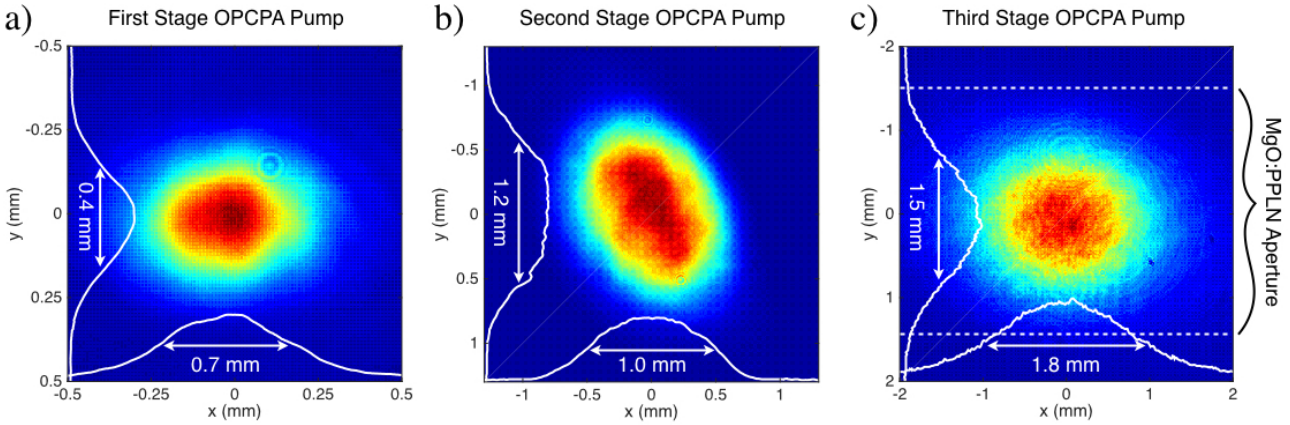

Figure 4.43: Mode images of the OPCPA pump at the crystal planes of each amplification stage. a) The first stage pump mode is shown and is converging at this point. b) The second stage pump mode is shown. It is a diverging beam with a nearly flat top. c) The third stage pump beam is shown. It is collimated and fills the crystal aperture. The white dashed lines indicate the edges of the third stage MgO:PPLN crystal.

4.5.2 Effect of Seed Stretch on Bandwidth

In what follows, we will devote some attention to the key factor that makes an OPCPA system different from an OPA: the stretch of the seed and pump. The relative pulse durations of the pump and seed are known to affect both amplified bandwidth and amplification efficiency. A seed pulse with temporal duration that is short compared to the pump will have all of its bandwidth amplified uniformly, but overlaps with only a small fraction of the pump pulse in the temporal domain. This will lead to optimal bandwidth amplification, but poor efficiency. On the other hand, if the seed is longer than the pump, it will be amplified efficiently, but its wings will be amplified less than its center, or not at all. Since a temporal stretch corresponds to a spread of the pulse's spectral components, this leads to a narrowed bandwidth in the amplified pulse. [101, 102]

The optimum stretch relationship between the seed and pump of an OPCPA stage depend on the specific value of the gain present. To avoid the added complication of applying a different stretch for each of the three stages in our OPCPA system, we simply tuned the stretch of the seed to optimize the bandwidth and power of the output. This was accomplished by varying the separation of the stretcher gratings between 7.5 and 46 mm, corresponding to seed pulse durations between 14 ps and 86 ps for a 70 nm-bandwidth pulse. The pump pulse duration was fixed at 120 ps.

We noticed a marked decrease in the amplified signal's bandwidth as we increased the seed's stretch. To isolate this effect in a controlled way, we measured the spectrum of the first stage's signal as a function of seed stretch. Over the course of this measurement, the seed spectrum and power were constant as were the pump power and relative timing between the pump and seed. The results are shown in Figure 4.44. The entire seed bandwidth (measured at full width and half maximum) was amplified with a 25 mm stretcher grating separation, corresponding to a 47 ps pulse. The amplified bandwidth decreased to about 80 percent of the original when the stretch was increased to 58 ps, and to just over 60 percent of the original when the stretch was increased further to 86 ps.

To see whether this trend was borne out in other situations, we put together a cumulative

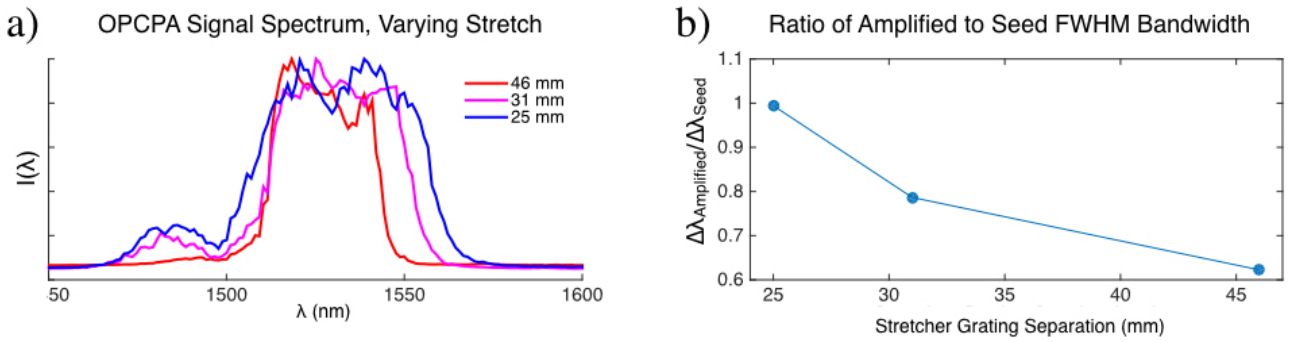

Figure 4.44: This dataset shows the effect of the seed stretch on the first stage OPCPA signal bandwidth. This measurement was performed with the same seed spectrum, seed power, pump power, and relative timing between the seed and pump in each case.

dataset consisting of other OPCPA spectra. These data were collected during the process of optimizing and exploring the parameter space of the OPCPA system, so there were variations in seed spectrum used in each case as well as the power for each amplification stage. These variations are expected to effect the optimum stretch for the seed. Even with these variations, however, we observe a relationship between the stretch and amplified bandwidth in the first stage OPCPA. Figure 4.45 a) and b) show these results plotted as a function of stretcher grating separation.

In spite of some variations due to the specifics of amplification, the general trend of better amplified bandwidth with smaller stretch is borne out. We see that the amplified bandwidth can be increased above that of the seed for small stretch factors. In one case, when the stretch was about 19 ps (about 10 percent of the pulse duration of the pump), the amplified signal's bandwidth was twice that of the seed. Also, the 86 ps stretch (46 mm grating separation), consistently results in amplified bandwidth smaller than 80 percent that of the seed.

We also looked at the final idler power and bandwidth as a function of seed stretch. Figure 4.45 c) shows the idler power for the cumulative dataset. This data shows no clear trend, indicating that the specifics of the amplification in the three OPCPA stages such as seed spectrum, seed power, and pump power for each stage are more important to the final idler power than the seed stretch alone. However, we do see a significant trend in idler bandwidth as a function of seed stretch, as illustrated by Figure 4.45 d). Our largest idler bandwidth of 240 nm was attained using 10 mm stretcher grating separation (19 ps seed pulses). As the seed stretch was increased to 86 ps, the idler bandwidth shrank to 100 nm. Though the specifics of OPCPA amplification also have a strong effect on the idler bandwidth, the seed stretch appears to be the dominant effect, at least in this limited data set.

4.5.3 Parasitic Processes and Noise Amplification

One of the advantages of OPCPA is its ability to give high gain in a single pass through a nonlinear crystal – in our system the gain can be as much as 2×10^4 in the first stage. Though generally a desirable trait, high gain also opens the door for amplification of Optical Parametric

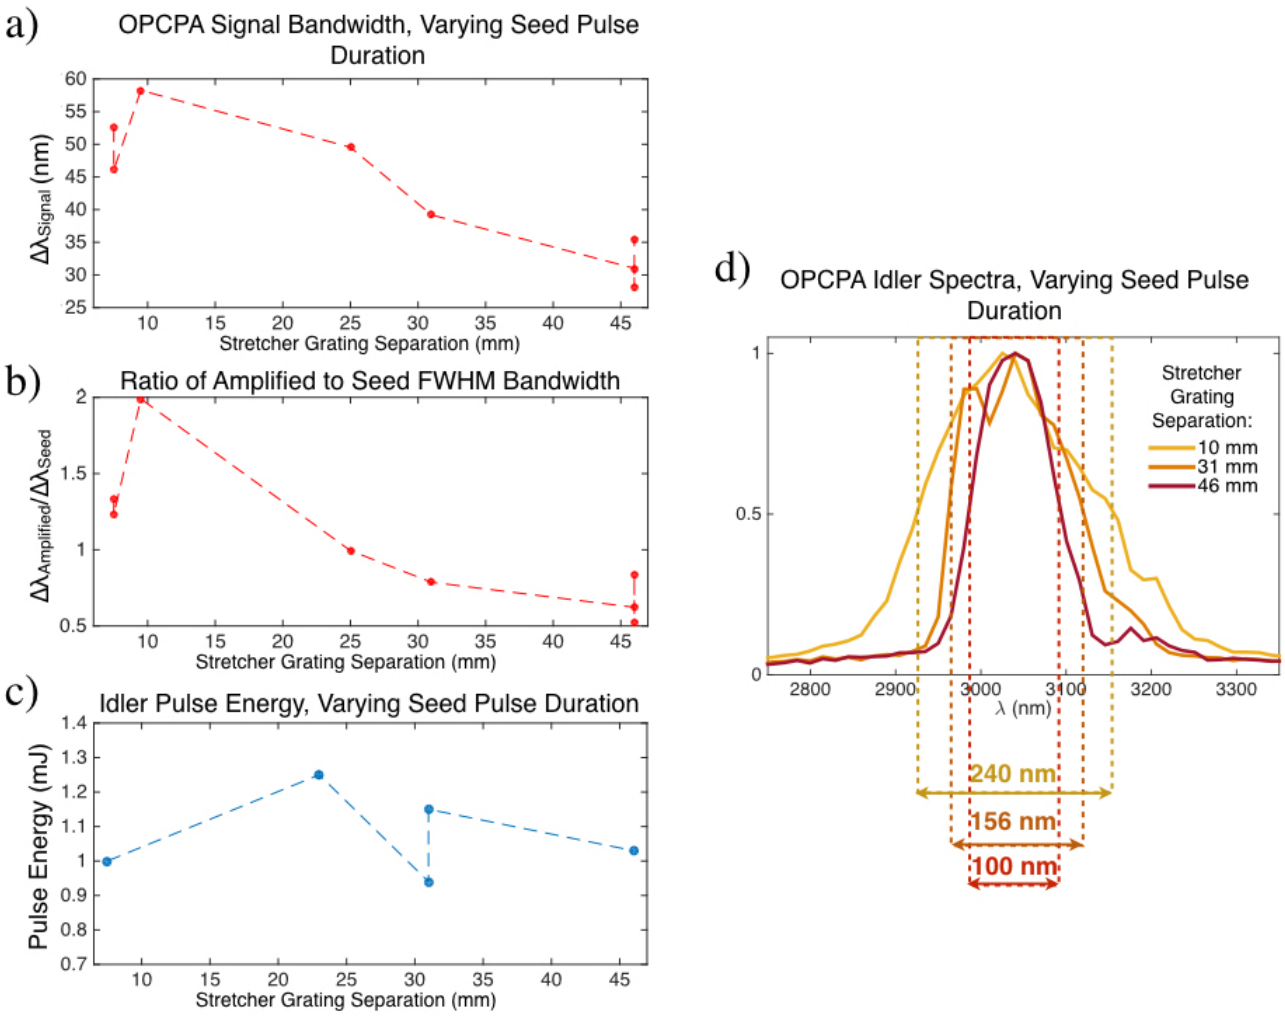

Figure 4.45: These plots show the effect of seed stretch on this first stage signal bandwidth in a) and b), on the final idler power in c), and on the idler bandwidth in d). These data points reflect a cumulative view of the OPCPA performance. The seed spectrum and power, the pump power for each stage, and the relative timings of the pump and seed were variable.

