

**A 3 μm wavelength optical parametric chirped pulse
amplifier for soft X-ray generation**

by

Drew William Morrill

B. Sc., Brown University, 2013

M. Sc., University of Colorado, 2019

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

2024

Committee Members:

Margaret Murnane, Chair

Henry Kapteyn

Scott Diddams

Michael Litos

Shu-wei Huang

Morrill, Drew William (Ph.D., Physics)

A 3 μm wavelength optical parametric chirped pulse amplifier for soft X-ray generation

Thesis directed by Prof. Margaret Murnane

Coherent light sources in the soft X-ray region (SXR: 0.6 - 8 nm, 2 - 0.15 keV) are vital to a number of scientific and industrial applications. The short wavelength and well-defined absorption features throughout this spectral region impart SXR imaging and spectroscopy techniques with excellent elemental specificity and high spatial resolution - enabling applications ranging from bio-imaging in the water window (284-532 eV) to semiconductor metrology. At present, coherent SXR sources are generally limited to less-accessible large-scale facilities such as synchrotrons and XFELs. In this work, we present a stable, tabletop source of coherent SXR light that uses high-harmonic generation (HHG) to convert pulses of mid-infrared laser light to pulses of spatially and temporally coherent SXR light. We describe the ground-up development of a 1 kHz, 3 μm wavelength optical parametric chirped pulse amplification (OPCPA) system generating ~ 650 μJ , 100 fs near-transform limited pulses used to drive HHG. Optical parametric amplification (OPA) is done in four-stages, each utilizing a periodically-poled lithium niobate nonlinear crystal and pumped by a ~ 20 mJ single-stage, uncompressed cryogenically-cooled Yb:YAG regenerative amplifier. We use numerical simulations to inform the optimization of the OPA process. The highly stable, all-fiber front-end is based on a single twin-output erbium oscillator. The 1 μm seed for the regenerative amplifier is formed via dispersive wave generation in a highly nonlinear fiber, and the broadband, 35 fs 1.5 μm OPA seed is generated in a dispersion-managed erbium fiber amplifier. A modular high-harmonic source was developed with an emphasis on differential pumping and versatility. Using a multi-atmosphere gas cell, we generated high-harmonics in argon reaching the argon absorption edge at 248 eV and to just above the carbon K-edge (284 eV) using nitrogen gas. The source is applications oriented by design, with an emphasis on stability and robustness. We present an avenue towards the generation of harmonics in the keV range using helium gas.

Dedication

To my parents, Ken and Susan, and my brother, Ryan.

Acknowledgements

I would first like to express my sincere gratitude to my advisors, Margaret Murnane and Henry Kapteyn, who have been consistently supportive and available throughout my PhD. The work described in this thesis was a highly collaborative effort, largely from a team at JILA consisting of Michaël Hemmer, Daniel Carlson, Will Hettel, Jeremy Thurston and Grzegorz Golba. It has been a pleasure and an honor learning from and working with you all.

I am indebted to the many scientific mentors from around the world who have supported me over the course of my education and career - in particular: Alan Doud of Squawamish High School - the teacher who introduced me to physics; Lai-Sheng Wang, Li-Qiong Wang and Domenico Pacifici of Brown University; Thomas Elsaesser, Michael Woerner and Klaus Reimann of the Max Born Institute in Berlin; Ambarish Ghosh of the Indian Institute of Science in Bangalore; and Michael Messerly and Leily Kiani at Lawrence Livermore National Lab. Gunnar Arisholm wrote the Sisyfos code package and patiently mentored me in its use. Scott Diddams and his research group taught our team how to build the unique fiber lasers in the front-end. Thank you to the Krell Institute and the DOE NNSA Stewardship Science Graduate Fellowship (award number DE-NA0003960), which helped fund my graduate education.

The collective intelligence, creativity and compassion of Margaret and Henry's research group, both past and present, is inspiring. Thank you to everyone who has been so generous with scientific guidance and personal friendship. In particular: Bin Wang, Nathan Brooks and Sinéad Ryan - your friendships help define who I am; current group members Albert Beardo, Iona Binnie, Eric Fang, Anya Grafov, Nicholas Jenkins, Tika Kafle, Joshua Knobloch, Rae Larson, Brendan McBennett,

Emma Nelson and Oliver Shao; past group members Amitava Adak, Charlie Bevis, Seth Cousin, Kevin Dorney, Jennifer Ellis, Travis Frazer, Ben Galloway, Christian Gentry, Michael Gerrity, Guan Gui, Peter Johnsen, Chen-Ting Liao, Quynh Nguyen, Daisy Raymondson, Michael Tanksalvala, Phoebe Tengdin and Dmitriy Zusin.

Our work was enabled by the skill and hard work of the JILA shops and support staff. In particular, James Uhrich: your instrument making talent and creativity enabled the high-harmonic source, but your friendship enabled this PhD. Other members of the JILA community that contributed strongly to this work include Hans Green, Calvin Schwadron, Kyle Thatcher, Kim Hagen and Adam Ellzey of the instrument shop; James Fung-A-Fat and Terry Brown of the electronics shop; David Alchenberger and Curtis Beimborn of the Keck lab; J.R. Raith, Dan Packman, Cory Keasling and Jim McKown of computing; Brian Lynch and the entire supply office; and Scott Brown and George Kavadas for keeping the lab running. Tin Nguyen generously offered his time in helping assemble our fiber lasers before we had a splicer of our own.

To my partner, Hannah Pittel, thank you for your endless empathy and patience. Your view of life as a process of constant growth paints the future in a beautiful light. I am extremely grateful to have friends who understand me on the deepest level. In particular, Peter Beck, Austin Czarnecki, Emily Tansey, Nikhil Parasher, Anustuv Pal, Anke Trojan, Ken Sanga and Noah Schlossberger - you have all helped me untangle the path worth pursuing from the easiest one.

Finally, I dedicate this thesis to my parents, Ken and Susan Morrill, and my brother, Ryan Morrill. Mom and dad, your parenting has always been guided by love and realized by selflessness. In some cases, the lessons you taught me early on turned out to be the bedrock of physics as a discipline - the necessity of understanding a problem before you can solve it. In other cases, those lessons turned out to be the ones that guide me as a physicist - to always try for a positive impact on the people and world around us. Ryan, you blazed the trail that I followed to get here. Your deep and unconventional way of thinking will always be an inspiration to me.

Contents

Chapter	
1	Introduction: compact, coherent soft X-ray sources and the lasers that drive them 1
1.1	Short wavelength light sources 3
1.2	Introduction to high-harmonic generation 7
1.3	History of fiber lasers 13
1.4	High average power solid-state lasers 15
1.5	Chirped pulse amplification to generate high peak powers 18
1.5.1	How group delay dispersion affects pulse duration 20
1.5.2	Optical parametric chirped pulse amplification 21
1.6	The current state of long wavelength lasers to drive high-harmonic generation 24
1.6.1	A survey of long wavelength OPCPAs producing femtosecond pulses 24
1.6.2	Direct amplification of long wavelength light to drive HHG 28
1.7	An introduction to the 3 μm wavelength OPCPA 30
2	Theory of parametric nonlinear optics 33
2.1	Electromagnetic wave equation 35
2.2	Spectral phase of ultrafast pulses 36
2.2.1	Absolute phase 40
2.2.2	Group delay 40
2.2.3	Group delay dispersion 42

2.2.4	Third and higher order dispersion	43
2.3	Extending the wave equation to describe nonlinear optics	44
2.4	The physical origin of nonlinear susceptibility	46
2.4.1	Lorentz model	46
2.4.2	The influence of centrosymmetry on nonlinear susceptibility	48
2.4.3	Wave mixing	50
2.5	Three-wave mixing processes ($\chi^{(2)}$)	51
2.5.1	Second-harmonic generation	53
2.5.2	Difference-frequency generation	56
2.5.3	Monochromatic optical parametric amplification	58
2.5.4	Why phase-matching bandwidth is not the same as amplification bandwidth in OPA	60
2.5.5	Extending the OPA coupled wave equations to the time-domain: the Sisyfos numerical simulation package	63
2.5.6	Quasi-phase matching in OPA	65
2.6	Four-wave mixing processes ($\chi^{(3)}$)	68
3	All-fiber front-end laser to generate 1 and 1.5 μm seed pulses	73
3.1	Introduction	73
3.1.1	Overview of the all-fiber front-end	74
3.1.2	Optical fibers in laser amplifiers	77
3.2	Erbium fiber oscillator	80
3.3	1.5 μm arm: Dispersion managed, broadband erbium-doped fiber amplifier	83
3.3.1	Background and theory	83
3.3.2	Design and construction	84
3.3.3	Pulse characterization	87
3.4	1.03 μm arm: dispersive-wave shifted seed for the regenerative amplifier	88

3.4.1	Dispersive wave generation in a highly nonlinear fiber	90
3.4.2	Ytterbium-doped fiber amplifier chain	95
4	A 3 μm, PPLN-based OPCPA pumped by a cryogenically-cooled Yb:YAG regenerative amplifier	101
4.1	Introduction	101
4.2	A review of PPLN OPCPA technology in the 2 - 3.5 micron range	104
4.3	Cryogenically cooled Yb:YAG regenerative amplifier	106
4.4	Dispersion control in the OPCPA	115
4.4.1	Programmable SLM-based pulse shaper	116
4.4.2	Grism stretcher	120
4.4.3	Grating pair compressor	122
4.5	Four stage PPLN optical parametric amplifier	126
4.5.1	Overview	126
4.5.2	Stage by stage performance	129
4.5.3	Characterization of the compressed 3 μm pulses	134
4.6	Distributed parametric amplification for higher pulse energies: numerical simulations of the strong seed regime, and the role of crystal length in conversion efficiency . . .	136
4.6.1	Numerical simulation of conversion efficiency and crystal length in OPA4 . .	140
4.7	Stability of the OPA chain	141
4.7.1	Numerical simulations of OPA stability: output energy sensitivity to pump and seed fluctuations	144
5	3 μm driven harmonic generation	147
5.1	Introduction	147
5.2	Theory of high-harmonic generation	147
5.2.1	The microscopic, single atom picture of HHG	147
5.2.2	The macroscopic picture of HHG: phase matching and efficiency	151

5.3	A versatile, modular HHG source	164
5.3.1	Introduction: context and design criteria	164
5.3.2	Design and fabrication of the source	166
5.3.3	Gas handling and differential pumping characterization	171
5.4	HHG results	175
5.4.1	Beamline	176
5.4.2	Modes and parameter scans	178
5.4.3	Spectra	181
5.5	Below-threshold harmonics	186
5.5.1	Introduction	186
5.5.2	Experiment	188
5.5.3	Perturbative effects in below-threshold harmonics	192
5.5.4	Below-threshold harmonic plateau: verifying intensity and interaction length	194
6	Concluding remarks	197
	Bibliography	200
	Appendix	
A	Numerical simulations of alternative $\chi^{(2)}$ crystals for 1 μm-pumped OPCPA	220
A.1	Introduction	220
A.1.1	Damage and pulse duration scaling	221
A.1.2	Summary of results	223
A.2	Comparative numerical simulations of periodically-poled crystals: PPLN, PPLT and PPKTP	225
A.3	Comparative numerical simulations of bulk crystals: KTA, LN and KN	233

B Two-photon beam profiling at telecom wavelengths with a standard silicon sensor	246
C Analytical studies of OPA gain scaling with crystal length, pump intensity and effective nonlinearity	251
D Technical drawings and user notes for the HHG source assembly	255

Tables

Table

3.1	A summary of the output parameters for each arm of the front-end.	75
3.2	A summary of key parameters of the HNLF used to generate the 1 μm dispersive wave.	92
4.1	Sign of the GDD and TOD for a grating pair compressor, Martinez-style grating stretcher and grism.	115
4.2	Simulation parameters used for strong seed simulations.	137
4.3	Simulation and experimental data on OPA4 conversion efficiency.	140
5.1	Calculated refractive indices and critical ionization levels at a driving wavelength of 3 μm	156
A.1	Simulation parameters used for comparison of PPLN, PPSLT and PPKTP as candidates for the final OPA stage.	226
A.2	Crystal parameters used in sisyfos simulations of periodically poled crystals.	226
A.3	Simulation parameters used for comparison of KTA, LN and KN as candidates for the final OPA stage.	235
A.4	Crystal parameters used in sisyfos simulations of bulk crystals.	236

Figures

Figure

1.1	The soft X-ray spectrum spans a number of strong, sharp elemental absorption lines relevant for biological imaging and materials spectroscopy.	5
1.2	Images of the fission yeast <i>Schizosaccharomyces pombe</i> using conventional fluorescence microscopy, differential interference contrast (DIC) light microscopy, and soft X-ray tomography.	6
1.3	Results from three early papers demonstrating high-harmonic generation.	8
1.4	The three-step model of high-harmonic generation.	10
1.6	An Alcatel Submarine Networks optical repeater being deployed in the ocean.	14
1.7	Solid-state laser gain media can be broadly classified by to their dimensionality, which significantly affects thermal management and average-power handling.	16
1.8	Single fiber chirped pulse amplification systems have amplified multiple GW of compressed peak power, with the potential for scaling via the combination of many fibers - potentially enabling both high peak and average power systems.	17
1.9	A depiction of chirped pulse amplification as it used in OPCPA.	19
1.10	How the application of group delay dispersion (GDD) from the stretcher influences pulse duration in the OPCPA.	21
1.11	Spatial mode distortion of a ~ 15 mJ, 30 fs (unchirped) OPA pump pulse due to severe nonlinearities.	23

1.12	A high harmonic spectrum extending to 1.6 keV, and the 3.9 μm wavelength OPCPA which drove the process.	25
1.13	A survey of few-cycle OPCPA systems.	26
1.14	A schematic overview of the entire laser system described in this thesis.	31
2.1	The influence of spectral phase on the temporal envelope of a pulse.	39
2.2	The carrier envelope phase, or absolute phase ϕ_0 , has measurable influence on high harmonic spectra in both experiment (a) and strong field approximation based simulation (b).	41
2.3	The Lorentz model treats electrons as part of a spring-damper-mass system. The nonlinear polarization arises from anharmonic potentials found in real materials - especially noticeable when electrons are strongly driven.	48
2.4	Comparison of centrosymmetric and noncentrosymmetric media and the time-dependent nonlinear polarizations they can support.	49
2.5	Second-harmonic generation displays qualitatively different behavior than difference-frequency generation.	53
2.6	The normalized intensity of the second-harmonic signal plotted against the product of wave-vector mismatch ΔK and length L	56
2.7	Difference frequency generation occurs when two pump beams generate a third beam at the difference frequency of the two pumps.	57
2.8	Wave-vector mismatch (phase-matching) in optical parametric amplification, and the influence of pump intensity and effective nonlinearity, compared with second harmonic generation.	61
2.9	Quasi-phase matching of nonlinear optical crystals.	66
2.10	The influence of quasi-phase matching on wave-vector mismatch for a 3 mm PPLN crystal.	67

2.11	Third-harmonic generation (THG) can occur when a sufficiently intense driving field is incident on a medium with a $\chi^{(3)}$ nonlinear susceptibility.	69
2.12	Self-phase modulation is a third-order nonlinear process that causes changes in the frequency spectrum of a pulse.	71
3.1	A simplified view of the all-fiber front-end laser.	73
3.2	A more detailed schematic of the front-end laser	75
3.3	A photo of the front-end laser.	77
3.4	An overview of optical fiber types used in laser amplifiers.	78
3.5	Dispersion properties of silica in the visible and near-IR.	80
3.6	Menlo Systems figure 9 oscillator diagram and measured spectra.	81
3.7	A detailed diagram of the dispersion managed, broadband EDFA used in the front-end.	86
3.8	SHG-FROG traces for 1.5 μm output of the front-end, as measured immediately after the fiber collimator.	88
3.9	35 fs pulse reconstruction of the 1.5 μm arm EDFA in the spectral and temporal domain using SHG-FROG.	89
3.10	Energy transfer from from a pump to a dispersive wave in a highly nonlinear fiber can be efficient due to a phase matched cascaded four-wave mixing process.	91
3.11	Numerical modeling of pulse propagation in the highly nonlinear fiber.	93
3.12	(a) A photo of the spliced and connectorized highly nonlinear fiber (HNLF). (b) Spectra before (shaded yellow) and after (black line) a ~ 3.5 cm segment of HNLF.	94
3.13	A detailed diagram of the ytterbium-doped fiber amplifier chain.	96
3.14	Output spectra of YDFA1 and YDFA2, with regenerative amplifier output spectrum overlaid in yellow.	97
3.15	(a,b) Photographs of ytterbium gain fiber damage associated with the high average power operation of YDFA2; (c) a photograph showing custom components in YDFA2 to improve output power and system reliability.	99

4.1	Schematic of the cryogenically cooled Yb:YAG regenerative amplifier.	106
4.2	Photographs of the regenerative amplifier showing, from left to right, the custom enclosure and auxiliary equipment for the cryocell, water cooled external beam dump and collimator for thermal stability, and the layout of the optics in the amplifier. . .	107
4.3	Output modes and spectrum of the regenerative amplifier.	108
4.4	Cryogenically cooling Yb:YAG depopulates the lower laser level, causing the gain medium to transition from quasi-three level to four level.	111
4.5	Output power of the regenerative amplifier over multiple days plotted with lab relative humidity and temperature.	112
4.6	An annotated photograph of the SLM-based 1.5 μm pulse shaper.	117
4.7	Algorithmic phase optimization of the compressed 3 μm pulse with the pulse shaper.	118
4.8	The 1.5 μm grism stretcher.	121
4.9	Calculations of the OPCPA pump and signal seed in the time domain.	123
4.10	A photograph of the 3 μm grating compressor with approximate representation of beams overlaid.	125
4.11	A photograph and annotated schematic of the OPA chain, with a “snapshot” of input and output parameters.	128
4.12	Optical parametric generation (or OPG) spectra can be used to determine the optimally phase matched signal wavelengths. Here, we see the phase matching shift to longer signal wavelengths as a 5 mm long, normal incidence PPLN with a 30.49 μm poling period is heated.	130
4.13	A 12 mm thick piece of AR coated NBK7 removes the idler light while transmitting the signal light after OPAs 1 and 2.	132
4.14	A series of signal spectra from the OPCPA using an improved configuration of the seed EDFA.	133
4.15	Unfocused and focused modes of the compressed 3 μm beam.	134

4.16	SHG FROG pulse characterization of the compressed 3 μm pulse, showing a near-transform limited 99 fs FWHM duration.	135
4.17	Numerical simulations of OPA in a 2 mm PPLN with variable idler seed pulse energy.	138
4.18	Analysis of the strong seed simulation set.	139
4.19	Pointing and power stability of the compressed 3 μm beam.	143
4.20	Numerical simulations of the sensitivity of OPA4 to fluctuations in pump and seed energy.	145
5.1	(a) A qualitative single-atom high harmonic spectrum. (b) Kinetic energy of returning electron as a function of ionization phase.	149
5.2	Visualization of a normalized, 2D projection of an electron wavefunction driven by a three cycle pulse of different wavelengths.	150
5.3	The generation of a bright, spatially coherent HHG beam requires the emitted harmonics and the driving laser to be well phase matched.	152
5.4	An illustration of phase matching in HHG.	153
5.5	(a) Critical ionization levels and (b) requisite laser intensities for driving laser wavelengths between 0.8 and 10 μm	158
5.6	Phase matching harmonics in a free focus and a waveguide.	159
5.7	The soft X-ray attenuation lengths and transmission of gases relevant for HHG.	162
5.8	The influence of reabsorption by gas on the optimum HHG flux for various phase matching conditions.	163
5.9	The full HHG source assembly is highly modular, enabling rapid optimization as well as reconfiguration for different driving laser parameters and gas types.	166
5.10	The full HHG source assembly on the pedestal, as it appears on the optical table.	168
5.11	A cross-sectional rendering of the modular HHG source driven by a high NA beam with a capillary cartridge.	170

5.12	A photograph showing a selection of chassis, inserts and targets, taken shortly after the units were machined.	171
5.13	A photograph of the HHG source configured with an argon filled gas cell being driven by the 3 μm driver, indicating jet-like behavior of the argon gas as well as density modulations thought to be shock diamonds.	172
5.14	(a) Microscope images of a CO_2 laser drilled capillary and (b) measured gas flow rates through a laser drilled hollow core capillary and a gas cell.	173
5.15	Differential pumping pressure measurements for multiple atmospheres applied to a laser drilled, 100 μm diameter hollow core capillary.	174
5.16	The gas cell used to generate 3 μm driven argon harmonics.	175
5.17	The soft X-ray beamline (partially to scale).	177
5.18	Raw X-ray CCD images of the HHG beam, including pressure and lens position scans.	180
5.19	Four measured HHG spectra from the argon filled gas cell using different filters. . . .	182
5.20	Low noise HHG spectrum from the argon filled gas cell with a silver filter, highlighting the stability of the harmonic spectrum over 10 minutes.	183
5.21	HHG transmission spectroscopy of a 75 nm boron film mounted on a 500 nm parylene substrate.	184
5.22	HHG spectra generated in nitrogen and argon under identical measurement conditions.	185
5.23	A photograph of 3 μm driven low-order harmonic generation in an argon filled gas cell.	188
5.24	A representative below-threshold harmonic spectrum measured from argon driven in a gas cell with no filtering optics.	189
5.25	Low-order harmonic generation spatial modes.	190
5.26	Measured below-threshold harmonic spectra.	191
5.27	High dynamic range reconstructed spectra of below-threshold harmonics.	193
5.28	Third-harmonic intensity versus driving laser intensity for below-threshold harmonics in argon gas driven by a Ti:sapphire laser.	195

5.29	Seventh-harmonic intensity versus driving laser intensity for below-threshold harmonics in argon gas driven by the 3 μm wavelength mid-IR OPCPA.	196
A.1	Expected scaling of PPLN's damage fluence with pulse duration for a 1030 nm laser, and the corresponding peak intensity.	222
A.2	Temporal envelopes of pulses in a PPLN simulation.	227
A.3	Transform-limit of the generated 1.5 μm signal pulse at the maximum of the gain curve.	228
A.4	Numerical simulation of OPA in 5% MgO-doped PPLN.	230
A.5	Numerical simulation of OPA in 0.5% MgO-doped PPLT.	231
A.6	Numerical simulation of OPA in PPKTP.	232
A.7	Transform-limit of the generated 1.5 μm signal pulse at the maximum of the gain curve, or the back facet of the crystal - whichever plane produces more energetic 3 μm pulses.	237
A.8	Numerical simulation of noncolinear OPA in bulk potassium niobate.	238
A.9	Total near field fluence spatial modes (intensity integrated over the duration of the pulse) from the simulation of noncolinear OPA in potassium niobate.	239
A.10	Numerical simulation of noncolinear OPA in bulk lithium niobate.	240
A.11	Total near field fluence spatial modes (intensity integrated over the duration of the pulse) from the simulation of noncolinear OPA in lithium niobate.	241
A.12	Numerical simulation of noncolinear OPA in bulk potassium titanyl arsenate.	242
A.13	Total near field fluence spatial modes (intensity integrated over the duration of the pulse) from the simulation of noncolinear OPA in potassium titanyl arsenate (KTA).	243
A.14	Numerical simulation of colinear OPA in bulk potassium titanyl arsenate.	244
A.15	Total near field fluence spatial modes (intensity integrated over the duration of the pulse) from the simulation of colinear OPA in potassium titanyl arsenate (KTA).	245
B.1	Raw image of an estimated 12 μJ , 1 kHz, 1.5 μm beam from a silicon CCD.	247

B.2	Raw CCD images (plotted using scaled colors) showing the simultaneous imaging of a 1.5 μm signal beam and a 1 μm pump beam.	248
B.3	Raw CCD images (plotted using scaled colors) showing amplification of a 1.5 μm signal beam in OPA1.	250
C.1	Calculated small-signal parametric gain of a 30.1 μm poling period PPLN at 80 C pumped at 4 GW/cm^2 , with front-end seed spectrum (arb. amplitude) overlaid. . . .	252
C.2	Dependence of the small signal parametric gain on crystal length for a 30.1 μm poling period PPLN at 100 C, with a $d_{\text{eff}} = 14.8 \text{ pm}/\text{V}$ and a pump intensity of 4 GW/cm^2 .253	253
C.3	Dependence of the small signal parametric gain on d_{eff} for a 3 mm long, 30.1 μm poling period PPLN at 100 C, with a pump intensity of 4 GW/cm^2	254
C.4	Dependence of the small signal parametric gain on pump intensity for a 3 mm long 30.1 μm poling period PPLN at 100 C, with a $d_{\text{eff}} = 14.8 \text{ pm}/\text{V}$	254

Chapter 1

Introduction: compact, coherent soft X-ray sources and the lasers that drive them

This thesis describes the development of a stable, compact source of coherent soft X-rays. We generate this radiation in a gaseous medium, through a laser-driven extreme nonlinear optical process known as high-harmonic generation (HHG). The precise boundary of the soft X-ray region of the electromagnetic spectrum is not universally agreed on, but we adopt a common definition that includes photon energies in the range of 0.15-2 keV, corresponding to wavelengths of 8-0.6 nm [1]. HHG produces a temporally and spatially coherent beam of short wavelength light with wide-ranging applications, from high resolution bio-imaging, to semiconductor metrology, to the study of attosecond dynamics in matter and much more. The generation of harmonics reaching keV photon energies requires the use of ultrafast mid-infrared¹ driving lasers [2, 3] - a relatively recent class of laser system that is an active area of scientific research. In this work, we present a 2.95 μm center wavelength optical parametric chirped pulse amplifier (OPCPA) to generate ~ 650 μJ , 100 fs pulses at a 1 kHz repetition rate. Using this laser, we have generated high harmonics in gas spanning to just above the carbon K-edge - with the ready possibility to extend those energies to the keV photon energy range. All told, the lightscape present on the optical table spans more than 9 octaves of optical frequencies: from the 3 μm wavelength driver (0.4 eV) to the 4 nm wavelength harmonics (290 eV).

Driving a high harmonic source is one of the most demanding applications of mid-IR lasers,

¹ The mid-IR spectral range according to ISO standard 20473:2007 includes wavelengths from 3 to 50 μm , though definitions in the literature vary.

requiring the laser to deliver peak intensities in the range of 10^{14} W/cm² from femtosecond pulses. Driving an HHG source intended for practical applications, such as soft X-ray imaging or spectroscopy, places even more stringent demands on the laser. This is the technology space that we are attempting to fill with the source presented here. We are motivated by the goal of developing a workhorse source that functions reliably and without constant intervention by a dedicated team of laser scientists. We hope that these developments will one day bring mid-IR driven soft X-ray HHG sources out of the specialized laser laboratory and into the wide range of application spaces to which they are suited. In order to develop a robust HHG source for applications, we have targeted a number of challenging goals for the driving laser. Optimally efficient harmonic conversion demands tight control of spectral phase and spatial modes. It also benefits from sufficiently energetic driving pulses, ideally in the mJ range, to enable phase-matched harmonic generation with a loose focus or, optimally, a hollow waveguide. Many applications benefit from a sufficiently high repetition rate, even though that can introduce challenging thermal issues, and a high degree of stability in pulse energy, average power, beam pointing, and spatial mode. Finally, we follow a design philosophy that prioritizes robustness, understanding that high system up-time and minimal user intervention are needed for a user-oriented source. In this thesis, we will describe the design, construction and characterization of both the driving laser and the high harmonic source, as well as the nonlinear optical theory, laser science and broad range of technology underpinning this work. We will see that the work presented here brings together cutting edge science from the fields of laser physics, nonlinear optics, frequency comb metrology, optical communications, crystal growth, vacuum science, numerical methods, X-ray science, and more.

In the remainder of this introduction, we will provide context and history for the most important pillars of this work: short wavelength light sources, high-harmonic generation, fiber lasers, chirped pulse amplification, high-average power lasers and mid-infrared lasers. At the end of this section, we will give a brief overview of the entire light source. In Chapter 2, we start with Maxwell's equations to derive and analyze spectral phase - a critical aspect of ultrafast laser pulses and chirped pulse amplification. We extend our classical theoretical treatment to nonlinear

optics, with an emphasis on three- and four-wave mixing processes relevant to optical parametric amplification and nonlinear fiber optics. In Chapter 3, we present a detailed overview of the design, construction and performance of the all-fiber front-end laser, which provides the first light of the system and harnesses fiber nonlinear optics to generate synchronized seed pulses at 1.5 and 1 μm wavelengths. In Chapter 4, we describe the remaining aspects of the OPCPA, including the 1 kHz, 20-mJ class cryogenically-cooled ytterbium:YAG regenerative amplifier pump laser, the periodically-poled lithium niobate (PPLN) based chain of optical parametric amplifiers, and the various instruments of dispersion control. We also present the results of numerical simulations to: address the cause of sub-optimal conversion efficiency; understand the influence of pump and seed fluctuations on amplified power stability; and demonstrate a route towards higher output power by distributing amplification over multiple crystals in series. In Chapter 5, we review the theory of high-harmonic generation from both the microscopic (single atom) and macroscopic (phase matched) picture, introduce a versatile, modular high-harmonic source, and present key results demonstrating below-threshold harmonic generation. We then present some spectra and basic characterization of the soft X-ray high harmonic beam generated in argon and nitrogen. Finally, in the appendix, we include some results that we hope will be useful for future work, including numerical simulations of alternative 1 μm pumped nonlinear crystals in an effort to further optimize bandwidth and pulse energy of the mid-IR driving laser - steps we believe are necessary to extend the photon energy of the high-harmonic beam to the keV level.

1.1 Short wavelength light sources

Scientific progress rests on the tools we have to observe the natural world. In the arsenal of experimental science, optical probes stand out for their tremendous power and versatility: from the ultrafast lasers that have enabled two-photon calcium imaging of neuronal activity [4], to the ultrastable lasers at LIGO that enabled the first observation of gravitational waves [5], tools that started out as experiments in optics labs have found their way into the broader scientific community to enable paradigm-shifting scientific observations. It's not so difficult to trace the origins of the vis-

ible and near-infrared laser wavelengths that enabled those two examples directly back to the first demonstrations of the laser in the early 1960s [6]. The large number of laser gain media with transitions in the visible and near-IR have made such sources ubiquitous - and with ready access to such sources, scientists continue to observe the world in new and clever ways. The visible and near-IR regions, however, represent a small (but important!) component of the electromagnetic spectrum. From the THz regime to hard X-rays, each spectral region offers unique application opportunities and source development challenges. While powerful, the currently available facility-scale sources of coherent soft X-ray light are complex, expensive and sparse. Despite the challenges in generation, coherent SXR light offers a dramatically different window into the natural world, enabling high spatial resolution image formation [7], high temporal resolution dynamical measurements [8] and well-defined elemental specificity [9].

