Development of a High Energy, kHz, Mid-Infrared OPCPA Laser for keV High Harmonic Generation

by

M. R. Gerrity

B.S., The University of Texas at Austin, 2007

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

2015

This thesis entitled:
Development of a High Energy, kHz, Mid-Infrared OPCPA Laser for keV High Harmonic Generation
written by M. R. Gerrity
has been approved for the Department of Physics

_____

Margaret Murnane

_____

Prof. Henry Kapteyn

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.

Gerrity, M. R. (Ph.D., Physics)

Development of a High Energy, kHz, Mid-Infrared OPCPA Laser for keV High Harmonic Generation

Thesis directed by Prof. Margaret Murnane

Coherent keV photon energy x-rays have many applications for materials science at the shortest length and time scales. Unfortunately, there are relatively few options for coherent x-ray generation. One of the most promising methods is high harmonic generation, wherein a femtosecond driving laser pulse is coherently upconverted to the x-ray region of the spectrum. Recent work has shown that the maximum x-ray photon energy that can be generated via high harmonic generation scales favorably with the wavelength of the driving laser pulse. This has sparked an interest in using mid-infrared (3-5μm) lasers to drive high harmonic generation.

However, high harmonic generation necessitates a mJ level, kHz repetition rate, femtosecond driving laser. At present, there are no such lasers in the mid-infrared region of the spectrum. This necessitates the development of new laser architectures for tabletop coherent x-ray generation.

OPCPA technology is one of the most promising avenues for high energy, high repletion rate lasers in the mid-infrared. This thesis reports on the design and development of a mJ level, kHz repetition rate, femtosecond OPCPA laser running at 3μm, optimized for tabletop coherent x-ray generation.

The system described here integrates and extends a variety of laser technologies towards this goal. The full laser is based upon an Yb:fiber oscillator and MgO:PPLN OPO front end. To pump our OPCPA system, we developed a four stage, cryogenic Yb:YAG laser running with >35mJ of output energy at 1kHz. We then use this to a pump a three stage OPCPA system, likewise running at 1kHz. We demonstrate over 3.4mJ of output energy at 1.55μm, along with 1.4mJ at 3μm. We then show compression of the 3μm output to <110fs. Finally, we conclude with the future directions for this laser, and discuss how it may be scaled to higher energies, shorter pulse lengths, and even further into the mid-infrared.

Dedication

To Terri.

Acknowledgements

There are many people I need to think for all of their help these past few years, as I've been fortunate to work with many wonderful people. First of all, to Margaret and Henry, thank you for all of the support you've given me throughout this work. Working in the KM group has been an incredible experience, and I've learned so much in my time here. Thank you for this one of a kind opportunity.

I've been fortunate to work with so many great people on this work, who all deserve thanks. In particular I want to thank Susannah Brown, who helped with every part of this laser right from the beginning, and Ming-Chang Chen, who taught me so much back when I first joined the group. I also want to express my thanks to Tenio Popmintchev, Chris Mancuso, Dimitar Popmintchev, Ben Galloway, Paul Arpin, Matt Seaberg, and Craig Hogle who worked with me on various projects. Thank you to all of those I worked with at KM Labs on this project, particularly Sterling Backus, Daisy Raymondson, and Matt Kirchner, but also Xiaoshi Zhang, Wil Estrada, and Chris Wood.

JILA would not be what it is without all of the dedicated people working hard to make it truly special. I want to thank Lois Jury, Lauren Mason, Hans Green and everyone in the machine shop who helped me along the way, Dave Alchenberger, and Brian, Jeff, and Randall in the supply office. To the many wonderful friends I've made, thank you for all of you encouragement and all of the fun we've had together. To all my family, you've supported me all the way and I am truly blessed to have you in my life. Mom and Dad, thank you for all of the love, kindness, and support you've given me. And finally, thank you Terri, for sharing this ride with me.

Contents

Chapter

1 Introduction 1

1.1 Coherent X-Ray Science 1

2 Nonlinear Optics and High Harmonic Generation 3

2.1 Nonlinear Optics 3

2.1.1 The Wave Equation: Linear propagation 5

2.1.2 The Wave Equation: Nonlinear Propagation 9

2.2 Observation of High Harmonic Generation 12

2.3 Three step model for high harmonic generation 12

2.3.1 Ionization 13

2.3.2 Propagation 15

2.3.3 Recombination 17

2.4 Macroscopic High Harmonic Generation 18

2.5 Mid Infrared High Harmonic Generation 22

2.5.1 Isolated Attosecond Pulses from mid-IR Driven High Harmonic Generation . 24

2.5.2 Group Velocity Effects on Mid-Infrared Driven High Harmonic Generation . 26

2.6 Conclusions 28

3 Ultrafast Mid-Infrared Lasers 30

3.1 The Infrared Region of the Spectrum 30

3.2 Parameters for High Energy, Mid-Infrared Laser for High Harmonic Generation . . . 32
3.3 Technologies for mJ-Level, Ultrafast, Mid-Infrared Lasers . . . 33
3.3.1 Mid-Infrared Laser Materials . . . 33
3.3.2 OPA Technology . . . 35
3.3.3 OPCPA Technology . . . 39
3.4 OPCPA Theory . . . 41
3.4.1 Three Wave Mixing . . . 41
3.4.2 Gain and Phase Matching . . . 42
3.4.3 Direct Phase Matching . . . 46
3.4.4 Quasi-Phase Matching . . . 47
3.4.5 Bandwidth Limitations . . . 52

4 Design of a High Energy, High Repetition Rate, Mid-Infrared OPCPA Laser 64
4.1 Full System Overview . . . 64
4.2 Front End . . . 66
4.2.1 Yb:fiber Oscillator . . . 67
4.2.2 MgO:PPLN Optical Parametric Oscillator . . . 68

5 Cryogenic Four Stage Yb:YAG Laser 73
5.1 Yb:YAG . . . 73
5.2 Cryogenic Yb:YAG Laser I: Stretcher and Yb:fiber Pre-Amplifier . . . 76
5.3 Cryogenic Yb:YAG Laser II: Yb:YAG Regenerative Amplifier . . . 78
5.3.1 Regenerative Amplifier Switching . . . 79
5.3.2 Regenerative Amplifier Dynamics . . . 80
5.3.3 Yb:YAG Regenerative Amplifier Performance . . . 81
5.4 Cryogenic Yb:YAG Laser III: Yb:YAG Second Stage Multipass Amplifier . . . 83
5.5 Cryogenic Yb:YAG Laser IV: Yb:YAG Third Stage Multipass Amplifier . . . 86
5.6 Cryogenic Yb:YAG Laser V: Yb:YAG Fourth Stage Multipass Amplifier . . . 88

5.6.1 Fourth Stage Efficiency 90

6 Broadband Mid-Infrared OPCPA Stages 92

6.1 Signal Stretching 92

6.1.1 First Stage 100

6.1.2 Second Stage 101

6.1.3 Third Stage 101

6.2 Compression 104

6.3 Future Work 106

6.3.1 OPA Front End 106

6.3.2 Idler Seeded OPCPA 109

6.3.3 Increased OPCPA Power Output 110

6.3.4 Prospects for Deep Mid-Infrared and Far-Infrared Pulse Generation 112

Bibliography 119

Appendix

A An Intuitive Explanation of the Relationship Between the Signal and Idler Phase 129

B Fourth Stage Crystal Mount 134

B.1 Stress Induced Birefringence 134

B.2 Low Stress Crystal Mounting I: Mounting Requirements 136

B.3 Low Stress Crystal Mounting II: Initial Mounting Attempts 137

B.4 Low Stress Crystal Mounting III: Detailed Models 140

B.4.1 Mounting Materials 140

B.4.2 Analytical Model I: Model Description 142

B.4.3 Analytical Model II: Model Results 144

B.4.4 Finite Element Analysis I: Model 146

B.4.5 Finite Element Analysis II: Results 148

B.5 CuW Crystal Mount 149

B.6 Conclusions 149

Tables

Table

3.1 The desired output parameters for a mid-infrared laser for high harmonic generation. 33

5.1 The physical, thermal, and optical properties of Yb:YAG at room temperature and cryogenic temperatures. 74

6.1 The signs of the terms in expansion of the phase of the idler compressor and stretched idler. 94

6.2 The GVD, TOD, and 4OD for the idler compressor with 300g/mm gratings, a 37°angle of incidence, and 10cm grating separation. The central wavelength is 3100nm. 95

6.3 The signs of the terms in expansion of the phase of the idler compressor, the stretched idler, and the stretched signal. 98

6.4 The GVD, TOD, and 4OD for the signal CFBG stretcher. Notice that all of the dispersion terms are negative, as require to compress the idler with a negative dispersion stretcher. 99

6.5 The parameters for each of the two stages for a 3.1μm pumped ZGP based OPA for mJ level, femtosecond pulse generation at 7.2μm. The seed parameters for the second OPA stage are based upon the first OPA stage output parameters. 115

Figures

Figure

2.1 The first experimental measurement of high harmonic generation. Up to the ninth harmonic the intensity falls as predicted by the perturbative theory, but the eleventh through the seventeenth harmonic are all of approximately equal intensity. No harmonics above the seventeenth were observed. From [77] 13

2.2 The potential seen by a bound electron in a laser field is the sum of the atomic Coulomb potential and the laser potential. At high enough laser field strengths, the laser will sufficiently distort the Coulomb potential to allow tunnel ionization. 14

2.3 The fractional population of neutral atoms (shown in blue) and singly ionized atoms (shown in green) for ADK ionization in helium. The electric field is overlayed in black. The calculation assumes a 10fs Gaussian pulse, with 3mJ pulse energy at 800nm, focused to a peak intensity of $\sim 2 \times 10^{15}$ W/cm^2. This corresponds to a 100μm focal radius. Notice that the ionization occurs in steps at the peaks of the electric field, and turns off when the amplitude is near zero. 15

2.4 The position of the electron oscillating in the laser electric field as a function of time, plotted for two periods of the laser electric field at ionization. The position is plotted for four different phases of the electric field (ϕ_0) at ionization, shown in blue, red, green, and black. The parent ion, being much more massive than the electron, is fixed at $\mathbf{x} = 0$. These four initial phase values will all bring the electron back to its parent ion. For ionization at phase $> \pi/2$, the electron will never return to its parent ion. 16

2.5 The kinetic energy of the electron upon its return to its parent ion, as a function of the laser phase at ionization. The maximum return kinetic energy for the electron is 3.17 times the ponderomotive energy U_p. 17

2.6 Double pinhole diffraction patterns for phase matched high harmonic generation. The different patterns reflect different pinhole spacing. The full fringe contrast demonstrates the full spatial coherence of the generated harmonics. Adapted from [11]. 22

2.7 A plot of the normalized harmonic emission spectrum from several different driving laser wavelengths in helium. On the far left in yellow is 800nm, with 1300nm in green, 2000nm in blue, and 3900nm, in purple on the far right. When driven at 3900nm the phase matched emission spectrum extends beyond 1 keV photon energy, all the way up to 1.6keV. Adapted from [95]. 23

2.8 Images demonstrating the spatial coherence of keV high harmonic generation. In **a** we see a beam profile for the keV harmonic in Helium, demonstrating a laser like beam profile. In **b** and **c** we have Young's double slit diffraction experiments for Neon (**b**) and Helium (**c**), showing the full spatial coherence. Adapted from [95]. . . 24

2.9 Experimental data showing the field autocorrelation signal of the harmonics from argon for three different driving laser wavelengths. Each trace was taken with a 10 cycle driving laser pulse and used the optimum phase matching pressure and intensity for the particular driving laser wavelength. With a near-infrared driving laser at 800nm, the harmonic emission shows several distinct x-ray bursts, leading to multiple peaks in the field autocorrelation. This is likewise true for 1300nm driven harmonic generation, although there are fewer distinct bursts than with 800nm. With mid-infrared driven harmonic generation however, the emission contains only a single isolated attosecond pulse, leading to a single peak in the field autocorrelation. Adapted from [20]. 25

2.10 The effect of the group velocity mismatch on output harmonic flux. Plotted is the output harmonic flux from Helium as a function of gas pressure for two 2.0μm driving laser pulses, one 1.4 cycles and one 5.8 cycles. Notice for the 1.4 pulse, the output initially grows with increasing gas pressure, but then decreases as the pressure induced group velocity mismatch destroys the phase matching. For the 5.8 cycle pulse, the group velocity mismatch does not destroy the phase matching, and the harmonic output continues to grow with increasing pressure. Adapted from [49]. 27

3.1 A coarse plot of the modeled atmospheric absorption from 1 through 5.5μm. The primary absorption is due to water and CO_2. Data calculated from [67] and provided by [84]. Modeled for the Cerro Pachon observatory, 4.3mm water column. 31

3.2 A typical layout for a single stage ti:sapphire pumped OPA. The input from the ti:sapphire is split into two arms. A small amount of the energy, usually on the order of microjoules, is used to generate a supercontinuum via white light generation. The white light is then used as the signal seed for an OPA stage, which is pumped by the remainder of the ti:sapphire input. 36

3.3 The signal gain in an OPCPA modeled with Eq. 3.7 for three different values of the phase mismatch Δk. This assumes a pump, signal, and idler wavelengths of 1.03 μm, 1.55μm, and 3.1μm, indices all equal to 1.5, a d_{eff} of 14.9 pm/V, and a pump intensity of 5 GW/cm^2. In red and blue are two cases where the phase mismatch is sufficiently large that γ is imaginary, which leads to oscillatory behavior. Larger phase mismatches leads to shorter oscillatory periods. In green $(\Delta k/2)^2$ is just slightly less than Γ^2, and gives exponential gain. Note that even though γ is still very small, the signal sees a gain $> 10^3$ in just 5mm. 45

3.4 The intensity growth in a nonlinear system for the three cases of no phase matching (shown in blue), quasi-phase matching (in red), and direct phase matching (in green). The arrows along with the alternating light and dark blue background indicate the crystal orientation and poling for the phase matched case. With no phase matching, the intensity initially grows, but after one coherence length the intensity begins to shrink again. This leads to an oscillation of the signal intensity with no significant buildup. The arrows along with the alternating light and dark blue background indicate the crystal orientation and poling for the phase matched case. With direct phase matching, the intensity continually grows with distance. This is likewise true for quasi-phase matching, although the growth is reduced when compared to the directly phase matched case. 49

3.5 A comparison of the temporal profile for the pump (shown in black), and the gain profiles for several different peak gain values between 10 through 10,000. When the peak gain is low, the gain temporal profile is essentially the same width as the pump pulse, but as the peak gain increases the gain temporal width narrows substantially. The gain profiles have been normalized and scaled to 80% the height of the pump profile for clarity. 57

3.6 A schematic of chirped quasi-phase matched crystal of length L. The periodicity increases from the $\Lambda(0)$ to $\Lambda(L)$ through the crystal, taking on the value Λ_0 at $z = z_0$. Λ_0 is assumed to be the periodicity which perfectly phase matches the signal and idler carrier frequencies ω_s and ω_i. 59

4.1 A diagram of the full laser system. The system can be thought of as being comprised of three subsystems - the front end (outlined in black), the cryogenic Yb:YAG laser (outlined in green), and the OPCPA laser (outlined in red). 65

4.2 A diagram of the components in the ANDi Yb:fiber oscillator. A 2m length of Yb doped fiber is pumped by a fiber coupled diode, emitting 3W at 980nm. The cavity contains no anomalously dispersive elements, and the mode locking is provided by a strong nonlinear phase shift provided a quarter waveplate ($\lambda/4$ is the diagram) and a birefringent plate (BP in the diagram). 67

4.3 On the left, the output spectrum from the Yb:fiber oscillator. Notice the spectrum covers 1030nm. This is important, since it means the Yb:fiber oscillator can be used to seed the Yb:YAG laser. In the center, the uncompressed temporal profile from the oscillator, with an approximate pulse width of 2200fs. On the right, the compressed output from the Yb:fiber oscillator, with a pulse width of 150s. 68

4.4 The phase matching range for MgO:PPLN, assuming a 1030nm pump and a crystal temperature of 100°C. Data calculated in SNLO. [107]. 69

4.5 On the left, the typical output spectrum of the Yb:fiber pumped MgO:PPLN OPO. The central wavelength is 1.55μm, and the spectrum has $\sim$25nm FWHM bandwidth. On the right, the spectrum after broadening in a section of HNLF. The spectrum has over 100nm of FWHM bandwidth, with almost 400nm of tail-to-tail bandwidth. 71

5.1 The fractional power transmission of a 20mm long, 2% at. doped Yb:YAG crystal as the crystal is cooled from room temperature to cryogenic temperatures. Notice that above ∼110K the crystal begins to absorb 1030nm, and that at room temperature approximately half the input power is absorbed. The temperature values used were the measured crystal mount temperatures, as the crystal temperature could not be measured directly in this setup. 75

5.2 A diagram of the Chirped Volume Bragg Grating used to stretch the seed for the Yb:YAG system. UV exposure is used to imprint a holographic grating pattern in the material, which reflects certain wavelengths of light. By varying the grating periodicity through the length of the material, different wavelengths are reflected at different depths in the material. This leads to a spectrally dependent path length difference in the reflected wavelengths, stretching the pulse in time. Wavelengths outside the reflected bandwidth simply pass through the entire grating. 77

5.3 On the left, a typical input spectrum for the Chirped Volume Bragg Grating stretcher. Notice that the input spectrum is much broader (approximately 25nm FWHM) than the operational range of the CVBG (0.6nm FWHM). On the right is the predicted temporal output from this input spectrum, using a 0.6nm FWHM square reflection bandwidth and 540ps/nm dispersion. The 0.6nm FWHM spectrum is stretched to approximately 285ps FWHM. The ringing on the predicted temporal output is an FFT artifact introduced by the square notch function used to model the finite reflection bandwidth of the CVBG. 78

5.4 The layout of the cryogenic Yb:YAG regenerative amplifier. An 8% doped Yb:YAG crystal is cooled to 80K, and pumped by a 940nm fiber coupled diode. The regen is seeded by the 60MHz pulse train at 1030nm. A Pockels Cell and $\lambda/4$ waveplate combination is used to to trap seed pulses inside the amplifier cavity, amplifying them to high energy. The amplifier output runs between 1 and 100kHz, and is set by the repetition rate of the Pockels Cell. 79

5.5 On the left, the output spectrum from the cryogenic Yb:YAG regenerative amplifier. The spectrum is centered at 1029.7nm with a FWHM bandwidth of 0.25nm. In the center, the modeled temporal output of the Yb:YAG regenerative amplifier when seeded with a 285ps, 0.6nm FWHM bandwidth seed pulse. The gain narrowing from the Yb:YAG amplifier cuts the pulse duration to 130ps. On the right, the output mode from the regenerative amplifier. 82

5.6 The layout of the cryogenic Yb:YAG second stage multiples amplifier. An 8% doped Yb:YAG crystal is cooled to 80K, and pumped with up to 41W by a 940nm fiber coupled diode. This stage is set up with a three pass configuration, with polarization used to couple the light in and out of the amplifier. Seeded by the 1kHz, 500μJ pulse train from the regenerative amplifier, this stage can run with over 8W of output power at 1kHz, though in practice it is typically run at 6W to mitigate the risk of damage. 83

5.7 The output mode immediately after the second Yb:YAG stage, showing a flat top beam profile. 85

5.8 A diagram of the third cryogenic Yb:YAG stage. This double pass stage is seeded by 6mJ from the second Yb:YAG stage. When pumped with 65W, it runs with 15mJ of output energy at 1kHz. 86

5.9 The far-field beam profile from the third cryogenic Yb:YAG stage. 87

5.10 A diagram of the fourth cryogenic Yb:YAG stage. This double pass stage is seeded by 15mJ from the third stage, and runs with 36mJ output energy at 1kHz. 89

5.11 The far-field beam profile from the fourth cryogenic Yb:YAG stage. 90

6.1 A ZEMAX ray trace model of the idler compressor, using two 300 g/mm gratings, a 37° angle of incidence, and a 10cm perpendicular grating separation. The ray trace model uses a 10mm beam diameter. 95

6.2 The modeled temporal output of the CFBG stretcher, and the modeled spectral output, based upon a measured output spectrum from the HNLF. The FWHM pulse length is 55ps. 99

6.3 The first OPCPA stage seed and output spectrum. The seed is provided by the broadened and stretched OPO output. The amplified bandwidth is approximately 35nm FHWM, with a pulse energy of 30μJ. 100

6.4 The second OPCPA stage seed and output spectrum. The second stage is seeded by the 30μJ, 35nm FWHM first stage output. This is amplified to 600μJ, with 28nmFWHM bandwidth. 102

6.5 The third OPCPA stage seed and output spectrum. The third stage is seeded by the 600μJ, 28nm FWHM second stage output. This is amplified to 3.4mJ, and the FWHM bandwidth is extended to approximately 62nm. 103

6.6 The third stage OPCPA idler spectrum, with a FWHM bandwidth of 275nm centered at 3.05μm. The transform limit is 57fs, showing more than sufficient bandwidth for <100fs pulse compression. 104

6.7 The spatial beam profiles for the OPCPA signal, shown on the left, and idler, shown on the right, demonstrating excellent spatial mode quality. 104

6.8 On the left, the retrieved temporal profile for the compressed idler, with a 105fs FWHM pulse length. On the right, the measured FROG trace for the compressed idler. 105

6.9 The output spectrum from the Y-Fi in the new front end. 107

6.10 Examples of the broadband output that can be obtained from the Y-Fi pumped OPA. These spectra were taken with YAG used for white light generation, and a 2mm fan out MgO:PPLN crystal for amplification. The signal spectrum on the left has a 70nm FWHM, and the idler on the right has 780nm FWHM. 108

6.11 The modeled input and output spectra and temporal profiles for the first ZGP OPA stage. The input pulse is a 400nm FWHM Gaussian centered at 7.2μm, with a FWHM pulse length of 190fs. At the output, the spectrum is broadened to 804nm FHWM, and the pulse length is 130s. The predicted gain in the first stage is 1.2×10^5, for an output pulse energy of 120μJ. 116

6.12 The modeled input and output spectra and temporal profiles for the second ZGP OPA stage. The input pulse is the 804nm FWHM output from the first OPA stage, with a pulse length of 130s. At the output, the spectrum is narrowed slightly to 790 FHWM, and the pulse length is 150s. The predicted gain in the second stage is 14, for an output pulse energy of 1.65mJ. 118

A.1 Pump and signal as function of frequency 130

A.2 Pump, signal, and idler as function of frequency 131

A.3 The pump, signal, and idler spectral intensity and phase as a function of frequency. Notice that the signal and idler phases are inverted mirror images of one another. That is, the idler phase is obtained from the signal phase by flipping it horizontally and vertically. Under this transformation odd functions (such as odd polynomials) retain their sign, while even functions (the even polynomials) flip sign. 132

B.1 The layout for measuring the depolarization loss due to stress birefringence. Linearly polarized input light enters through an input polarizer, then passes through the crystal. Stress birefringence causes a polarization shift in the light, which is then separated into its components parallel and perpendicular to the input polarization, which can be studied separately. 135

B.2 The original crystal mount. Shown are the two halves of the copper mount, and a sleeve with a gold coated crystal partially inserted. 137

B.3 The geometry of the analytical model used to estimate the stress on the crystal. The model consists of a crystal of radius r_c, a sleeve of radius r_s, and a mount of radius r_m. On the crystal-sleeve boundary is radial pressure p_{cs}, and on the sleeve-mount boundary is radial pressure p_{sm}. The model assumes that all layers are mounted such that $p_{cs} = p_{sm} = 0$ at room temperature. 142

B.4 The predicted pressure on the crystal in MPa for a Cu mount and a CuW sleeve as a function of sleeve thickness, for several different integrated CTE mismatch values between YAG and CuW. 151

B.5 The predicted pressure on the crystal in MPa for a Cu mount and a Si sleeve as a function of sleeve thickness. 151

B.6 The predicted pressure on the crystal for a CuW mount in MPa as a function of the integrated CTE mismatch between YAG and CuW, ranging from 0 to 500ppm. . . 152

B.7 The predicted pressure on the crystal in MPa for a CuW mount and a Si sleeve as a function of sleeve thickness, for several different integrated CTE mismatch values between YAG and CuW. 152

B.8 A cross sectional schematic of the mount design used for finite element modeling. The dark grey material is the mount material - CuW for the planned design. In the case of a single piece mount, the light grey material is the same as the dark grey material. In the case of a sleeved mount, the light grey material is the sleeve. The blue material is the crystal, which has been divided into 10 sections along the crystal depth. The central radial section of the crystal represents the pump area, and this is where heat to due to the quantum defect is deposited. 153

B.9 The finite element analysis modeled temperature distributions for: a) a Cu mount with no sleeve, b) a Cu mount with a Si sleeve, c) a CuW mount with no sleeve, and d) a CuW mount with a Si sleeve. 154

B.10 The finite element analysis modeled von Mises stress in: a) a Cu mount with no sleeve, b) a Cu mount with a Si sleeve, c) a CuW mount with no sleeve, and d) a CuW mount with a Si sleeve. 155

B.11 The finite element analysis modeled von Mises stress on the crystal for the Cu and CuW mounts with the Si sleeve. 156

B.12 The measured depolarization loss for the new CuW, and one of the previous best results from a copper mount with a 1mm thick copper sleeve. The CuW mount reduced the measured single pass depolarization loss to <0.2%, only a factor of 2 greater than the 1000:1 extinction ratio of the polarizer used in the measurement. . . 156

Chapter 1

Introduction

1.1 Coherent X-Ray Science

Coherent x-ray radiation has many uses for probing both the shortest length and time scales. The short wavelength, on the order of angstroms or nanometers, allows for incredibly high resolution imaging, while the capability of producing pulses that approach zeptosecond timescales allows for the resolution of processes that occur on the shortest of timescales. In addition many elements have absorption edges in the x-ray band, allowing for elemental specificity. An area interest are the "water window", the region for 300-500eV where water is transparent, but numerous other elements of biological interest have absorption edges. Further on, elements of interest for material and magnetic behavior such as Iron and Nickel have absorption edges between 750 and 850eV.

Unfortunately, despite its usefulness, there are not many ways to generate ultrashort, coherent x-ray radiation. Large synchrotron facilities such as the Advanced Light Source in Berkeley, CA, can provide high power pulses with photon energies >10keV, but the pulse lengths are on the order of 100fs, too long for probing the longest time scales. Free electron lasers such as the Linac Coherent Light Source can generate shorter keV photon energy pulses, but these are still limited to femtosecond pulse lengths [33]. Furthermore, while these facilities enable world-class science, they are limited in number, and it would be an incredible boost to x-ray material science to bring the experiments that can traditionally only be done at large facilities to the tabletop.

High harmonic generation is an extreme nonlinear process that allows for tabletop coherent x-ray generation. This process begins with a femtosecond laser, running anywhere from the ultraviolet

to the mid-infrared, and then coherently combining tens, hundreds, or even thousands of these photons to produce photons from the Vacuum Ultraviolet (VUV) all the way to soft (and possibly even hard) x-ray regions of the spectrum. High harmonic sources have been used for numerous application experiments, including magnetic studies, coherent diffractive imaging, and nanoscale acoustic studies [102, 73, 65].

This thesis will proceed as follows. First, in chapter 2, I will cover nonlinear optics in general, and high harmonic generation in particular. I will describe high harmonic generation at both the microscopic and macroscopic scales, with a particular emphasis on why it is a viable tabletop coherent x-ray source. The result of this section will be that for generating large fluxes of coherent soft x-rays, it is advantageous to use driving wavelengths in the mid-infrared region of the spectrum.

Next, in chapter 3 I will give a brief overview of the mid-infrared region of spectrum. This will include some of the challenges of working here, as well as some of the strategies for generating ultrashort pulses in the mid-infrared region. This will include a discussion of conventional laser sources, OPA laser technology, and OPCPA laser technology. It will be seen that OPCPA laser technology is the most promising avenue for generating high energy, high repetition rate, femtosecond pulses in the mid-infrared. This section will conclude with an overview of many of the theoretical concerns of OPCPA technology.

Chapter 4 will introduce the design of the laser, beginning with a high level overview of the full laser architecture. Next, it will give a detailed description of the front end.

The next two chapters, 5 and 6, will cover the design and construction of this system in detail. Chapter 5 will focus on the cryogenic Yb:YAG amplifier used to pump the three stage OPCPA system. This will include an overview of Yb:YAG as a material, and challenge of using it in a cryogenic high energy system. Chapter 6 will cover the three stage OPCPA system, demonstrating >1mJ output at 1kHz at 3μm, and show the initial compression results. Finally, it will conclude with a discussion of the future directions for this project, showing how the laser can be improved to allow for higher energy operation, as well as the prospects for moving further into the mid-infrared, and possibly into the far-infrared.

Chapter 2

Nonlinear Optics and High Harmonic Generation

Though high harmonic generation wasn't observed until 1987, its story begins back in 1961, when Franken reported the first observation of second harmonic generation [39]. In this first experiment, the output of a ruby laser (or optical maser, as it was known at the time), a technology itself less than one year old, was focused into a quartz crystal. The spectrum was measured after the focus, both with and without the quartz crystal in place. Without the quartz crystal in place, the spectrum following the focus was identical to that of the ruby laser - a bright peak at 694.3nm. With the quartz crystal in place however, a second, small peak was observed at 347.2nm - exactly half the wavelength of the light from the ruby laser. This was the first demonstration of optical harmonic generation, and the first experiment in the field of nonlinear optics. As we will see, high harmonic generation is simply nonlinear optics taken to the extreme. So, before understanding high harmonic generation, we must first understand nonlinear optics.

2.1 Nonlinear Optics

Nonlinear optics begins with a simple observation: at high enough optical intensities, the polarization of a material no longer responds linearly to the electric field from the light. That is, whereas in linear optics the instantaneous optically induced polarization can be written

$$\mathbf{P} = \epsilon_0 \chi \mathbf{E}, \tag{2.1}$$

now the polarization must be written as a power series in higher order terms of the susceptibility χ,

$$\mathbf{P} = \epsilon_0 \left(\chi^{(1)}\mathbf{E} + \chi^{(2)}\mathbf{E}^2 + \chi^{(3)}\mathbf{E}^3 + \cdots \right) \tag{2.2}$$

where $\mathbf{E}$ is the electric field [1]

There are several things to note about this expansion. The first is that this is a perturbative theory - the magnitude of each subsequent $\chi^{(n)}$ falls by roughly e/a_o^2, where e is the electron charge and a_0 is the Bohr radius [13]. This is why the laser was necessary for the discovery of nonlinear optics - it enabled the high field strengths required to observe effects beyond the linear response. Furthermore, as the magnitude of the susceptibility falls rapidly with each subsequent term, we expect the magnitudes of the effects of succeeding terms to likewise fall rapidly.

The second thing to note is that, as a perturbative expansion, this assumes that the electric field from the driving laser is small compared to the Coulomb field seen by the electron. As the atomic electric field seen by the electron is typically on the order of 10^{11}V/m, this is often a reasonable assumption. However, for an intense enough optical field this is no longer the case, and the the driving field can no longer be treated as a perturbation to the Coulomb field. This can result in qualitatively different behaviors from that predicted by the perturbative expansion. An obvious example is ionization. If the driving optical field is intense enough, it can suppress the Coulomb barrier to allow the electrons to escape the atom. The perturbative theory, however, always treats the electrons as bound in the atom, no matter how high the optical field strength. It simply does not account for this type of behavior.

The final important point is that n^{th} term in the expansion allows for interactions involving $n+1$ total photons. Therefore, for a monochromatic driving field, the second order susceptibility

[1] In this section, I will neglect the effects of the polarization of the electric field to avoid overly cluttered equation. More properly, each term $\chi^{(n)}$ is a rank $n+1$ tensor. So, for example, the second order susceptibility would be written as $\chi^{(2)}_{ijk}$, where i, j, and k, can take on the values of the crystal axes x, y, or z. With this, the second order polarization in the i^{th} direction is given by

$$P_i = \sum_{jk} \chi^{(2)}_{ijk} E_j E_k, \tag{2.3}$$

and similar summations can be written for the higher order terms. A more detailed account of the effects of polarization and the tensor nature of the nonlinear susceptibility can be found in [13].

allows for the generation of the second harmonic ($\omega + \omega = 2\omega$), the third order susceptibility allows for the generation of the third harmonic ($\omega + \omega + \omega = 3\omega$), and so forth. Combining this with the fact that the magnitudes the nonlinear susceptibility as the terms get higher, we expect it to be increasingly difficult to generate higher and higher harmonic orders.

With these factors in mind, we now return to see how the expanded polarization comes into play. This will proceed in the following manner. First, we will introduce the standard wave equation for an electromagnetic wave in a material. We will first focus on the linear polarization response, and we will derive the parabolic equation for linear propagation. Then, the nonlinear response will be added back in, resulting in a propagation equation for nonlinear interactions.

2.1.1 The Wave Equation: Linear propagation

In this section we will derive the parabolic wave equation for linear propagation. This is a differential equation describing the temporal and spatial evolution of a pulse propagating through a material, and will be very useful as we move to nonlinear propagation and nonlinear interactions.

The derivation begins with the general equation for an electromagnetic wave propagating in material, given by

$$\nabla^2 \mathbf{E} - \epsilon_0 \mu_0 \frac{\partial^2 \mathbf{E}}{\partial t^2} = \mu_0 \frac{\partial^2 \mathbf{P}}{\partial t^2}, \tag{2.4}$$

where $\mathbf{E} = \mathbf{E}(\mathbf{r}, t)$ is the electric field, $\mathbf{P}$ is the material polarization, and ϵ_0 and μ_0 are the vacuum permittivity and permeability, respectively. Notice that when the polarization term is removed, this simplifies into the well known wave equation for vacuum propagation. Our goal is to determine how the field, $\mathbf{E}$, evolves in such a material. We will assume solutions of the form

$$\mathbf{E}(\mathbf{r}, t) = A(\mathbf{r}, t) \left(e^{-i(\mathbf{k_0} \cdot \mathbf{r} - \omega_0 t)} + c.c. \right), \tag{2.5}$$

where $A(\mathbf{r}, t)$ is a time and spatially dependent amplitude function, ω_0 is the carrier frequency, and $\mathbf{k_0}$ is the wave vector for the carrier frequency. The wave vector is given by

$$\mathbf{k_0} = \frac{n(\omega_0)\omega_0}{c} \hat{\mathbf{k}}, \tag{2.6}$$

where $n(\omega)$ is the frequency dependent refractive index. We will assume the material is lossless, with no attenuation or gain, so that the refractive index is strictly real. By taking the Fourier transform of Eq. 2.5, we can also inspect the field in the frequency domain,

$$\begin{aligned}
\mathbf{E}(\mathbf{r},\omega) &= \mathscr{F}\left[\mathbf{E}(\mathbf{r},t)\right] && (2.7)\\
&= \int_{-\infty}^{\infty} \mathbf{E}(\mathbf{r},t)e^{-i\omega t}\mathrm{d}t && (2.8)\\
&= \int_{-\infty}^{\infty} A(\mathbf{r},t)e^{-i(\mathbf{k_0}\cdot\mathbf{r}-\omega_0 t)}e^{-i\omega t}\mathrm{d}t && (2.9)\\
&= e^{(i\mathbf{k_0}\cdot\mathbf{r})}\int_{-\infty}^{\infty} A(\mathbf{r},t)e^{i(\omega_0-\omega)t}\mathrm{d}t && (2.10)\\
&= e^{(i\mathbf{k_0}\cdot\mathbf{r})}A(\mathbf{r},\omega). && (2.11)
\end{aligned}$$

In the frequency domain, we have a spectral amplitude function $A(\mathbf{r},\omega)$. We will assume that our pulse has a reasonably narrow spectrum, so that $A(\mathbf{r},\omega)$ is fairly narrowly constrained around the carrier frequency ω_0. In addition, we will also assume that the refractive index $n(\omega)$ is likewise reasonably close to the refractive index at the carrier frequency $n(\omega_0)$. Together, these two assumptions also mean we can take $|\mathbf{k}| \approx |\mathbf{k_0}|$.

In the first step of our derivation, we can divide the polarization into linear and nonlinear terms,

$$\mathbf{P} = \mathbf{P}_L + \mathbf{P}_{NL}. \tag{2.12}$$

For the remainder of the this section we will focus only on the linear term, P_L, and omit the nonlinear polarization. In the next section we will add the nonlinear term back into the result from this section. With this, the wave equation now reads

$$\nabla^2\mathbf{E} - \epsilon_0\mu_0\frac{\partial^2\mathbf{E}}{\partial t^2} = \mu_0\frac{\partial^2\mathbf{P_L}}{\partial t^2}. \tag{2.13}$$

As in many cases, this analysis is easier done in the frequency domain, after which it can be transformed back into the time domain. As such we can take the Fourier transform of Eq. 2.13,

$$\mathscr{F}\left[\nabla^2\mathbf{E}(\mathbf{r},t) - \epsilon_0\mu_0\frac{\partial^2\mathbf{E}(\mathbf{r},t)}{\partial t^2} = \mu_0\frac{\partial^2\mathbf{P_L}(\mathbf{r},t)}{\partial t^2}\right], \tag{2.14}$$

which can be evaluated term by term. The first term on the left hand side is straightforward enough

$$\mathscr{F}\left[\nabla^2\mathbf{E}(\mathbf{r},t)\right] = \nabla^2\mathscr{F}\left[\mathbf{E}(\mathbf{r},t)\right] \tag{2.15}$$

$$= \nabla^2\mathbf{E}(\mathbf{r},\omega). \tag{2.16}$$

The second term on the left and the term on the right use the following Fourier identity [4]

$$\mathscr{F}\left(\frac{\partial^n f(t)}{\partial t}\right) = (i\omega)^n\mathscr{F}f(t). \tag{2.17}$$

Evaluating the second term on the left, we find

$$\mathscr{F}\left[\epsilon_0\mu_0\frac{\partial^2\mathbf{E}(\mathbf{r},t)}{\partial t^2}\right] = \epsilon_0\mu_0(i\omega)^2\mathscr{F}\left[\mathbf{E}(\mathbf{r},t)\right] \tag{2.18}$$

$$= -\epsilon_0\mu_0\omega^2\mathbf{E}(\mathbf{r},\omega). \tag{2.19}$$

For the polarization term on the right hand side

$$\mathscr{F}\left[\mu_0\frac{\partial^2\mathbf{P_L}(\mathbf{r},t)}{\partial t^2}\right] = \mu_0(i\omega)^2\mathscr{F}\left[\mathbf{P_L}(\mathbf{r},t)\right] \tag{2.20}$$

$$= -\mu_0\omega^2\mathbf{P_L}(\mathbf{r},\omega). \tag{2.21}$$

Here we make use of the fact that for an instantaneous, linear polarization in a lossless medium, we can simply use [103]

$$\mathbf{P_L}(\mathbf{r},\omega) = \epsilon_0\chi_L(\omega)\mathbf{E}(\mathbf{r},\omega), \tag{2.22}$$

where $\chi_L(\omega)$ is the frequency dependent linear susceptibility. Putting these three together, we find for the wave equation in the frequency domain

$$\left[\nabla^2 + \omega^2\epsilon_0\mu_0 + \omega^2\chi(\omega)\epsilon_0\mu_0\right]\mathbf{E}(\mathbf{r},\omega) = 0. \tag{2.23}$$

The next step is to evaluate the derivative in the first term on the left hand side. Prior to this, however, we may without any loss of generality define the beam to be propagating in the $+z$ direction, so that $\mathbf{k_0} = k_0\hat{\mathbf{z}}$, where $k_0 = |\mathbf{k_0}|$. Doing this, we may replace ∇^2 with $\partial^2/\partial z^2$, and our wave equation now reads

$$\left[\frac{\partial^2}{\partial z^2} + \mu_0\omega^2\epsilon(\omega)\right]\mathbf{E}(\mathbf{r},\omega) = 0, \tag{2.24}$$

where we have substituted in the dielectric constant $\epsilon(\omega) = \epsilon_0(1+\chi(\omega))$. With these simplifications, the second order derivative can be evaluated, giving

$$\left[\frac{\partial^2}{\partial z^2} - 2ik_0\frac{\partial}{\partial z} - k_0^2 + \mu_0\omega^2\epsilon(\omega)\right] A(z,\omega)e^{-ik_0 z} = 0. \tag{2.25}$$

Here we can make an approximation known as the slowing varying envelope approximation, which assumes that the amplitude of the pulse envelope varies much slower than the carrier frequency ω_0 [103]. Formally, this says

$$\left|\frac{\partial^2 A(z,\omega)}{\partial z^2}\right| << \left|2ik_0\frac{\partial A(z,\omega)}{\partial z}\right|. \tag{2.26}$$

With this approximation, we can drop the first term from Eq. 2.25. We can also take advantage of the fact that $k^2 = \omega^2\epsilon(\omega)/c^2$ to recast Eq. 2.25 as

$$\left[-2ik_0\frac{\partial}{\partial z} + (k^2 - k_0^2)\right] A(z,\omega)e^{-ik_0 z} = 0. \tag{2.27}$$

Remembering that our end goal is an equation in the time domain, we need to take the Fourier transform Eq. 2.26. However, as in general we don't know the functional form of $k(\omega)$, it is helpful to make a second approximation. As we have previously specified that $|\mathbf{k}| \approx |\mathbf{k_0}|$, we can use this to write

$$\begin{aligned} (k^2 - k_0^2) &= (k + k_0)(k - k_0) && (2.28)\\ &\approx 2k_0(k - k_0), && (2.29) \end{aligned}$$

and now $k - k_0$ can be expanded in a Taylor series around the central frequency ω_0

$$k - k_0 = k_0 + \left.\frac{\partial k}{\partial \omega}\right|_{\omega_0}(\omega - \omega_0) + \frac{1}{2}\left.\frac{\partial^2 k}{\partial \omega^2}\right|_{\omega_0}(\omega - \omega_0)^2 + \cdots - k_0. \tag{2.30}$$

Again, remembering we have restricted ourselves to situations with only a relatively narrow range of frequencies around ω_0, we may here drop all the terms beyond $(\omega-\omega_0)^2$. Making this approximation and substituting into Eq. 2.27, this gives

$$\left[-2ik_0\frac{\partial}{\partial z} + 2k_0\left.\frac{\partial k}{\partial \omega}\right|_{\omega_0}(\omega - \omega_0) + k_0\left.\frac{\partial^2 k}{\partial \omega^2}\right|_{\omega_0}(\omega - \omega_0)^2\right] A(z,\omega)e^{-ik_0 z} = 0. \tag{2.31}$$

At this point we are finally ready to Fourier transform back to the time domain. The first term is again trivial

$$\mathscr{F}^{-1}\left[-2ik_0\frac{\partial}{\partial z}A(z,\omega)e^{-ik_0z}\right] = -2ik_0\frac{\partial}{\partial z}\mathscr{F}^{-1}\left[A(z,\omega)e^{-ik_0z}\right] \tag{2.32}$$

$$= -2ik_0\frac{\partial}{\partial z}A(z,t)e^{-i(k_0z-w_0t)}. \tag{2.33}$$

For the next two terms, we again use the Fourier identity given in Eq. 2.17. The second term, when evaluated, gives

$$\mathscr{F}^{-1}\left[2k_0\frac{\partial k}{\partial\omega}\bigg|_{\omega_0}(\omega-\omega_0)A(z,\omega)e^{-ik_0z}\right] = -2ik_0\frac{\partial k}{\partial\omega}\bigg|_{\omega_0}\frac{\partial A(z,t)}{\partial t}e^{-i(k_0z-w_0t)}. \tag{2.34}$$

The final term evaluates as

$$\mathscr{F}^{-1}\left[k_0\frac{\partial^2 k}{\partial\omega^2}\bigg|_{\omega_0}(\omega-\omega_0)^2A(z,\omega)e^{-ik_0z}\right] = -k_0\frac{\partial^2 k}{\partial\omega^2}\bigg|_{\omega_0}\frac{\partial^2 A(z,t)}{\partial t^2}e^{-i(k_0z-w_0t)} \tag{2.35}$$

Putting this all together and simplifying a bit, we find

$$\left[\frac{\partial}{\partial z}+\frac{\partial k}{\partial\omega}\bigg|_{\omega_0}\frac{\partial}{\partial t}-\frac{i}{2}\frac{\partial^2 k}{\partial\omega^2}\bigg|_{\omega_0}\frac{\partial^2}{\partial t^2}\right]A(z,t)=0. \tag{2.36}$$

Finally, we use the fact that $\partial k/\partial\omega$ is simply the reciprocal of the group velocity v_g, and introduce $k^{(2)}$ for the group velocity dispersion $\partial^2 k/\partial\omega^2$, and we get

$$\left[\frac{\partial}{\partial z}+\frac{1}{v_g}\frac{\partial}{\partial t}-\frac{i}{2}k^{(2)}\frac{\partial^2}{\partial t^2}\right]A(z,t)=0. \tag{2.37}$$

This is our final result for this section, giving the parabolic equation for pulse propagation. So long as the approximations we made are valid, this provides an accurate description for pulse propagation in a lossless medium with a linear, instantaneous, susceptibility. In the next section, we will see how this is modified by the reintroduction of the nonlinear polarization term.

2.1.2 The Wave Equation: Nonlinear Propagation

In the previous section we derived the parabolic wave equation, which describes pulse propagation in a material using only the linear polarization. Here we will bring the nonlinear polarization term back in, and see how it affects pulse propagation.

The nonlinear polarization P_{NL} enters the parabolic equation as a source term, and the parabolic equation becomes [28]

$$\left[\frac{\partial}{\partial z}+\frac{1}{v_g}\frac{\partial}{\partial t}-\frac{i}{2}k^{(2)}\frac{\partial^2}{\partial t^2}\right]A(z,t)=\frac{i\mu_o}{k_0}\frac{\partial^2 P_{NL}(z,t)}{\partial t^2}. \tag{2.38}$$

Here we have again assumed that the light is linearly polarized and propagating in the $+z$ direction, with a narrow frequency band around ω_0, and a narrow wave vector band around k_0.

How the nonlinear polarization actually ends up affecting the propagation depends on the exacts of the situation. In most (but not all) cases, only the lowest nonzero term in the nonlinear polarization is important. In non-centrosymmetric materials (e.g. KTA, BBO, quartz) this is the $\chi^{(2)}$ term. In centrosymmetric materials (e.g. sapphire, beryl, diamond), the third order term $\chi^{(3)}$ is the lowest nonvanishing higher order term. As the processes of interest for this laser are primarily $\chi^{(2)}$ processes, this is where we will focus our analysis.