Fluorescence (OPF, or superfluorescence) and parasitic amplification of colors other than the signal and idler. These can cause problems such as reduced pulse contrast and loss of power. In this section, we will discuss the parameters of the OPCPA system which affect these processes, the extent to which they are present in our design, and how we address them.

Superfluorescence Superfluorescence is the result of quantum noise being amplified in the presence of the pump beam. A spontaneously generated photon at the signal or idler wavelength can seed the difference frequency generation process, producing a fluorescent beam at the signal and idler wavelengths that temporally overlaps with the stretched pump and can be further amplified in following stages. This light is different from the seeded signal and idler in that it cannot be temporally compressed since it is generated from random events rather than a coherent seed beam. Also, the fluorescent photons are emitted at random angles and wavelengths so they are not dependent on direction of the crystal for phase matching as the seed beam photons are. If superfluorescence is amplified to a significant fraction of the OPCPA signal or idler wavelength, it will result in a large pedestal surrounding the compressed pulse. The reduced contrast ratio of the pulse can be a problem for high harmonic generation because the pulse's pedestal may increase ionization levels in the gas before the main pulse arrives to create high-energy photons.

Fortunately, we do not see significant superfluorescence in the current configuration of our OPCPA system. The primary reason for this is that we have substantial seed energy of 6 nJ. This seed beam is strong enough to suppress the amplification of spontaneously generated photons. This effect is called gain quenching and is well known to suppress superfluorescence for seed pulses of nanojoule-level energies. [103, 104] This suppression effect can be best seen by blocking the seed before it reaches the first OPCPA stage. When this is done, we often observe superfluorescent signal and idler beams which can have significant power – we have observed 500 mW, 0.5 mJ fluorescent signal beams, for example. The amplified superfluorescence can be tuned to have a slightly different spectrum than the seeded signal so that it is obvious whether or not this beam is still present when the seed is let through. With a 6 nJ seed, we observe that the contribution to the spectrum from superfluorescence disappears to below what we are able to detect with our spectrometer.

In the past, when using smaller pulse energies in the 10-100 pJ range, we were not able to automatically suppress the superfluorescence with the seed. In this case the OPF can still be minimized by tuning other parameters of the OPCPA. Adjustments which will reduce superfluorescnece include reducing the parametric gain by decreasing the pump intensity, changing the temporal overlap between the seed and pump so that the leading edge of the pump is amplifying the seed, and spatial filtering between the OPCPA stages. Additionally, two factors that can worsen the problem of superfluorescence are gain saturation in a high-gain OPCPA stage and angular misalignment of the nonlinear crystal so that the phase matching condition is not perfectly satisfied. [104] In practice, we were able to suppress OPF without decreasing the pump intensity. Instead, we make sure the crystals were well aligned, optimize the timing between the seed and pump, and spatially filter the beam between the first and second OPCPA stages.

Parasitic Processes Another side effect of the OPCPA's high gain is the amplification of unwanted light due to coincidentally phase matched processes in the PPLN crystal. Any combination of the photons present have some probability of being amplified. These processes include sum frequency generation (SFG) from the pump and signal, second harmonic generation (SHG) from the pump, SHG from the signal or idler, third harmonic generation (THG) of the signal, and many others. [60] The strength of the amplification depends primarily on whether the products of amplification can build up coherently over the length of the crystal. That is, whether the process in question is phase matched.

In a periodically poled material such as MgO:PPLN, we see evidence of many processes in the form of visible red, green, and purple light. Sum frequency generation between signal and pump, for example, is observed in the form of red visible light at about 620 nm. All of these processes are relatively weak, generating only a few milliwatts of power, with the exception of the second harmonic of the pump light. This beam shows up as a bright green beam with wavelength of 515 nm. Under certain circumstances, this parasitic green beam can be bright enough to deplete the pump and limit the amplification of the desired signal and idler beams. Additionally, the green light can be reabsorbed in the MgO:PPLN by two-photon absorption, leading to heating. This

heating can cause average power-related damage. [105]

At room temperature, the first-order quasi phase matched second harmonic generation process driven by 1030 nm requires a poling period of 6.3 μm. Since this is far from the 28-31 μm poling periods used in our OPCPA laser, at first glance one would expect the second harmonic generation process to be so poorly phase matched that the green light would not be able to build up coherently at all. This is not always the case, however, due to higher-order quasi phase matching (QPM). The fifth order QPM period of the second harmonic generation (SHG) process (five times that of the first order period) is in fact 31.5 μm at room temperature, close to that of our crystals.

In a fifth order QPM process, the first two coherence lengths result in zero average intensity of the second harmonic, but the next poling boundary will be crossed right in the center of the following coherence length. This corresponds to a maximum in the field magnitude and, instead of decreasing again, the change in poling direction causes the light to continue to increase for another half coherence length. Thus, the second harmonic will grow over the length of the crystal, however slowly compared to a first-order QPM process. There are no shortage of pump photons in the OPCPA, so the slow growth does not preclude significant power being siphoned off to the parasitic green light. The growth of the second harmonic over the length of a crystal with the correct period is illustrated in Figure 4.46 c).

If the poling period is slightly longer or shorter than 31.5 μm, the second harmonic buildup will be reversed after some distance, as illustrated by Figure 4.46 a), b), and d). However, it is worth noting that this analysis assumes room temperature operation and a perfectly collimated pump beam. In order to optimize phase matching of the DFG process, the temperature of the MgO:PPLN crystals are often tuned between room temperature and 200°C. The QPM period of the second harmonic is also tuned by this temperature change. At 100°C, for example, the period decreases to 6.17 μm compared to 6.3 μm at room temperature, corresponding to a fifth-order QPM period of 30.9 μm. At 200°C, the QPM period shrinks further to 5.97 μm, so that fifth-order QPM will be experienced by a MgO:PPLN crystal with 29.9 μm poling period. This tells us that, for a given crystal poling period, we need to be careful to avoid the temperature that corresponds

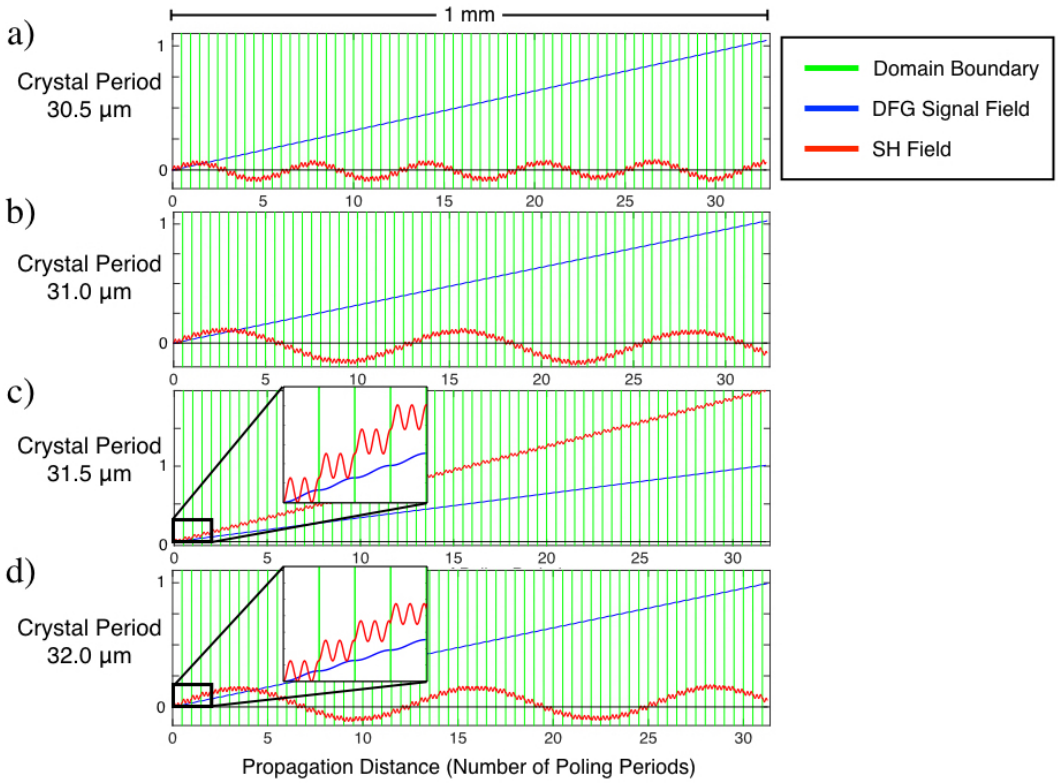

Figure 4.46: Illustration of the buildup of the 1030 nm pump's second harmonic (having a 6.3 μm coherence length) over the length of a 1 mm-long crystal with varied poling period at room temperature. The blue line shows the amplitude of the difference frequency generation process (which is first-order QPM for comparison). The y-axes of these plots represent the field amplitude of the light. The relative amplitudes of the red and blue lines are arbitrary and were chosen to be able to represent both on one plot. The green lines show the domain boundaries of the periodically poled crystal. a) Amplification in a crystal with 30.5 μm poling period. After building up for a short length, the Second Harmonic (SH) amplitude decreases again. b) Amplification in a crystal with 31.0 μm poling period, showing an increased length of SH buildup before the process reverses. c) Amplification in a crystal with 31.5 μm poling period, corresponding to exactly five times the first-order QPM period of the SH. The SH light is therefore fifth-order quasi phase matched and builds up in amplitude over the length of the crystal. d) Amplification in a crystal with 32.0 μm poling period, again showing a reverse in the amplification after a certain propagation distance.

to fifth-order QPM of the pump's second harmonic.

Another important consideration is the fact that the pump beam is not always collimated as it passes through the crystal. If it has some divergence, it therefore contains a spread of k-vectors. With each different propagation direction, the light experiences slightly different phase matching conditions. Therefore, even when a collimated beam would not result in significant SH buildup, we may see green rings around the center of the pump beam. Tight focal spots are particularly prone to causing this since they have a larger spread of k-vectors than a beam with a looser focus. In our experience, these green rings usually do not significantly affect the amplification of the desired signal and idler since most of the pump beam's light is not parasitized. However, if the conversion to green is particularly strong the signal and idler mode quality may be affected.