Beginning with Roentgen's discovery in 1895 [10], incoherent radiation from X-ray tubes has had incalculable impacts on medicine and physics. Despite those many early achievements, scientific research with X-rays experienced massive growth following the first experimental demonstration of synchrotron radiation in 1946 [11]. Compared to radiation from an X-ray tube, synchrotron radiation is well-collimated, linearly polarized and very intense. Among many other achievements, according to a 2016 report, synchrotron radiation was responsible for around 70,000 protein structures added to the protein data bank [12]. The usefulness of synchrotron radiation is more limited for ultrafast dynamical studies, due to the typical electron bunch duration of a few 10s of ps [13]. In contrast, a more recently developed type of facility-scale source known as an X-ray free electron laser, or XFEL, is capable of delivering coherent, high brightness radiation in sub-100 fs pulse durations [14]. In an XFEL, relativist bunches of electrons travel through a periodically-poled magnetic field, resulting in acceleration that causes the radiation of coherent X-ray pulses with high brightness. In one recently reported spectroscopy experiment, a temporal resolution of 16 +/- 2 fs was reported - a record for XFELs [15]. While facility-scale coherent X-ray sources are extremely powerful tools, accessibility is limited: only a handful of such sources currently exists worldwide, and may cost $>10^9$ USD to construct. In an effort to address these shortcomings,

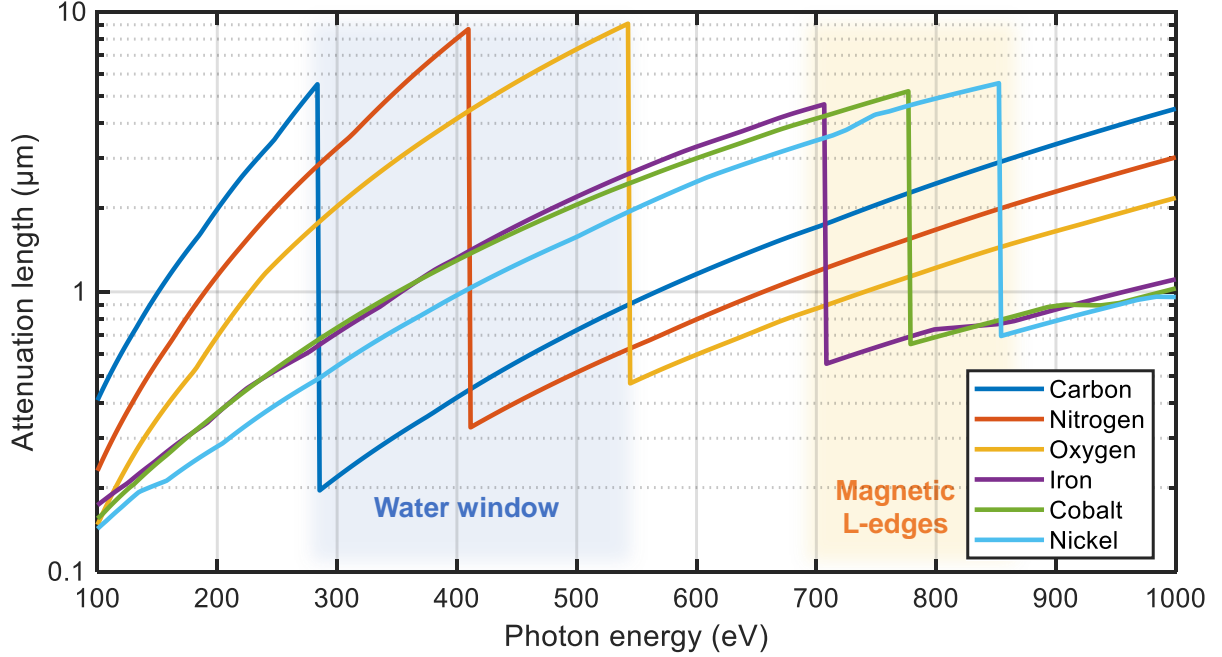


Figure 1.1: **The soft X-ray spectrum spans a number of strong and sharp elemental absorption lines relevant for biological imaging and materials spectroscopy.** The water window spans the K-edges of carbon and oxygen. The relative transparency of water and opacity of carbon make this region of particular interest for biological imaging. Further out, many relevant materials exhibit sharp L-edges - enabling, for instance, element resolved magnetic spectroscopy on femtosecond time scales. The attenuation length is the depth into the material where the intensity falls to $1/e$ of its value at the surface. All materials are assumed to have a uniform density of 1 g/cm^3 . Data was compiled from the Center for X-ray Optics database [9].

scientists are actively developing more compact and economical sources based on laser wakefield acceleration (LWFA) [16] and high-harmonic generation (HHG) - with exciting results. A recent experiment used a 12 m long LWFA-based FEL to generate $\sim 10^{10}$ photons per shot at a central wavelength of 27 nm [17]. High harmonic sources offer a remarkably compact, complementary approach to the generation of coherent short wavelength radiation. Over the past two decades, high harmonic sources of extreme ultraviolet (EUV) radiation in the $\sim 10\text{-}150$ eV spectral range have transitioned from optics experiments to applications oriented sources in widespread laboratory use, largely enabled by the maturation of Ti:sapphire laser technology. As near- and mid-IR ultrafast laser technology develops, we expect further maturation of compact, laser-driven SXR sources. While HHG sources deliver significantly less average power than facility scale sources, they offer




Figure 1.2: Images of the fission yeast *Schizosaccharomyces pombe* using conventional fluorescence microscopy, differential interference contrast (DIC) light microscopy, and soft X-ray tomography. SXR images were taken in the water window, using the ALS synchrotron tuned to deliver 517 eV photons. The organic matter was reported to be about an order of magnitude more absorbing than the surrounding water. Individual organelles were identified by shape and their characteristic X-ray absorption coefficients. Three different SXR tomographic slices are shown next to the resultant reconstruction. The scale bar represents 1 μm . Figure from Reference [19] using work reported in Reference [20].

distinct advantages. HHG sources - comfortably fitting on an optical table - may be the most viable candidate for the widespread dissemination of coherent SXR light into applications-oriented labs and industry. HHG SXR beams retain a high degree of spatial coherence [3] and exhibit temporally coherent supercontinua spanning many octaves, supporting the generation of pulses just 10s of as in duration [18], while maintaining near-perfect synchronization with driving pulses for pump-probe type experiments.

In Figure 1.1, we plot the attenuation length of various elemental materials as a function of photon energy in the range of 100-1000 eV. In the so-called water window, spanning from the carbon (284 eV) to oxygen (532 eV) K-edges, water exhibits relatively high transparency, while the biologically relevant elements carbon and nitrogen exhibit sharp absorption features. Since the 1990s, soft X-ray microscopy driven by synchrotron sources has been able to deliver stunning images

of entire hydrated cells, up to 10 μm in thickness with high contrast and a resolution approaching 30 nm [21, 22]. Figure 1.2 presents an example of soft X-ray tomography of a yeast cell performed at the Advanced Light Source (ALS) synchrotron in Berkeley, CA, alongside images from more standard light microscopy. In that work, individual organelles were identified not only by their shape, but also by their distinctive composition-specific X-ray absorption profiles [19]. The high resolution of SXR microscopy stems from the diffraction-limit, first formulated by Abbe, which tells us that the minimum resolvable length scale in a microscope is proportional to the illumination wavelength [7]. Without refractive optics, soft X-ray microscopy techniques commonly rely on either scanning techniques (e.g., STXM) or coherent diffractive imaging (CDI) [23, 24] - the latter technique relying on the spatial coherence of the SXR beam in order to form diffraction patterns from the sample. Beyond the biological realm, SXR light sources are widely applicable to the metrology challenges of the semiconductor industry, from non-destructive defect detection with ptychographic imaging [25] to near-edge spectroscopy of next generation photoresist materials [26]. Furthermore, SXR radiation from synchrotrons has been dominant in informing our understanding of magnetic materials - with important transition metals such as cobalt and iron exhibiting L-edge transitions in the 700-800 eV range. SXR probes at the L-edges offer improved elemental specificity in comparison with EUV probes at the lower energy M-edges. Nevertheless, the study of magnetic dichroism in the EUV has grown in recent years as a result of the growing accessibility of HHG sources [27].

1.2 Introduction to high-harmonic generation

The first recorded observation of high-harmonic generation was reported in 1987 by an American group operating at the University of Illinois at Chicago [28]. In the work, McPherson et. al focused their 248 nm wavelength KrF* laser, capable of producing ~ 20 mJ pulses with durations as low as 350 fs, into a collection of rare gases delivered into a vacuum chamber by a pulsed valve backed by $1.03\text{-}3.10 \times 10^4$ Torr (200–600 psi). The team sought to characterize perturbative low-order harmonics generated in the gases. In perturbative harmonic generation, as will be shown




Figure 1.3: **Results from three early papers demonstrating high-harmonic generation.** (a) The first reported experimental demonstration of HHG (1987). Relative conversion efficiency is plotted for the first 17 harmonics of a 248 nm laser in neon. The two guiding lines indicate the transition from the perturbative to the newly discovered non-perturbative regime. From Reference [28]. (b) Another demonstration of HHG (1988) from a 1064 nm laser in argon, from Reference [29]. The high relative intensity of the third harmonic is considered erroneous. (c) Results from a computer simulation reported in 1989 showing the first theoretical explanation of HHG. The authors recreate the experimental result shown in (b). From Reference [30].

in Chapter 2, the intensity of successively higher harmonic orders falls off rapidly. The harmonic spectrum that the researchers observed was strikingly different. An original spectrum from the 1987 paper is shown in 1.3(a), showing the generation of odd-ordered harmonics all the way up to the 17th harmonic, corresponding to radiation with a remarkably short wavelength of 14.6 nm. Shortly after, in 1988, Ferray et. al reported generating the 33rd harmonic of an Nd:YAG laser in argon, generating light with a wavelength as short as 32.2 nm. In one of their data sets, reproduced in Figure 1.3(b), the intensity of the 5th harmonic was an order of magnitude weaker than the 3rd - as would be expected from perturbative harmonic generation - but from the 5th to the 21st, the intensity fell only by one further order of magnitude. We now refer to this characteristic feature of the HHG spectrum as the plateau [29]. It was likely no surprise to the researchers that only odd-ordered harmonics were generated; this can be derived from the centrosymmetry of the gaseous medium with arguments that will be made in 2. The harmonic plateau, on the other hand, was something could not be explained by the well-known process of perturbative harmonic generation.

At around the same time, an increase in scientific computing power was enabling detailed

quantum mechanical numerical simulations of the time-dependent Schrödinger equation to describe above-threshold multiphoton ionization in the presence of a strong laser field. The theory of HHG began to take form in work published in 1989 by Kulander and Shore, working at Lawrence Livermore National Lab, who used numerical simulations to reproduce the high-harmonic spectrum reported just a year earlier by Ferray et. al [30]. The original results of Kulander and Shore’s numerical simulations are also included in Figure 1.3(c). In the years following, an intuitive semiclassical picture of HHG emerged that is generally referred to as the *three-step model* [31, 32], illustrated in Figure 1.4. While HHG is fundamentally quantum mechanical, the semiclassical model provides useful insight into the microscopic picture of HHG. The steps according to the model are:

- (1) A bound electron is exposed to a strong laser field (on the order of 10^{14} W/cm²), causing the Coulomb potential holding the electron to the nucleus to distort asymmetrically and lowering the potential energy barrier between the electron and the continuum. This vastly increases the electron’s probability of tunneling through the lowered barrier and escaping towards the continuum.
- (2) If an electron tunnel ionizes, it is accelerated away from the parent ion by the electric field of the laser. When the electric field changes direction, the electron changes direction and is accelerated back towards the parent ion.²
- (3) The electron may recombine with the parent ion, releasing its excess kinetic energy in the form of high-energy photons.

The three-step model makes it clear that the maximum energy of a generated photon - a quantity we will refer to as the single-atom cut-off energy - is determined by the kinetic energy acquired by the electron during its excursion. What can we deduce about the single atom cut-off energy from this simple picture? A free electron of charge e under the influence of an electric field of strength E will experience a force $F = eE$, and therefore an acceleration $a = eE/m_e$, where m_e

² Note that the amount of time the electron undergoes acceleration in part determines the final kinetic energy of the electron, and in turn, the energy of the photons that can be emitted upon recombination. This is an intuitive way to understand the scaling of $E_{\text{cut-off}} \propto I \lambda^2$.




Figure 1.4: **The three-step model of high-harmonic generation.** (a) An electron is bound to the nucleus of an atom by a Coulomb potential (black lines). Under the influence of a strong laser field (on the order of 10^{14} W/cm²), the Coulomb potential distorts, significantly lowering the potential energy barrier between the bound electron and the continuum and increasing the likelihood that the electron will tunnel out (tunnel ionization). (b) The electron undergoes an excursion from the nucleus, acquiring kinetic energy under the influence of the laser field. (c) When the electric field of the laser pulse changes direction, the electron is accelerated back towards the nucleus. (d) Some percentage of the liberated electrons will recombine with their parent ion and emit high energy photons. Figure adapted from Reference [33].

is the mass of the electron. The final velocity of the electron is simply qET/m_e , where T is the amount of time the electron is exposed to the field. In an oscillatory electric field from a laser, the amount of time the electron spends in the electric field before recombining is gated by the electric field cycle, and is proportional to the wavelength λ . In this way, the electron's final velocity after its excursion is $v_{final} \propto E\lambda$. The kinetic energy of the electron, and therefore the cut-off energy, is simply $E_{SA \text{ cut-off}} = m v_{final}^2 / 2$, therefore $E_{SA \text{ cut-off}} \propto E^2\lambda^2 = I_L\lambda^2$. The suggestion of this simple argument that the single-atom cut-off energy is proportional to the product of the laser intensity I_L and λ^2 has been verified experimentally and through quantum simulations [34, 35]. The single atom cut-off energy can be more rigorously written as:

$$E_{SA \text{ cut-off}} = I_p + 3.17U_p \propto I_L\lambda_L^2, \quad (1.1)$$

where I_p is the ionization potential of the gas and U_p is the ponderomotive energy $U_p = e^2I_L\lambda_L^2/16\pi^2m_e c^2$

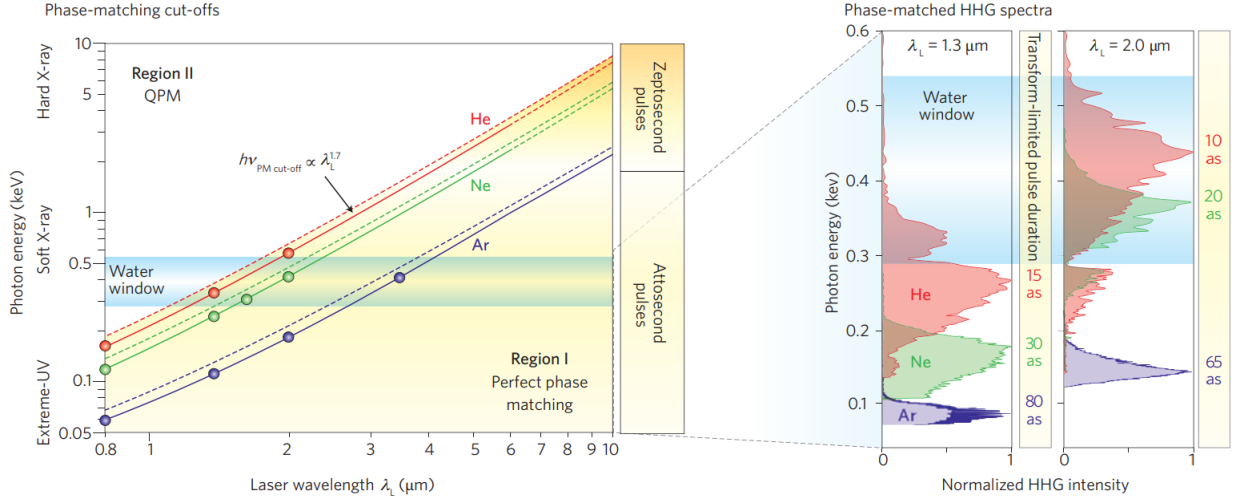


Figure 1.5: **The phase matched cut-off energy for high-harmonic generation scales as $I_L \lambda_L^{1.6-1.7}$, motivating the development of ultrafast mid-infrared lasers.** The higher ionization potential of lighter atomic species allows critical ionization and therefore phase matching to be realized with higher laser intensities, pushing the cut-off energy to higher energies. Experimental HHG spectra validate the scaling of HHG photon energies with atomic species and driving laser wavelength. Figure adapted from Reference [36].

(the cycle averaged energy of a free-electron in an oscillating electric field). The scaling of the photon energy of emitted harmonics with the driving laser wavelength is illustrated in Figure 1.5, where it can be seen that a 3 μm driving wavelength can in principle generate phase matched harmonics in helium gas with photon energies up to about 1 keV ($\lambda = 1.2 \text{ nm}$). The wavelength scaling of high-harmonic generation has motivated us to develop the 3 μm laser system described in this thesis.

As with other nonlinear processes described in this thesis, the development of a *bright*, high-flux harmonic source requires the generated harmonics to be *phase matched* with the driving laser. When phase matching is achieved, the harmonics travel with nearly the same phase velocity as the driving field, allowing all components of the harmonic field generated over the phase matched length of the interaction region to add constructively. The percentage of atoms which can be ionized in a phase matched condition is known as the *critical ionization*. For helium and argon driven by a 3 μm wavelength driver, the critical ionization is 0.034% and 0.26%, respectively - preventing one from simply turning up the laser intensity to increase the cut-off energy. In addition to the

contribution from ionized atoms, the effective refractive index of the medium is also influenced by gas pressure and geometry. The phase matching constraints have been experimentally determined to slightly reduce the wavelength scaling of the cut-off energy to $E_{\text{PMcut-off}} \propto \lambda^{1.6-1.7}$ [37].

In practice, the short wavelength light from HHG is often strongly absorbed by the gases which enable its generation. The balance between harmonic generation and absorption in the medium puts limits on the achievable brightness of a source in two key ways: (1) reabsorption in the gas target itself causes diminishing returns on brightness as the laser-gas interaction length is increased, and (2) reabsorption from residual gas in the region after the gas target can also diminish useable flux. For the first effect, a simple 1D analysis by Constant et. al has been reported that considers local emission, phase mismatch and reabsorption in a single expression. Even for the case of perfect phase matching, the harmonic flux saturates after about ten absorption lengths, with only about three absorption lengths needed for half the optimal efficiency [38]. Nevertheless, especially for the case of helium harmonics driven by mid-IR lasers, those length scales can be quite long - generally much longer than the Rayleigh range accessible by available laser systems - necessitating waveguide geometries for sustained high peak intensity over long lengths [39]. For the second effect, in most HHG geometries, gas emitted from the interaction region continues to absorb the harmonic light after the interaction region, demanding carefully designed differentially pumped vacuum systems. Adding the challenges associated with reabsorption in achieving high fluxes, the single atom efficiency of HHG has been shown to scale poorly with increase in the driving laser wavelength. Depending on how the scaling is defined, single atom efficiency scales as between λ_L^{-5} and λ_L^{-9} [32, 40, 37, 41]. Nevertheless, improved phase matching conditions with long-wavelength drivers and the higher transparency of helium in the soft X-ray region may compensate for some of these losses in efficiency.

The single atom picture, phase matching, reabsorption, wavelength scaling and HHG source design will all be discussed in greater detail in Chapter 5. This thesis, however, is largely about the development of a novel laser system to drive HHG. In the subsequent sections, we introduce the laser technology that drives this unique light source, beginning with fiber lasers.

1.3 History of fiber lasers

The earliest recorded discussions of the guiding of light in media through total internal reflection occurred in the early 1840s, when researchers D. Colladon at the University of Geneva and J. Babinet in France observed optical waveguide effects in water jets. Decades later, in the 1920s, researchers explored the use of thin glass rods to guide light for medical applications. The 1960s marked a period of rapid progress in field of optics and lasers, including research into optical transmission media. By 1961, E. Snitzer had published high quality theory on optical modes in dielectric waveguide [42]. But the early optical fibers were extremely lossy, with a typical attenuation of 1000 dB/km. C. Kao helped resolve this issue through the use of high-purity fused silica glass.³

Just one year after T. Maiman published about stimulated radiation in a ruby crystal [6], E. Snitzer published an experimental demonstration of optical maser action in a 300 micron diameter Nd³⁺ doped glass fiber [44]. It is often overlooked that fiber laser development occurred nearly simultaneously with the development of free space lasers. Like many physicists of the era, Snitzer experienced a major interruption to his career when, in 1958, he was called to testify before the House Un-American Activities Committee [45]. Nevertheless, Snitzer persisted, and in 1964, he published the first demonstration of a fiber laser amplifier with his colleague C. Koester [46]. Together, they showed gain as large as 5×10^5 in a 1 meter long neodymium-doped glass fiber.

In the subsequent decades, the rapid growth of global communications further motivated research in long distance transmission of optical signals. In 1988, Bell Labs brought online the first transatlantic fiber-optic cable. Known as TAT-8, it was designed to carry 565 MBit/s - less than some household internet connections today, but an order-of-magnitude improvement over the previous transatlantic coaxial cable. These early fiber-optic systems operated at a wavelength of 1310 nm driven by InGaAsP semiconductor lasers. It was known that optical attenuation was lower at longer wavelengths, but no suitable 1.5 μm laser source existed at the time [48]. The lack

³ This brief historical overview was derived from the scientific background published with the 2009 Nobel Prize in Physics [43], which was awarded in part to Charles K. Kao for his work in optical communication.



Figure 1.6: **An Alcatel Submarine Networks optical repeater being deployed in the ocean.** EDFA technology revolutionized global telecommunications in the 1980s and 1990s. The 976 nm, 950 mW fiber Bragg grating stabilized laser diodes used to pump the EDFAs described in Chapter 3 are qualified to the highest Telcordia reliability rating and are certified for use in submarine applications. The front-end laser benefits from the reliability, low cost and availability of fiber-optic components in the 1.5 μm band - although we still avoid flooding in the lab. Photo from Reference [47].

of a 1.5 μm source was not the only issue that plagued the early fiber-optic networks. Repeater stations, which were needed to periodically boost the signal of long-distance communication lines, relied on the conversion of the optical signal to an electronic signal before the optical signal could be regenerated by another laser [49].

In the 1980s, erbium-doped glass was developed as an effective laser gain medium, able to produce light around the low-loss 1.5 μm band. The first erbium-doped fiber amplifiers (EDFAs) were developed in the late 1980s by groups at the University of Southampton and Bell Laboratories [50, 51]. EDFAs not only allowed for operation in a better region of the spectrum of waveguiding in silica glass, they also eliminated the need for inefficient and limiting optical-electronic converters, instead allowing for the direct amplification of light in repeater stations. This spurred massive investment into fiber-optic technology in the 1.5 μm region, some of which was utilized in this work to develop the front-end (including wavelength division multiplexers, polarization-maintaining fiber, and high-quality fiber-coupled diode pump lasers operating around 976 nm).

1.4 High average power solid-state lasers

Many demanding scientific and industrial applications of lasers require not just high peak intensity, but also high average power. For the application discussed in this thesis, we have multiple compelling reasons to scale our systems to higher average power. High-harmonic generation intrinsically exhibits a low conversion efficiency. One way to increase HHG flux and the usefulness for experimental applications is to increase the repetition rate. Due to their statistical nature, some measurement techniques - such as HHG-based molecular dynamics studies using coincidence detection [52] - may only be feasible with repetition rates $>$ kHz. Another factor contributing to the need for high-average power lasers in nonlinear parametric amplifiers is the generally low conversion efficiency of pump energy in optical parametric amplification. The system presented here exhibits approximately 5-10% conversion from pump energy to 3 μ m wavelength light - roughly in line with commonly reported values. In practice, lower repetition rates are widely used for multi-mJ nonlinear frequency conversion based sources: for instance, the record-setting 2012 demonstration of 1.6 keV HHG photons was set using a 20 Hz laser system [3]. Setting the bar even higher, the laser-plasma accelerator community has called for the development of ultrafast lasers that operate in the multi-kW or even hundreds of kW average power range [53]. Laser technology has come a long way in answering these calls. The last few decades has seen significant improvement in the average power handling of solid state lasers, with two contributing factors being improvements in diode pump lasers and amplifier architectures. In this section, we discuss the latter, emphasizing the role that simple geometry of the gain medium can play in thermal management of high-average power laser systems.

A solid-state laser is a laser which utilizes a solid-state gain medium, such as a crystal or doped glass. Energy is added to the gain medium, typically through optical pumping, and is stored in the form of excited electrons. Later, these excited electrons decay to a lower energy level through a process known as stimulated emission. The energy which is absorbed by the gain medium and not converted to laser light is typically converted to heat. In lasers which operate at a high average




Figure 1.7: **Solid-state laser gain media can be broadly classified by to their dimensionality, which significantly affects thermal management and average-power handling.** Bulk crystal or doped-glass media have three large spatial dimensions and no small dimension for heat to easily be extracted. Crystals are usually surrounded on all non-optical surfaces with a heat sink (not shown). Thin-disk media have two large spatial dimensions and one small dimension for easy heat extraction to a heat sink (shown in grey). Fiber gain media have one large spatial dimension (length). Heat can easily be extracted from the outer surface of the fiber through the remaining two dimensions to a heat sink (not shown).

power, this heat load on the gain medium can cause detrimental thermal effects, such as thermal lensing. As a result, a key consideration in designing high average power laser amplifiers is thermal management. In recent years, with the development of high average power thin-disk amplifiers and fiber amplifiers, it has become useful to consider thermal issues in the context of the surface area to volume ratio of the gain medium (Figure 1.7).

The earliest technology used for high average power solid-state laser amplifiers utilizes a bulk crystal gain medium, which we consider thick in all three spatial dimensions. This architecture exhibits the largest distance between the heat-generating optically active ions and the surface of the medium. Typical laser crystals or slabs in this arrangement might have all three spatial dimensions on the order of the beam diameter. Typically, heat is extracted transversely to the beam propagation direction, causing thermal gradients across the beam profile which can lead to problematic thermal lensing. Most laser crystals, including Yb:YAG, exhibit an increase in thermal conductivity at lower temperatures [54]. In Ti:sapphire amplifiers, KMLabs reported that cooling the crystal to 50-80 K using a helium gas cryo-cooling system results in a greater than 200x decrease in thermal beam distortions [55]. While cryogenic cooling can mitigate some thermo-optic effects, high average power is problematic for bulk crystals.

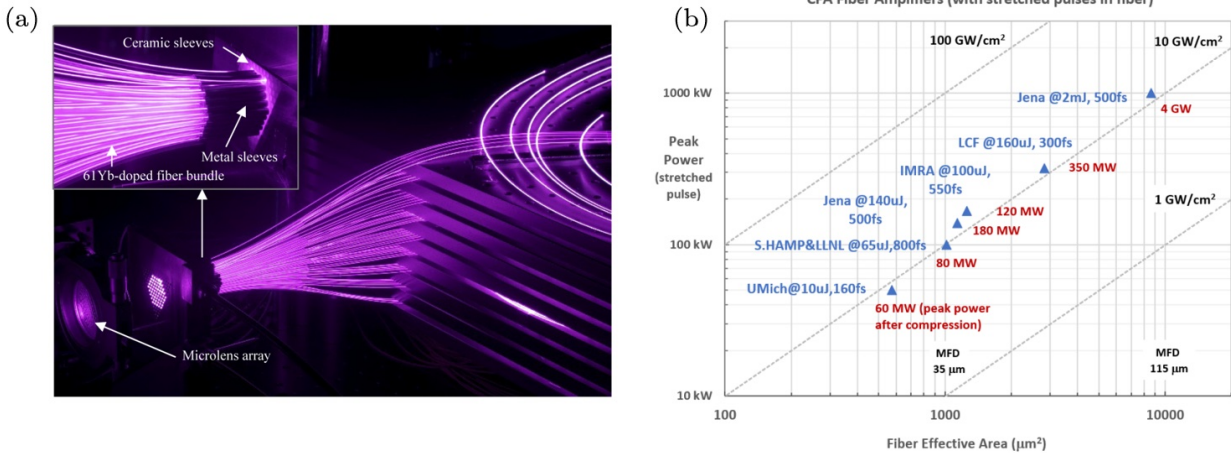


Figure 1.8: **Single fiber chirped pulse amplification systems have amplified multiple GW of compressed peak power, with the potential for scaling via the combination of many fibers - potentially enabling both high peak and average power systems.** (a) Coherent beam combination can be thought of as a spatial analog to chirped pulse amplification, allowing for the spatial distribution of amplification across many large mode area fibers. Here, a photograph of the XCAN laser is shown, which combines beams from 61 Yb fiber amplifiers with relative phase errors as low as $\lambda/90$ RMS [58]. (b) A limited 2022 survey of research work in power scaling of large mode area fiber chirped pulse amplifiers. Blue notations are the pulse energy and duration after compression, and red notations are the peak powers after compression [59].

More recent technology includes thin-disk amplifiers [56], in which the gain medium comprises a disk which is thin relative to the beam diameter (often a couple hundred μm) in the direction of laser propagation. The back surface of the medium is reflective, leading these systems to also be known as active-mirror lasers. The thin, quasi-two dimensional disk gain medium is mounted to a heat sink, allowing efficient heat extraction along one dimension. Such media exhibit improved thermal management, not only because of shorter paths between the heat source and the heat sink, but also because the thermal gradients are in the longitudinal direction, contributing less to transverse thermal lensing. However, the beam path through such media is very short, leading to very low single pass gains. To achieve high gain, these systems can involve hundreds of passes through the crystal, greatly increasing system complexity [57].