The second order term in the nonlinear polarization is given by

$$P^2=\epsilon_0\chi^{(2)}E(z,t)^2. \tag{2.39}$$

Here, we will introduce a slight modification, and replace $\chi^{(2)}$ with d_{eff}, an effective nonlinear susceptibility which takes into account the crystal orientation, the direction of propagation through the crystal, and the polarization of the light. Details on how d_{eff} can be calculated are found in [13].

Let us assume that the initial electric field is composed of three pulses, of three distinct frequencies, ω_p, ω_s, and ω_i, with the condition that $\omega_p=\omega_s+\omega_i$. In this case, the full electric field is given by

$$E(z,t)=A_p(z,t)e^{i(k_pz-\omega_pt)}+A_s(z,t)e^{i(k_sz-\omega_st)}+A_i(z,t,)e^{i(k_iz-\omega_it)}+c.c.. \tag{2.40}$$

When this field is squared, we find the result has terms of several different frequencies, formed by different combinations ω_p, ω_s, and ω_i. The general strategy at this point is to expand the full second order nonlinear polarization term in Eq. 2.38, then group terms on the left and right hand

sides of the same frequency. When this is done, we end up with three equations, one for each of the frequencies ω_p, ω_s, and ω_i. Specifically, we find

$$\left[\frac{\partial}{\partial z}+\frac{1}{v_{g,p}}\frac{\partial}{\partial t}-\frac{i}{2}k_p^{(2)}\frac{\partial^2}{\partial t^2}\right]A_p(z,t) = -id_{eff}\frac{\omega_p^2}{4c^2k_p}A_s(z,t)A_i(z,t)e^{i\Delta kz} \quad (2.41)$$

$$\left[\frac{\partial}{\partial z}+\frac{1}{v_{g,s}}\frac{\partial}{\partial t}-\frac{i}{2}k_s^{(2)}\frac{\partial^2}{\partial t^2}\right]A_s(z,t) = -id_{eff}\frac{\omega_s^2}{4c^2k_s}A_p(z,t)A_i^*(z,t)e^{-i\Delta kz} \quad (2.42)$$

$$\left[\frac{\partial}{\partial z}+\frac{1}{v_{g,i}}\frac{\partial}{\partial t}-\frac{i}{2}k_i^{(2)}\frac{\partial^2}{\partial t^2}\right]A_i(z,t) = -id_{eff}\frac{\omega_i^2}{4c^2k_i}A_p(z,t)A_s^*(z,t)e^{-i\Delta kz} \quad (2.43)$$

where we have introduced a phase mismatch term $\Delta k = k_p - k_s - k_i$.

What we have here are three coupled differential equations, each one describing the evolution of one of the three applied fields. These equations can be used for any three wave mixing $\chi^{(2)}$ process, including sum-frequency generation $\omega_s+\omega_i \rightarrow \omega_p$, difference frequency generation $\omega_p-\omega_s \rightarrow \omega_i$, or optical parametric amplification $\omega_p \rightarrow \omega_s + \omega_i$.[2]

The subscripts p, s, and i have been used with optical parametric amplification in mind, where they refer to the pump wave, signal wave, and idler wave, respectively [3] . In this scheme, which will form the basis of this laser, the inputs are generally a strong pump wave, and a weaker wave at either the signal or the idler frequency. The output is a reduced pump wave, an amplified signal/idler, and a new wave at the idler/signal frequency (depending on whether the signal or idler is used as an input wave). In the photon picture of optical parametric amplification, a pump photon is being split into two photons, one at the signal frequency and one at the idler frequency. This process will be more fully covered later on in Chapter 3.

This gives a basic description of the nonlinear optical processes which will form the basis for

[2] In actuality, difference frequency generation and optical parametric amplification are identical processes. Each consists of one pump photon being split into two photons, one at the signal frequency, and one at the idler frequency. Experimentally, both cases typically utilize a strong input wave at the pump frequency ω_p, and a secondary input wave, usually at the signal frequency ω_s. In this case, the resultant output is a weakened pump wave, an amplified signal wave, and a new idler wave. (If the input consists of the pump and idler waves, the output will instead have an amplified idler wave and a new signal wave). In common usage, the term optical parametric amplification is more often used when the amplified signal wave is of importance, and the term difference frequency generation is more often used when the new idler wave is of importance.

[3] It should be noted there is no physical distinction that determines which of the resultant photons is the signal and which is the idler. There are two conventions for deciding what a particular wave is called. In the first convention, which will be used for the remainder of this thesis, the signal refers to the shorter wavelength of the two, and the idler refers to the longer wavelength. In the other convention, the signal is whichever one is used as an input along with the pump beam. Very often, it is the shorter wavelength of the two which is used as an input, and the conventions agree.

this laser. In the next section we will cover the intended use for this laser - an extreme nonlinear process called high harmonic generation.

2.2 Observation of High Harmonic Generation

High harmonic generation is an extreme nonlinear process that was first reported in 1987 by McPherson et al. [77]. In these experiments, a laser was focused into a gas, and it was seen that a fraction of the light was upconverted to odd higher order harmonics of the driving field. This alone would not have been a remarkable discovery - after all, the ability to upconvert light to higher harmonics was by now over two decades old. What was new was the relative intensities of the higher order harmonics. Recall that in the perturbative theory, the magnitude of the nonlinear susceptibility falls by a factor of $\sim e/a_o^2$ with each subsequent term in the expansion. The consequence of this is that the magnitude of the effects from any term should likewise fall rapidly with increasing terms. For this experiment, the standard perturbative theory predicted that the intensity of the generated harmonics would fall rapidly at increasing harmonic orders. And, for the first few harmonics, up to the ninth order, this was the case. However, above the ninth order harmonic, the harmonic intensity no longer fell rapidly with each subsequent harmonic. Instead, the harmonics after the ninth harmonic were all of roughly equal intensity up to the seventeenth harmonic, above which no harmonics were seen. The spectral measurement is shown in Fig. 2.1. This observation indicated a breakdown of the perturbative nonlinear theory. Whatever process was responsible for generating the harmonics was not accounted for by the expansion given in Eq. 2.2. Indeed, as we will see, high harmonic generation can only be described with a qualitatively different picture.

2.3 Three step model for high harmonic generation

A semi-classical model, known as the three step model, is often used to describe high harmonic generation. This model, first proposed by Corkum in 1993, consists of three distinct steps: ionization, propagation, and recombination.

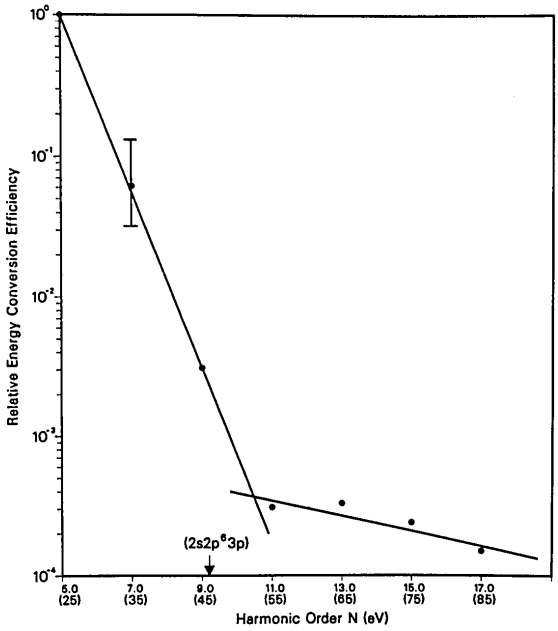

Figure 2.1: The first experimental measurement of high harmonic generation. Up to the ninth harmonic the intensity falls as predicted by the perturbative theory, but the eleventh through the seventeenth harmonic are all of approximately equal intensity. No harmonics above the seventeenth were observed. From [77]

2.3.1 Ionization

The three step model begins by considering a laser pulse focused into a gas. Within the atoms of the gas, the electrons are bound to the nucleus in a Coulomb potential,

$$V(r) = -\frac{Ze^2}{4\pi\epsilon_0 r}. \tag{2.44}$$

The potential of the laser, given by $e\mathbf{E}\cdot\mathbf{r}$, will distort the Coulomb potential seen by the electron. Fig. 2.2 shows the Coulomb potential, the laser potential, and the sum of the two. At high enough laser peak intensities, around $\sim 5\times10^{14}\mathrm{W/cm}^2$, the laser will distort the Coulomb field sufficiently to allow electrons to tunnel ionize away from their parents atoms [4] . The tunnel ionization rate was described by Ammosov, Delone, and Krainov in 1986 [6, 61], and in atomic units ($\hbar = m_e = e = 1$)

[4] Different ionization methods can occur at lower or higher peak intensities. Multiphoton ionization occurs at lower intensities, and barrier suppression ionization occurs at higher peak intensities. The peak intensities for these fall outside the range typically used in high harmonic generation.

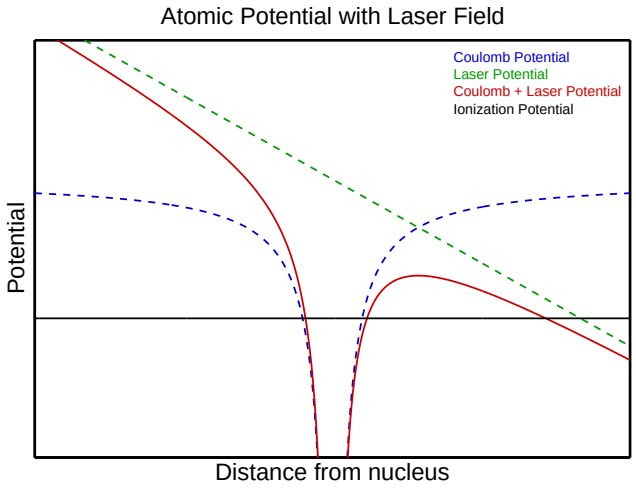

Figure 2.2: The potential seen by a bound electron in a laser field is the sum of the atomic Coulomb potential and the laser potential. At high enough laser field strengths, the laser will sufficiently distort the Coulomb potential to allow tunnel ionization.

is given by

$$\omega(t) = C^2_{n^*l^*}\left(\frac{3}{\pi}\right)^{1/2} I_p \frac{(2l+1)(l+|m|))!}{2^{|m|}(|m|)!(l-|m|)!}\left(\frac{2(2I_p)^{1/2}}{E_0}\right)^{2n^*-|m|-3/2}\exp\left(-\frac{2(2I_p)^{1/2}}{3E_0}\right). \quad (2.45)$$

Here I_p is the ionization potential, E_0 is the laser electric field magnitude, l and m are the orbital and magnetic quantum numbers. The factor $C^2_{n^*l^*}$ is given by

$$C^2_{n^*l^*} = \frac{2^{2n^*}}{n^*\Gamma(n^*+l^*+1)\Gamma(n^*-l^*)}. \quad (2.46)$$

The factors n^* and l^* are the effective principal and orbital quantum numbers, and are given by

$$n^* = \frac{Z}{\sqrt{2I_p}}, \quad (2.47)$$

and

$$l^* = n^* - 1, \quad (2.48)$$

where Z is the charge of the ion following ionization.

In a gas with initial atomic density N_0, and ion density N, the time-dependent fractional ionization $\eta = N(t)/N_0$ can be calculated

$$\eta(t) = \exp\left[-\int_{-\infty}^{t}\omega(t')dt'\right]. \quad (2.49)$$

As an example, the fractional neutral and singly ionized populations are plotted in Fig. 2.3. This calculation assumes a 3mJ, 10 fs, 800nm Gaussian pulse focused to a 100μm radius spot size in helium. This corresponds to a peak intensity of approximately $\sim 2 \times 10^{15}$ W/cm^2. An important

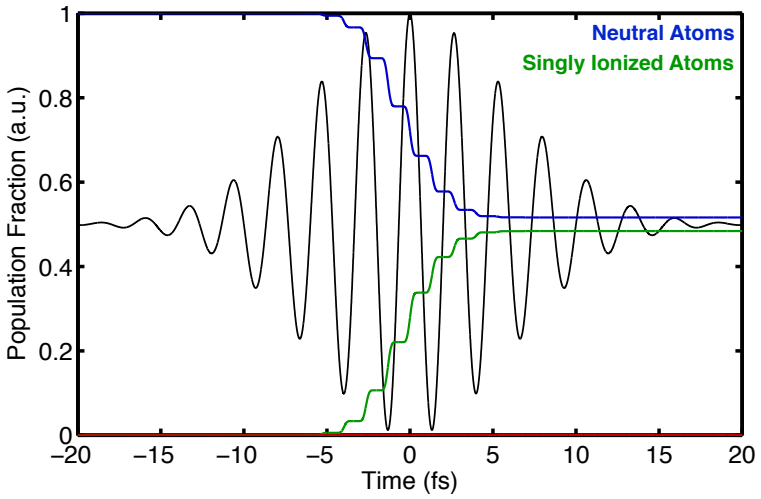

Figure 2.3: The fractional population of neutral atoms (shown in blue) and singly ionized atoms (shown in green) for ADK ionization in helium. The electric field is overlayed in black. The calculation assumes a 10fs Gaussian pulse, with 3mJ pulse energy at 800nm, focused to a peak intensity of $\sim 2 \times 10^{15}$ W/cm^2. This corresponds to a 100μm focal radius. Notice that the ionization occurs in steps at the peaks of the electric field, and turns off when the amplitude is near zero.

point to notice is that the ionization occurs in a series of steps at the peaks of the laser electric field. When the electric field amplitude is close to zero, the ionization shuts off.

2.3.2 Propagation

The next step in the three step model is the propagation step, which occurs after an electron has been freed from its parent atom through tunnel ionization. The newly free electron now finds itself in the presence of a strong, oscillating laser electric field. Assuming the laser is linearly polarized along the **x** direction, the laser electric field will accelerate the electron according to

$$m_e \ddot{\mathbf{x}} = -e\mathbf{E}\cos(\omega t + \phi_0), \tag{2.50}$$

where m_e is the electron mass, e is the electron charge, E is the magnitude of the laser electric field, ω is the laser's angular frequency, and ϕ_0 is the phase of the electric field at ionization, which is assumed to occur at time $t = 0$. The instantaneous velocity of the electron is obtained by integrating Eq. 2.50,

$$\dot{\mathbf{x}} = \frac{-e\mathbf{E}}{m_e\omega}(\sin(\omega t + \phi_0) - \sin(\phi_0)), \tag{2.51}$$

and by integrating again we find the instantaneous position as a function of time

$$\mathbf{x} = \frac{-e\mathbf{E}}{m_e\omega^2}(-\cos(\omega t + \phi_0) + \cos(\phi_0) - \omega t \sin(\phi_0)). \tag{2.52}$$

Fig. 2.4 shows the instantaneous electron position plotted as a function of time for four different values of the electric field phase ϕ_0 at ionization. As can be seen, the electron is first accelerated

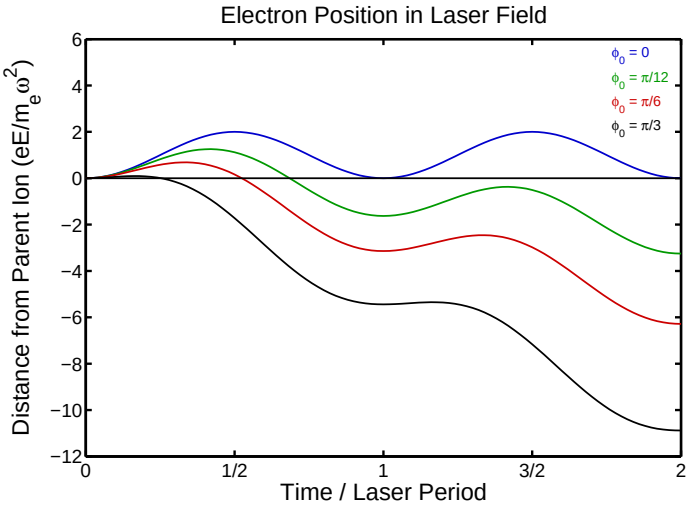

Figure 2.4: The position of the electron oscillating in the laser electric field as a function of time, plotted for two periods of the laser electric field at ionization. The position is plotted for four different phases of the electric field (ϕ_0) at ionization, shown in blue, red, green, and black. The parent ion, being much more massive than the electron, is fixed at $\mathbf{x} = 0$. These four initial phase values will all bring the electron back to its parent ion. For ionization at phase $> \pi/2$, the electron will never return to its parent ion.

away from its parent ion, but, when the electric field switches direction, the electron can be accelerated back towards its parent ion. If the phase at ionization is between 0 and $\pi/2$, the electron trajectory will bring it back to its parent ion at some time. For ionization with initial phase between

$\pi/2$ and π, the electron will never return to its parent ion. In cases where the electron does return to its parent ion, the final step, recombination, can occur.

2.3.3 Recombination

After propagating in the continuum, an electron which returns to its parent ion has a chance to recombine with its parent ion. When this happens, the electron will release its accumulated kinetic energy in the form of a single photon. By solving Eq. 2.52, we can find the electron's return time t as a function of the phase at ionization, ϕ_0. Then, plugging this into Eq. 2.51, we obtain the electron's velocity at recombination, which gives the electron's return kinetic energy. The return kinetic energy for the electron is plotted in Fig. 2.5 as a function of the driving laser's phase at ionization, ϕ_0. The return kinetic energy is given in terms of the ponderomotive enegry, U_p, given

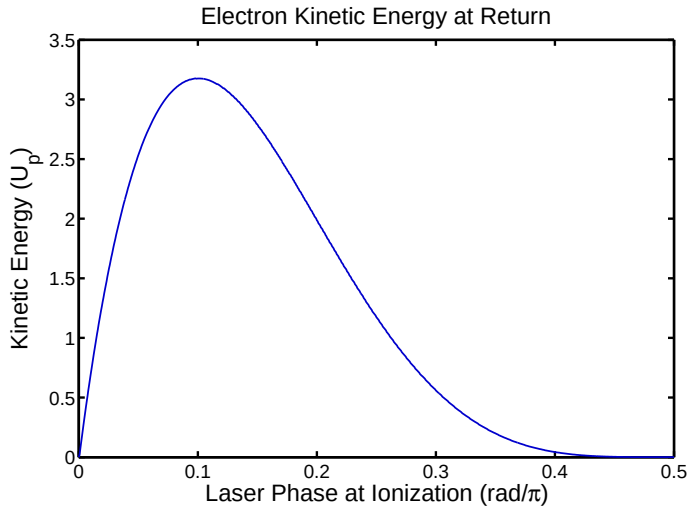

Figure 2.5: The kinetic energy of the electron upon its return to its parent ion, as a function of the laser phase at ionization. The maximum return kinetic energy for the electron is 3.17 times the ponderomotive energy U_p.

by

$$U_p = \frac{e^2 E_a^2}{4 m_e \omega_0^2}. \tag{2.53}$$

The ponderomotive energy is equal to the electron's time averaged kinetic energy, averaged for one cycle of the driving laser field.

As can be seen in Fig. 2.5, the electron return kinetic energy depends upon the phase of the driving laser field at ionization. Specifically, the electron returns with the maximum kinetic energy when ionization occurs at $\phi_0 \approx 0.09\pi$, where it returns with a kinetic energy of $3.17U_p$.

As the recombination step converts the electron's accumulated kinetic energy into a high energy photon, knowing the maximum return kinetic energy of the electron gives the maximum photon energy of the emitted x-ray. Specifically, the maximum photon energy x-ray that can be emitted is equal to the maximum return kinetic energy of the electron, plus the ionization potential of its parent atom. That is,

$$\hbar\omega_{max} = I_p + 3.17U_p. \tag{2.54}$$

No harmonics are seen above this energy.

Here we also see the first benefit to driving high harmonic generation with longer driving wavelengths. As the ponderomotive energy scales with the square of the driving wavelength, a longer wavelength laser focused to the same field strength can generate higher photon energy harmonics. This, however, is not the main reason for using longer driving wavelengths. Up until this point, the picture of high harmonic generation presented has only considered a single atom in isolation in the laser field. For high harmonic generation to serve as practical x-ray source though, we require a large, macroscopic, harmonic flux. The next section will transition from the single atom picture to the macroscopic picture, to show how bright, coherent x-ray beams can be generated. It is in the macroscopic picture that we will find the primary benefit to driving high harmonic generation with longer laser wavelengths.

2.4 Macroscopic High Harmonic Generation

Previously we'd covered the single atom picture of high harmonic generation, and presented a semi-classical three step model that can accurately predict some of the observed phenomena. However, generating an experimentally useful harmonic flux require large numbers of emitters, and so we will now transition to the macroscopic picture of high harmonic generation, and see how such

a large harmonic output may be created.

In general, there are two requirements for generating a large harmonic flux. The first requirement is rather straightforward: as a single atom emits a single photon, a large output flux necessarily requires a large number of atoms. The second requirement is more subtle: the individual harmonic contributions of all of the emitting atoms must all be emitted in phase with one another, so that the harmonics constructively interfere with each other.

Fortunately, there is a simple technology which accomplishes both of these goals. As Rundquist et al. demonstrated in 1998, by coupling the driving laser into a hollow, gas filled, dielectric waveguide, one can simultaneously greatly increase the number of emitting atoms, and ensure that they all emit in phase with one another [97]. It is easy to see how this allows for a larger number of emitting atoms: by coupling the laser into the waveguide at its focus, the waveguide allows the light to maintain the high peak intensities necessary for tunnel ionization for a much longer distance than a free space focus would allow.

To understand how the waveguide geometry satisfies the second requirement - that all the emitting atoms emit in phase with one another - we first need to know that the phase of any individual harmonic emission depends only on the phase the electron accumulates during the propagation step [64]. This means that the harmonic phase depends only on the laser phase at ionization. With this, we can simplify the requirement that all the individual atoms emit in phase with one another and say that the harmonic field must propagate with the same phase velocity as the driving field.

The requirement that the two fields propagate with the same phase velocity is called phase matching, and we can introduce a phase mismatch term Δk defined as

$$\Delta k = q k_\omega = k_q, \tag{2.55}$$

where k_ω is the wave vector of the driving laser field and k_q is the wave vector of the q^{th} harmonic, such that $q\omega_0 = w_q$. Phase matching is accomplished by setting $\Delta k = 0$.

When this is the case, as the driving pulse propagates through the waveguide, all of the harmonic emission from any given point will be emitted in phase with the existing harmonic field,

and the harmonic intensity will build up coherently over the length of the waveguide. When the phase mismatch is not zero, then the x-ray emission at different points along the waveguide will destructively interfere with one another. In this case, the harmonic intensity will initially build up, similar to the phase matched case, but, as the light propagates along the waveguide, the phase difference between the existing radiation and the emitted radiation will grow, and eventually they will destructively, rather than constructively interfere. At this point the harmonic intensity will begin to decrease. This oscillation will continue for the length of the waveguide. The period of the oscillations depends on the phase mismatch, and we can define a coherence length l_c, which is the distance over which the phase between the existing and emitted radiation slips by π,

$$l_c = \frac{\pi}{\Delta k}. \tag{2.56}$$

Practically speaking, this mean that two points in the waveguide separated by the coherence length will emit radiation π out of phase from each other, leading to destructive interference. A large phase mismatch between the driving field and the generated field will result in a short coherence length.

In a waveguide, there are several contributions to the dispersion of the fundamental driving field and the harmonic radiation, including neutral atoms, free electrons, and waveguide dispersion. When all of these are taken into account, the net phase mismatch in a waveguide is

$$\Delta k = q\frac{u_{11}^2\lambda_L}{4\pi a^2} - qp(1-\eta)\frac{2\pi}{\lambda_L}(\Delta\delta + n_2) + qp\eta N_a r_e \lambda_L, \tag{2.57}$$

where q is the harmonic order, u_{11} is the waveguide mode factor, a is the waveguide radius, λ_L is the driving laser wavelength, p is the pressure, η is the ionization fraction, r_e is the classical electron radius, N_a is the number density of atoms in the gas, $\Delta\delta$ is the difference in the refractive index of the gas between the driving laser wavelength λ_L and the harmonic wavelength λ_L/q, and $n_2 = \tilde{n}_2 I_L$ is the intensity dependent refractive index of the driving laser wavelength per atmosphere [1, 93]. The first term in Eq. 2.57 takes into account the waveguide dispersion, the second term contains the dispersion from the free atoms, and the final term takes the ionized electrons into account.

Phase matching high harmonic generation requires setting $\Delta k = 0$. To do this, we notice several things. First, in Eq. 2.57 the waveguide and ionized electron terms are strictly positive,

while the free atom term is strictly negative. Next, both the free atom and the ionized electron term depend linearly on the gas pressure p. These two facts suggest a strategy for phase matching: first, ensure the magnitude of the free atom term is larger than the magnitude of the ionized electron term. This ensures the sum of the two terms will be negative. Then, adjust the gas pressure, p, so that the sum of the two terms exactly offsets the waveguide dispersion term. In this case the net phase mismatch will be equal to 0.

The requirement that the magnitude of the free atom term is greater than the magnitude of the electron dispersion term leads to a critical ionization fraction η_{CR}, given by

$$\eta_{CR}(\lambda_L) = \left(\frac{\lambda_L^2 N_a r_e}{2\pi\Delta\delta + 1} \right)^{-1}. \tag{2.58}$$

When the ionization exceeds the critical ionization, the magnitude of the electron dispersion term is greater than the magnitude of the free atom dispersion term, and phase matching is impossible.

When phase matched, high harmonic generation has many attractive qualities. First, as we just showed it leads to a coherent buildup in the harmonic intensity over the length of the waveguide. As this is a coherent process, the output intensity scales as N_{emit}^2, where N_{emit} is the number of emitting atoms in the waveguide. The quadratic growth of the harmonics for phase matched emission allows for macroscopic x-ray output.

In addition to a large output flux, phase matched harmonic generation also has excellent spatial coherence properties. This can be seen in Fig. 2.6, showing a double pinhole diffraction pattern for phase matched harmonic generation. The 100% fringe contrast demonstrates the full spatial coherence for high harmonic generation.

This shows how phase matched high harmonic generation can serve as a practical source for coherent x-ray generation. In the next section, we will explore how high harmonic generation benefits from mid-infrared driving laser wavelengths.

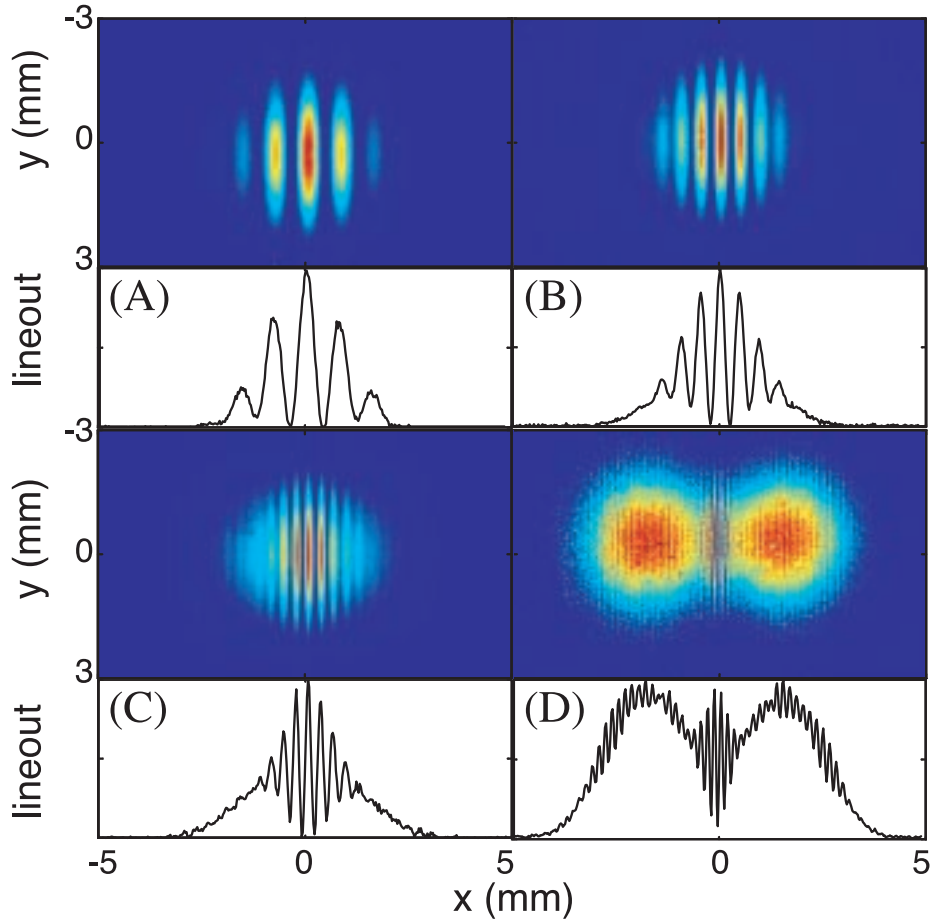

Figure 2.6: Double pinhole diffraction patterns for phase matched high harmonic generation. The different patterns reflect different pinhole spacing. The full fringe contrast demonstrates the full spatial coherence of the generated harmonics. Adapted from [11].

2.5 Mid Infrared High Harmonic Generation

The macroscopic picture of high harmonic generation gives the motivation for driving the process with mid-infrared lasers. The reason is that the harmonic photon energy phase matching cutoff scales favorably with the driving laser wavelength, so that longer driving wavelengths allow for higher photon energy x-ray beams. Specifically, the phase matching cutoff is expected to scale as $\lambda^{1.7}$ [93]. This scaling has has been experimentally demonstrated at multiple wavelengths in the near and mid-infrared. The first experiments, in 2008, showed the extension of the phase matching cutoff in Argon from 45eV to 105eV by increasing the driving laser wavelength from 800nm to

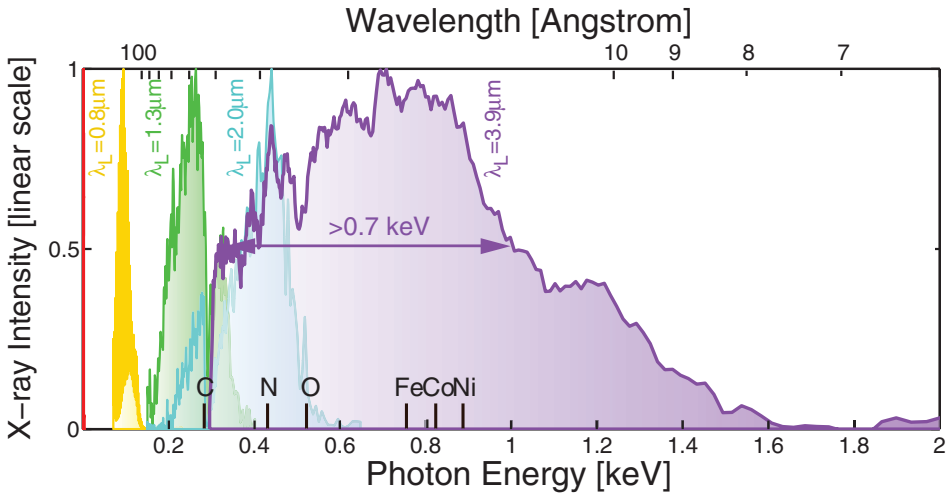

Figure 2.7: A plot of the normalized harmonic emission spectrum from several different driving laser wavelengths in helium. On the far left in yellow is 800nm, with 1300nm in green, 2000nm in blue, and 3900nm, in purple on the far right. When driven at 3900nm the phase matched emission spectrum extends beyond 1 keV photon energy, all the way up to 1.6keV. Adapted from [95].

1300nm [94]. Further work continued to push the driving laser further into the infrared, first at 1600nm, then 2100nm, and finally all the way into the mid-infrared at 3900nm [19, 93].

Fig. 2.7 shows a comparison of the generated harmonic spectra in helium as a function of harmonic photon energy for several different driving laser wavelengths. On the far left in yellow is a plot from an 800nm laser, then a 1.3μm in green, 2μm in blue, and finally a 3.9μm in purple. As can be seen, with a 3.9μm driving laser, the harmonic spectrum extends well beyond 1keV, with the phase matching cutoff around 1.6keV. It should be noted that with a 3.9μm driving laser and 1.6keV harmonic photon energy, this represents the coherent combination of >5000 driving laser photons, showing that high harmonic generation is truly an extreme nonlinear process. Furthermore, if we assume a flat spectral phase for the keV harmonic spectrum, we find a transform limit pulse length of approximately 2.5as.

The harmonics from mid-infrared driving lasers likewise demonstrate the full spatial coherence properties of harmonics driven with shorter wavelengths. This is shown in Fig. 2.8. Fig. 2.8a shows the harmonic beam profile from helium, demonstrating an excellent spatial mode, while Fig. 2.8b-c

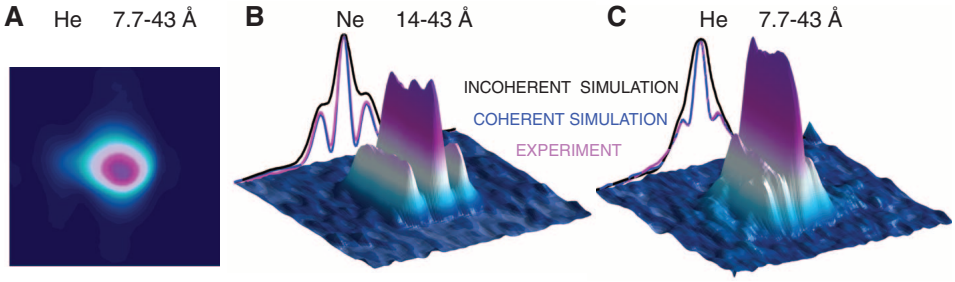

Figure 2.8: Images demonstrating the spatial coherence of keV high harmonic generation. In **a** we see a beam profile for the keV harmonic in Helium, demonstrating a laser like beam profile. In **b** and **c** we have Young's double slit diffraction experiments for Neon (**b**) and Helium (**c**), showing the full spatial coherence. Adapted from [95].

show Young's double slit diffraction patterns from neon (b) and helium (c), with a slit width of 5μm and slit spacing of 10μm.

This is the primary benefit of using mid-infrared driving lasers for high harmonic generation. The ability to phase match the process at keV photon energies makes high harmonic generation the only practical tabletop coherent soft x-ray source.

2.5.1 Isolated Attosecond Pulses from mid-IR Driven High Harmonic Generation

In addition to extending the phase matching cutoff to keV photon energies, mid-infrared driven high harmonic generation also provides perhaps the most direct path for generating isolated attosecond pulses. In fact, it turns out that isolated attosecond pulses are the intrinsic result of phase matched mid-infrared driven high harmonic generation [20, 48, 69]. This is due to the fact that the phase mismatch is not a constant across the temporal profile of the laser, but rather changes due to the changing ionization levels across the temporal profile. It turns out that for longer driving wavelengths, the phase mismatch varies much more rapidly across adjacent half cycles of the driving laser, so that true phase matching can only be achieved in a single half cycle. While for shorter wavelengths, several half cycles may have a coherence length l_c greater than the length of the interaction medium (ensuring a coherent harmonic buildup over the entire length of the medium), for mid-infrared driven harmonic generation only a single half cycle will have a

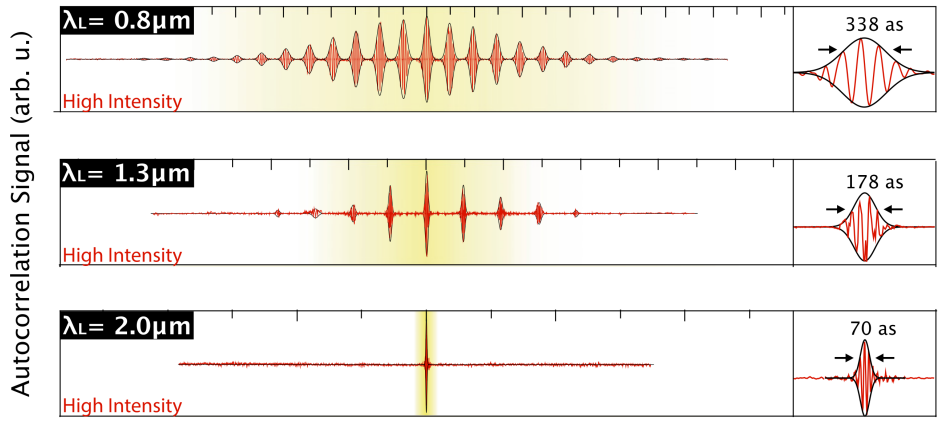

Figure 2.9: Experimental data showing the field autocorrelation signal of the harmonics from argon for three different driving laser wavelengths. Each trace was taken with a 10 cycle driving laser pulse and used the optimum phase matching pressure and intensity for the particular driving laser wavelength. With a near-infrared driving laser at 800nm, the harmonic emission shows several distinct x-ray bursts, leading to multiple peaks in the field autocorrelation. This is likewise true for 1300nm driven harmonic generation, although there are fewer distinct bursts than with 800nm. With mid-infrared driven harmonic generation however, the emission contains only a single isolated attosecond pulse, leading to a single peak in the field autocorrelation. Adapted from [20].

coherence length long enough to allow a coherent buildup over the entire length of the medium.

This has been demonstrated experimentally, as shown in Fig. 2.9, which shows the field autocorrelation signal of the generated harmonics in argon from three different driving laser wavelengths: 800nm, 1300nm, and 2000nm. Each of the field autocorrelation traces was taken with a 10 cycle driving laser pulse, and each was taken at the optimum phase matching pressure and intensity for the particular driving wavelength. As can be seen, with the shorter driving wavelength around 800nm, the field autocorrelation shows many distinct peaks, indicating numerous x-ray pulses. Each of the distinct bursts occurs during the phase matching window of a distinct half cycle in the driving laser pulse. In the autocorrelation trace for 1300nm, there are fewer distinct x-ray bursts, as the harmonic generation is phase matched across fewer half cycles of the driving laser. Notice also that the spacing of the individual bursts has increased, due to the longer driving laser wavelength. Finally, at 2000nm, there is only a single peak, indicating an isolated attosecond pulse. In this case harmonic emission is only phase matched over a single half cycle of the driving

laser pulse, so only a single pulse sees a coherent buildup for the length of the waveguide. As this is an intrinsic property of mid-infrared driven high harmonic generation, this provides a very simple route for generating isolated attosecond pulses.

2.5.2 Group Velocity Effects on Mid-Infrared Driven High Harmonic Generation

Further theoretical work has shed light on another surprising aspect of mid-infrared driven high harmonic generation. For some time the common wisdom has been that for high harmonic generation it is advantageous to use the shortest driving pulse possible. This is a reasonable supposition, as shorter driving pulses have higher peak intensities at any given energy and spot size. However, it turns out that this is not the case for mid-infrared driven high harmonic generation [20].

To understand this, the first thing to know is that mid-infrared driven high harmonic generation requires much higher gas pressures when compared to near-infrared or visible high harmonic generation [93]. For example, while high harmonic generation driven with a ti:sapphire laser at 800nm may use gas pressures on the order of tens of torr, high harmonic generation driven at 3μm may require gas pressures of several atmospheres [94]. This high gas pressure results in a large mismatch between the phase and group velocities in the driving laser pulse. The effect of this mismatch is a temporal walkoff between the pulse envelope and the electric field beneath it. Therefore, as the driving pulse propagates, the pulse envelope reshapes the electric field beneath it. For very short, few cycle pulses, where the envelope varies the significantly with time, even a small temporal slip will cause a significant reshaping of the electric field. In modifying the electric field, the phase matching is destroyed, so that the harmonic generation is no longer phase matched for even a single half cycle of the driving laser pulse. This prevents a coherent buildup, limiting the output flux.

The story is different for longer driving pulses, of around 8-10 cycles. While longer driving pulses still require very high gas pressures (as the required gas pressure depends on the driving laser wavelength), and thus still experience the exact same group velocity mismatch (as the group

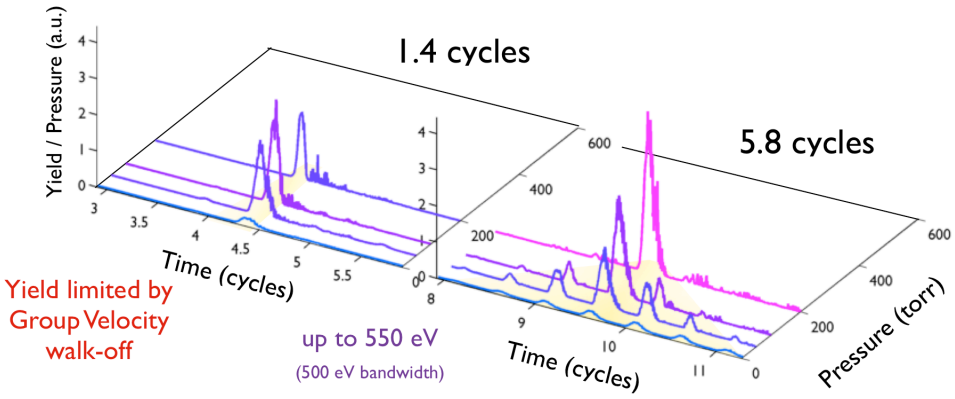

Figure 2.10: The effect of the group velocity mismatch on output harmonic flux. Plotted are the output harmonic flux as a function of gas pressure for two 2.0μm driving laser pulses, one 1.4 cycles and one 5.8 cycles. Notice for the 1.4 pulse, the output initially grows with increasing gas pressure, but then decreases as the pressure induced group velocity mismatch destroys the phase matching. For the 5.8 cycle pulse, the group velocity mismatch does not destroy the phase matching, and the harmonic output continues to grow with increasing pressure. Adapted from [49].

velocity mismatch depends only on the gas pressure), the difference is that the field envelope now varies much slower with time. As such the electric field is not reshaped to the same degree as it is for shorter pulses, and the phase matching can be maintained for the entire length of the waveguide, allowing for a coherent buildup.

Theoretical work has demonstrated this, shown in Fig. 2.10. This shows the predicted harmonic output in helium as a function of the gas pressure for both a 1.4 cycle and 5.8 cycle driving laser pulse, each at 2.0μm. For the 1.4 cycle pulse, the output harmonic flux initially grows as the pressure increases, as expected for phase matched conversion. However, as the gas pressure continues to increase, the group velocity mismatch continues to grow, and the reshaping of the field envelope grows more severe. Eventually, this reshaping destroys the phase matching, and the harmonic output begins to fall with pressure. For the 5.8 cycle pulse, the reshaping due to group velocity mismatch is not so severe, and the phase matching is preserved. In this case the output harmonic flux continues to grow with increasing pressure.

2.6 Conclusions

High harmonic generation is an extreme nonlinear process that can coherently combine thousands of driving photons to generate coherent x-ray radiation. While a full understanding requires a full quantum picture, a general understanding can be found in a simple three step model. In the three step model, a driving laser field first causes an electron to tunnel ionize out of an atom. After ionization, the electron finds itself in a strong, oscillating laser electric field. The field first accelerates the electron away form its parent ion, only to send it back towards the ion when the field switches direction. Upon returning to its parent ion, the electron can recombine with the ion, releasing its accumulated kinetic energy in the form a high energy photon. This simple model can predict the maximum photon energy that will be observed.

Generating a useful output flux requires phase matching the process over macroscopic length scales. To do this, the driving laser is coupled into a gas filled, hollow dielectric waveguide. The waveguide pressure can be tuned to ensure the driving laser and the generated harmonics travel with the same phase velocity. When this is the case, the harmonic emission from any one point in the waveguide will be emitted in phase with the emission from any other point in the waveguide. This allows for a coherent buildup over the entire length of the waveguide, and a large output flux.

Pressure tuned phase matching can generate large output fluxes up to the phase matching cutoff photon energy. The phase matching cutoff scales favorably with driving laser wavelength, so longer driving wavelengths can generate higher photon energy output. This has been demonstrated with phase matched conversion up to 1.6keV in high harmonic generation driven by a 3.9μm laser. This is the primary motivation to develop a mid-infrared laser for high harmonic generation. Additionally, mid-infrared driven high harmonic generation is perhaps the simplest way to generate isolated attosecond pulses. This is because in mid-infrared driven high harmonic generation the phase mismatch varies rapidly with time across the driving laser pulse, so full phase matching can only be achieved for a single driving half-cycle.

Theoretical work has shown that for mid-infrared driven high harmonic generation, using too

short of a driving laser pulse will also destroy the phase matching. Rather, it is expected that the optimal pulse length is on the order of 8 to 10 cycles, corresponding to 80-100fs for a 3μm laser.

With this understanding of the benefits of mid-infrared driven high harmonic generation, the next step is to understand how to build a high energy, high repetition rate mid-infrared laser. This will be the subject of the next chapter.

Chapter 3

Ultrafast Mid-Infrared Lasers

In the previous chapter we discussed the benefits of driving high harmonic generation with wavelengths in the mid-infrared region of the spectrum. With this motivation, we will now cover the mid-infrared region of the spectrum, and discuss different methods for generating ultrashort mid-infrared pulses. Three methods will be reviewed: conventional lasers sources, OPA sources, and OPCPA sources. This section will show why OPCPA technology is the most attractive path towards mJ-level mid-infrared laser systems. Following this, the theory behind OPCPA technology will be presented in detail.

3.1 The Infrared Region of the Spectrum

Traditionally the infrared region of the spectrum is defined as the region from the red end of the visible spectrum, around 700nm, all the way up to wavelengths of approaching 1mm. As this is an enormous spectral range, it is typically subdivided into smaller regions . There are numerous different schemes for dividing the infrared, which vary in how many subdivisions they contain, and where the boundaries are, but for laser science it is usually divided into three regions: the near-infrared, the mid-infrared, and the far-infrared. The exact boundaries vary amongst different disciplines, but for our purposes the near-infrared will refer to the region from 700nm through 3μm, the mid-infrared will refer to the atmospheric window from 3-5μm, and the far-infrared will refer to anything beyond 5μm.