Finally, the presence of parasitic nonlinear processes in MgO:PPLN can be exacerbated by imperfections in the lengths of the poling periods which originate from the lithographic manufacturing process. These imperfections, known as Random Duty Cycle Errors (RDCE), can lead to more green buildup than would otherwise be expected for a given poling period and temperature. This effect is particularly pronounced for chirped crystals which have a range of poling periods. [105]

Fortunately, the current design of our OPCPA laser is not limited by pump deletion or thermal loading due to reabsorption of parasitic green light. In all configurations involving MgO:PPLN, we observe a bright parasitic green beam, but its power remains low compared to the signal and idler. The primary reason that the effect of parasitic conversion is minor in our setup appears to be that the crystal poling periods and temperatures that we used are comfortably far from the fifth-order QPM point for the pump's second harmonic in most of the crystals used. We use two single-period crystals with 30.1 μm period with temperatures between 70 and 100°C, and two chirped crystals with period between 28.1 and 30.4 μm at 100°C. A portion of each chirped crystal is expected to overlap with the fifth-order QPM period, but this occurs over a relatively small portion of its total length.

4.5.4 Results

In this section we will describe the typical output of the OPCPA in its current configuration including the pulse energies and spectra of the signal and idler, as well as the idler mode.

The first stage is seeded with about 30 mW of 1.6 μm light at 5 MHz which is met by a 1 kHz pump pulse train at 1.0 μm. The first stage is pumped by 500 mW (0.5 mJ) and the signal is amplified to 75 mW (75 μJ) at 1 kHz. This signal seeds the second stage OPCPA which is pumped by 4.0 W (4 mJ), generating 920 mW of signal and 460 mW of idler. Due to the fact that the second stage crystal is anti-reflection coated for the signal and pump, but not the idler, the idler is attenuated by a reflection at back face of the crystal. The measured idler power after the second stage is 270 mW (270 μJ). The second stage's idler is then used to seed the third stage. The third stage is pumped by 10 W (10 mJ) and generates 1.25 W of idler and 2.0 W of signal. The third stage crystals are both coated for the pump, signal, and idler so no significant reflection losses are measured.

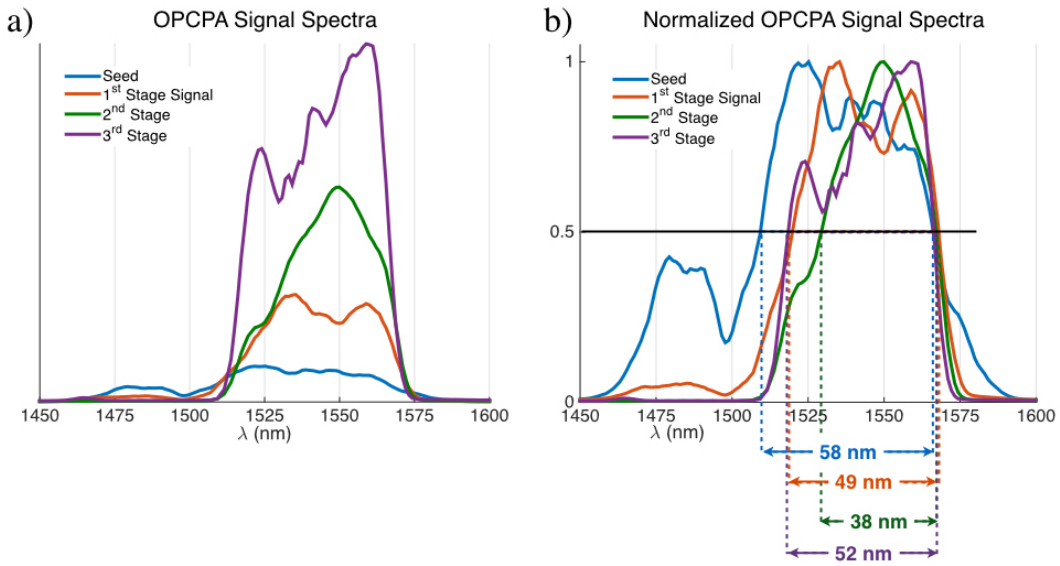

Figure 4.47: The seed and signal spectra from the first, second, and third stage OPCPA are plotted together. a) The spectra are plotted with relative amplitudes chosen to let the shape of each be visible. b) The same spectra are plotted with normalized amplitude and their FWHM bandwidths are labeled.

These results were obtained using a 23 mm stretcher grating separation, corresponding to a 31 ps pulse duration for a 50 nm signal bandwidth. The seed and signal spectra are shown in Figure 4.47. The seed spectrum has about 60 nm of bandwidth, measured at full width and half maximum. This is largely (but not entirely) maintained in the first stage, where the gain is 12,500. The bandwidth is further narrowed in the second stage which has a gain of 12.2 and, like the first stage, uses a single-period crystal. The bandwidth is increased to 52 nm in the third stage which uses two chirped crystals and has a gain of 4.6. The exact shape, bandwidth, and center wavelength of the signal spectra can be changed by adjusting the timing between the three stages. The signal spectra shown in Figure 4.47 are representative of spectra obtained when optimizing the final idler output power.

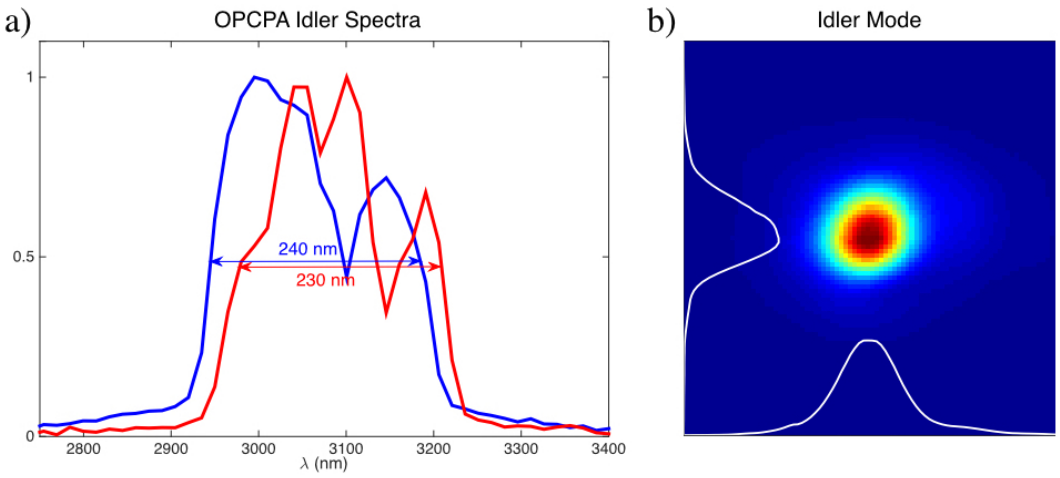

Figure 4.48: a) Two idler spectra are plotted. Their FWHM bandwidths are shown. b) An image of the idler mode in the far field is shown. The white lineouts show the beam profile through the peak of the beam in the x- and y-directions.

The idler spectrum corresponding to this signal is shown in blue in Figure 4.48 a). The other idler spectrum shown was obtained by tuning the relative timing of the three OPCPA stages to shift the idler toward the red. The spectrum can be changed further if the seed wavelength is altered and/or if the OPCPA crystals' temperatures are changed, affecting the phase matching.

The two idler spectra shown in Figure 4.48 a) have enough bandwidth to support 59 fs and

61 fs transform-limited gaussian pulses, respectively. This is well below our maximum desired pulse duration of 100 fs.

An image of the idler's beam profile in the far field is shown in Figure 4.48 b). Under conditions of good mode matching, the beam emerges as a focusable beam with near gaussian shape. The ratio of the x- to y-profiles of the beam is 0.94, showing a high degree of roundness and even an improvement compared to the pump beam emerging from the fourth Yb:YAG amplifier stage. The beam maintains a round shape as it is focused to within the resolution of our camera (80 μm). This focusability is important for high harmonic generation because the beam must be efficiently coupled into a gas-filled waveguide with an inner core diameter that may need to be as small as 50 to 100 μm depending on the available pulse energy and pulse duration.

4.6 Idler Compression

The final step in preparing the 3 μm light from the OPCPA for high harmonic generation is compressing the pulse to obtain maximum peak intensity. In this section, we will describe our compressor design, show some measurements of compressed pulses, and discuss the current limits to compression and potential solutions.

4.6.1 Compressor Design

The compressor is designed to be as efficient as possible in imparting negative group delay dispersion to the 3 μm idler beam. It uses two silver-coated, ruled, reflection-mode diffraction gratings. They have 300 grooves per millimeter and the optimal angle of incidence is 37° with respect to the grating normal. They are arranged as shown in Figure 4.49. A roof mirror and pickoff mirror are used to separate the outgoing beam from the incident beam. This simple compressor configuration imparts negative GDD with magnitude that is tunable by changing the grating separation. The GDD, TOD, and FOD can be calculated from Equations 4.38, 4.39, and 4.40. The amount of third order dispersion can also be tuned by changing the angle of incidence on the gratings, though it will always have a positive value.

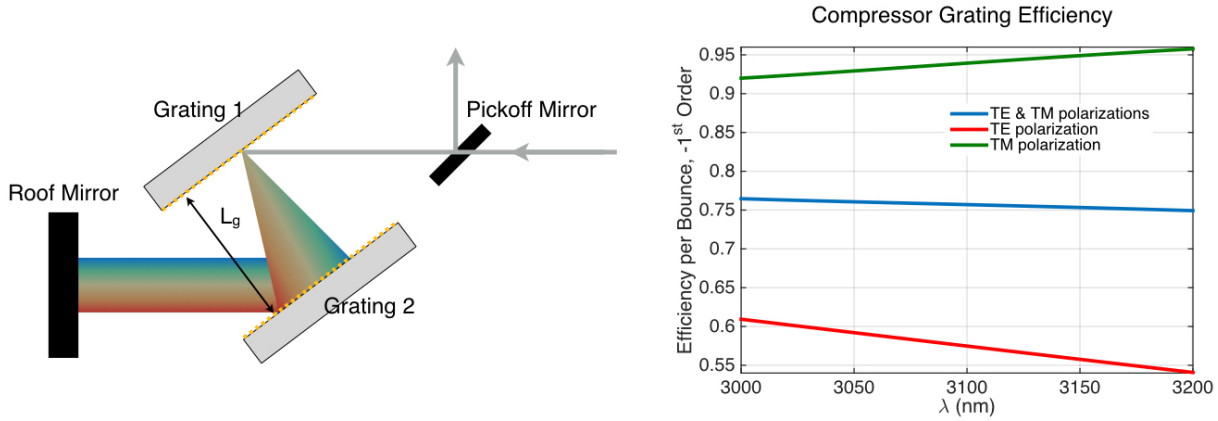

Figure 4.49: Left: Diagram of the idler compressor. Right: Theoretical grating efficiency curves showing the single-reflection efficiency performance of the compressor gratings at 37° incidence angle. Best performance is realized with TM polarization (polarization perpendicular to the direction of the grating grooves).