Lastly, we consider the case of fiber amplifiers. Since the gain medium is contained in a very thin (often few-hundred μm) filament, it can be considered quasi-one dimensional. Heat can therefore easily be extracted over a short path along two spatial dimensions, leading to impressive

thermal management and the ability for fiber amplifiers to amplify high average powers. Modal distortions due to nonlinearities can to some degree be mitigated by the waveguiding effects of the fiber. As an example, the most recent generation of the aeroGAIN rod amplifier from NKT Photonics specifies the amplification of 250 W of single-mode power from a single pass through a Yb-doped 85 μm core diameter large mode area fiber [60]. Multi-mode fiber laser amplifiers can exhibit even higher average powers, and are now commonly used in industrial applications such as materials processing, cutting and welding. However, owing to nonlinearities and even damage that can occur due to the tight confinement of the amplified laser beam and the large amount of material that the beam must pass through, fiber amplifiers are severely limited in their peak power handling, and amplification of high peak power pulses requires chirped pulse amplification. Figure 1.8(b) provides a survey as of 2022 of chirped pulse fiber amplifiers showing a general limit of stretched pulse peak power in fiber of around $10 \text{ GW}/\text{cm}^2$ due to nonlinearities. To some extent, this issue can be addressed through the multiplexing of multiple fiber amplifiers, which has been proposed as a scalable architecture for high peak power, high average power applications such as laser plasma accelerators [61]. Figure 1.8(a) shows a photograph of the multiplexing of 61 large mode area fiber lasers in the XCAN laser at the Institut Polytechnique de Paris. The front-end described in Chapter 3 has been used to seed a large-mode area ytterbium fiber amplifier, and investigations are ongoing for their use to pump optical parametric amplification.

1.5 Chirped pulse amplification to generate high peak powers

In the 1970s and early 1980s, laser physicist found themselves up against a wall as they tried to drive up peak intensity by increasing pulse energy. If a short laser pulse seeds a high gain amplification system, such as an optical parametric amplifier or a regenerative amplifier, the pulse can easily develop sufficiently high intensity to induce nonlinear responses or damage in optical components, such as crystalline or fiber gain media, polarizers, or mirrors. Although mJ level pulses amplified from a mode-locked oscillator had been demonstrated, the only route forward to increase pulse energy without damaging the amplifier was to increase the mode size, resulting in

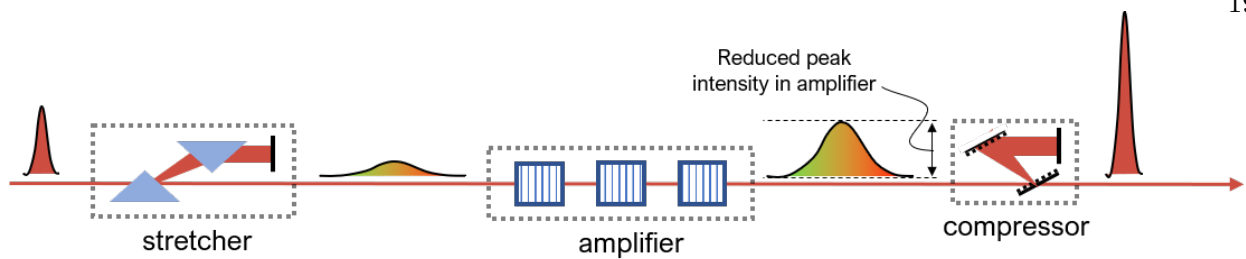


Figure 1.9: **A depiction of chirped pulse amplification as it used in OPCPA.** In chirped pulse amplification, a short seed pulse is first temporally stretched in a dispersive optical component known as a stretcher. The elongated pulse contains comparable pulse energy but much lower peak intensity, reducing the effect of nonlinearities and the likelihood of damage as it is amplified. Finally, a dispersive optic known as a compressor, which introduces the opposite spectral phase as the stretcher, reduces the pulse duration to dramatically increase the peak intensity. High peak intensities are useful for many applications, such as nonlinear optical experiments.

large and expensive installations [62].

Attainable pulse energies, which had nearly plateaued by the early 1980s, began to increase again with the advent of chirped pulse amplification (CPA) - a technique borrowed from radar researchers and first applied to short pulse lasers by Strickland and Mourou in 1985, who demonstrated the technique with an Nd:glass amplifier [63]. Strickland and Mourou were awarded the 2018 Nobel Prize in physics for this work. In CPA, a short seed pulse is first stretched in time through the addition of a known and compensable spectral phase to reduce its peak intensity, then amplified, and finally compressed to a high peak-intensity, energetic, short pulse. CPA is illustrated in Figure 1.9. Every modern amplification scheme to drive high-harmonic generation that the author is aware of utilizes chirped pulse amplification in some capacity.

Chirped pulse amplification can be applied to direct laser amplification (as in solid-state laser amplifiers) or in nonlinear parametric amplification to create an optical parametric chirped pulse amplifier, or OPCPA. In an OPCPA, a (generally) large bandwidth signal is stretched, amplified in a nonlinear crystal(s) by a three-wave mixing process called optical parametric amplification, and finally recompressed. The physics of optical parametric amplification will be reviewed in Chapter 2. Stretching and compressing laser pulses primarily involves the application of group delay dispersion

(GDD, also commonly referred to as second order dispersion)⁴ with some sort of dispersive optical element, like a grating, optical fiber, or prism.

1.5.1 How group delay dispersion affects pulse duration

It is useful to be able to easily calculate how the pulse duration of a Gaussian pulse changes when some amount of GDD is applied by some dispersive optic. Here, we consider a transform limited pulse with a duration of τ_{TL} . When we apply GDD, the pulse is stretched to τ_{str} . The relationship is defined by the following formula (from [64]):

$$\tau_{str} = \tau_{TL} \sqrt{1 + \left(\frac{4 \ln 2 GDD}{\tau_{TL}^2} \right)^2} \quad (1.2)$$

As an example, the stretched signal and idler pulses in the OPCPA described here have a GDD of around $-2.85 \times 10^6 \text{ fs}^2$ (alternatively, -2.85 ps^2), resulting in stretched pulses of around 200 ps.⁵ Equation 1.2 is plotted in Figure 1.10 for a large range of transform limited input pulse duration. We highlight the approximate operating regime of the signal and idler pulses in the OPCPA.

We note some important takeaways regarding pulse compression from Figure 1.10. The influence of a fixed amount of GDD on a stretched pulse duration decreases as the pulse duration of the transform limited pulse increases. This has practical consequences for the laser described in this thesis. The cryogenically cooled Yb:YAG pump laser has a narrow gain bandwidth, producing pulses with a transform limit of around 6 ps. These pulses are stretched to around 275 ps as the result of GDD applied in the chirped fiber Bragg gratings in the stretcher. The narrow bandwidth means that compressing these pump pulses would require the application of a very large amount of GDD - on the order of $6 \times 10^8 \text{ fs}^2$, which is about 2 orders of magnitude larger than the GDD applied by the mid-IR pulse compressor. While higher peak intensity pulses can be useful to pump some OPCPA systems, such compressors are generally very large and can be difficult to build. Divided

⁴ Due its importance in ultrafast laser physics and chirped pulse amplification, pulse dispersion and spectral phase will be covered in depth in Chapter 2

⁵ This was the amount of GDD in our experimental configuration that first drove HHG. More recent upgrades increased the bandwidth of the front end and require less GDD to reach the same pulse duration.

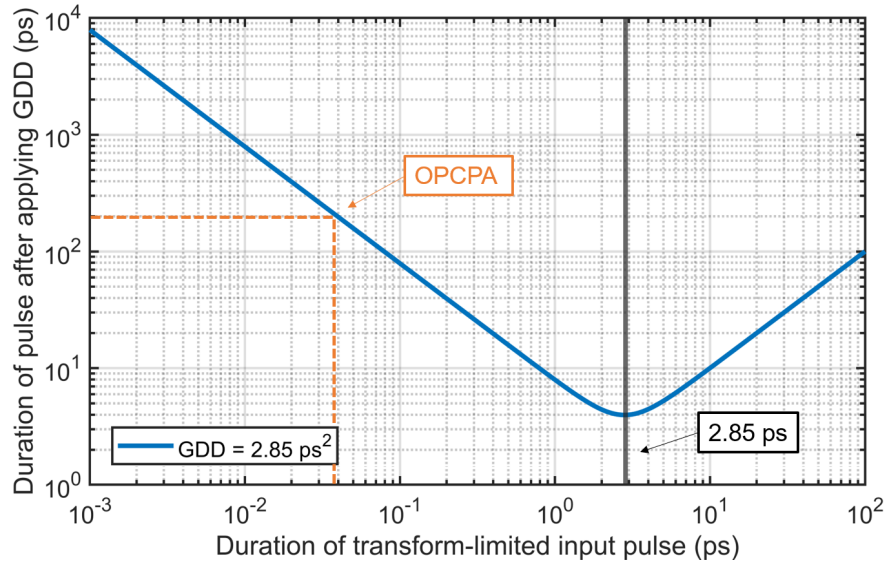


Figure 1.10: **How the application of group delay dispersion (GDD) from the stretcher influences pulse duration in the OPCPA.** The grism stretcher applies -2.85 ps^2 of GDD to the signal pulse, stretching it to around 200 ps (indicated in orange). Note that if the same amount of GDD were applied to the narrowband, 6 ps transform limit pump pulses, the pulse duration would be all but unaffected.

pulse amplification may be suitable in such cases [65].

1.5.2 Optical parametric chirped pulse amplification

OPCPAs are a type of laser amplifier which utilize chirped pulse amplification to achieve high peak intensity pulses through a nonlinear amplification technique known as optical parametric amplification (OPA), which will be discussed in detail in Chapters 2 (theory) and 4 (experiment). Compared to direct amplification of light through stimulated emission, OPCPAs offer advantages and drawbacks. In OPA, no energy is transferred to the gain medium - rather, energy is quasi-instantaneously transferred between the pump and the signal/idler through a mixing process in a transparent, nonlinear medium. Barring (generally) unwanted absorption, this process creates very little thermal load on the crystal, which can simplify implementation and power handling. Extremely high single pass gains (for example, around 10,000 in the first OPA of the system presented in this work) vastly simplify architecture, which improves system stability and avoids

temporal problems such pre-pulse formation, which is a common issue for regenerative amplifiers. Further, the spectral regions that can be amplified are limited only by the transparency of the crystal and phase matching, and therefore can be tuneable over a broad range, unlike with direct amplification, where the spectral gain is set by electronic transitions in the gain media.

On the other hand, pump lasers for OPA/OPCPA are often considerably more complex than the lasers which excite traditional gain media. A laser gain medium retains the energy of the pump for time scales on the order of the lifetime of the medium. This means that for broad bandwidth, direct amplification systems such as Ti:sapphire or Cr:ZnSe, ns pulse duration Q-switched lasers can often be used; for media with longer lifetime, such as Yb, very robust and simple diode-array CW lasers can be used as pumps. In OPA, efficient transfer of energy from the pump to the signal and idler waves requires that the pulse duration of all three be comparable, demanding sophisticated ps or even fs lasers to pump OPAs. Further, since OPA is a nonlinear process, it generally requires high laser intensities on the crystals, forcing operation near the damage threshold. This can drive problematic effects such as photorefraction and self-focusing in nonlinear crystals. OPA can also cause significant distortions in the temporal, spectral and spatial structure of the pulse, which are generally all coupled (such as in back-conversion). Lastly, without resorting to sophisticated pulse shaping methods, a commonly accepted conversion efficiency from pump to the combined energy of the signal and idler is 30%. The energy distribution between the signal and idler depends on their relative frequencies - in the case where the idler is much longer wavelength than the signal, a relatively small percentage of pump energy can be transferred to the idler. Improvement in conversion efficiency has been demonstrated by absorbing or depleting one of the generated waves [66, 67, 68]. Fundamentally, this interesting effect occurs because back-conversion requires the combination of both a signal and an idler photon to form a pump photon. Eliminating either the signal or the idler prevents that process. The same effect can be seen in our own numerical simulations exploring very strongly seeded OPA (found in Chapter 4), where we find that a strong seed does not herald the earlier onset of back-conversion and a decrease efficiency - but rather, back-conversion occurs only with sufficient intensity of *both* the signal and the idler.




Figure 1.11: **Spatial mode distortion of a ~ 15 mJ, 30 fs pulse due to severe nonlinearities.** Prior to developing the Yb:YAG pumped OPCPA, we developed a Ti:sapphire pumped (un-chirped) OPA. Chirping the energetic pump pulses in an OPCPA vastly relieves high peak power issues associated with un-chirped, high energy OPA systems.

We conclude this section by providing an anecdote from our own research which motivates the OPCPA architecture. Prior to focusing our attention on the development of a Yb:YAG pumped OPCPA, we developed a BBO based (un-chirped) OPA to generate multi-mJ, sub-2 μm wavelength, femtosecond pulses to drive HHG. The OPA was pumped by an ~ 20 mJ, 30 fs two-stage Ti:sapphire amplifier. The extremely high peak intensities of the pump laser lead to severe nonlinearities due to propagation through transmission gratings in the compressor and even through air. A spatial mode of this pump laser showing the effects these nonlinearities is presented in Figure 1.11. Increasing the energy of the pump pulses without stretching would require dramatic solutions, such as very large beam profiles or housing the entire compressor in a large vacuum chamber. In (un-chirped) OPAs, the so-called *pulse-splitting length* - the propagation length for which the signal or idler temporally walk off the pump - can also become prohibitively short, a problem readily solved by stretching the pulses. Further, the two-stage Ti:sapphire pump itself required three separate Q-switched pump lasers. The considerable complexity of the system limited robustness and long-term stability, partially motivating the pursuit of OPCPA technology.

1.6 The current state of long wavelength lasers to drive high-harmonic generation

In the past decade, we have witnessed impressive advances in the field of long wavelength, ultrafast lasers. Much of this progress has been at least partially motivated by the $E_{\text{HHG}} \propto I_L \lambda_L^2$ wavelength scaling of high-harmonic generation. As will be described in Chapter 5, the generation of higher photon energies in HHG requires the use of longer wavelength (λ_L) driving lasers - simply increasing driving laser intensity I_L would eventually over-ionize the gas and prevent the phase matching that is necessary for the generation of a bright, coherent beam. Here, we review the most important recent results in ultrafast, long wavelength lasers, to provide context to the work presented in this thesis. Of the laser architectures surveyed, optical parametric chirped pulse amplification (OPCPA) dominates the field. Nevertheless, transition metal and rare earth doped bulk laser gain media, as well as thulium doped fibers, have shown exciting promise as sources of long wavelength, femtosecond pulses.

1.6.1 A survey of long wavelength OPCPAs producing femtosecond pulses

The first experimental demonstration of OPCPA was by Dubietis et al. in 1992 [69], which utilized a utilized an 8 ps, second harmonic pulse from an Nd:glass regenerative amplifier to pump a BBO crystal. The system operated at degeneracy, being seeded by an Nd:glass oscillator lasing at 1055 nm. Since then, the progress of OPCPA technology can be pegged to the progress of its underpinning technologies: pump lasers with sufficiently high average power and pulse energy, nonlinear crystals with appropriate transparency and a high nonlinearity, and sufficiently broad bandwidth front-end seed sources. A survey of few cycle OPCPA systems over the past two decades is presented in Figure 1.13(a). Systems operating in the NIR tend to rely on Ti:sapphire oscillators as seed sources, and Nd or Yb amplifiers as pumps, and early systems favored BBO as a nonlinear crystal. Some recent systems operating in the mid-infrared are utilizing 2 μm Ho pump lasers and non-oxide semiconductor crystals (such as zinc germanium phosphide, or ZGP), and often have

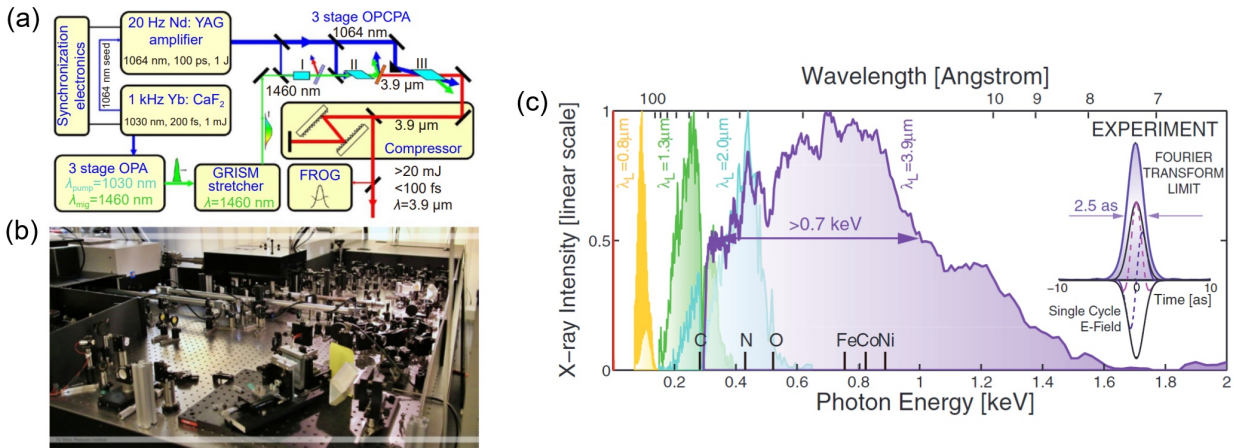



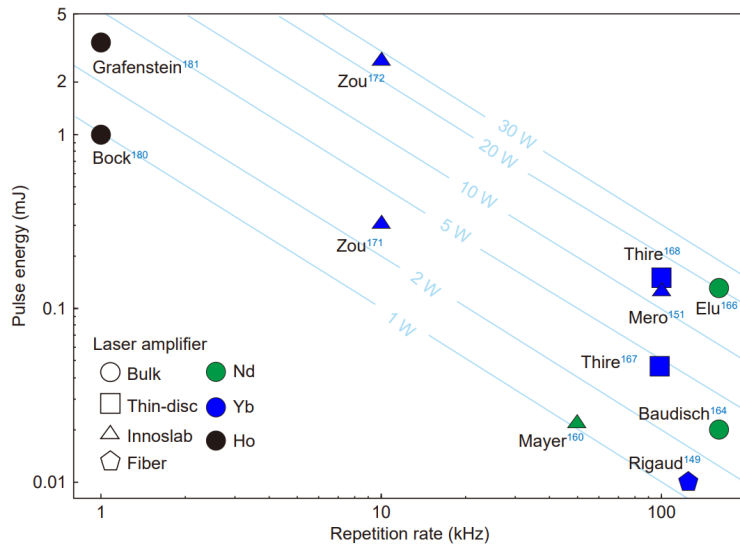
Figure 1.12: **A high harmonic spectrum extending to 1.6 keV, and the 3.9 μm wavelength OPCPA which drove the process.** (a) The layout of the OPCPA, which used a series of three KTA crystals pumped by a 20 Hz Nd:YAG amplifier, and delivered 10 mJ, 80 fs pulses at 3.9 μm . (b) A photograph of the back end of the system. (c) Experimental HHG spectra emitted driven by various wavelengths, including the 3.9 μm OPCPA. (a) and (b) were adapted from Reference [70]; (c) is from Reference [3].

sophisticated front-ends based on nonlinear frequency conversion. A survey of mid-IR OPCPAs (Figure 1.13(b)) shows that these systems are currently limited to a few tens of watts in average power output, but are rapidly scaling up with the high average power pump laser technology described in the previous section.

There are only two mid-infrared ($\lambda_L \geq 3 \mu\text{m}$) laser sources to date that have successfully driven high-harmonic generation in a gas: the 1 kHz system presented in this thesis and the 20 Hz 3.9 μm OPCPA at TU Wien, described in [71] and illustrated in Figure 1.12. The 3.9 μm system is perhaps most famous for the 2012 generation of 1.6 keV photons - the highest energy HHG photons ever generated, corresponding to harmonic orders greater than 5000 [3]. The OPCPA employed in the experiment underwent a number of modifications over the years, but is generally based on a 20 Hz Nd:YAG pump laser delivering joule level pulses. The pump energy was split between three KTA crystals, which amplified a signal at 1.46 μm with a corresponding idler at 3.9 μm . The signal seed was generated in a three stage OPA and stretched in a grism stretcher, to enable compression of the idler in a grating compressor to 80 fs with 10 mJ of energy. The OPCPA was later used for



(a) A survey of published OPCPA systems, plotted against pulse width and wavelength. The marker fill color denotes the type of nonlinear crystal used according to the legend, and the width of the horizontal bars through each marker represents the full width of the amplified spectrum. The years at the top of the figure mark the year that the corresponding wavelength was first experimentally generated with OPCPA.



(b) A survey of high average power mid-infrared OPCPA systems. The marker fill color and marker shape denote the pump laser gain medium and architecture, respectively.

Figure 1.13: A survey of few-cycle OPCPA systems. Both figures are from a 2022 survey in Reference [70].

K- and L-edge near- and extended-edge X-ray absorption spectroscopy at energies near 1 keV [72].

A number of other mid-IR ultrafast OPCPA have been reported in the literature. The Biegert group at ICFO has developed a 3.25 μm OPCPA that generates 131 μJ , 97 fs (sub-9 cycle) pulses at 160 kHz (21 W average power), reported in 2017 [73] and later in 2020 [74]. Using this source, harmonics up to the 7th order were generated by reflecting the compressed pulses from a 100 nm thick YCOB film.⁶ Spectroscopy was performed by analyzing the harmonics. Despite the low order and relatively low driving intensity (90 GW/cm²), the harmonics were considered to be generated in the nonperturbative regime [75]. The same group reported a 7 μm OPCPA in 2019, which delivers 750 μJ pulse energy. It is based on a series of ZGP crystals pumped by a 260 mJ, 2052 nm amplifier, delivering 16 ps pulses at 100 Hz. The system generated harmonics up to the 13th order in ZnSe [76].

In 2018, Thiré et al. from Fastlite published a 3.2 μm , 152 μJ , 38 fs (sub-4 cycle) OPCPA delivering an average power of 15.2 W at a repetition rate of 100 kHz. The system is pumped by a Yb thin disk amplifier delivering 1.1 ps pulses and utilizes PPLN for the first two OPA stages, and unpoled lithium niobate for the final two stages. The shorter pulse duration of the pump allows the use of lower d_{eff} unpoled lithium niobate while still maintaining bandwidth, the final crystal being 1.5 mm long and driven at 90 GW/cm². The system was seeded at 1.53 μm by a supercontinuum generated in YAG [77]. In 2023, harmonics up to the 13th order in solid thin films were reported in a pre-print using this system at its new home at ELI-ALPS in Hungary [78]. In 2022, a group at CREOL presented a 3.1 μm , 3.2 mJ, 92 fs OPCPA delivering an average power of 3.2 W at a repetition rate of 1 kHz. The system is composed of a chain of ZGP crystals pumped by a 2 μm Ho:YLF amplifier. The front end utilizes a 2 mJ Ti:sapphire amplifier. Harmonics up to the seventh order were observed by focusing the laser into air [79].

The Max-Born-Institut (MBI) in Berlin has developed several impressive OPCPA systems. In 2020, a very well thought out 4.9 μm OPCPA was reported that delivers 3.4 mJ pulses as short as 89 fs at a 1 kHz repetition rate. The system comprises a chain of ZGP crystals pumped by 5

⁶ YCOB is a high temperature superconductor.

ps pulses from a Ho:YLF amplifier. The regenerative amplifier followed by a booster delivered 36 mJ of pump energy. Although the final 3.5 μm signal power was not reported, it can be estimated to be around 5 mJ - an ideal range for SXR HHG experiments, suggesting that a booster amplifier will likely be necessary for 2 μm pumped OPCPAs for HHG. No high harmonics were reported, but the peak intensity of 33 GW at 4.9 μm is a record for OPCPAs operating above 4 μm [80]. In 2022, the MBI reported an extremely long wavelength 11.4 μm OPCPA, delivering 65 μJ of pulse energy in 185 fs (sub-5 cycle). The long wavelength light is generated as the idler in a three stage GaSe OPCPA pumped by a 2 μm Ho:YLF amplifier. Notably, the seeds for both the pump laser and the 2.4 μm signal originated from the same pulse, generated in a femtosecond Cr:ZnS laser [81].

Several high harmonic sources, some with cut-off energies extending into the soft X-ray, have been driven by OPCPAs at wavelengths in the 2-3 μm range. In 2014, a group at MIT reported a 2.1 μm OPCPA generating 2.6 mJ, 40 fs pulses at 1 kHz. Similar to the system presented here, the MIT group used a cryogenically cooled Yb:YAG amplifier as a pump laser [82]. In 2020, the MBI presented a 2.1 μm OPCPA delivering 2.7 mJ, 30 fs pulses at 10 kHz. The higher average power of 27 W was likely aided by their use of a Yb:YAG thin disk laser as a pump [83]. Finally, also in 2020, a group at ETH Zurich published a 2.2 μm OPCPA, generating 250 μJ , 16.5 fs pulses with an average power of 25 W and a repetition rate of 100 kHz. The system was pumped by an Innoslab Yb amplifier [84]. All three of these systems have successfully driven HHG.

1.6.2 Direct amplification of long wavelength light to drive HHG

Laser gain media with broad emission at wavelengths of 2 μm and above are limited. Nevertheless, there has been significant recent progress in developing direct amplification at these long wavelengths, with some systems already driving HHG.

Chromium and iron doped chalcogenide gain media offer broad emission in the range of 1.8-3 μm and 3-5 μm , respectively. While these materials offer broad emission bandwidth and large emission cross-sections, and can be grown in large sizes, they also have some notable drawbacks. Namely, they exhibit a generally poor quantum defect, and have short excited state lifetimes,

requiring the use of a complicated pulsed (generally Q-switched) pump laser. Furthermore, the chalcogenide host media often exhibits a large nonlinear index, increasing the likelihood of problematic self-focusing.

Two significant results using Cr:ZnSe lasers were published in 2020. A group at CREOL generated 2.5 μm pulses bearing 4 mJ of pulse energy at 44 fs (sub-6 cycle) and 1 kHz repetition rate from a single stage Cr:ZnSe amplifier [85]. The same year, a group at The Ohio State University generated 2.4 μm 7 mJ, 100 fs pulses at 1 kHz from a Cr:ZnSe laser. The pulses were used to generate high-harmonics in a gas jet. Using argon gas with a backing pressure of 2.5 bar, they reached photon energies of 280 eV; in neon, that was extended to 600 eV. Post-compression was also demonstrated, leading to 39 fs, 6.2 mJ and 115 GW peak power [86].

Holmium doped materials, such as Ho:CALGO and Ho:CaF₂, also exhibit broad emission spectra around 2 micron wavelength. These materials have been primarily exploited for ps pulse production (for instance, for pumping OPCPAs) - however, pulses around the 100 fs duration have been demonstrated from oscillators [87], and these gain media are actively being studied.

Lastly, we review some exciting results from thulium-doped fiber lasers. Thulium has broadband emission centered around 1.9 μm [88], and despite a number of potential pump absorption bands, is commonly pumped with readily available high power 793 nm diodes. This can lead to large amounts of thermal loading which can limit high average powers [89]. Nevertheless, ultrafast thulium fiber lasers have been demonstrated at kW average powers. This impressive average power can likely be attributed to a few factors. Firstly, the fiber architecture is intrinsically well suited for thermal management. Thulium doped fiber is also considered more robust to thermal effects than, for instance, ytterbium. Lastly, despite the high quantum defect, pumping at 793 nm can theoretically reach efficiencies of up to 82% by exploiting transitions which allow a single 793 nm pump photon to excite two thulium ions [90].

The Limpert group at Jena has performed pioneering work with thulium fiber lasers, including the 1 kW system described above. In a 2016 experiment, they generated 1.95 μm , 460 μJ pulses with a duration of 200 fs. The repetition rate was 61 kHz and the system delivered 28.7 W of

average power. The system comprised four fiber lasers in series, each separated by free space optics [91]. The same system was later modified to include a hollow core fiber pulse compression stage, which reduced the duration to 50 fs. When focused into an argon gas jet backed by 7 bar, the post-compressed pulses generated high-harmonics up to 160 eV [92]. A particularly clever experiment was published in 2021, in which 100 fs, 450 μJ pulses from a thulium fiber amplifier were coupled into a helium-filled antiresonant hollow core fiber. The pulses underwent self-compression as they propagated through the fiber, which had an increasing gradient of gas pressure. Near the outlet of the fiber, soft X-ray high-harmonics were generated with photon energies up to about 350 eV, with enhanced brightness due to favorable waveguide phase matching [93].