Working in the mid-infrared, there are several factors that must be considered that are less

important when working in the visible or the near-infrared. Chief amongst these is atmospheric absorption, primarily due to water and CO_2. A modeled plot of the atmospheric transmission is shown in Fig. 3.1. This plot is calculated for the absorption seen by the Cerro Pachon observatory,

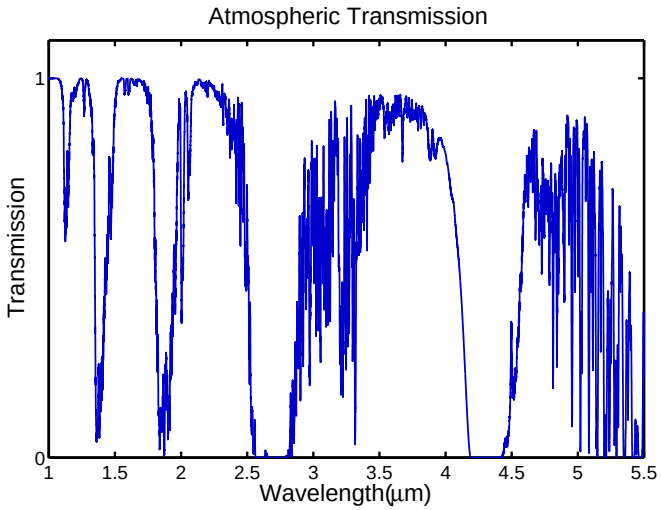

Figure 3.1: A coarse plot of the modeled atmospheric absorption from 1 through 5.5μm. The primary absorption is due to water and CO_2. Data calculated from [67] and provided by [84]. Modeled for the Cerro Pachon observatory, 4.3mm water column.

with a 4.3mm water column. As can be seen, there is significant absorption by the atmosphere in the mid-infrared region of the spectrum. On this plot, the water absorption is primary source of absorption, and is responsible for all of the structure, except for CO_2 absorption at 2.6μm and at 4.3μm.

In actual laser systems, water absorption tends to be a much bigger concern than that from CO_2. This is due to the tendency for optical coatings to absorb water from the environment. This is particularly true for lasers around 3μm, as the strong water absorption at 2.9μm can cause significant absorption in transmissive optics. Further, the power absorbed by water in the coating can heat the coating, leading to long term degradation, or thermally induced damage.

In addition to atmospheric and water absorption, many of the more common optical materials used for the visible and near-infrared, (e.g. BK7 or fused silica) are absorptive in the mid-infrared. This is not typically an issue for common items such as lenses, as there are plenty of optical

glasses that do work well in the mid-infrared, such as CaF_2, but can be an issue for such optics as waveplates, or nonlinear crystals. Typically there are materials that will work in the mid-infrared, but it is not uncommon to find serious drawbacks to these materials, most commonly increased price, decreased damage threshold, decreased efficiency, or size limitations.

3.2 Parameters for High Energy, Mid-Infrared Laser for High Harmonic Generation

We will next cover ultrafast, mid-infrared sources. Before this however, it is useful to briefly cover the desired output of the laser, informed by the requirements for high harmonic generation.

The first requirement is the wavelength of the light. As seen previously, the phase matching cutoff is set by the driving wavelength, with longer wavelengths allowing higher photon energy phase matching. Generating phase matched emission above 1 keV photon energy requires a driving wavelength of around 3μm. The next requirement is the repetition rate. High stability experiments necessarily require high repetition rate lasers. For this laser, we would like a repetition rate of at least 1kHz, and scalable to 10kHz.

For the pulse length, we recall the discussion in Chapter 2.5.2, where we saw that phase matched mid-infrared high harmonic generation requires relatively long pulses, of around 8-10 cycles. At the 3μm driving wavelengths, this corresponds to a FWHM pulse length of 80-100fs.

The desired pulse duration and the operating wavelength set the necessary bandwidth. To estimate the required bandwidth, let's assume a Gaussian temporal pulse, with an 80fs FWHM pulse length, with a central wavelength of 3.1μm, compressed to the transform limit. For a Gaussian pulse, the minimum time-bandwidth product is

$$\Delta\nu\Delta\tau \approx 0.44. \tag{3.1}$$

For an 80fs pulse, this sets the required FWHM bandwidth at 5.5×10^6MHz, or ≈ 176nm.

Next is the pulse energy. This is determined by the the necessity of ionizing the medium, and informed by the pulse length and the desired waveguide geometry. The waveguide geometry

benefits from the largest focal spot possible, as the attenuation length of the EH_{11} mode scales as r^{-3}, where r is the waveguide radius [70]. For a waveguide diameter of 150μm, and a pulse length of 80fs, reaching a peak intensity of 10^{14}W/cm^2 requires a pulse energy of approximately 750μJ, assuming a Gaussian temporal profile and a TEM_{00} Gaussian spatial mode[1] . So, in order to reach the high peak intensities required, we need pulses on the order of 1mJ.

The desired parameters for a mid-infrared laser for high harmonic generation are shown in Table 3.1. To summarize, we need a central wavelength at or above 3μm, running at 1kHz, with

Table 3.1: The desired output parameters for a mid-infrared laser for high harmonic generation.

Parameter	
Wavelength	$> 3\mu$m
Bandwidth	$>$ 300nm FWHM
Pulse Length	80-100fs
Pulse Energy	$>$ 1mJ
Repetition Rate	1kHz

pulse energies greater than 1mJ, compressed to <100fs. Next, we will discuss the various laser technologies that could be utilized for such a laser, beginning with conventional laser materials. Note that I will limit the overview to sources that can be utilized in a tabletop setting, and not cover free electron lasers and other facility type systems.

3.3 Technologies for mJ-Level, Ultrafast, Mid-Infrared Lasers

3.3.1 Mid-Infrared Laser Materials

In our review of mid-infrared sources that could be used for an ultrafast, high energy laser, we will first examine conventional laser materials, specifically limiting our review to materials with enough bandwidth to enable femtosecond pulse generation. This eliminates many potential mid-

[1] In a waveguide geometry the spatial mode is ideally the EH_{11} mode, which may be generally approximated as a TEM_{00} Laguerre-Gaussian mode for rough calculations.

infrared gain media, such as fluorine fiber lasers [2] , or quantum cascade lasers. In fact, the entire range of potentially viable mid-infrared sources with the necessary bandwidth is limited to a few solid state materials.

Currently, commercially available mid-infrared solid state gain materials are based upon one of two active ions, either Cr^{2+} or Fe^{2+} [79]. Cr^{2+} has been primarily used with ZnSe and ZnS as the host material, although CdSe and CdMnTe have also been used. To date, Fe^{2+} has only been used in ZnSe [76, 50].

Thus far, the majority of the work on mid-infrared solid state lasers has utilized Cr^{2+}. When used in ZnSe or ZnS, it has a broad gain spectrum, ranging from around 1.6μm to just over 3μm, while in CdSe the gain spectrum ranges from 2μm to approximately 3.5μm [108]. When used in CdMnTe, the gain spectrum runs from about 2μm to about 2.7μm [50].

In Cr:ZnS and Cr:ZnSe these broad gain bandwidths have been used to generate sub-100 femtosecond pulses, although femtosecond systems have thus far been limited to nJ level pulse energies, and wavelengths around 2.5μm, [23, 109]. At pulse lengths above 100fs, a 300μJ, 300fs regenerative amplifier has been reported [81]. At still longer pulse lengths, nanosecond level systems have been reported with mJ-level pulses in Cr:ZnSe, and sub-mJ level pulses in Cr:ZnS [79]

In Cr:CdSe, no femtosecond systems have been reported, but it has been used in an 815μJ, kHz, nanosecond level system, as well as a 17mJ, 300μs system [76, 5].

More recently, Fe:ZnSe has gained attraction as a mid-infrared laser material. This is due to its broad emission spectrum, with about 1600nm of bandwidth centered at 4.35μm [79]. Despite this large bandwidth, it has not yet been used in any femtosecond laser systems. It has been used in some extremely high energy systems, producing microsecond level pulses with over 400mJ of pulse energy, and a tunable system with over 100mJ pulse energy [35, 36]. At nanosecond pulse lengths, it has been used to generate pulses with just under 5mJ pulse energy [83].

Thus far neither Cr^{2+} or Fe^{2+} has been used to generate mJ level pulses at sub-100fs pulse

[2] Fiber lasers have been used in various ultrashort lasers around 2-2.5μm, but not yet been demonstrated at the 3μm range of interest [110].

lengths. The closest that has been demonstrated is the previously mentioned 1GW peak power, 0.3mJ, 300fs system running at 2.5μm [81]. While these materials still hold promise for producing high peak power, mid-infrared pulses for high harmonic generation, they as of yet have not been shown capable of generating the high energy, ultrashort pulses necessary. Further, the host materials used primarily thus far, ZnSe and ZnS, have numerous disadvantages. They are both quite brittle and easily damaged, which is problematic for high energy, high power systems which may experience large thermal gradients. They can also only be doped via diffusion doping, which necessarily limits the thickness of the crystals that can be made to no more than a few mm in one direction, and can lead to inaccurate and inconsistent doping levels.

The dearth of attractive conventional laser materials for ultrashort mid-infrared pulse generation necessitates the use of other means to generate these pulses. Next, I will cover Optical Parametric Amplification (OPA) technology, which has been used in numerous systems to extend ultrashort pulse generation further into the infrared.

3.3.2 OPA Technology

A more practical means of ultrashort mid-infrared light is Optical Parametric Amplification (OPA) technology. In contrast to conventional laser technology, OPA technology utilizes a nonlinear interaction to amplify the desired pulse. It has been successfully used to generate ultrashort near-infrared and mid-infrared pulses, and has even been used to generate ultrashort pulses used for high harmonic generation [30].

Optical Parametric Amplifiers utilize three waves, a pump, a signal, and an idler, and are described by the parabolic equations derived in Chapter 2.1.2. In an OPA, at the photon level, one pump photon is split into one signal and one idler photon. Conservation of energy demands that

$$\hbar\omega_p = \hbar\omega_s + \hbar\omega_i. \tag{3.2}$$

By convention, the signal has shorter wavelength than the idler, so that $\omega_p > \omega_s > \omega_i$.

A typical OPA setup is rather simple, and a general layout is shown in Fig. 3.2. Typically, an

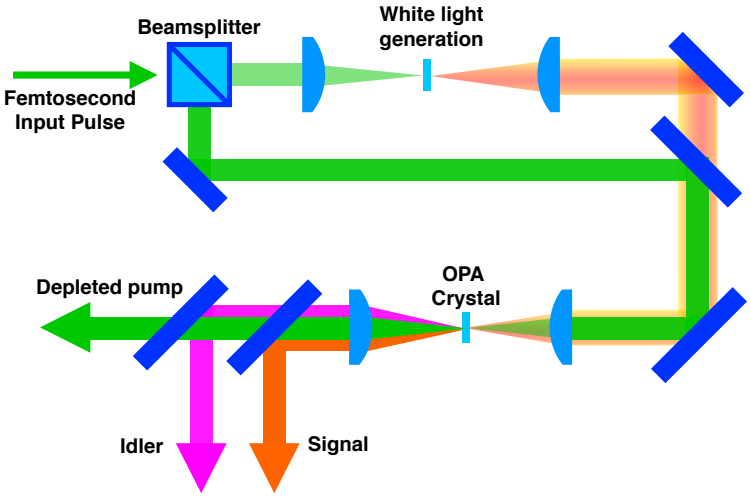

Figure 3.2: A typical layout for a single stage ti:sapphire pumped OPA. The input from the ti:sapphire is split into two arms. A small amount of the energy, usually on the order of microjoules, is used to generate a supercontinuum via white light generation. The white light is then used as the signal seed for an OPA stage, which is pumped by the remainder of the ti:sapphire input.

OPA to be used for infrared driven high harmonic generation begins with a mJ level, femtosecond ti:sapphire laser. The ti:sapphire input is split into two beams. A small amount of energy is used to generate a supercontinuum, via white light generation, typically by focusing into a dielectric such as fused silica, YAG, or sapphire, or by coupling it into a photonic crystal fiber [8, 117, 104, 99, 25]. The supercontinuum spectrum can extend well into the infrared, and provides the signal seed for subsequent amplification. Following the supercontinuum generation, the white light and the remainder of the ti:sapphire input are recombined, and spatially and temporally overlapped in a nonlinear crystal such as BBO, with the ti:sapphire input serving as the pump.

OPA technology has numerous benefits that make it attractive for use in ultrashort, infrared pulse generation. First, as it is most often used with a ti:sapphire amplifier, it relies on existing, well developed laser technology. In conjunction with this, it allows for relatively simple, small scale optical layouts. The simplicity of the layout and the robustness of the pump lasers allows for high stability pulse generation.

OPA technology also allows for relatively simple tunablility of the signal and idler wave-

lengths. As the actual output wavelengths are set by the phase matching conditions, the wavelengths can be tuned by adjusting the phase matching. In many systems this is as simple as adjusting the crystal angle (for angle tuned phase matching, such commonly used with BBO or KTP), or adjusting the crystal temperature (for temperature tuned phase matching, such as commonly used with LBO or PPLN). This allows a single system to cover a tremendous range.

Another benefit of OPA technology is that it typically does not suffer from the thermal effects that can limit conventional laser systems. Because it utilizes a nonlinear interaction in which energy is transferred directly from one pulse to another, only negligible energy is transferred to the nonlinear crystal. This means there is no significant heat load that must be extracted from the crystal, nor any thermal gradients in the crystal that lead to problems such as thermal lensing, stress induced birefringence, or thermal stress induced damage.

There is an additional benefit realized in OPA systems which use white light generation to generate the initial signal pulse. On these systems, because the same input pulse is split to provide both the pump, and the signal (generated via white light generation) for amplification, the carrier envelope phase (CEP) of the pump and signal pulses are necessarily locked to each other[3] [10]. As the idler carrier envelope phase depends on the relative difference between the pump and signal carrier envelope phase, this means the idler has intrinsic CEP stability. This can be useful for experiments using few cycle pulses, including high harmonic generation [52].

While OPA technology has its benefits, there are some limitations in its use as a practical source for high energy, ultrafast, mid-infrared pulses. The first one is a purely practical matter: many of the most commonly available nonlinear materials used in the visible and near-infrared are absorptive in the mid-infrared. For example, the transmission of BBO drops after 2μm, and is almost entirely absorptive above 3μm. [107]. While materials can still be used in regions where

[3] It is commonly said here that the signal and pump have the same carrier envelope phase. This is not strictly true. The carrier envelope phase of the pump and signal each depend on the carrier envelope phase of the input pulse, and any material each one passes through (which does not change shot to shot). In a system that does not have CEP-stable input pulses, the absolute carrier envelope phase of the pump and signal will each change shot to shot. However, the relative carrier envelope phase between the pump and signal will remain constant for each shot. As the idler carrier envelope phase depends only on the difference between the pump and signal carrier envelop phase, the idler carrier envelope phase does remain stable shot to shot.

they are partially absorptive, this both reduces the efficiency of the system, and deposits heat in the crystal that must be extracted.

Secondly, reaching mJ levels pulses around 3μm requires prohibitively high pump energies. To illustrate, suppose we want to generate a 1mJ, 3μm pulse from an 800nm ti:sapphire driving laser. For an 800 nm pump and 3μm idler, the signal wavelength is 1090nm. With these photon energies, generating 1mJ of idler requires generating 2.75mJ of signal. Therefore, we need to generate 3.75mJ total energy in the combined signal and idler. The total conversion efficiency of an OPA is often around 30%. So, generating the required 3.75mJ needs a pump pulse energy of 12.5mJ. Now this is certainly in the range of commercially available ti:sapphire systems, but if we need more than 1mJ at 3μm, then we very quickly require more pump energy than current state of the art ti:sapphire systems can deliver [4] .

Another significant limitation is simply that as we go to higher and higher pump energies, we need larger and larger crystals. This is particularly problematic for mid-infrared OPA systems, as several of the preferred nonlinear material that do work in the mid-infrared can only be made with very small aperture sizes. As an example, one of the materials used in this system is Magnesium Oxide doped Periodically Poled Lithium Niobate (MgO:PPLN), which is presently limited to aperture thicknesses of 3mm for commercially available products. For reference, to achieve an average intensity of 10GW/cm^2, with a 50fs pulse duration, and a 3mm mode diameter, only requires 35μJ of pulse energy.

Ultimately, it is these practical concerns which end up limiting OPA technology's usefulness for high energy, mid-infrared pulse generation. The energy requirements for the pump laser, and the size of the crystals required with high energy femtosecond pump pulses simply exceeds the current technology. In the next section we will cover a related technology, Optical Parametric Chirped Pulse Amplification (OPCPA), which does not suffer from these limitations. As we will see, OPCPA technology is the best option for generating high energy, ultrashort, mid-infrared pulses.

[4] Presently, 30mJ at 1kHz is the approximate upper bound for commercial systems.

3.3.3 OPCPA Technology

Optical Parametric Chirped Pulse Amplification (OPCPA) is a combination of the previously discussed OPA technology, and Chirped Pulse Amplification (CPA), which has been used for generating high energy, femtosecond pulses in the visible and near-infrared for several decades [111, 32]. Chirped Pulse Amplification, first developed by Strickland and Mourou in 1985, is a process where a pulse is first stretched in time, then amplified to high energies, and finally compressed in time to produce a high energy, ultrashort pulse. This basic scheme is likewise followed in OPCPA technology, with the difference being that the amplification step occurs in an OPA, rather than a conventional laser material.

Combing OPA technology with CPA technology provides a number of benefits that make OPCPA technology a more attractive choice for mid-infrared, mJ level, femtosecond laser systems. First, as a practical matter, by stretching the pulses in time the peak power for any given pulse energy is cut. As stretch factors in mJ level systems typically range from 10^3 to 10^5, the reduction in peak power can be quite substantial. This allows for a smaller mode size, and smaller optics in the system. Recall that one of the limitations of high energy, mid-infrared OPA systems is that many of the better nonlinear crystals for the mid-infrared can only be produced in relatively small sizes. By stretching the pulse and cutting the peak power, small aperture crystals can be used with higher energy systems while still keeping the peak intensity below the damage threshold.

In addition, reducing the peak intensity has the benefit of reducing the total nonlinear phase accumulated during amplification. The nonlinear phase B, also known as the B-integral, is given by the integrated product of the intensity $I(z)$ and the nonlinear refractive index $n_2(z)$,

$$B = \int I(z) n_2(z) \mathrm{d}z. \tag{3.3}$$

Since nonlinear phase is difficult to compensate, excessive nonlinear phase (usually defined as $B > 1$) makes recompression difficult [87]. By cutting the peak intensity, large stretch factors work to reduce this effect.

An additional benefit of OPCPA technology over OPA technology relates to the pump laser.

As an OPA requires a pump pulse of the same approximate pulse length as the desired final output pulse length, femtosecond OPA systems necessarily require femtosecond pump lasers. This limits the pump laser to sources with sufficient bandwidth to support femtosecond pulses. In OPCPA technology however, the pump pulse need not be compressible to the final pulse length. This allows for other pump sources, particularly Nd:YAG, Yb:YAG, and Yb:CaF_2, which can be run with very high pulse energies and pulse lengths between 10 and 100ps. In our system, the pump laser is based upon Yb:YAG.

While OPCPA technology has numerous benefits, it does have some drawbacks. Primarily, OPCPA technology requires compressing the final output pulse, whereas in OPA technology the output pulse comes out compressed. This is disadvantageous for three reasons. First, compression is always lossy to some degree. For a typical grating compressor with four reflections off the grating, even a 95% efficient grating will only run with ≈80% efficiency. Second, if there are any sources of uncompensated phase, whether they be from B-integral, or the amplification process itself, these can make compression difficult. The final drawback is one particularly important for our system. Traditionally, in a CPA system, the same pulse is both stretched and then compressed. For mid-infrared systems however, it is sometimes simpler to stretch and seed with a near-infrared signal pulse (in our case at 1.55μm), and then later on compress the mid-infrared idler (at 3μm for our system). This presents an additional challenge, as it is no longer the same pulse being stretched and compressed. The implications of this will be covered in greater detail in Chapter 6.

Despite the added inefficiencies and difficulties of compressing the pulse, OPCPA technology is still the most promising path for generating mJ level mid-infrared pulses. For conventional laser materials, there are simply not any robust materials in the mid-infrared region, while OPA technology is limited by the pump laser sources, and the size of the available crystals for the mid-infrared. OPCPA technology gets around this issue with stretched pulses, which allows the use of smaller aperture nonlinear crystals, and additionally allows for simpler pump technology.

3.4 OPCPA Theory

In the previous section we discussed the various technologies available for femtosecond, mJ-level, mid-infrared pulse generation. Of the three technologies covered, OPCPA technology is the most straightforward and promising. In this section I will give a detailed overview of the theory behind OPCPA technology. This will begin with an overview of three wave mixing in nonlinear media, the basis of which was introduced in Chapter 2.1.2. From here, I will introduce a simplified solution set for these couple equations valid under certain conditions. With this solution set I will analyze the gain properties of OPCPA systems, which will lead to a discussion of phase matching. This discussion will cover both direct phase matching and quasi-phase matching, which is of particular importance to this system. In the case of quasi-phase matching, the use of chirped crystals to increase the bandwidth will be discussed. Lastly, I will cover the bandwidth limitations of OPCPA technology. This will cover bandwidth limitations from both the phase matching bandwidth and temporal limitations arising from the chirped signal pulses.

3.4.1 Three Wave Mixing

The basics behind three wave mixing were covered in Chapter 2.1.2. The main result was a set of three coupled differential equations describing the evolution of three fields of frequencies ω_p, ω_s, and ω_i, with $\omega_p = \omega_s + \omega_i$. The equations are:

$$\left[\frac{\partial}{\partial z} + \frac{1}{v_{g,p}}\frac{\partial}{\partial t} - \frac{i}{2}k_p^{(2)}\frac{\partial^2}{\partial t^2}\right] A_p(z,t) = -id_{eff}\frac{\omega_p^2}{4c^2k_p}A_s(z,t)A_i(z,t)e^{i\Delta kz} \tag{3.4}$$

$$\left[\frac{\partial}{\partial z} + \frac{1}{v_{g,s}}\frac{\partial}{\partial t} - \frac{i}{2}k_s^{(2)}\frac{\partial^2}{\partial t^2}\right] A_s(z,t) = -id_{eff}\frac{\omega_s^2}{4c^2k_s}A_p(z,t)A_i^*(z,t)e^{-i\Delta kz} \tag{3.5}$$

$$\left[\frac{\partial}{\partial z} + \frac{1}{v_{g,i}}\frac{\partial}{\partial t} - \frac{i}{2}k_i^{(2)}\frac{\partial^2}{\partial t^2}\right] A_i(z,t) = -id_{eff}\frac{\omega_i^2}{4c^2k_i}A_p(z,t)A_s^*(z,t)e^{-i\Delta kz}, \tag{3.6}$$

where $\Delta k = k_p - k_s - k_i$ is the phase mismatch, and $A_n(z,t)$ is the spatially and temporally dependent electric field envelope for frequency ω_n. Here, some of the assumptions behind these equations should be recalled. First, we assumed that they are propagating in a lossless medium with instantaneous polarization response. Second, we assumed that each wave contains a relatively

narrow band of frequencies around its central frequency. Finally, we assumed that each of the pulse envelopes varies relatively slowly in space when compared to the carrier frequency of that field. In general, all of these assumptions work well for OPCPA systems[5] .

For OPCPA systems, the fields indicated by the subscripts p, s, and i refer to the the pump, signal, and idler, respectively. However, these equations can describe any three wave mixing process for which $\omega_p = \omega_s + \omega_i$.

3.4.2 Gain and Phase Matching

Typically, for a system described by Eqs. 3.4 - 3.6 the goal is to predict the signal and idler output, given the pump, signal, and idler inputs. In many situations, both the temporal and spectral profiles are of interest for the signal and idler outputs. Unfortunately, the coupled Eqs. 3.4 - 3.6 can not be solved analytically, and must be solved numerically. Often this is done with a Split-Step Fourier Algorithm, which treats the nonlinear and dispersion terms in Eqs. 3.4 - 3.6 separately [105]. By doing so, the coupled equations are broken into two sets of equations, one set describing the nonlinear interaction, and the other describing material dispersion. The method iterates between solving each of the two equation sets for a small propagation step through the crystal. This method will be used later on to evaluate several far-infrared OPA designs in Chapter 6.

While in general, the gain of an OPCPA system can not be calculated analytically from the coupled field equations, closed form solutions are available in certain approximate cases. For the remainder of this section, we will utilize a frequently used set of solutions that assume a strong pump input, a weak signal input, and no idler input. With a couple of approximations, we can find general solutions for the signal and idler of the form

$$I_s(z) = I_s(0)\left(1 + \frac{\Gamma^2}{\gamma^2}\sinh^2(\gamma z)\right) \tag{3.7}$$

$$I_i(z) = I_s(0)\frac{\omega_i}{\omega_s}\frac{\Gamma^2}{\gamma^2}\sinh^2(\gamma z) \tag{3.8}$$

[5] While all three of these assumptions tend to hold for OPCPA systems, the last assumption may not be valid for OPA systems with few-cycle pulses. In these cases, the slowly varying envelope approximation should not be used.

where I_s and I_i indicate the signal and idler intensity, respectively, and the factor γ is given by

$$\gamma = \sqrt{\Gamma^2 - (\Delta k/2)^2}, \tag{3.9}$$

with

$$\Gamma^2 = \frac{2\omega_s \omega_i d_{eff}^2 I_p(z)}{n_s n_i n_p \epsilon_0 c^3}, \tag{3.10}$$

and $I_p(z)$ is the pump intensity [17].

As mentioned, these solutions depend on two approximations. First, the temporal derivatives in the coupled field equations given in Eqs. 3.4-3.6 have been neglected. Physically, this means that dispersion and its effects are being ignored. This can be a good approximation for long pulses used in conjunction with short crystals, where dispersion effects do not have the opportunity to markedly reshape the pulses, or cause temporal walkoff. Specifically, we can define a pulse splitting length for a crystal as

$$l_{split} = \frac{\tau_p}{\delta_{ps/i}} \tag{3.11}$$

where τ_p is the pump pulse length and $\delta_{ps/i} = 1/v_{gs/i} - 1/v_{gp}$ is the group velocity difference between the pump and the signal/idler. Physically, this is the crystal length over which the pump and the signal/idler pulses will temporally walk off from each other. For interaction lengths shorter than the pulse splitting length, temporal dispersion may often be neglected. The second approximation is that the pump pulse in never depleted in the interaction, so that $I_p(z)$ is assumed to be constant.

Unfortunately, these approximations do not hold for many OPCPA systems. In particular, for the system described here, the pump experiences significant depletion in the latter two amplification stages. To properly model these stages, numerical methods must be used. Despite this limitation, the approximate solutions given by Eqs. 3.7 and 3.8 can provide many useful insights for our system, specifically in regards to the phase matching conditions, the amplification bandwidth, and the factors that limit the bandwidth. Therefore we will use these solutions, recognizing that they can not be trusted for exact predictions, which must instead come from numerical models.

With the solutions provided in Eqs. 3.7 and 3.8, we can define the gain for the signal or the

idler as

$$G_{s/i}(z) = \frac{I_{s/i}(z)}{I_{s/i}(0)}. \tag{3.12}$$

With this, we now turn to the question of how OPCPA systems can achieve significant gain in the signal and the idler

The biggest factor in determining the output energy is the phase mismatch Δk. Here the phase mismatch operates the same as the phase mismatch introduced for high harmonic generation in Chapter 2. To see this, we return to the solutions introduced in Eqs 3.7 and 3.8. Notice that both depend on the factor $\gamma = \sqrt{\Gamma^2 - (\Delta k/2)^2}$. In the event $|\Delta k/2| > |\Gamma|$, the term γ is imaginary. In this case, the solutions for the signal and idler are oscillatory in z. Physically, this is similar to the effect of the phase mismatch in high harmonic generation. The pump, as it propagates, continuously drives new signal and idler emission via the nonlinear polarization. However, the pump and signal / idler do not travel with the same phase velocity. Therefore, once emitted, the signal and idler slip out of phase with the pump as they travel through the interaction medium. At first the phase difference is small, and the emitted signal and idler adds constructively with the existing signal and idler. However, as they continue to propagate, the phase difference grows until this is no longer the case. At this point, the pump is driving the emission of signal and idler that is out of phase with the existing signal and idler, and they destructively interfere. The net effect is the signal and idler intensity oscillate with distance, and that neither sees significant gain.

On the other hand, when $|\Gamma| > |\Delta k/2|$, then γ is real, and the signal and idler see exponential growth with distance. In the undeleted pump approximation, this growth continues indefinitely. In reality, this is of course never the case, but for a weak signal input, this exponential growth can be used to achieve very high gains - easily exceeding 10^4 in many systems, and exceeding 10^6 in others [118].

These two cases are shown in Fig. 3.3, where the signal gain is plotted for three different values of Δk. In red and blue are two values which give an imaginary γ, leading to oscillatory behavior. Notice that the a larger phase mismatch (in blue) has a shorter period of oscillation. In

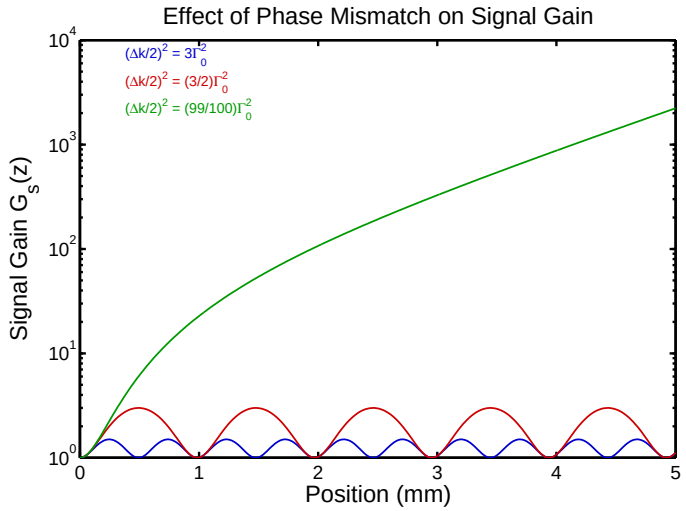

Figure 3.3: The signal gain in an OPCPA modeled with Eq. 3.7 for three different values of the phase mismatch Δk. This assumes a pump, signal, and idler wavelengths of 1.03 μm, 1.55μm, and 3.1μm, indices all equal to 1.5, a d_{eff} of 14.9 pm/V, and a pump intensity of 5 GW/cm^2. In red and blue are two cases where the phase mismatch is sufficiently large that γ is imaginary, which leads to oscillatory behavior. Larger phase mismatches leads to shorter oscillatory periods. In green $(\Delta k/2)^2$ is just slightly less than Γ^2, and gives exponential gain. Note that even though γ is still very small, the signal sees a gain $> 10^3$ in just 5mm.

green we have the case where $(\Delta k/2)^2$ is just slightly less than Γ^2, leading to exponential growth. Note that even with the values this close, the signal sees over three orders of magnitude of gain. For smaller phase mismatches the gain can be even higher. However, again at this point we must remember that these solutions assume a pump that sees no significant depletion. For very high gains, this may not be a valid assumption, at which point only numerical methods will give accurate predictions.

The challenge, therefore, is to ensure that $|\Gamma| > |\Delta k/2|$. Typically this must be done by minimizing Δk, as the only adjustable factor in Γ is the pump intensity, which is limited by the material damage threshold. In the next two sections we will look at two different types of phase matching, by which the phase mismatch Δk is set to zero.

3.4.3 Direct Phase Matching

The preferred scheme for phase matching is to directly phase match the process by finding a method to naturally set the phase mismatch Δk to zero. The most common method for directly phase matching a process is birefringent phase matching, which takes advantage of the polarization dependence of the refractive index in a birefringent material to set the phase mismatch to zero.

Birefringent materials are categorized by their number of optical axes. An optical axis is a direction of propagation within the crystal where the index of refraction shows no polarization dependence. Birefringent materials may have either one or two optical axes. Crystals with a single optical axis are called uniaxial, while crystals with two are called biaxial.

As a example, we can consider two pulses propagating in a crystal, with frequencies ω_1 and ω_2. They have wave vectors $\mathbf{k_1}$ and $\mathbf{k_2}$, and we will assume they are propagating collinearly so that $\mathbf{k_1} \parallel \mathbf{k_2}$. Let us assume the goal is to set $\mathbf{k_1} - \mathbf{k_2} = 0$. For simplicity, we'll look at the case of a uniaxial crystal.

In uniaxial crystals, the crystal z-axis is by definition set to be parallel to the optical axis. For light propagating so that its propagation direction makes an angle θ with the optical axis, we can decompose the polarization into two orthogonal polarization states. The first polarization state lies in the plane made by $\mathbf{k}$ and optical axis, and the second is perpendicular to the plane made by $\mathbf{k}$ and the optical axis.

The perpendicular polarization is known as the ordinary polarization, and it sees refractive index n_o. For this polarization, the index does not depend on the direction of propagation: it is n_o regardless of θ. The other polarization, lying in the plane, is known as the extraordinary polarization. For this polarization, the refractive index does depend on the direction of propagation. Specifically, the refractive index ranges from n_o when the $\theta = 0^\circ$ and the light is propagating along the optical axis, to n_e, when $\theta = 90^\circ$ and the light is propagating perpendicular to the optical axis. At an arbitrary angle θ, the index for the extraordinary ray is given by [13]

$$\frac{1}{n_e(\theta)^2} = \frac{\sin^2\theta}{n_e^2} + \frac{\cos^2\theta}{n_o^2}. \tag{3.13}$$

With this we can see how we can set $\mathbf{k_1} - \mathbf{k_2} = 0$ for two different frequencies. As the refractive indices also depend upon frequency, this can be accomplished by polarizing the two frequencies orthogonally to each other, and setting θ so that

$$\frac{1}{n_o(\omega_1)^2} = \frac{1}{n_e(\omega_2)^2} \tag{3.14}$$

$$= \frac{\sin^2\theta}{n_e(\omega_2)^2} + \frac{\cos^2\theta}{n_o(\omega_2)^2}. \tag{3.15}$$

This general method can be extended to more complicated situation involving more waves, non-collinear propagation, or biaxial crystals [121, 12], but the general strategy remains the same. By taking advantage of natural crystal birefringence, the phase mismatch may be set to zero for many nonlinear interactions. However, it is not always possible. In some materials, it is simply impossible to satisfy Eq. 3.15. In other situations, it may be possible, but the nonlinear interaction strength d_{eff} is significantly weakend[6] . For this reason, alternative phase matching schemes have been developed. Most notably is quasi-phase matching, which we use significantly in this laser.

3.4.4 Quasi-Phase Matching

There are many materials for which direct phase matching is either impossible or impractical. One example is lithium niobate, which has many attractive properties for use in the mid-infrared. It is transparent up to 4.5μm, and has a particularly high nonlinear coefficient at $d_{33} = 25$ pm/V. However, to access d_{33} requires all the waves are polarized in the same direction, which makes birefringent phase matching impossible. For this reason, quasi-phase matching is used with lithium niobate and other materials which can not be directly phase matched. Our discussion of quasi-phase matching will have two parts, looking at the microscopic and macroscopic pictures individually. In the microscopic picture, we will see how specially engineered crystals can allow for a coherent buildup even with a nonzero phase mismatch. On the macroscopic scale, we will see how these crystals affect the final output parameters.

[6] Recall that in Chapter 2 we neglected the directional dependence of the nonlinear polarization expansion. In actuality, the interaction strength depends on the direction of propagation and the polarization of the light in the crystal. Practically, for birefringent phase matching to be feasible, not only must Eq. 3.15 be satisfied, but the nonlinear strength must be sufficiently strong as well.

3.4.4.1 Microscopic Quasi-Phase Matching

In quasi-phase matching, rather then try to set the phase mismatch Δk equal to zero, we instead engineer the crystal so that the crystal orientation periodically flips directions. This has the effect of periodically flipping the sign of d_{eff}, which in turn flips the sign of the phase of the generated signal[7] . The strategy is to set the periodicity of the poling such that the sign flips whenever the phase between the signal and the pump has slipped by π. Remembering that the coherence length l_c is defined as the length over the which the accumulated phase difference is equal to π, this means the poling periodicity Λ is given by

$$\Lambda = 2l_c. \tag{3.16}$$

Recalling that the coherence length can be expressed in terms of the phase mismatch, $l_c = \pi/\Delta k$, the periodicity can also be expressed in terms of the phase mismatch

$$\Lambda = \frac{2\pi}{\Delta k}. \tag{3.17}$$

The poling periodicity usually ranges from a few microns to several 10s of microns for common nonlinear materials and processes. Two examples are useful for illustration, each in lithium niobate. First, we can consider second harmonic generation of 1030nm to generate 515nm. For this process, the coherence length is approximately 3μm, which gives a poling periodicity of 6μm. Second, we have the OPCPA process which will form the basis for this laser, 1030nm $\rightarrow$ 1515nm + 3070nm. Here the coherence length is approximately 15μm, so the poling periodicity is 30μm [107]

Fig. 3.4 shows how quasi-phase matching allows for a coherent signal / idler buildup. This plots the signal intensity growth for the three cases of no phase matching (shown in blue), direct matching (shown in green), and quasi-phase matching (shown in red). The black arrows and the alternating light and dark blue backgrounds indicate the poled crystal orientation for the quasi-phase matched case. As expected, with no phase matching the intensity simply oscillates, and never appreciably builds up. When directly phase matched the intensity continually grows. Quasi-phase

[7] The reason d_{eff} changes sign is because flipping the crystal relative to the fields flips the sign of the second order susceptibility $\chi^{(2)}$.

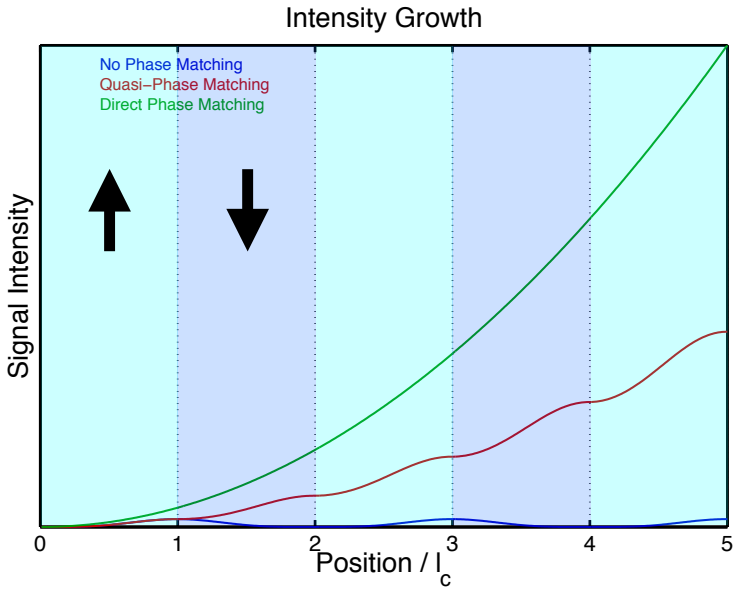

Figure 3.4: The intensity growth in a nonlinear system for the three cases of no phase matching (shown in blue), quasi-phase matching (in red), and direct phase matching (in green). With no phase matching, the intensity initially grows, but after one coherence length the intensity begins to shrink again. This leads to an oscillation of the signal intensity with no significant buildup. With direct phase matching, the intensity continually grows with distance. This is likewise true for quasi-phase matching, although the growth is reduced when compared to the directly phase matched case.

matching likewise leads to continual growth, although the growth is reduced when compared to the directly phase matched growth. This can be taken into account by using a modified effective nonlinear coefficient d_{eff}, and this will be further investigated in the next section which covers the macroscopic effects of quasi-phase matching

In this discussion we started with the assumption that the crystal orientation flipped at each coherence length. This is the simplest method of quasi-phase matching, but in reality the crystal orientation does not need to be flipped at every coherence length. In general, any poling period that is an integer multiple of $2l_c$ will also phase match that particular process. The case where $\Lambda = 2l_c$ is called first order quasi-phase matching, and we can defined the m^{th} order quasi-phase matching periodicity as

$$\Lambda_m = m2l_c. \tag{3.18}$$

For example, we had previously used the example of quasi-phase matching second harmonic generation in lithium niobate with a 1030nm pump. The first order QPM periodicity was about 6μm, however, with higher order quasi-phase matching, we could also use periodicities of 12μm, 18μm, and so forth.

This also brings up a practical consideration that should be taken into account whenever quasi-phase matching is used: coincidentally phase matched processes. A coincidentally phase matched process is a nonlinear process, other than the desired process, that just happens to be phase matched by the poling periodicity needed for the desired process. To see this, let's return to the other example process given: OPCPA with a 1030nm pump, 1550nm signal, and 3070nm idler. The first order QPM period for this is about 30μm. However, notice this is also the fifth order quasi-phase matching periodicity for second harmonic generation of the pump at 1030nm. This means that in addition to the OPCPA process, a 30μm period can also phase match second harmonic generation of the 1030nm pump.

In some situations coincidental phase matching can be advantageous, and indeed several systems have been reported which taken advantage of coincidental processes for various desired effects [8] [56, 88, 68, 26]. More often though, coincidentally phase matched processes are undesirable, and divert energy away from the desired process and into one or more undesired processes (these are sometimes called parasitic processes). Understanding possible coincidentally phase matched is important in the design of nonlinear systems, and particularly those based on quasi-phase matching. Fortunately, the effects of these competing processes can often be suppressed. Understanding this requires a macroscopic picture of quasi-phase matching, which we will next cover.

3.4.4.2 Macroscopic Quasi-Phase Matching

On macroscopic scales, quasi-phase matching is very similar to directly phase matched systems, although there are some differences. The first difference involves the output signal energy and

[8] In these systems the term "simultaneous phase matching" is more often used, rather than "coincidental phase matching", as these systems are usually designed to specifically take advantage of the ability to phase match multiple nonlinear processes.

conversion efficiency. As we saw in Fig. 3.4, the quasi-phase matched intensity growth is reduced when compared to the directly phase matched case. On the macroscopic scale, this reduced output can be expressed via a modified nonlinear coefficient. To see this, we first write the nonlinear coefficient as a Fourier series

$$d_{QPM}(z) = d_{eff} \sum_{m=-\infty}^{\infty} G_m e^{-ik_m z} \tag{3.19}$$

where d_{eff} is the material nonlinear coefficient, m is the quasi-phase matching order, k_m is the grating vector given by $k_m = 2\pi m/\Lambda$, and G_m are the Fourier coefficients. If the poling structure is that of a square wave (usually the case), then the Fourier coefficients are given by

$$G_m = \frac{2}{m\pi} \sin(m\pi D) \tag{3.20}$$

where D is the duty cycle of the poling (usually 0.5) [82].

For n^{th} order quasi-phase matching, only the $m = n$ term in Fourier expansion matters on the macroscopic scale. For first order $m = 1$ quasi-phase matching, this gives

$$d_{QPM} = \frac{2}{\pi} d_{eff}. \tag{3.21}$$

Quasi-phase matching thus results in a reduced nonlinear coefficient as compared to the crystal's effective nonlinear coefficient, with the first order coefficient being reduced by a factor of $2/\pi$. For the cases of higher order quasi-phase matching, or crystals with duty cycles other than 50%, the effective nonlinear coefficient will be even further reduced [82]. This can be important in reducing the effects of coincidentally phase matched processes, which we covered in the previous section. This can be seen by referring back to Eq. 3.20. We can immediately see that the nonlinear coefficient depends inversely on the phase matching order. So, referring back to our previous example, we saw first order quasi-phase matching an OPCPA process with a 1030nm pump, 1550nm signal, and 3070nm idler, also gave fifth order quasi-phase matching for 1030nm second harmonic generation. Here, the effective nonlinear coefficient for the fifth order quasi-phase matched second harmonic generation will be 1/5 of that for the first order quasi-phase matched OPCPA process.

In addition to the simple inverse dependence on order, the $\sin(m\pi D)$ term in Eq. 3.20 also means the effective nonlinear coefficient will completely vanish for certain higher orders. In these cases, the coincidentally phase matched process will have no effect. In the most common case, with a duty cycle of 50%, this occurs for the even higher orders. Other times, the duty cycle can be adjusted to suppress specific higher orders. Returning to our previous example, the sixth order second harmonic generation could be suppressed with a duty cycle of 60%. Here, it should be noted that analytical work has shown that control over random duty cycle errors in the crystal is important for the suppression of undesired higher order processes [86].

In addition to modifying the nonlinear coefficient, quasi-phase matching also modifies the phase mismatch on the macroscopic scale. For a grating vector $K = 2\pi/\Lambda$, the effective phase mismatch can be written as [91]

$$\Delta k_{QPM} = k_p - k_s - k_i - mK. \tag{3.22}$$

This modified phase mismatch should be used for any analytic calculations for the macroscopic gain in a system based on quasi-phase matched crystals.

With these two modifications, quasi-phase matched crystals behave substantially similar to directly phase matched crystals on the macroscopic scale. In the next section we will discuss the bandwidth limitations in OPCPA systems, and strategies for optimizing the bandwidth.

3.4.5 Bandwidth Limitations

In this section, we will look at the bandwidth of OPCPA processes, and the limiting factors. There are two primary factors that can limit the amplified bandwidth in an OPCPA system. The first is the phase matching bandwidth. This arises because the phase mismatch $\Delta k = k_p - k_s - k_i$ is not simply a constant value. This is because the signal and idler each contain a range of frequencies, and therefore Δk will in general vary across the signal and idler spectra. If Δk varies across the signal and idler spectra, this necessarily limits the bandwidth over which the process can be phase matched.

The second factor that will be covered are temporal bandwidth limitations. This limitation arises from the fact that the input signal is a chirped pulse - which means the spectrum has been mapped onto the temporal profile. Therefore, if there is any temporal structure in how the signal is amplified, this will in turn translate into selective spectral amplification.

Finally, following our discussion on the bandwidth limitations, we will cover one strategy to work around these limitations and amplify a large bandwidth - chirped quasi-phase matching.

3.4.5.1 OPCPA Phase Matching Bandwidth

To understand the phase matching bandwidth, we again start with the solutions described in Eqs. 3.7 and 3.8. We previously showed how the solutions depend on the phase mismatch Δk, and how the maximum output is achieved with a phase mismatch of zero (or an effective phase mismatch Δk_{QPM} of zero for quasi-phase matching). However, we must remember that the phase mismatch is calculated using the wave vectors of the carrier frequencies of the signal and idler, when in reality we have a range of frequencies around the carrier frequencies. Therefore, even if $\Delta k = 0$ for the central frequencies of the signal and idler, it is not necessarily zero across the entire spectrum for the signal and idler. As we move further away from the carrier frequency of the signal or idler, we expect the magnitude of the phase mismatch to increase. As the phase mismatch increases, the gain drops. We can define a phase matching bandwidth as the FWHM bandwidth over which the gain falls by a factor of 2 when compared to the gain at the carrier frequency. In this section, we will give an analytical estimate for the phase matching bandwidth[9] . Also, it should be noted that this section is equally valid for a quasi-phase matched OPCPA, with the exception that the effective phase mismatch Δk_{QPM} should be used rather than Δk.