These gratings have wavelength, polarization, and incident angle-dependent efficiency, as plotted in Figure 4.49. At the designed angle of incidence (37° with respect to the grating normal), the theoretical average efficiency over the idler's wavelengths for a horizontally polarized beam should be about 94% per bounce. This corresponds to a total efficiency of 78%. The idler emerges from the OPCPA with vertical polarization, so we rotate the polarization by 90 degrees using a periscope before the compressor.

As discussed in the previous section about the OPCPA seed stretcher, this compressor will compensate for the GDD imparted by the stretcher, but the TOD from the current stretcher design and this compressor will compound. Table 4.4 shows the compounded TOD that will result for several different stretch factors. There will also be residual fourth order dispersion due to the fact that, though the stretcher and compressor FOD amounts will have opposite sign, they will not necessarily have the same magnitude. Table 4.4 also shows us that the FOD dispersion alone will cause pulse broadening that is significant compared to a transform-limited pulse but much less than that caused by the compounded TOD.

L_g, Stretcher (cm)	L_g, Compressor (cm)	Compounded TOD, (fs^3)	Residual FOD (fs^4)	Compressed Pulse Duration (fs)	Pulse Duration FOD only (fs)
0.0	0.0	0	0	59	59
1.0	3.3	4.95×10^6	7.37×10^6	448	84
1.5	4.9	7.43×10^6	1.10×10^7	518	91
1.0	6.6	9.91×10^6	1.47×10^7	571	97
2.5	8.2	1.24×10^7	1.84×10^7	618	102
3.0	9.9	1.49×10^7	2.21×10^7	657	106
4.0	13.3	1.99×10^7	2.94×10^7	726	114
5.0	16.5	2.48×10^7	3.68×10^7	782	121

Table 4.4: For a given stretcher grating separation, L_g, the compressor grating separation that cancels out the GDD of the stretcher is given. The resulting compounded third order dispersion and residual fourth order dispersion are also shown, assuming that the stretcher and compressor are operated at their incident angles of highest efficiency. Finally, the compressed idler pulse duration (assuming 240 nm FWHM bandwidth and 3.1 μm center wavelength) with these values of TOD and FOD (assuming GDD is zero) is given as is the pulse duration with only the residual FOD applied. A transform-limited (flat phase) pulse of this bandwidth will have a pulse duration of 59 fs.

From these calculations, we expect that TOD will prevent the OPCPA idler with 240 nm bandwidth from being compressed below about 450 fs using the current stretcher at its designed angle of incidence. The TOD can be mitigated slightly by increasing the angle of incidence as illustrated in Figure 4.39. This gain in compression comes at the expense of grating efficiency, however. We also expect that smaller stretch factors will allow better compression since the total TOD will be smaller. Also, with high TOD, idler pulses with narrower bandwidths turn out to be more compressible in spite of the fact that their transform limited pulse duration is larger. A 100 nm bandwidth pulse with 3.1 μm center wavelength, for example, is compressible down to 243 fs with the dispersion due to a 1.0 cm stretcher grating separation. Its transform limited pulse duration is 140 fs. On the other hand, a 240 nm bandwidth pulse of the same wavelength and stretch has a transform limited pulse duration of 60 fs, but is only compressible to 450 fs.

4.6.2 Compressed Pulses

In this thesis, we have aimed to amplify an idler with as large a spectral bandwidth as possible with an eye to the future when we will use a stretcher that does not result in compounding TOD. However, this does not result in the best pulse compression for the current stretcher and compressor designs. In what follows we will show the results of our preliminary idler pulse compression using a 150 nm bandwidth 3 μm pulse. With this narrowed bandwidth, we demonstrated pulse compression to less than 500 fs in spite of the high TOD.

To measure our compressed pulses, we use a Frequency Resolved Optical Gating (FROG) device with second harmonic generation of the idler as the optical gate. This technique collects second harmonic spectra generated by the crossing of two idler beams in a Gallium Selenide (GaSe) crystal at a range of delays between the two beams. The resulting dataset contains information both about the pulse's amplitude and phase, fully characterizing the pulse's electric field. We use a FROG device with a 30 μm-thick GaSe crystal. A diagram of the FROG device is shown in Figure 4.50 a).

We used the FROG measurement to determine the compressor grating separation yielding

optimal pulse compression, as shown in Figure 4.50 b). Without tuning the angles of the gratings with respect to the incident beam, the point of optimum pulse compression corresponds to the point where GDD is closest to zero. Comparing the measured compressor grating separation for minimum GDD to that expected from Equation 4.38, we found that the measured grating separations were slightly larger than those predicted. This reflects a positive GDD on the idler that is in addition to the dispersion from the stretcher. This extra GDD factor represents the effect of the OPCPA, including material dispersion. It does not appear to be a nonlinear effect because it is apparently stretch-independent (and therefore intensity-independent). We see this in that the magnitude of the extra stretch remains constant for a wide range of seed stretch factors. The chirp on the 1.0 μm pump beam also contributes to this extra GDD, but this contribution is very small since the pump is essentially monochromatic compared to the signal and idler. The amount of extra stretch is equivalent to 16 mm extra compressor grating separation, or 7 ps extra idler pulse duration.

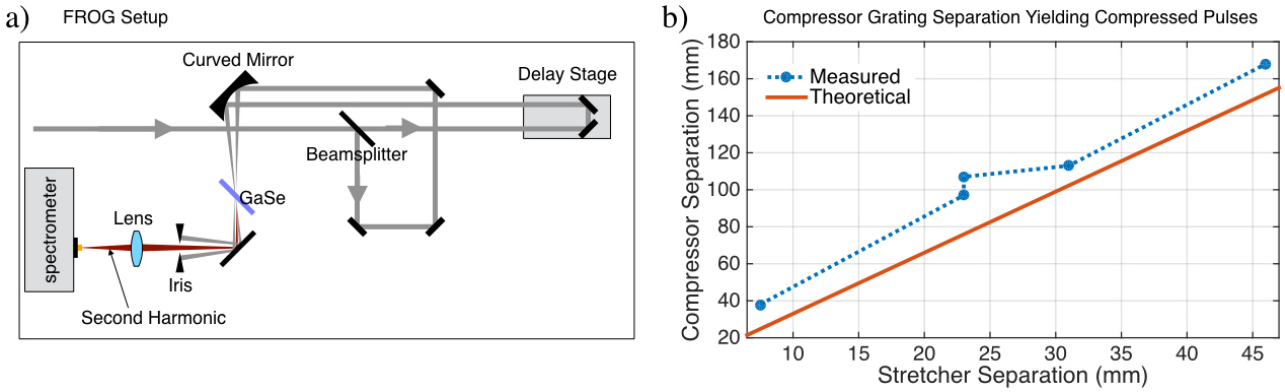

Figure 4.50: a) Diagram of our second harmonic FROG for measuring the pulse duration and phase of the compressed idler pulses. b) The compressor grating separation as a function of the stretcher grating separation. The measured grating separations show the same slope as those we theoretically expect, but have an additional constant offset.

The compressor's total efficiency was measured to be about 65 percent at the optimal incidence angle of 37°, corresponding to 90% of the intensity reflected per bounce. This is below the theoretical efficiency of 78% (94% per bounce). The reduced efficiency may be partially due to the polarization not being perfectly horizontal. This could happen either because of an impure

polarization state emerging from the final stage OPCPA or imperfect polarization rotation by the periscope (or both). Another possible source of the reduced efficiency is imperfect reflectivity of the silver grating surface. With 1.25 mJ of idler before the compressor, the measured compressor efficiency corresponds to a 0.81 mJ compressed pulse.

Finally, we used the FROG to measure the compressed pulse durations and residual phase. Figure 4.51 shows a FROG trace of the compressed idler (a) as well as the reconstructed trace (b), spectrum (d), temporal shape (c), and phase (c and d). The stretcher grating separation used was 46 mm and the idler's spectral bandwidth was 150 nm, corresponding to a 48 ps seed pulse duration. The retrieved idler pulse duration is 478 fs with a FROG error of 2.8 %. The FROG algorithm faithfully reconstructed the bandwidth of the measured spectrum of the OPCPA idler, though it added some modulations to the short-wavelength side that does not match the measured FROG trace or spectrum. In spite of this, the measured and reconstructed FROG traces show reasonable agreement. Also, the reconstructed spectral and temporal phases show strong third-order contributions that we expect with the current stretcher and compressor designs. The fact that the idler bandwidth is less than 240 nm allows compression below 750 fs (the expected limit to compression for a 240 nm bandwidth pulse of the same stretch factor) in the presence of compounded TOD.

These results clearly show that the compounded TOD must be eliminated or at least strongly mitigated to obtain pulses shorter than 100 fs. We have also shown that there is a non-negligible contribution to the total phase from the OPCPA itself in addition to the stretcher and compressor. These results point to the value of a stretcher with tunable dispersion, particularly tunable third- and higher-order dispersion components which are not easy to measure precisely. For this reason, the new stretcher design which incorporates a programmable SLM is particularly attractive. The grism stretcher design may also prove effective.

In summary, we have demonstrated preliminary pulse compression of the 3 μm light down to 478 fs and anticipate that a new stretching method will allow pulse compression to near the transform limit. We have generated 1.25 mJ pulses with 240 nm FWHM bandwidth. If optimally