1.7 An introduction to the 3 μm wavelength OPCPA

We conclude this chapter by offering a cursory overview of the entire laser system presented in this thesis. With the exception of the oscillator, all optical sub-systems shown in the schematic (Figure 1.14) were assembled at JILA. The first light of the system comes from a 1.5 μm , 100 MHz erbium fiber oscillator (Menlo System) with two identical fiber outputs. One output is used to seed an erbium-doped fiber amplifier (EDFA) that directly pumps a highly nonlinear fiber (HNLF) to generate a dispersive wave spectral component near 1 μm . The 1 μm component is amplified in a series of two ytterbium-doped fiber amplifiers (YDFA). Before each YDFA, a chirped fiber Bragg grating stretches and spectrally filters the pulse to a narrow window around the emission spectrum of Yb:YAG at cryogenic temperatures. The YDFA chain seeds a cryo-cooled Yb:YAG regenerative amplifier through a pulse-picking Pockels cell and a free space isolator. The 1 kHz repetition rate regenerative amplifier output is split into four delay lines, each of which pumps a separate PPLN crystal in the four-stage OPA. The second output of the erbium oscillator is used to seed a second EDFA, which is carefully dispersion balanced to produce a broad spectrum with a well controlled phase. The second EDFA output is sent into a free space spatial light modulator (SLM) based pulse shaper, which applies a user programmeable spectral phase. The 1.5 μm pulse is then temporally stretched by a grism-type stretched and used to seed the collinear four-stage OPA chain, which uses

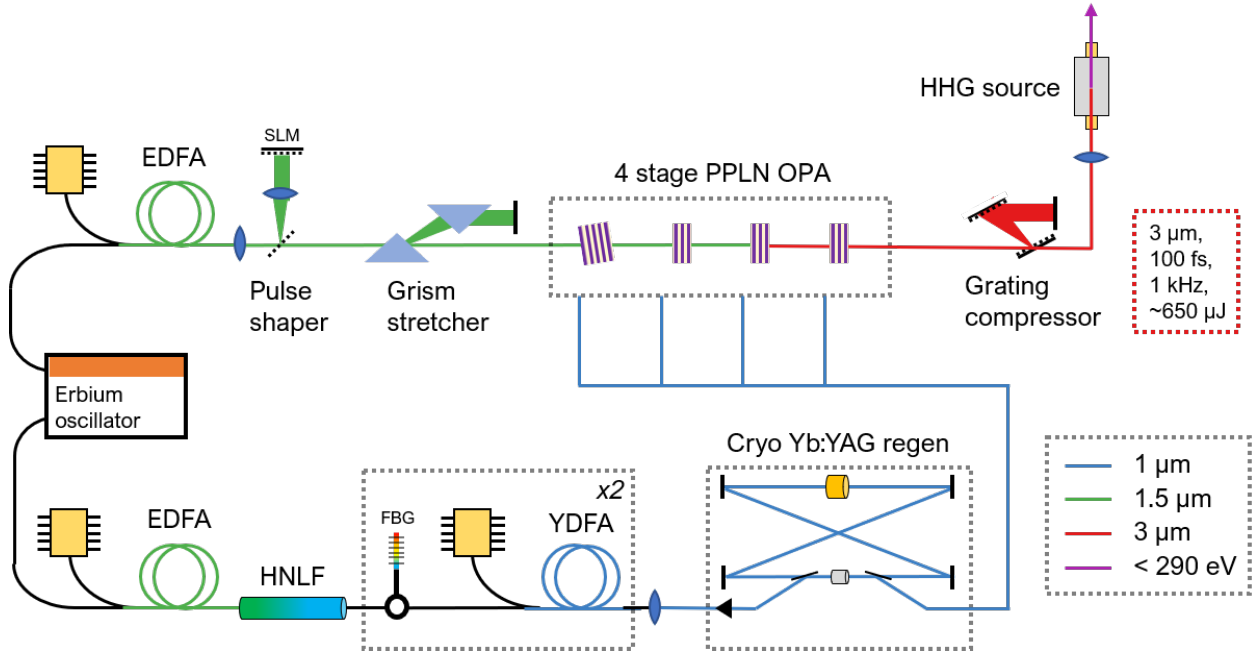


Figure 1.14: **A schematic overview of the entire laser system described in this thesis.** EDFA: erbium-doped fiber amplifier; YDFA: ytterbium-doped fiber amplifier; HNLf: highly non-linear fiber; FBG: fiber Bragg grating; SLM: spatial light modulator; PPLN: periodically poled lithium niobate. Drawing is not to scale.

periodically poled lithium niobate (PPLN) as a nonlinear crystal for each stage. The first three OPAs are seeded with 1.5 μm pulses (signal), each generating an idler at 3 μm wavelength. The final stage is seeded with 3 μm pulses to enhance the total outputted 3 μm power. The 3 μm pulses are compressed using a double-pass grating compressor, and immediately sent to the HHG source assembly. The HHG source assembly features interchangeable gas targets, including a gas jet, gas cell, and gas-filled hollow waveguide. The OPCPA in its current form produces 100 fs pulses at a 2.95 μm center wavelength with ~650 μJ of pulse energy at a 1 kHz repetition rate. At the time of writing, high harmonics have been generated in argon and nitrogen gases using a gas cell. In argon, harmonics have been observed up to the argon absorption edge at 248 eV; in nitrogen, we have observed harmonics up to ~290 eV, corresponding to the 689th harmonic of the driving laser.

As a final note, we'd like to clarify some definitions of important terms that you will encounter throughout this thesis. The quantities referred to as *peak intensity*, I_{peak} , *peak fluence*, F_{peak} , and

peak power, P_{peak} , are calculated by:

$$I_{\text{peak}} = \frac{2E}{A\tau} \quad (1.3)$$

$$F_{\text{peak}} = \frac{2E}{A} \quad (1.4)$$

$$P_{\text{peak}} = \frac{0.94E}{\tau} \quad (1.5)$$

where E is the total pulse energy, A is the area calculated from a $1/e^2$ beam radius, and τ is the FWHM pulse duration. The factors of 2 and 0.94 [94] stem from the assuming a Gaussian spatial and temporal profile, respectively. Unfortunately, the literature is replete with ambiguities as to whether the factor of two was applied, or which conventions were used for area or duration - so we'd like to set this straight for the reader at the start!

Chapter 2

Theory of parametric nonlinear optics

The earliest laboratory demonstration of nonlinear optics occurred in 1961, when Franken, Hill, Peters and Weinreich focused a 694 nm beam from the recently developed “pulsed ruby optical maser” [6] into a piece of crystalline quartz and observed second harmonic generation [95]. Their laser delivered an impressive 3 joules of energy in a 1 ms pulse. Shortly after, in 1962, optical mixing was reported [96]. Those researchers used two ruby lasers - one at room temperature and the other bathed in liquid nitrogen - to mix waves separated by just 1 nm wavelength to generate a sum-frequency wave. By 1965, a modest parametric gain of 1 dB had been demonstrated from difference frequency generation [97], establishing the technique of optical parametric amplification.

Parametric nonlinear optical effects occur when light is sufficiently intense to cause a *non-linear* relationship between the applied optical field and the material response. However, not all nonlinear relationships between measurable quantities in optics need be dependent on the electric field amplitude [98]. Non-parametric nonlinear processes often involve absorption and can be flux dependent - that is, they need not depend on the instantaneous electric field amplitude, but rather the cumulative exposure to radiation over time, as in saturable absorption. Nonparametric nonlinear optics plays an important role in nonlinear fiber optics, which is utilized in the front-end laser described in Chapter 3. For instance, Raman scattering can induce a soliton self-frequency shift of pulses propagating in optical fiber. In contrast to a parametric nonlinear process, in which the medium passively mediates the interaction, the Raman shift involves an alteration of the quantum state of the medium via phononic vibrations of the silica [99]. The remainder of this section

will exclusively consider parametric nonlinear optics, for which the nonlinear susceptibilities are real-valued.

We build our theoretical study of nonlinear optics on classical electrodynamics. We begin by deriving the electromagnetic wave equation, and then move on to the description of optical pulses in the frequency domain. While not necessarily the purview of nonlinear optics, a thorough understanding of spectral phase is foundational to ultrafast laser physics. We attempt to give the reader some intuition and mathematical foundation for the various orders of dispersion that underlie chirped pulse amplification. We then extend the wave equation to include a polarization source term that can depend nonlinearly on the driving field. We then ask the question: what is it about the optical medium that enables this nonlinear response? - and find that structure and symmetry of the medium dictate the type and magnitude of the response. We then consider second-order ($\chi^{(2)}$) nonlinear effects in detail. We derive formulae for the strength of second harmonic generation, and find that wave-vector mismatch plays a singular role in determining the spectral bandwidth of the process. We follow a similar derivation for difference frequency generation/optical parametric amplification, deriving the coupled wave equations, and from them, the equations governing small-signal gain. What we find is that the spectral dependence of the small-signal gain is quite different from the case of second harmonic generation, being determined by both wave-vector mismatch and the gain itself. This result is important for the design of short pulse OPCPAs. We then extend the coupled wave equations to include time-dependent terms - enabling a much more rigorous description of OPA at the cost of difficulty in solving the equations. We briefly describe the Sisyfos code package [100], which is used throughout this thesis to solve those equations and simulate OPA. The long pump pulses that are used in the OPCPA result in a reduced peak intensity. To accommodate this, we use quasi-phase matched (QPM) periodically poled crystals as a nonlinear medium. In this chapter, we describe QPM and how the bulk approximation can be used to describe those crystals in calculations and simulations. Finally, we close the chapter by describing third-order ($\chi^{(3)}$) effects such as self-phase modulation, which enables the broad spectrum and short pulses from the OPCPA.

2.1 Electromagnetic wave equation

We first seek to derive a wave equation to describe optical propagation in a transparent dielectric medium. We begin with Maxwell's equations, which are given here in their differential form using the convention of SI units:

$$\nabla \cdot \mathbf{D} = \rho, \quad (2.1)$$

$$\nabla \cdot \mathbf{B} = 0, \quad (2.2)$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \quad (2.3)$$

$$\nabla \times \mathbf{H} = \frac{\partial \mathbf{D}}{\partial t} + \mathbf{J} \quad (2.4)$$

We can simplify these equations by noting that a transparent dielectric medium will in general contain no free charges or currents and will be nonmagnetic. We therefore make the following assumptions:

$$\rho = 0 \quad \text{no free charges,} \quad (2.5)$$

$$\mathbf{J} = 0 \quad \text{no free currents,} \quad (2.6)$$

$$\mathbf{B} = \mu_0 \mathbf{H} \quad \text{nonmagnetic} \quad (2.7)$$

The vector \mathbf{D} is the electric displacement field, which is a superposition of the polarization vector \mathbf{P} and the electric field: $\mathbf{D} = \epsilon_0 \mathbf{E} + \mathbf{P}$. The polarization is defined as the dipole moment per unit volume. In the next chapter, we will see that the relationship between the time and frequency dependent polarization and the driving electric field will be the key to distinguishing linear and nonlinear optics.

We can readily derive the electromagnetic wave equation from Maxwell's equations and the stated assumptions. Taking the curl of the Maxwell-Faraday equation (2.3), applying the identity $\nabla \times (\nabla \times \mathbf{A}) = \nabla(\nabla \cdot \mathbf{A}) - \nabla^2 \mathbf{A}$ to the resultant left hand side of the equation, and interchanging the order of the spatial and temporal derivatives on the right hand side, gives:

$$\nabla(\nabla \cdot \mathbf{E}) - \nabla^2 \mathbf{E} = -\frac{\partial}{\partial t}(\nabla \times \mathbf{B}) \quad (2.8)$$

The assumption that the medium is nonmagnetic allows us to rewrite the right hand side of this equation as $-\mu_0 \frac{\partial}{\partial t} (\nabla \times \mathbf{H})$, which can be further reduced using Eq.'s 2.4 and 2.6:

$$\nabla(\nabla \cdot \mathbf{E}) - \nabla^2 \mathbf{E} = -\mu_0 \frac{\partial^2}{\partial t^2} \mathbf{D} \quad (2.9)$$

The first term on the left hand side is generally taken to be zero for isotropic source-free media. In the linear regime, the linear dependence of \mathbf{D} on \mathbf{E} allows us to directly apply Eq.'s 2.1 and 2.5 to disregard this term without any loss of generality. The situation is more complicated when \mathbf{P} may depend nonlinearly on \mathbf{E} , since $\nabla \cdot \mathbf{D} = 0$ does not necessarily imply $\nabla \cdot \mathbf{E} = 0$. Nevertheless, in the cases relevant to this section, we can consider the contribution of this term to be negligible: it is zero for transverse, infinite plane waves and considered negligible under the slowly varying envelope approximation [101]. We also substitute the magnetic permeability with electric permittivity, $\mu_0 = \frac{1}{\epsilon_0 c^2}$ via the relation $c = \frac{1}{\sqrt{\epsilon_0 \mu_0}}$, resulting in the following wave equation:

$$\nabla^2 \mathbf{E} - \frac{1}{\epsilon_0 c^2} \frac{\partial^2 \mathbf{D}}{\partial t^2} = 0 \quad (2.10)$$

For simplicity, we now consider this equation in the scalar approximation for a field evolving in the z-coordinate only:

$$\frac{\partial^2 E(z, t)}{\partial z^2} - \frac{1}{\epsilon_0 c^2} \frac{\partial^2 D(z, t)}{\partial t^2} = 0 \quad (2.11)$$

Equation 2.11 is applicable to both linear and nonlinear optics, depending on how the polarization P is defined, and will serve as a starting point in Chapter 2. For the remainder of the section, the displacement field $D(z, t)$ considers only the linear response of the medium, $D(z, t) = \epsilon_0 E(z, t) + P_L(z, t)$, where P_L is the linear component of the polarization.

2.2 Spectral phase of ultrafast pulses

Among the most fundamental observations from Equation 2.11 is that it admits plane wave solutions of the form

$$E(z, t) = B(z, t) \cos(\omega_0 t - k_0 z + \phi(t)), \quad (2.12)$$

where $B(z, t)$ represents a slowly varying envelope function, ω_0 is the angular frequency, k_0 is the propagation constant and $\phi(t)$ is the phase. It is somewhat more useful to represent the solution in complex notation:

$$E(z, t) = \frac{1}{2}A(z, t)e^{i(\omega_0 t - k_0 z)} + c.c. \quad (2.13)$$

$$= \Re[A(z, t)e^{i(\omega_0 t - k_0 z)}]. \quad (2.14)$$

The time-dependent phase $\phi(t)$ has been absorbed in a complex temporal amplitude $A(z, t) = B(z, t)e^{i\phi(t)}$, where $B(z, t)$ is a real-valued function. The Fourier transform of the complex time-dependent amplitude $A(z, t)$ allows us to more easily and intuitively represent pulses of arbitrary structure. We will observe the following notation for the Fourier transform: $\tilde{A}(\omega) = \mathcal{F}[A(t)]$. The Fourier transform of the complex time-dependent amplitude results in the complex frequency-dependent amplitude, which is generally broken into its real and imaginary components as the spectral amplitude and spectral phase, respectively. We therefore consider the Fourier transform of Equation 2.13, having dropped the dependence on the z -coordinate for simplicity, to solve for the complex spectral amplitude, $\tilde{E}(\omega)$:

$$\tilde{E}(\omega) = \mathcal{F}\left[\frac{1}{2}A(t)e^{i\omega_0 t} + c.c.\right] \quad (2.15)$$

$$= \frac{1}{2}\tilde{A}(\omega - \omega_0) + \frac{1}{2}(\tilde{A} * (-\omega - \omega_0)) \quad (2.16)$$

In practice, the second term of Equation 2.16 is generally dropped, leaving only the positive frequency components. It is important to note that the complex spectral amplitude contains all the information of the temporal electric field waveform itself - that is to say, it tells us all the classical information there is to know about a pulse [102]. This is of particular importance, since the complex temporal amplitude is very difficult - if not practically impossible - to measure directly in most cases, due to the extremely short time scales typically associated with field oscillations. Nevertheless, a variety of experimental techniques exist to directly sample waveforms in the time-domain, most prominently electro-optic sampling, which is commonly applied to pulses in the deep mid-infrared to terahertz regions of the spectrum, since their field oscillations can be slow compared

to the duration of accessible probe laser pulses in the lab (10's of fs) [103, 104].

We can again consider the complex spectral amplitude as a decomposition into a slowly varying envelope and a phase term: $A(\omega) = B(\omega)e^{i\phi(\omega)}$. The spectral intensity is generally readily accessible using a spectrometer, yielding $|B(\omega)|^2$. In cases where the pulse chirp is already well known,¹ the spectral amplitude alone can be informative of the temporal envelope of the pulse. Generally, however, it is necessary to also quantify the spectral phase to completely characterize the pulse. The standard technique for complete pulse characterization is known as frequency-resolved optical gating (FROG), and will be utilized throughout this thesis. The best resource for understanding FROG is Rick Trebino's book [105].

The spectral phase $\phi(\omega)$ is among the fundamental currencies of ultrafast laser physics. It is manipulated throughout chirped pulse amplification systems via stretching, compression, pulse shaping, propagation through material, and nonlinear processes such as self phase modulation, and sets the limit on achievable pulse duration. As the particularly shape of $\phi(\omega)$ can be arbitrary, it is very useful to describe the spectral phase $\phi(\omega)$ in terms of the various orders of a Taylor expansion about a central frequency ω_0 :

$$\phi(\omega) = \phi \Big|_{\omega_0} + \frac{d\phi}{d\omega} \Big|_{\omega_0} (\omega - \omega_0) + \frac{1}{2} \frac{d^2\phi}{d\omega^2} \Big|_{\omega_0} (\omega - \omega_0)^2 + \frac{1}{6} \frac{d^3\phi}{d\omega^3} \Big|_{\omega_0} (\omega - \omega_0)^3 + \frac{1}{24} \frac{d^4\phi}{d\omega^4} \Big|_{\omega_0} (\omega - \omega_0)^4 + \dots \quad (2.17)$$

The terms of the expansion of $\phi(\omega)$ are defined according to:

$$\phi_0 = \phi \Big|_{\omega_0} \quad \text{Absolute phase,} \quad (2.18)$$

$$\text{GD} = \frac{d\phi}{d\omega} \Big|_{\omega_0} \quad \text{Group delay,} \quad (2.19)$$

$$\text{GDD} = \frac{1}{2} \frac{d^2\phi}{d\omega^2} \Big|_{\omega_0} \quad \text{Group delay dispersion,} \quad (2.20)$$

$$\text{TOD} = \frac{1}{6} \frac{d^3\phi}{d\omega^3} \Big|_{\omega_0} \quad \text{Third order dispersion,} \quad (2.21)$$

$$\text{FOD} = \frac{1}{24} \frac{d^4\phi}{d\omega^4} \Big|_{\omega_0} \quad \text{Fourth order dispersion.} \quad (2.22)$$

It is important to understand how each of these dispersion terms manifest in the laser laboratory. Figure 2.1 illustrates the influence of three different spectral phases applied to the same

¹ See, for instance, the regenerative amplifier discussed in Chapter 4.

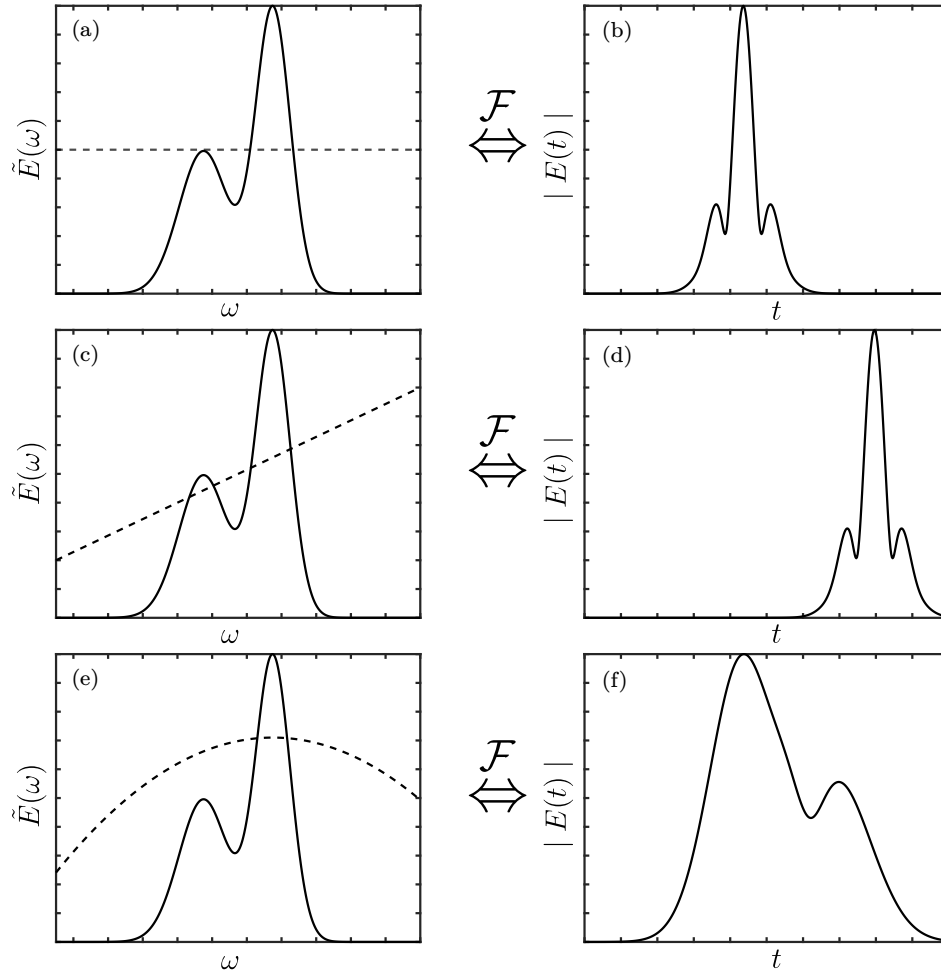


Figure 2.1: **The influence of spectral phase on the temporal envelope of a pulse.** The complex spectral amplitude (left side) is a complete description of a laser pulse and is related to the time domain electric field waveform through a Fourier transform. An arbitrary spectral amplitude is shown in the left column (solid line), with an isolated order of spectral phase shown in a dotted line: (a) absolute phase, ϕ_0 ; (c) group delay, GD; (e) group delay dispersion, GDD. The real-valued temporal amplitude envelope is shown in the right column. Note that while the complex spectral amplitude is technically double sided, since Equation 2.16 is a complex conjugate pair, all information is given by the positive frequency components. A temporal phase is also produced by the Fourier transform, but is not shown in the right column, since it is in general a less useful and intuitive quantity.

spectral amplitude on the resultant temporal pulse envelope. In the left hand column of Figure 2.1, an arbitrary real spectral amplitude $|B(\omega)|$ is plotted in a solid line, which could represent, for instance, the (square root) of the spectral density of a laser pulse. The dashed lines represent the spectral phase $\phi(\omega)$. Note that the negative frequency components of $\tilde{E}(\omega)$ are omitted for practical reasons.

2.2.1 Absolute phase

$\phi(\omega)$ which is constant with respect to frequency but is non-zero is shown in Figure 2.1(a). This represents a phase given only by the zeroth-order ϕ_0 term in Equation 2.17. It is often referred to as the *absolute phase*; despite the name, it is the relative phase of the carrier wave with respect to the pulse envelope [105], and therefore often referred to as the carrier envelope offset phase. In order to distinguish the influence of this term in the time domain, it would be necessary to resolve the actual electric field oscillations; therefore, the temporal envelope in Figure 2.1(b) appears at its *transform limit* - that is, under the slowly varying amplitude approximation for the temporal envelope, it is the least dispersed, shortest temporal pulse which can be generated from the spectrum. In cases of few-cycle pulses, the carrier-envelope offset phase can have practical implications for frequency metrology using frequency combs [106] and high-harmonic generation [107]. In HHG, electrons are tunnel ionized by individual half-cycles of the electric field; depending on the carrier envelope phase (CEP), very short pulses might have differing numbers of half cycles with sufficient intensity to cause high-harmonic generation. Furthermore the distribution of amplitude across the individual half cycles may change with the CEP. Both of these effects can manifest clearly in HHG experiments and simulations, as is shown in Figure 2.2. On the other hand, the longer ~ 10 cycle pulses used for HHG in this thesis are less effected by CEP.

2.2.2 Group delay

Group delay (GD) is a term in the spectral phase with linear dependence on frequency. GD causes a temporal shift of a pulse, which is illustrated in Figure 2.1(c) and (d). This can be

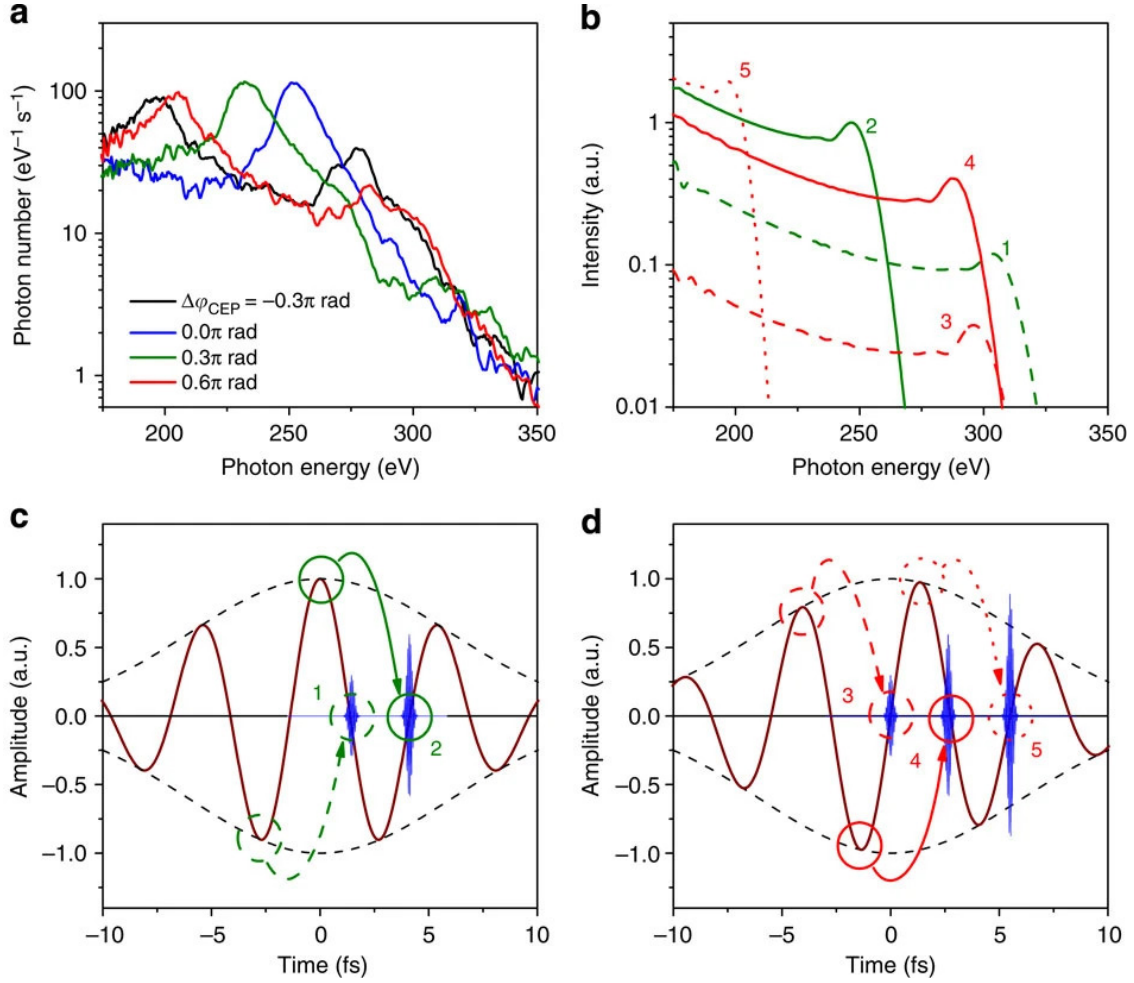


Figure 2.2: **The carrier envelope phase, or absolute phase ϕ_0 has measurable influence on high harmonic spectra in both experiment (a) and strong field approximation (SFA) based simulation (b).** In this experiment, sub 2-cycle optical pulses at 1.6 μm wavelength from an OPCPA are focused into a neon gas cell to generate high harmonics. The pulses represented in (c) and (d) each show a different offset in the phase relationship between the carrier and the envelope. Shifting the CEP of the driving pulse causes a change in the predominant trajectories which generate harmonics (shown with arrows in (c) and (d)), which in turn changes the HHG spectrum. The experiment was conducted by the Itatani group at the University of Tokyo, and the figure is from their publication [107].

understood in terms of the Fourier shift theorem [108], which states that a linear phase in one domain introduces a translation in the reciprocal domain:

$$\mathcal{F}[E(t - \tau)] = \tilde{E}(\omega)e^{-i\omega\tau} \quad (2.23)$$

2.2.3 Group delay dispersion

The way in which spectral phase influences the temporal distribution of frequency components in a pulse is best understood by considering the *instantaneous frequency*:

$$\omega_{ins} = \frac{d\phi(t)}{dt} \quad (2.24)$$

In words, a differential relationship exists between the temporal phase and the instantaneous frequency. It can be shown that the Fourier transform that relates the temporal and spectral phases preserves the functional order of the phase term [105]. For instance, a pulse with a quadratic temporal phase $\phi(t)$ also has a quadratic spectral phase $\phi(\omega)$, and so forth for other orders. As a result, we can use this differential relationship to understand how spectral phase relates to instantaneous frequency across a pulse's temporal envelope. The effect of a quadratic spectral phase (GDD) is a linear frequency ramp in time, cubic spectral phase (TOD) yields a quadratic dependence of the instantaneous frequency in time, and so forth for higher orders.

When pure GDD is applied to a pulse, the frequency components are ordered linearly based on their frequency. If the pulse contained the visible spectrum, an observer taking note of when various colors arrived in time would see the blue part of the spectrum arriving before the green part of the spectrum, followed lastly by the red part of the spectrum. The ordering in time of the colors can be reversed based on the sign of the GDD. In this way, the application of pure GDD to a pulse approximately imprints the spectrum onto the temporal envelope, which can be seen in the similarity between the shape of the envelopes in Figure 2.1(e) and (f). This effect is of particular usefulness when a pulse has a known GDD, but cannot easily be FROGed or characterized with a fast photodiode. This is the case for the few hundred picosecond pulses from the regenerative amplifier discussed in Chapter 4, which are stretched by chirped fiber Bragg gratings of known dispersion. Note that the flipping of the envelope in Figure 2.1(f) is caused by a negatively valued GDD.

2.2.4 Third and higher order dispersion

The application of pure TOD to a pulse results in a quadratic dependence of group delay vs frequency. Returning to our example of a pulse containing the full visible spectrum, an observer watching the pulse arrive in time would see the central, green part of the spectrum arriving first, since the group delay there is at a minimum. The red and blue parts of the spectrum would arrive later (or earlier depending on the sign of the TOD), but they would arrive around the same time. The presence of the two offset colors (red and blue) at the same time would cause interference, which would manifest as beats in the brightness of the pulse in the time domain. For this reason, the temporal profiles of pulses with TOD tend to exhibit oscillations before or after the pulse.

The same functional relationships described for lower order phase apply to higher order phase. In practical chirped pulse amplification systems, fourth order and higher phases are often less significant than lower order phases; nevertheless, when a truly flat phase is sought, higher order phases often need to be addressed. Most dispersive optical elements and processes introduce some amount of higher order phase. For instance, chirped mirrors are often designed to apply pure GDD to a pulse for stretching or compression, but in practice, the applied spectral phase is modulated across the spectrum due to the interferometric nature of the phase delay [109]. This modulation in $\phi(\omega)$ can only be described using higher order terms of the Taylor expansion. Another example comes from the common nonlinear process of self-phase modulation, which introduces a temporal phase $\phi(t)$ proportional to the temporal intensity envelope of the pulse [105]. Fourth and higher order dispersion terms are generally difficult to compensate with standard components in a chirped pulse amplification system; for instance, in aligning a standard grating compressor, the common approach is to change the absolute value of the GDD by adjusting the grating spacing and the ratio of GDD/TOD by adjusting the angle of incidence. These adjustments also affect higher order phase, but no independent knobs exist to compensate these higher orders. The most promising approach to compensating higher order phase is with a programmable dispersive pulse shaper, which can take various forms in both fiber and free space arrangements. One pulse shaping approach, in which a

spectrally dispersed beam is incident on a programmable spatial light modulator, will be described in Chapter 4. Before we move on to discussing nonlinear optics, it is important to note that the importance of spectral phase is weighted by the spectral amplitude. In other words, components of the spectral phase in regions where the amplitude is low contribute little to the shape of the pulse in the time domain.