We start by assuming an interaction length of L, so that we can set $z = L$ in Eqs. 3.7 and 3.8. We further assume that we are operating in the high gain regime, so that $\Gamma L \gg 1$. This allows us to replace $\sinh(\gamma L)$ with $\exp(\gamma L)$ in Eqs. 3.7 and 3.8. With these assumptions the expression

[9] Here we will assume that the pump spectrum can be approximated by a delta function without any concern for the finite bandwidth. For this laser, which uses a narrowband pump, this is a fine approximation, but for OPCPA systems with broadband pumps this is not a valid assumption. In these cases the bandwidth of the pump must be taken into account.

for the signal gain given in Eq. 3.12 now reads

$$G_s(L) \approx \frac{1}{4} e^{2[\Gamma^2 - (\Delta k/2)^2]^{(1/2)} L}. \tag{3.23}$$

When the phase mismatch $\Delta k = 0$, the gain takes on its maximum value, which we shall call G_0, and is given by

$$G_0 = \frac{1}{4} e^{\Gamma L}. \tag{3.24}$$

To find the phase matching bandwidth, we first find the phase mismatch for which the signal gain falls by a factor of 2 from its maximum value G_0. This is found by setting Eq. 3.23 equal to $G_0/2$, and solving for the phase mismatch. Doing this we find

$$\Delta k_{FWHM} = \pm \frac{2}{L} \sqrt{\Gamma L \ln 2 - \left(\frac{\ln 2}{2}\right)^2}. \tag{3.25}$$

This is how much the phase mismatch needs to change in order for the signal gain to fall to one half of its maximum value. The next step is to find what change in frequency for the signal and idler will provide this change in the phase mismatch.

To do this, we expand this phase mismatch around the carrier frequencies ω_s and ω_i. Keeping only the lowest order derivatives, we find that for a deviation $\delta\omega$,

$$\Delta k(\delta\omega) \approx k_p - k_s(\omega_s) - \left.\frac{\partial k}{\partial \omega}\right|_{\omega_s} \delta\omega - k_i(\omega_i) + \left.\frac{\partial k}{\partial \omega}\right|_{\omega_i} \delta\omega. \tag{3.26}$$

It may help to pause here and consider what Eq. 3.26 physically means. We started with the fact that the phase mismatch Δk is zero at the signal and idler carrier frequencies, ω_s and ω_i. Now suppose we detune each of the signal and idler away from their carrier frequency by equal and opposite amounts, $\delta\omega$, so that we have frequencies $\omega_s + \delta\omega$ and $\omega_i - \delta\omega$ (or vise versa). Eq. 3.26 gives the amount that the phase mismatch will change with this change in the signal and idler frequencies.

Returning to our calculation, we next use $\partial k/\partial\omega = 1/v_g$, where v_g is the group velocity. With this Eq. 3.26 becomes

$$\Delta k(\delta\omega) = k_p - k_s - k_i + \left(\frac{1}{v_{gs}} - \frac{1}{v_{gi}}\right) \delta\omega. \tag{3.27}$$

Since we have already specified that the phase mismatch is zero for the carrier frequencies, the first three terms in Eq. 3.27 sum to zero. This leaves

$$\Delta k(\delta\omega) = \left(\frac{1}{v_{gs}} - \frac{1}{v_{gi}}\right)\delta\omega. \tag{3.28}$$

Here, let us define the group velocity mismatch δv_g

$$\frac{1}{\delta v_g} \equiv \frac{1}{v_{gs}} - \frac{1}{v_{gi}}, \tag{3.29}$$

which is substituted into Eq. 3.28 to give

$$\Delta k(\delta\omega) = \frac{\delta\omega}{\delta v_g}. \tag{3.30}$$

We see here that the rate at which the phase mismatch changes in response to a small change in the signal/idler frequency depends on the group velocity mismatch between the signal and idler frequencies. Specifically, with a large group velocity mismatch, a small change in frequency away from the carrier frequencies will likewise cause a large change in the phase mismatch. On the other hand, if the group velocity mismatch is zero, then to first order, the phase mismatch does not vary at all with frequency around the carrier frequencies. In reality, even if the group velocity mismatch is zero, the phase mismatch will still change with frequency. To analyze this we would need to keep even higher order terms in the expansion given in Eq. 3.27. The general principle, however, is that maximizing the phase matching bandwidth requires minimizing the group velocity mismatch between the signal and idler.

This is explicitly seen by combining Eqs. 3.25 and 3.28, which gives

$$\Delta\omega_{FWHM} = \frac{4|\delta v_g|}{L}\sqrt{\Gamma L \ln 2 - \left(\frac{\ln 2}{2}\right)^2}. \tag{3.31}$$

This gives the FWHM phase matching bandwidth, around either the signal or idler, in the large gain regime. There are several things to notice here. First, as explained before it depends inversely on the group velocity mismatch, so that a larger mismatch leads to a smaller bandwidth. Second, it also depends inversely on the square root of the interaction length L. Therefore, a second strategy for increasing the bandwidth is to use as short an interaction length as possible. Finally, the

phase matching bandwidth scales favorably with the nonlinear coupling strength Γ. This is doubly beneficial, as for any desired signal gain, increasing the nonlinear coupling strength Γ allows one to shorten the interaction L. Fortunately, Γ depends on the pump intensity, so it is readily adjusted experimentally. So, to optimally maximize the phase matching bandwidth, this suggests one should pump as hard as possible (while staying below intensities that cause damage, self phase modulation, or other detrimental effects) and use the shortest possible interaction length.

3.4.5.2 OPCPA Temporal Bandwidth Limitations

In addition to the phase matching bandwidth, OPCPA systems can suffer from temporal bandwidth limitations. These are caused by the temporal profiles of the pump and signal / idler. To understand this, we again start with a system described by Eqs. 3.7 and 3.8. As in the previous section, we assume perfect phase matching so that $\Delta k = 0$, and work in the large gain limit $\Gamma L \gg 1$, where L is the interaction length. With these conditions, the signal gain again takes on its maximum value G_0,

$$G_0 = \frac{1}{4} e^{\Gamma L}. \tag{3.32}$$

Now, we remember that Γ itself is a function of the pump temporal profile, and that specifically $\Gamma \propto \sqrt{I_p(t)}$. This means that unless the pump has a square temporal profile with constant intensity, the gain profile will also have some temporal structure, specifically $G \propto \exp(\sqrt{I_p(t)})$.

Let us assume that the input pump pulse has a Gaussian intensity temporal profile, so that

$$I_p(t) = I_0 e^{-4\ln(2)\left(\frac{t}{\tau}\right)^2}, \tag{3.33}$$

where τ is the FWHM pulse length. With this, the gain takes a temporal profile

$$G = \frac{1}{4} e^{\Gamma_0 L \sqrt{I_0 \exp(-4\ln(2)(t/\tau)^2}}, \tag{3.34}$$

having defined $\Gamma = \Gamma_0 I_p(t)$. The pump intensity profile and the gain profile are plotted together in Fig. 3.5 for several different values of the peak gain G_0. Notice that higher peak gains correspond to shorter gain windows.

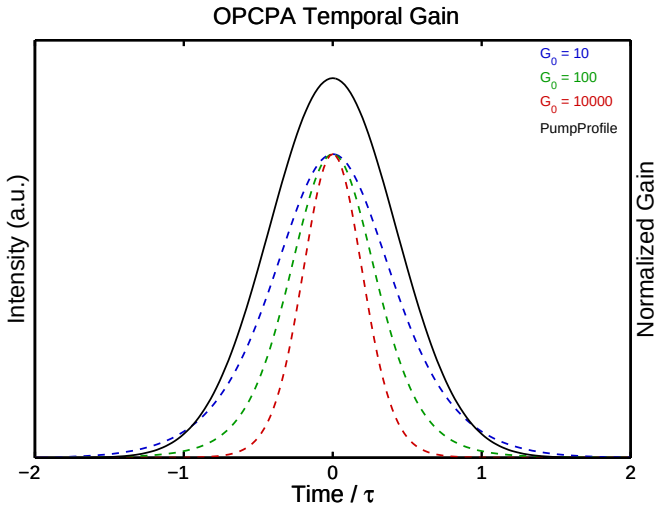

Figure 3.5: A comparison of the temporal profile for the pump (shown in black), and the gain profiles for several different peak gain values between 10 through 10,000. When the peak gain is low, the gain temporal profile is essentially the same width as the pump pulse, but as the peak gain increases the gain temporal width narrows substantially. The gain profiles have been normalized and scaled to 80% the height of the pump profile for clarity.

If we define $\pm t_g$ as the time interval over which the gain falls to $1/e$ of its maximum value G_0, then we find we can express the length of the temporal gain window as a function of the pump FWHM pulse length τ and the peak gain G_0 [80]

$$t_g = \frac{\tau}{2\sqrt{\ln 2}}\sqrt{-2\ln[1 - 1/\ln(4G_0)]}. \tag{3.35}$$

This matches what we saw in Fig. 3.5, with higher peak gains leading to a shorter temporal window over which amplification takes place. In OPCPA, this can lead to a reduced bandwidth, as the signal pulse is stretched in time. As the stretching maps the signal spectrum to the signal temporal profile, only the signal frequency components which fall into the gain window will see significant gain. Any signal frequencies outside the gain temporal profile will not be amplified.

To illustrate this, let's suppose we are amplifying a chirped signal pulse at carrier frequency ω_s. Let us further assume that the pulse has been chirped with pure GVD, and no higher order phase. This corresponds to a spectral phase of

$$\phi_s(\omega) = \frac{1}{2}\beta(\omega - \omega_s)^2, \tag{3.36}$$

where β is the total Group Delay Dispersion (GDD). Taking the Fourier transform, we can find the temporal phase, given by

$$\phi_s(t) = \omega_s t + \frac{1}{2\beta} t^2 \tag{3.37}$$

In this case the instantaneous frequency is

$$\begin{aligned} \omega(t) &= \frac{\mathrm{d}}{\mathrm{dt}} \phi_s(t) && (3.38) \\ &= \omega_s + \frac{1}{\beta} t && (3.39) \end{aligned}$$

As expected, this frequency of the signal pulse varies linearly with time, with the rate depending on the GDD β. If our signal pulse has a FWHM bandwidth of $\Delta\omega_{FWHM}$, then the stretched pulse has a FWHM pulse length of $\tau_s = \beta \Delta\omega_{FWHM}$.

With this, we can check the bandwidth that will fit in a time window defined by $\pm t_g$, the region of significant gain. Assuming that the signal is centered on the pump so that at time $t = 0$, $\omega(0) = \omega_s$, then the amplified bandwidth $\Delta\omega_{amp}$ is given by

$$\begin{aligned} \Delta\omega_{amp} &= \omega(t_g) - \omega(-t_g) && (3.40) \\ &= = \frac{2}{\beta} t_g && (3.41) \\ &= \frac{\tau}{\beta\sqrt{\ln 2}} \sqrt{-2\ln[1 - 1/\ln(4G_0)]} && (3.42) \end{aligned}$$

We see here that the amplified bandwidth depends inversely on the GDD β, so that the further the signal is stretched in time, the less amplified bandwidth there is. To preserve the entirety of the signal seed bandwidth, the signal should be stretched such that the $\Delta\omega_{FWHM} < \Delta\omega_{amp}$.

Both the phase matching bandwidth and the temporal amplification bandwidths limit the amplification bandwidth of OPCPA systems. While the temporal amplification bandwidth can be worked around by simply stretching the signal / idler less, the phase matching bandwidth is not so simple, and in many cases provides a firm limit on the total bandwidth, However, in systems based upon quasi-phase matched crystals, there is one relatively straightforward way to extend the phase matching bandwidth. As we previously discussed, in quasi-phase matching the spectral amplification region is set by the poling periodicity. So, to allow a crystal to amplify a

larger bandwidth, one can used chirped crystals, where the periodicity is not a constant, but varies through the length of the crystal. In these case different sections of the crystal will amplify different spectral regions. This method, Chirped Quasi-Phase Matching, is covered next.

3.4.5.3 Chirped Quasi-Phase Matching

Chirped quasi-phase matching is a strategy that can be used in quasi-phase matched crystals to extend the amplification bandwidth beyond what the phase matching bandwidth would normally allow. Recall that the phase matched wavelength depends on the poling periodicity. By changing the poling periodicity through the length of the crystal, the crystal can amplify a larger bandwidth than it could with a constant periodicity.

Whereas before for quasi-phase matched crystals we assumed a constant grating periodicity Λ, with chirped crystals we now have a periodicity that depends on the position within the crystal, $\Lambda(z)$. A diagram can be seen in Fig. 3.6. For a crystal of length L, the periodicity varies from

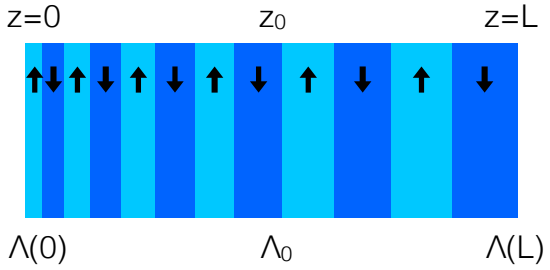

Figure 3.6: A schematic of chirped quasi-phase matched crystal of length L. The periodicity increases from the $\Lambda(0)$ to $\Lambda(L)$ through the crystal, taking on the value Λ_0 at $z = z_0$. Λ_0 is assumed to be the periodicity which perfectly phase matches the signal and idler carrier frequencies ω_s and ω_i.

$\Lambda(0)$ to $\Lambda(L)$. The phase mismatch now also depends on the position within the crystal, and again assuming collinear propagation normal to the crystal grating, we can write

$$\kappa(z) = k_p - k_s - k_i - mK(z). \tag{3.43}$$

Here we have introduced $\kappa(z)$, which is the spatially dependent overall phase mismatch that takes

into account both the intrinsic material phase mismatch Δk, and the spatially dependent crystal grating $K(z)$. We can further define a periodicity Λ_0, occurring at position z_0, so that $\Lambda(z_0) = \Lambda_0$, where the overall phase mismatch κ is zero for the signal and idler carrier frequencies ω_s and ω_i. That is,

$$\kappa(z_0) = k_p - k_s - k_i - mK(z_0) \tag{3.44}$$

$$\kappa(z_0) = 0. \tag{3.45}$$

For the remainder of this section we will work with $\kappa(z)$, rather than $\Lambda(z)$ or $K(z)$. This is because translating any particular grating pattern $\Lambda(z)$ into $\kappa(z)$ can be quite difficult, as doing so analytically requires an analytic expression for the spectrally dependent intrinsic material phase mismatch, $\Delta k(\omega)$, which typically can only be approximated via the Sellemeier Equations. For analytical evaluations, it is simpler to not bother with the actual grating pattern, but instead start directly with the effective phase mismatch.

While the overall phase mismatch κ can take on any number of functional forms, the simplest is a linear chirp through the length of the crystal. That is, we can write

$$\kappa(z) = \kappa(0) + \frac{\mathrm{d}\kappa}{\mathrm{d}z} z, \tag{3.46}$$

$$\frac{\mathrm{d}\kappa}{x} = \text{constant}, \tag{3.47}$$

where $\mathrm{d}\kappa/\mathrm{d}z$ is the chirp rate. We next wish to see how a linear chirp will affect the phase matching bandwidth. To do this, we will roughly follow the same procedure as in section 3.31, where we found the phase matching bandwidth for a directly phase matched crystal, or a quasi-phase matched crystal with a constant grating. In this procedure, we analyzed how small deviation in frequency $\delta\omega$ away from the signal / idler carrier frequency affected the intrinsic material phase mismatch, Δk. Here, we wish to see how a small frequency deviation $\delta\omega$ affects the spatially dependent overall phase mismatch, $\kappa(z, \delta\omega)$.

To start, we will first recast the overall phase mismatch as

$$\kappa(z) = \kappa(z_0) + \frac{\mathrm{d}\kappa}{\mathrm{d}z}(z - z_0) \tag{3.48}$$

to center it around z_0, the location in the chirped grating where the grating periodicity perfectly phase matches the signal and idler carrier frequencies. We next write the overall phase mismatch as a function of position z and the frequency $\delta\omega$ [18]

$$\kappa(z, \delta\omega) = -m\frac{\mathrm{d}K}{\mathrm{d}z}(z - z_0) - \frac{\delta\omega}{\delta v} \tag{3.49}$$

where we have made use of Eq. 3.47. What we have in this equation is the total phase mismatch for any point in the crystal, for any frequency deviation $\delta\omega$ from signal / idler carrier frequencies. For any given frequency deviation $\delta\omega$, we can determine where in the crystal it will be phase matched by setting Eq. 3.49 to zero. Alternatively, we can likewise use Eq. 3.49 to determine which frequencies will be phase matched at any desired position within the crystal.

At this point, we should recall that the overall phase mismatch need not be exactly zero for there to be gain. If this were the case, any given frequency would only see amplification at a single infinitesimally thin slice of the crystal. However, because this is not the case, any given frequency will see exponential gain over some finite length within the crystal. Let us call this length L_g. Next, we will determine how the grating chirp affects this length.

To do this, first remember that in Eq. 3.31 we saw that the in the high gain regime, the output signal gain can be expressed as

$$G_s(L) \approx \frac{1}{4}e^{[\Gamma^2 - (\kappa/2)^2]^{(1/2)}L}, \tag{3.50}$$

having replaced the intrinsic phase mismatch Δk with the total phase mismatch κ. Exponential gain will occur so long as $|\kappa| < 2\Gamma$. Next, suppose we know that frequency component $\delta\omega$ is perfectly phase matched at location z_m within the crystal, so that $\kappa(z_m, \delta\omega) = 0$. We want to determine what distance, Δz, we need to move away from z_m, such that $|\kappa(z_m \pm \Delta z, \delta\omega)| = 2\Gamma$. At this distance the frequency component $\delta\omega$ will no longer experience exponential gain. This distance can expressed simply by

$$\left|\frac{\mathrm{d}\kappa}{\mathrm{d}z}\Delta z\right| = 2\Gamma. \tag{3.51}$$

Finally, as $L_g = 2\Delta z$, we arrive at our end result,

$$L_g = \frac{4\Gamma}{\left|\frac{\mathrm{d}\kappa}{\mathrm{d}z}\right|}. \tag{3.52}$$

Here we have the length over which any single frequency component sees exponential gain.

At this point we can determine the total amplification bandwidth for a linearly chirped crystal. We will define the amplification bandwidth as the full range of frequencies which experience exponential gain for the full length L_g. Note that this definition exclude some frequencies which still see exponential gain, but for a distance less than L_g. In particular, notice that this definition excludes the frequencies which are phase matched exactly on the front and back edges of the crystal, as they would only see exponential gain for a distance of $L_g/2$. With this definition, the amplification bandwidth is

$$\Delta\omega_{CQPM} = \left|\frac{\mathrm{d}\kappa}{\mathrm{d}z}\delta v\right|(L - L_g). \tag{3.53}$$

There are several factors to notice here. First, the amplification bandwidth is linearly proportional to the crystal length. In contrast, earlier we found the amplification bandwidth scaled as $L^{-1/2}$ for direct phase matching or constant periodicity quasi-phase matching. Second, the bandwidth also depends linearly on the chirp rate. This makes intuitive sense, as for a given crystal length a higher chirp rate means a larger range of grating periodicities. Finally, the amplification bandwidth still depends on the group velocity mismatch between the signal and the idler.

There is one further wrinkle to understand about chirped quasi-phase matching: the increased bandwidth comes a price of decreased peak gain. Specifically, the peak gain is now given by [18]

$$G_{0,CQPM} = e^{\pi\Gamma^2/|\mathrm{d}\kappa/\mathrm{d}z|}. \tag{3.54}$$

Larger chirp rates, which give border amplification bandwidths, correspond to smaller peak gains.

This shows how the used of linearly chirped quasi-phase matched crystal can be used to extend the amplification bandwidth beyond what the phase matching bandwidth would normally allow. In addition to linear chirp patterns, more complicated grating patterns can be designed with specific goals in mind, such as broader amplification bandwidths than can be achieved with

linear chirped gratings [40, 89]. Note again, here, that chirped crystals don't affect any temporal bandwidth limitations - only adjusting the relative pump and signal / idler stretch can do that.

Chapter 4

Design of a High Energy, High Repetition Rate, Mid-Infrared OPCPA Laser

Previously, we've shown why mid-infrared lasers are useful for driving keV high harmonic generation, and why OPCPA technology is the most promising avenue for building a mid-infrared laser optimized for high harmonic generation. In this chapter we will turn to the actual design of our high energy, femtosecond, mid-infrared OPCPA laser. This will begin with a high level overview of the entire system, which is composed of three individual subsystems. The first of these subsystems, the front end, will then be covered in greater detail. The remaining subsystems will be covered individually in Chapters 5 and 6.

4.1 Full System Overview

A diagram of the full laser system is shown in Fig. 4.1. Broadly, the full system can be thought of as being composed of three individual subsystems: the front end, a cryogenic Yb:YAG laser, and an OPCPA laser. The front end is the foundation for the system, and serves two roles. First, it provides a 1.03μm seed pulse for the cryogenic Yb:YAG laser, and second, it provides a 1.55μm seed for the OPCPA laser.

The second component is the cryogenic Yb:YAG laser. This is a four stage laser running at 1.03μm, and it seeded by the front end. The final output from this laser is a 1kHz, >35mJ pulse train with 130ps pulses at 1.03μm. The output from the cryogenic Yb:YAG laser is used to pump the final component, the OPCPA system.

The OPCPA laser, like the cryogenic Yb:YAG laser, is seeded by the front end. Here, the

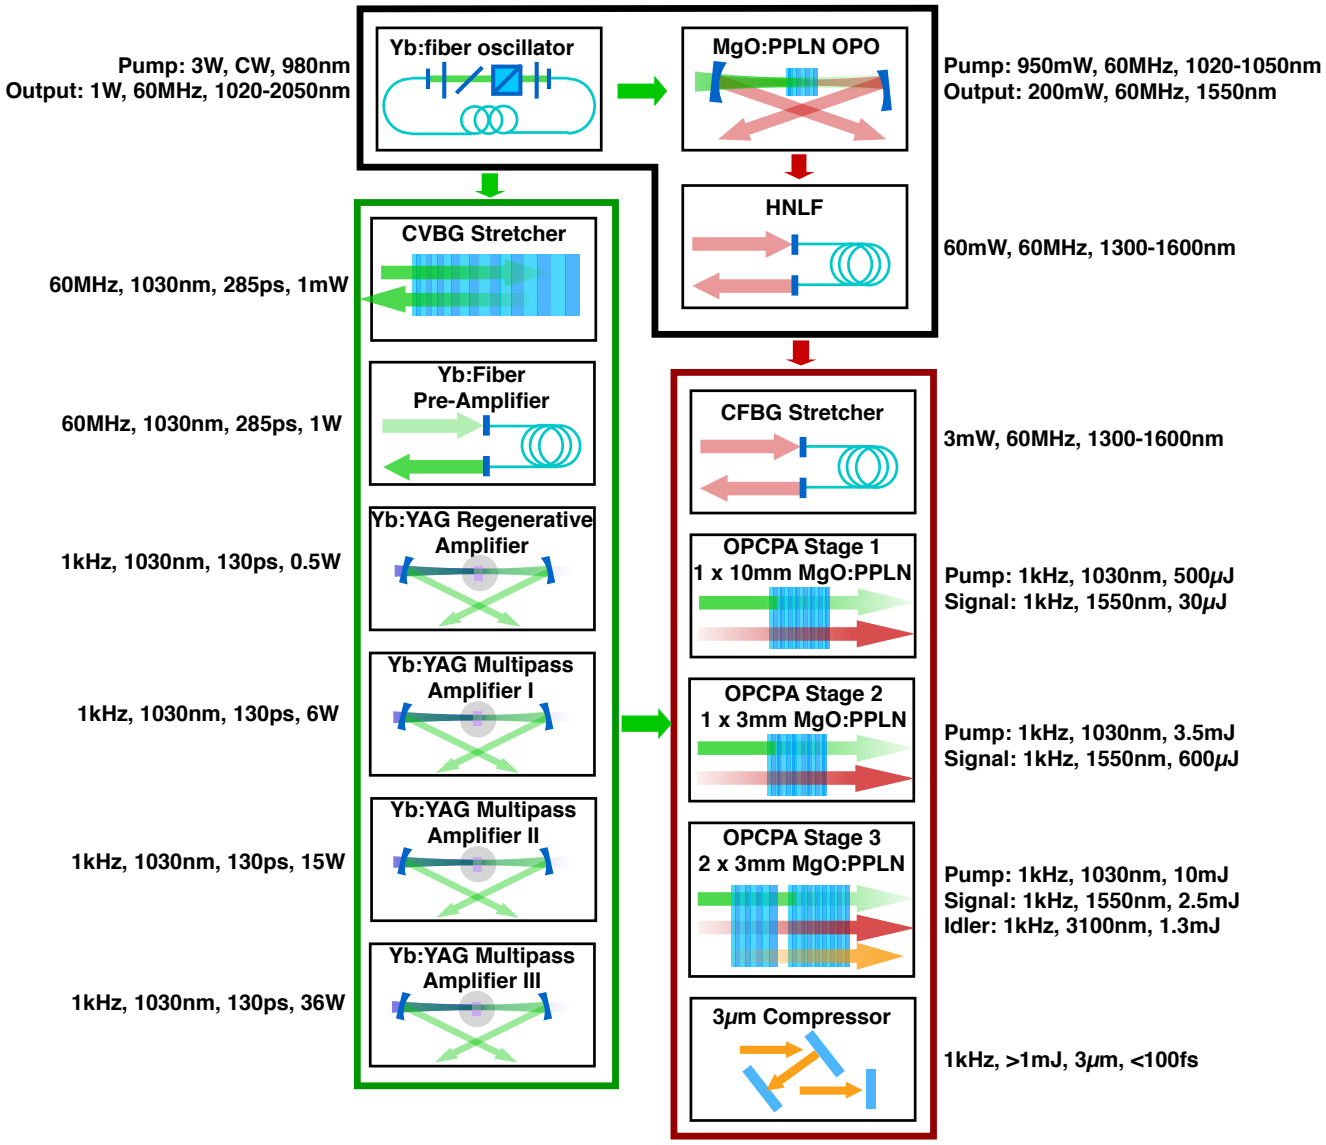

Figure 4.1: A diagram of the full laser system. The system can be thought of as being comprised of three subsystems - the front end (outlined in black), the cryogenic Yb:YAG laser (outlined in green), and the OPCPA laser (outlined in red).

seed is a 60MHz pulse train at 1.55μm. These are amplified in three OPCPA stages, each of which is pumped by the cryogenic Yb:YAG laser. The 1.55μm seed is amplified to >3mJ in these three stages, and the final OPCPA stage also gives a >1mJ pulse at 3μm, our final desired output. The 3μm output is compressed in a negative dispersion compressor, ideally to 100fs to be used for driving high harmonic generation.

4.2 Front End

The front end begins with a single Yb:fiber oscillator. The remainder for the front end, and indeed the remainder of the entire system, is based upon the single oscillator. Every other component in the system is either directly or indirectly seeded or pumped by this laser.

For an OPCPA system, this is hugely advantageous. Recall that in an OPCPA system, a pump and seed pulse are overlapped, spatially and temporally, in a nonlinear crystal. There are several requirements for maintaining this overlap. First, the pump and seed laser must run at the exact same repetition rate. Any discrepancy between the two repetition rates will cause a gradual temporal walkoff between the two pulse trains. Furthermore, in addition to locking the repetition rate of the seed and pump lasers, they must also run with a stable temporal relationship. That is, there must be a very well defined and controllable relationship between when the pump and the seed lasers fire. If this is not the case, there is no way to ensure that the pulses will overlap temporally in the crystal, even if the two lasers are running at the exact same repetition rate.

Fortunately, both of these requirements are intrinsically satisfied by basing the entire system upon a single modelocked Yb:fiber oscillator. As every component is based on this one laser, all components in the system naturally run at the same repetition rate. Similarly, the timing offset is likewise automatically locked.

The output from the Yb:fiber oscillator is split into two branches. The first branch, with approximately 50mW of average power, is used to seed the cryogenic Yb:YAG laser, which will be covered in Chapter 5. The second branch, with the remaining 950mW of average power from the oscillator, is used as the pump for an Optical Parametric Oscillator (OPO). The OPO is what

generates the 1.55μm seed that will be amplified in the OPCPA sections.

Next, we will cover the Yb:fiber oscillator and OPO in further detail, beginning with the Yb:fiber oscillator.

4.2.1 Yb:fiber Oscillator

A diagram of the Yb:fiber oscillator is shown in Fig.4.2. The oscillator begins with a 980nm

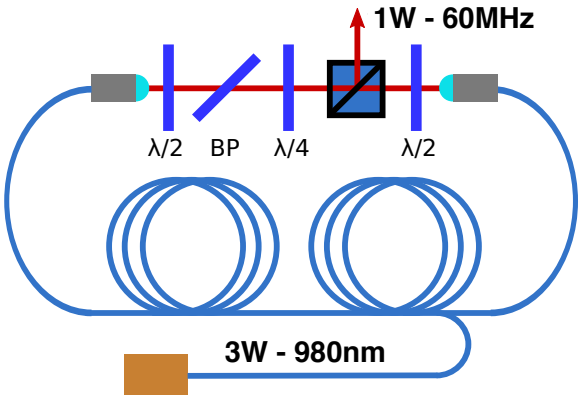

Figure 4.2: A diagram of the components in the ANDi Yb:fiber oscillator. A 2m length of Yb doped fiber is pumped by a fiber coupled diode, emitting 3W at 980nm. The cavity contains no anomalously dispersive elements, and the mode locking is provided by a strong nonlinear phase shift provided a quarter waveplate ($\lambda/4$ is the diagram) and a birefringent plate (BP in the diagram).

fiber coupled diode laser. The diode laser runs with 3W of CW power, which is coupled into a 2m section of double clad Yb doped fiber, with a 10μm core diameter as the active laser medium. This is an "All Normal Dispersion" (ANDi) laser, meaning that there is no intracavity dispersion compensation. Instead, the laser is modelocked via a nonlinear polarization provided by a $\lambda/4$-waveplate and a birefringent plate [21, 22].

The oscillator runs at 60MHz (set by the cavity round trip time), with 1W of average output power, or 16.6nJ pulse energy. The output spectrum is shown in Fig. 4.3, and consists of ~27nm FWHM bandwidth centered at 1040nm. Here, it should be noted that the spectrum from the oscillator covers 1030nm. This will be important later on in the system, because 1030nm is the peak of the emission spectrum for cryogenic Yb:YAG. As the spectrum from the Yb:fiber oscillator

covers the emission peak of Yb:YAG, the Yb:fiber oscillator can be used to seed the cryogenic Yb:YAG laser. The pulses are positively chirped coming out of the oscillator, with a FWHM

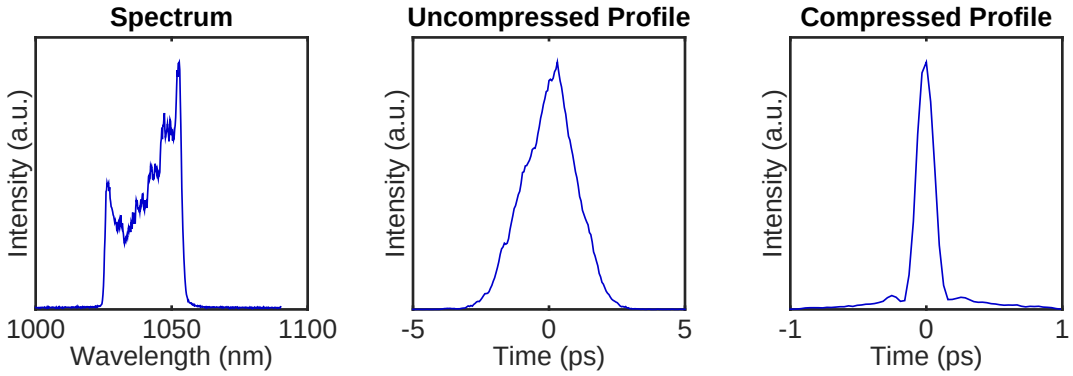

Figure 4.3: On the left, the output spectrum from the Yb:fiber oscillator. Notice the spectrum covers 1030nm. This is important, since it means the Yb:fiber oscillator can be used to seed the Yb:YAG laser. In the center, the uncompressed temporal profile from the oscillator, with an approximate pulse width of 2200fs. On the right, the compressed output from the Yb:fiber oscillator, with a pulse width of 150s.

pulse length around 2.2ps, but can be compressed with a simple negative dispersion stretcher. The uncompressed and compressed temporal profiles are shown in Fig. 4.3, with a compressed pulse length of 150fs.

The output from the Yb:fiber oscillator is split into two arms. One arm, with 50mW of average power, is used to seed the cryogenic Yb:YAG laser. The second arm, with the remaining 950mW of average power, is used to pump the next component of the front end, an Optical Parametric Oscillator.

4.2.2 MgO:PPLN Optical Parametric Oscillator

The front end utilizes an Optical Parametric Oscillator (OPO) to generate the 1.5μm seed for the OPCPA stages. This begins with the second of the two arms from the Yb:fiber oscillator, with 950mW average power. This is first compressed with a transmission grating compressor (remembering that the Yb:fiber oscillator output is positively chirped), and then used to synchronously pump an OPO to generate the seed for the OPCPA stages.

An OPO is similar to a conventional laser oscillator, with the difference that rather than a conventional laser material, such as Ti:sapphire, a nonlinear crystal is used for the gain medium. In this case the gain medium is Magnesium Oxide doped Periodically Poled Lithium Niobate (MgO:PPLN).

Lithium niobate is a very attractive nonlinear material. It has a very high d_{eff} at 25 pm/V, and is highly transmissive up to 4500nm. However, Lithium Niobate does not allow for direct phase matching, necessitating the use of quasi-phase matching techniques, as discussed in Chapter 3. For this reason, Lithium Niobate is typically grown in a periodically poled manner. The phase matching range of the crystal is set by the periodicity and the crystal temperature. To prevent the crystal temperature from drifting and negatively affecting the phase matching, crystals are typically housed in a temperature controlled oven. Furthermore, PPLN is also frequently doped with up to 5% mol. MgO, as this increases the photorefractive damage threshold [92, 29, 16]. The phase matching wavelength as a function of the poling periodicity is shown in Fig. 4.4 [107]. This assumes a 1030nm pump, and a crystal temperature of 100°C. As can be seen, a periodicity around

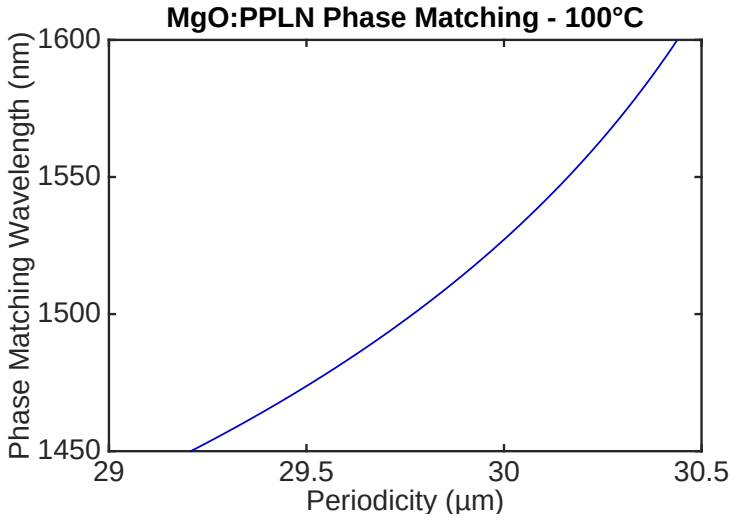

Figure 4.4: The phase matching range for MgO:PPLN, assuming a 1030nm pump and a crystal temperature of 100°C. Data calculated in SNLO. [107].

30μm is ideal for running with a signal output between 1500 and 1550nm.

While optical parametric oscillators are similar in many ways to conventional laser oscillators, there are important difference. First, as in all nonlinear interactions, an OPO requires a rather high pump intensity (usually on the order of $10^9 \mathrm{W/cm}^2$) in order to efficiently amplify the signal. This precludes CW pumping, as this would require very high average power levels to reach these intensities. For example, even assuming a 50μm focal diameter, reaching an intensity of $10^9 \mathrm{W/cm}^2$ would require an average power approaching 40kW. For practical purposes, pulsed sources are the only viable means for pumping an OPO.

The next complication is that because an OPO depends upon a nonlinear interaction to transfer energy directly from the pump pulse to the circulating intracavity signal pulse, the pump and signal pulses must overlap both spatially and temporally in the nonlinear crystal. For this reason, the round trip time of the OPO must exactly match the temporal spacing between pump pulses. To understand this requirement intuitively, consider the very first pump pulse that arrives at the crystal. When it arrives there is no signal, but the pump pulse produces a small amount of OPG. The OPG travels one round trip through the cavity, and returns to the crystal just as the second pump pulse arrives. The second pump pulse amplifies the OPG, which again circulates through the cavity, the whole process repeating itself until a stable, circulating pulse develops. However, should the cavity round trip time not match the time spacing between pump pulses, then the OPG from the first pulse will not return to the crystal when the next pump pulse arrives, and the OPG will not be amplified any further. This prevents the buildup of a stable circulating intracavity pulse.

Further complicating matters is the fact that the repetition rate of the Yb:fiber oscillator is also subject to temporal drift, and therefore the OPO cavity length must be actively stabilized to match the cavity length of the Yb:fiber oscillator. The locking mechanism utilizes a small amount of parasitic 623nm sum frequency generation between the pump and signal pulses [58]. The parasitic light is first sent off a diffraction grating to angularly disperse the spectrum. A fast silicon photodiode is positioned to read the spectral intensity at one of the tail ends of the spectrum. A PID lock reads the photodiode signal, and adjusts a piezo actuated mirror to lock the

OPO cavity length to that of the Yb:fiber oscillator.

The MgO:PPLN OPO runs with up to 200mW of output power, also at 60MHz, as the OPO repetition must match the repetition rate Yb:fiber oscillator. This gives an output pulse energy of 3.33nJ. An advantage of this OPO is that the output spectrum is easily tunable by adjusting either the crystal poling period, or by adjusting the crystal temperature. The OPO typically runs with around 25nm FWHM bandwidth, and can be tuned between 1.3μm and 1.6μm. The tuning range is limited by the reflection bandwidth of the cavity mirrors. The OPO typically runs with a central wavelength of 1.55μm, and the approximately 25nm FWHM bandwidth. The output spectrum is shown in Fig. 4.5. The output bandwidth can be pushed further, to around 45nm FWHM, but

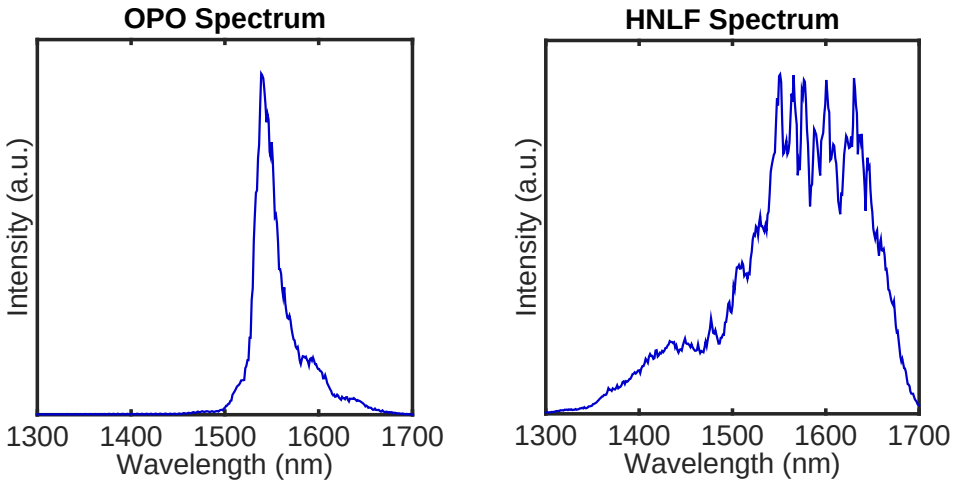

Figure 4.5: On the left, the typical output spectrum of the Yb:fiber pumped MgO:PPLN OPO. The central wavelength is 1.55μm, and the spectrum has ~25nm FWHM bandwidth. On the right, the spectrum after broadening in a section of HNLF. The spectrum has over 100nm of FWHM bandwidth, with almost 400nm of tail-to-tail bandwidth.

doing so reduces the output power and the stability of the timing lock.

The pulse length from the OPO is on the order of 150fs. Like the output of the Yb:fiber oscillator, the output of the OPO is positively chirped.

The OPO output is next sent through a 30cm section of Highly Nonlinear Fiber (HNLF), used to broaden the seed spectrum prior to amplification in the OPCPA stages. The HNLF is

a fiber with a very small effective mode area - usually between 10 and 20 μm^2. The small mode allows for a high peak intensity, even with a relatively modest pulse energy. This allows for spectral broadening due to self-phase modulation from the intensity dependent refractive index, n_2. The nonlinear properties of optical fiber are usually described with a nonlinear coefficient, γ, given by

$$\gamma = \frac{2\pi}{\lambda}\frac{n_2}{A_{eff}}, \tag{4.1}$$

where λ is the wavelength and A_{eff} is the effective mode area. In our case the effective mode area is 11.6μm^2, and the nonlinear coefficient γ is 11.5$W^{-1}km^{-1}$. An output spectrum from the HNLF fiber is shown in Fig. 4.5.

The input spectrum of approximately 25nm FWHM is broadened to ∼140nm FWHM, with almost 400nm bandwidth tail-to-tail. The power throughput of the HNLF is around 60mW, with the majority of the power loss occurring on the input coupling into the HNLF. The output from the HNLF provides the seed for the OPCPA stages.

To review, the front end begins with an Yb:fiber oscillator, which runs with 1W of output power at 60MHz, with a central wavelength of 1.04μm. Of the 1W of average power, 50mW are used to seed the cryogenic Yb:YAG laser, which will be covered in detail in Chapter 5. The output of the cryogenic Yb:YAG laser will subsequently be used as the pump for the OPCPA stages. The remaining 950mW are compressed, then used to synchronously pump an MgO:PPLN OPO. The OPO also runs at 60MHz, with 200mW of output power and a central wavelength of 1.55μm. The OPO output has 25nm of FWHM bandwidth, but the output of the OPO is broadened in an HNLF to greater than 100nm tail-to-tail bandwidth. This provides the signal seed for subsequent OPCPA amplification.

Chapter 5

Cryogenic Four Stage Yb:YAG Laser

The pump for the OPCPA system is provided by a four stage, cryogenically cooled Yb:YAG laser running at 1kHz. This laser is designed to deliver more than 30mJ of pump energy in 130ps pulses, at 1030nm. This chapter will cover the design and performance of this high energy, high repetition rate laser. First, I will discuss Yb:YAG and why it is an attractive option for high energy, high repetition rate lasers. Next I will go over the front end seed for the Yb:YAG system, which is based upon the ANDi Yb:fiber oscillator. Following the front end I will move in depth through each of the four stages in the system.

5.1 Yb:YAG

Cryogenic Yb:YAG is one of the most attractive materials for high average power lasers. While Yb:YAG has been used in room temperature systems, it is most often used in cryogenic systems as many of Yb:YAG's thermal and optical properties change favorably when cooled to cryogenic temperatures. High power CW systems have been developed with kilowatt level average powers [15, 96, 14]. In pulsed operation, it has been used for both actively and passively modelocked oscillators, with watt level average powers at repetition rates in the tens of MHz [44, 47]. Amplifier systems for pulsed operations have been developed for repetition rates ranging from kHz levels to the MHz levels, with average powers of hundreds of watts at higher repetition rates, to tens of mJ output energies at the kHz level [98, 51, 43].

A summary of Yb:YAG's mechanical, thermal, and optical properties are shown in Table 5.1

at both room temperature and cryogenic temperatures.

Table 5.1: The physical, thermal, and optical properties of Yb:YAG at room temperature and cryogenic temperatures.

	300K	77K	Source
Density	4.56 g/cm^3	4.56 g/cm^3	[101]
Young's Modulus	335 GPa	Not reported	[101]
Vickers Hardness	14.5 GPa	Not reported	[54]
Specific Heat	0.55 J / gK	0.1 J / gK	[3]
dn/dT	7.8	0.9	[3]
Thermal Conductivity	17	47	[34]
Linear Thermal Expansion Coefficient	6.14 ppm	1.95 ppm (100K)	[3]
Pump Wavelength	940nm	940nm	[34]
Emission Wavelength	1030nm	1030nm	[34]
Quantum Defect	9.3%	9.3%	N/A
Pump Absorption Cross Section	8×10^{-21} cm^2	1.8×10^{-20} cm^2	[60, 34]
Emission Cross Section	2.5×10^{-20} cm^2	11×10^{-20} cm^2	[31]
Upper State Lifetime	1.3ms	1ms	[31]
Fractional Thermal Population of Lower Laser Level	5.3%	0.001%	[59]
Laser System	Quasi-three level	Four level	N/A

As can be seen, Yb:YAG's thermal conductivity increases by a factor of roughly 3, and the thermal coefficient for the refractive index dn/dT decreases by a factor of 9 when cooled from room temperature to cryogenic temperatures. The coefficient of thermal expansion also decreases from 6.14ppm to just under 2ppm when cooled to 100K.

Optically, Yb:YAG's broad absorption band at 940nm allows for direct diode pumping, while its long upper states lifetime allows for CW pumping for pulsed laser operation, provided the laser repetition rate is on the order of 1kHz or higher. While it is pumped with 940nm, it lases at 1030nm, meaning it has a small quantum defect of only 9.3%. The low quantum defect allows for very high efficiency operation, with CW systems regularly reporting greater than 60% efficiency. Yb:YAG is an isotropic material, which eliminates the need to orient the crystal in a particular direction when mounting. The relatively large emission cross section of 11×10^{-20}cm^2 allows for very high gains.

Several of Yb:YAG's optical properties also improve when cooled to cryogenic temperatures. Most importantly, at room temperature Yb:YAG is a quasi-three level system. This is because the lower laser level is only separated by the terminal level by 612cm^{-1} [59]. This gives a 5.5% fractional thermal population of the lower laser level at room temperature. At 100K, the fractional population of the lower laser level is cut to 0.015%, and this is further cut to 0.001% at 77K. So, while it is a quasi-three level system at room temperature, it freezes out to a four level system at cryogenic temperatures [34].