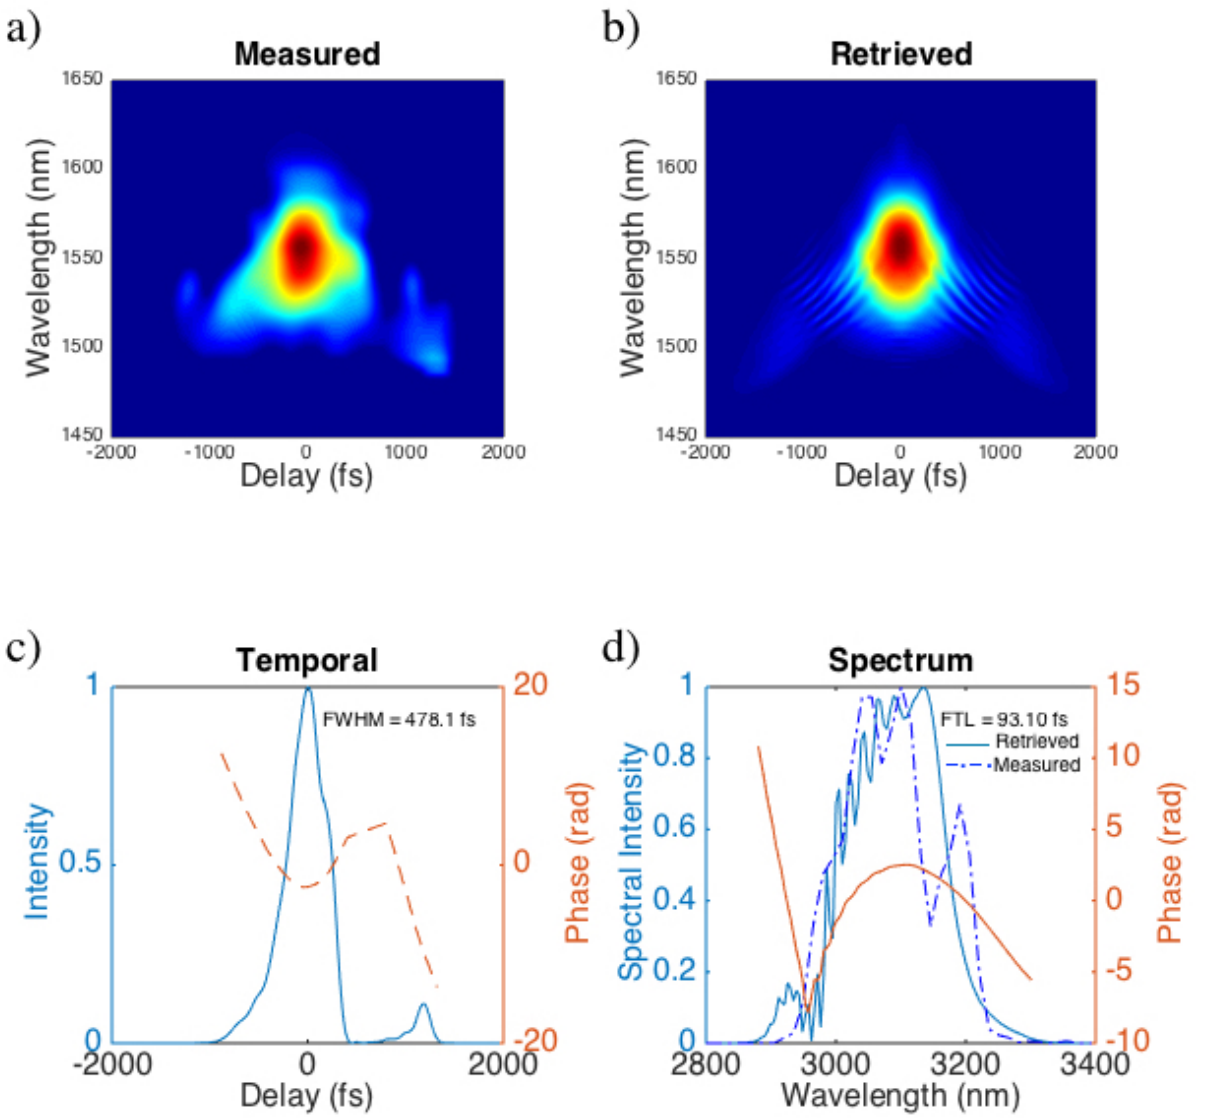

Figure 4.51: Measured and retrieved data from a second-harmonic FROG scan of the OPCPA idler. In this case, the idler had a center wavelength of 3.1 μm and a FWHM spectral bandwidth of 150 nm. a) The measured FROG trace of the OPCPA idler. b) The retrieved FROG trace obtained using a second-harmonic FROG reconstruction algorithm. c) The retrieved pulse's temporal profile (blue) and temporal phase (orange). d) The pulse's measured spectrum (dashed-dotted dark blue), retrieved spectrum (solid blue), and retrieved spectral phase (orange).

compressed with the current stretcher efficiency of 65%, this would yield 60 fs, 0.8 mJ pulses. For high harmonic generation in helium, this pulse energy and duration would yield a maximum photon energy of 1.0 keV with a 30 μm diameter focal spot.

Chapter 5

Conclusions and Future Work

5.1 Summary and Conclusion

In this thesis, we described the demonstration of the first (to our knowledge) mid-infrared OPCPA laser system capable of producing millijoule-level pulse energies at 1 kHz repetition rate. It has bandwidth supporting 60 fs pulse durations, making it a suitable source to drive high harmonic generation. We addressed peak and average-power scaling issues associated with the kilohertz, multi-milijoule regime while maintaining excellent beam quality in the pump, signal, and idler. This laser system is capable of operating stably for extended periods of time with low pulse-to-pulse amplitude or spectral jitter. The 3 μm beam is usable for high harmonic generation due to its stability, energy, potential for sub-100 fs pulse compression, and excellent mode. Furthermore, its kilohertz repetition rate will enable sufficient harmonic flux to be useful for application experiments such as time-resolved spectroscopy or high-resolution imaging.

In addition to its utility for keV high harmonic generation, this laser design is flexible enough to be used for various other applications. It has been demonstrated to be scalable to the high-repetition rate, low energy regime for molecular spectroscopy, for example. [97] It is worth noting that this multi-wavelength source provides not only broadband mid-infrared light, but also a broadband near-infrared beam at 1.6 μm and a high-energy narrowband beam at 1.0 μm. Simultaneous compression of the 1.6 and 3.1 μm beams could allow this laser to be used for applications such as two-color-driven high harmonic generation of circularly polarized light. This is useful for measurements of magnetic dichroism and has been demonstrated with the output of a Ti:sapphire-driven

OPA (providing wavelengths up to 2.1 μm). [106] The use of 3 μm light as one of the driving wavelengths has yet to be demonstrated. Alternatively, the 1.6 μm signal light could be used on its own for driving water-window high harmonic generation. This laser source may be advantageous over others in that the signal and idler wavelengths are tunable over a wide range. The signal wavelength is currently tunable from 1450 nm to 1600 nm and this range could be increased to span 1300 to 1700 nm by tuning the white light source and using optical coatings that have good reflectivity over that range.

In summary, this laser system is unique for its wavelength, pulse energy, and repetition rate. By nature of using nonlinear frequency mixing rather than laser amplification to generate the mid-infrared beam, it is fundamentally flexible in wavelength, spectral bandwidth, pulse duration, and repetition rate. One of the main challenges in reaching the high peak and average power regime lies in building a pump laser that has these properties while maintaining excellent beam quality. We have accomplished this in that our pump laser provides more than enough energy to pump the OPCPA in its current design at 1 kHz. Another challenge in OPCPA design is the added complexity of idler compression when the seed has a different wavelength. We have tested the limits of pulse compression with a simple negative dispersion stretcher design and got to within a factor of five of our desired pulse duration (we have a 48- rather than 10-cycle pulse). This result confirms the need for a different stretcher design.

In addition to reaching record pulse energy at this wavelength and repetition rate, the work reported in this thesis also represents progress in laser science toward the goal of developing an OPCPA system that is compact, robust, and scalable. The knowledge gained will help inform designs of other mid-infrared laser systems in the future.

5.2 Future Work

There are several areas in which this laser system can be improved in the near future as well as near-term possibilities to put it to work in performing experiments. In this section, we will discuss the next steps in improving its performance as well as our plans to demonstrate its utility

in generating high harmonics.

Pulse Compression The next step for this laser system as it is currently designed is to improve the compression of the mid-infrared light, reaching pulse durations less than 100 fs. This can be done by eliminating or at least mitigating the compounded third-order dispersion from the stretcher and compressor. We have two new stretcher designs which should be able to achieve this goal. One incorporates a liquid crystal spatial light modulator that can be programmed to give the desired sign and amount of third- and higher-order dispersion. The second is a grism stretcher – a combination of prisms and diffraction gratings which can be used to tune the third order dispersion to have the desired sign.

High Harmonic Generation Once the desired pulse compression is reached, this laser system will be ready to be used as a driving laser for high harmonic generation. In a first experiment, we could use this laser to generate water-window harmonics (300-550 eV). For a given laser intensity, we should see higher-energy harmonics from the 3 μm driving laser than from a shorter-wavelength source. From there, the parameters of HHG setup including the capillary diameter and coating, gas pressure, and propagation length can be optimized to reach toward the keV energy region. Once harmonics are generated, this source will provide a platform for experiments requiring high energy light and/or isolated attosecond pulses.

Idler Pulse Energy To aid the high harmonic generation process, we also have plans to further increase the 3 μm pulse energy, ideally to several millijoules. Currently, the idler pulse energy is limited by the 3 mm aperture of the third stage MgO:PPLN crystal. This is the largest aperture that is currently commercially available, but this may change in the near future as fabrication technology improves. An increase to a 4 mm aperture crystal would allow 18 mJ of pump power to be applied rather than 10 mJ. With the current pump-to-idler conversion efficiency, this would result in a 2.2 mJ idler beam (1.4 mJ after compression). By the same logic, an increase to a 5 mm aperture would allow 28 mJ of pump light, at which point the system would be limited by the pump laser energy rather than the crystal size.

A second method for effectively increasing the aperture of the PPLN crystal is to use an

oval-shaped pump beam. The aperture of the MgO:PPLN crystals is limited to 3 mm in only one dimension, so the pump mode area could be increased by expanding its profile in one dimension. The crystals we currently use in the third OPCPA stage measure 3 mm in height and 10 mm in width, so an elliptical pump beam could have a mode area equivalent to that possible using a 5.5 mm square crystal. The main challenge for this design is achieving and maintaining good mode matching between the pump and idler over the length of the crystal. In addition, the ellipticity of the beam must be corrected for after amplification so that the idler will be focusable to a round spot.

Alternatively, other nonlinear crystals can be used which are available with larger apertures. One of these is Potassium Titanyle Arsenate (KTA) which is commercially available with apertures greater than 6 mm and has a high damage threshold. One disadvantage of KTA is that its phase matching conditions require a slight angle between the pump and seed. This limits the distance over which the two beams can be overlapped without spatial walk-off and can lead to degradation of amplified beam quality if the phase matching conditions are not well satisfied over the whole diameter of the amplified beam. Certain regions of the amplified beam may be brighter than others and/or exhibit spatial chirp. In spite of these challenges, however, KTA may prove to be a good path forward, especially if the idler energy is large enough that HHG capillary coupling losses due to reduced beam quality are tolerable.

Pump Laser The 1.0 μm pump laser also has room for improvement in the future. Improvements could address several aspects of the laser's current operation: pulse energy, compactness and robustness, and repetition rate scaling. Though the pump pulse energy does not currently limit the OPCPA system performance, this may become a limiting factor if larger aperture PPLN crystals are obtained and/or the repetition rate is increased. Currently, the pulse energy is limited by thermal loading in the final two amplification stages, by low-level birefringence that degrades the beam quality when the number of passes is increased, and by the damage threshold of the optics which prevents the beam diameter from being reduced.

One strategy for pushing the pulse energy limits of the pump laser is to increase the pulse

stretch while the light is being amplified. This would reduce the peak intensity and therefore allow smaller beam diameters in each stage. Decreasing the beam diameter has several desirable consequences. First, the 940 nm pump beam diameter could also be scaled down, increasing the gain per pass without increasing the thermal load on the crystal. Also, the beam would pass through a smaller section of the laser crystal, making it less sensitive to birefringence at the crystal's edges. The requirement for low-stress mounting would therefore be relaxed, and higher-thermal conductivity mounting materials could be considered. Specifically, a copper mount might become usable in the fourth stage amplifier, increasing the thermal loading limit. Alternatively, decreased sensitivity to birefringence may allow the number of passes to be increased without significant beam degradation using to current low-stress CuW mount design.