2.3 Extending the wave equation to describe nonlinear optics

Up to now, we have assumed that the displacement field $D(z, t)$ depends linearly on the electric field $E(z, t)$ - that is, the medium responds linearly to the driving field. We now extend our discussion to nonlinear optics by redefining the polarization as a power series expansion:

$$P(z, t) = \epsilon_0[\chi^{(1)}E(z, t) + \chi^{(2)}E^2(z, t) + \chi^{(3)}E^3(z, t) + \dots] \quad (2.25)$$

$$\equiv P^{(1)}(z, t) + P^{(2)}(z, t) + P^{(3)}(z, t) + \dots \quad (2.26)$$

$$\equiv P^L(z, t) + P^{NL}(z, t) \quad (2.27)$$

where in the final equation, we have grouped all nonlinear contributions to the polarization into a single term P^{NL} . In Equation 2.25, $\chi^{(n)}$ represents the n^{th} order susceptibility. The majority of the parametric nonlinear optical effects described in this thesis, such as optical parametric amplification and self-phase modulation, can be accounted for by considering the second- and third-order susceptibilities, $\chi^{(2)}$ and $\chi^{(3)}$ - however, the role of higher-order Kerr effects in below threshold harmonics is an active area of research, and is considered in the context of results presented in Chapter 4. In this derivation, we treat quantities such as \mathbf{P} and \mathbf{E} as scalars for simplicity, but they are more generally vector quantities. In this case, the n^{th} order susceptibility becomes an n^{th} rank tensor.

In order to develop intuition for the intensity of laser fields needed to observe a nonlinear polarization response, it is useful to develop order of magnitude estimates of the lower order susceptibilities. The linear susceptibility $\chi^{(1)}$ is a dimensionless quantity that is on the order of unity. Dimensional analysis suggests that the units of $\chi^{(n)}$ must be $(\text{m/V})^{n-1}$. We know that $\chi^{(1)}$ is the

dominant term when the applied electric field is much smaller than the electric field binding an electron to an atom. It would be reasonable to predict that the contribution to the polarization from the higher order terms would approach that of the $\chi^{(1)}$ term when the applied field strength approaches the atomic electric field strength [101]. For the hydrogen atom, the atomic electric field strength is $E_{at} \approx 5 \times 10^{11} \text{ V/m}$.² It is therefore reasonable to expect $\chi^{(2)} \approx \frac{1}{E_{at}} \approx 2 \times 10^{-12} \text{ m/V}$, $\chi^{(3)} \approx \frac{1}{E_{at}^2} \approx 4 \times 10^{-24} \text{ m}^2/\text{V}^2$ and so forth. These estimates are correct to within a couple orders of magnitude for most solid dielectrics.

We now return to the wave equation (2.11) to understand the influence of the nonlinear polarization on the electric field - and the observable radiation from a nonlinear interaction. In our derivation of the linear wave equation, we defined the linear, scalar electric displacement field as $D^L(z, t) = \epsilon_0 E(z, t) + P^L(z, t)$. We now redefine the electric displacement field to include contributions from both the linear and nonlinear polarization:

$$D(z, t) = \epsilon_0 E(z, t) + P^L(z, t) + P^{NL}(z, t) \quad (2.28)$$

$$\equiv D^L(z, t) + P^{NL}(z, t) \quad (2.29)$$

By substituting the total electric displacement field from Equation 2.29 into the wave equation (2.11), we can now describe the evolution of the field including the nonlinear polarization as an inhomogenous differential equation:

$$\frac{\partial^2 E(z, t)}{\partial z^2} - \frac{1}{\epsilon_0 c^2} \frac{\partial^2 D^L(z, t)}{\partial t^2} = \frac{1}{\epsilon_0 c^2} \frac{\partial^2 P^{NL}(z, t)}{\partial t^2}. \quad (2.30)$$

Writing Equation 2.30 in terms of both D and E somewhat obfuscates the intuition that can be gained from the nonlinear wave equation. Therefore, we consider light propagation through an isotropic, dispersionless medium which is characterized by a unitless relative permittivity ϵ_r . In this case, the linear displacement field can be expressed as

$$D^L(z, t) = \epsilon_0 \epsilon_r E(z, t). \quad (2.31)$$

² The characteristic atomic electric field strength E_{at} equates to a very high optical intensity of $3.5 \times 10^{16} \text{ W/cm}^2$ in vacuum.

Substituting this into the more general nonlinear wave equation (2.30), we arrive at a slightly more accessible form of the wave equation:

$$\frac{\partial^2 E(z, t)}{\partial z^2} - \frac{\epsilon_r}{c^2} \frac{\partial^2 E(z, t)}{\partial t^2} = \frac{1}{\epsilon_0 c^2} \frac{\partial^2 P^{NL}(z, t)}{\partial t^2} \quad (2.32)$$

The term on the right hand-side of Equation 2.32 acts a source term for the generated electric field. More intuitively, since the term $P^{NL}(z, t)$ represents the motion of charges in the medium, the second derivative of $P^{NL}(z, t)$ with respect to time represents an acceleration of charges. We know from Larmor's result for nonrelativistic, accelerated charges that the instantaneous power radiated into an electromagnetic field is proportional to the square of the acceleration of the charge [110]. Therefore, understanding the material response to an applied field in terms of the nonlinear polarization is fundamental to understanding which frequencies get generated, amplified or attenuated in a nonlinear optical process.

2.4 The physical origin of nonlinear susceptibility

We have not yet made a case for *why* the polarization need be expanded in terms of a power series in the first place. Before delving into the plethora of parametric nonlinear effects and how we can use them in the lab, it is worthwhile to consider the microscopic origin of the nonlinear response of materials.

2.4.1 Lorentz model

The Lorentz model is perhaps the simplest and best known model to describe the optical response of bound charges. The Lorentz model considers the motion of electrons under the influence of an oscillatory applied field. The electronic motion is derived by treating the electrons as bound to the nucleus of the atom by a spring and damper. Since the proton mass is about 2×10^3 greater than the mass of the electron, the nucleus is considered fixed while the electron oscillates under the influence of an oscillatory driving field. Since the polarization is the dipole moment per unit volume, this simple model can deliver some intuition on the origin of the nonlinear time-dependent

polarization.

In a harmonic oscillator, such as a mass on a spring, the system experiences a restoring force that is linearly proportional to the displacement. Since the force is related to the potential energy via a differential relationship $F = -\frac{dU}{dx}$, the harmonic oscillator potential is quadratic. The harmonic motion is readily shown to be sinusoidal, as illustrated in Figure 2.3. The Lorentz model is effective in describing the linear optical response of dielectrics because the actual potential felt by a bound electron is well approximated as a parabola, especially for small displacements from equilibrium. Nevertheless, the actual potential felt by an electron in a condensed matter system is, in general, considerably more complicated than a parabola, and the deviations from a parabolic potential become more pronounced with greater displacement from the equilibrium position. Since most potentials are well approximated by parabolas in the small amplitude regime, weak driving fields tend to induce a primarily linear response. The harder the system is driven, the greater the effect of the nonparabolic potential, accounting for why nonlinear optical effects tend increase with the application of stronger driving fields. In Figure 2.3(a), a parabolic potential (red) is shown on top of a cartoon version of a potential that more closely approximates a real potential in a noncentrosymmetric medium, such as lithium niobate. Note that in the small amplitude regime, the parabolic and actual potential agree well. The crystal structure of lithium niobate is shown in Figure 2.3(c) as an example of a dielectric material in which the crystal structure considerably distorts the potential felt by an electron.

The distortions to the potential caused by the structure of the medium manifest in nonlinearities of the time dependent restoring force. Under this nonlinear restoring force, the motion of the electron can no longer be described by an approximately single frequency sinusoid. The driven nonlinear wave equation (2.32) shows that if a new frequency component ω appears in the time dependent polarization, it can result in the generation of electromagnetic radiation at the new frequency ω . The periodic motion associated with the distorted potential in 2.3 can be described by Fourier decomposition as the superposition of waves of multiple frequencies. In crystal nonlinear optics, these generated frequencies can include harmonics of the driving field, differences and




Figure 2.3: **The Lorentz model treats electrons as part of a spring-damper-mass system. The nonlinear polarization arises from anharmonic potentials found in real materials - especially noticeable when electrons are strongly driven.** Harmonic motion in a parabolic potential [red line, (a)], such as the up and down motion of a mass on a spring, can be well described by a single frequency sine wave [red line, (b)]. In contrast, time-dependent motion in a nonparabolic potential is described by more complicated waveforms, with Fourier decompositions including multiple frequencies (green lines). In this model, $x(t)$ is analogous to the polarization $P(z, t)$. According to the driven nonlinear wave equation (2.32), the frequency components of $P(z, t)$ are represented in the radiated electric field $E(z, t)$, accounting for the generation of frequencies distinct from the frequency of the driving field in nonlinear media. Figure inspired by [101]. (c) The crystal structure of lithium niobate, which exhibits high nonlinear susceptibility and noncentrosymmetric potentials (from Ref. [111]).

summations of driving fields, and zero-frequency DC fields.

2.4.2 The influence of centrosymmetry on nonlinear susceptibility

The structural symmetry of a nonlinear medium determines which terms are allowed in the expansion of the nonlinear polarization, and therefore which nonlinear processes will be allowed. A centrosymmetric medium refers to a medium possessing inversion symmetry. In a structure with inversion symmetry, changing the sign of a spatial coordinate does not change the response of that structure to an applied field. A common example of a centrosymmetric material is glass, possessing an amorphous structure does not depend on spatial direction. Figure 2.4(c) illustrates an example of a centrosymmetric potential, which looks the same even if the positive and negative components of x are swapped. As a consequence, the response $P(t)$ is also symmetric about $P(t) = 0$, since each direction of the electric field (up and down) must induce an identical response. Centrosymmetric materials lack even-ordered nonlinear susceptibilities, and therefore cannot facilitate $\chi^{(2)}$, three-wave mixing processes such as second harmonic generation or difference frequency generation. Intuition gained from the Lorentz model can be applied here as well: the restoring force that




Figure 2.4: **Comparison of centrosymmetric and noncentrosymmetric media and the time-dependent nonlinear polarizations they can support.** In the figure above, a single frequency field is applied to electrons inhabiting media with different structurally-dependent potentials. Why can glass support self-focusing but not second harmonic generation? Symmetry in a medium can silence even-ordered terms of the nonlinear polarization. (a) Single frequency applied electric field. (b) Electrons in a medium with a purely parabolic potential respond linearly to the applied electric field. (c) A centrosymmetric medium's potential cannot not be described with odd-ordered terms. As a consequence, the polarization lacks even-ordered susceptibilities, generating only odd-ordered harmonics. (d) A noncentrosymmetric medium's potential can be described with all terms of an expansion, allowing a nonlinear response with both even- and odd-ordered harmonics. Figure adapted from [101].

acts on bound charges is proportional to the derivative of the potential. If the expansion that describes the potential lacks odd-ordered terms, the restoring force must lack even-ordered terms. As a consequence, the polarization response also lacks even-ordered terms. This relationship is illustrated in the expansions below Figure 2.4(c).

In contrast, a noncentrosymmetric potential (Figure 2.4(d)) can be described by an expansion with even- and odd-ordered terms, and therefore, the polarization response of the medium can also be described with even- and odd-ordered susceptibilities. Noncentrosymmetric media include some (but not all) crystals, and are capable of facilitating both $\chi^{(2)}$ and $\chi^{(3)}$ effects.

2.4.3 Wave mixing

In general, an n -th order susceptibility allows for the mixing for $n+1$ waves. This result emerges from a simple characteristic of polynomial expansions. One can, of course, apply an arbitrary number of frequency components to a nonlinear material; however, the largest number of those components that can mix for any given process is determined by the order of the susceptibility. For instance, if three fields are applied to a material with a second-order susceptibility, only two of those fields will ever be multiplied in the same term of the polarization:

$$(E_1 + E_2 + E_3)^2 = E_1^2 + 2E_2E_1 + 2E_3E_1 + E_2^2 + E_3^2 + 2E_2E_3 \quad (2.33)$$

Simply adding more applied fields will not increase the number that interact in a given term. A second-order process is considered three wave mixing because the interaction of the two applied fields in a given term will, in general, create a third field.

The laser system described in this thesis exploits the large $\chi^{(2)}$ susceptibility of lithium niobate to amplify mid-IR light via a three-wave mixing process known as optical parametric amplification. It also exploits the $\chi^{(3)}$ susceptibility of fused silica glass to spectrally broaden the output of an Erbium-doped fiber amplifier seed laser via a four-wave mixing process known as self-phase modulation. The purpose of this section is to derive these processes, and others, from the nonlinear polarization.

2.5 Three-wave mixing processes ($\chi^{(2)}$)

We first consider the class of nonlinear interactions that result from the $\chi^{(2)}$, second-order nonlinearity. Consider an applied field containing two distinct frequency components, ω_1 and ω_2 . As described above, more fields could be added, but they would not change the types of nonlinear processes that occur. The total applied field is represented as

$$\tilde{E}(t) = E_1 e^{-i\omega_1 t} + E_2 e^{-i\omega_2 t} + E_1^* e^{i\omega_1 t} + E_2^* e^{i\omega_2 t}. \quad (2.34)$$

where the tilde (\sim) indicates a quantity that is varying quickly in time, and the asterisk ($*$) indicates an amplitude associated with negative frequency components.

Recall that the second order polarization has the form

$$\tilde{P}^{(2)} = \epsilon_0 \chi^{(2)} \tilde{E}(t)^2. \quad (2.35)$$

Plugging the applied field into the polarization term results in ten separate terms, but four of the terms can be clustered as a complex conjugate, *c.c.*, and ignored. The expanded second-order polarization is given by

$$\begin{aligned} \tilde{P}^{(2)} = \epsilon_0 \chi^{(2)} [& E_1^2 e^{-2i\omega_1 t} + E_2^2 e^{-2i\omega_2 t} + 2E_1 E_2 e^{-i(\omega_1 + \omega_2)t} \\ & + 2E_1 E_2^* e^{-i(\omega_1 - \omega_2)t} + c.c.] + 2\epsilon_0 \chi^{(2)} [E_1 E_1^* + E_2 E_2^*]. \end{aligned} \quad (2.36)$$

Since the nonlinear polarization is a source term in the nonlinear wave equation, each term in the above expansion represents a particular physical nonlinear process associated with a particular frequency. It is useful to separate each term of the polarization based on its associated frequency, which will be done according to the convention

$$\tilde{P}^{(2)} = \sum_n P(\omega_n) \exp^{-i\omega_n t}, \quad (2.37)$$

where $P(\omega_n)$ is an amplitude associated with a particular process. We list below the amplitudes of

each process alongside the name that describes it:

$$P(2\omega_1) = \epsilon_0 \chi^{(2)} E_1^2, \quad \text{Second-harmonic generation (SHG)} \quad (2.38)$$

$$P(2\omega_2) = \epsilon_0 \chi^{(2)} E_2^2, \quad \text{Second-harmonic generation (SHG)} \quad (2.39)$$

$$P(\omega_1 + \omega_2) = 2\epsilon_0 \chi^{(2)} E_1 E_2, \quad \text{Sum-frequency generation (SFG)} \quad (2.40)$$

$$P(\omega_1 - \omega_2) = 2\epsilon_0 \chi^{(2)} E_1 E_2^*, \quad \text{Difference-frequency generation (DFG)} \quad (2.41)$$

$$P(0) = 2\epsilon_0 \chi^{(2)} (E_1 E_1^* + E_2 E_2^*), \quad \text{Optical rectification (OR).} \quad (2.42)$$

The significance of the above expansion is that if a material has a $\chi^{(2)}$ susceptibility and is driven by an applied field containing frequencies ω_1 and ω_2 , the induced polarization will contain frequencies $2\omega_1$, $2\omega_2$, $\omega_1 + \omega_2$ and $\omega_1 - \omega_2$, in addition to the driving frequencies and a DC offset.

The strength of the $\chi^{(2)}$ response in a material is generally expressed in terms of the quantity d . In full tensor notation,

$$d_{ijk} = \frac{1}{2} \chi_{ijk}^{(2)}. \quad (2.43)$$

When allowed by symmetry arguments, the tensor d_{ijk} is generally reduced to a 3 x 6 matrix given by d_{il} . Each term of d_{il} represents the strength of the $\chi^{(2)}$ interaction for a particular polarization and propagation direction. For simplicity, d_{il} is often described by a single scalar value called the effective nonlinearity, d_{eff} . For instance, a nonlinear term frequently used for second order order interactions in β -BBO is $d_{22} = 2.2 \text{ pm/V}$, giving a second order susceptibility of $\chi^{(2)} = 2d = 4.4 \text{ pm/V}$.

We have developed a description of parametric processes at the level of a single dipole emitter. While all of these frequencies may be present in the time dependent polarization of an individual dipole, in order to calculate the evolution of the amplitude of a generated frequency component over the length of the interaction, we must also consider the spatial dependence of the fields. In doing so, we will see that each frequency component has a medium-dependent wave-vector. The macroscopic evolution of any one of these new frequencies over the length of a medium depends strongly on the wave-vector mismatch of the interacting fields (widely known as *phase-matching*). We will now

explore two three-wave mixing processes in greater depth: second-harmonic generation (SHG) and difference-frequency generation (DFG). SHG was the first (1961) experimentally observed parametric nonlinear optical process [95] and is critical for widespread photonics applications, from laser pointers to surface-spectroscopy. The mathematical approach to developing the coupled wave equations for SHG bears similarities to the slightly more complicated case of DFG, and will therefore be a useful tool in clarifying the derivation. The influence of wave-vector mismatch and the quadratic gain characteristics of SHG are qualitatively distinct from DFG, and are an important point of comparison. DFG is the more general case of optical parametric amplification (OPA), which is the fundamental mechanism of the mid-IR amplifier described in this thesis. We will derive formulas for the gain of an optical parametric amplifier, which will be used to discover important, and perhaps unexpected, insights into optimizing gain bandwidth for broadband OPCPAs.

2.5.1 Second-harmonic generation




Figure 2.5: **Second-harmonic generation displays qualitatively different behavior than difference-frequency generation.** In the undepleted pump approximation, the intensity of the second-harmonic signal grows quadratically with the intensity of the fundamental driver and the length of the medium. In the particular case of perfect wave-vector matching and no second-harmonic seed, all incident radiation can be converted into second-harmonic. In cases of imperfect wave-vector matching, the intensity of the generated second-harmonic signal oscillates, with a periodicity and maximum conversion efficiency determined by the wave-vector mismatch. Figure adapted from Reference [101].

We first seek an expression for the spatial evolution of the three waves involved in SHG, for which we follow an approach generally known as coupled-wave analysis. The polarization associated

with the second harmonic $\omega_2 = 2\omega_1$ is

$$P_{\omega_2}(z, t) = \epsilon_0 \chi^{(2)} E_1(z)^2 e^{-i\omega_2 t} \quad (2.44)$$

$$= \epsilon_0 \chi^{(2)} A_1(z)^2 e^{ik_1 z} e^{-i\omega_2 t}. \quad (2.45)$$

In the second line, we have separated the spatial dependence of the electric field into amplitude and phase components, $E_n(z) = A_n(z)e^{ik_n z}$.

We are now ready to plug the polarization into the driving term (right-hand side) of the nonlinear wave equation (2.32). Since we seek an expression governing the second harmonic, E_{ω_2} , the expression we use for the electric field on the left-hand side of 2.32 is:

$$E_2(z, t) = A_2(z)e^{i(k_2 z - \omega_2 t)} \quad (2.46)$$

Plugging $P_{\omega_2}(z, t)$ and $E_2(z, t)$ into Equation 2.32 and performing the derivatives results in:

$$\left[\frac{\partial^2 A_2}{\partial z^2} + 2ik_2 \frac{\partial A_2}{\partial z} - k_2^2 A_2 - \frac{n_2^2 \omega_2^2 A_2}{c^2} \right] e^{i(k_2 z - \omega_2 t)} = \frac{\omega_2^2}{c^2} \chi^{(2)} A_1^2 e^{2ik_1 z - i\omega_2 t}, \quad (2.47)$$

where the spatial derivatives are easily calculated using the product rule. We simplify this expression in two ways. Firstly, we note that $k_2^2 = n_2^2 \omega_2^2 / c^2$, allowing us to combine the rightmost two terms in the brackets. Secondly, we apply an important approximation called the *slowly varying envelope approximation*, or SVEA, which assumes that the spatial or temporal envelope of a pulse varies slowly in space or time compared to the wavelength or period of the carrier wave. The SVEA applied to Equation 2.47 tells us that:

$$\left| \frac{\partial^2 A_2}{\partial z^2} \right| \ll \left| k_2 \frac{\partial A_2}{\partial z} \right|. \quad (2.48)$$

Applying the two simplifications, the wave equation becomes

$$2ik_2 \frac{\partial A_2}{\partial z} = \frac{-\omega_2^2}{c^2} \chi^{(2)} A_1^2 e^{i\Delta k z}, \quad (2.49)$$

where we have introduced an important new term called the wave-vector mismatch (also known as the phase mismatch), defined here as $\Delta k = k_1 + k_1 - k_2$. The wave-vector mismatch Δk determines whether a parametric process is phase-matched, and therefore if it will result in appreciable changes

to the composition of the fields. Of course, the quantity A_1 will decrease as energy is transferred from the fundamental wave to the second-harmonic; to avoid the considerable complications of the math that this causes, we make the undepleted pump approximation, which treats A_1 as a constant.

Deriving an expression for the amplitude A_2 after interaction over some length L is now simply a matter of integrating both sides of Equation 2.49 over some length L . We express the result in terms of $|A|^2$, which is proportional to the more experimentally useful quantity of intensity [112]:

$$|A_2(L)|^2 = \frac{4\omega_1^2}{n_2^2 c^2} \left(\chi^{(2)} \right)^2 |A_1|^4 L^2 \operatorname{sinc}^2 \left(\frac{\Delta k L}{2} \right) \quad (2.50)$$

$$= \frac{8\omega_1^2 d_{eff}^2}{n_2^2 c^2} |A_1|^4 L^2 \operatorname{sinc}^2 \left(\frac{\Delta k L}{2} \right) \quad (2.51)$$

In the second line, we have substituted $\chi^{(2)} = 2 d_{\text{eff}}$.

Equation 2.51 shows that in the undepleted pump approximation, the intensity of the second-harmonic signal grows quadratically with both the intensity of the fundamental beam and the length of the medium. The $\operatorname{sinc}(\Delta k L/2)$ term is critical for understanding the influence of wave-vector mismatch, and is plotted in Figure 2.6. We find that as the term $\Delta K L/2$ reaches $+/- \pi$, the efficiency of the conversion plummets to zero, with negligible and quickly decaying resurgences as $\Delta K L/2$ continues to increase. It is common, but incorrect, to apply this same intuition of wave-vector mismatch to optical parametric amplification. As we will see in the derivation of coupled wave equations for DFG, the influence of wave-vector mismatch on the efficiency of the conversion process can be significantly altered by the gain, which is largely determined by the pump intensity and the d_{eff} of the crystal.

A more thorough coupled-wave treatment of SHG was done by Armstrong in 1962 [113], and shows that in the case of both zero wave-vector mismatch and no incident second harmonic light (no seed), the conversion of the fundamental to the second-harmonic can be complete. This is qualitatively different behavior than DFG, where even perfect phase-matching will not yield a conversion efficiency of unity. It is also worth noting that SHG does not require a seed at the second-harmonic frequency, because the second-order polarization already contains that frequency.




Figure 2.6: **The normalized intensity of the second-harmonic signal plotted against the product of wave-vector mismatch ΔK and length L .** Figure from Reference [112].

DFG, on the other hand, will not occur unseeded, since the nonlinear polarization driven by a single-frequency fundamental wave will not contain frequencies less than the fundamental. Nevertheless, frequencies less than the fundamental can occur at very low levels due to the effect of parametric fluorescence, often referred to as optical parametric generation (OPG), which arises from quantum mechanical noise. OPG can seed DFG and cause measurable amplified output, with a spectrum often determined by phase-matching conditions.

2.5.2 Difference-frequency generation

We now turn our attention to DFG, the $\chi^{(2)}$ process most relevant to this thesis, for which we closely follow the derivation described by Manzoni and Cerullo [114]. We first derive the process for the interaction of three monochromatic waves with frequencies ω_1 (signal), ω_2 (signal) and ω_3 (pump), which are all propagating in the z -direction (see Figure 2.7). Later, we will discuss how these results can readily be applied to the case of three interacting broad-band pulses. We also note that the labels *signal* and *idler* are a matter of convention, not physical difference: the signal generally refers to the frequency component that seeds the process, while the idler is typically generated.




Figure 2.7: **Difference frequency generation occurs when two pump beams generate a third beam at the difference frequency of the two pumps.** Optical parametric amplification is a particular case of DFG - generally referring to the case where one of the incident beams (pump, ω_3) is significantly stronger than the other (signal, ω_1), causing amplification of the signal and generation of an idler at the difference frequency ω_2 . The Manley-Rowe relations dictate that the creation of a signal photon must be accompanied by the creation of an idler photon and the destruction of a pump photon, as can be intuited from conservation of energy principles and the inspection of (b). Figure adapted from Reference [101].

The incident electric field including all three waves is:

$$E(z, t) = \frac{1}{2} \left[A_1(z) e^{i(\omega_1 t - k_1 z)} + A_2(z) e^{i(\omega_2 t - k_2 z)} + A_3(z) e^{i(\omega_3 t - k_3 z)} \right] + c.c. \quad (2.52)$$

We now use Equation 2.35 to calculate the second-order nonlinear polarization $P^{NL}(z, t)$,

$$P^{NL}(z, t) = \epsilon_0 d_{eff} \left[A_2^*(z) A_3(z) e^{i[\omega_1 t - (k_3 - k_2)z]} \right. \\ \left. + A_1^*(z) A_3(z) e^{i[\omega_2 t - (k_3 - k_1)z]} + A_1(z) A_2(z) e^{i[\omega_3 t - (k_3 - k_2)z]} \right] + c.c. \quad (2.53)$$

As before in our analysis of SHG, we calculate the right-hand side of the nonlinear wave equation (2.32) by taking the second derivative of Equation 2.53, which acts as the driving term for DFG.

$$P^{NL}(z, t) = \epsilon_0 d_{eff} \left[\omega_1^2 A_2^*(z) A_3(z) e^{i[\omega_1 t - (k_3 - k_2)z]} \right. \\ \left. + \omega_2^2 A_1^*(z) A_3(z) e^{i[\omega_2 t - (k_3 - k_1)z]} + \omega_3^2 A_1(z) A_2(z) e^{i[\omega_3 t - (k_3 - k_2)z]} \right] + c.c. \quad (2.54)$$

We note that the amplitude of the nonlinear polarization at one frequency is proportional to the product of the electric field amplitudes at the other two frequencies. We continue to follow the same

prescription as in our derivation of SHG and plug the electric field and the nonlinear polarization into the nonlinear wave equation. We then apply the slowly varying envelope approximation to eliminate the second derivative of the electric field amplitude with respect to z . Finally, we divide the nonlinear wave equation into three separate equations based on frequency to get the coupled wave equations:

$$\frac{\partial A_1}{\partial z} = \frac{-id_{eff}\omega_1}{cn_1} A_2^* A_3 e^{-i\Delta kz}, \quad (2.55)$$

$$\frac{\partial A_2}{\partial z} = \frac{-id_{eff}\omega_2}{cn_2} A_1^* A_3 e^{-i\Delta kz}, \quad (2.56)$$

$$\frac{\partial A_3}{\partial z} = \frac{-id_{eff}\omega_3}{cn_3} A_1 A_2 e^{i\Delta kz}, \quad (2.57)$$

where $\Delta k = k_3 - k_2 - k_1$ is the wave-vector mismatch, and n_j is the refractive index at ω_j .

2.5.3 Monochromatic optical parametric amplification

Optical parametric amplification commonly refers to a difference frequency generation process that occurs when the pump (ω_3) field is strong relative to the seed field (ω_1), and energy is transferred from the pump to the signal and idler, causing amplification. However, it should be noted that after an optical parametric amplifier reaches its maximum conversion efficiency, energy begins to flow back from the pump to the signal and idler, a process known as back-conversion. In this regimes, the generated fields are not necessarily much weaker than the pump field. Further, we have numerically simulated OPA seeded by a 3 μm field that is stronger than the 1 μm pump field, but due to the absence of 1.5 μm light in the seed, parametric amplification of the 3 μm still occurs with high efficiency (see Chapter 4).

Integrating the coupled wave equations to calculate the evolution of the actual intensity of the signal and idler fields requires the application of some boundary conditions. Here, we again use the undepleted pump approximation, which assumes that A_3 is constant. This results in a solution which is applicable to the *small signal regime*, in which the signal is small relative to the pump; nonetheless, we can derive significant understanding of the OPA process in this approximation. We also make the assumption that the idler field is zero at the beginning of the process ($A_2(z = 0) = 0$).