When Yb:YAG is a three level system, thermal population of the lower laser level causes absorption at the 1030nm lasing wavelength. This can be seen in Fig. 5.1, which plots the transmission of a 6W, 1030nm input through a 20mm long, 2% at. doped Yb:YAG crystal as the crystal is cooled from room temperature to cryogenic temperatures. It should be noted here that the temperatures recorded are that of the crystal mount, as the crystal temperature could not be measured directly. It is seen that the power transmission begins to decline when the temperature exceeds ~110K, and that at room temperature approximately half of input power is absorbed.

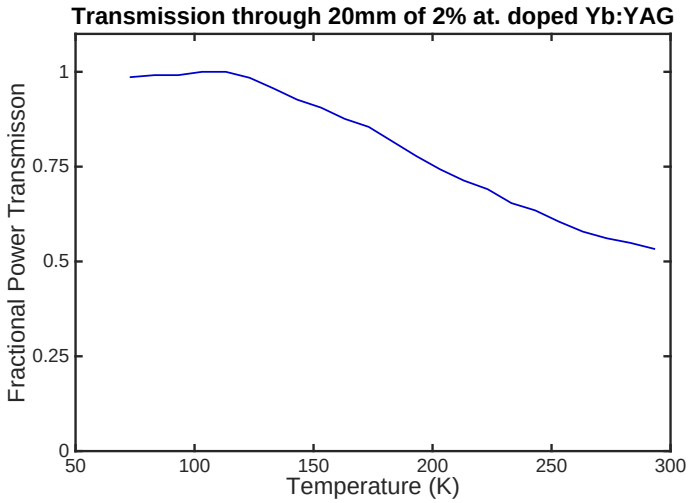

Figure 5.1: The fractional power transmission of a 20mm long, 2% at. doped Yb:YAG crystal as the crystal is cooled from room temperature to cryogenic temperatures. Notice that above ~110K the crystal begins to absorb 1030nm, and that at room temperature approximately half the input power is absorbed. The temperature values used were the measured crystal mount temperatures, as the crystal temperature could not be measured directly in this setup.

In addition to becoming a four level system, the stimulated emission cross section increases by a factor of 4 when cooled to cryogenic temperatures. The emission cross section peaks at 1029.7nm, and has approximately 0.25nm FWHM bandwidth.

Though cryogenic Yb:YAG is a very attractive material for use in high power laser systems, it does present some unique challenges. First, though the thermal conductivity increases by a factor of 3 when cooled to cryogenic temperatures, it still remains fairly low at 47 W/mK. For comparison, sapphire has a thermal conductivity around 500 W/mK at cryogenic temperatures. This limits the rate at which heat can be extracted from the crystal, and causes the crystal to warm up under heavy heat loads. Remembering that Yb:YAG transitions to a quasi-three level system as it warms, even a small temperature rise can cause significantly reduced power output. Thermal management ended up being one of the more challenging aspects in the development of this laser, and will be discussed in detail in Appendix B.

5.2 Cryogenic Yb:YAG Laser I: Stretcher and Yb:fiber Pre-Amplifier

As discussed in Chapter 4, the entire system is based upon a single Yb:fiber oscillator. Prior to compressing the oscillator output, approximately 50mW are picked off - this will be used for seeding the cryogenic Yb:YAG laser.

After this initial seed is picked off, it is first sent through a Chirped Volume Bragg Grating (CVBG) to stretch pulse in time. The CVBG is a reflecting optic approximately 5cm in length, and a diagram of it shown in Fig 5.2. The CVBG is made from photo-thermo-refractive (PTR) glass, which allows a holographic index pattern to be written into it via UV exposure. By writing a grating pattern into the glass, the CVBG acts as a wavelength dependent mirror, as it will reflect only wavelengths which interfere constructively upon reflection off each layer in the grating. Specifically, the reflected wavelength is given by

$$\lambda_r = 2n\Lambda, \tag{5.1}$$

where n is the refractive index of the PTR glass, and Λ is the periodicity of the grating pattern.

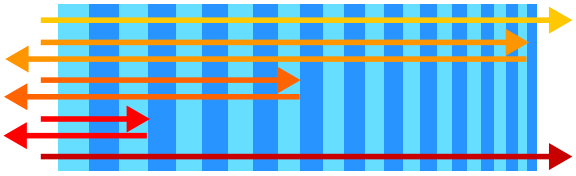

Figure 5.2: A diagram of the Chirped Volume Bragg Grating used to stretch the seed for the Yb:YAG system. UV exposure is used to imprint a holographic grating pattern in the material, which reflects certain wavelengths of light. By varying the grating periodicity through the length of the material, different wavelengths are reflected at different depths in the material. This leads to a spectrally dependent path length difference in the reflected wavelengths, stretching the pulse in time. Wavelengths outside the reflected bandwidth simply pass through the entire grating.

By changing the grating periodicity through the length of the grating, the CVBG reflects different wavelength at different depths. This spectrally dependent path length difference is used to stretch the pulse in time.

The CVBG reflects a square band centered at 1029.7, with approximately 0.6nm FWHM bandwidth. The operational band can be adjusted slightly by heating or cooling the CVBG, and shifts by approximately 0.1nm for every 10K change in temperature. This particular CVBG is cooled to 10°C with a PID controlled Peltier cooler. The CVBG has a dispersion of 540ps/nm, which, assuming a flat input spectrum across the 0.6nm operational bandwidth, translates into a stretched pulse length of approximately 285ps. A typical input spectrum and a modeled temporal output are shown in Fig. 5.3.

Following the CVBG stretcher, the seed pulse is sent through an Yb:fiber pre-amplifier, to boost the power of the seed pulse train prior to seeding the Yb:YAG system. This amplifier runs with approximately 1W of output power, which is about 17nJ at 60MHz. Due to the long pulse length (285ps) no self-phase modulation (SPM) spectral broadening is observed [66, 120]. If the seed pulse is not stretched, and so remains at ∼2ps, we see significant broadening due to SPM, and can easily broaden the spectrum from 0.6nm FWHM to 150nm tail-to-tail. This is undesirable, as Yb:YAG can only amplify a narrow ∼0.25nm FWHM spectral region around 1029.7nm, so any energy in the seed pulse outside of this spectral window goes to waste. After the Yb:fiber pre-

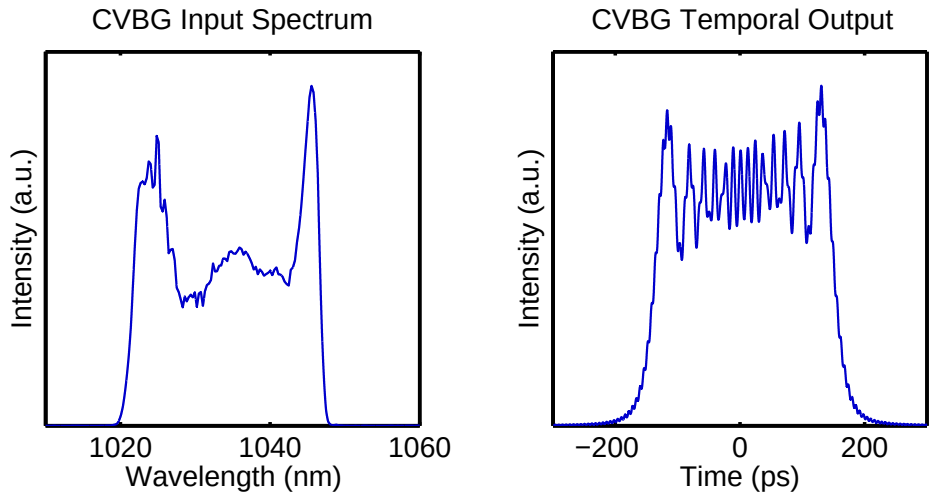

Figure 5.3: On the left, a typical input spectrum for the Chirped Volume Bragg Grating stretcher. Notice that the input spectrum is much broader (approximately 25nm FWHM) than the operational range of the CVBG (0.6nm FWHM). On the right is the predicted temporal output from this input spectrum, using a 0.6nm FWHM square reflection bandwidth and 540ps/nm dispersion. The 0.6nm FWHM spectrum is stretched to approximately 285ps FWHM. The ringing on the predicted temporal output is an FFT artifact introduced by the square notch function used to model the finite reflection bandwidth of the CVBG.

amplifier, the seed pulse train is sent into the first component of the Yb:YAG laser, a cryogenic Yb:YAG regenerative amplifier.

5.3 Cryogenic Yb:YAG Laser II: Yb:YAG Regenerative Amplifier

The first component of the Yb:YAG laser is a cryogenically cooled, Yb:YAG regenerative amplifier, a diagram of which is shown in Fig. 5.4. The regenerative amplifier begins with an 8mm diameter, 5mm long, 8% doped, flat cut Yb:YAG crystal, which is mounted in a copper crystal mount in a cryogenic vacuum chamber. The cooling is provided by a Cryomech PT-60 cryocooler, which provides 60W of cooling capacity at 80K. The Yb:YAG crystal is pumped by a 940nm fiber coupled diode laser, which can provide up to 50W of pump power. The pump fiber is a 400μm core diameter, 0.22NA multimode fiber. The fiber output is reimaged onto the Yb:YAG crystal with an approximate magnification of 1. As the fiber output is unpolarized, a flat crystal and normal incidence windows must be used, rather than a Brewster cut crystal and Brewster windows.

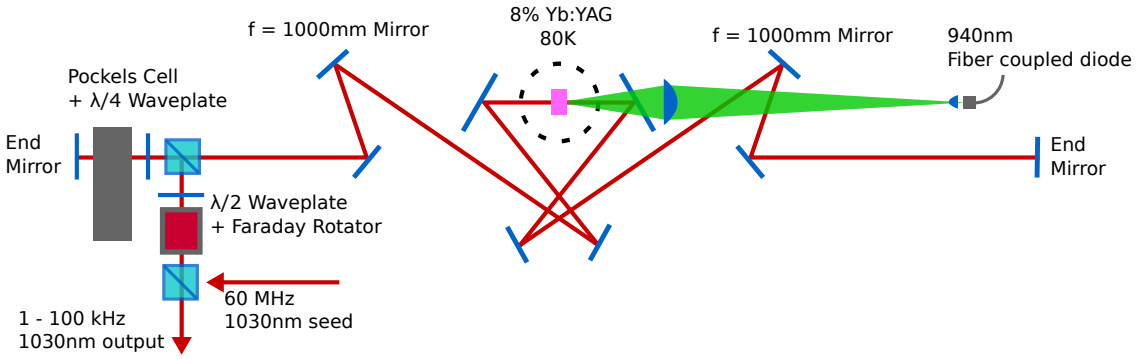

Figure 5.4: The layout of the cryogenic Yb:YAG regenerative amplifier. An 8% doped Yb:YAG crystal is cooled to 80K, and pumped by a 940nm fiber coupled diode. The regen is seeded by the 60MHz pulse train at 1030nm. A Pockels Cell / $\lambda/4$ waveplate combination is used to to trap seed pulses inside the amplifier cavity, amplifying them to high energy. The amplifier output runs between 1 and 100kHz, and is set by the repetition rate of the Pockels Cell.

Outside of the cryocell, the regenerative amplifier consists of a cavity formed between two end mirrors, and an electronically controlled switch, consisting of a polarizer, a quarter waveplate, and a Pockels Cell. This switch is used to control the cavity, is covered in detail next.

5.3.1 Regenerative Amplifier Switching

Regenerative amplifiers utilize an amplifier cavity that be alternately "opened" or "closed" via a combination of a polarizer, a quarter waveplate, and a quarter wave Pockels Cell, located near one of the end mirrors. To understand how these work, consider what happens to a small amount of unpolarized spontaneous emission from the Yb:YAG crystal[1] . The light travels from the crystal, through the cavity until it reaches the polarizer. The polarizer is set so that horizontally polarized light passes through it and remains in the cavity, while vertically polarized light is sent out of the cavity. Following the polarizer, it passes through the waveplate, the Pockels Cell, then reflects off the end mirror, and passes back through the Pockels Cell and waveplate before it encounters the polarizer again.

[1] In regenerative amplifiers based upon other materials, such as ti:sapphire, the crystal and the cryocell windows may be brewster cut, in which case the spontaneous emission that leaves the cryocell is already strongly polarized. As this regenerative amplifier uses a flat cut crystal and flat windows, the spontaneous emission is essentially unpolarized

When the Pockels Cell is turned off, it has no affect on the polarization. Therefore, the horizontally polarized light that passed through the polarizer experiences the net effect of two quarter waveplates. These combine to produce a 90° polarization rotation, so that the light is now vertically polarized. Now, when the light encounters the polarizer on the return trip, it is sent out of the cavity. This means when the Pockels Cell is turned off, any emission from the crystal is immediately coupled out of the cavity, preventing the buildup of a large photon flux. This is the "open" configuration.

On the other hand, when the Pockels Cell is turned on, it also provides a quarter wave polarization rotation. Now, the horizontally polarized light that passed through the polarizer experiences a 180° polarization rotation, and is still horizontally polarized when it returns to the polarizer. This means that when the Pockels Cell is turned on, no light can leave the cavity[2] . This is the "closed" configuration.

This switch provides the basis for the regenerative amplified dynamics, which are covered in the next section.

5.3.2 Regenerative Amplifier Dynamics

Regenerative amplifier dynamics are in general quite complicated, and a full quantitative treatment will not be given here. Instead, a brief qualitative description of the how the regenerative amplifier operates will be presented

In the first step, the cavity is completely open. As any spontaneous emission from the crystal is coupled out of the cavity, none of it returns to the crystal to stimulate further emission Since the crystal is being continuously pumped, the population inversion grows continuously. If allowed to grow indefinitely, it will eventually reach its steady state inversion density.

In the next step, the Pockels Cell is switched, and the cavity is closed. At this point, any horizontally polarized spontaneous emission emitted along the beam axis is completely trapped

[2] Technically, vertically polarized spontaneous emission still leaves the cavity. However, when compared to the amplified horizontally polarized emission this is negligible

in the cavity. All of this emission travels through the cavity and returns to the crystal, where it stimulates further emission. This continues, and rapidly depletes the inversion density, leading to a large photon density in the cavity.

In the final step, the Pockels Cell is switched back off, and the cavity opens. All of the light that had built up within the cavity now leaves the cavity, and the photon density in the cavity falls back to zero. The inversion density begins to build up again, and the whole process is repeated.

A regenerative amplifier can be run either seeded or unseeded. In the unseeded configuration, all of the cavity emission builds up from spontaneous emission. This leads to an output pulse length equal to the cavity round trip time. To seed the cavity, a small input pulse is sent into the cavity while the cavity is open. While the pulse is in the cavity, the cavity is closed, and the pulse is trapped within the cavity. The pulse circulates in the cavity, being amplified each time it passes through the crystal, until the gain is depleted. When seeded, the output pulse length depends predominantly on the seed pulse length.

5.3.3 Yb:YAG Regenerative Amplifier Performance

While the seed pulse train for the regenerative amplifier runs at 60MHz, the regenerative amplifier can run between 1 and 100kHz, set by the Pockels Cell repetition rate. When running at 1kHz, the regenerative amplifier runs with an average power of 0.5W, corresponding to a pulse energy of 500μJ, this is an energy gain of approximately 30,000. At 100kHz, the regenerative amplifier runs with slightly lower pulse energy - only 300μJ, but can deliver 30W of average power[3] . The regenerative amplifier output is highly stable, with a measured RMS shot-to-shot noise level of 1.58%.

As mentioned previously, cryogenic Yb:YAG's emission spectrum is centered on 1029.7nm, with approximately 0.25nm FWHM bandwidth. The measured output spectrum from the Yb:YAG regenerative amplifier is shown in Fig. 5.5. Here, remember that the seed pulses have been stretched

[3] The 30W output at 100kHz required approximately 70W of pump power, and for these experiments a 250W fiber coupled pump diode was used in place of the 50W unit. Both the 50W and the 250W pump diodes couple into the same pump fiber, so these units can be interchanged without affecting the pump geometry on the Yb:YAG crystal.

by the CVBG to 285ps, with a square bandwidth of 0.6nm FWHM. Since Yb:YAG cannot amplify the full seed bandwidth, the spectrum is narrowed in the Yb:YAG regenerative amplifier. Consequently, as stretching the pulses maps the spectrum to time, this leads to a reduction in the pulse length, and the 285ps seed pulses are shortened to 130ps. This can be seen in Fig. 5.5, showing the modeled output pulse duration.

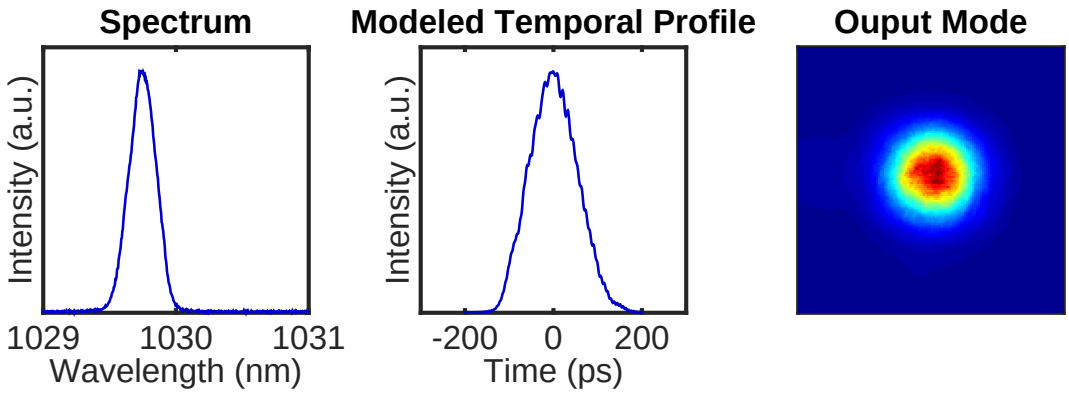

Figure 5.5: On the left, the output spectrum from the cryogenic Yb:YAG regenerative amplifier. The spectrum is centered at 1029.7nm with a FWHM bandwidth of 0.25nm. In the center, the modeled temporal output of the Yb:YAG regenerative amplifier when seeded with a 285ps, 0.6nm FWHM bandwidth seed pulse. The gain narrowing from the Yb:YAG amplifier cuts the pulse duration to 130ps. On the right, the output mode from the regenerative amplifier.

The efficiency of the regenerative amplifier can vary, and depends on the pump power, the seed power, the repetition rate, and the number of round trips the pulse makes within the amplifier cavity. Typically, the regenerative amplifier runs in a configuration with 14 round trips inside the cavity. Here, at 1kHz, 500μJ output energy requires 4W of pump power, an efficiency of 20%.

Following the Yb:YAG regenerative amplifier, the light is sent through an isolation arm consisting of a polarizer and a double passed quarter wave Pockels Cell. The isolation arm only allows light through when the Pockels Cell is switched on, and is used to remove any prepulses or

postpulses from the main 1kHz pulse train, as well as remove the residual unamplified pulses from the 60MHz pulse train that seeded the Yb:YAG regenerative amplifier. The output of the isolation arm is a clean 1kHz pulse train. The output mode from the regenerative amplifier following the isolation arm is shown in Fig. 5.5.

After the isolation arm, the beam is collimated with a telescope consisting of an f = 100mm lens and an f = 400mm lens to a $1/e^2$ mode diameter of 5mm, and sent into the remaining stages for further amplifications.

The next three stages will be covered sequentially in the next sections.

5.4 Cryogenic Yb:YAG Laser III: Yb:YAG Second Stage Multipass Amplifier

Following the Yb:YAG regenerative amplifier, the light is sent through a cryogenically cooled Yb:YAG multipass amplifier, a diagram of which is shown in Fig. 5.6. This amplifier also begins with an 8mm diameter, 5mm long, 8% at. doped Yb:YAG crystal. The crystal is mounted in a

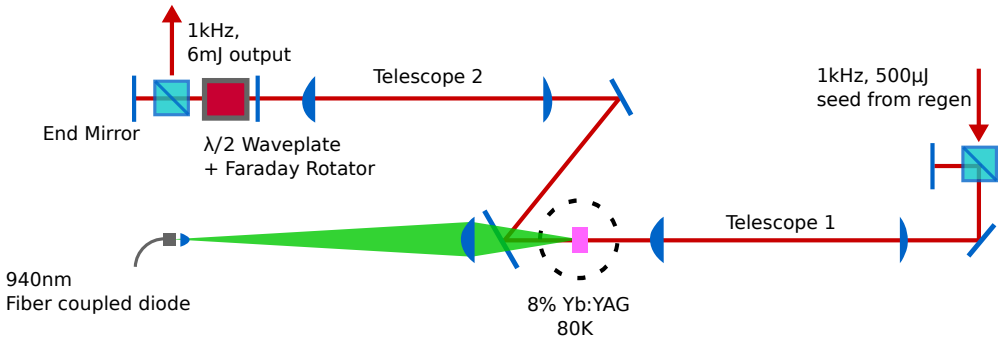

Figure 5.6: The layout of the cryogenic Yb:YAG second stage multiples amplifier. An 8% doped Yb:YAG crystal is cooled to 80K, and pumped with up to 41W by a 940nm fiber coupled diode. This stage is set up with a three pass configuration, with polarization used to couple the light in and out of the amplifier. Seeded by the 1kHz, 500μJ pulse train from the regenerative amplifier, this stage can run with over 8W of output power at 1kHz, though in practice it is typically run at 6W to mitigate the risk of damage.

copper mount, and fixed to a Cryomech PT-90 cryogenic cooler delivering 90W of cooling capacity at 80K. The crystal is CW pumped by a 940nm fiber coupled laser diode capable of delivering up to

250W average power. The pump fiber has a 400μm core diameter, and 0.22NA, and is imaged onto the crystal with an approximate magnification of 4, for a pump spot diameter just over 1.5mm.

Outside of the cryocell, the mulitpass amplifier consists of two relay imaging telescopes, labeled telescope 1 and telescope 2, two polarizers, one at the input of the amplifier and one at the output, a $\lambda/2$ waveplate and a Faraday rotator which provide a directionally dependent polarization rotation, and two end mirrors. The Faraday rotator and waveplate are set so that light traveling from the cryocell will experience no polarization rotation, while light traveling towards the cryocell experiences a 90° rotation.

This stage is seeded by the 1kHz, 500μJ pulse train from the Yb:YAG regenerative amplifier. The input beam has some slight divergence, with a $1/e^2$ mode diameter of approximately 5mm.

The light makes three passes in total through the amplifier. On the first pass, the horizontally polarized light enters through a polarizing beamsplitter and passes through telescope 1. Telescope 1 consists of an f = 300mm lens and an f = 100mm lens. As the seed has some slight divergence, these are set 460mm apart to down collimate the beam to a 1.5mm mode diameter. Here, the telescope again gives the beam some slight divergence to counteract the thermal lens in the second stage crystal.

The light passes through the crystal, is amplified for the first time, and then passes through telescope 2. Telescope 2 uses an f = 100mm lens and an f = 400mm lens, set 50mm apart to collimate the beam at 5mm. Following telescope 2 the light passes through the $\lambda/2$ waveplate and the Faraday rotator, but as it is traveling away from the cryocell this has no effect on the polarization. The light, still horizontally polarized, passes through the exit polarizer, and reflects off end mirror 1.

After reflecting off end mirror 1, the light passes first through the exit polarizer, and then back through the waveplate and Faraday rotator. This time, since the light is traveling towards the cryocell, these rotate the polarization from horizontal to vertical. Telescope 2 now down collimates the light to 1.5mm, and the light is amplified in the crystal for the second time. Telescope 1 now serves as an up collimating telescope, to bring the diameter back to 5mm. Finally, the now

vertically polarized light reflects off of the entrance polarizer, only to be immediately reflected back on itself by end mirror 2.

The third pass proceeds much the same as the first pass, being down collimated by telescope 1, amplified in the crystal for the final time, and up collimated by telescope 2. On this pass however, the light is vertically polarized when it encounters the $\lambda/2$ waveplate and Faraday rotator. As the light is traveling away from the cryocell, they have no effect on the polarization, and the light, still vertically polarized, is sent out of the second amplification stage by the exit polarizer.

When pumped with 41W, the second Yb:YAG stage can amplify the pulse energy from an input energy of 500μJ to 8.09mJ, an extraction efficiency of 19.7%. The measured slope efficiency is 26.9%. The second stage amplifier runs with a slightly more intensity noise than the Yb:YAG regenerative amplifier, with 2.46% RMS shot-to-shot noise. The second stage output mode is shown in Fig. 5.7.

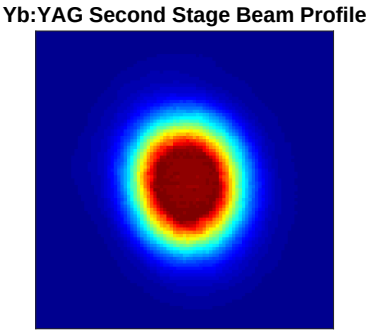

Figure 5.7: The output mode immediately after the second Yb:YAG stage, showing a flat top beam profile.

This high output energy risks damaging optics in the system, primarily the Yb:YAG crystal and cryocell windows which experience the highest peak intensity. Through practice, we have found that the risk of damage becomes significant when the peak intensity exceeds 5GW/cm^2. For this reason, this stage is often run at lower energies, at 5mJ output. This caps the peak intensity in the second stage amplifier at just over 4GW/cm^2, and allows the second stage to run without significant

risk of damage.

Following this amplification stage is the third cryogenic Yb:YAG amplification stage, which is another multipass amplifier.

5.5 Cryogenic Yb:YAG Laser IV: Yb:YAG Third Stage Multipass Amplifier

Much like the second Yb:YAG amplification stage, the third stage is another multipass amplifier, a diagram of which is shown in Fig. 5.8. Unlike the first two stages, the third stage uses a 12mm diameter, 7mm long, 8% at. doped Yb:YAG crystal. Similarly, it is also mounted in a copper crystal mount, and fixed to a Cryomech PT-90 cryocooler.

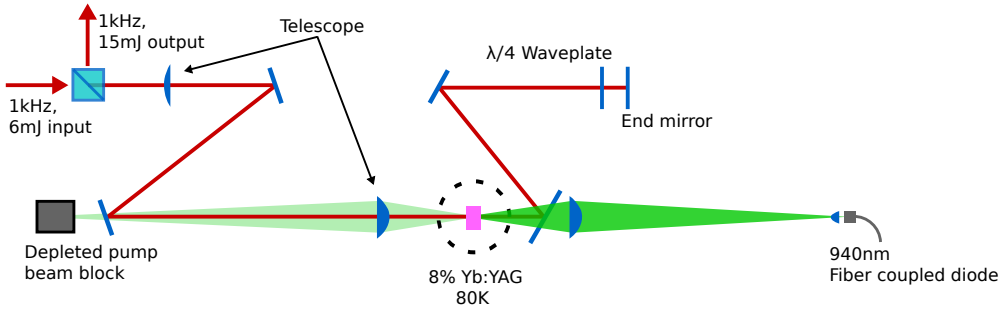

Figure 5.8: A diagram of the third cryogenic Yb:YAG stage. This double pass stage is seeded by 6mJ from the second Yb:YAG stage. When pumped with 65W, it runs with 15mJ of output energy at 1kHz.

The third Yb:YAG stage is CW pumped by a 940nm fiber coupled diode capable of providing 400W average power. The fiber has a 400μm core diameter and a 0.22NA, and the output is imaged on to the crystal with a spot diameter of 2.5mm.

Outside of the cryocell the third Yb:YAG stage consists of one relay imaging telescope, one $\lambda/4$ waveplate, a polarizing beamsplitter, and an end mirror. The third stage is set up so that the light makes two passes through the crystal.

The light from the second amplification stage enters through the polarizer, and is down collimated by the relay imaging telescope to a diameter of 2.5mm, again matched to the pump spot diameter. After passing through the crystal, the beam pass through the $\lambda/4$ waveplate, reflects off

the end mirror, and back through the $\lambda/4$ waveplate. The two passes through the waveplate combine to produce a 90° rotation in the beam, rotating the input horizontal polarization to vertical. The light agains passes through the Yb:YAG crystal, and the telescope up collimates the beam back to a 5mm diameter. The amplified light, now vertically polarized, is sent out of the third amplification stage by the polarizing beamsplitter.

When seeded by the full power (6mJ) second stage, the third stage can run with output energies approaching 20mJ. However, this puts the peak intensity on the crystal at approximately 6.3GW/cm^2, and much as in second stage the risk of damage becomes significant when the intensity is around 5GW/cm^2. Again, much as in the second stage, the optics most at risk for damage are the Yb:YAG crystal and the cryocell windows, which experience the highest intensity. To prevent damage, the third stage is run no higher than 15mJ, corresponding to 15W of average power. At 15W output, the stage is pumped by 65W for a relatively modest efficiency of 15.4%. This keeps the peak intensity at 4.7GW/cm^2, and allows the system to run without a significant risk of damage.

The noise level on the third stage is only slightly higher than the second stage, with 3.34% RMS shot-to-shot variation. The beam maintain a nice spatial mode through the third stage as well, with the far field beam profile shown in Fig. 5.9.

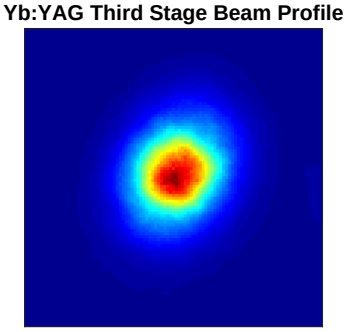

Figure 5.9: The far field beam profile from the third cryogenic Yb:YAG stage.

5.6 Cryogenic Yb:YAG Laser V: Yb:YAG Fourth Stage Multipass Amplifier

The fourth stage, while sharing some commonalities in architecture with the second and third stages, differs in many significant details. In fact, the fourth Yb:YAG stage proved to be one of the most challenging components in the project.

The fourth stage, as in previous three, is based upon on a single Yb:YAG crystal. In the fourth stage however, the crystal is 12mm in diameter, and either 15 or 20mm in length. The 15 mm long crystal are doped at 3.2% at. doping, and the 20mm long crystals are doped at 2.0% at, doping.

Either crystal is mounted in a single piece Copper-Tungsten (90%W /10% Cu) crystal mount. Copper-Tungsten is used because its coefficient of thermal expansion is well matched to that of Yb:YAG, and this significantly reduces the stress on the crystal as it is cooled to cryogenic temperatures. In the construction of the laser, the design and optimization of the fourth stage crystal mount ended up being a significant engineering challenge, that required analytical and numerical modeling, as well as significant testing before achieving an acceptable configuration. The details of this process are given in B. The crystal mount is fixed to the copper cold head 200W Cold-Helium Circulation Cryocooler (Cryomech CHCS-200), which provides 200W of cooling capacity at 80K.

As the fourth stage is designed to eventually provide average output powers of hundreds of watts, it requires a very high power pump. This is given by four fiber coupled 940nm diode lasers, each of which can deliver up to 250W of CW pump power. These are all connected to a beam combiner, which combines them with >90% efficiency into a single, 1mm diameter, 0.22NA fiber. The output from this fiber is imaged onto the Yb:YAG crystal with an approximate magnification of 5, for a 5mm diameter pump spot size.

A full diagram of the fourth stage is shown in Fig. 5.10. Outside of the cryocell, the fourth stage consists of two telescopes, two $\lambda/4$ waveplates, one $\lambda/2$ waveplate, a polarizing beamsplitter, and an isolator consisting of a $\lambda/2$ waveplate, a Faraday Rotator, and a polarizing beamplitter.

The 15mJ horizontally polarized input enters through the isolator, and passes through the

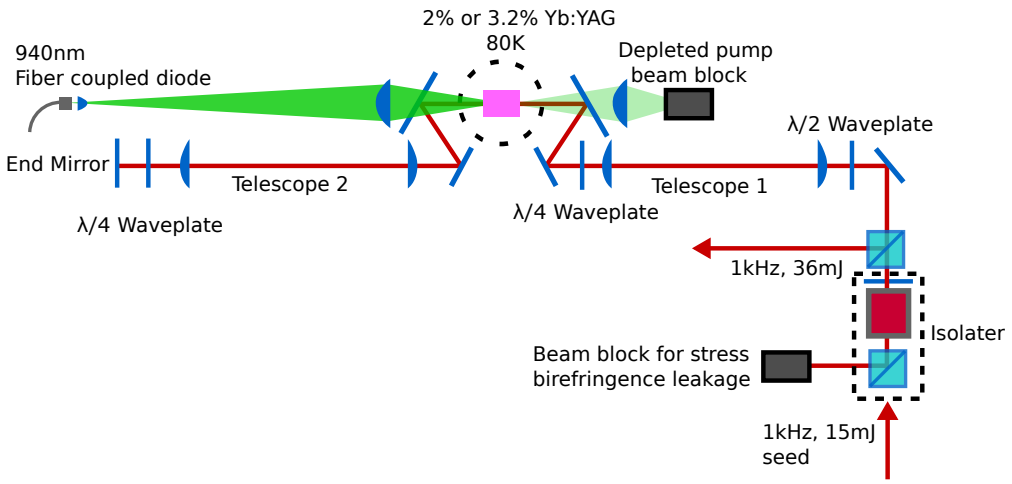

Figure 5.10: A diagram of the fourth cryogenic Yb:YAG stage. This double pass stage is seeded by 15mJ from the third stage, and runs with 36mJ output energy at 1kHz.

first telescope on the way to the cryocell. The first telescope consists of two f = 20cm lenses, with the lenses placed approximately 46cm apart from each other. This adds a small amount of divergence to the beam before it enters cryocell to help compensate for the thermal lens. Following the crystal, the beam passes through the second telescope, which consists of two f = 25cm lenses. Just as in the first telescope, these lenses are set to add a small amount of divergence to the beam on each pass. Following the telescope is a $\lambda/4$ waveplate, followed by the end mirror.

After the end mirror, the light retraces its path back through the amplification stage. However, as it is now vertically polarized due to the double pass through the $\lambda/4$ waveplate, it is reflected off of the output polarizer.

Also in the cavity are two other waveplates, a $\lambda/2$ waveplate before telescope 1, and a $\lambda/4$ waveplate after telescope 1. These are used to reduce depolarization loss, by slightly tuning the polarization of the light to better match the effective optic axis of the crystal [38, 24]. Without the two waveplates, the system has a double pass depolarization loss of 1.7%. With the two waveplates, the double pass depolarization loss is reduced to 0.6%, a reduction by more than a factor of two. The issue of depolarization will be covered in more detail in Appendix B.

At 1kHz, the fourth stage is pumped with 200W from of pump power from one of the diodes. When seeded with 15mJ from the third stage, the final amplified power is approximately 36mJ, for a gain of 2.4, and an extraction efficiency of 10.5%. The mode profile is shown in Fig. 5.11. The

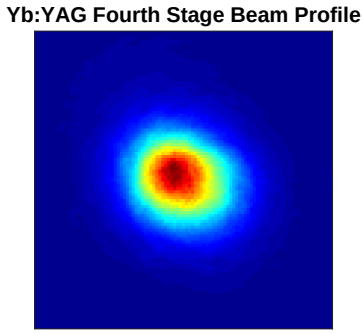

Figure 5.11: The far field beam profile from the fourth cryogenic Yb:YAG stage.

fourth stage is significantly less efficient than either of the previous stages, and the reason for this decreased efficiency will be discussed next.

5.6.1 Fourth Stage Efficiency

There are several reasons the fourth stage runs with less efficiency than the previous three stages. First, the high energies (up to 36mJ) mandate a large spot size on the crystal in order to keep the peak intensity in check. As the gain depends on the population inversion, larger mode sizes require higher pump powers to achieve the same gain, reducing the system efficiency.

In addition to this, because the required pump power scales unfavorably with the pump radius, the total heat deposited in the crystal also scales unfavorably with the pump spot size. This has the effect of raising the steady state operational temperature of the crystal. In the fourth stage, with the cryocooler running and no pump or signal input (i.e. no thermal load), the measured final crystal mount temperature is approximately 45K (as we can not measure the crystal temperature directly). When run at full power, the crystal mount warms to a steady state temperature around

93K[4] . From our finite element thermal modeling (which will be discussed further in Appendix B), we expect the crystal is between 10 and 35K warmer than the measured crystal mount temperature (the variation in crystal temperature is due to non-uniform pumping). So, we expect the crystal temperature varies from 103K to 128K. Recall from Fig. 5.1 showed that Yb:YAG begins to become absorptive above 110K, and we expect that this heating of the crystal will lead to some absorption at 1030nm, further degrading the amplifier performance.

Despite these limitations, the fourth stage runs with sufficient energies to pump all three of the OPCPA stages and achieve mJ level amplification at 3μm. There are plans to further boost the output energy to >100mJ, some of the details of which will covered in the section on future work in Chapter 6.

Following the fourth stage, the output is split via three $\lambda/2$ waveplate / polarizer combinations, each one controlling the pump energy for one of the OPCPA stages. The OPCPA stages will be covered in detail in the next chapter.

[4] The final operational temperatures can vary by a few Kelvin in day-to-day operation. This is thought to be due to the vacuum within the cryocell and the 200W cold helium cryocooler. The pressures in both of these tend to rise slowly over time, and they are periodically pumped out.

Chapter 6

Broadband Mid-Infrared OPCPA Stages

Now, having covered the front end of the system in Chapter 4, and the cryogenic Yb:YAG laser in Chapter 5, we reach the actual OPCPA stages, where the desired high energy, mid-infrared light is generated, which will be covered in this chapter. This chapter will proceed as follows. First, we will cover how the input signal pulse is stretched prior to amplification. Here we will discuss one of the complications in this OPCPA system - we stretch the signal, but compress the idler. As such, our discussion of stretching the signal will cover the relation between the signal and idler phases, and how the signal must be stretched in order for the idler to be later compressed.

Following this, we will next go through the three OPCPA stages individually, which together produce mJ level, 3μm pulses with bandwidths supporting <100fs compression. After this I will show our initial compressions results. Finally, the chapter will conclude with a discussion of the future directions for the OPCPA project.

6.1 Signal Stretching

Prior to any OPCPA amplification, the signal seed must be temporally stretched to match the pump pulse, which itself has been stretched to 130ps in a Chirped Volume Bragg Grating stretcher, covered in Chapter 5. This is typical of many high intensity laser systems, for reasons covered in Chapter 3. However, there is an unusual wrinkle that differentiates this system from the majority of previous systems. Usually, the pulse being stretched is the same pulse that will be compressed later on. When this is the case, stretching and recompressing is relatively straightforward. This

is because stretching and recompressing a pulse requires that the stretcher and compressor impart equal and opposite opposite phase functions to the pulse, up to an arbitrary linear phase function. That is,

$$\phi_{stretcher}(\omega) = -\phi_{compressor}(\omega) + \alpha\omega + \phi_0, \tag{6.1}$$

where $\phi_{stretcher}(\omega)$ is the spectral phase from the stretcher, $\phi_{compressor}(\omega)$ is the spectral phase from the compressor, α is an arbitrary constant, and ϕ_0 is an arbitrary phase offset. There are several well developed methods for ensuring the stretcher and compressor satisfy this condition. Commonly used methods include a Martinez positive dispersion grating stretcher with a Treacy negative dispersion grating compressor, or a grism stretcher with material compression [113, 72, 55, 45, 41, 42].

In this OPCPA system however, where we seed with the signal and compress the idler, no longer are we stretching and compressing the same pulse. This has only been done in a limited number of systems [7, 74, 75]. The challenge is to determine how we stretch the signal, such that we will be able to compress the idler later on. This challenge will be broken intro three questions, which will be discussed individually in the next three sections. As is common in laser design, the simplest approach to this problem is to begin with the end goal in mind, and works backwards from there. Therefore, we will first decide how we want to compress the idler, and this will let us know what phase the stretched idler must have prior to compression. The next step is to determine the relationship between the signal phase and the idler phase in an OPCPA process. This will give us the stretcher phase that must be applied to the signal so that the idler may be compressed. In the final step, we must determine how to actually apply the necessary phase to the signal.

6.1.0.1 Required Idler Phase for Compression

The first step to figuring out how to stretch the signal is to determine how the idler will be compressed. Here, we decided early on that the best way to compress the idler was with a standard negative dispersion grating compressor [113]. There are several reasons to use a grating compressor. First, high quality diffraction gratings are commercially available for the mid-infrared region of the spectrum. Second, a grating compressor uses all reflective optics, which is important

as many materials are absorptive in the mid-infrared. Most importantly though is that it is a well developed technology already commonly used in CPA amplifier systems [9].

A negative dispersion compressor with groove separation d, angle of incidence γ, and grating separation l will give a spectrally dependent phase

$$\phi_c(\omega) = \frac{2\omega l}{c}\sqrt{1 - \left(\frac{2\pi nc}{\omega d} - \sin(\gamma)\right)^2}. \tag{6.2}$$

Knowing the phase the compressor will impart on the idler, we know that for this phase to compress the pulse, the idler phase before compression must be given by $\phi_i(\omega_i) = -\phi_c(\omega_i)$,

$$\phi_i(\omega_i) = -\frac{2\omega_i l}{c}\sqrt{1 - \left(\frac{2\pi nc}{\omega_i d} - \sin(\gamma)\right)^2}. \tag{6.3}$$

Eq. 6.3 gives the phase the idler must have for it to be compressed by a grating compressor.

This may be further illuminated by looking at the Taylor expansions of the compressor and idler, in the standard form

$$\phi(\omega) = \phi_0 + \frac{d\phi}{d\omega}(\omega - \omega_0) + \frac{1}{2}\frac{d^2\phi}{d\omega^2}(\omega - \omega_0)^2 + \cdots. \tag{6.4}$$

Both the compressor phase and the idler phase (Eqs. 6.2 and 6.3) may be expanded out as such. As the full derivatives are not particularly illuminating, they will be omitted here, but may be easily calculated in Mathematica, or found in [9]. What is illuminating is to consider the signs of the various terms in the expansion. The results are shown in Table. 6.1.

Table 6.1: The signs of the terms in expansion of the phase of the idler compressor and stretched idler.

	GVD	TOD	$4OD$	$5OD$
Idler Compressor	-	+	-	+
Stretched Idler Phase	+	-	+	-

As expected, the terms in the expansion of the stretched idler and the idler compressor have opposite signs from each other. Additionally, we see that in the compressor the even terms in the

expansion (GVD, 4OD) are negative, whereas the odd terms are positive, and vice versa for the idler.

The exact magnitude of the terms depends on the idler compressor design, a ray trace model of which is shown in Fig. 6.1. The idler compressor utilizes two 300g/mm gratings, as highly

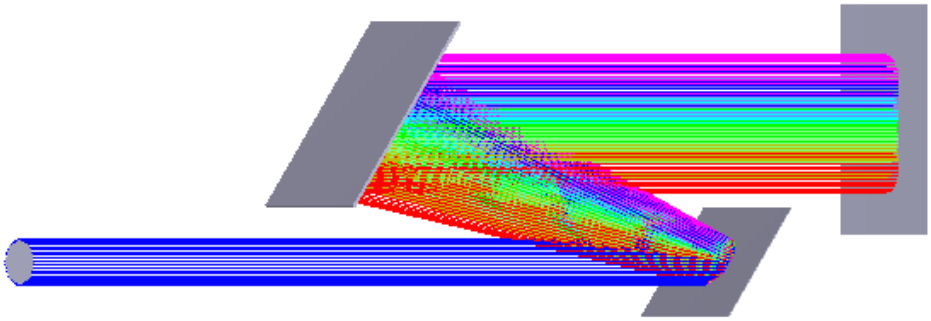

Figure 6.1: A ZEMAX ray trace model of the idler compressor, using two 300 g/mm gratings, a 37° angle of incidence, and a 10cm perpendicular grating separation. The ray trace model uses a 10mm beam diameter.

efficient gratings for 3-4μm are available with sufficient aperture sizes. The angle of incidence is 37°, and the designed grating separation is 10cm (measured normal to the gratings). Using a central wavelength of 3100nm, the dispersion terms for this configuration are given in Table 6.2. Also given

Table 6.2: The GVD, TOD, and 4OD for the idler compressor with 300g/mm gratings, a 37°angle of incidence, and 10cm grating separation. The central wavelength is 3100nm.

	GVD	TOD	$4OD$	TOD/GVD	$4OD/GVD$
Idler Compressor	-1.128 $\times 10^6 fs^2$	$7.473 \times 10^6 fs^3$	$-7.601 \times 10^7 fs^4$	-6.628fs	$67.410 fs^2$

in Table 6.2 are the ratios TOD/GVD and $4OD/GVD$. As all of the higher order dispersion terms depend linearly on the grating separation, they can all be adjusted in tandem by moving the grating further apart or closer together. The ratios, however, are fixed by the groove density and angle of incidence, so these remain the same for any grating separation.

This gives the full dispersion from the compressor, and the required stretch of the idler. The next section will present the exact relationship between the signal and idler phases, to see how the signal must be stretched to give the idler this phase.

6.1.0.2 Relationship Between Signal and Idler Phase

The simplest way to understand the phase relation between the signal and the idler in a parametric process is to return to the microscopic picture concerning individual photons. In this picture, each pump photon at ω_p is split into two photons: one signal photon at ω_s, and one idler photon at ω_i. Conservation of energy demands that

$$\hbar\omega_p = \hbar\omega_s + \hbar\omega_i, \tag{6.5}$$

while for the phase we can write

$$\phi_p(\omega_p) = \phi_s(\omega_s) + \phi_i(\omega_i). \tag{6.6}$$

where ϕ_p, ϕ_s, and ϕ_i are the pump, signal, and idler phases respectively [10]. In the case of a flat pump phase we can set $\phi_p = 0$, which simplifies the relation to

$$\phi_s(\omega_s) = -\phi_i(\omega_i). \tag{6.7}$$

So, at the individual photon level, corresponding signal and idler photons are generated with opposite phase.

The next step is to move from the microscopic photon picture to the macroscopic picture. Here, we need to find a relation between the signal and idler phase functions, $\phi_s(\omega_s)$ and $\phi_i(\omega_i)$. To do, we simply combine Eq. 6.5, and Eq. 6.7, giving the simple result

$$\phi_s(\omega_s) = -\phi_i(\omega_p - \omega_s). \tag{6.8}$$

This simple result allows one to determine the phase of the signal from the phase of the idler. A similar equation could also be written that allowed one to determine the idler phase from the signal phase.