Increasing the pulse stretch would add some complication to the system in that it would require recompression down to the 100-150 ps level before the OPCPA to maintain the current level of temporal overlap between the pump and seed. This compression would require diffraction gratings due to the high energy of the amplified beam, and, depending on the stretch factor used, may require a long propagation distance between gratings. This would necessitate careful alignment and control of air currents to maintain the current beam quality and directional beam stability. This approach would therefore be technically challenging, but may provide a reward in performance that justifies the effort.

Another strategy for increasing the pump pulse energy is to increase the thermal loading capacity of the third and/or fourth stage amplifiers so the pump power can be increased. This requires more effective heat removal from the crystal. The most conceptually simple way to accomplish this is by improving the thermal conductivity of the crystal mounts while maintaining their low-stress properties. This could be accomplished using a mount material such as sapphire or by shortening the path length between the crystal and cold head. The later strategy could be accomplished by redesigning the vacuum chambers to accommodate shorter mounts, by using a high thermal conductivity spacer between the mount and heat exchanger, or by circulating the coolant around the crystal within the mount itself. Alternatively, we may be able to reduce the

heat load by more effectively managing heating effects other than quantum defect heating, such as trapped fluorescence. Rather than a mount that encloses the crystal around its whole diameter, it may be advantageous to allow all but one side to be open so that fluorescence can escape.

We also have a few strategies for improving the pump laser's design for compactness and robust operation. The main problem of the current design is that there are four stages of amplification. The multipass amplifiers are simple in their design, but require significant table space to fit in the telescopes needed for up- and down-collimating the beam. In a four-stage amplifier with a large amount of path length, a small misalignment in the beginning is compounded in the later stages, making the system more difficult to operate on a consistent basis. To mitigate this issue, we are exploring options for replacing the regenerative amplifier with simpler lasers that can produce similar or higher energy pulses. One option is a high-gain rod amplifier which may be able to produce the same energy as the regenerative amplifier in a single pass through a large-aperture fiber-like rod. Another option is to use a thin disk Yb:YAG amplifier. Though the thin disk laser is not a significant improvement over the regenerative amplifier in complexity, single-stage thin disk amplifiers have recently been demonstrated to produce 220 mJ per pulse at 1 kHz. [90] This type of laser could reduce the number of amplification stages from four to one.

Finally, the pump laser could be improved by scaling to higher repetition rates. This would ultimately result in more high harmonic flux for experiments – a feature that would be particularly useful for high-resolution diffractive imaging. Since only 15 mJ are currently used to pump the OPCPA, some repetition rate scaling can be accomplished without much change in the amplifier design. This would lead to a compromise in pulse energy, but that is tolerable down to 15 mJ without compromising the idler pulse energy. As a first step, we should assess the extent to which we can scale the repetition rate before the pulse energy drops below 15 mJ. Beyond this, repetition rate scaling without pulse energy compromise requires more pump power. Currently, the pump power is limited by thermal loading in the third and fourth stage amplifiers. Breaking this barrier is yet another incentive to tackle the problem of thermal loading.

In conclusion, pulse compression below the 100 fs level is needed for this laser to drive high

harmonic generation. After that, further improvements to the laser system could go in several directions depending on the demands of the experiment to be performed. There is room for increased idler pulse energy as well as repetition rate scaling for experiments that require high flux. Improvements in compactness and robustness are also being planned which will bring this laser away from the state of a prototype, toward a more easily usable and robust scientific tool.

Bibliography

[1] Tenio Popmintchev, Ming-Chang Chen, Oren Cohen, Michael E. Grisham, Jorge J. Rocca, Margaret M. Murnane, and Henry C. Kapteyn. Extended phase matching of high harmonics driven by mid-infrared light. Optics Letters, 33(18):2128 – 2130, September 2008.

[2] Eiji J. Takahashi, Tsuneto Kanai, Kenichi L. Ishikawa, Yasuo Nabekawa, and Katsumi Midorikawa. Coherent water window x ray by phase-matched high-order harmonic generation in neutral media. Physical Review Letters, 101(253901), December 2008.

[3] Tenio Popmintchev, Ming-Chang Chen, Alon Bahabad, Michael Gerrity, Pavel Sidorenko, Oren Cohen, Ivan P. Christov, Margaret M. Murnane, and Henry C. Kapteyn. Phase matching of high harmonic generation in the soft and hard x-ray regions of the spectrum. PNAS, 106(26):10516–10521, June 2009.

[4] M.-C. Chen, P. Arpin, T. Popmintchev, M. Gerrity, B. Zhang, M. Seaberg, D. Popmintchev, M. M. Murnane, and H. C. Kapteyn. Bright, coherent, ultrafast soft x-ray harmonics spanning the water window from a tabletop light source. Physical Review Letters, 105(173901), October 2010.

[5] Tenio Popmintchev, Ming-Chang Chen, Dimitar Popmintchev, Paul Arpin, Susannah Brown, Skirmantas Alisauskas, Giedrius Andriukaitis, Tadas Balciunas, Oliver D. Mucke, Andrius Pugzlys, Andrius Baltuska, Bonggu Shim, Samuel E. Scgrauth, Alexander Gaeta, Carlos Hernandez-Garcia, Luis Plaja, Andreas Becker, Agnieszka Jaron-Becker, Margaret M. Murnane, and Henry C. Kapteyn. Bright coherent ultrahigh harmonics in the kev x-ray regime from mid-infrared femtosecond lasers. Science, 336:1287–1291, June 2012.

[6] Sterling Backus, Charles G. Durfee III, Margaret M. Murnane, and Henry C. Kapteyn. High power ultrafast lasers. Review of Scientific Instruments, 69(3):1207–1223, March 1998.

[7] Sterling Backus, Randy Bartels, Sarah Thompson, Robert Dollinger, Henry C. Kapteyn, and Margaret M. Murnane. High-efficiency, sing-stage 7-khz high-average-power ultrafast laser system. Optics Letters, 26(7):465–467, April 2001.

[8] Dmitriy Zusin, Ronny Knut, Patrik Grychtol, Ofer Kfir, Christian Gentry, Hans Nembach, Justin Shaw, Tom Silva, Avner Fleischer, Oren Cohen, Henry Kapteyn, and Margaret Murnane. Probing ultrafast magnetization dynamics using bright circularly polarized high harmonics. In CLEO: 2015 OSA Technical Digest, page STu4N.5. Optical Society of America, 2015.

[9] Kathleen M. Hoogeboom-Pot, Jorge N. Hernandez-Charpak, Xiaokun Gu, Travis D. Frazer, Erik H. Anderson, Weilun Chao, Roger W. Falcone, Ronggui Yang, Margaret M. Murnane, Henry C. Kapteyn, and Damiano Nardi. A new regime of nanoscale thermal transport: Collective diffusion increases dissipation efficiency. PNAS, 112(16):4846 – 4851, March 2015.

[10] Matthew D. Seaberg, Daniel E. Adams, Ethan L. Townsend, Daisy A. Raymondson, William F. Schlotter, Yanwei Liu, Carmen S. Menoni, Lu Rong, Chien-Chun Chen, Jianwei Miao, Henry C. Kapteyn, and Margaret M. Murnane. Ultrahigh 22 nm resolution coherent diffractive imaging using a desktop 13 nm high harmonic source. Optics Express, 19(23):22470 – 22479, October 2011.

[11] P. B. Corkum. Plasma perspective on strong-field multiphoton ionization. Physical Review Letters, 71(13):1994–1997, September 1993.

[12] Andy Rundquist, Charles G. Durfee III, Zenghu Chang, Catherine Herne, Sterling Backus, Margaret M. Murnane, and Henry C. Kapteyn. Phase-matched generation of coherent soft x-rays. Science, 280(5368):1412–1415, May 1998.

[13] Charles G. Durfee III, Andy R. Rundquist, Sterling Backus, Catherine Herne, Margaret M. Murnane, and Henry C. Kapteyn. Phase matching of high-order harmonics in hollow waveguides. Physical Review Letters, 83(11):2187 – 2190, September 1999.

[14] E. Constant, D. Garzella, P. Breger, E. Mevel, Ch. Dorrer, C. Le Blanc, F. Salin, and P. Agostini. Optimizing high harmonic generation in absorbing gases: Model and experiment. Physical Review Letters, 82(8):1668 – 1671, February 1999.

[15] Kenneth C. Kulander. Dynamics of laser excitation, ionization and harmonic conversion in inert gas atoms. Progress in Crystal Growth and Characterization of Materials, 33(1-3):193–199, 1996.

[16] J. Tate, T. Auguste, H. G. Muller, P. Salieres, P. Agostini, and L. F. DiMauro. Scaling of wave-packet dynamics in an intense midinfrared field. Physical Review Letters, 98(013901):1 – 4, January 2007.

[17] M. V. Frolov, N. L. Manakov, and Anthony F. Starace. Wavelength scaling of high-harmonic yield: Threshold phenomena and bound state symmetry dependence. Physical Review Letters, 100(173001):1 – 4, May 2008.

[18] A. L'Huillier, K. J. Schafer, and K. C. Kulander. High-order harmonic generation in xenon at 1064 nm: The role of phase matching. Physical Review Letters, 66(17):2200 – 2203, April 1991.

[19] M. D. Perry, D. Pennington, B. C. Stuart, G. Tietbohl, J. A. Britten, C. Brown, S. Herman, B. Golick, M. Kartz, J. Miller, H. T. Powell, M. Vergino, and V. Yanovsky. Petawatt laser pulses. Optics Letters, 24(3):160–162, February 1999.

[20] M. Aoyama, K. Yamakwa, Y. Akahane, J. Ma, N. Inoue, H. Ueda, and H. Kiriyama. 0.85-pw, 33-fs ti:sapphire laser. Optics Letters, 28(17):1594 – 1596, September 2003.

[21] Amy L. Lytle, Xiaoshi Zhang, Richard L. Sandberg, Oren Cohen, Henry C. Kapteyn, and Margaret Murnane. Quasi-phase matching and characterization of high-order harmonic generation in hollow waveguides using counterpropagating light. Optics Express, 16(9):6544–6566, April 2008.

[22] Emily A. Gibson, Ariel Paul, Nick Wagner, Ra'anan Tobey, David Gaudiosi, Sterling Backus, Ivan P. Christov, Andy Aquila, Eric M. Gullikson, David T. Attwood, Margaret M. Murnane, and Henry Kapteyn. Coherent soft x-ray generation in the water window with quasi-phase matching. Science, 302(5642):95–98, October 2003.

[23] Xiaoshi Zhang, Amy L. Lytle, Tenio Popmintchev, Xibin Zhou, Henry C. Kapteyn, Margaret M. Murnane, and Oren Cohen. Quasi-phase-matching and quantum-path control of high-harmonic generation using counterpropagating light. Nature Physics, 3:270–275, April 2007.

[24] Xiaoshi Zhang, Amy L. Lytle, Oren Cohen, Margaret M. Murnane, and Henry C. Kapteyn. Quantum-path control in high-order harmonic generation at high photon energies. New Journal of Physics, 10:1–9, 2008.