Manipulating the coupled wave equations under these assumptions results in a description of the signal evolution [114]:

$$\frac{\partial^2 A_1}{\partial z^2} = -i\Delta k \frac{\partial A_1}{\partial z} + \Gamma^2 A_1, \quad (2.58)$$

where

$$\Gamma^2 = \frac{2d_{eff}^2 \omega_1 \omega_2}{c^3 \epsilon_0 n_1 n_2 n_3} I_3, \quad (2.59)$$

and I_3 is the pump intensity, given by $I = n\epsilon_0 c |A|^2 / 2$. Integrating Equation 2.58 over some length L results in an expression for the signal intensity $I_1(L)$; a similar approach can be taken for the idler $I_2(L)$, resulting in:

$$I_1(L) = I_1(0) \left[1 + \left[\frac{\Gamma}{g} \sinh(gL) \right]^2 \right] \quad (2.60)$$

$$I_2(L) = I_1(0) \frac{\omega_2}{\omega_1} \left[\frac{\Gamma}{g} \sinh(gL) \right]^2, \quad (2.61)$$

where $I_j(0)$ is incident intensity at $z = 0$, and we have defined the *small-signal gain* as

$$g = \sqrt{\Gamma^2 - \frac{\Delta k^2}{4}} \quad (2.62)$$

The product of refractive indices in the denominator of the Γ^2 term has significant consequences for nonlinear crystals: high indices of refraction reduce the small signal gain. For instance, for the wavelengths considered in this thesis, the product of indices is around 6 for KTA and around 11 for LNbO₃ - for ZGP in the mid-IR, this product can be around 30! A commonly used figure of merit for characterizing the strength of a nonlinear crystal is:

$$\text{FOM} = \frac{d_{eff}^2}{n_1 n_2 n_3}. \quad (2.63)$$

Another useful quantity is the signal *parametric gain*, G , which is the ratio of the signal intensity at position $z = L$ in the crystal to the initial intensity at $z = 0$:

$$G(L) = \frac{I_1(L)}{I_1(0)} = 1 + \left[\frac{\Gamma}{g} \sinh(gL) \right]^2 \quad (2.64)$$

We will frequently refer back to Equation 2.64 to understand the scaling of the OPA process with quantities such as crystal nonlinearity and pump intensity. This equation also has significant

consequences for phase matching, which allow the amplification bandwidth to exceed the phase matching bandwidth. To gain more qualitative insight into the OPA process, we consider the large gain limit of $gL \gg 1$. We note that for large and positive values of x , the function $\sinh(x)$ behaves approximately as e^x . In this case, the parametric gain reduces to:

$$G(L) \approx \left(\frac{\Gamma}{g}\right)^2 \frac{e^{2gL}}{4} \quad (2.65)$$

Therefore, in the large gain limit, the signal and idler intensities in OPA grow exponentially over the length of the crystal. This contrasts with SHG, for which the second harmonic intensity grows quadratically over the length of the crystal. We further consider the case of perfect wave-vector matching, $\Delta k = 0$. In this case, we see that the small-signal gain reduces to $g = \Gamma$. Equation 2.65 reduces even further in this limit to:

$$G(L) \approx \frac{e^{2\Gamma L}}{4} \quad (2.66)$$

Equation 2.66 suggests that in the case of perfect wave-vector matching, large gain, and an undepleted pump, the gain of an optical parametric amplifier is exponential with a growth rate that is proportional to $\sqrt{\text{FOM}}$ and $\sqrt{I_{\text{pump}}}$. This scaling is very important for the selection of nonlinear crystals to be used in OPA, for which there is often a trade-off between the effective nonlinearity of the crystal and the damage threshold. If one wishes to maximize the power gain (as well as gain bandwidth, which will be discussed later), it is generally preferable to choose a crystal with twice the d_{eff} than twice the damage threshold. Even a small increase in Γ can result in a dramatic improvement of the gain properties of the OPA. These will be born out in comparative numerical simulations in Appendix A.

2.5.4 Why phase-matching bandwidth is not the same as amplification bandwidth in OPA

One of the more striking conclusions we can draw from inspecting the parametric gain (Equation 2.64) of the OPA process is that the dependence on wave-vector mismatch is quite different than in second-harmonic generation, often leading to confusion. In second harmonic generation, one

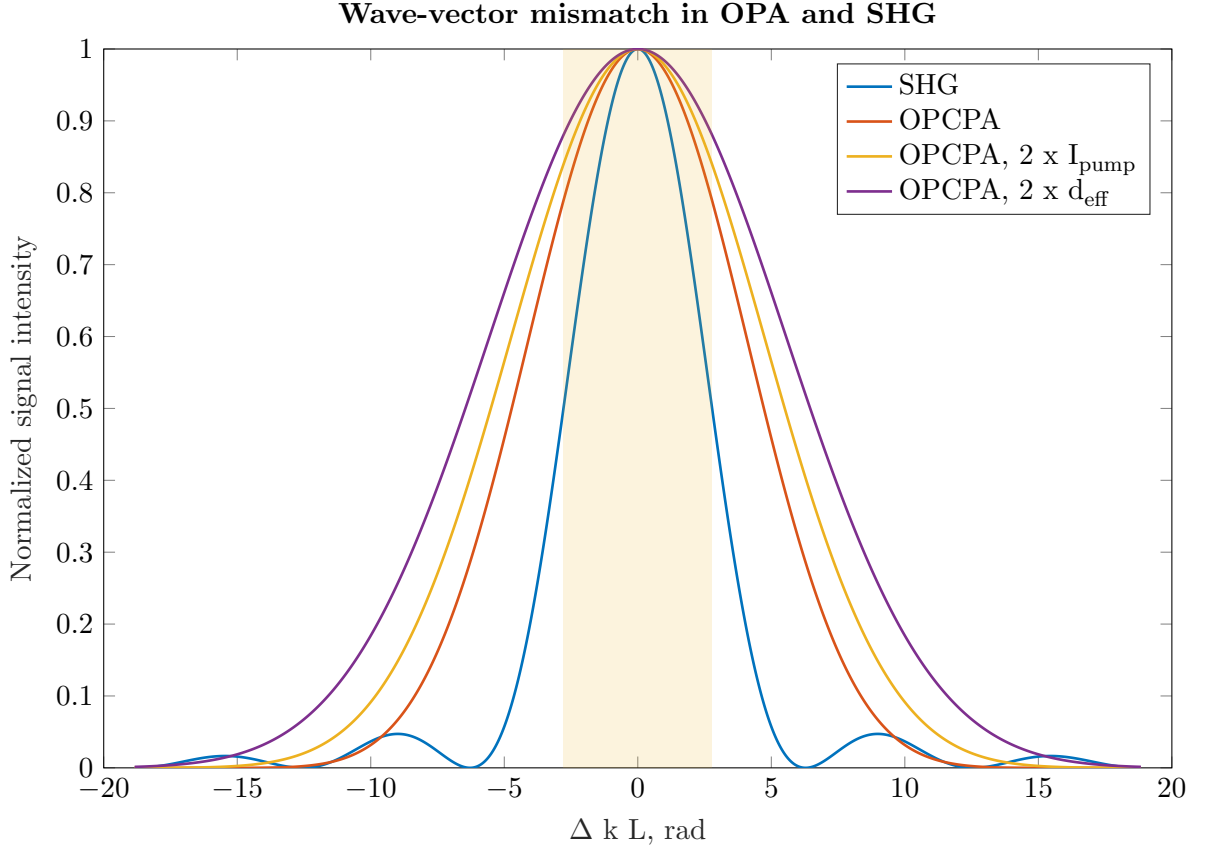


Figure 2.8: **Wave-vector mismatch (phase-matching) in optical parametric amplification, and the influence of pump intensity and effective nonlinearity, compared with second harmonic generation.** Both calculations are done in the small-signal regime. The curve labeled ‘SHG’ is a normalized plot of Equation 2.51, which gives the intensity of the second harmonic beam as a function of fundamental intensity, effective nonlinearity, crystal length L , and wave-vector mismatch Δk . The normalized SHG curve is not affected by changes in the fundamental intensity nor the effective nonlinearity; these parameters only scale the intensity. The curves labeled OPCPA plot the normalized intensity of a signal beam in an OPA process, following Equation 2.64. The parameters used in the calculation reflect the experimental conditions of the OPCPA described in this thesis: 1030 nm pump, 1550 nm signal, pump intensity = 4 GW/cm², crystal length = 3 mm, $d_{\text{eff}} = 14.8$ pm/V (for PPLN). The curve labeled ‘OPCPA’ uses these parameters exactly, while the curves labeled ‘2 x I_{pump} ’ and ‘2 x d_{eff} ’ describe the same parameters but with double the pump intensity and effective nonlinearity, respectively. It can be seen that the sensitivity of OPA to wave-vector mismatch in this regime is decreased with harder pumping and higher nonlinear coefficients, which is distinctly different than in SHG. The shaded region highlights the accumulated phase mismatch falling within the commonly defined *phase-matching bandwidth* (set by ≈ 2.78 radians). The extremes of the phase-matching bandwidth represent the point where the generated second-harmonic intensity will always reach 1/2 of its maximum value, regardless of crystal nonlinearity or pump intensity.

cannot precisely define the concept of *gain*, since it is generally assumed that the second-harmonic wave is unseeded, causing the gain $I_{2\omega}(L) / I_{2\omega}(0)$ to tend towards infinity. Nevertheless, we can consider how the intensity of a generated second harmonic signal compares to physical quantities. In Figure 2.8, we plot the normalized intensity of a second harmonic wave generated at different wave-vector mismatches, which follows a familiar *sinc*² dependence. Neither increasing the intensity of the fundamental driving beam I_2 , nor increasing the effective nonlinearity d_{eff} will change the normalized second harmonic intensity versus Δk . Similarly, the SHG bandwidth supported by a crystal is determined wholly by the Sellmeier equations, which determine Δk .

In contrast, in an OPA process, changing either the pump intensity or the effective nonlinearity *does* change the dependence of the normalized signal intensity on the wave-vector mismatch. This is illustrated in Figure 2.8, where we have overlaid the normalized signal intensity as a function of wave-vector mismatch for three different scenarios. For the OPA curves in Figure 2.8, we consider experimental parameters relevant to this thesis. The parameters are from an OPA process pumped by a 1029.5 nm laser, with a signal at 1550 nm. The length of the medium is $L=3\text{mm}$, which is a common length of the PPLN crystals used in the OPCPA. For the curve labeled OPCPA, we consider an intensity of $4 \text{ GW} / \text{cm}^2$. This corresponds to the approximate intensity associated with a 300 ps long, 200 μJ pulse focused to a $1/e^2$ diameter of 200 μm , which are the parameters associated with the first stage of the OPCPA. The value used for d_{eff} is 14.8 picometers/volt, which is the effective nonlinearity of PPLN. In Figure 2.8, we also plot the normalized signal intensity for the same conditions described above, but with either: (1) twice the pump intensity ($I_{\text{pump}} = 8 \text{ GW}/\text{cm}^2$) or (2) twice the effective nonlinearity of the crystal ($d_{\text{eff}} = 29.6 \text{ pm}/\text{volt}$). In Appendix C, we extend these calculations to explore the small signal gain as a function of wavelength, driving laser intensity, d_{eff} and crystal length.

We see that increasing pump intensity and/or effective nonlinearity makes optical parametric amplification less sensitive to wave-vector mismatch. Furthermore, we see that the particular wave-vector mismatch values of $\pm 2\pi$, which effectively bookend the efficiency of second-harmonic generation, have no particular meaning in optical parametric amplification. A commonly used term

is *phase-matching bandwidth*. This is commonly defined as bandwidth for which the phase-mismatch ΔkL , accumulated over the length of the interaction, varies by 2.78 radians. The reason for this definition is that the conversion efficiency curve for second-harmonic generation (which follows as sinc^2 curve) drops to half of its maximum value when the phase mismatch reaches approximately 2.78 radians. This is illustrated in Figure 2.8, where the phase-matching bandwidth (as commonly defined) is shaded and overlaid. While the commonly defined phase-matching bandwidth is clearly a useful quantity for second harmonic generation, it is not particularly useful in optical parametric amplification. Rather, OPA amplification bandwidth is gain dependent, and must take into account not just the dispersion of the crystal, but also the pump intensity and the effective nonlinearity. Intuitive understanding of this phenomenon is relatively straightforward. In the absence of gain, when the medium length exceeds the coherence length, the complex exponential term in the relevant coupled wave equation changes sign. This results in a sign change of the derivative $\frac{\partial A}{\partial z}$, resulting in reconversion of the generated wave amplitude back to pump. However, in the case of a high gain OPA, the growth of the signal amplitude can actually out-compete the reconversion process [98]. It is important to note that this behavior fundamentally depends on the existence of gain, and is not applicable to conversion processes such as SHG or HHG. Confusion in defining the relationship between wave-vector mismatch and the gain bandwidth for OPAs exists in the literature [115].

2.5.5 Extending the OPA coupled wave equations to the time-domain: the Sisyfos numerical simulation package

The coupled wave equations derived in the monochromatic case described the evolution of a field amplitude $A(z)$ at a single frequency. In order to describe a temporal pulse envelope, we must introduce a time-dependent amplitude $A(z, t)$. This considerably complicates solving the nonlinear wave equation, because we can no longer disregard the temporal derivatives of the amplitude coefficients. If we choose a frame of reference moving at the group velocity of the pump pulse v_{g3} ,

the equations can be expressed in terms of $\tau = t - z/v_{g3}$, resulting in coupled nonlinear equations:

$$\frac{\partial A_1}{\partial z} + \frac{1}{2i}GVD_1 \frac{\partial^2 A_1}{\partial \tau^2} + \delta_{13} \frac{\partial A_1}{\partial \tau} = -i\sigma_1 A_2^* A_3 e^{-i\Delta kz}, \quad (2.67)$$

$$\frac{\partial A_2}{\partial z} + \frac{1}{2i}GVD_2 \frac{\partial^2 A_2}{\partial \tau^2} + \delta_{23} \frac{\partial A_2}{\partial \tau} = -i\sigma_2 A_1^* A_3 e^{-i\Delta kz}, \quad (2.68)$$

$$\frac{\partial A_3}{\partial z} + \frac{1}{2i}GVD_3 \frac{\partial^2 A_3}{\partial \tau^2} = -i\sigma_3 A_1 A_2 e^{i\Delta kz}. \quad (2.69)$$

where $\delta_{ij} = 1/v_{gi} - 1/v_{gj}$ is the group velocity mismatch between two pulses, GVD_i is the group-velocity dispersion, and $\sigma_i = d_{eff}\omega_i/c_0 n_i$ [114].

An even more complete analysis of the differential equations governing DFG in a single transverse direction, which includes both pump depletion and diffraction effects, was done in 1989 by Nieto-Vesperinas and Lera [116]. Later, the solutions were expanded to include all three spatial dimensions by Gunnar Arisholm [100]. Arisholm incorporated these solutions into a powerful simulation model known as SISYFOS (Simulation System for Optical Science), which includes a nearly exhaustive set of physical effects. Sisyfos consists of a C++ library that contains the actual simulation code and a library of MatLab/Python functions for interacting with the simulation. Users are not generally able to modify the C++ code, but can readily change predefined parameters, such as crystal type, propagation directions, etc. The electric field in the two transverse coordinates and in time can be extracted at a user-defined plane, from which MatLab or Python functions can be used to extract useful quantities, such as spectra and spatial modes. Beam propagation can be handled by a split-step method or by solving the coupled equations for the scalar mode amplitudes in spatial or temporal Fourier space, where each mode represents a monochromatic plane wave component. The Fourier space method is chosen by default. Noncollinear OPA can be handled without excessively high spatial resolution by factoring out the transverse k-vector components.

The Sisyfos package was utilized heavily in this thesis to simulate various nonlinear amplification schemes as well as to understand the behavior of various nonlinear crystals. In Chapter 4, we present Sisyfos simulations of PPLN OPCPA. In Appendix A, we offer simulations of alternative 1 μm pumped $\chi^{(2)}$ crystals, including quasi-phase matched PPLT and PPKTP, as well as bulk crystals, including KTA, LN and KN.

2.5.6 Quasi-phase matching in OPA

Evaluation of the wave-vector mismatch depends on the chosen phase matching technique. In general, the process of phase-matching involves minimizing the quantity Δk for a particular OPA process. The refractive index of a lossless medium exhibits a property known as normal dispersion,³ which causes the refractive index to increase monotonically with frequency. For this reason, it is generally not possible for two (or three) waves of different frequency and the same polarization to travel with the same phase velocity. The difference in refractive index for the pump, signal and idler waves at a given polarization in a typical OPA configuration is so large that, even in the narrowband case, significant parametric gain would be impossible due to an overwhelmingly large Δk . As an example, in the top plot of Figure 2.10, we plot Δk for three waves with the same polarization (at 1, 1.5 and 3 μm), and find wave-vector mismatches of greater than 600 radians after a 3 mm lithium niobate crystal - Figure 2.8 indicated that significant gain requires magnitudes of less than 10 or 15 radians. To solve this problem, there are two common approaches to matching the wave-vectors in an OPA process: birefringent phase-matching and quasi-phase matching.

In a birefringent medium, the index of refraction depends on the polarization direction. We can generally simplify the situation and consider only two polarization, which correspond to the ordinary and extraordinary refractive index. In birefringent phase matching, we choose the polarization of the pump, signal and idler to make the real part of index of refraction experienced by each of the three waves as similar as possible. In contrast, in quasi-phase matching (QPM), the polarization of the beams is chosen to maximize the effective nonlinearity, generally resulting in all three beams having the same polarization. Under normal circumstances, Δk would become large over a very short interaction length, preventing significant amplification (Figure 2.10(top)). In QPM, the crystal is fabricated in such a way that one of the crystal axes changes sign in a periodic fashion, causing the sign of d_{eff} to also switch in a periodic fashion, as illustrated in Figure

³ Confusion often arises over this terminology. The refractive index relates the phase velocity to the speed of light, and generally exhibits normal dispersion in lossless media. In contrast, the group index, which is the ratio of the group velocity and the speed of light, can exhibit either normal or anomalous dispersion. Fused silica famously switches between the two at around 1.3 μm wavelength.





Figure 2.9: **Quasi-phase matching of nonlinear optical crystals.** Left: a cartoon of OPA in a QPM crystal, showing the periodic modulation of the sign of d_{eff} (from Reference [114]). Right: the evolution of the field amplitude of a second harmonic generated wave in a nonlinear crystal, with (a) perfect phase matching - showing linear field growth/quadratic intensity growth; (b) QPM - with a monotonic but slightly sub-optimal growth; (c) a wave vector mismatch - showing the amplitude oscillate as energy flows back and forth between the pump and signal/idler waves. We note that the linear field amplitude growth is not representative of OPA, which exhibits exponential field amplitude growth in the small-signal regime with perfect phase matching. From Reference [101].

2.9. If the modulation period is chosen to be twice the coherent build-up length, the effective value of Δk can be dramatically reduced, causing amplification over a much larger interaction length.

We define the modulation period as Λ and the corresponding grating wave-vector as $K_g = 2\pi/\Lambda$. In QPM, the wave-vector mismatch is defined as:

$$\Delta k = k_3 - k_2 - k_1 - K_g \quad (2.70)$$

where k_1 , k_2 and k_3 refer to the wave-vectors of the signal, idler and pump, defined as $k_i = 2\pi n_i/\lambda_i$.

In practice, the particular value of the modulation period Λ is chosen such that the grating wave-vector cancels the wave-vector mismatch of the central frequencies of the pump, signal, and idler - in other words, K_g is chosen to make $\Delta k = 0$ for the central frequencies and the central frequencies only. For crystals with a single grating period, perfect wave-vector matching can therefore only be achieved for a single set of central-frequency pump, signal and idler wave-vectors. For broadband pulses, we can easily use Equation 2.70 to calculate a particular $\Delta k(\lambda)$ value for all

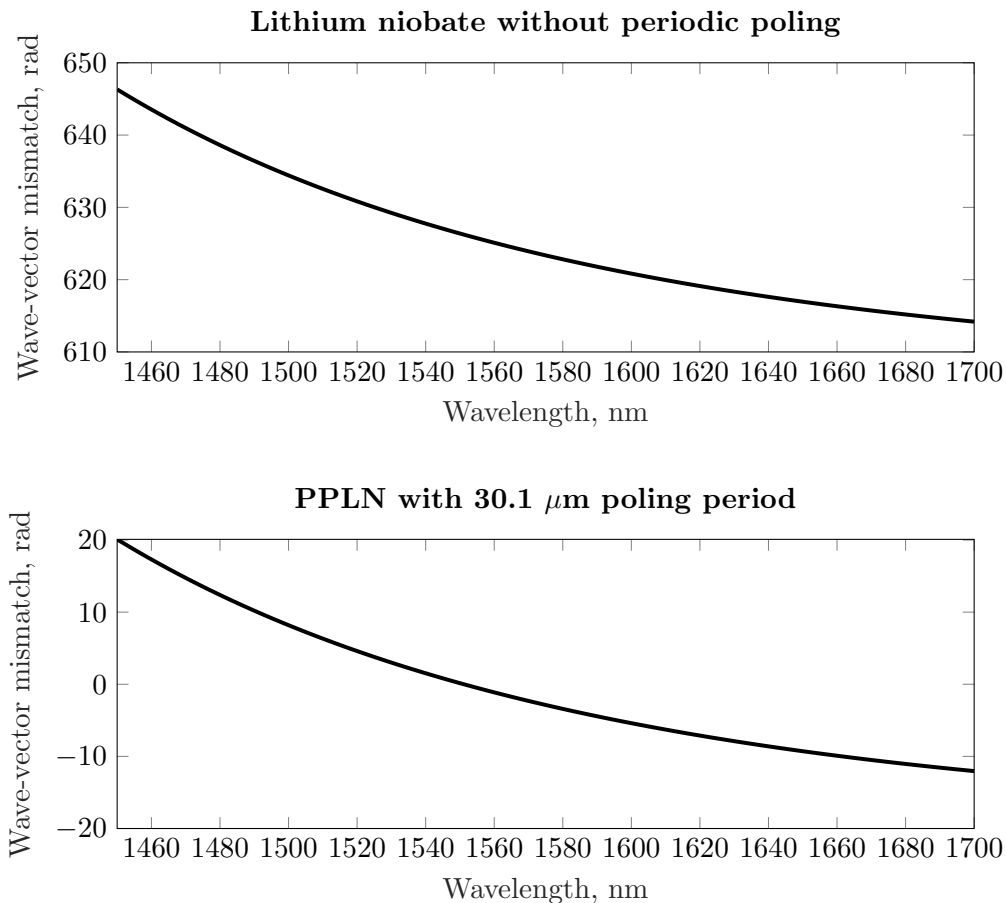


Figure 2.10: **The influence of quasi-phase matching on wave-vector mismatch for a 3 mm PPLN crystal.** After propagating through a 3 mm long PPLN at 100°C, the signal, idler and pump accumulate a total wave-vector mismatch, shown above. The top plot illustrates the calculated wave-vector mismatch for non-periodically poled lithium niobate, by removing the K_g term in the calculation of Δk . The bottom plot shows the wave-vector mismatch in the case of quasi-phase matching, with the $K_{grating}$ term included for a 30.1 μm poling period. The frequency dependent wave vectors were calculated using Sellmeier equations from the erratum of reference [117]. All three beams have extraordinary polarization, which is necessary to access the highest nonlinear coefficient d_{33} .

wavelength components present in a generated pulse, so long as the frequency dependent refractive index is known for the pump, signal and idler. The result of this calculation for a 3 mm PPLN is shown in Figure 2.10(bottom). In the context of Sisyfos simulations, this condition will be referred to as the *bulk approximation* for quasi-phase matching, and will allow us to simulate propagation through periodically poled media without discretizing the medium into separately poled domains.

Reference [118] investigated the bulk approximation and found that it is effective unless the number of periods is very small. However, we caution that parasitic effects, such as SHG or SFG can often incidentally be phase-matched in periodically poled crystals - as happens in our experiment. This behavior is not accurately captured by the bulk approximation, nor are aperiodically poled crystals. Simulations aside, the bulk approximation applies to the analytical calculation of spectral parametric small-signal gain as well, which is then a simple matter of substituting $\Delta k(\lambda)$ into the equation for parametric gain (see Appendix C).

2.6 Four-wave mixing processes ($\chi^{(3)}$)

Four-wave mixing processes arise from the third term in the expansion of the nonlinear polarization:

$$\tilde{P}^{(3)} = \epsilon_0 \chi^{(3)} \tilde{E}(t)^3. \quad (2.71)$$

We can understand all four-wave mixing processes by calculating the third-order polarization induced by a field containing three frequency components: ω_1 , ω_2 and ω_3 . This results in a lengthy expression for the nonlinear polarization that contains 22 distinct positive-frequency components. For simplicity, we briefly turn our attention to two important effects: third-harmonic generation (THG) and self-phase modulation (SPM).

For mathematical convenience, we consider a real, monochromatic electric field of the form

$$\tilde{E}(t) = E \cos(\omega t) \quad (2.72)$$

Through the use of a common identity,⁴ we cube the electric field and calculate the third-order polarization from Equation 2.71:

$$\tilde{P}^{(3)}(t) = \frac{1}{4} \epsilon_0 \chi^{(3)} E^3 \cos(3\omega t) + \frac{3}{4} \epsilon_0 \chi^{(3)} E^3 \cos(\omega t) \quad (2.73)$$

The first term on the right-hand side of Equation 2.73 represents third-harmonic generation (THG), which is illustrated in Figure 2.11. Due to the observation that $\chi^{(3)}$ effects are generally less

⁴ $\cos^3(\omega t) = \frac{1}{4} \cos(3\omega t) + \frac{3}{4} \cos(\omega t)$

pronounced than $\chi^{(2)}$ effects, this effect is less commonly employed in frequency conversion devices than $\chi^{(2)}$ effects. An illustrative example of this is the National Ignition Facility in Livermore, CA, which produces megajoule-energy laser pulses at the 351 nm third harmonic of a 1053 nm driving laser. Instead of utilizing direct third harmonic generation, it is more efficient to utilize cascaded $\chi^{(2)}$ effects in KDP crystals: the second harmonic is generated in the first stage, and $3\omega = \omega + 2\omega$ sum-frequency generation is used in the second stage [119].




Figure 2.11: **(a) Third-harmonic generation (THG) can occur when a sufficiently intense driving field is incident on a medium with a $\chi^{(3)}$ nonlinear susceptibility.** (b) An energy-level representation of THG, where the dashed lines represent virtual states. Figure from Ref. [101].

The second term on the right-hand side of Equation 2.73 causes an intensity-dependent refractive index. In order to understand this effect, we should first consider the origin of the refractive index in the linear regime. The first term of the expansion of the polarization is $P^{(1)} = \epsilon_0 \chi^{(1)} E$, where $\chi^{(1)}$ is a unitless quantity that can readily be related to the refractive index of the material according to

$$n^2 = 1 + \chi^{(1)}, \quad (2.74)$$

where we note that n^2 is also equal to the relative permittivity $\epsilon_r = \frac{\epsilon}{\epsilon_0}$.

By comparing expressions for the first and third order nonlinear polarization, we can develop a clear physical picture as to how the refractive index can be intensity dependent in $\chi^{(3)}$ materials. In general, the refractive index can be thought of as a term which relates the polarization at a frequency ω to an incident field also at frequency ω . We have seen that while the second-order polarization

generates harmonics and other frequencies, it does not generate new light at the fundamental (driving) frequency. Thus, for a polarization containing only $\chi^{(1)}$ and $\chi^{(2)}$ susceptibilities, the refractive index is governed entirely by $\chi^{(1)}$, and the refractive index is not intensity dependent.

When a $\chi^{(3)}$ susceptibility is included in the polarization, a new contribution to the polarization is made at frequency ω :

$$P(\omega) = \epsilon_0\chi^{(1)}E(\omega) + 3\epsilon_0\chi^{(3)}E(\omega)^3 \quad (2.75)$$

$$= \epsilon_0\chi^{(1)} + 3\chi^{(3)}|E(\omega)|^2E(\omega) \quad (2.76)$$

$$= \epsilon_0\chi_{\text{eff}}E(\omega) \quad (2.77)$$

We can calculate the refractive index as before, but now we must use χ_{eff} instead of $\chi^{(1)}$:

$$n^2 = 1 + \chi_{\text{eff}} \quad (2.78)$$

where χ_{eff} is dependent on the intensity of the incident field, $|E(\omega)|^2$. A more common, and practically useful, representation of the nonlinearity of the refractive index is

$$n = n_0 + n_2I \quad (2.79)$$

where n_0 is the commonly encountered linear refractive index, and n_2 is the nonlinear index given by

$$n_2 = \frac{3}{4n_0^2\epsilon_0c}\chi^{(3)}. \quad (2.80)$$

The strength of $\chi^{(3)}$ nonlinearity varies tremendously for different materials. For instance, $\chi^{(3)}$ of silicon is $2.8 \times 10^{-18} \text{ m}^2/\text{V}^2$, while it is as low as $6.2 \times 10^{-23} \text{ m}^2/\text{V}^2$ for lithium fluoride [101]. In practice, the magnitude of $\chi^{(3)}$ in a material is generally expressed in terms of its nonlinear index, n_2 . This is likely a product of convenience: many $\chi^{(3)}$ effects from a pseudo-monochromatic driving field, such as self-focusing and self-phase modulation, can be intuitively understood by an intensity dependent refractive index, also known as the Kerr effect.

The intensity dependence of the refractive index in materials with a positive and non-zero n_2 can cause light to self focus as it propagates through the medium. Since the intensity of a




Figure 2.12: **Self-phase modulation is a third-order nonlinear process that causes changes in the frequency spectrum of a pulse.** In SPM, the most intense part of the pulse generally experiences a higher refractive index than the tails, inducing a phase shift between the different parts of the pulse. Here, spectra are shown of pulses that have experienced various amounts of SPM. The SPM is quantified by the maximum phase displacement at the peak of the pulse. Figure from Ref. [120].

beam is typically higher at the center of the spatial mode, the refractive index felt by the beam decreases radially, causing a focusing action. A similar effect occurs in the time domain for short pulses, and is known as self-phase modulation, or SPM. In a simple example of SPM, the tails of the temporal pulse envelope may experience a lower refractive index than the peak of the pulse. The change of the temporal phase in time manifests as new instantaneous frequencies, as described by 2.24. The formation of new instantaneous frequencies can either broaden or compress the spectrum of the pulse depending on the initial chirp of the pulse. This effect is best understood by considering the Fourier decomposition of the pulse into frequency components with complex spectral amplitudes each with a defined spectral phase. Depending on the phase relationship between the added frequency components and the initial frequency components, the spectrum can either broaden, compress or be unaffected. The classic spectral broadening most commonly associated with SPM is illustrated in Figure 2.12, and occurs when the incident pulse has an

up-chirp. Spectral broadening in this configuration can be understood intuitively by considering that in a self-phase modulated pulse, the same instantaneous frequency can occur at two different temporal positions in the pulse, leading to the interference spectra shown in Figure 2.12 [121]. SPM is widely exploited in ultrafast optics because it enables the generation of broad spectral bandwidths that would otherwise be impossible due to limitations such as gain-medium emission bandwidth and gain narrowing. The spectral broadening capabilities of SPM are fundamental to the fiber front-end laser described in Chapter 3, allowing the generation of 35 fs pulses directly from an Erbium-doped fiber amplifier.

Chapter 3

All-fiber front-end laser to generate 1 and 1.5 μm seed pulses

3.1 Introduction

The *front-end* of a laser system generally describes the component of the system which generates the first light, and often does preliminary reshaping and amplification of that first light to optimize it as a seed for later stages of amplification. The primary purpose of the front-end described in this thesis is to produce two synchronized pulses of light: one with a wavelength centered near 1.5 μm to seed optical parametric amplification, and the other near 1 μm to seed the regenerative amplifier which generates pulses that pump the OPA process.

All other components of the optical system described in this thesis depend on the front-end. For this reason, in designing and building the front-end, we placed particular emphasis on robustness, modularity, serviceability and stability of all key metrics, including average power, pulse energy, beam pointing, spatial mode and spectrum. Additional design considerations included the delivery of: sufficient pulse energy to minimize the noise associated with subsequent amplification




Figure 3.1: **A simplified view of the all-fiber front-end laser.** One arm delivers a linearly chirped, approximately 2 ns pulse which seeds the regenerative amplifier. The other arm delivers a 35 fs pulse centered around 1.5 μm wavelength, which acts as a seed for the OPA stages.

stages; sufficiently large bandwidth to sustain ultrashort pulses from the OPA; well-characterized and compensable spectral phase; and near-Gaussian beam profiles.