With this results, we can now write a closed form solution for the required signal phase. By combining combining Eq. 6.8 and Eq. 6.3, we get an expression for the full signal phase

$$\phi_s(\omega_s) = \frac{2(\omega_p - \omega_s)l}{c}\sqrt{1 - \left(\frac{2\pi nc}{(\omega_p - \omega_s)d} - \sin(\gamma)\right)^2}. \tag{6.9}$$

At this point it is again illuminating to consider the Taylor expansions of the signal and idler phase, and see how the signs of individual terms compare. Both sides of Eq.6.8 are expanded in this manner. The left side gives

$$\phi_s(\omega_s) = \phi_0 + \frac{d\phi_s}{d\omega_s}(\omega_s - \omega_0) + \frac{1}{2}\frac{d^2\phi_s}{d\omega_s^2}(\omega_s - \omega_0)^2 + \cdots, \tag{6.10}$$

while the right side gives

$$-\phi_i(\omega_p - \omega_s) = -\phi_0 + (-1)^2\frac{d\phi_i}{d\omega_s}(\omega_s - \omega_0) + (-1)^3\frac{1}{2}\frac{d^2\phi_i}{d\omega_s^2}(\omega_s - \omega_0)^2 + \cdots. \tag{6.11}$$

At this point we equate terms of the same power. Comparing the Taylor expansions of the phase for the two sides, we see that the even terms have the opposite signs, while the odd terms have the same sign. This means that the signal and idler phases will have opposite signs for the even terms, like GVD and Fourth Order Dispersion, but they will have the same sign for odd terms, Third Order and Fifth Order Dispersion [1] .

This gives the full relation between the signal phase and the idler phase for an OPCPA process, assuming negligible pump phase. Odd terms - Third Order, Fifth Order, etc. - are the same between the signal and the idler, while even terms - GVD, Fourth Order, etc. - have opposite signs but equal magnitude. We now combine this result with Table 6.1, to get the required signs of the higher order phase terms in the signal pulse. These are given in Table 6.3. As can be seen, to compress the idler, each term in the expansion of the signal phase must be negative. In the next section, we will cover how to stretch the signal pulse in this manner.

[1] While this gives a full account of the signal / idler phase relation, it may not be physically obvious why only the even terms flip sign. An intuitive, physical explanation for this is given in Appendix A.

Table 6.3: The signs of the terms in expansion of the phase of the idler compressor, the stretched idler, and the stretched signal.

	GVD	TOD	$4OD$	$5OD$
Idler Compressor	-	+	-	+
Stretched Idler Phase	+	-	+	-
Stretched Signal Phase	-	-	-	-

6.1.0.3 Signal Chirped Fiber Bragg Grating Stretcher

The previous two sections demonstrated that for an OPCPA system seeded by the signal, compressing the idler with a negative dispersion compressor requires that the signal be stretched with negative dispersion in all of the higher order terms. This is first of two requirements for stretching the signal. The second requirement relates to the relative magnitudes of the higher order phase terms. It is not enough that the signal stretcher provide the correct signs for the higher order dispersion terms, it also must get the magnitudes correct. Specifically, to compress the idler, we require[2]

$$\left|\frac{TOD_{ic}}{GVD_{ic}}\right| = \left|\frac{TOD_{ss}}{GVD_{ss}}\right| \tag{6.12}$$

$$\left|\frac{4OD_{ic}}{GVD_{ic}}\right| = \left|\frac{4OD_{ss}}{GVD_{ss}}\right|. \tag{6.13}$$

Of all the ways a pulse may be stretched, we only found three ways that give all negative higher order terms: a prism pair, certain grism configurations, and certain tailored Bragg reflectors. Prism pairs are ruled out, as the magnitude of the dispersion needed is greater than they practically allow. Similarly, we decided against a grism stretcher, as the matching the ratios $|GVD/TOD|$ and $|GVD/4OD|$ with the idler compressor proved impractical.

A Chirped Fiber Bragg Grating (CFBG) is used to stretch the signal. This operates in much the same manner as the volume Bragg grating which is used to stretch the pump prior to the

[2] The reason it is sufficient that the ratios $|TOD/GVD|$ and $|4OD/GVD|$ match is because in the idler compressor, all of the terms in the expansion scale linearly with the grating separation. Therefore, as long as the relative ratios are correct, the actual magnitudes of each can be adjusted to compress the idler by changing the grating separation.

cryogenic Yb:YAG laser, except that the Bragg grating is written on a length of photorefractive fiber. While the fiber coupling is advantageous for ease of use and alignment, the main benefit is that commercially available fiber Bragg gratings can be manufactured to give exact control over all the dispersion terms up to fifth order. This allows a fiber Bragg grating to be written to exactly compensate for the designed compressor.

The CFBG has measured phase values given in Table 6.4. Fig. 6.2 shows a modeled temporal

Table 6.4: The GVD, TOD, and 4OD for the signal CFBG stretcher. Notice that all of the dispersion terms are negative, as require to compress the idler with a negative dispersion stretcher.

	GVD	TOD	$4OD$	TOD/GVD	$4OD/GVD$
Signal Stretcher	-1.137 $\times 10^6 fs^2$	-7.64$\times 10^6$ fs^3	-7.0$\times 10^7 fs^4$	6.7194fs	61.5655fs^2

output profile, and the modeled spectral output profile, based upon a typical measured spectrum from the HNLF. The sharp cutoff in both is due to modeling the sharp spectral cutoff from the circulator in the CFBG. The FWHM temporal length is 55ps, ideal for use with a 130ps pump pulse.

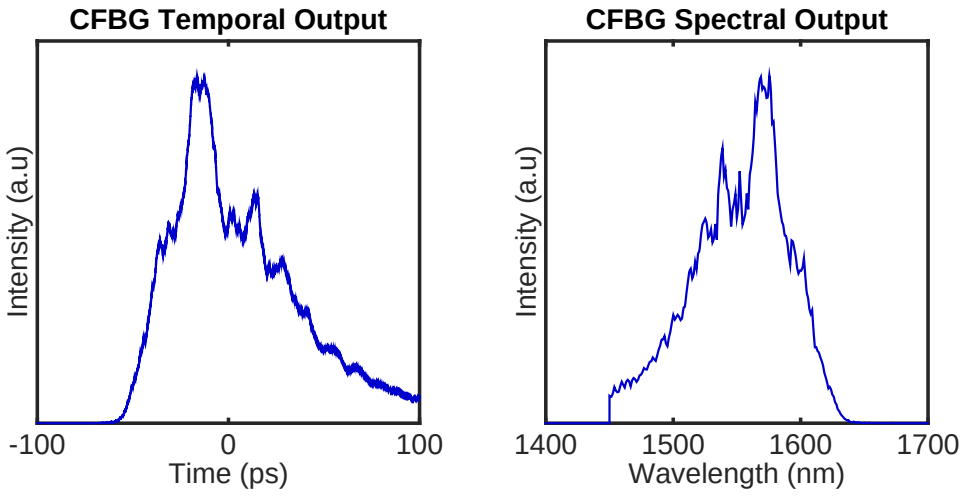

Figure 6.2: The modeled temporal output of the CFBG stretcher, and the modeled spectral output, based upon a measured output spectrum from the HNLF. The FWHM pulse length is 55ps.

6.1.1 First Stage

The first OPCPA stage is based on a 1mm thick, 10mm long MgO:PPLN crystal, with a periodicity of 29.98μm, operated at 100°C. The broadened and stretched signal from the OPO is combined with 700μJ of energy from the cryogenic Yb:YAG laser, then together they are focused onto the MgO:PPLN with an f = 35cm lens, to an approximate pump mode diameter of 410μm. This gives a peak pump intensity of 8×10^9 W/cm^2. A delay stage in the signal line is used to optimize the temporal overlap between the signal and the pump.

When seeded with 500pJ at 1.55μm from the broadened and stretched OPO output, the signal is amplified to over 30μJ output energy - a single pass gain of >45dB. The input and output signal spectra are shown in Fig. 6.3, demonstrating the amplification of $\sim$ 35nm FWHM bandwidth in the first amplification stage.

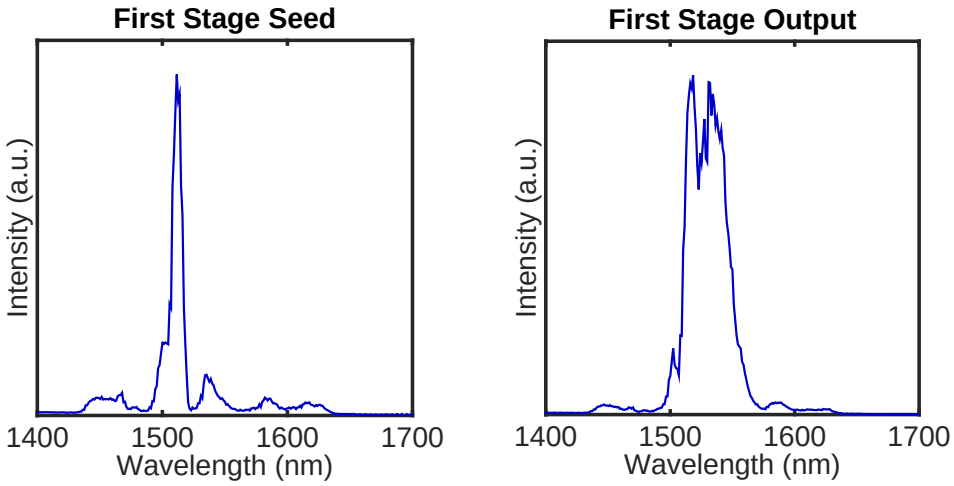

Figure 6.3: The first OPCPA stage seed and output spectrum. The seed is provided by the broadened and stretched OPO output. The amplified bandwidth is approximately 35nm FHWM, with a pulse energy of 30μJ.

Following the first amplification stage, the signal and pump are separated with a dichroic mirror, and the signal is sent to the second OPCPA stage.

6.1.2 Second Stage

The second amplification stage is based upon a 3mm thick, 3mm long MgO:PPLN crystal, this one with a periodicity of 30.1μm, also operated at 100°C. This stage is pumped with 3.5mJ pulse energy. Unlike the first amplification stage, in the second stage the pump and signal focusing is done separately, before the two beams are recombined. The pump is operated in a very nearly-collimated geometry, with only some slight convergence on the beam as it goes through the second stage crystal. The pump mode diameter is approximately 1.6mm on the crystal, which gives a peak intensity of 2.6×10^9W/cm^2. The signal spot from the first stage crystal is reimaged onto the second stage crystal with an f = 125mm lens with a magnification of 3:1.

The second OPCPA stage is seeded with the 30μJ output from the first amplification stage. This is amplified to 600μJ, a gain of 20. From the 3.5mJ pump input energy, this is a pump to signal conversion efficiency of 16%. While the idler from the second stage is not used, we know from conservation of energy that the second stage also has approximately 300μJ of idler output at 3μm, based upon an approximate photon energy ratio of $\hbar\omega_s/\hbar\omega_i \approx 2$. With the combined signal + idler energy of 900μJ, this is a total conversion efficiency of just under 25%.

The second stage input and output signal spectra are shown in Fig. 6.4. The second stage amplification bandwidth is around 28nm FWHM.

After the second OPCPA stage the signal is separated from the depleted pump, and the signal is collimated with an f = 125mm lens.

6.1.3 Third Stage

The third OPCPA stage utilizes two 3mm thick, 3mm long MgO:PPLN crystals. The first of these has a constant grating periodicity of 30.1μm, while the second is a chirped crystal, with periodicity ranging from 29.1 - 30.4μm. This is a linearly chirped grating, and the periodicity (in μm) is given by $\Lambda(z) = 29.1 + z(30.4 - 29.1)/3$, where z is the depth in the crystal, in mm. Note that the grating structure itself has a linear chirp, which is different from the situation used in the

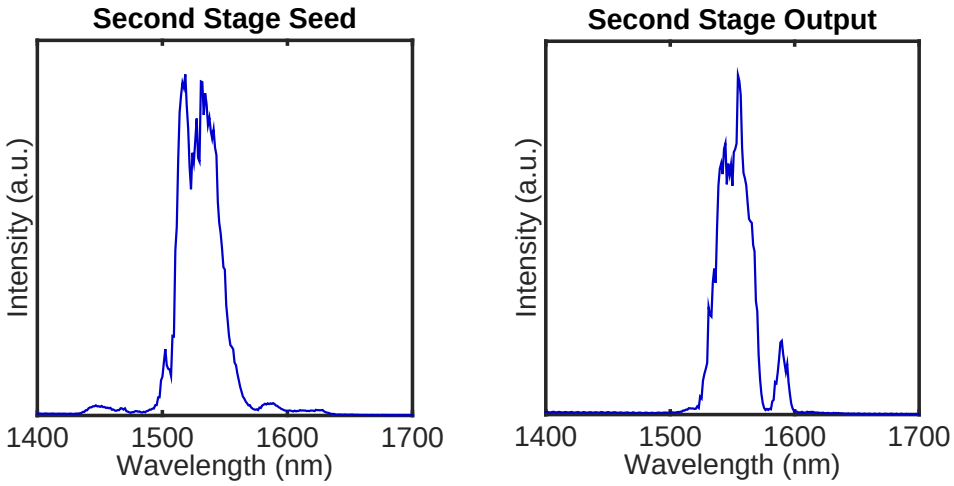

Figure 6.4: The second OPCPA stage seed and output spectrum. The second stage is seeded by the 30μJ, 35nm FWHM first stage output. This is amplified to 600μJ, with 28nmFWHM bandwidth.

analysis in Chapter 3.4.5.3, which assumed that the total phase mismatch, κ, had a linear chirp through the crystal.

In a geometry similar to the second stage, the pump and signal modes are focused separately. The pump beam is roughly collimated, with a mode diameter of approximately 2.5mm, giving a peak pump intensity of 3×10^9 W/cm^2. The third stage is seeded by the 600μJ second stage signal output. The third stage signal is amplified to 3.4mJ, a gain of 5.7, and a pump-to-signal conversion efficiency of 28%.

The third stage has an idler output of approximately 1.4mJ. This gives a total power output of 4.8mJ and a conversion efficiency of 42% in the third OPCPA stage. It should be noted though that in the present third stage configuration much of this energy is lost, as the MgO:PPLN crystals used for the third stage are not AR coated for the idler. In addition, after the MgO:PPLN an uncoated CaF_2 is lens is used to expand the the pump, signal and idler beams, which gives an additional 5% power loss. This lens is necessary to reduce the peak intensity on a sapphire dichroic beamsplitter. The beamsplitter is needed to separate the idler from the pump and signal, but has a low damage threshold.

The improperly coated crystals plus the uncoated lens results in approximately ≈ 40% of the idler energy being lost due to back reflections, so that the measured idler output is 900mW.

The signal bandwidth is greatly increased in the third amplification stage, as shown in Fig. 6.5. The input signal bandwidth from the second stage is approximately 28nm FWHM, while the output bandwidth is approximately 62nm FWHM. The idler bandwidth is shown in Fig. 6.6, and

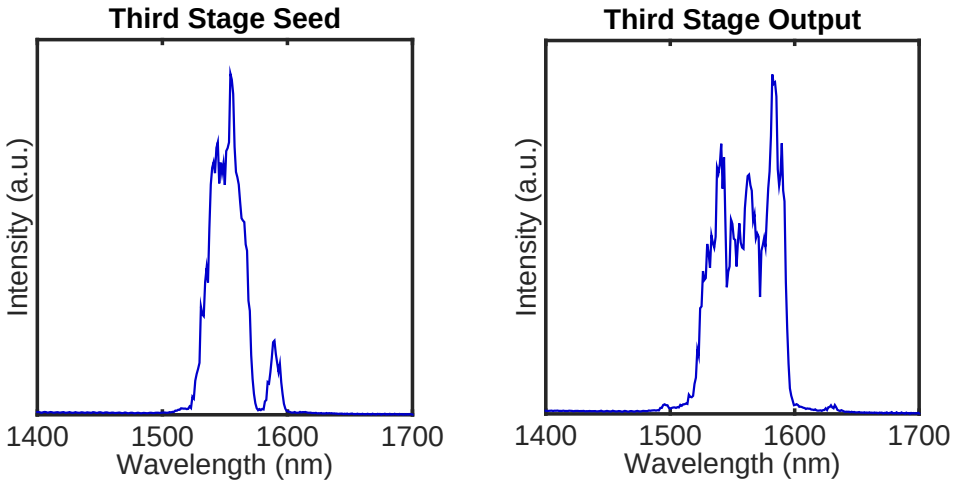

Figure 6.5: The third OPCPA stage seed and output spectrum. The third stage is seeded by the 600μJ, 28nm FWHM second stage output. This is amplified to 3.4mJ, and the FWHM bandwidth is extended to approximately 62nm.

is approximately 275 FWHM, centered at 3.05μm. This spectrum has a transform limit of 57fs, demonstrating that the third stage idler is capable of running with mJ level out with more than sufficient bandwidth to compress to sub-100fs pulse lengths.

In addition to having good amplification energy and broad signal and idler amplification bandwidths, the third OPCPA stage also has very nice spatial mode properties. Mode images from the third stages signal and idler are shown in Fig. 6.7, demonstrating the excellent spatial mode of both the signal and the idler.

Following the third OPCPA stage, the idler output is sent into the compressor, which will be covered next.

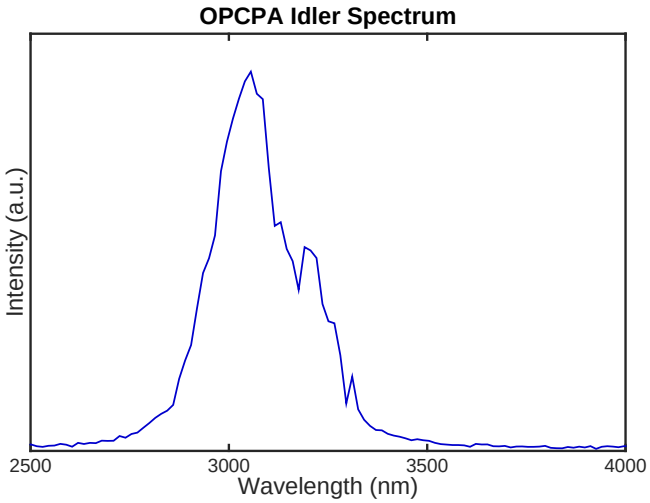

Figure 6.6: The third stage OPCPA idler spectrum, with a FWHM bandwidth of 275 centered at 3.05μm. The transform limit is 57fs, showing more than sufficient bandwidth for <100fs pulse compression.

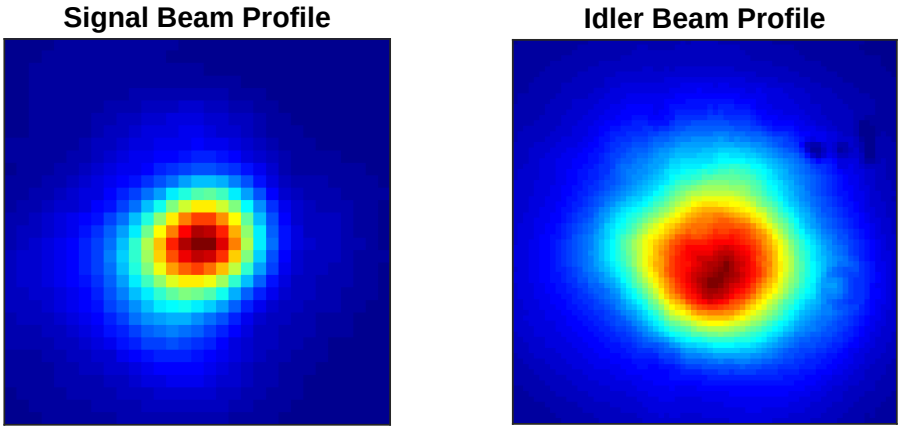

Figure 6.7: The spatial beam profiles for the OPCPA signal, shown on the left, and idler, shown on the right, demonstrating excellent spatial mode quality.

6.2 Compression

Following the third OPCPA stage the idler is compressed with a negative dispersion compressor, as discussed previously in Chapter 6.1.0.1. Thus far, the idler pulses have been compressed to

<110fs. A FROG trace and the retrieved temporal profile are shown in Fig. 6.8, showing a 105fs FWHM pulse length.

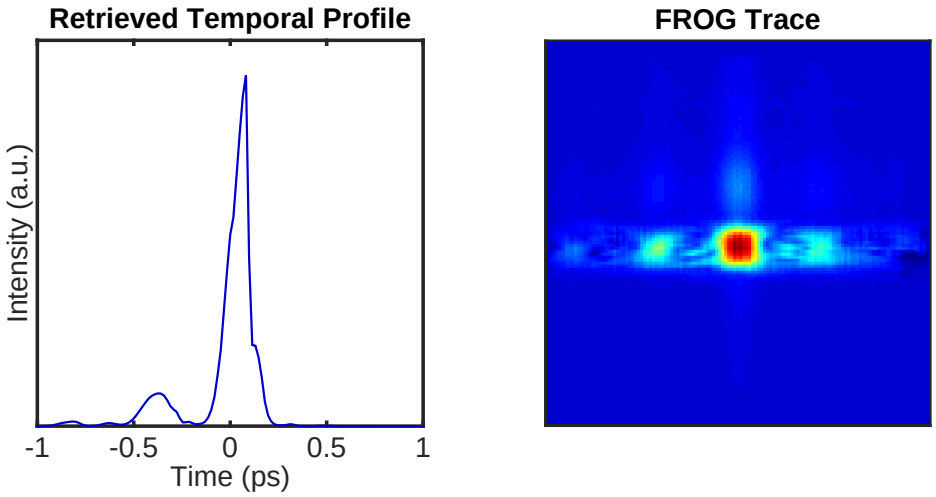

Figure 6.8: On the left, the retrieved temporal profile for the compressed idler, with a 105fs FWHM pulse length. On the right, the measured FROG trace for the compressed idler.

A few factors currently limit the ability to compress the idler to <100fs. The first is phase from the the pump mapping onto the idler. Our calculations show that with our pump and signal current stretching, we should expect the pump to add approximately $+4.22\times10^5 \text{fs}^2$ of GVD to the idler. This can be compensated by moving the compressor grating further apart, however, this results in a small amount of uncompensated third order and fourth order phase. All in all, the uncompensated phase from the pump slightly change the ratios $|GVD/TOD|$ and $|GVD/FOD|$ of the idler, so that they no longer exactly match our compressor. The manner in which this will be addressed will be covered in section 6.3.

The compressor at present runs with only 41% efficiency for a total output power of 370mW, or a pulse energy of 0.37mJ. This is primarily due to the gratings, which at 80% efficiency per reflection will give a 41% total efficiency with four passes. More efficient gratings, with over 95% efficiency per pass, are expected to boost the total compressor efficiency to >80%. With the 0.9mJ current available energy for the compressor, this would give 0.73mJ of compressed pulse energy. If

the full 1.4mJ were available for the compressor, this would give 1.15mJ of compressed output.

Despite these limitations, this system demonstrates the proof of concept for a kHz repetition rate, mJ pulse energy, mid-infrared OPCPA system. The system can run with 1.4mJ at 3μm, with sufficient bandwidth to allow for $<$ 100fs pulse compressor. The final output energy is cut due to improperly coated crystal and a low efficiency compressor, but has been compressed to $<$ 110fs with 0.35mJ of output energy.

6.3 Future Work

Thus far we have shown a proof of principle OPCPA system that can deliver mJ level mid-infrared pulses. These pulses have thus far been compressed to 110fs, with bandwidths supporting sub-100fs compression. However, there is much more to be done to extend this to even higher pulse energies to make this system a practical tabletop keV x-ray source.

At present, there are numerous losses in the system that reduce the overall efficiency - namely improperly coated crystals and low efficiency compressor gratings. Together, these two factors reduce the total output from 1.4mJ to 0.37mJ. Properly coated crystals and improved efficiency gratings are expected to increase the output energy to >1mJ.

Beyond these simple changes, some larger redesigns and modification are planned to optimize the system with the knowledge gained in this proof of concept work. These planned modifications and extensions will be described next.

6.3.1 OPA Front End

The first upgrade to the OPCPA system, which has already been put in place, replaces the Yb:fiber oscillator and OPO front end. The new front end also utilizes an Yb:fiber oscillator, but this oscillator is commercially packaged with a Large Mode Area Yb:fiber amplifier (Y-Fi, KMLabs). Currently, running at 10MHz, this units delivers 500nJ, <150fs pulses with a broadband spectrum around 1040nm, shown in Fig. 6.9. Plans are in the works to further upgrade this unit to allow for pulse energies exceeding 1μJ, although perhaps at a reduced reputation rate between 1 and 5MHz.

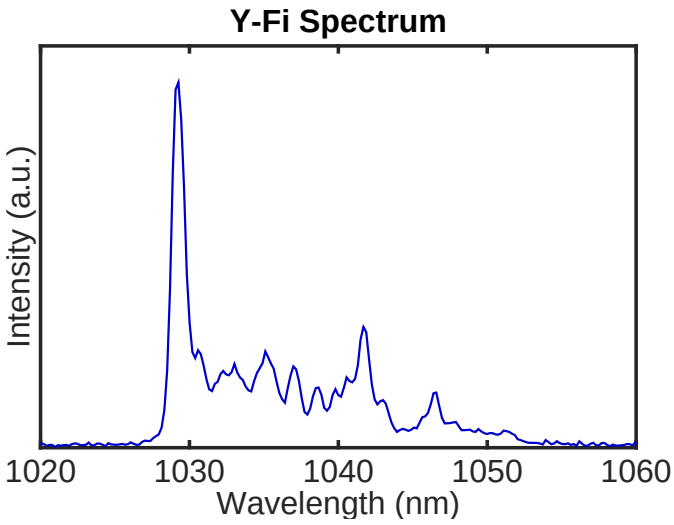

Figure 6.9: The output spectrum from the Y-Fi in the new front end.

This new Y-Fi based front end seeds the cryogenic Yb:YAG system in exactly the same manner as the Yb:fiber oscillator. Approximately 50mW of the Y-Fi output is picked off, stretched in the CVBG, and then amplified in an Yb:fiber pre-amplifier. Presently, the fiber pre-amplifier output is up to 1W - a pulse energy of 100nJ at 10MHz. This has been used to seed the Yb:YAG regenerative amplifier and the entire Yb:YAG system. Seeded with the new front end the Yb:YAG laser runs with identical performance parameters as with the previous front end.

To generate the signal seed for the OPCPA stages, the new front end design uses an OPA in place of the OPO. The OPA is pumped by the bulk of the Y-Fi output, approximately 4.95W. The OPA has two main components: a white light generation stage and an amplification stage. The white light generation stage is used to generate the initial supercontinuum seed, which is then amplified in the amplification stage.

Two configurations have been tested for the OPA. In the first configuration, the entire input is focused into a 3mm thick YAG crystal for white light generation. A lens after the YAG crystal collimates both the white light and the depleted pump, which are then split with a dichroic mirror. The depleted white light pump is sent through a delay arm, then recombined with the white light in an MgO:PPLN crystal, where the pump amplifies the white light. Numerous dif-

ferent MgO:PPLN crystals have been used in the system, including a 10mm single grating crystal (periodicity 29.98μm), a 2mm fan out crystal (periodicity ranging from 21-35μm), or combinations of 3mm flat (29.98μm periodicity) and chirped crystals (29.1-30.4μm) periodicity.

It has been found that the OPA signal power output does not depend significantly on the specific crystal used, but the signal bandwidth does. Fan out and chirped crystal allow for broadband amplification, as can be seen in example spectra shown for the signal and idler in Fig. 6.10, which have 70nm and 780nm FWHM bandwidths, respectively. These were both taken in a configuration using YAG for the white light generation, and a 2mm fan out MgO:PPLN crystal. It should be noted that these spectra were not obtained simultaneously, and that typically the OPA is tuned to optimize one or the other. However, it should also be noted that in practice this limitation is not important, as only one of the signal or idler is used at any given time.

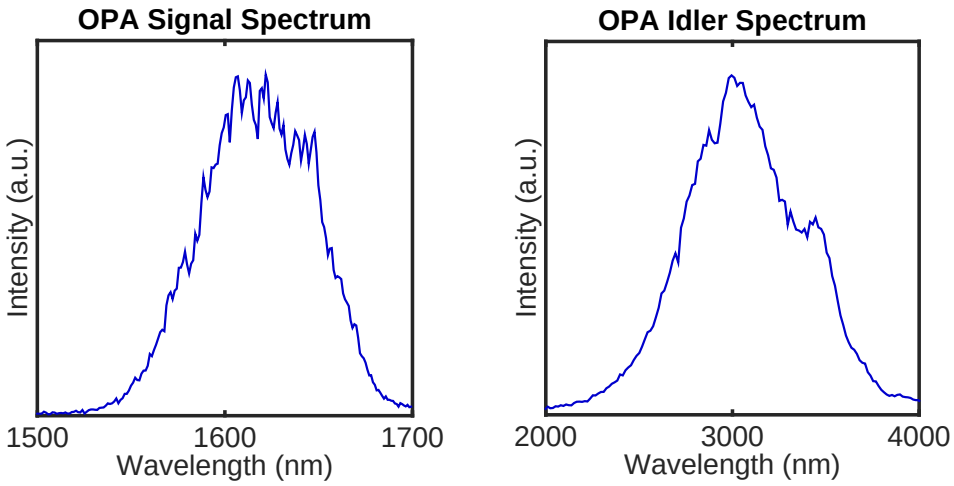

Figure 6.10: Examples of the broadband output that can be obtained from the Y-Fi pumped OPA. These spectra were taken with YAG used for white light generation, and a 2mm fan out MgO:PPLN crystal for amplification. The signal spectrum on the left has a 70nm FWHM, and the idler on the right has 780nm FWHM.

While this configuration can result in broadband, high power output, we have found that it suffers from amplitude noise on the output signal. This is due to the fact that the depleted white light pump is also used to pump the amplification stage. As the white light generation stage

imparts significant phase on the pump due to self phase modulation, any amplitude noise on the input is mapped to temporal noise on the output. This causes an unstable temporal structure of the depleted pump, which in turn leads to further noise when it is used to pump the amplification stage.

The second OPA configuration has been designed to alleviate this issue. Here, the YAG crystal has been replaced with a photonic crystal fiber (PCF) for the white light generation stage. As the PCF needs significantly less energy for the white light generation, different pulses can be used for pumping the white light stage and amplification stage, reducing the total noise. This OPA configuration isn't yet complete, but it is expected to provide similar output power and bandwidths as the previous design, with increased stability.

Replacing the OPO with an OPA will improve the OPCPA system in a number of ways. First, the OPO is the source of much of the instability in the OPCPA system. As the OPO is very sensitive to alignment and the chirp of the pump pulse, any temporal or spectral drift in the Yb:fiber oscillator is amplified in the OPO. Further, the OPA will provide sufficient idler output to allow the idler to be stretched and used for seeding of the OPCPA system. This will be discussed in the next section.

6.3.2 Idler Seeded OPCPA

Replacing the OPO with an OPA will allow for direct idler seeding of the OPCPA system. This will provide a number of benefits. First, in the present configuration, approximately 300μJ of idler pulse energy are simply lost after the second stage. By amplifying the idler directly, we will be able to keep this energy and subsequently boost expected idler output by the same amount. Furthermore, another advantage of the direct idler seeding is that the idler from the OPA is intrinsically Carrier Envelope Phase (CEP) stable [10]. In turn this means that the OPCPA output itself is intrinsically CEP stable.

The biggest benefit however is that in this configuration, the idler will be both stretched and compressed. This new arrangement means the idler can be stretched with a standard positive

dispersion stretcher, and compressed with a Treacy negative dispersion compressor. Use of a standard positive dispersion stretcher has three benefits. First, the amount the OPCPA seed is stretched is now tunable, whereas before the stretch was fixed by the Chirped Fiber Bragg Grating. This means the idler stretch can now be tuned to optimize OPCPA output for bandwidth and power. Second, an adjustable idler stretcher gives another knob that can be used to compress the idler output, which make compression easier. Finally, because we are seeding with the idler, it is the signal pulse which will take on undesirable phase from the pump pulse.

Seeding the OPCPA with the idler is expected to increase the output energy, and enable easier optimization of the OPCPA, and easier compression of the idler.

6.3.3 Increased OPCPA Power Output

Eventually, the goal for the OPCPA laser is to be able to deliver 10mJ compressed pulses at 3μm. With a fairly conservative assumption of 33% compression losses, this translates into 15mJ pre-compression pulse energy, roughly an 11$\times$ increase from the current idler energy. In this section, I will discuss the future plans to allow for these higher pulse energies. This will begin with a discussion of the current factors that limit the OPCPA pulse energy.

Currently, the primary factor that limits the output pulse energy is the aperture size of the third OPCPA crystals. As the crystals are only 3mm thick, the necessarily limits the beam diameter to less than 3mm. The small beam diameter limits the amount of pump energy that can be used before the crystal damages. This represents the most immediate challenge to scaling to higher pulse energies.

There are several ways to address this issue. First, cylindrical optics could be used to create an elliptical beam on the last stage MgO:PPLN crystals. As the crystal are 10mm wide, this could enlarge the mode area and lower the intensity on the crystal. However, this also requires reshaping the signal / idler input, and then reshaping the idler back to a circular beam. As the idler mode quality is very important for coupling into the waveguide for high harmonic generation, and cylindrical optics require very precise alignment, we've decided that any solution which compromises

the idler mode quality is not logical given the goals of the project.

Another option is to increase the crystal size. At the moment, 3mm is the largest aperture commercially available, but 4mm crystals are expected to be available in the near future. Keeping the same peak intensity, this will increase the maximum pump energy by a factor of ~ 1.75 - only a small fraction of the desired $11\times$ increase. Even larger aperture crystals may become available, but this is not expected to happen in the near future. By itself, simply using larger MgO:PPLN is not a sufficient solution.

For this reason, other nonlinear materials which can be made in larger sizes are being considered, chiefly KTA. KTA can be easily manufactured with apertures greater than 1cm^2, and lengths up to several cm long. Its damage threshold is also higher than MgO:PPLN's. The downside to KTA is that it has a modest d_{eff} at 2 pm/V, necessitating the use of long crystals, which correspondingly decreases the phase matching bandwidth [107]. However, some preliminary work has shown that it can be used to increase the output power. A 10mm long, Brewster cut KTA crystal was tried in the third OPCPA stage. In this configuration, the pump and idler are p-polarized, while the signal is s-polarized, so the signal experiences a $\sim$27% loss on the input and exit faces. With approximately 25mJ of pump energy, we measured 3.7mJ in the signal and 2.5mJ in the idler from third stage. Taking into account the 27% loss the signal experiences on the exit face, this means there was roughly 5.07mJ of signal in the crystal - which matched the signal energy you would expect from conservation of energy, given we were generating 2.5mJ in the idler. All told, this system produced 7.5mJ of output, of which 6.2mJ could be used taking losses into account. However, this was done at the cost of substantially lower bandwidth - only 7nm FHWM in the signal. Further, it should be pointed out that for this measurement the pump was stretched to 260ps by double passing the CVBG on the seed for the Yb:YAG laser.

This brings us to the final way to increase the third stage pump energy - further stretching of the pump pulse. In our operating regime, peak intensity is the main driver of optically induced damage, so by stretching the pulse further, we can use a higher energy with the same spot size. As mentioned, this has already been tried in limited experiments by double passing the CVBG pump

stretcher to stretch the pump to 260ps FWHM. Further stretching of the pump, potentially up to 1ns, could be achieved with a CFBG.

This final method holds a lot of promise, since it also helps to address a future limitation in further scaling the OPCPA idler energy: the pump energy available from the Yb:YAG laser. Currently, the maximum stable output is 36mJ at 1kHz, of which presently only ∼14mJ are used to pump the OPCPA. With the current OPCPA conversion efficiency, 36mJ of pump could conceivably deliver around 3.5mJ at the 3μm. Further scaling the output will require more pump energy, and reaching 15mJ in the idler will require 125mJ, again using our current conversion efficiency.

Scaling the Yb:YAG system to these output energies, while maintain kHz repetition rates will be a major challenge as this project continues. As discussed in Chapter 5.6.1, the large mode size in the fourth stage presently limits the efficiency.

Further stretching the pump pulse will allow a smaller mode size for any given power output, but this alone is not expected to be enough to scale to >100mJ pulse energies. This will instead require better mounting and thermal management of the fourth stage crystal. The prospects for this are discussed in more detail in Appendix B, which covers the work that has been done to date, and how this can be extended.

With the expected increased power output, the laser will also soon enable another exciting possibility: using this OPCPA laser as pump for another OPA that runs at even longer wavelengths, potentially up to 10μm. Recently we have begun numerical work to study how this can be done, and this will be discussed next.

6.3.4 Prospects for Deep Mid-Infrared and Far-Infrared Pulse Generation

One of the most exciting potential uses for this system is to use either the signal or the idler output as the pump for another OPA / OPCPA to generate high energy, fs level pulses further into the mid or far-infrared, with wavelengths up to 10μm. Remembering that the high harmonic generation phase matching cutoff energy scales favorably with the driving wavelength, by moving further into the infrared we have the potential to explore coherent x-ray generation

at multi-keV photon energies. There are several nonlinear materials that could be used of this purpose, most importantly Zinc Germanium Phosphide (ZGP), Cadmium Silicon Phosphide (CSP), and Orientation-Patterned Gallium Arsenide (OP:GaAs) (orientation patterning is a quasi-phase matching method functionally identical to periodic poling).

These three materials have all recently found use in mid-infrared nonlinear systems. ZGP has a transparency range from $\sim 2\mu$m through $\sim 11\mu$m, and a large effective nonlinear coefficient d_{eff} of ~ 75 pm/V [122]. It has been used in an OPO tunable from 3.8-12μm, as well as a high energy 2μm pumped OPA with 30mJ of output at 3.4μm with 40nm of bandwidth [114, 27]. Even higher energy OPA use has been reported, with over 100mJ output for ns pulses tunable from 3.5-5μm, although the bandwidths were not reported [53].

CSP is relatively new material, with a transparency range from 660nm - 6.5μm, and an effective nonlinear coefficient d_{eff} of 85 pm/V [107, 57, 100]. Importantly, CSP allows for 1μm pumping, so well developed laser technologies like Yb:YAG and Nd:YAG can be used. As it is a relatively new material (first fabricated in large apertures sizes in 2008) there are few systems using it to date, though it has been used in a 1μm pumped OPO, with 6.2μm output with a transform limit pulse duration of 160fs [123].

Finally, OP:GaAs is very attractive new quasi-phase matched material. GaAs is transmissive from $\sim 1\mu$m - 16μm, and also has a very high effective nonlinear coefficient d_{eff} at 95 pm/V [106]. While partially transmissive at 1μm, absorption still limits the potential for 1μm pumping, and no 1μm pumped systems have been reported. Rather, OP:GaAs is usually pumped at either 2μm, by sources such as Tm:fiber, or Tm,Ho:YLF, or at 3μm by a PPLN based OPO. For example, Tm:fiber pumping has been used in OPO systems with 7-13μm tunable output, as well as systems with 2.6-6.1μm simultaneous output [90, 63]. Systems pumped at 3μm have demonstrated Optical Parametric Generation (OPG) yielding a 4.5-10.7μm mid-infrared supercontinuum, and OPO systems with a tunable 4-14.2μm output [62, 115, 116].

To understand how these materials may best be employed for mJ level, femtosecond pulse generation in the 5-10μm range, we recently conducted some numerical studies to simulate and

evaluate the performance of proposed OPA and OPCPA systems.

6.3.4.1 Infrared Pumped OPA / OPCPA: Numerical Model

In this model, we evaluate the amplification of a seed pulse centered 7.2μm, with 400nm FWHM bandwidth, and a Gaussian spectral profile. The transform limit of this spectrum is roughly $\sim$ 185fs FWHM, which at 7.2μm corresponds to an $\sim$ 8 cycle pulse. The pump pulse was either at 1.5μm, (corresponding to our OPCPA signal output) or at 3.1μm (corresponding to our OPCPA idler output). For each pump wavelength, three different FWHM pump bandwidths were evaluated: 1nm, 30nm, and 100nm for 1.55μm, and 10nm, 100nm, and 300nm for 3.1μm. Like the seed, the pump was also assumed to have a Gaussian spectral profile.

In these simulations, we studied two types of systems. The first type is called the "high gain design". This design is similar to our first OPCPA stage - a modest pump pulse amplifies a very low energy seed pulse, and achieves very high gains, potentially greater than 10,000. In our simulations we assume a pump pulse energy of 500μJ, and a seed energy of 1nJ. With these high gains, the expected output energy is on the order of 1-10μJ.

The second type of system is called the "low gain design", and this similar to our third OPCPA stage. Here, a high energy pump amplifies a moderate energy seed pulse, and achieves modest gains, but significant energy output. In our models, we assumed a 5mJ pump pulse and a 10μJ seed pulse. Even with moderate gain, typically on the order 10-150, this can give a mJ level output pulse. An actual system would likely employ one "high gain" stage, and one "low gain" stage.

At this stage, we decided to focus our theoretical work on ZGP, as it can be made with sufficient apertures to be used in OPA configurations. OP:GaAs is currently available with apertures that would necessitate OPCPA usage, and it is preferable to avoid the need to stretch and compress a 7.2μm pulse.

The numerical model uses a linearized split step Fourier method, a well developed method for modeling nonlinear pulse propagation [37, 119]. In the split step method, the propagation distance

is first divided into a large number of small, discrete steps. For each step through the material, the nonlinear propagation equations given by Eqns. 3.6 are split into two parts: a nonlinear term describing the relevant interaction being modeled, and a dispersion term describing linear pulse propagation. At each step, the nonlinear term is linearized, then solved in the time domain, while the linear term is solved in the frequency domain. The method has the advantage of producing accurate results, while allowing for fast computation times [112].

6.3.4.2 Infrared Pumped OPA / OPCPA: Numerical Results

The results of these simulations indicate that both OP:GaAs and ZGP hold potential for mJ level amplification at 7.2μm. In particular, we have found that a two stage, 3.1μm pumped ZGP based OPA is particularly promising architecture for mJ level, femtosecond pulse generation at 7.2μm. The parameters for both stages are summarized in Table. 6.5. The first OPA stage uses

Table 6.5: The parameters for each of the two stages for a 3.1μm pumped ZGP based OPA for mJ level, femtosecond pulse generation at 7.2μm. The seed parameters for the second OPA stage are based upon the first OPA stage output parameters.

	1^{st} OPA Stage	2^{nd} OPA Stage
Material	ZGP	ZGP
Cut	$\theta = 46.1^{\circ}$	$\theta = 46.1^{\circ}$
Length	2.50mm	0.88mm
Spot Diameter	5mm	15mm
Pump Wavelength	3.1μm	3.1μm
Pump Bandwidth	100nm FWHM	100nm FWHM
Pump Pulse Length	200fs	200fs
Pump Energy	500μJ	5mJ
Seed Wavelength	7.2μm	7.2μm
Seed Bandwidth	400nm FWHM	800nm FWHM (1^{st} Stage Output)
Output Bandwidth	~805nm	~790nm
Seed Pulse Length	~190fs	~130fs (1^{st} Stage Output)
Output Pulse Length	~ 130fs	~150fs
Seed Energy	1nJ	120μJ (1^{st} Stage Output)
Gain	120 $\times 10^6$	14
Output Energy	120μJ	1.68mJ

a 500μJ pump pulse, assumed to be 200fs long, with a spot size of 5mm. The pump spectrum is assumed to be a Gaussian with 100nm FWHM bandwidth centered around 3.1μm. The seed is assumed to be a 1nJ pulse, with 400nm FWHM bandwidth centered at 7.2μm. With a 2.5mm long ZGP crystal, we calculate a gain of $\sim 1.2 \times 10^5$, giving an output energy of 120μJ. The bandwidth is substantially broadened, and the output bandwidth is approximately 800nm FWHM. The input and output spectra and temporal profiles are shown in Fig. 6.11. Temporally, the 7.2μm output

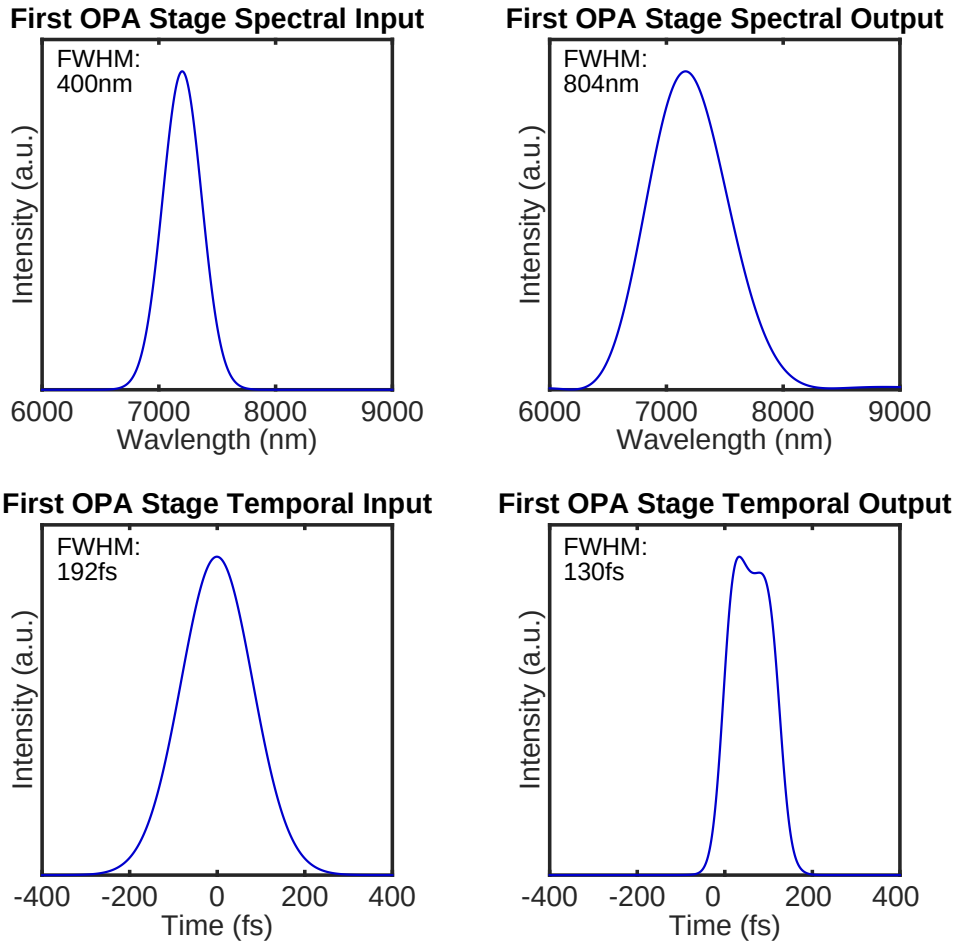

Figure 6.11: The modeled input and output spectra and temporal profiles for the first ZGP OPA stage. The input pulse is a 400nm FWHM Gaussian centered at 7.2μm, with a FWHM pulse length of 190s. At the output, the spectrum is broadened to 804nm FHWM, and the pulse length is 130s. The predicted gain in the first stage is 1.2×10^5, for an output pulse energy of 120μJ.

pulse is narrowed from approximately 190s to 130fs FWHM.