[25] Tenio Popmintchev, Ming-Chang Chen, Alon Bahabad, Margaret M. Murnane, and Henry C. Kapteyn. Phase-matched generation of coherent soft and hard x-rays using ir lasers. US Patent US8462824 B2, June 2013.

[26] J. Stohr, Y. Wu, B. D. Hermsmeier, M. G. Samant, G. R. Harp, S. Koranda, D. Dunham, and B. P. Tonner. Element-specific magnetic microscopy with circularly polarized x-rays. Science, 259:658–661, January 1993.

[27] Ming-Chang Chen, Christopher Mancuso, Carlos Hernandez-Garcia, Franklin Dollar, Ben Galloway, Dimitar Popmintchev, Pei-Chi Huang, Barry Walker, Luis Plaja, Agnieszka Jaron-Becker, Andreas Becker, Margaret M. Murnane, Henry C. Kapteyn, and Tenio Popmintchev. Generation of bright isolated attosecond soft x-ray pulses driven by multicycle midinfrared lasers. PNAS, pages E2361–E2367, May 2014.

[28] C. Hernández-García, M.-C. Chen, C. Mancuso, F. Dollar, B. Galloway, D. Popmintchev, P.C. Huang, B.C. Walker, T. Popmintchev, M.M. Murnane, H.C. Kapteyn, L. Plaja, A. Jaron-Becker, and A. Becker. Theory of time-gated phase-matching for isolated attosecond soft x-ray pulse generation using mid-infrared lasers. In CLEO: 2014, OSA Technical Digest. Optical Society of America, 2014.

[29] Dimitar Popmintchev, Carlos Hernandez-Garcia, Bonggu Shim, Ming-Chang Chen, Franklin Dollar, Christopher A. Mancuso, Jose Pérez-Hernández, Xiaohui Gao, Amelia Hankla, Alexander L. Gaeta, Maryam Tarazkar, Dmitri Romanov nad Robert Levis, Agnieszka Jaron-Becker, Andreas Becker, Luis Plaja, Margaret Murnane, Henry Kapteyn, and Tenio Popmintchev. Bright high order harmonic generation in a multiply ionized plasma up to the water window. In CLEO: 2014 Postdeadline Paper Digest, OSA Technical Digest, page FTh5A.9. Optical Society of America, 2014.

[30] P. F. Moulton. Spectroscopic and laser characteristics of ti:al2o3. Journal of the Optical Society of America B, 3(1):125–133, January 1986.

[31] Irina T. Sorokina and Konstantin L. Vodopyanov, editors. Solid-State Mid-Indrared Laser Sources, volume 89 of Topics in Applied Physics. Springer Berlin Heidelberg, Tiergartenstr. 17, 69121 Heidelberg, Germany, May 2003.

[32] S. Backus, J. Peatross, C. P. Huang, M. M. Murnane, and H. C. Kapteyn. Ti:sapphire amplifier producing millijoule-level, 21-fs pulses at 1 khz. Optics Letters, 20(19):2000–2002, October 1995.

[33] P. P. Sorokin and M. J. Stevenson. Stimulated infrared emission from trivalent uranium. Physical Review Letters, 5(12):557–559, December 1960.

[34] L. F. Mollenauer and R. H. Stolen. The soliton laser. Optics Letters, 9(1):13–15, January 1984.

[35] Evgeni Sorokin, Sergey Naumov, and Irina T. Sorokina. Ultrabroadband infrared solid-state lasers. IEEE Journal of Selected Topics in Quantum Electronics, 11(3):690–712, 2005.

[36] Sergey B. Mirov, Vladimir V. Fedorov, Dmitry Martyshkin, Igor S. Moskalev, Mike Mirov, and Sergey Vasilyev. Progress in mid-ir lasers based on cr and fe-doped ii-vi chalcogenides. IEEE Journal of Selected Topics in Quantum Electronics, 21(1), February 2015.

[37] Christopher Marshall, editor. Advanced Solid-State Lasers, volume 46 of OSA Trends in Optics and Photonics, chapter Active and passive mode-locking of the Cr^{2+}:ZnSe laser, pages 157–161. Optical Society of America, Washington D. C., 2001.

[38] Nikolai Tolstik, Evgeni Sorokin, and Irina T. Sorokina. Graphene mode-locked cr:zns laser with 41 fs pulse duration. Optics Express, 22(5):5564 – 5571, February 2014.

[39] Peter Moulton and Evgeni Slobodchikov. 1-gw-peak-power, cr:znse laser. In CLEO:2011 - Laser Applications to Photonic Applications, OSA Technical Digest, page PDPA10. Optical Society of America, 2011.

[40] Sterling Backus, Charles G. Durfee III, Gerard Mourou, Henry C. Kapteyn, and Margaret M. Murnane. 0.2-tw laser system at 1 khz. Optics Letters, 22(16):1256–1258, August 1997.

[41] Y. Jeong, J. K. Sahu, D. N. Payne, and J. Nilsson. Ytterbium-doped large-core fiber laser with 1.36 kw continuous-wave output power. Optics Express, 12(25):6088–6092, December 2004.

[42] Stuart D. Jackson. Towards high-power mid-infrared emission from a fibre laser. Nature Photonics, 6:423–431, July 2012.

[43] Vincent Fortin, Martin Bernier, Souleymane T. Bah, and Réal Vallée. 30 w fluoride glass all-fiber laser at 2.94 um. Optics Letters, 40(12):2882 – 2885, 2015.

[44] Ori Henderson-Sapir, Stuart D. Jackson, and David J. Ottaway. Versatile and widely tunable mid-infrared erbium doped zblan fiber laser. Optics Letters, 41(7):1676 – 1679, 2016.

[45] Vincent Fortin, Frédéric Maes, Martin Bernier, Souleymane Toubou Bah, Marc D'Auteuil, and Réal Vallée. Watt-level erbium-doped all-fiber laser at 3.44um. Optics Letters, 41(3):559 – 562, 2016.

[46] Jutta Schneider, Carsten Carbonnier, and Udo B. Unrau. Characterization of a ho3+-doped fluoride fiber laser with a 3.9-μm emission wavelength. Applied Optics, 36(33):8595–8600, November 1997.

[47] Yu Yao, Anthony Hoffman, and Claire F. Gmachl. Mid-infrared quantum cascade lasers. Nature Photonics, 6:432–439, July 2012.

[48] Claire Gmachl, Deborah L. Sivco, Raffaele Colombelli, Federico Capasso, and Alfred Y. Cho. Ultra-broadband semiconductor laser. Letters to Nature, 415:883–887, February 2002.

[49] Christine Y. Wang, Lyuba Kuznetsova, V. M. Gkortsas, L. Diehl, F. X. Kärtner, M. A. Belkin, A. Belyanin, X. Li, D. Ham, H Schneider, P. Grant, C. Y. Song, S. Haffouz, Z. R. Wasilewski, H. C. Liu, and Federico Capasso. Mode-locked pulses from mid-infrared quantum cascade lasers. Optics Express, 17(15):12929–12943, July 2009.

[50] Pranalytica Inc. Quantum cascade lasers. http://www.pranalytica.com/core-technologies/quantum-cascade-lasers.php, September 2016.

[51] Robert W. Boyd. Nonlinear Optics. Academic Press, Burlington, third edition, 2008.

[52] Erich P. Ippen. Nonlinear optics. Class notes in 6.634, MIT, Spring Semester 2009.

[53] Giulio Cerullo and Sandro De Silvestri. Ultrafast optical parametric amplifiers. Review of Scientific Instruments, 74(1):1 – 18, January 2003.

[54] Irving H. Malitson. Refraction and dispersion of synthetic sapphire. Journal of the Optical Society of America, 52(12):1377–1379, December 1962.

[55] Eugene Hecht and A. R. Ganesan. Optics. Pearson education. Addison-Wesley, 4 edition, 2002.

[56] L. Hongjun, Z. Wei, C. Guofu, W. Yishan, C Zhao, and R. Chi. Investigation of spectral bandwidth of optical parametric amplification. Applied Physics B, 79(5):569–576, June 2004.

[57] Stefan Witte and Kjeld S. E. Eikema. Ultrafast optical parametric chirped-pulse amplification. IEEE Journal of Quantum Electronics, 18(1):296–307, January 2012.

[58] J. Armstrong, N. Bloembergen, J. Ducuing, and P. Pershan. Interactions between light waves in a nonlinear dielectric. Physical Review Letters, 127(6):1918–1939, September 1962.

[59] Christopher B. Ebert, Loren A. Eyres, Martin M. Fejer, and James S. Harris Jr. Mbe growth of antiphase gaas films using gaas/ge/gaas heteroepitaxy. Journal of Crystal Growth, 201(202):187–193, 1999.

[60] C.R. Phillips, B. W. Mayer, L. Gallmann, M. M. Fejer, and U. Keller. Design constraints of optical parametric chirped pulse amplification based on chirped quasi-phase-matching gratings. Optics Express, 22(8):9627–9658, April 2014.

[61] M.M. Fejer, G. A. Magel, Dieter H. Jundt, and R. L. Byer. Quasi-phase-matched second harmonic generation: tuning and tolerances. IEEE Journal of Quantum Electronics, 28(11):2631–2654, November 1992.

[62] Mathieu Charbonneau-Lefort, Bedros Afeyan, and M. M. Fejer. Optical parameteric amplifiers using chirped quasiphase-matching gratings i: practical design formulas. Journal of the Optical Society of America, 25(4):463–480, April 2008.

[63] A. V. Smith. Snlo nonlinear optics code. available from AS-Photonics, Albuquerque, NM.

[64] C. Heese, C. R. Phillips, L. Gallmann, M. M. Fejer, and U. Keller. Role of apodization in optical parametric amplifiers based on aperiodic quasi-phasematching gratings. Optics Express, 20(16):18066–18071, July 2012.

[65] Haim Suchowski, Peter R. Krogen, Shu-Wei Huang, Franz X. Kärtner, and Jeffrey Moses. Octave-spanning coherent mid-ir generation via adiabatic difference frequency generation via adiabatic difference frequency conversion. Optics Express, 21(23):28892–28901, November 2013.

[66] Donna Strockland and Gerard Mourou. Compression of amplified chirped optical pulses. Optics Communications, 55(6):447–449, October 1985.

[67] Gerard Mourou, Christopher P. J. Barty, and Micheal D. Perry. Ultrahigh-intensity lasers: Physics of the extreme on a tabletop. Physics Today, 51(1):22–28, January 1998.

[68] B.C. Stuart, M. D. Feit, A. M. Rubenchik, B.W. Shore, and M. D. Perry. Laser-induced damage in dielectrics with nanosecond to subpicosecond pulses. Physical Review Letters, 74(12):2248–2252, March 1995.

[69] B. W. Mayer, C. R. Phillips, L. Gallmann, and U. Keller. Mid-infrared pulse generation via achromatic quasi-phase-matched opcpa. Optics Express, 22(17):20798 – 20808, August 2014.