An all-fiber architecture has been found to be particularly effective in meeting all of these design criteria. The key limitation associated with fiber amplifiers is low-peak power, which is generally not problematic for the generally low power front-end. On the other hand, the very high single pass gain, alignment-free operation, compactness, robustness and stability of fiber amplifiers has made the technology an ideal choice for the front-end. The low electrical power consumption of the front-end allows us to power the entire system by a rack-mountable, double-conversion, ‘always on’ uninterrupted power supply (part SU1500RTXLCD2U from Tripp Lite). In the absence of building power, the front-end can run for hours on battery without any change in performance, and is completely isolated from building power quality issues that can damage sensitive electronics.

We gratefully acknowledge Scott Diddams and his group at NIST / CU Boulder for sharing designs and technical expertise used in the construction of the dispersion managed EDFAs, as well as direct assistance in the challenging splice between the highly nonlinear fiber and the transport fiber. We further acknowledge Tin Nguyen for his direct assistance with splicing the EDFAs before we acquired our own splicer. A conference proceeding describing the front-end can be found in Reference [122].

3.1.1 Overview of the all-fiber front-end

The front-end begins with a 100 MHz Erbium fiber oscillator (ELMO, Menlo Systems GmbH) customized to have two identical fiber output ports, each producing about 3 mW of average power. The outputs of the oscillator are connectorized with SC/APC connectors; adapters are used to connect them with subsequent SC/APC connectors, which are used in all other connections. Each of the oscillator output ports seed a homemade Erbium-doped fiber amplifier (EDFA). The two EDFAs share a nearly identical design and both produce around 200 mW of average power and sub-50 fs. The two EDFAs differ slightly in the length of transport fiber before and after the amplifier, and the currents used to drive the pumps. The EDFA in the 1.5 μm arm is optimized for

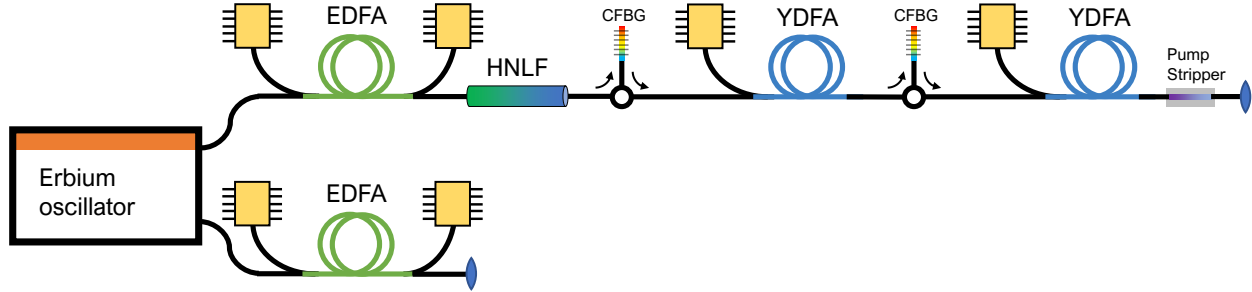


Figure 3.2: **A more detailed schematic of the front-end laser.** EDFA: erbium-doped fiber amplifier; YDFA: ytterbium-doped fiber amplifier; HNLF: highly-nonlinear fiber; CFBG: chirped fiber Bragg grating.

broad bandwidth and flat spectral phase in order to deliver a good seed for the OPCPA. Its output is connectorized with an FC/APC connector and is connected to a fiber collimator, which launches a free-space beam towards the pulse shaper, pulse stretcher and finally the first OPA crystal.

	1.5 μm branch	1 μm branch
Central λ	1550 nm	1029.5 nm
Bandwidth (tail-tail)	200 nm	1.2 nm
Pulse duration	35 fs	2 ns
Spectral phase	flat	1.5 ns/nm linear chirp
Pulse energy	2.2 nJ	20 nJ
Repetition rate	100 MHz	100 MHz
Average power	220 mW	2 W
Peak power	49 kW	10 W

Table 3.1: **A summary of the output parameters for each arm of the front-end.** Values are for day-to-day operation at the time of writing.

The output of the EDFA in the 1 μm arm is connected to the highly nonlinear fiber (HNLF) assembly via FC/APC connectors. The HNLF assembly consists of a short segment of PM1550 fiber, spliced to an ~ 3 cm segment of HNLF, spliced to a short segment of FC/APC connectorized

PM980 fiber. A strong third-order nonlinear response of the HNLF generates a dispersive wave with spectral amplitude around a 1 μm wavelength. We typically observe a few mW of optical power in the ~ 4 nm wide spectral band around 1030 nm. The entire spectrum from the HNLF is fed into a circulator, which first directs the light to a chirped fiber Bragg grating (CFBG). The CFBG has two purposes: it acts as a spectral notch filter, transmitting a narrow band around 1029.5 nm, and it applies a linear chirp to the pulse (quadratic spectral phase). The stretched and spectrally narrowed 1 μm component seeds a ytterbium-doped fiber amplifier which amplifies the signal to about 60 mW. Following the first YDFA is another identical circulator and CFBG, which further stretches the signal and cuts out any ASE from the first YDFA. A second YDFA amplifies the signal further. After the second YDFA, a pump stripper removes residual pump light, and sends the approximately 2.3 W signal to a high-power free space collimator. The free-space beam is then sent to a Pockels cell to reduce the repetition rate from 100 MHz to 1 kHz, and eventually seeds the regenerative amplifier. A schematic of the front-end laser is shown in Figure 3.2. The output parameters of both branches of the front-end are summarized in Table 3.1.

A photograph of the front-end laser is shown in Figure 3.3. The laser system is made more compact by stacking four different levels, with a 12" x 18" footprint on the optical table. The bottom level contains the oscillator and the EDFA that pumps the HNLF. The second level contains the HNLF assembly, the first YDFA, and the pump for the second YDFA. The third level was added to accommodate large, home-built aluminum spools which contain the ytterbium gain fiber of the second YDFA and the collimator pigtail. The motivation for these larger spools will be detailed later in this chapter. The top level contains the EDFA for the 1.5 μm arm. Under normal operation, the top level is covered with a clear acrylic cover to prevent falling objects from damaging exposed fibers, as well as to mitigate dust (cover not shown in the photo).

All of the electronics associated with the front-end, including power supplies, laser diode controllers, and thermo-electric cooler controllers are housed in a single standard 19" wide electronics rack which is shared with the regenerative amplifier.




Figure 3.3: **A photo of the front-end laser.** The front-end laser is composed of a stack of four aluminum breadboards, and has a compact 12" x 18" footprint.

3.1.2 Optical fibers in laser amplifiers

In Chapter 1, we reviewed the history of fiber lasers up to the development of the EDFA in the late-1980's, and illustrated the advantage of fiber-based gain media in high average power handling. In this section, we will explore some important aspects of fiber laser technology as it pertains to the front-end laser. Fiber lasers are a sub-category of solid-state lasers, and are distinguished from free-space lasers in that the gain medium is contained in an optical fiber, typically in the form of a rare-earth doped glass. Optical fibers typically guide light through a radial modulation in the refractive index. Generally, an optical fiber waveguide consists of a slightly higher index core surrounded by a slightly lower index cladding. This allows for the possibility of total internal reflection at the interface between the core and cladding, so long as the incident angle is shallow enough for a given index gradient. Several types of optical fibers are illustrated in Figure 3.4.

Optical fibers may be designed to support a single-mode (SM) or multiple modes (MM). The seeding of free-space laser amplifiers requires high mode quality and modal-stability; therefore, all

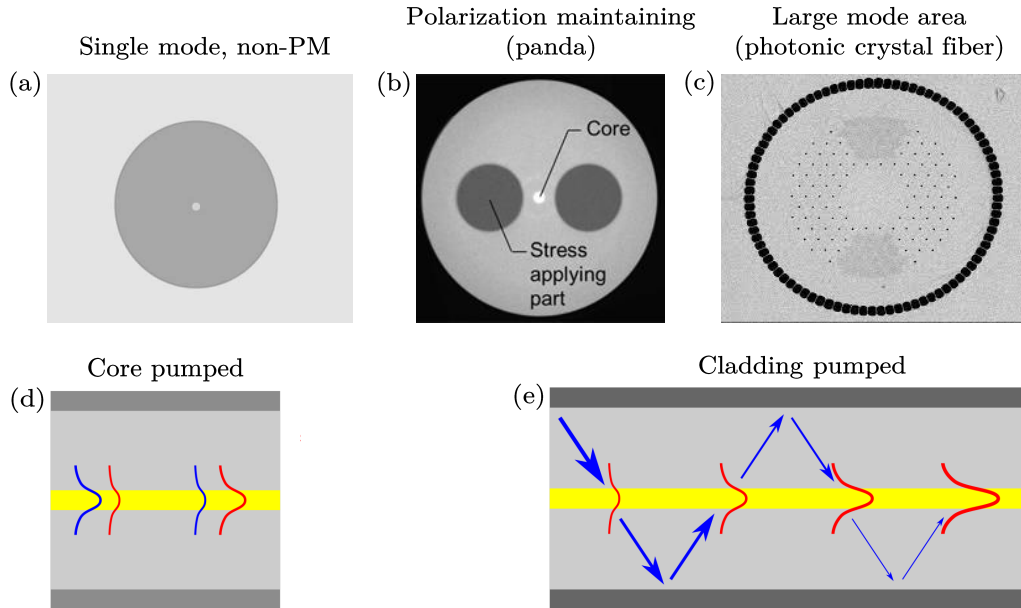


Figure 3.4: **An overview of optical fiber types used in laser amplifiers.** (a) a single mode, non-PM fiber [123]. (b) A polarization maintaining, panda-style single mode fiber using birefringence inducing stress rods [124]. (c) An 85 μm single-mode, ytterbium doped core large mode area photonic crystal fiber, from the NKT Photonics aeroGAIN rod data sheet. (d),(e) An illustration of the difference between core pumping and cladding pumping, with the blue lines represent the pump and red lines representing the signal. Adapted from Reference [125].

fibers in the front-end laser are single mode. In contrast, the regenerative amplifier is relatively insensitive to the mode of the pump and requires high average power; furthermore, high powered fiber coupled diode bar arrays do not emit single mode radiation, and therefore would have low coupling efficiency into a single mode fiber. For these reasons, the regenerative amplifier uses a 400 μm diameter multi-mode fiber to deliver light from a 200 W, 940 nm diode bar array. In general, it becomes more challenging to support single modes as the diameter of the core increases. The largest single-mode core for a Yb fiber laser that the author is aware of is about 85 μm diameter [60], found in the NKT Photonics aeroGAIN rod amplifier, which must be held straight to avoid bend losses. Such extremely large core sizes typically housed within photonic crystal fibers, which enable precision engineered effective index profiles. Large improvements have been made in recent decades in drawing such sophisticated fibers.

In low power single mode fiber amplifiers, one can typically couple the pump directly into the core, such that both the signal and pump are guided in the doped active core of the fiber. As the mode area increases, this becomes more problematic, as single mode, high brightness pumps which would be required to couple into a core are less available and more expensive, especially at higher powers. This issue has been resolved through the use of cladding pumped fiber amplifiers. Cladding pumped gain fibers have a secondary cladding, much larger than the cladding which guides the signal, which guides the (generally multimode) pump. During propagation, the pump light interacts with the active ions in the core, causing population inversion. The two EDFAs in the front-end are core-pumped, whereby a single-mode 976 nm laser diode is coupled into the a 4 μm erbium doped core. In contrast, the two YDFAs are cladding pumped, whereby a multimode pump interacts with a 10 μm core.

Polarization may in general drift and become scrambled in a radially symmetric fiber; to address this, polarization-maintaining (PM) fibers can be made to maintain a particular polarization through the introduction of birefringence. When birefringence is sufficiently strong, there is strong modal mismatch between modes associated with different polarizations. While other techniques exist, the PM fibers used in the front-end utilize birefringence induced through the application of stress rods which run parallel to the core (see Figure 3.4); owing to the appearance, these fibers are often called panda fibers.

Dispersion is another important parameter of optical fibers. Figure 3.5 illustrates the dispersion properties of silica glass in the visible and near-IR spectral regions. The refractive index of a material determines the phase velocity. As can be seen, silica exhibits normal dispersion in its refractive index across the visible and near-IR. By analogy with the refractive index, the group index of a material determines the group velocity, and here the behavior of silica is more interesting. The group index of silica switches from normal to anomalous dispersion for vacuum wavelengths greater than 1.3 μm (the zero-dispersion wavelength). In simpler terms, for wavelengths in the normal dispersion regime (below 1.3 μm for silica), the red wavelengths travel faster than the blue wavelengths. If an unchirped pulse propagates through a normally dispersive fiber, it will impart a

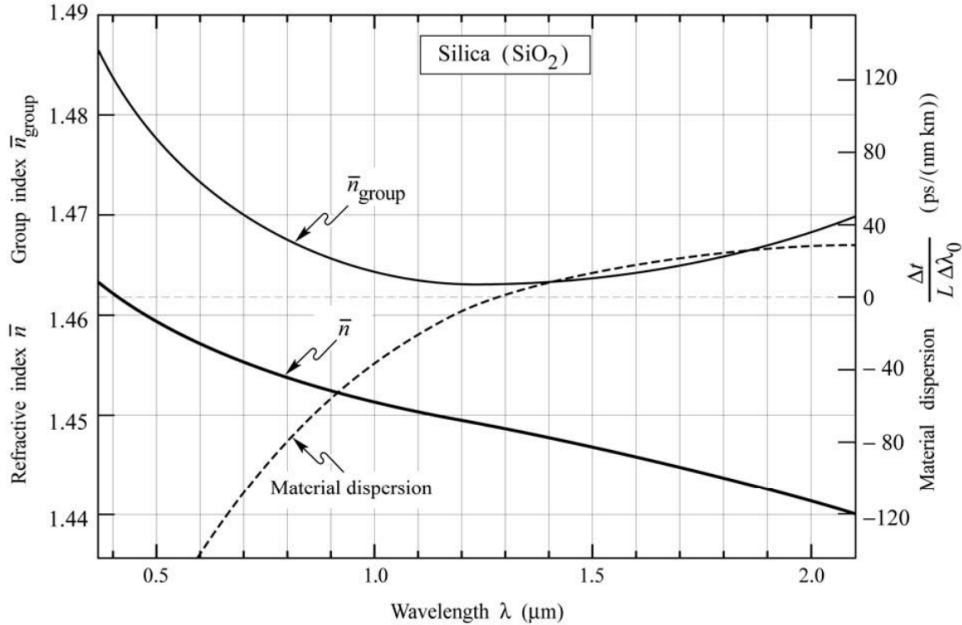


Figure 3.5: **Dispersion properties of silica in the visible and near-IR.** From Reference [126].

so-called *down-chirp* on the pulse, as the frequencies out in front will be lower than the frequencies at the back of the pulse. In the anomalous dispersion regime (above 1.3 μm for silica), the blue wavelengths travel faster than the red wavelengths. If an unchirped pulse propagates through an anomalously dispersive fiber, it will impart a so-called *up-chirp* on the pulse. These distinctions are critical to the spectral-broadening capabilities of self-phase modulation and the behavior of dispersion controlled EDFAs.

3.2 Erbium fiber oscillator

The high-performance of the front-end system can largely be attributed the 100 MHz, modulated, dual output erbium fiber oscillator (ELMO from Menlo Systems). The oscillator delivers high stability and low amplitude and phase noise, and is composed of PM-fiber components. Menlo Systems attributes these characteristics to a patented modelocking mechanism known as "figure 9." All modelocked oscillators require a modelocking mechanism, whether active or passive, in order to lock the phase relationship of the Fourier components of the spectrum in such a way that they

interfere to produce a stable pulse train. A common approach to modelock fiber oscillators is to put a saturable absorber in the oscillator cavity; since the behavior of this optic depends on the optical intensity, it can enforce pulsed operation over CW operation. This modelocking mechanism has been associated with excessive phase noise [127] and long-term degradation.




Figure 3.6: **Menlo Systems figure 9 oscillator diagram and measured spectra.** (a) Menlo Systems attributes the high stability and low noise of their fiber oscillators to the patented figure 9 fiber oscillator architecture, which involves modified nonlinear amplifying loop mirror modelocking mechanism. Diagram from Menlo Systems [128]. (b) Spectra measured from the two output ports of the oscillator.

The figure 9 oscillator architecture, which is illustrated in Figure 3.6(a), utilizes a modified nonlinear amplifying loop mirror (NALM). A fiber loop mirror is a loop of fiber connected to the two output ports of a 2x2 coupler. The phase relationship between the waves travelling either clockwise or anti-clockwise in the loop determines which of the two input ports the outgoing waves will favor. In a NALM, the fiber loop contains a segment of pumped, active gain fiber and a segment of passive transport fiber; in this way, the two directions of propagation around the loop will accrue different amounts of phase due to $\chi^{(3)}$ nonlinearities: in one case, the light will be amplified before traveling through the long passive fiber, inducing more nonlinear phase than if the weaker, unamplified pulse first travelled through the passive fiber. In this way, the NALM acts an intensity dependent reflector which can favor pulsed operation and therefore act as a modelocking mechanism. The figure 9 architecture modified the NALM architecture through the introduction

of a directional phase-shifter in the loop, and allows for full-PM construction [129].

Figure 3.6(b) shows measured spectra from the two ports of the oscillator, each centered around 1570 nm and containing enough bandwidth for an approximately 90 fs transform limited pulse. Each port delivers approximately 3 mW of average power (30 pJ of pulse energy), and the pulses have a linear chirp of -22000 fs^2 , stretching them to a duration of around 700 fs at the SC/APC fiber connectors.

It is worth commenting on the decision to start with an erbium oscillator and shifting to shorter wavelengths, instead of a ytterbium oscillator and shifting to longer wavelengths. Earlier iterations of the front-end for this experiment started with a ytterbium fiber oscillator using the all-normal dispersion (ANDi) architecture developed by the group of F. Wise [130]. The most recent approach involved generating a broadband “white light” spectrum which contained components near 1.5 μm by focusing a 1 μm beam through a piece of sapphire. This broadband spectrum was then used to seed a difference frequency generation process in a PPLN crystal which was wave-vector matched for the amplification of a suitable 1.5 μm spectrum. This system was unusable due to instability and frequent loss of modelocking, potentially arising from the nonlinear polarization evolution in the oscillator. Furthermore, this system involves significant free-space optics and recoupling from free-space into fiber. This is counter to the general design principle of our group: to achieve greater stability, keep as many components in fiber as possible. In contrast to earlier approaches beginning with a 1 μm oscillator, operating with a center wavelength above the 1.3 μm zero-dispersion wavelength of bulk silica glass allows for careful balancing and material dispersion and nonlinearities [131], and for the design of highly nonlinear fibers capable of generating tuneable, broadband and highly coherent radiation [132]. Lastly, starting with a 1.5 μm oscillator allows us to take better advantage of the tremendous investment from the telecommunications industry in high-quality, affordable components in the 1.5 μm band. An example of this is the dispersion managed EDFA design which will be described in the next section.

3.3 1.5 μm arm: Dispersion managed, broadband erbium-doped fiber amplifier

The erbium oscillator described in the previous section immediately seeds two parallel erbium-doped fiber amplifiers. These two amplifiers are tuned slightly differently: the 1.5 μm arm EDFA is tuned to deliver the broadest spectrum centered near 1.5 μm with a well-behaved phase, and the 1 μm arm EDFA is tuned for the largest spectral amplitude in a dispersive wave near 1 μm after a highly nonlinear fiber (HNLF). Despite minor differences, both amplifiers are in general tuned to deliver broadband, compressed pulses at their output. In the following section, we highlight the specific design of the 1.5 μm arm EDFA, noting that specific values of fiber lengths and pump diode currents do not necessarily refer to the 1 μm arm.

3.3.1 Background and theory

Before delving into the details of the design and operation, we can take a moment to put these amplifiers into the historical context that we began to establish in the introduction. Just a few years after the invention of the first EDFAs [50, 51], sub-picosecond erbium fiber amplifiers were first reported in 1989 [133] and 1990 [134, 135]. It would take another decade or so before the type of ultra-broadband, all-fiber EDFAs employed in this work would be developed. By designing the EDFA with a careful balance between material dispersion and $\chi^{(3)}$ nonlinearities, it is possible to amplify pulses in the few nJ energy regime with sub-40 fs durations [131, 136]. In 2003, the group of A. Leitenstorfer was the first to demonstrate this careful dispersion balancing in an EDFA. By adjusting the amount of single mode fiber in front of the EDFA, they demonstrated significant spectral broadening immediately out of the EDFA. The pulse was then launched into free space, compressed with a bulk silicon prism compressor, and recoupled into an HNLF to demonstrate supercontinuum generation [131]. In 2004, a group at OFS labs demonstrated an all-fiber version of this system, which had the tremendous benefit of being able to splice the EDFA directly to an HNLF. This was done by utilizing the material dispersion of single mode fiber to directly

compress the pulse to 34 fs without the need for a bulk compressor. After the HNLF, a 400 mW supercontinuum between 0.85 and 2.6 μm was demonstrated [136]. Later, it was demonstrated that the output of the HNLF itself could be compressed to sub-10 fs pulse durations [137, 138].

The delicate balance between material dispersion, nonlinear phase and spectral broadening is enabled by a unique characteristic of operating near the zero-dispersion wavelength of silica glass. Standard single mode silica fiber at 1.5 μm exhibits anomalous dispersion, whereas the 6.5 μm mode-field diameter erbium doped gain fiber exhibits normal dispersion¹ (see Section 3.1.2 for clarification of these terms). The balancing effect of these two dispersive regimes has a number of useful consequences for broadband, short pulse EDFAs. Firstly, by temporally pre-stretching the pulse in the transport fiber, and then allowing it to recompress over the course of amplification, one can prevent nonlinearities from occurring until the final segments of the amplifier and avoid runaway nonlinearities and pulse breakup. Secondly, as was described in Section 2.6, an initial up-chirp on a pulse can enhance the spectral broadening effects of self-phase modulation [121]; the anomalous dispersion of undoped silica fiber at 1.5 μm provides this up-chirp. Thirdly, comparison of the output of dispersion balanced EDFAs with theoretical models has indicated that these amplifiers act in the parabolic regime, which is enabled by the normal dispersion of the erbium gain fiber [136]. In parabolic amplification, the asymptotic solution of the temporal envelope shape of the amplified pulse is parabolic. A number of distinct advantages arise from operating in this regime, including avoidance of pulse breakup and relative insensitivity to input parameters. Furthermore, the amplified pulses have a nearly linear chirp, which can easily be compensated to achieve clean pulse compression [139].

3.3.2 Design and construction

The design of the dispersion managed EDFA in the 1.5 μm arm of the front-end is shown in Figure 3.7. All cleaves were performed with a Fujikura CT52 cleaver, and all splices were done

¹ We note that the normal dispersion seems to arise from the tight confinement of the mode in the 4 μm core. For instance, a similar gain fiber with an 8 μm core exhibits anomalous dispersion. This poses a challenge for scaling the power of dispersion managed ultrabroadband EDFAs.

with a Fujikura FSM-100p fusion splicer. Beginning at the oscillator, we will now detail the design and construction of the amplifier, and then discuss the optical pumping and power electronics.

An FC/APC connectorized PM1550 patch cable (part PMP-15-R-L-1-2SA-2SA from AFW Technologies or similar) was cleaved in the middle. The connectorized end was joined to the oscillator output via an adapter sleeve, and the bare end was spliced to the pigtail of a hybrid 980nm/1550nm wavelength division multiplexer (WDM) and optical isolator (part HB-IWDM-PM-1598-1-1-L-S-S-2W-FB from AFW Technologies). The isolator protects the oscillator from potentially disruptive or damaging backwards travelling pulses, and is incorporated into the same microoptic as the WDM. Approximately 123 cm of anomalous dispersion PM1550 fiber lies between the oscillator and the splice point to the gain fiber. The common port of the WDM was fusion spliced to an approximately 1.7 m long segment of single mode, single clad erbium-doped fiber (part Er80-4/125-HD-PM from Liekki) with an absorption of 80 dB/m at 1530 nm and panda geometry for polarization maintenance. The gain fiber was then spliced a very short, 5 cm pigtail of the output side WDM. The FC/APC connectorized pass port of the WDM was left with its 25 cm pigtail intact. Through guidance from collaborators and experimentation, we found that minimizing the length of passive, anomalous dispersion transport fiber after the gain fiber delivered shortest pulses. In principle, the distribution of PM1550 before and after the output side WDM is not critical for dispersion balancing - however, some length of pigtail eased beam delivery to the collimator and ultimately the free-space pulse shaper. For compactness and organization, most of the fibers are coiled in stacking plastic trays with a coil bend radius of 30.3 mm (part BFCT from Thorlabs).

The EDFA is configured to enable forward pumping by one laser diode, and backward pumping by two laser diodes. More pump power is devoted to backward pumping because the majority of pump extraction occurs with the higher energy pulses towards the back of the gain fiber. The pumps are single mode, 976nm, 950mW fiber bragg grating stabilized laser diodes with a PM fiber pigtail (part 1999CVB from 3SP Technologies), packaged in a 14-pin butterfly module. The two backward pumping diodes are joined using a 1x2 polarization combiner (part HPBC-98-P-1-




Figure 3.7: **A detailed diagram of the dispersion managed, broadband EDFA used in the front-end.** The fusion splices between the pump diodes and the WDMs are not shown. Approximate lengths (± 5 cm) between splices and components are given on the bottom line. Note that the pulse has a linear chirp of -22000 fs^2 at the SC/APC to FC/APC mating sleeve.

03-NNN-BBB-3-1 from Optizone Technology Limited). In total, the front-end has six laser diode modules, divided evenly between the 1.5 and 1 μm arms.

Each laser diode module is mounted on a universal butterfly laser diode mount (part LM14S2 from Thorlabs) configured for a type 1 pin configuration. The mount is connected to the controllers through two home-built 9-pin D-type connectors - one for the laser diode (LD) and one for the thermoelectric cooler (TEC). The controllers are housed in a 19" rack near the optical table, with the cabling passing underneath a cable bridge. Both the LD and TEC for a single diode module are controlled by single controller (part LDTC2/2E from Wavelength Electronics), capable of delivering 2.2 amps separately to the LD and TEC. The controller was operated in constant current mode and the currents were set manually with trimpots, and when necessary, were read out using a removeable USB computer interface. All six LD/TEC controllers were housed in a single, home-built rack shelf with fan cooling.


The electrical power supplies to convert AC wall power to low-voltage DC for the controllers were also homebuilt at JILA. In order to avoid electrical noise from a switching power supply

imprinting on the pump laser power, we chose to use low-noise and highly reliable linear power supplies often used by the medical device industry (part HBB5-3-OV-A+G from Condor/SL Power). To mitigate any cross-talk between the controllers and provide the cleanest possible electrical power, each LD/TEC controller has its own dedicated two channel power supply, with each channel supplying 3 amps at 5 volts. One channel of a single power supply powers the laser diode, and the other powers the thermoelectric cooler. This design affords a high degree of modularity, with a single power supply powering a single laser diode. The six power supplies are divided into two compact, rack-mounted shelves: one shelf contains the three power supplies for the 1.5 μm arm EDFA, and the other contains the three supplies for the 1 μm arm EDFA.


3.3.3 Pulse characterization

The 1.5 μm arm EDFA delivers 2.2 nJ pulses with a near-transform limited pulse duration of 35 fs. Through experimental optimization of the 1.5 μm arm, we settled on a pumping configuration which delivers approximately 800 mA and 1000 mA of current to the two backward pumping diodes, corresponding to an estimated 560 mW and 720 mW (1.28 W combined) cw optical power, and no current to the forward pumping diode. We characterized the resulting pulse immediately after the fiber collimator using a home-built SHG-FROG with a BBO crystal. Measured and experimental FROG traces are shown in Figure 3.8. When the experimental trace was mapped to a 1024x1024 grid, the retrieval showed a low error of 0.31%.

The reconstructed pulse in the spectral and temporal domain is shown in Figure 3.9. The small amount of spectral phase on the pulse upon exiting the collimator, comprising about two radians over the width of the spectrum, can readily be compensated by the SLM-based programmable pulse shaper. To highlight the remarkable spectral broadening and compression properties of the amplifier, it is worth noting that the seed pulse for the amplifier was a 700 fs pulse with a 90 fs transform limit.



(a) Experimentally measured SHG-FROG trace.




(b) Reconstructed SHG-FROG trace with an error of 0.31%.


Figure 3.8: SHG-FROG traces for 1.5 μm output of the front-end, as measured immediately after the fiber collimator.

3.4 1.03 μm arm: dispersive-wave shifted seed for the regenerative amplifier

In this section, we review the branch of the front-end which generates an ~ 20 nJ, ~ 2 ns, 100 MHz repetition rate, linearly chirped seed pulse for the Yb:YAG regenerative amplifier. This branch begins at the second output of the erbium oscillator, which is in turn amplified in a dispersion managed EDFA very similar to the unit described in the previous section. Instead of the compressed and amplified pulse being coupled into free-space, as is the case for the other EDFA, the output of this EDFA is spliced immediately to an HNLF, which generates a dispersive wave with significant amplitude near 1.03 μm . This component is then amplified in a series of two linear ytterbium-



(a) Spectral domain EDFA pulse characterization. Grey shaded region: normalized experimentally measured spectrum from an Ando optical spectrum analyzer; blue line: retrieved normalized power spectral density; orange line: spectral phase in radians of the retrieved pulse. Fine agreement between the measured and retrieved spectra reinforce the integrity of the reconstruction.



(b) Temporal domain EDFA pulse characterization. Blue line: retrieved time domain pulse, with FWHM of 35 fs and a small pre-pulse characteristic of minor third-order spectral phase; grey shaded region: Fourier transform of the measured spectrum with a 32.2 fs FWHM (transform limit).

Figure 3.9: **35 fs pulse reconstruction of the 1.5 μm arm EDFA in the spectral and temporal domain using SHG-FROG.** Analysis and plotting tool courtesy of Will Hettel.

doped fiber amplifiers (YDFAs). A similar approach was demonstrated in 2016 to seed a thin-disk regenerative amplifier with a couple of key differences: the EDFA signal went free-space to a bulk compressor before recoupling into the HNLF, and again went free-space into a bulk stretcher before recoupling into the YDFA; and the front-end produced a single output [57]. Although not used for seeding a high power free-space amplifier, the general sequence of an erbium oscillator, EDFA,

HNLf, and two YDFAs was also demonstrated in 2016 [140].