In the second stage, the pump is a 5mJ pulse with the same spectral and temporal properties as in the first stage. The 7.2μm seed is provided by the first stage 7.2μm output. With a spot size of 15mm, and a crystal length of 0.88mm, the seed is expected to see a gain for ~ 14, which gives an output energy of $\sim$1.65mJ. The output bandwidth isn't expected to change significantly, only narrowing slightly from the $\sim$800nm FWHM input to approximately $\sim$ 790nm FWHM. Temporally the pulse is expected to lengthen from approximately 130fs on the input to 150fs FWHM at the output. The input and output spectra for the second OPA stage can be seen in Fig. 6.12.

These simulations show that a 3.1μm pumped ZGP OPA provides a relatively straightforward path towards generating mJ level femtosecond pulses further into the infrared. This is a particularly interesting use for the 3μm OPCPA laser, as the main goal for this laser is to enable keV high harmonic generation. The theoretical phase matching cutoff for high harmonic emission driven at 7.2μm is greater than 4keV [93].

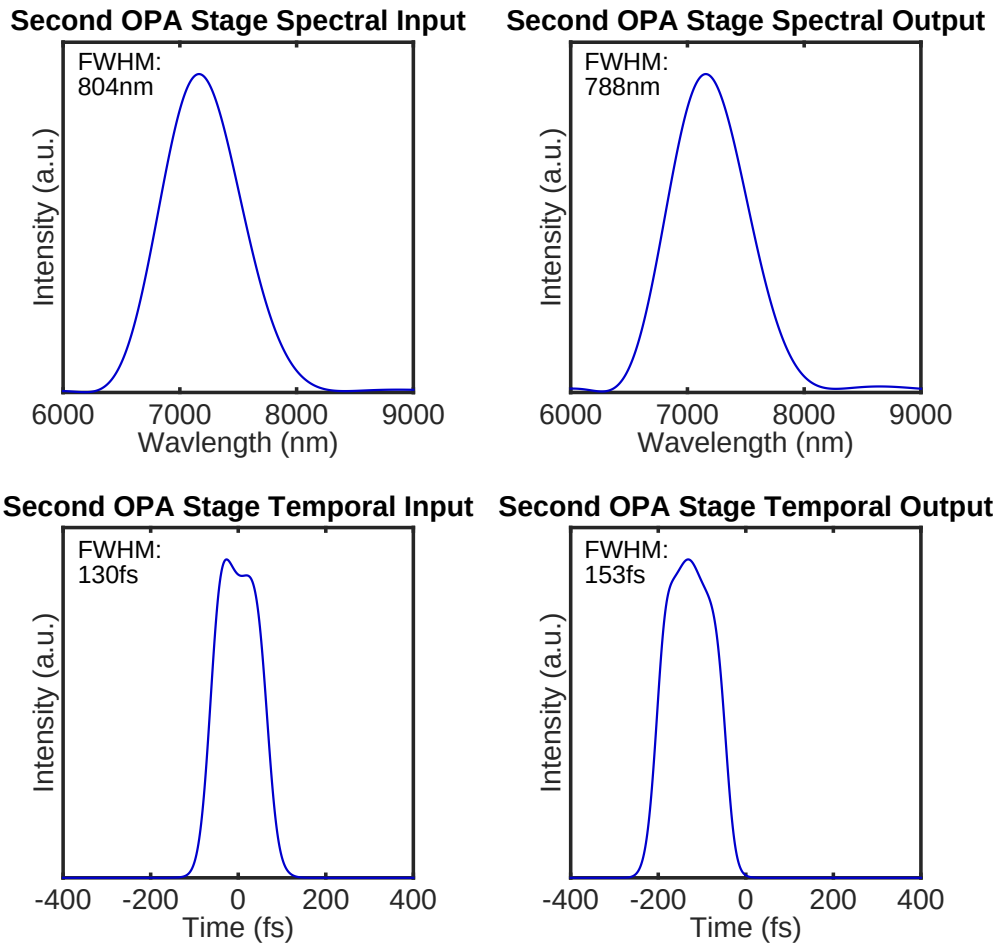

Figure 6.12: The modeled input and output spectra and temporal profiles for the second ZGP OPA stage. The input pulse is the 804nm FWHM output from the first OPA stage, with a pulse length of 130s. At the output, the spectrum is narrowed slightly to 790 FHWM, and the pulse length is 150s. The predicted gain in the second stage is 14, for an output pulse energy of 1.65mJ.

Bibliography

[1] Phase matching of high-order harmonics in hollow waveguides. Phys. Rev. Lett., 83:2187–2190, Sep 1999.

[2] Ossama R. Abdelsalam and Ramin Sedaghati. Design optimization of compound cylinders subjected to autofrettage and shrink-fitting processes. J. Pressure Vessel Technol., 135:021209, 2013.

[3] R. L. Aggarwal, D. J. Ripin, J. R. Ochoa, and T. Y. Fan. Measurement of thermo-optic properties of $Y_3Al_5O_{12}$, $Lu_3Al_5O_{12}$, $YAIO_3$, $LiYF_4$, $LiLuF_4$, BaY_2F_8, $KGd(WO_4)_2$, and $KY(WO_4)_2$ laser crystals in the 80-300K temperature range. Journal of Applied Physics, 98(10):103514, 2005.

[4] Govind Agrawal. Nonlinear Fiber Optics. Academic Press, fifth edition, 2013.

[5] V A Akimov, V I Kozlovskii, Yu V Korostelin, A I Landman, Yu P Podmar'kov, Ya K Skasyrskii, and M P Frolov. Efficient pulsed Cr^{2+}:CdSe laser continuously tunable in the spectral range from 2.26 to 3.61μm. Quantum Electronics, 38(3):205, 2008.

[6] M. V. Ammosov, N. B. Delone, and V. P. Krainov. Tunnel ionization of complex atoms and of atomic ions in an alternating electromagnetic field. Soviet Physics - JETP, 64(6):1191–1194, 1986.

[7] 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. Opt. Lett., 36(15):2755–2757, Aug 2011.

[8] Jonathan B. Ashcom, Rafael R. Gattass, Chris B. Schaffer, and Eric Mazur. Numerical aperture dependence of damage and supercontinuum generation from femtosecond laser pulses in bulk fused silica. J. Opt. Soc. Am. B, 23(11):2317–2322, Nov 2006.

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

[10] Andrius Baltuška, Takao Fuji, and Takayoshi Kobayashi. Controlling the carrier-envelope phase of ultrashort light pulses with optical parametric amplifiers. Phys. Rev. Lett., 88:133901, Mar 2002.

[11] Randy A. Bartels, Ariel Paul, Hans Green, Henry C. Kapteyn, Margaret M. Murnane, Sterling Backus, Ivan P. Christov, Yanwei Liu, David Attwood, and Chris Jacobsen. Generation of spatially coherent light at extreme ultraviolet wavelengths. Science, 297(5580):376–378, 2002.

[12] N. Boeuf, D. Branning, I. Chaperot, E. Dauler, S. Guerin, G. Jaeger, A. Muller, and A. Migdall. Calculating characteristics of noncollinear phase matching in uniaxial and biaxial crystals. Optical Engineering, 39(4):1016–1024, 2000.

[13] Robert W. Boyd. Nonlinear Optics. Academic Press, second edition, 1967.

[14] J.K. Brasseur, AK. Abeeluck, AR. Awtry, L.S. Meng, K.E. Shortoff, N.J. Miller, R.K. Hampton, M.H. Cuchiara, and D.K. Neumann. 2.3 kW cryogenically cooled Yb:YAG laser. In Lasers and Electro-Optics, 2009 and 2009 Conference on Quantum electronics and Laser Science Conference. CLEO/QELS 2009. Conference on, pages 1–2, June 2009.

[15] David Brown, Joseph Singley, Katie Kowalewski, James Guelzow, and Victoria Vitali. High sustained average power cw and ultrafast Yb:YAG near-diffraction-limited cryogenic solid-state laser. Optics express, 18(24):24770–24792, 2010.

[16] D. A. Bryan, Robert Gerson, and H. E. Tomaschke. Increased optical damage resistance in lithium niobate. Applied Physics Letters, 44(9):847–849, 1984.

[17] Giulio Cerullo and Sandro De Silvestri. Ultrafast optical parametric amplifiers. Review of Scientific Instruments, 74(1), 2003.

[18] Mathieu Charbonneau-Lefort, Bedros Afeyan, and M. M. Fejer. Optical parametric amplifiers using chirped quasi-phase-matching gratings i: practical design formulas. J. Opt. Soc. Am. B, 25(4):463–480, Apr 2008.

[19] 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. Phys. Rev. Lett., 105:173901, Oct 2010.

[20] Ming-Chang Chen, Christopher Mancuso, Carlos Hernndez-Garca, Franklin Dollar, Ben Galloway, Dimitar Popmintchev, Pei-Chi Huang, Barry Walker, Luis Plaja, Agnieszka A. Jaro?-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. Proceedings of the National Academy of Sciences, 111(23):E2361–E2367, 2014.

[21] Andy Chong, Joel Buckley, Will Renninger, and Frank Wise. All-normal-dispersion femtosecond fiber laser. Opt. Express, 14(21):10095–10100, Oct 2006.

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

[23] M. N. Cizmeciyan, J. W. Kim, S. Bae, B. H. Hong, F. Rotermund, and A. Sennaroglu. Graphene mode-locked femtosecond Cr:ZnSe laser at 2500nm. Opt. Lett., 38(3):341–343, Feb 2013.

[24] W. A. Clarkson, N. S. Felgate, and D. C. Hanna. Simple method for reducing the depolarization loss resulting from thermally induced birefringence in solid-state lasers. Opt. Lett., 24(12):820–822, Jun 1999.

[25] Stéphane Coen, Alvin Hing Lun Chau, Rainer Leonhardt, John D. Harvey, Jonathan C. Knight, William J. Wadsworth, and Philip St. J. Russell. White-light supercontinuum generation with 60-ps pump pulses in a photonic crystal fiber. Opt. Lett., 26(17):1356–1358, Sep 2001.

[26] M. Conforti, F. Baronio, C. De Angelis, M. Marangoni, and G. Cerullo. Theory and experiments on multistep parametric processes in nonlinear optics. J. Opt. Soc. Am. B, 28(4):892–895, Apr 2011.

[27] Alexey Dergachev, Darrell Armstrong, Arlee Smith, Thomas Drake, and Marc Dubois. High-power, high-energy ZGP OPA pumped by a 2.05μm Ho:YLF MOPA system, 2008.

[28] Jean-Claude Diels and Wolfgang Rudolph. Ultrashort Laser Pulse Phenomena. Academic Press, second edition, 2006.

[29] M. Digonnet, M. Fejer, and R. Byer. Characterization of proton-exchanged waveguides in MgO:$LiNbO_3$. Opt. Lett., 10(5):235–237, May 1985.

[30] Chengyuan Ding, Wei Xiong, Tingting Fan, Daniel D. Hickstein, Tenio Popmintchev, Xiaoshi Zhang, Mike Walls, Margaret M. Murnane, and Henry C. Kapteyn. High flux coherent supercontinuum soft x-ray source driven by a single-stage, 10mj, ti:sapphire amplifier-pumped opa. Opt. Express, 22(5):6194–6202, Mar 2014.

[31] Jun Dong, Michael Bass, Yanli Mao, Peizhen Deng, and Fuxi Gan. Dependence of the Yb^{3+} emission cross section and lifetime on temperature and concentration in yttrium aluminum garnet. JOSA B, 20(9):1975–1979, 2003.

[32] A Dubietis, G Jonušauskas, and A Piskarskas. Powerful femtosecond pulse generation by chirped and stretched pulse parametric amplification in BBO crystal. Optics Communications, 88(4):437–440, 1992.

[33] S Düsterer, P Radcliffe, C Bostedt, J Bozek, A L Cavalieri, R Coffee, J T Costello, D Cubaynes, L F DiMauro, Y Ding, G Doumy, F Grüner, W Helml, W Schweinberger, R Kienberger, A R Maier, M Messerschmidt, V Richardson, C Roedig, T Tschentscher, and M Meyer. Femtosecond x-ray pulse length characterization at the linac coherent light source free-electron laser. New Journal of Physics, 13(9):093024, 2011.

[34] Tso Yee Fan, Daniel J Ripin, Roshan L Aggarwal, Juan R Ochoa, Bien Chann, Michael Tilleman, and Joshua Spitzberg. Cryogenic Yb^{3+}-doped solid-state lasers. Ieee J Sel Top Quant, 13(3):448–459, Jan 2007.

[35] V.V. Fedorov, D.V. Martyshkin, M. Mirov, I Moskalev, S. Vasyliev, and S.B. Mirov. High energy 4.1-4.6μm Fe:ZnSe laser. In Lasers and Electro-Optics (CLEO), 2012 Conference on, pages 1–2, May 2012.

[36] V.V. Fedorov, S.B. Mirov, A Gallian, D.V. Badikov, M.P. Frolov, Y.V. Korostelin, V.I Kozlovsky, AI Landman, Y.P. Podmar'kov, V.A Akimov, and AA Voronov. 3.77-5.05μm tunable solid-state lasers based on Fe^{2+}-doped ZnSe crystals operating at low and room temperatures. Quantum Electronics, IEEE Journal of, 42(9):907–917, Sept 2006.

[37] Robert A. Fisher and William K. Bischel. Numerical studies of the interplay between self-phase modulation and dispersion for intense plane-wave laser pulses. Journal of Applied Physics, 46(11), 1975.

[38] R. Fluck, M. R. Hermann, and L. A. Hackel. Birefringence compensation in single solid-state rods. Applied Physics Letters, 76(12):1513–1515, 2000.

[39] P. A. Franken, A. E. Hill, C. W. Peters, and G. Weinreich. Generation of optical harmonics. Physical Review Letters, 7:118–120, 1961.

[40] L. Gallmann, G. Steinmeyer, U. Keller, G. Imeshev, M. M. Fejer, and J.-P. Meyn. Generation of sub-6-fs blue pulses by frequency doubling with quasi-phase-matching gratings. Opt. Lett., 26(9):614–616, May 2001.

[41] David M. Gaudiosi, Etienne Gagnon, Amy L. Lytle, Julie L. Fiore, Emily A. Gibson, Steve Kane, Jeff Squier, Margaret M. Murnane, Henry C. Kapteyn, Ralph Jimenez, and Sterling Backus. Multi-kilohertz repetition rate Ti:sapphire amplifier based on down-chirped pulse amplification. Opt. Express, 14(20):9277–9283, Oct 2006.

[42] David M. Gaudiosi, Amy L. Lytle, Pat Kohl, Margaret M. Murnane, Henry C. Kapteyn, and Sterling Backus. 11-w average power ti:sapphire amplifier system using downchirpedpulse amplification. Opt. Lett., 29(22):2665–2667, Nov 2004.

[43] Theodor Seuss Geisel. The Cat in the Hat. Random House, 1957.

[44] Michael Gerrity, Susannah Brown, Tenio Popmintchev, Ming-Chang Chen, Stefan Witte, Margaret M. Murnane, Henry C. Kapteyn, and Sterling Backus. High power, 60MHz, cryogenically cooled, mode-locked, Yb:YAG oscillator. In CLEO:2011 - Laser Applications to Photonic Applications, page CThAA4. Optical Society of America, 2011.

[45] Emily A. Gibson, David M. Gaudiosi, Henry C. Kapteyn, Ralph Jimenez, Steve Kane, Rachel Huff, Charles Durfee, and Jeff Squier. Efficient reflection grisms for pulse compression and dispersion compensation of femtosecond pulses. Opt. Lett., 31(22):3363–3365, Nov 2006.

[46] C. J. Glassbrenner and Glen A. Slack. Thermal Conductivity of Silicon and Germanium from 3°K to the Melting Point. Phys. Rev., 134:A1058–A1069, May 1964.

[47] Juliet T. Gopinath, Kevin F. Wall, John Hybl, Peter F. Moulton, and T. Y. Fan. High-power, actively modelocked cryogenic Yb:YAG laser. In Lasers, Sources and Related Photonic Devices, page AWB22. Optical Society of America, 2010.

[48] C. Hernández-García, J. A. Pérez-Hernández, T. Popmintchev, M. M. Murnane, H. C. Kapteyn, A. Jaron-Becker, A. Becker, and L. Plaja. Zeptosecond high harmonic kev x-ray waveforms driven by midinfrared laser pulses. Phys. Rev. Lett., 111:033002, Jul 2013.

[49] Carlos Hernandez-Garcia, Tenio Popmintchev, Margaret M. Murnane, Henry C. Kapteyn, Luis Plaja, Agnieszka A. Jaron-Becker, and Andreas Becker. Efficient generation of isolated attosecond soft x-ray pulses. In Frontiers in Optics 2014, page JTu5G.3. Optical Society of America, 2014.

[50] U. Hömmerich, X. Wu, V. R. Davis, S. B. Trivedi, K. Grasza, R. J. Chen, and S. Kutcher. Demonstration of room-temperature laser action at 2.5 μm from Cr^{2+}:$Cd_{0.85}Mn_{0.15}Te$. Opt. Lett., 22(15):1180–1182, Aug 1997.

[51] Kyung-Han Hong, Juliet T. Gopinath, Darren Rand, Aleem M. Siddiqui, Shu-Wei Huang, Enbang Li, Benjamin J. Eggleton, John D. Hybl, Tso Yee Fan, and Franz X. Kärtner. High-energy, kHz-repetition-rate, ps cryogenic Yb:YAG chirped-pulse amplifier. Opt. Lett., 35(11):1752–1754, Jun 2010.

[52] Nobuhisa Ishii, Keisuke Kaneshima, Kenta Kitano, Teruto Kanai, Shuntaro Watanabe, and Jiro Itatani. Carrier-envelope phase-dependent high harmonic generation in the water window using few-cycle infrared pulses. Nat Commun, 5 SP -, Feb 2014.

[53] A.G. Kalintsev, U.V. Katsev, A.F. Kornev, A.S. Narivonchik, D.O. Oborotov, A.L. Pavlova, V.P. Pokrovskyi, V.A. Serebryakov, V.K. Stupnikov, and S.S. Terekhov. 100 mJ/100 Hz Mid-IR laser source. In Laser Optics, 2014 International Conference, pages 1–1, June 2014.

[54] A.A. Kaminskii, M.Sh. Akchurin, V.I. Alshits, K. Ueda, K. Takaichi, J. Lu, T. Uematsu, M. Musha, A. Shirikawa, V. Gabler, H.J. Eichler, H. Yagi, T. Yanagitani, S.N. Bagayev, J. Fernandez, and R. Balda. New data on the physical properties of $Y_3Al_5O_12$-based nanocrystalline laser ceramics. Crystallography Reports, 48(3):515–519, 2003.

[55] S. Kane and J. Squier. Grism-pair stretcher–compressor system for simultaneous second- and third-order dispersion compensation in chirped-pulse amplification. J. Opt. Soc. Am. B, 14(3):661–665, Mar 1997.

[56] Tolga Kartaloğlu, Z. Gürkan Figen, and Orhan Aytür. Simultaneous phase matching of optical parametric oscillation and second-harmonic generation in aperiodically poled lithium niobate. J. Opt. Soc. Am. B, 20(2):343–350, Feb 2003.

[57] Vincent Kemlin, Benoit Boulanger, Valentin Petrov, Patricia Segonds, B. Ménaert, Peter G. Schunneman, and Kevin T. Zawilski. Nonlinear, dispersive, and phase-matching properties of the new chalcopyrite CdSiP2. Opt. Mater. Express, 1(7):1292–1300, Nov 2011.

[58] Matthew Kirchner, Andrew Niedringhaus, Charles G. Durfee, Frank W. Wise, Daisy Raymondson, Lora Nugent-Glandorf, Henry C. Kapteyn, Margaret M. Murnane, and Sterling Backus. Ultrafast optical parametric oscillator pumped by an all normal dispersion (andi) yb:fiber oscillator. In Conference on Lasers and Electro-Optics 2012, page CM1B.1. Optical Society of America, 2012.

[59] Walter Koechner. Solid State Laser Engineering. Springer New York, 2006.

[60] Joerg Koerner, Christian Vorholt, Hartmut Liebetrau, Martin Kahle, Diethard Kloepfel, Reinhard Seifert, Joachim Hein, and Malte C. Kaluza. Measurement of temperature-dependent absorption and emission spectra of Yb:YAG, Yb:LuAG, and Yb:CaF_2 between 20°C and 200°C and predictions on their influence on laser performance. J. Opt. Soc. Am. B, 29(9):2493–2502, Sep 2012.

[61] V. P. Krainov. Ionization rates and energy and angular distributions at the barrier-suppression ionization of complex atoms and atomic ions. J. Opt. Soc. Am. B, 14(2):425–431, Feb 1997.

[62] P. S. Kuo, K. L. Vodopyanov, M. M. Fejer, D. M. Simanovskii, X. Yu, J. S. Harris, D. Bliss, and D. Weyburne. Optical parametric generation of a mid-infrared continuum in orientation-patterned GaAs. Opt. Lett., 31(1):71–73, Jan 2006.

[63] Nick Leindecker, Alireza Marandi, Robert L. Byer, Konstantin L. Vodopyanov, Jie Jiang, Ingmar Hartl, Martin Fermann, and Peter G. Schunemann. Octave-spanning ultrafast OPO with 2.6-6.1µm instantaneous bandwidth pumped by femtosecond Tm-fiber laser. Opt. Express, 20(7):7046–7053, Mar 2012.

[64] M. Lewenstein, Ph. Balcou, M. Yu. Ivanov, Anne L'Huillier, and P. B. Corkum. Theory of high-harmonic generation by low-frequency laser fields. Phys. Rev. A, 49:2117–2132, Mar 1994.

[65] Qing Li, Kathleen Hoogeboom-Pot, Damiano Nardi, Margaret M. Murnane, Henry C. Kapteyn, Mark E. Siemens, Erik H. Anderson, Olav Hellwig, Elizabeth Dobisz, Bruce Gurney, Ronggui Yang, and Keith A. Nelson. Generation and control of ultrashort-wavelength two-dimensional surface acoustic waves at nanoscale interfaces. Phys. Rev. B, 85:195431, May 2012.

[66] J. Limpert, T. Gabler, A. Liem, H. Zellmer, and A. Tnnermann. Spm-induced spectral compression of picosecond pulses in a single-mode yb-doped fiber amplifier. Applied Physics B, 74(2):191–195, 2002.

[67] Steven D. Lord. A new software tool for computing earth's atmospheric transmission of near- and far-infrared radiation. NASA Technical Memorandum, (103957), 1994.

[68] G. Z. Luo, S. N. Zhu, J. L. He, Y. Y. Zhu, H. T. Wang, Z. W. Liu, C. Zhang, and N. B. Ming. Simultaneously efficient blue and red light generations in a periodically poled litao3. Applied Physics Letters, 78(20), 2001.

[69] Christopher A. Mancuso, Ming-Chang Chen, Carlos Hernandez-Garcia, Franklin Dollar, Benjamin Galloway, Dimitar Popmintchev, Benjamin Langdon, Amelie Auger, P. C. Huang, Barry C. Walker, Luis Plaja, Agnieszka Jaron-Becker, Andreas Becker, Margaret Murnane, Henry Kapteyn, and Tenio Popmintchev. Generation of bright isolated attosecond soft x-ray pulses driven by multi-cycle mid-infrared lasers. In CLEO: 2014, page FTu3B.5. Optical Society of America, 2014.

[70] E. A. J. Marcatili and R. A. Schmeltzer. Hollow metallic and dielectric waveguides for long distance optical transmission and lasers. Bell Systems Technical Journal, 43:1783–1809, 1964.

[71] O.E. Martinez. 3000 times grating compressor with positive group velocity dispersion: Application to fiber compensation in 1.3-1.6 μm region. Quantum Electronics, IEEE Journal of, 23(1):59–64, Jan 1987.

[72] O.E. Martinez. Design of high-power ultrashort pulse amplifiers by expansion and recompression. Quantum Electronics, IEEE Journal of, 23(8):1385–1387, Aug 1987.

[73] Stefan Mathias, Chan La o vorakiat, Justin M. Shaw, Emrah Turgut, Patrik Grychtol, Roman Adam, Dennis Rudolf, Hans T. Nembach, Thomas J. Silva, Martin Aeschlimann, Claus M. Schneider, Henry C. Kapteyn, and Margaret M. Murnane. Ultrafast element-specific magnetization dynamics of complex magnetic materials on a table-top. Journal of Electron Spectroscopy and Related Phenomena, 189(0):164 – 170, 2013.

[74] 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 optical parametric chirped-pulse amplifier at 3.4μm. Opt. Lett., 38(21):4265–4268, Nov 2013.

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

[76] Jason McKay, Kenneth L. Schepler, and Gary Catella. Kilohertz, 2.6-μm Cr^{2+}:CdSe Laser. In Advanced Solid State Lasers, page WD1. Optical Society of America, 1999.

[77] A. McPherson, G. Gibson, H. Jara, U. Johann, T. S. Luk, I. A. McIntyre, K. Boyer, and C. K. Rhodes. Studies of multiphoton production of vacuum-ultraviolet radiation in the rare gases. J. Opt. Soc. Am. B, 4:595–601, 1987.

[78] A. A. Miraje1 and S. A. Patil. Minimization of material volume of three layer compound cylinder having same materials subjected to internal pressure. Int. J. Eng. Sci. and Tech., 3:26–40, 2012.

[79] S. Mirov, V. Fedorov, I. Moskalev, D. Martyshkin, and C. Kim. Progress in Cr^{2+} and Fe^{2+} doped mid-IR laser materials. Laser Photonics Reviews, 4(1):21–41, 2010.

[80] Jeffrey Moses, Cristian Manzoni, Shu-Wei Huang, Giulio Cerullo, and Franz X. Kaertner. Temporal optimization of ultrabroadband high-energy OPCPA. Opt. Express, 17(7):5540–5555, Mar 2009.

[81] Peter Moulton and Evgeni Slobodchikov. 1-GW-peak-power, Cr:ZnSe laser. In CLEO:2011 - Laser Applications to Photonic Applications, page PDPA10. Optical Society of America, 2011.

[82] L. E. Myers, R. C. Eckardt, M. M. Fejer, R. L. Byer, W. R. Bosenberg, and J. W. Pierce. Quasi-phase-matched optical parametric oscillators in bulk periodically poled LiNbO3. J. Opt. Soc. Am. B, 12(11):2102–2116, Nov 1995.

[83] NoSoung Myoung, M.S. Mirov, V.V. Fedorov, and S.B. Mirov. High-energy gain-switched mid-IR lasers based on Cr and Fe doped ZnSe. In 2011 Conference on Lasers and Electro-Optics (CLEO), pages 1–2, May 2011.

[84] Gemini Observatory. IR Transmission Spectra. Downloaded from website.

[85] National Institute of Standards and Technology. NIST Standard Reference Material 736 - Copper - Thermal Expansion.

[86] J. S. Pelc, C. R. Phillips, D. Chang, C. Langrock, and M. M. Fejer. Efficiency pedestal in quasi-phase-matching devices with random duty-cycle errors. Opt. Lett., 36(6):864–866, Mar 2011.

[87] M. D. Perry, T. Ditmire, and B. C. Stuart. Self-phase modulation in chirped-pulse amplification. Opt. Lett., 19(24):2149–2151, Dec 1994.

[88] O. Pfister, J. S. Wells, L. Hollberg, L. Zink, D. A. Van Baak, M. D. Levenson, and W. R. Bosenberg. Continuous-wave frequency tripling and quadrupling by simultaneous three-wave mixings in periodically poled crystals: application to a two-step 1.19-10.71μm frequency bridge. Opt. Lett., 22(16):1211–1213, Aug 1997.

[89] C. R. Phillips, L. Gallmann, and M. M. Fejer. Design of quasi-phasematching gratings via convex optimization. Opt. Express, 21(8):10139–10159, Apr 2013.

[90] C. R. Phillips, J. Jiang, C. Mohr, A. C. Lin, C. Langrock, M. Snure, D. Bliss, M. Zhu, I. Hartl, J. S. Harris, M. E. Fermann, and M. M. Fejer. Widely tunable midinfrared difference frequency generation in orientation-patterned GaAs pumped with a femtosecond Tm-fiber system. Opt. Lett., 37(14):2928–2930, Jul 2012.

[91] 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. Opt. Express, 22(8):9627–9658, Apr 2014.

[92] I. Pipinyt, R. Grigonis, K. Stankeviit, S. Kias, R. Drazdys, R. C. Eckardt, and V. Sirutkaitis. Laser-induced-damage threshold of periodically poled lithium niobate for 1030 nm femtosecond laser pulses at 100 kHz and 75 MHz , 2013.

[93] 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. Proceedings of the National Academy of Sciences, 106(26):10516–10521, 2009.

[94] 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. Opt. Lett., 33(18):2128–2130, Sep 2008.

[95] Tenio Popmintchev, Ming-Chang Chen, Dimitar Popmintchev, Paul Arpin, Susannah Brown, Skirmantas Aliauskas, Giedrius Andriukaitis, Tadas Bal?iunas, Oliver D. Mcke, Audrius Pugzlys, Andrius Baltuka, Bonggu Shim, Samuel E. Schrauth, Alexander Gaeta, Carlos Hernndez-Garca, 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(6086):1287–1291, 2012.

[96] DJ Ripin, JR Ochoa, RL Aggarwal, and TY Fan. 165-W cryogenically cooled Yb:YAG laser. Optics letters, 29(18):2154–2156, Jan 2004.

[97] Andy Rundquist, Charles G. Durfee, 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, 1998.

[98] P Russbueldt, T Mans, G Rotarius, J Weitenberg, HD Hoffmann, and R Poprawe. 400W Yb: YAG Innoslab fs-amplifier. Optics express, 17(15):12230–12245, 2009.

[99] Philip Russell. Photonic crystal fibers. Science, 299(5605):358–362, 2003.

[100] Peter G. Schunemann, Kevin T. Zawilski, Thomas M. Pollak, Valentin Petrov, and David E. Zelmon. Cdsip2: A new nonlinear optical crystal for 1- and 1.5-micron-pumped mid-ir generation. In Advanced Solid-State Photonics, page TuC6. Optical Society of America, 2009.

[101] Scientific Materials. Yb:YAG Properties. Downloaded from website.

[102] Matthew D. Seaberg, Bosheng Zhang, Dennis F. Gardner, Elisabeth R. Shanblatt, Margaret M. Murnane, Henry C. Kapteyn, and Daniel E. Adams. Tabletop nanometer extreme ultraviolet imaging in an extended reflection mode using coherent fresnel ptychography. Optica, 1(1):39–44, Jul 2014.

[103] Anthony E. Siegman. Lasers. University Science Books, 1986.

[104] F. Silva, D. R. Austin, A. Thai, M. Baudisch, M. Hemmer, D. Faccio, A. Couairon, and J. Biegert. Multi-octave supercontinuum generation from mid-infrared filamentation in a bulk crystal. Nat Commun, 3:807, 05 2012.

[105] O.V. Sinkin, R. Holzlohner, J. Zweck, and C.R. Menyuk. Optimization of the split-step fourier method in modeling optical-fiber communications systems. Lightwave Technology, Journal of, 21(1):61–68, Jan 2003.

[106] T. Skauli, K. L. Vodopyanov, T. J. Pinguet, A. Schober, O. Levi, L. A. Eyres, M. M. Fejer, J. S. Harris, B. Gerard, L. Becouarn, E. Lallier, and G. Arisholm. Measurement of the nonlinear coefficient of orientation-patterned GaAs and demonstration of highly efficient second-harmonic generation. Opt. Lett., 27(8):628–630, Apr 2002.

[107] Arlee Smith. SNLO. Downloaded from http://www.as-photonics.com/snlo, 2010. Version 56.

[108] E. Sorokin, S. Naumov, and IT. Sorokina. Ultrabroadband infrared solid-state lasers. Selected Topics in Quantum Electronics, IEEE Journal of, 11(3):690–712, May 2005.

[109] Evgeni Sorokin, Nikolai Tolstik, and Irina T. Sorokina. 1 watt femtosecond mid-IR Cr:ZnS laser, 2013.

[110] IT. Sorokina, V.V. Dvoyrin, N. Tolstik, and E. Sorokin. Mid-IR ultrashort pulsed fiber-based lasers. Selected Topics in Quantum Electronics, IEEE Journal of, 20(5):1–12, Sept 2014.

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

[112] Thiab R Taha and Mark I Ablowitz. Analytical and numerical aspects of certain nonlinear evolution equations. ii. numerical, nonlinear schrodinger equation. Journal of Computational Physics, 55(2):203 – 230, 1984.

[113] E. Treacy. Optical pulse compression with diffraction gratings. Quantum Electronics, IEEE Journal of, 5(9):454–458, Sep 1969.

[114] K. L. Vodopyanov, F. Ganikhanov, J. P. Maffetone, I. Zwieback, and W. Ruderman. Zngep2 optical parametric oscillator with 3.8—12.4-μm tunability. Opt. Lett., 25(11):841–843, Jun 2000.

[115] K. L. Vodopyanov, O. Levi, P. S. Kuo, T. J. Pinguet, J. S. Harris, M. M. Fejer, B. Gerard, L. Becouarn, and E. Lallier. Optical parametric oscillation in quasi-phase-matched gaas. Opt. Lett., 29(16):1912–1914, Aug 2004.

[116] K. L. Vodopyanov, I. Makasyuk, and P. G. Schunemann. Grating tunable 4 - 14 µm GaAs optical parametric oscillator pumped at 3 µm. Opt. Express, 22(4):4131–4136, Feb 2014.

[117] Ke Wang, Liejia Qian∗, Hang Luo, Peng Yuan, and Heyuan Zhu. Ultrabroad supercontinuum generation by femtosecond dual-wavelength pumping in sapphire. Opt. Express, 14(13):6366–6371, Jun 2006.

[118] L. J. Waxer, V. Bagnoud, I. A. Begishev, M. J. Guardalben, J. Puth, and J. D. Zuegel. High-conversion-efficiency optical parametric chirped-pulse amplification system using spatiotemporally shaped pump pulses. Opt. Lett., 28(14):1245–1247, Jul 2003.

[119] S. Witte, R.T. Zinkstok, W. Hogervorst, and K.S.E. Eikema. Numerical simulations for performance optimization of a few-cycle terawatt nopcpa system. Applied Physics B, 87(4):677–684, 2007.

[120] E. Yahel, O. Hess, and A.A. Hardy. Ultrashort-pulse high-power yb3+-doped fiber amplifiers. Quantum Electronics, IEEE Journal of, 43(9):824–832, Sept 2007.

[121] Xin Yin, Shaojun Zhang, and Zhaobing Tian. Refractive indices of biaxial crystals evaluated from the refractive indices ellipsoid equation. Optics Laser Technology, 39(3):510 – 513, 2007.

[122] Kevin T. Zawilski, Peter G. Schunemann, Scott D. Setzler, and Thomas M. Pollak. Large aperture single crystal zngep2 for high-energy applications. Journal of Crystal Growth, 310(7–9):1891 – 1896, 2008. the Proceedings of the 15th International Conference on Crystal Growth (ICCG-15) in conjunction with the International Conference on Vapor Growth and Epitaxy and the {US} Biennial Workshop on Organometallic Vapor Phase Epitaxy.

[123] Zhaowei Zhang, Derryck T. Reid, S. Chaitanya Kumar, Majid Ebrahim-Zadeh, Peter G. Schunemann, Kevin T. Zawilski, and Christopher R. Howle. Femtosecond-laser pumped CdSiP2 optical parametric oscillator producing 100MHz pulses centered at 6.2μm. Opt. Lett., 38(23):5110–5113, Dec 2013.

Appendix A

An Intuitive Explanation of the Relationship Between the Signal and Idler Phase

A crucial factor in the design of this system was determining how to stretch the signal pulse in a manner that would allow the idler to be easily compressed. This is not an issue for traditional CPA systems, where the same pulse is being both stretched, and the later compressed. For that case, there are numerous stretcher / compressor combinations that give equal and opposite phase, the most commonly used being a Martinez stretcher and a Treacy compressor (also known as positive and negative dispersion stretchers, respectively) [72, 71, 113].

For this OPCPA system, because we would be stretching the signal and compressing the idler, it was important to understand exactly how the idler phase depends on the input signal phase. The mathematics behind this were covered in Chapter 3, which led to

$$\phi_s(\omega_s) = -\phi_i(\omega_p - \omega_s), \tag{A.1}$$

and similarily,

$$\phi_i(\omega_i) = -\phi_s(\omega_p - \omega_i). \tag{A.2}$$

These two equation contain the entirety of the relation between the phase and the idler, but they are rather opaque, and not very useful for understanding the physics of what is happening. Fortunately, there is a very simple way to understand exactly how the phase from the signal is transferred to the idler.

To begin this intuitive explanation, let's first return to the photon picture of optical parametric amplification. At this level, an individual pump photon is split into two photons. This gives

us two equations, one for energy

$$\hbar\omega_p = \hbar\omega_s + \hbar\omega_i, \tag{A.3}$$

and one for the phase

$$\phi_p(\omega_p) = \phi_s(\omega_s) + \phi_i(\omega_i). \tag{A.4}$$

These are the only two equations we need for this explanation.

For this example, we will assume that we're running an OPCPA process with two input pulses. The first is a monochromatic pump pulse at 1030nm. The second is a broadband signal pulse, centered at 1525nm. This explanation is easier to follow if the red and blue sides of the signal pulse can be visually distinguished, so we'll assume that the signal spectrum is asymmetric. These are plotted as a function of frequency in Fig. A.1. With these two inputs, the pump

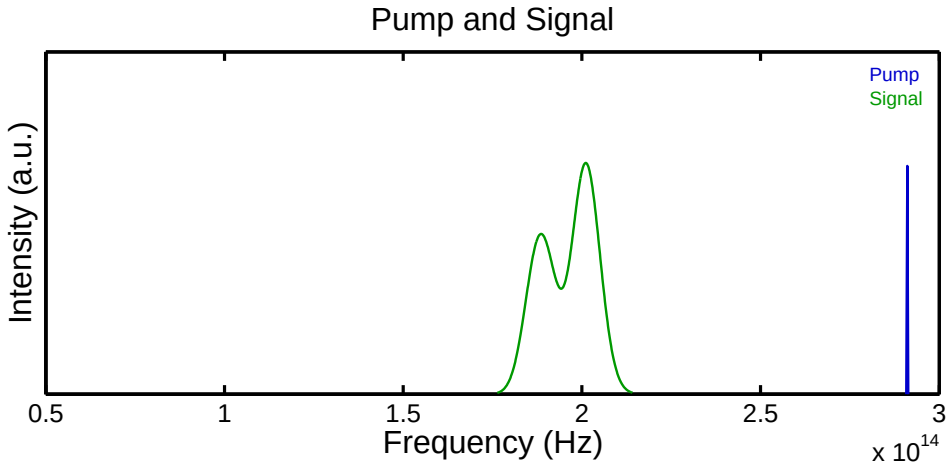

Figure A.1: A 1030nm monochromatic pump and the broadband 1525nm signal plotted as a function of frequency.

will amplify the signal, and also produce a third pulse, the idler. In this case need to make one further simplifying assumption. Namely, we must assume the nonlinear conversion efficiency of the OPCPA process is uniform across the signal spectrum. This is never completely true, but it is a reasonable simplification in cases where the phase matching bandwidth is greater than the input signal bandwidth[1] .

[1] Even absent phase matching considerations, there is always some spectral variation of the conversion efficiency

With these starting assumptions, the next step is to understand how the shape of the signal spectrum influences the shape of the idler spectrum. We can already state that the functional shape of the idler must match that of the signal. This is directly due to our assumption that the nonlinear conversion efficiency is uniform across the signal spectrum. However, there is a small wrinkle. For this, we return to Eq. A.3. The first observation is that at the photon level, the signal and idler frequencies must sum to the pump frequency. A consequence of this is that if we make the frequency of either the signal or the idler larger, the frequency of the other must get correspondingly smaller. On the macroscopic picture, this means that highest (lowest) frequencies in the signal are responsible for generating the lowest (highest) frequencies in the idler. In other words, the blue side of the signal generates the red side of the idler, and vise versa. The net result is that the spectrum of the idler is identical to that of the signal, but flipped horizontally about the idlers central frequency. This is shown in Fig.A.2.

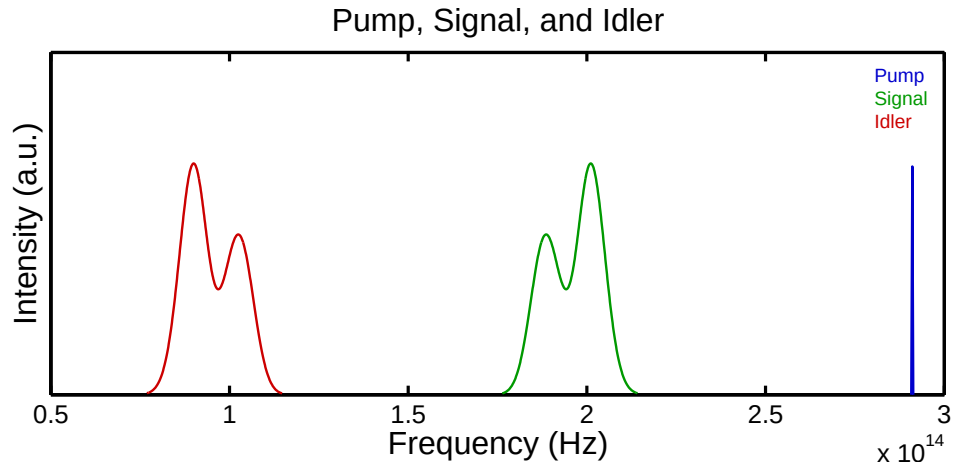

Figure A.2: The 1030nm pump, the 1525nm signal, and the generated idler. Notice the idler is a mirror image of the signal.

Now that we know how the OPCPA process affects the functional shape of the idler, the next step is to see how the process affects the phase. The previous analysis for the shape of the spectrum also holds for the phase between the signal and the idler, so we know that the idler spectral phase

in a parametric process due to the the frequency dependence of the the nonlinear coupling constant κ. This will also be neglected for this section.

will likewise maintain the same shape as the signals spectral phase, only flipped horizontally about the central frequency. But, going back to Eq.A.4, we see that corresponding signal and idler photons must have the opposite phase, assuming the pump phase is zero[2] . Therefore, the idler phase will also be flipped vertically about the horizontal axis. This can be seen in Fig. A.3, which plots the spectral intensity and spectral phase for the signal and the idler.

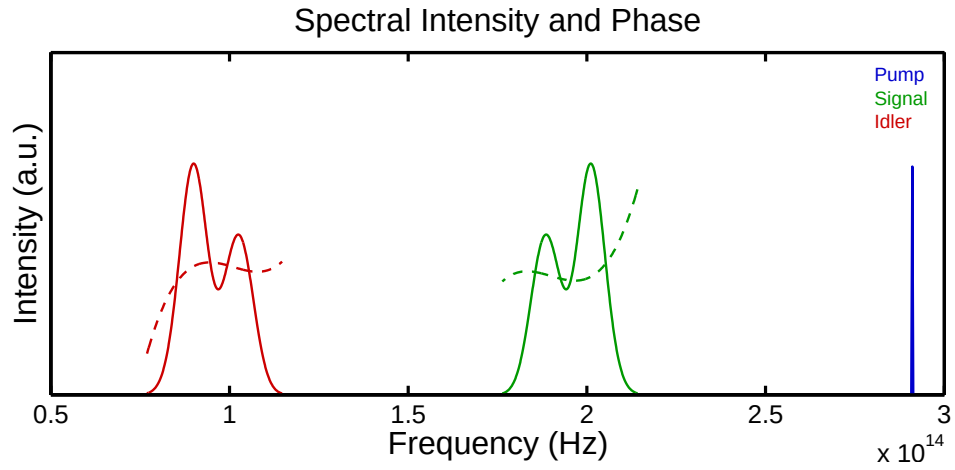

Figure A.3: The pump, signal, and idler spectral intensity and phase as a function of frequency. Notice that the signal and idler phases are inverted mirror images of one another. That is, the idler phase is obtained from the signal phase by flipping it horizontally and vertically. Under this transformation odd functions (such as odd polynomials) retain their sign, while even functions (the even polynomials) flip sign.

Now, with this in mind, we once again consider the Taylor expansions of the signal and idler phases,

$$\phi(\omega) = \phi_0 + \frac{d\phi}{d\omega}(\omega - \omega_0) + \frac{1}{2}\frac{d^2\phi}{d\omega^2}(\omega - \omega_0)^2 + \frac{1}{6}\frac{d^3\phi}{d\omega^3}(\omega - \omega_0)^3 + \cdot. \tag{A.5}$$

The previous analysis for the spectral phase of the signal and the idler holds for each term in the Taylor expansions as well. That is, each term in the Taylor expansion of the idler is identical to the corresponding term in the Taylor expansion of the signal phase, except flipped horizontally about the central frequency, and vertically about the horizontal axis. Now, as the Taylor expansion is

[2] Remember here that absolute spectral phase has no meaning, and that a linear phase corresponds to a temporal shift. In an OPCPA process where all the pulses are overlapped in time, the phase of any pulse is only determined up to the addition of a linear term.

simply a polynomial expansion, it is easy to see why the even terms have opposite signs between the signal and the idler, while the odd terms have the same sign. Even functions, when flipped horizontally and vertically, flip sign, while odd functions do not.

This gives a straightforward way to understand the phase relation between the signal and idler. The fact that odd dispersion terms retain the same sign, while even terms flip sign, is simply a result of the fact that the signal and idler phases are inverted, mirror images of one another. As odd functions keep the same sign under this transformation, all the odd terms in the Taylor expansion likewise keep the same sign. And, for even functions, which flip signs under this transformation, the even terms in the expansion likewise flip sign.

Appendix B

Fourth Stage Crystal Mount

One of the biggest engineering challenges in the development of the OPCPA laser was the development of the fourth cryogenic Yb:YAG stage. Developing a successful crystal mount required both analytical and numerical modeling, in addition to testing numerous different designs. In this Appendix I will give an overview of the problems encountered, and how they were solved.