[70] Matthias Baudisch, Alexandre Thai, Michael Hemmer, Hideki Ishizuki, Takunori Taira, and Jens Biegert. 5-cycle, 160-khz, 20-μj mid-ir opcpa. In G. Huber and P. Moulton, editors, Advanced Solid-State Lasers Congress Technical Digest, number paper AF1A.7 in OSA Technical Digest. Optical Society of America, 2013.

[71] B. W. Mayer, C. R. Phillips, L. Gallmann, M. M. Fejer, and U. Keller. Sub-four-cycle laser pulses directly from a high-repetition-rate optica parametric chirped-pulse amplifier at 3.4 um. Optics Letters, 38(21):4265–4268, November 2013.

[72] Giedrius Andriukaitis, Tadas Balčiūnas, Skirmantas Ališauskas, Audrius Pugžlys, Andrius Baltuška, Tenio Popmintchev, Ming-Chang Chen, Margaret M. Murnane, and Henry C. Kapteyn. 90 gw peak power few-cycle mid-infrared pulses from an optical parametric amplifier. Optics Letters, 36(15):2755–2757, August 2011.

[73] Yunpei Deng, Alexander Schwarz, Hanieh Fattahi, Moritz Ueffing, Xun Gu, Marcus Ossiander, Thomas Metzger, Volodymyr Pervak, Hideki Ishizuki, Takunori Taira, Takayoshi Kobayashi, Gilad Marcus, Ferenc Krausz, Reinhard Kienberger, and Nicholas Karpowicz. Carrier-envelope-phase-stable, 1.2 mj, 1.5 cycle laser pulses at 2.1 um. Optics Letters, 37(23):4973–4975, December 2012.

[74] Chien-Jen Lai, Kyung-Han Hong, Jonathas P Siqueira, Peter Krogen, Chun-Lin Chang, Gregory J Stein, Houkun Liang, Phillip D Keathley, Guillaume Laurent, Jeffrey Moses, Luis E Zapata, and Franz X. Kärtner. Multi-mj mid-infrared khz opcpa and yb-doped pump lasers for tabletop coherent soft x-ray generation. Journal of Optics, 17(094009), September 2015.

[75] Tso Yee Fan, Daniel J. Ripin, Roshan L. Aggarwal, Juan R. Ochoa, Bien Chann, Michael Tilleman, and Joshua Spitzberg. Cryogenic yb3+-doped solid-state lasers. IEEE Journal of Selected Topics in Quantum Electronics, 13(3):448 – 459, May 2007.

[76] David C. Brown, Rufus L. Cone, Yongchen Sun, and Randy W. Equall. Yb:yag absorption at ambient and cryogenic temperatures. IEEE Journal of Selected Topics in Quantum Electronics, 11(3):604–612, May 2005.

[77] Falgun D. Patel, Eric C. Honea, Joel Speth, Stephen A. Payne, Ralph Hutcheson, and Randy Equall. Laser demonstration of yb3al5o12 (ybag) and materials properties of highly doped yb:yag. IEEE Journal of Quantum Electronics, 37(1), January 2001.

[78] David C. Brown. The promise of cryogenic solid-state lasers. IEEE Journal of Selected Topics in Quantum Electronics, 11(3):587–599, May 2005.

[79] Andy Chong, Joel Buckley, Will Renninger, and Frank Wise. All-normal-dispersion femtosecond fiber laser. Optics Express, 14(21), October 2006.

[80] Andy Chong, William H. Renninger, and Frank W. Wise. All-normal-dispersion femtosecond fiber laser with pulse energy above 20 nj. Optics Letters, 32(16):2408 – 2410, August 2007.

[81] M. R. Gerrity, S. Brown, T. Popmintchev, M. M. Murnane, H. C. Kapteyn, and S. Backus. Multi-mj, high repetition rate, mid-ir opcpa system. In Advanced Solid-State Lasers, page AM3A.5. Optical Society of America, 2014.

[82] M. R. Gerrity, S. Brown, T. Popmintchev, M.-C. Chen, S. Witte, M. M. Murnane, H. C. Kapteyn, and S. Backus. High power, 60mhz, cryogenically cooled, mode-locked, yb:yag oscillator. In CLEO: 2011 OSA Technical Digest. CLEO, 2011.

[83] Jun Dong, Michael Bass, Yanli Mao, Peizhen Deng, and Fuxi Gan. Dependence of the yb3+ emission cross section and lifetime on temperature and concentration in yttrium aluminum garnet. Journal of the Optical Society of America, 20(9):1975–1979, September 2003.

[84] Jochen Dörring, Alexander Killi, and Uwe Morgner. Period doubling and deterministic chaos in continuously pumped regenerative amplifiers. Optics Express, 12(8):1759–1768, April 2004.

[85] Mikhail Grishin, Vidmantas Gulbinas, and Andrejus Michailovas. Bifurcation suppression for stability improvement in nd:yvo4 regenerative amplifier. Optics Express, 17(18):15700–15708, August 2009.

[86] Joseph T. Verdeyen. Laser Electronics. Series in Solid State Physical Electronics. Prentice Hall, Upper Saddle River, NJ 07458, third edition edition, 1995.

[87] Adolf Giesen and Jochen Speiser. Fifteen years of work on thin-disk lasers: Results and scaling laws. IEEE Journal of Selected Topics in Quantum Electronics, 13(3):598 – 609, May 2007.

[88] Tina Gottwald, Christian Stolzenburg, Dominik Bauer, Jochen Kleinbauer, Vincent Kuhn, Thomas Metzger, Sven-Silvius Schad, Dirk Sutter, and Alexander Killi. Recent disk laser development at trumpf. In Willy L. Bohn Harro Ackermann, editor, High-Power Lasers 2012: Technology and Systems, volume 8547, pages 1–8. SPIE, 2012.

[89] Thomas Metzger, Alexander Schwarz, Catherine Yuriko Teisset, Dirk Sutter, Alexander Killi, Reinhard Kienberger, and Ferenc Krausz. High-repetition-rate picosecond pump laser based on a yb:yag disk amplifier for optical parametric amplification. Optics Letters, 34(14):2123–2125, July 2009.

[90] Sandro Klingebiel, Marcel Schultze, Catherine Y. Teisset, Robert Bessing, Matthias Häfner, Stephan Prinz, Martin Gorjan, Dirk Sutter, Knut Michel, Helena G. Barros, Zsuzsanna Major, Ferenc Krausz, and Thomas Metzger. 220mj ultrafast thin-disk regenerative amplifier. In CLEO: 2015 OSA Technical DigestOSA Technical Digest, page STu4O.2. Optical Society of America, 2015.

[91] Y.-F. Chen. Design criteria for concentration optimization in scaling diode end-pumped lasers to high powers: Influence of thermal fracture. IEEE Journal of Quantum Electronics, 35(2):234 – 239, February 1999.

[92] David C. Brown. Ultrahigh-average-power diode-pumped nd:yag and yb:yag lasers. IEEE Journal of Quantum Electronics, 33(5):861 – 873, May 1997.

[93] R. L. Aggarwal, D.J. Ripin, J.R. Ochoa, and T. Y. Fan. Measurement of thermo-optic properties of y3al5o12, lu3al5o12, yaio3, liyf4, liluf4, bay2f8, kgdwo4...2, and kywo4...2 laser crystals measurement of thermo-optic properties of y3al5o12, lu3al5o12, yaio3, liyf4, liluf4 bay2f8, kgd(wo4)2, and ky(wo4)2 laser crystals in the 80 – 300 k temperature range. Journal of Applies Physics, 98(103514):1–14, November 2005.

[94] David C. Brown. Nonlinear thermal distortion in yag rod amplifiers. IEEE Journal of Quantum Electronics, 34(12):2383–2392, December 1998.

[95] M. J. Weber, D. Milam, and W. L. Smith. Nonlinear refractive index of glasses and crystals. Optical Engineering, 17(5):463 – 469, September 1978.

[96] Damian N. Schimpf, Tino Eidam, Enrico Seise, Steffen Hadrich, Jens Limpert, and Andreas Tünnermann. Circular versus linear polarization in laser-amplifiers with kerr-nonlinearity. Optics Express, 17(21), October 2009.

[97] Bradley M. Luther, Kathryn M. Tracy, Michael Gerrity, Susannah Brown, and Amber T. Krummel. 2d ir spectroscopy at 100 khz utilizing a mid-ir opcpa laser source. Optics Express, 24(4), February 2016.

[98] John M. Dudley, Goery Genty, and Stephane Coen. Supercontinuum generation in photonic crystal fiber. Reviews of Modern Physics, 78:1135 – 1184, October 2006.

[99] Scott R. Domingue and Randy A. Bartels. Overcoming temporal polarization instabilities from the latent birefringence in all-normal disperion, wave-breaking-extended nonlinear fiver supercontinuum generation. Optics Express, 21(11):13305 – 13321, 2013.

[100] Charles G. Durfee, Jeff A. Squire, and Steve Kane. A modular approach to the analytic calculation of spectral phase for grisms and other refractive/diffractive structures. Optics Express, 16(22):18004–18016, October 2008.

[101] Jeffrey Moses, Cristian Manzoni, Shu-Wei Huang, Guilio Cerullo, and Franz X. Kärtner. Temporal optimization of ultrabroadband high-energy opcpa. Optics Express, 17(7):5540–5555, March 2009.

[102] J.A. Fülöp, Zs. Major, B. Horváth, F. Tavella, A. Baltuška, and F. Krausz. Shaping of picosecond pulses for pumping optical parametric amplification. Applied Physics B, 87(1):79–84, March 2007.

[103] Kiminori Kondo, Hirohito Maeda, Yoshikazu Hama, Satoshi Morita, Arnaud Zoubir, Ryosuke Kodama, Kazuo A. Tanaka, Yoneyoshi Kitagawa, and Yasukazu Izawa. Control of amplified optical parametric fluorescence for hybrib chirped-pulse amplification. Journal of the Optical Society of America B, 23(2):231 – 235, February 2006.

[104] F Tavella, A Marcinkevicius, and F Krausz. Investigation of the superfluorescence and signal amplification in an ultrabroadband muliterawatt optical parametruc chirped pulse amplifier system. New Journal of Physics, 8(219):1 – 11, October 2006.

[105] C. R. Phillips, J. S. Pelc, , and M. M. Fejer. Parametric processes in quasi-phasematching gratings with random duty cycle errors. Journal of the Optical Society of America, 30(4):982–993, April 2013.

[106] Tingting Fan, Patrik Grychtol, Ronny Knut, Carlos Hernández-García, Daniel D. Hickstein, Dmitriy Zusin, Christian Gentry, Franklin J. Dollar, Christopher A. Mancuso, Craig W. Hogle, Ofer Kfir, Dominik Legut, Karel Carva, Jennifer L. Ellis, Kevin M. Dorney, Cong Chen, Oleg G. Shpyrko, Eric E. Fullerton, Oren Cohen, Peter M. Oppeneer, Dejan B. Milošević, Andreas Beckera, Agnieszka A. Jaroń-Becker, Tenio Popmintchev, Margaret M. Murnane, , and Henry C. Kapteyn. Bright circularly polarized soft x-ray high harmonics for x-ray magnetic circular dichroism. PNAS, 112(46):14206–14211, November 2015.