B.1 Stress Induced Birefringence

Stress induced birefringence is a commonly encountered problem in high energy, cryogenic solid state lasers, and overcoming this was the primary engineering challenge in the fourth Yb:YAG stage. Stress induced birefringence reduces the output power via depolarization loss, which is measured as shown in Fig. B.1. In the simple setup, light enters through a polarizing beamsplitter, which ensures the light is linearly polarized. The light passes through the crystal, which causes a polarization shift because of stress birefringence. Following the crystal an output polarizer separates the light into its component polarizations parallel and perpendicular to the input polarization. Measuring the relative power between the two polarizations, $P_{\perp}$ and $P_{\parallel}$, gives the magnitude of the depolarization loss,

$$D = \frac{P_{\perp}}{P_{\perp} + P_{\parallel}}. \tag{B.1}$$

In the case of our cylindrical crystals, the measured depolarization loss can be expressed as a function of the input laser intensity profile $I(r, \phi)$, and a spatially dependent depolarization function

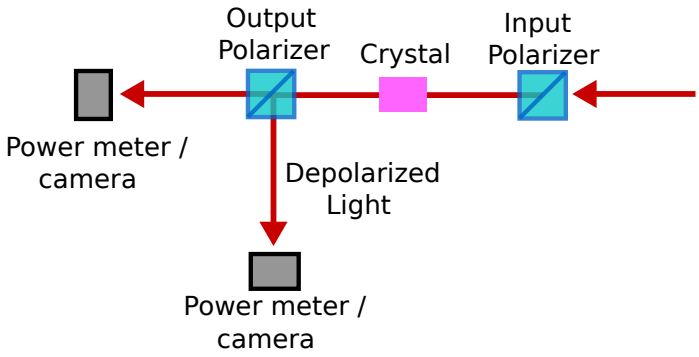

Figure B.1: The layout for measuring the depolarization loss due to stress birefringence. Linearly polarized input light enters through an input polarizer, then passes through the crystal. Stress birefringence causes a polarization shift in the light, which is then separated into its components parallel and perpendicular to the input polarization, which can be studied separately.

$f(r,\phi)$, [38],

$$D = \frac{\int_0^{2\pi}\int_0^{r_0} I(r,\phi)f(r,\phi)rdrd\phi}{\int_0^{2\pi}\int_0^{r_0} I(r,\phi)rdrd\phi}, \tag{B.2}$$

where r_0 is the crystal radius. As the stress is generally nonuniform throughout the crystal, the depolarization function $f(r,\phi)$ is likewise nonuniform across the aperture. In this case a camera can be used to analyze the spatially dependent depolarization.

Assuming there are no defects causing intrinsic stress within the the crystal, there are in general two sources of stress, and thus two sources of birefringence. The first is the crystal mount. Oftentimes, the crystal mount is clamped tightly around the crystal. For cylindrical crystal rods used in this system, this results in an inward radial pressure around the barrel of the crystal. In the most general case, the pressure on the barrel of the crystal is not uniform, and varies both azimuthally around the crystal and along the length of the crystal. This pressure profile depends on the mount geometry, and can not be calculated exactly. The pressure may be estimated either with simplified analytic models, or numerically with finite element modeling, but these cannot account for all the potential issues that can cause pressure, such as imperfections in the machining or assembly.

An additional factor that must be considered in cryogenic systems is that the crystals are typically mounted at room temperature, but operated at cryogenic temperatures. When cooled to cryogenic temperatures, both the mount and the crystal contract according to

$$\Delta L = L_0 \int_{T_0}^{T_f} \alpha_L(T) \mathrm{d}T, \tag{B.3}$$

where T_0 and T_f are the initial and final temperatures, $\alpha_L(T)$ is the temperature dependent linear expansion coefficient (often called the coefficient of thermal expansion, or just CTE), and L_0 is the linear length at the initial temperature. As the crystal and mount are usually made of different materials, they do not in general contract the same amount. If the mount contracts further than the crystal when cooled, then the stress on the crystal will likewise increase when cooled. Unfortunately, this often is the case with cryogenically cooled lasers, as metals usually have significantly higher coefficients of thermal expansion than crystalline materials. Copper, for example has a CTE of 16 ppm at room temperature, whereas YAG's CTE is only 6.2 ppm at room temperature. A useful metric here is the integrated linear expansion from room temperature to cryogenic temperatures,

$$\frac{\Delta L}{L_{293}} = \int_{293}^{80} \alpha_L(T) \mathrm{d}T, \tag{B.4}$$

which gives the total fractional change in the material length when cooled form 293K to 80K. Copper's total contraction by this metric is approximately 3100ppm, while YAG's is 917ppm [85, 3].

The second source of stress in the crystal is a thermal gradient within the crystal due to nonuniform pumping and cooling. To minimize the total stress induced birefringence requires minimizing the contribution of each of these factors. These factors will be covered individually next, beginning with stress from the crystal mount.

B.2 Low Stress Crystal Mounting I: Mounting Requirements

Before going into the details of the performance of the crystal mount, I will give a brief overview of the performance requirements for the crystal mount. The primary job of the crystal mount is to transfer the heat deposited by the pump from the crystal to the cryogenic cold head.

Several factors influence how well the mount can do this. The first is the thermal contact between the mount and the crystal. Simply put, the crystal and mount need to be in good thermal contact to get efficient heat flow from the crystal to the mount. The other factors are related to how well the mount itself conducts heat, as seen in Fourier's law

$$\mathbf{q} = -k\nabla T, \tag{B.5}$$

where $\mathbf{q}$ is the heat flux in W/m^2, k is thermal conductivity, and T is the temperature field. This gives three factors that affect the ability to transport heat: the thermal conductivity, the linear distance the heat must be transported over, and the cross sectional area of the mount. A good mount will have a large cross sectional area, keep the conduction length as short as possible, and be made of a high thermal conductivity material.

B.3 Low Stress Crystal Mounting II: Initial Mounting Attempts

The initial crystal mount, shown in Fig. B.2, was a "sleeved-clamshell" type mount. In this

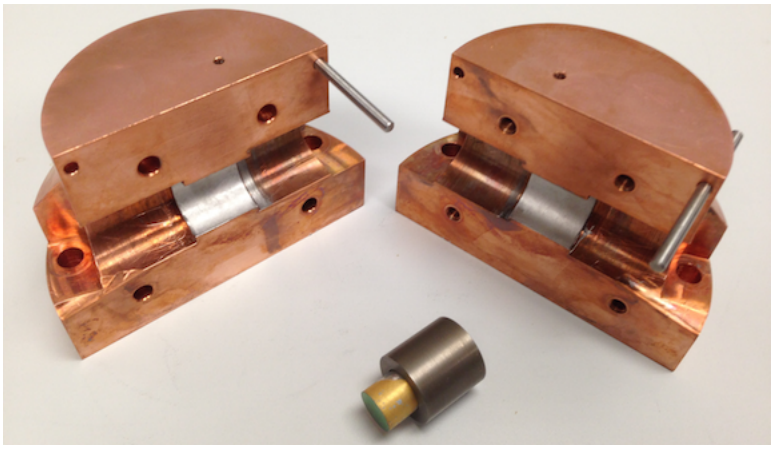

Figure B.2: The original crystal mount. Shown are the two halves of the copper mount, and a sleeve with a gold coated crystal partially inserted.

type of mount, the crystal itself is first soldered into a hollow OFHC copper cylindrical "sleeve". The inner diameter of the sleeve is 12.1mm, leaving 50μm of clearance on the radius between the sleeve and the crystal. The crystal is soldered into the sleeve with pure indium solder. Next, the

crystal-sleeve assembly is soldered into an OFHC copper mount which consists of two halves which bolt together to clamp down around the sleeve. The sleeve is 3mm thick, giving an outer diameter of 18.1mm, and the bore of the mount has an inner diameter of 18.2mm. This likewise gives 50μm of clearance on the radius between the sleeve outer diameter and the mount inner diameter.

The sleeve assembly is soldered into the mount using Field's Metal. As indium melts at 156°C, and Field's Metal melts at 62°C, the use of two different solders allows the sleeve assembly to be soldered into the mount without melting the indium layer between the crystal and sleeve.

The purpose of the sleeve in this mount is to help decrease stress on the crystal in two manners. First, by placing the crystal in a cylindrical sleeve, we can mitigate the effect of any imperfections in the machining or assembly of the mount. In mounts without a sleeve, the crystal sits in a bore formed by the two halves of the mount. Any imperfection in how the mount is made, or fits together, will result in the crystal sitting a nonuniform bore, which can create stress points on the crystal. By putting the crystal in a sleeve these can be mitigated. Second, the sleeve is designed to reduce the stress on the crystal from the CTE mismatch of the mount material and the crystal.

While this type of mount had worked well in previous systems, we found that it did not work well in this system. Specifically, it was found that there was significant depolarization even at room temperature, with approximately 8.3% total depolarized power. As expected, the depolarization grew when the cooled to cryogenic temperatures, reaching 12.6% at 80K. Ideally, the single pass depolarization loss should be less than 0.5%.

When the initial mount had unacceptable levels of depolarization loss, we tried several modified designs to better understand the source of the depolarization. The first variable we investigated was the crystal-sleeve clearance, to see if 50μm of indium was simply not a large enough layer. To this end, the crystal was mounted and tested with sleeves 100, 150, and 200μm of clearance on the radius around the crystal. These sleeves were also made from OFHC copper, and had the same outer diameter and length as the initial sleeves - only the inner diameter was changed.

These tests provided gave unexpected results. While it was expected that a larger indium

layer would reduce the depolarization loss, we found that in general the opposite was usually true, and that lower clearances typically performed better than larger clearances, especially at cryogenic temperatures. For example, one test of the 200μm clearance sleeve produced greater than 70% depolarization loss when cooled to cryogenic temperatures.

However, the main result from these tests was that they were not repeatable. That is, the same crystal, sleeve, and mount could be assembled following the same procedure multiple times, and each test would give different amounts of depolarization loss. Further investigation showed that the amount depolarization was highly sensitive to the assembly process. For example, slightly tightening or loosening the bolts holding the mount together could significantly change the depolarization at room temperature. From this we concluded that a two piece, clamshell type mount would not work, as it required tighter assembly tolerances than could be reliably delivered in a laboratory setting.

Following this, a few other mounting configurations were tried. These included a single piece copper mount designed to hold the crystal in a 1mm thick sleeve, and the use of another sleeve material, specifically copper-tungsten (CuW), which was used on the advice of a collaborator. CuW, a sintered composite material, is known to have a close CTE match to Yb:YAG at room temperature. However, we still found that none of these produced acceptable results, with large levels of depolarization loss regardless of which mount or sleeve material was used.

With these results in mind, we decided to undertake a more serious study of the crystal mounting, with the goal of designing a new mount that would allow for repeatable, low stress cryogenic operation. There were several specific questions we hoped to answer with the study: 1) Does the sleeve actually reduce stress on the crystal? We have some previous data that indicates it does, but our recent work casts doubt on this. 2) What are the optimum mount materials, taking into account thermal performance, crystal stress, price, and machinability.

In addition to these questions, we imposed several design requirements on the new mount: 1) It must be a single piece, rather than a clamshell. This requirements was made to simplify the mounting procedure, and hopefully allow for reliable and repeatable mountings. 2) The total

mount must weigh less than 20 pounds - the load limit of our cryocooler.

The next section will cover these models and present the results.

B.4 Low Stress Crystal Mounting III: Detailed Models

This study worked in several steps. The first was identify materials of interest. Then, a simple analytical model was developed with which we could quickly evaluate numerous different designs. Following this, a specific mount was designed and further analyzed via finite element modeling. Finally, with the results of the finite element modeling, a new mount was designed and built. Each of these steps will be discussed individually summarize

B.4.1 Mounting Materials

The first step in our detailed study was to narrow the field of potential mounting materials to a few highly promising materials worth future consideration. As all of our previous crystal mounts had used copper, it was included both as a potential material, and a baseline by which to judge other materials.

Materials were evaluated against the following criteria:

- Price. The cost of the mounting material must be reasonable.
- Machinability. We would prefer to be able to machine the mount at JILA in a conventional machine shop. However, materials which require specialized equipment are still an option providing there are vendors who can machine them, and the machining is not cost prohibitive.
- Toxicity. This is mainly a consideration for materials which we would machine ourselves.
- Thermal conductivity at cryogenic temperatures. For reference, at 77K copper's thermal conductivity is $\sim$ W/mK.

- Integrated thermal expansion from 293K to 80K. The closer the match to Yb:YAG, the better.

A variety of different types of materials were considered, including metals, composites / alloys, ceramics, and crystalline materials. However, amongst all the materials investigated, only three were found to worth future consideration. These were copper, copper-tungsten composite, and mono-crystalline silicon.

Copper was included because of its excellent thermal conductivity, reasonable cost and ease of machinability. The biggest drawback is that we already know it has a large CTE mismatch with Yb:YAG that induces stress birefringence.

Copper-Tungsten (CuW) is a sintered composite of copper and tungsten, typically 80-90% tungsten. Though cryogenic data is not available, this has the advantage of a reasonably high thermal conductivity and close CTE match with YAG at room temperature. It is expensive and more difficult to machine, but not prohibitively so.

Mono-crystalline silicon was included because it has a very high thermal conductivity at cryogenic temperatures, greater than 1000W/mK [46]. While machining silicon requires specialized equipment not found in typical machine shops, machined piece can be procured at low cost. The main drawback is silicon's CTE is less than YAG's, so a YAG crystal inside a silicon mount would pull away from the mount when cooled to cry temperatures. However, we decided to look for a way around this problem in hopes of taking advantage of the high thermal conductivity. It should also be noted that mono-crystalline germanium has a high cryogenic thermal conductivity, and is a very close CTE match to YAG. However, at present time it is cost prohibitive, both for the material and for machining.

With these materials in mind, we next move to the simple analytic model for evaluating general mounting geometries.

B.4.2 Analytical Model I: Model Description

The simplified analytical model is a two dimensional model that assumes the mounting geometry consists of a series of concentric cylinders. The stresses in the cylinders must satisfy Lamé's equation for thick walled cylinders[1] . A diagram of this model is shown in Fig. B.3

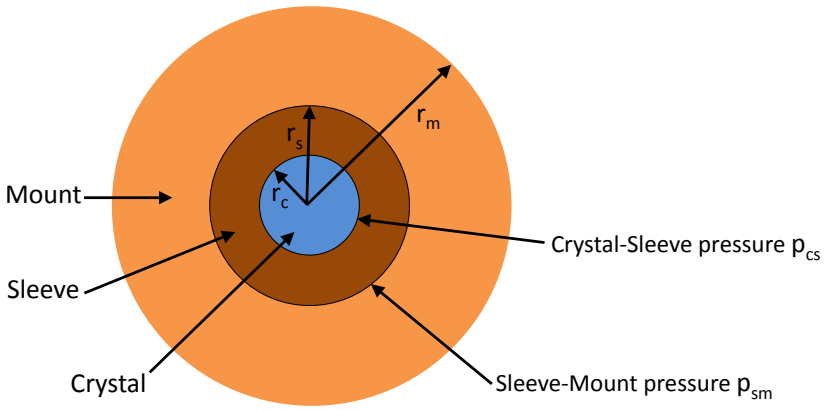

Figure B.3: The geometry of the analytical model used to estimate the stress on the crystal. The model consists of a crystal of radius r_c, a sleeve of radius r_s, and a mount of radius r_m. On the crystal-sleeve boundary is radial pressure p_{cs}, and on the sleeve-mount boundary is radial pressure p_{sm}. The model assumes that all layers are mounted such that $p_{cs} = p_{sm} = 0$ at room temperature.

The model consists of three distinct sections: the crystal, the sleeve, and the mount. The crystal is always assumed to be Yb:YAG, while the sleeve and mount can be either Cu, CuW, or Si. Note that using the same material for the mount and sleeve is the same as using a single piece mount.

Within the model, there are three radii of interest: the crystal radius r_c, the sleeve out radius r_s, and the mount out radius r_m. In addition, there are two pressures: the crystal-sleeve interface pressure p_{cs}, and the sleeve-mount interface pressure p_{sm}. The external pressure on the mount is assumed to be zero.

All three layers are assumed to fit together perfectly at room temperature, taken at 300K, so that $p_{cs} = p_{sm} = 0$. From here, the total thermal contraction of each layer from 300K to 80K

[1] Generally, thick walled cylinder theory is used for hollow cylinders where the thickness is greater than $1/20^{th}$ the cylinder radius.

is calculated for each layer individually. This gives an interference for each boundary: δ_{cs} for the crystal-sleeve boundary, and δ_{sm} for the sleeve-mount boundary. The crystal sleeve interference is the difference in the thermal contraction between the crystal outer radius and the sleeve inner radius, and similarly for the sleeve-mount interference. From here, we have two linear equations for the pressures and interferences, [2, 78]

$$\begin{pmatrix} \frac{r_c}{E_s}(\gamma_s + \nu_s) + \frac{r_c}{E_c}(\gamma_c - \nu_c) & -\frac{2r_c r_s^2}{E_s(r_s^2 - r_c^2)} \\ -\frac{2r_s r_c^2}{E_s(r_s^2 - r_c^2)} & \frac{r_s}{E_m}(\gamma_m + \nu_m) + \frac{r_s}{E_s}(\gamma_s - \nu_s) \end{pmatrix} \begin{pmatrix} p_{cs} \\ p_{sm} \end{pmatrix} = \begin{pmatrix} \delta_{cs} \\ \delta_{sm} \end{pmatrix} \tag{B.6}$$

where $E_{c,s,m}$ and $\nu_{c,s,m}$ are Young's modulus and the Poisson ratio for the crystal, sleeve, or mount as indicated, and

$$\gamma_c = 1, \quad \gamma_s = \frac{r_s^2 + r_c^2}{r_s^2 - r_c^2}, \quad \gamma_m = \frac{r_m^2 + r_s^2}{r_m^2 - r_s^2}.$$

The solution to this system of equations gives the pressures the interface boundary layers, p_{cs} and p_{sm}, with the goal of minimizing p_{cs}. As presented, a positive pressure for either of the boundaries indicates that the external layer is pressing radially inwards on the layer (and vice versa), whereas a negative pressure indicates that the external layer is pulling the inner layer radially outwards.

There are several limitations to this model that must be kept in mind, or else it may yield unphysical results. First, it requires the adjacent layers always remain in contact, when this is not necessarily the case. For example, if an interior layer shrinks more than the surrounding layer when cooled, it may pull away and separate, leaving a gap between the two layers. Whether or not this actually occurs will depend on whether there is an chemical bond between the layers (as may be, but is not necessarily, the case with solder layers), the strength of any bond, and the tensile yield strength of the materials. The model, however, assumes that the layers never separate.

The second limitation, is that Young's modulus and the Poisson ratio are taken to be independent of temperature, when in actuality they change with temperature. This is mainly due to the lack of published data for these values at cryogenic temperatures for materials of interest. Third, any indium or Field's metal solder or foil are neglected for these calculations, although it is a relatively straightforward extension to add these to the model. Appendix B gives presents the model with these layers added back in, and presents a justification for why they may be neglected.

Finally, the model assumes the layers always behave elastically, and never deform permanently. On these lines, the pressures the model gives should be checked against the compressive and tensile yield strengths of the materials, to verify that they do indeed remain in the elastic deformation regime.

B.4.3 Analytical Model II: Model Results

Five different configurations were tested in the analytical model, and the radial pressure on the crystal was recorded for each configuration. The configurations tested were: pure Cu, Cu mount + CuW sleeve, Cu mount + Si sleeve, pure CuW, and CuW + Si sleeve. In all configurations, a 6mm radius YAG crystal was used, and the mount outer radius was held constant at 40nm. For the sleeved configurations, the pressure is evaluated for sleeve thicknesses ranging from 1mm to 30mm (or, outer sleeve radii between 7 and 37mm).

As there is no temperature dependent CTE data available for CuW at cryogenic temperatures, the total thermal contraction of CuW when cooled from 300K to 80K is presently unknown. For this reason, in the calculations with CuW, we introduce an integrated CTE mismatch parameter for CuW. This parameter $\Delta L_{YAG,CuW}$, is defined as the difference between the total thermal contractions of CuW and YAG when cooled from 300 to 80 K, That is,

$$\Delta L_{YAG,CuW} = \int_{300}^{80} (\alpha_{YAG} - \alpha_{CuW}) \mathrm{d}T. \tag{B.7}$$

Each configuration with was evaluated six different values of $\Delta L_{YAG,CuW}$, from 0 to 500 ppm. These limits were used because a mismatch of 0 ppm represents the case where CuW and YAG are perfectly CTE matched for the entire range - the best case scenario. On the other hand, 500 ppm is the approximate value calculated if the CTE of CuW is assumed to equal constant over form 300K to 80K - a reasonable worst case scenario.

The model results are summarized as follows:

- Configuration 1: Pure Cu, no sleeve.

For a pure Cu mount with no sleeve, the model predicts an inward radial pressure of approximately 37MPa.

- Configuration 2: Cu mount, CuW sleeve.

  The pressure on the crystal as a function of sleeve thickness for several different CTE mismatches is shown in Fig. B.4.

  Perhaps surprisingly, the model predicts that in most situation a CuW sleeve actually increases the pressure on the crystal when compared to a pure Cu mount. The only exception is if $\Delta L_{YAG,CuW} = 0$ ppm, where a sleeve of any thickness reduces the pressure, or $\Delta L_{YAG,CuW} < 100$ ppm, where a sleeve can help, but must be over $\sim$10mm thick.

- Configuration 3: Cu mount, Si sleeve.

  Fig. B.5 shows the pressure on the crystal for a Cu mount and an Si sleeve as a function of sleeve thickness. It is seen that a Si sleeve reduces the pressure regardless of thickness, but must be quite thick to substantially reduce the pressure.

- Configuration 4: Pure CuW, no sleeve.

  In the case of a pure CuW mount, the inward radial pressure depends linearly on the CTE mismatch $\Delta L_{YAG,CuW}$. For the range of 0 to 500 ppm, the inward pressure ranges form 0MPa to just under 75MPa, as shown in Fig. B.6.

  It may be surprising that even a relatively small CTE mismatch gives such high pressures. After all, the total CTE mismatch between YAG and Cu is about 2100 ppm, and that configuration only gives 37MPa. This is because CuW has a much higher Young's modulus then Cu - 280GPa compared to 115GPa. From this we conclude that unless the total integrated CTE mismatch between YAG and CuW is less than $\sim$ 200ppm, CuW is not likely to perform better than Cu.

- Configuration 5: CuW mount, Si sleeve.

The results from a CuW mount and a Si sleeve are shown in Fig. B.7. As can be seen a Si sleeve of an pressure decreases the pressure on the crystal as compared to a pure CuW mount, though thicker sleeve do more to reduce the pressure. With a thick enough sleeve and $\Delta L_{YAG,CuW} < 200$ppm, the pressure may be reduced to zero.

From these results we decided to move forward with two different configurations: a pure CuW mount with no sleeve, and a CuW mount with a Si sleeve. While we did not know the exact integrated CTE mismatch between YAG and CuW, a couple factors convinced us that it was close enough to YAG that CuW would work for our purposes. The first was a simple experiment we did to measure the depolarization loss of an Yb:YAG crystal held only in a 4mm wall thickness CuW sleeve. In this test, we saw no increase in stress birefringence as the unit was cooled to cryogenic temperatures. The second factor were some conversations with a collaborator who had experience with CuW used in mounting cryogenic thin disk lasers. Together, these gave us confidence CuW was the most promising material.

With this in mind, we next moved to finite element analysis to study the thermal and mechanical properties of actual mount designs. This will be discussed next.

B.4.4 Finite Element Analysis I: Model

Following the results of analytical model, finite element analysis was used to evaluate two specific mount designs - the first a monolithic (single piece) CuW mount, the second a CuW mount that incorporates a large Si sleeve. The cross sectional design for these mounts is shown in Fig. B.8. As both mounts have the same general geometry, they can be described with the same model by simply using different materials for different sections of the mount. For single piece mounts, the dark grey and light material are both CuW. In the sleeved design, the dark grey is CuW, while the light grey is Si. In either case, the blue is the crystal - Yb:YAG.

Both of these mounts were evaluated in a finite element analysis in COMSOL. In each situation, the crystal a 12mm diameter, 20mm long 2% at. doped Yb:YAG crystal with 400W of incident CW pump power. The crystal is divided into 10 discrete sections through the length of the

crystal, each 2mm in length. Each section absorbs the incident pump power according to Beer's law.

From the total power deposited in the crystal, two heat sources are modeled. The first is heating due to the quantum defect. This is set to be equal to 9.3% of the total absorbed power for each section, and the heat is assumed to be deposited in the pump spot on the crystal. In the mount shown in Fig. B.8, the inner radial section of the crystal represent the pump area - and this is where quantum defect heat is deposited.

The second is heating due to trapped fluorescence. All of the absorbed power not lost through the quantum defect is assumed to be converted into fluorescence. The fluorescence is modeled as a being emitted from a line source along the crystal axis, that emits equally in directions. The fluorescence emission is modeled as being constant through the length of the crystal, though this is not actually the case due to the exponential dependence of the pump absorption.

The total heat deposited by the fluorescence is assumed to be equal to fraction of the fluorescence that strikes the crystal barrel. A simple geometrical estimate says that for an end pumped cylindrical crystal of length l and radius r, the fraction, f_t, of the fluorescence that hits the barrel of the crystal is equal to

$$f_t = 1 - \frac{l - c - \sqrt{l^2 + r^2}}{l}, \tag{B.8}$$

which for a 12mm diameter, 20mm long crystal is equal to 0.744. Therefore, of the 90.7% of the absorbed power not lost through the quantum defect, 74.4% is assumed to be converted to heat. This heat is modeled as being deposited equally along the barrel of the crystal, with no length or azimuthal dependence.

In addition to the heat sources, one cooling source is modeled. This is modeled as a cooling source that provides with a cooling capacity linearly dependent on temperature. The cooling capacity is assumed to be 0W at 50K, and 200W at 80K, approximating the 200W unit used on the fourth stage.

With these heating and cooling parameters, the temperature field throughout the mount and

crystal is calculated. The temperature distribution is used to calculate the stress due to thermal expansion.

Though we planned to use CuW for the mount material, we ran additional simulations with identical geometries using Cu for the mount material. This was done because we have much more experience with Cu crystal mounts, and we hoped the results from the Cu mount simulations would help provide context for our CuW mount simulations.

B.4.5 Finite Element Analysis II: Results

The finite element analysis gave both temperature and stress field for the mounts tested. The temperature distributions for the different mounts are shown in Fig. B.9 The temperature scale and range are the same for each plot, ranging from 90K to 150K. As can be seen, Cu mounts perform better than CuW mounts, and sleeved mounts perform better than mounts with no sleeve. For each configuration, the highest temperature on the crystal is recorded in the upper right. With no sleeve, this suggests that the crystal is approximately 6K warmer in a CuW mount as compared to a Cu mount. In the sleeved mounts, this difference drops to 2K.

The stress fields predicted by the model are shown in Fig. B.10. The plot range and scale are the same for each plot, ranging from 0MPa to 500MPa. The benefits of a CuW mount are clearly visible here, as the stress is drastically reduced. In particular, the modeled stress on the crystal from a Cu mount predicts a maximum stress of of 470MPa, while for a CuW mount it is less around 10MPa. With a Si sleeve, the pressure on the crystal is approximately 153MPa with a Cu mount, and close to 25MPa for a CuW mount. This is easier seen Fig. B.11, which shows the crystal in isolation, with a plot range from 0MPa to 200MPa for the sleeved configurations.

This model clearly predicts a reduced stress on the crystal with a CuW mount as compared to a Cu mount, agreeing with our analytic model. With these results, we made a CuW mount to use for our fourth stage. The results from this mount are summarized next.

B.5 CuW Crystal Mount

Based on our analytic and finite element models, we decided to construct a monolithic CuW mount for our fourth stage, believing the reduced stress would be more important than the decreased thermal performance. The mount design was identical to that used in our finite element models, except that the real mount was machined with a 100μm gap between the crystal and mount for Indium solder. This mount was tested in a manner similar to our Cu mounts, measuring the depolarization as a function of crystal temperature. This is shown in Fig. B.12, which plots the single pass depolarization loss as a function of temperature for the CuW mount, plus one of our best earlier configurations, a single piece Cu mount with a 1mm thick Cu sleeve. With the CuW mount, the final temperature depolarization loss was reduced by a factor of 35 when compared with the Cu mount, with a final single pass depolarization loss <0.2% - a figure approaching the extinction ratio for the polarizer of 1000:1.

This single piece CuW mount was a substantial improvement over the previous mounts, and allowed the fourth stage to run with output powers up to 36W. Pumping does induce some further depolarization by non-uniformly heating the crystal, but we have found that intracavity wave plates can mitigate this issue with relative ease [38, 24]. In particular, when pumped, the double pass depolarization loss increases to 1.7%, but by adjusting the waveplates we are able to reduce the double pass loss to 0.6%, an acceptable loss factor.

B.6 Conclusions

With the combination of the CuW mount and intracavity waveplates, we were able to successfully operate the fourth stage at the necessary powers and energies. The low stress CuW mount solved the issue of stress birefringence form the crystal mount, while the waveplates address depolarization loss caused by the nonuniform pumping and cooling. Together, they were able to reduce the measured double pass depolarization loss to <1%.

In addition to our CuW mount, we have also begun work on a CuW mount with a Si sleeve.

As the temperature plots in Fig. B.9 show, these are expected to greatly improve the thermal performance of the fourth stage amplifier. These mounts will be tested in the future, to bring the Yb:YAG energy close to 100mJ.

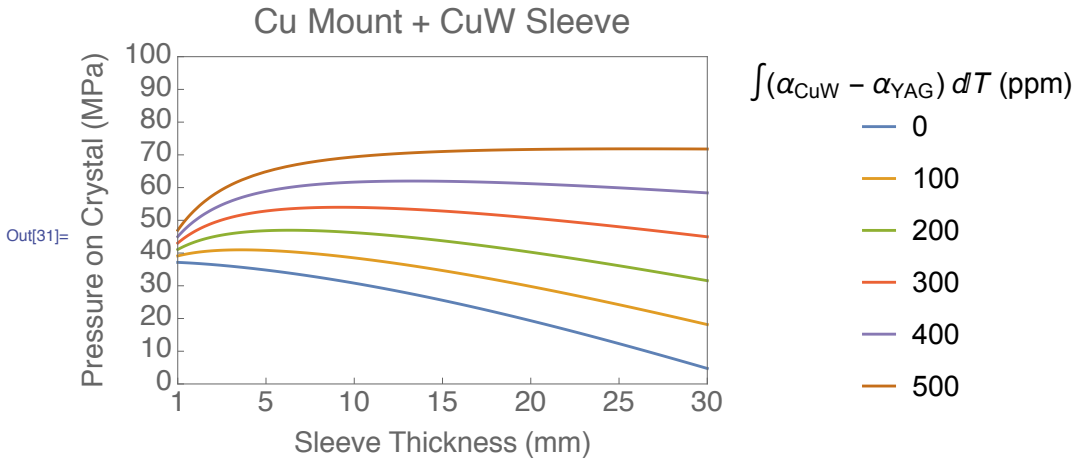

Figure B.4: The predicted pressure on the crystal in MPa for a Cu mount and a CuW sleeve as a function of sleeve thickness, for several different integrated CTE mismatch values between YAG and CuW.

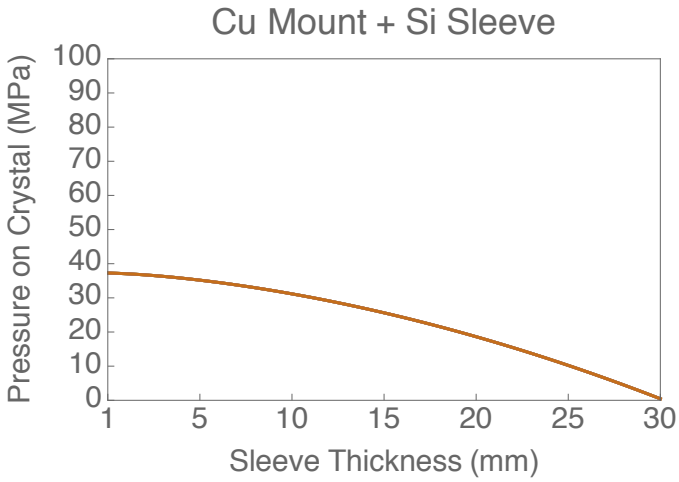

Figure B.5: The predicted pressure on the crystal in MPa for a Cu mount and a Si sleeve as a function of sleeve thickness.

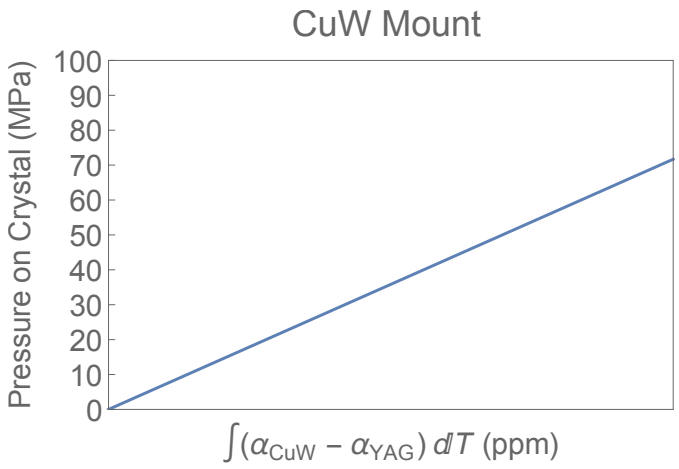

Figure B.6: The predicted pressure on the crystal in MPa for a CuW mount as a function of the integrated CTE mismatch between YAG and CuW, ranging from 0 to 500ppm.

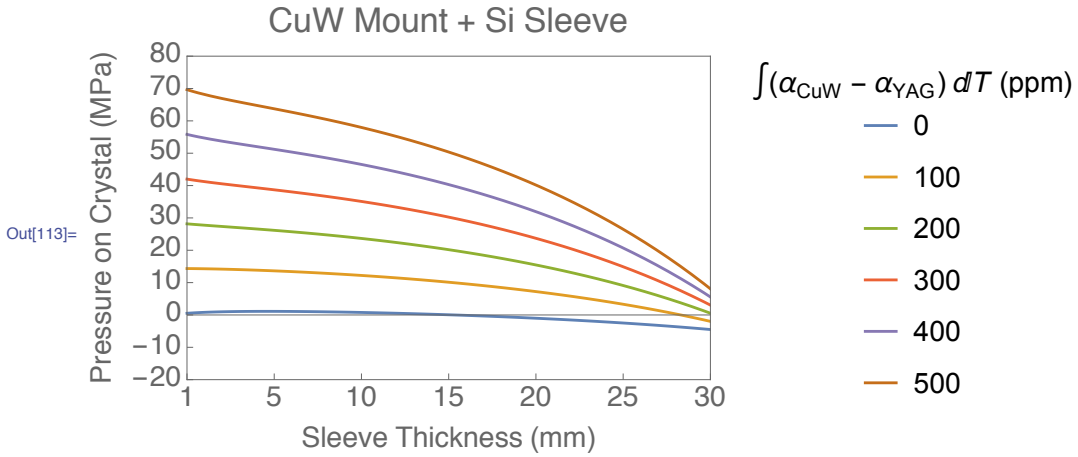

Figure B.7: The predicted pressure on the crystal in MPa for a CuW mount and a Si sleeve as a function of sleeve thickness, for several different integrated CTE mismatch values between YAG and CuW.

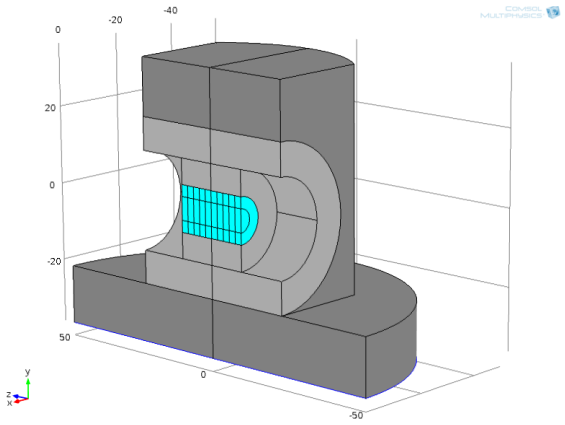

Figure B.8: A cross sectional schematic of the mount design used for finite element modeling. The dark grey material is the mount material - CuW for the planned design. In the case of a single piece mount, the light grey material is the same as the dark grey material. In the case of a sleeved mount, the light grey material is the sleeve. The blue material is the crystal, which has been divided into 10 sections along the crystal depth. The central radial section of the crystal represents the pump area, and this is where heat to due to the quantum defect is deposited

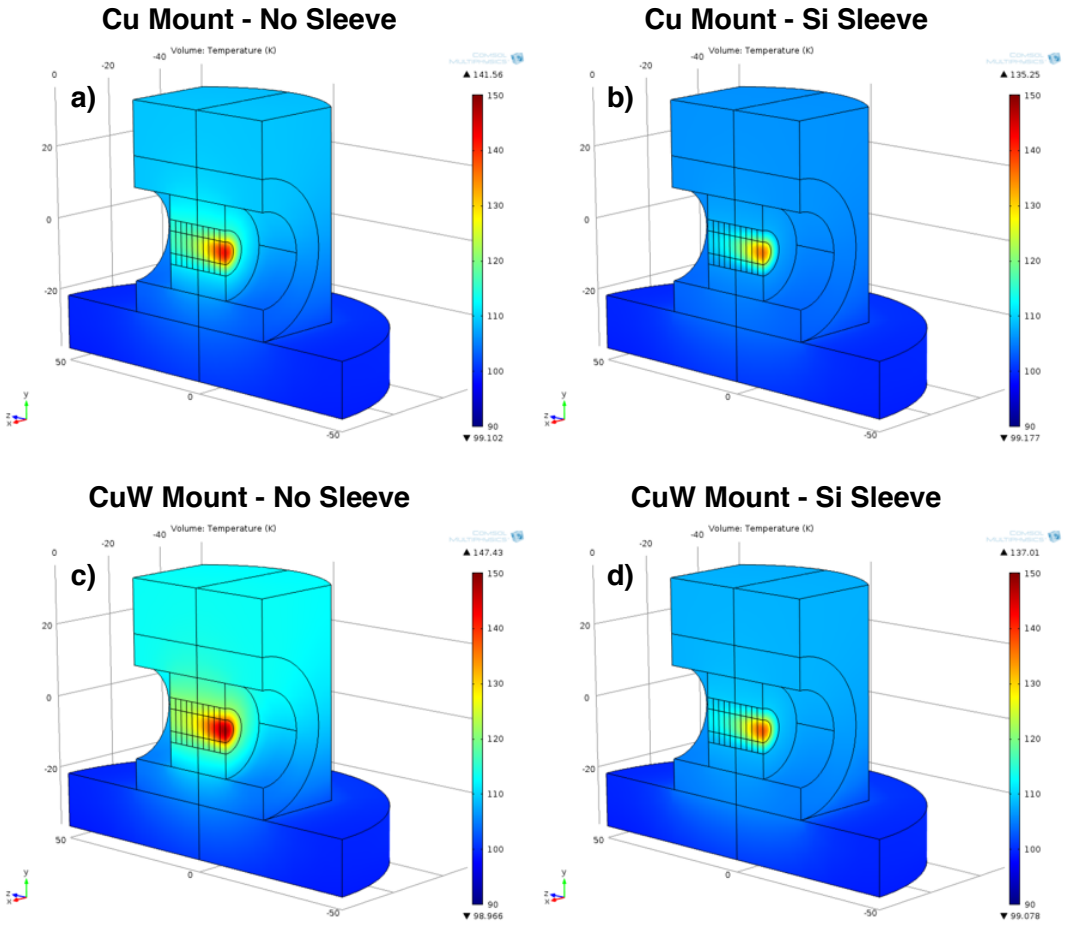

Figure B.9: The finite element analysis modeled temperature distributions for: a) a Cu mount with no sleeve, b) a Cu mount with a Si sleeve, c) a CuW mount with no sleeve, and d) a CuW mount with a Si sleeve.

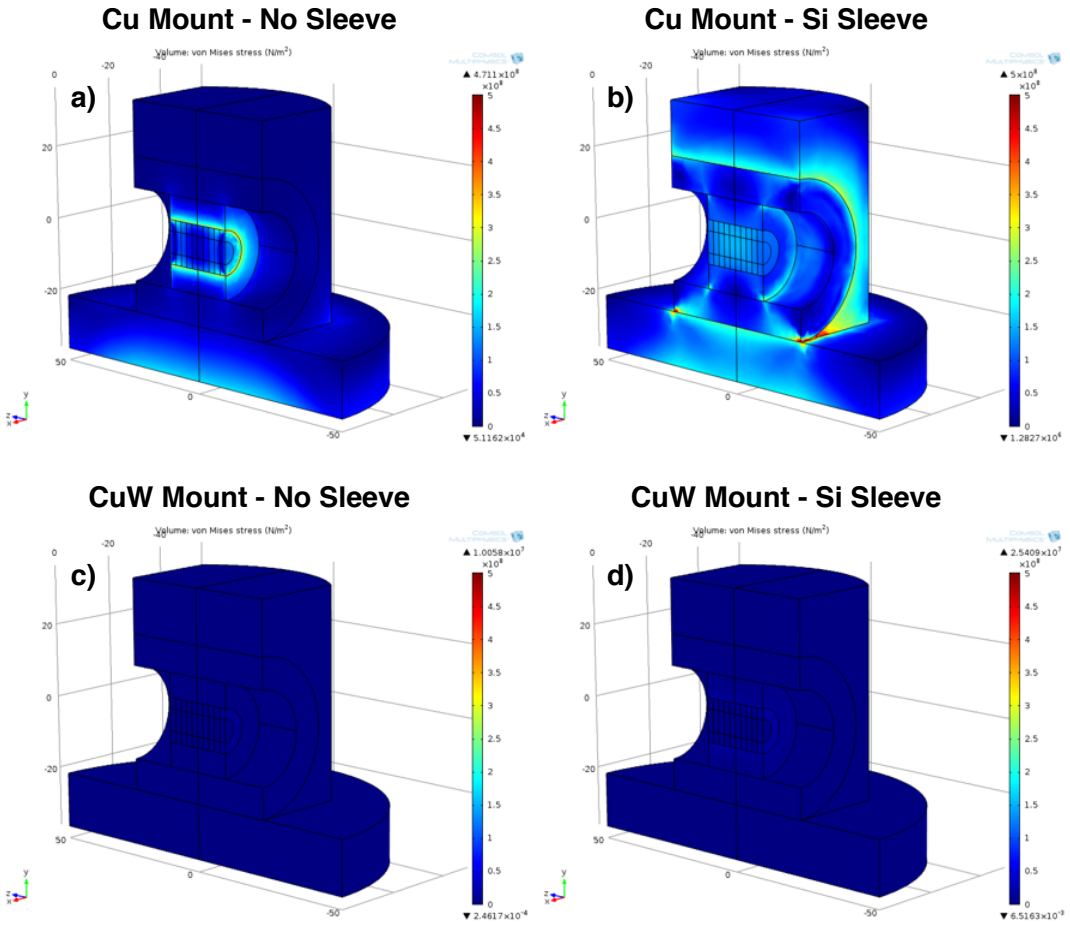

Figure B.10: The finite element analysis modeled von Mises stress in: a) a Cu mount with no sleeve, b) a Cu mount with a Si sleeve, c) a CuW mount with no sleeve, and d) a CuW mount with a Si sleeve.

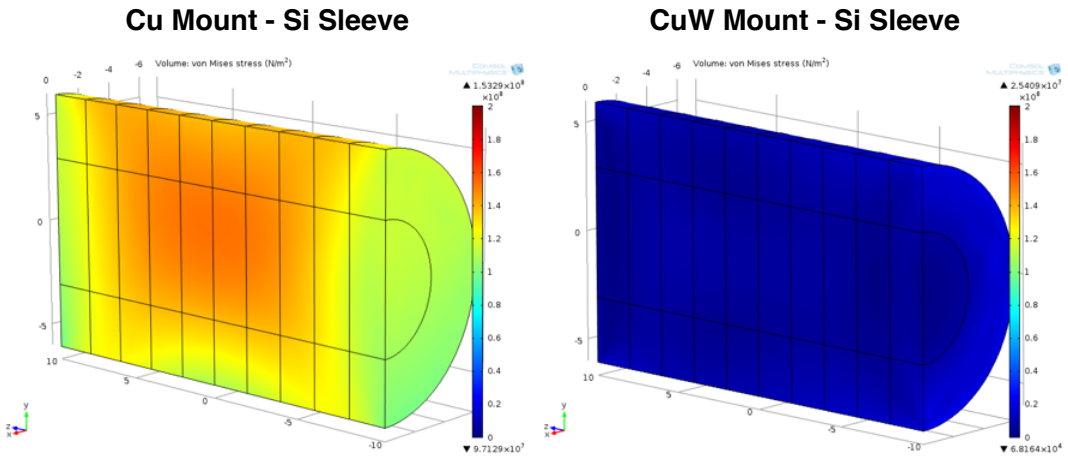

Figure B.11: The finite element analysis modeled von Mises stress on the crystal for the Cu and CuW mounts with the Si sleeve.

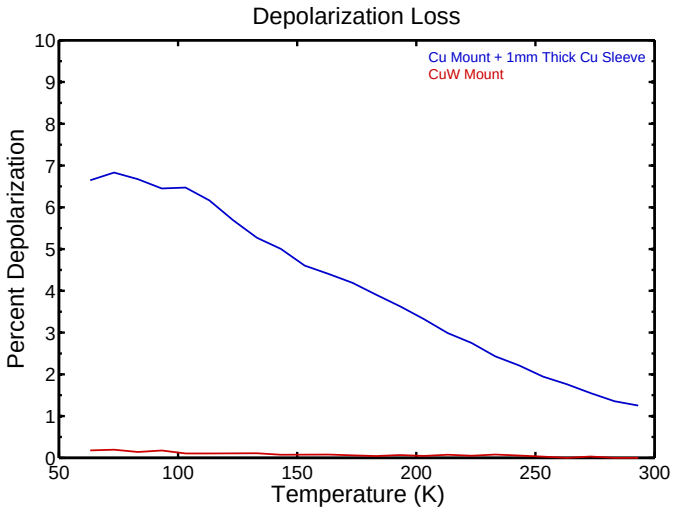

Figure B.12: The measured depolarization loss for the new CuW, and one of the previous best results from a copper mount with a 1mm thick copper sleeve. The CuW mount reduced the measured single pass depolarization loss to <0.2%, only a factor of 2 greater than the 1000:1 extinction ratio of the polarizer used in the measurement.