3.4.1 Dispersive wave generation in a highly nonlinear fiber

The great challenge of the front-end system described here is the generation of synchronized pulses at disparate frequencies. This challenge is elegantly solved in robust, all-fiber packaging through the use of a highly-nonlinear fiber. Very high $\chi^{(3)}$ nonlinearities can be achieved in engineered HNLfs through reduced mode areas and/or the addition of dopants such as germanium. An all-fiber, octave-spanning supercontinuum from an HNLf was reported in 2003 [141]. Shortly after, those authors demonstrated a high-degree of optical coherence across the broad spectrum [132].

Here, we have optimized our HNLf and seed pulse not for supercontinuum generation but for the generation of a *dispersive wave* near a wavelength of 1 μm . The mechanism behind the generation of dispersive waves has been an active area of research in the past two decades. A conventional explanation of dispersive wave generation begins with a pulse with sufficient intensity to drive nonlinearities being launched into a fiber with anomalous chromatic dispersion. When the effects of $\chi^{(3)}$ nonlinearities and chromatic dispersion exactly balance each other, a soliton is formed, which will propagate for long distances without distortion of the spectrum or temporal envelope. Inevitably, the pulse shape and power of the launched pulse will not exactly match that of the soliton. As the launched pulse evolves into a soliton under the perturbation of third or higher order nonlinearities, energy is shed by the inchoate soliton into a background commonly referred to as a dispersive wave. The observation that the dispersive wave tends to be blue-shifted relative to the soliton is attributed to the positive dispersion slope at the soliton wavelength [143]. Indeed, the HNLf chosen for dispersive wave generation in the front-end has anomalous dispersion across the soliton spectrum.

More recent insight into the process, in part from a group at the University of Auckland, demonstrates two key points about the generation of dispersive waves. Firstly, they demonstrated that pulses need not be soliton-like to shed energy to a dispersive wave, by showing that pulses in

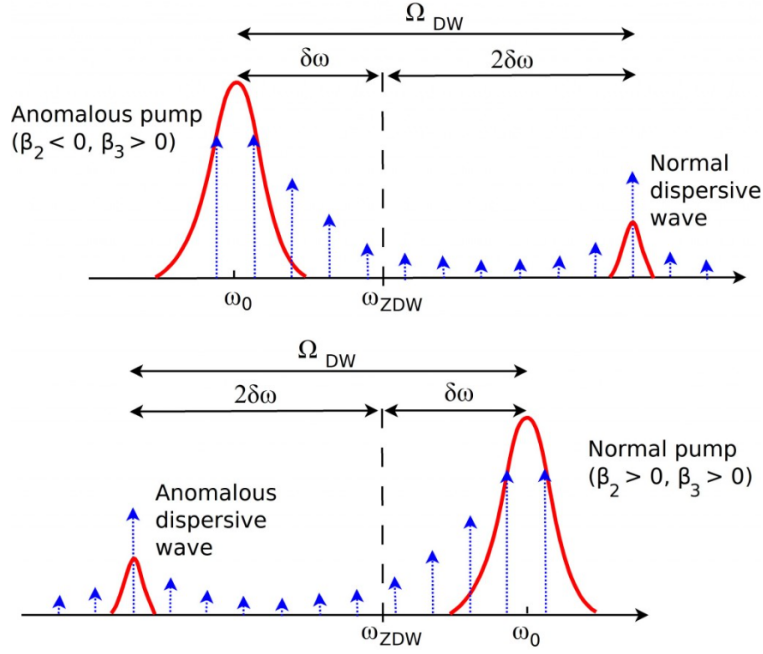


Figure 3.10: **Energy transfer from from a pump to a dispersive wave in a highly nonlinear fiber can be efficient due to a phase matched cascaded four-wave mixing process.** Although the pump wavelength in the work presented here is near $1.5 \mu\text{m}$ and experiences anomalous dispersion, generating a dispersive wave in the normal dispersion region, it is also possible for energy to be transferred the other direction across the zero-dispersion wavelength. The vertical blue arrows depict spectral components which may be involved in the intermediate steps of the cascaded process. Phase matching is enabled by the Kerr nonlinearity experienced by the pump beam. This figure was originally published on the blog of the Nonlinear Optics and Photonics Group at The University of Auckland and is reproduced here with their permission; this work is associated with Reference [142].

the normal dispersion regime can excited dispersive waves across the zero-dispersion wavelength [144]. Secondly, they demonstrated that the surprisingly large energy transfer often observed between the pump pulse and the dispersive wave can be accounted for by a phase-matched, cascaded four-wave mixing process [142], illustrated in Figure 3.10. The amplitude of dispersive wave generation generally remains negligible until the phase-matching condition is satisfied. A cascaded parametric nonlinear process can mimic higher-order nonlinearities while only employing lower order nonlinearities. For instance, a the third harmonic can be generated using entirely $\chi^{(2)}$ processes:

$$\omega + \omega = 2\omega \text{ (SHG)} \rightarrow 2\omega + \omega = 3\omega \text{ (SFG)} \quad (3.1)$$

In a cascaded parametric process, strong conversion efficiency from an effective higher-order process can be achieved with proper phase matching, even if the intermediate processes in the cascade are not well phase matched [145]. It can be shown that in the region of the zero-dispersion wavelength (ZDW), a pump detuned by $\delta\omega$ on the anomalous dispersion side of the ZDW will be phase matched with a dispersive wave detuned by $2\delta\omega$ on the normal dispersion side of the ZDW [142], as is illustrated in Figure 3.10. This phase matching is what enables significant energy transfer from the driving pulse in the anomalous dispersion region to the dispersive wave in the normal dispersion region. In the absence of the nonlinear phase accrued by the pump due to the intensity dependent Kerr nonlinearity, this process could not be phase matched - a simple observation of the material properties of fused silica in Figure 3.5 reminds us of that. For a thorough discussion of dispersive wave generation, we refer the reader to Section 12.1 of Reference [146]. Cascaded parametric processes have also been observed in gas-phase systems to generate VUV light with high efficiency [147].

Length	~ 3.5 cm
Dispersion of slow axis @ 1550 nm	5.7 ps/(nm km)
Dispersion slope of slow axis @ 1550 nm	0.027 ps/(nm ² km)
Mode field diameter @ 1550 nm	4 μ m
Estimated nonlinear coefficient	10.5 1/(W km)

Table 3.2: **A summary of key parameters of the HNLF used to generate the 1 μ m dispersive wave.**

The propagation of pulses in highly nonlinear fibers is often described by the nonlinear Schrödinger equation (NLSE), which can be solved by a variety of readily available numerical tools. We use pyNLO, a freely-available nonlinear optics modeling tool written in the python language, which uses a fourth-order Runge-Kutta solver to solve the NLSE for propagation through $\chi^{(3)}$ materials [148]. The results of a pyNLO simulation using the split-step Fourier model are shown in Figure 3.11. Simulating approximate experimental conditions, a 50 fs duration pulse centered




Figure 3.11: **Numerical modeling of pulse propagation in the highly nonlinear fiber.** A 1.3 nJ, 50 fs pulse centered at 193.5 THz (1.55 μm) is launched into an HNLF with parameters from the fiber in the front-end, generating a bright dispersive wave centered just above 300 THz (1 μm). The bottom map shows the spectral evolution over the length of the fiber, while the top plot compares the spectral intensity at the beginning of the fiber (blue curve) and after 35 mm of propagation (red curve). The intensity is plotted on a dB scale in both plots. Results are obtained from the python tool pyNLO [148].

at 1550 nm and containing 1.3 nJ of energy is launched into an HNLF with similar parameters to what is used in the front-end. The soliton in the anomalous spectral region can be seen Raman shifting to lower frequencies over the first few cm of fiber, while shedding a dispersive wave that accumulates energy around a wavelength of 1 μm (300 THz). The parameters of the polarization maintaining HNLF (from OFS) used in this work are summarized in Table 3.2.

The HNLF assembly is the first component following the dispersion managed EDFA in the 1 μm arm of the front-end. The fiber segments are composed of PM1550, HNLF and PM980. In order to ensure thermal and mechanical stability, as well as protect the fragile splices around the HNLF,

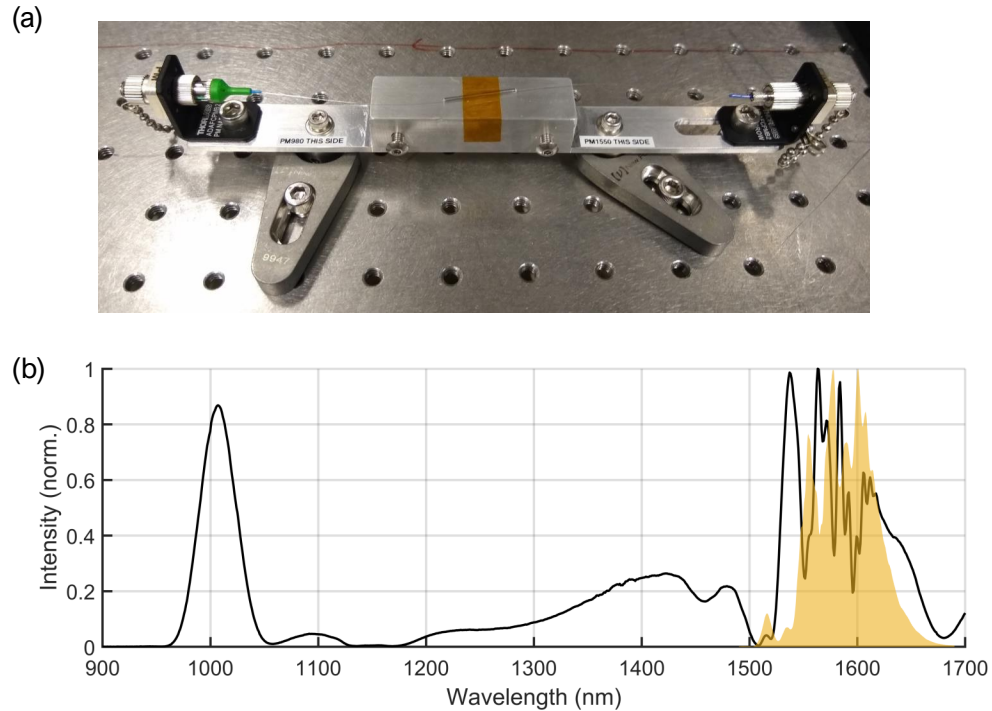


Figure 3.12: (a) A photo of the spliced and connectorized highly nonlinear fiber (HNLF). (b) Spectra before (shaded yellow) and after (black line) a ~ 3.5 cm segment of HNLF. From right to left in the photo: an FC/APC connectorized PM1550 patch fiber; spliced to the ~ 3.5 cm long highly nonlinear fiber; spliced to an FC/APC connectorized PM980 patch cable. Note that the spectra were recorded in a single sweep of an optical spectrum analyzer.

the spliced fiber segments are held in a custom machined aluminum assembly, photographed in Figure 3.12(a). The assembly is then covered with an anodized, formed sheet metal cover for added protection, allowing the assembly to be incorporated into the front-end as a robust, connectorized module.

Spectra taken before and after the HNLF using an optical spectrum analyzer are shown in Figure 3.12(b). The pulse that drove this spectrum originated from an EDFA with an additional 30 cm PM1550 patch cable in front of the input side WDM and a 40 cm patch cable after the output WDM. At the time of recording the spectra, the total power out of the HNLF was measured at 125 mW, but with modifications to the system, is 135 mW at the time of writing.

3.4.2 Ytterbium-doped fiber amplifier chain

After the HNLF, a series of two ytterbium-doped fiber amplifiers (YDFAs) amplify, stretch and spectrally clip the dispersive wave into a stable, energetic seed pulse for the cryogenically-cooled, 1 kHz Yb:YAG regenerative amplifier (“regen”). Although the regen is operated in the saturation regime, amplifier instabilities of various types can be reduced by a sufficiently stable and energetic seed pulse. In particular, a more energetic seed allows for few passes in the regen, making the system less sensitive to misalignment or other issues throughout the beam path. A more energetic seed also reduces the likelihood of a destabilizing and potentially catastrophic effect in the regen known as bifurcation [149], which is discussed in greater detail in Chapter 4. Therefore, there is strong motivation produce more energetic pulses from the YDFAs, and without the added complexity of pulse picking in fiber, the more energetic pulses result in relatively high average power. Design considerations for this higher average power will be discussed throughout the rest of this chapter as we review the design, construction and performance of the YDFA amplifier chain.

3.4.2.1 Design and construction

The design of the YDFA chain is illustrated in detail in Figure 3.13. The PM980 output of the HNLF is connected with FC/APC connectors to the input port of a three-port circulator (part HPM CIR-03-0.5-S-B-S-P from Advanced Fiber Resources). As light travels between ports, insertion losses are between 0.4-0.7 dB in one direction, and 37 dB in the other direction - allowing the circulator to also function as an isolator. The second port of the circulator is spliced to a chirped fiber Bragg grating, which has an ~ 1.3 nm reflectivity band centered about 1029.3 nm. A linear chirp of 750 ps/nm is applied to the reflected pulse, stretching it in time to around 1 ns in duration. This pulse seeds the first YDFA via a PM multimode combiner (part PMC02112A71A from ITF Technologies). The signal input port on the combiner is PM 10/125 single mode fiber; the pump input port is 105/125 multimode fiber. The pump port is connected to a single 10 W, 974 nm multimode fiber coupled laser diode (part PLD-10-974 from IPG Photonics). The laser diode is

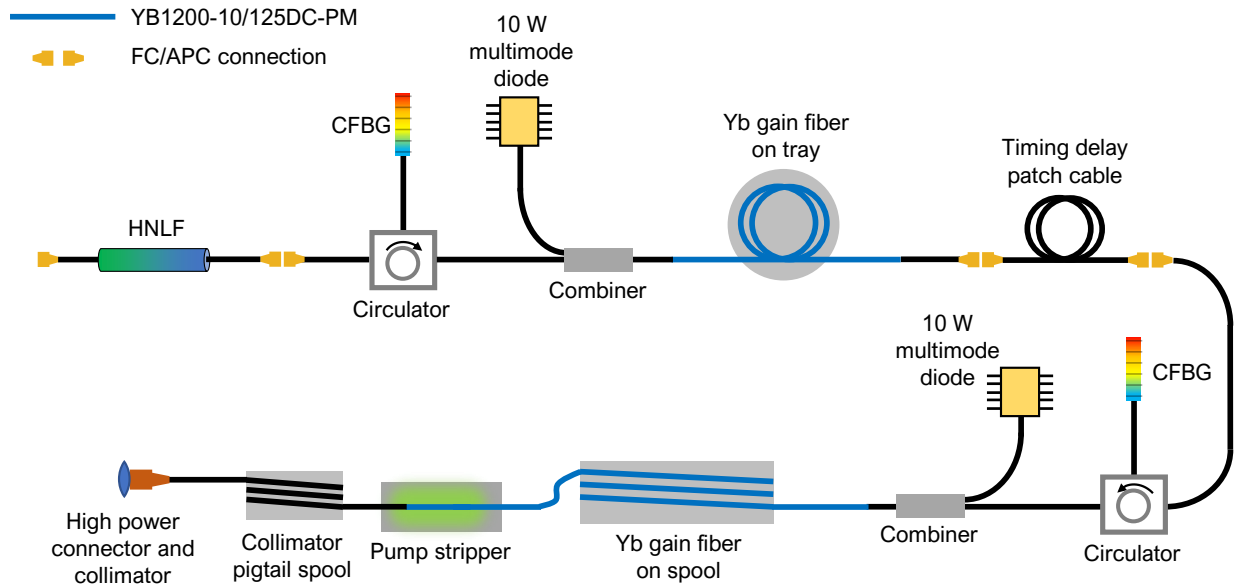


Figure 3.13: **A detailed diagram of the ytterbium-doped fiber amplifier chain.** HNLF: highly-nonlinear fiber; CFBG: chirped-fiber Bragg grating. The HNLF output pigtail and the timing delay patch cable are 6/125 PM980. The rest of the single mode fiber is 10/125, and the multimode laser pigtail is 105/125.

mounted to a homebuilt, copper plate cooled by a thermoelectric cooler to 30°C. Each pump laser diode has a dedicated 8 A Arroyo Instruments 4308 power supply and an Arroyo Instruments 5305 thermoelectric controller. The electronics are housed in the same 19" rack as the EDFA electronics.

The multimode pump light is coupled into the cladding the output fiber of the combiner (10/125, NA=0.08/0.46). This output fiber is spliced to around 3 m of large mode area, double clad, cladding pumped ytterbium gain fiber, with a core diameter of 10 μm and a cladding diameter of 125 μm (part YB1200-10/125DC-PM from Thorlabs). The core is estimated to have an absorption at 976 nm of around 6 dB/m. The gain fiber is coiled and lays flat on an approximately 10 cm diameter round aluminum tray. The output side of the gain fiber of YDFA1 is spliced to an FC/APC connectorized segment of PM980, which is connected to a PM980 timing delay patch cable. By simply connecting in a patch cable of arbitrary length, one can conveniently adjust the timing delay between the 1.5 and 1 μm outputs of the front-end.

Following the timing delay patch cable, YDFA2 is constructed nearly identically to YDFA1:




Figure 3.14: **Output spectra of YDFA1 and YDFA2, with regenerative amplifier output spectrum overlaid in yellow.** Spectra were measured using an optical spectrum analyzer. The narrowing that occurs between the regenerative amplifier seed and output results from the narrow gain-bandwidth of cryogenically cooled Yb:YAG.

the circulator, CFBG, pump, combiner and gain fiber are all the same. The YDFA2 gain fiber is spooled on the outside of a home-made, 20 cm diameter aluminum spool. Following the gain fiber is a home-built pump stripper to remove residual pump light. The gain fiber is spliced to the 10 μm core high-power connector pigtail immediately after the unabsorbed pump light has been removed, and the splice is protected in the epoxy on the stripper. The high-power connector is connected to a collimator, sending the beam free-space towards the regenerative amplifier.

3.4.2.2 Operating parameters

The first YDFA pump is driven by 3.2 A of current and maintained at 30°C, delivering 2.35 W of power. It produces about 60 mW of average signal power. The relatively low conversion efficiency is primarily due to the excessively large 10 μm core of the gain fiber, which effectively reduces the seed intensity of the amplifier and allows amplified spontaneous emission (ASE) to grow even for rather low pump powers. This feature was inherited, and we have plans to switch YDFA1 to a core pumped amplifier more similar to the EDFA described earlier. The spectrum from YDFA1 is shown in Figure 3.14(a). The sharp edges are set by the CFBG and the relative

lack of gain narrowing.

Nevertheless, 60 mW of average power is a sufficient seed for YDFA2, for which cladding pumping and a large mode area is much more appropriate. YDFA2 is pumped by 4.8 W (pump diode driven by 6 A at 30°C) and produces 2.15 W of usable signal after the free-space collimator, for an effective efficiency of around 45%. The spectrum at the output of YDFA2 is shown in 3.14(a), and exhibits some reshaping due to gain. The ASE background is around 40 dB below the signal peak as measured by an optical spectrum analyzer.

3.4.2.3 High power handling

YDFA2 steps into the exciting territory of higher average power fiber laser amplifiers. Although Yb fiber amplifiers can produce hundreds of watts of average power [60], even YDFA2 required some interesting modifications in order to perform well at a few watts. We briefly review some of these design considerations here.

We estimate that there is about 300 mW of absorbed pump light contained in the cladding at the end of the YDFA2 gain fiber. If the gain fiber were spliced directly to the output collimator pigtail, all of this light would be incident on the splice - likely much would pass through and pollute the output beam, and it may accelerate the degradation of the splice. To address this, we constructed a pump stripper - a device which removes residual pump from a cladding pumped gain fiber. An ~4 cm segment gain fiber near the end is stripped of its protective acrylic covering and laid in a machined groove on an aluminum block. The bare fiber is then covered in Norland 61 UV cureable epoxy, which has a relatively high index of 1.56. The high index layer destroys the guiding properties of the inner cladding, causing the pump light to rapidly leak out of the fiber before the next splice. This device is indicated in Figure 3.15(c)(i).

Even when a standard fiber connector is perfectly clean, when high power laser radiation leaves a fiber through a standard connector, heat generated at the fiber tip can cause the epoxy used to hold the fiber in the ferrule to degrade and off-gas. These gases can adhere to the end of the fiber, scattering and absorbing light and causing runaway burning of the fiber tip. These problems




Figure 3.15: (a,b) Photographs of ytterbium gain fiber damage associated with the high average power operation of YDFA2; (c) a photograph showing custom components in YDFA2 to improve output power and system reliability. The scale bar in (b) is 200 μm . Visible parts in (c) include the pump stripper (i), the large diameter gain fiber spool (ii), and the high power connector spool (iii). Also visible is the high power connector (brass, to right) and the two 10 W MM pump diodes mounted on copper plates. Under normal operation, the 1.5 μm arm EDFA sits on top of YDFA2.

can begin to occur in the few watt regime. To address this, we used a specialty connector in which the fiber tip is surrounded by an air gap (part PMJ-A3AHPMX-1060-10/125-1-5-1-PM1060L from OZ Optics). During the installation of this component, while coiling the 10 μm core PM pigtail, it was observed that the output power depended strongly on the bend radius of the fiber. While coiled in a 6 cm diameter plastic tray (part BFCT from Thorlabs), the output power was 0.882 W at 4.8 W pump; when coiled in a custom made ~ 10 cm diameter aluminum spool (indicated in Figure 3.15(c)(iii)), the output power increased by almost 60% to 1.38 W at the same pump power.

In the first generation of YDFA2, the gain fiber was coiled loosely in ~ 8 cm loops and rested on an aluminum puck. On a time scale of about once per year, the gain fiber would fail catastrophically with a break or kink in the fiber about 30 cm after the splice to the combiner,

where pump intensity is relatively high. Photography and microscopy of this damage is shown in Figures 3.15(a) and (b). We hypothesize that this damage was caused by a couple of factors: due to relatively large thermal loads from the 4.8 W of pump light with relatively low extraction near the front of the gain fiber, the fiber likely experienced significant heating. Without careful heat extraction, this heating may have contributed to the damage. Further, the fibers rested on top of each other, potentially giving rise to microbends that may have contributed to the fracture. To solve this problem, we carefully machined an ~ 150 μm deep groove in a spiral pattern along the outer surface of an ~ 20 cm diameter aluminum ring. The gain fiber was laid in this groove and held in place with kapton tape (Figure 3.15(c)(ii)). The precision machined groove minimizes any transverse stress on the fiber, prevents overlap of the fiber, and improves thermal management. Spooling the gain fiber in this way resulted in an another greater than 50% increase of the output power from 1.38 W to 2.15 W. No failures have been observed in the Yb gain fiber since the installation of the custom spool.

Chapter 4

A 3 μm , PPLN-based OPCPA pumped by a cryogenically-cooled Yb:YAG regenerative amplifier

4.1 Introduction

Due to the wavelength scaling of the high-harmonic generation phase matching cut-off energy, the generation of a bright (phase matched) harmonic beam with photon energies reaching 1 keV requires the use of ultrafast lasers with a wavelength of $\sim 3 \mu\text{m}$. We generate and amplify these pulses in a home-built, optical parametric chirped pulse amplifier (OPCPA) using four stages of MgO-doped periodically poled lithium niobate crystals. The OPCPA in its current form produces $\sim 620 \mu\text{J}$ compressed pulses centered at a wavelength of $2.95 \mu\text{m}$, with spectral intensity spanning from 2.8 to $3.1 \mu\text{m}$ (tail-to-tail). The compressed pulses were characterized using SHG FROG to have a FWHM duration of 99 fs - near the transform-limit of 95 fs - resulting in a peak power of 5.9 GW. The system delivers high pointing and power stability and an excellent mode. The OPCPA has been used to generate high harmonics in argon gas, as described in Chapter 5. Conference proceedings describing the OPCPA can be found in [150, 151, 152, 153].

The OPCPA described here evolved from an earlier iteration, described in the PhD theses of Michael Gerrity [154] and Susannah Wang [155], as well as Reference [156]. The prior version of the system employed a $1 \mu\text{m}$ fiber front-end, which drove a white-light seeded PPLN OPA to produce a $1.5 \mu\text{m}$ seed. The seed was stretched by a grating pair stretcher, and the OPA chain was composed of four PPLN crystals. The system was pumped by a chain of four cry-Yb:YAG amplifiers consisting of a regenerative amplifier and three booster amplifiers, which resulted in 26 mJ, 1 kHz

at 130 ps. The compressed 3.1 μm idler pulse had a pulse energy of 0.85 mJ and a duration of 420 fs. The bandwidth supported 93 fs pulses, with the long duration attributed to stretcher/compressor mismatch. A number of major changes have been made to the system since this first iteration. The front-end is entirely different (see Chapter 3), with improvements in seed bandwidth, stability and robustness. Critically, the stretcher was replaced by a grism stretcher, with the second and third order dispersion on the phase conjugated idler pulse compensable by a grating compressor. An SLM-based pulse shaper was incorporated for fine, user-programmable dispersion control. The four-amplifier pump chain was replaced with a ~ 19 mJ single stage regenerative amplifier, which offers advantages in terms of robustness, compactness, stability and mode quality, but has also resulted in a decrease in the available pump energy - ultimately limiting the achievable 3 μm pulse energy without the addition of a booster. The pulse duration of the pump was increased by a factor of ~ 2 (to about 275 ps FWHM), enabling higher pulse energy from a single-stage amplifier before damage and/or nonlinearities set in. In principle, the longer pulses enable higher pump pulse energy within the limited aperture size of PPLN, but they also decrease the peak intensity achievable before damaging the crystals.

The requirement of phase matching in HHG has informed our target pulse energy and duration for the OPCPA: ~ 8 -cycle (80 fs) pulses in the single to few mJ pulse energy range. What follows is a summary of the aspects of HHG phase matching that are relevant to determining the optimum drive laser parameters. We direct the reader to Chapter 5 for more thorough coverage of this material. The condition of perfect phase matching puts an upper-bound on the fraction of gas atoms that should be ionized (called the critical ionization). This fact addresses the often asked question of why the cut-off photon energy cannot be extended arbitrarily by increasing peak intensity: doing so could prevent phase-matching and result in a dim source. Rather, there is an optimum on-target driving laser intensity that is needed to reach, but not exceed, critical ionization: for the case of a 3 μm driver, that is expected to be around 1.3, 2.7 and 3.8 W/cm^2 for Ar, Ne and He, respectively¹ [41]. With a sufficiently tight focus and short pulse duration, these intensities can

¹ In practice, having some overhead on peak intensity is useful, especially with high pressure gases.

in principle be realized with low, few-hundred μJ or less pulse energies [84]. Here again, the constraints of phase matching must be considered. The tight focusing HHG geometry severely limits the volume of emitters that can be phase matched - not just because the short confocal parameter, but also because the Gouy phase and the intensity dependent atomic dipole phase impede phase matching. Phase matching can be improved through the use of a looser focus, requiring higher peak powers to achieve critical ionization, but the density-volume product of phase-matched emitters is still significantly less than what can be achieved with a gas-filled waveguide - even more so in the mid-IR than with shorter wavelength drivers. The use of gas-filled waveguides demands mode matching between the focused driving laser and the waveguide mode. In the mid-IR, due to the unfortunate scaling of propagation losses from unstructured glass capillaries with wavelength and waveguide diameter, unstructured fibers are typically required to be at least $150\ \mu\text{m}$ in diameter for few mm interaction lengths and larger for longer interaction lengths. Again, we find that higher peak power benefits mid-IR HHG in waveguides. It is then reasonable to ask if higher peak power should be sought with shorter pulses or more pulse energy. The HHG process can in principle be scaled up to larger interaction diameters, resulting in more harmonic flux, if more pulse energy is available. However, the situation for pulse duration scaling is more complicated. Single-cycle drivers tend to exhibit group velocity walkoff effects that limit the viable interaction length to generate HHG - in this case, dispersion of the very short pulse impedes phase matching after a short interaction length [157]. In the opposite extreme, if a driving pulse is too long, the ionization fraction builds slowly, reaching critical ionization before the peak of the pulse, and thereby reducing efficiency as well as the photon energy cut-off, which scales linearly with the intensity. An 8-cycle pulse is thought to offer an acceptable compromise between these two competing effects. A good review of this topic can be found in Reference [158].

In the remainder of this chapter, we will review the literature on existing PPLN OPCPA technology in the 2- $3.5\ \mu\text{m}$ range, describe the cryo-cooled Yb:YAG regenerative amplifier, explain how dispersion control is managed with a grism stretcher, grating-pair compressor and an SLM-based programmeable pulse shaper, and finally describe the 4-stage PPLN OPA chain, with an

emphasis on practical considerations of the amplifier, pulse characterization and stability.

4.2 A review of PPLN OPCPA technology in the 2 - 3.5 micron range

In this section, we review existing periodically-poled lithium niobate (PPLN) based OPCPA systems that are relevant to the work presented here. PPLN is able to access a high effective nonlinearity (d_{eff}) of around 15 pm/V, is broadly transparent in the range of 0.5 to 4 μm , and can phase match many different nonlinear optical processes depending on the selected poling period. On the other hand, PPLN suffers from a relatively low damage threshold compared to alternative crystal candidates and is extremely difficult to acquire with clear apertures larger than 3 mm in one dimension. For these reasons, PPLN is commonly used in low pulse energy systems, or in the early stages of high pulse energy systems. The OPCPA presented in this thesis is rather unique in that, at the time of writing, PPLN is utilized in all four stages of amplification to generate (uncompressed) pulse energies of greater than 1 mJ. Nevertheless, PPLN has been a primary limiting factor of the OPCPA, and alternative crystals will be discussed in detail in Appendix A.

There are two systems of note generating femtosecond pulses in the 2-3 μm range. The Max-Planck Institute of Quantum Optics has developed an OPCPA that utilizes PPLN for the first two OPA stages followed by a bulk lithium niobate stage [159]. The system is pumped by a 1.6 ps Yb:YAG thin disk amplifier and generates 10.5 fs (1.5 cycle), 1.2 mJ, 3 kHz pulses at a wavelength of 2.1 μm without nonlinear post-compression. We note that the relatively short pump pulses allow for the use of lower d_{eff} bulk crystals - a challenge with the long pump pulses from the pump laser presented here. An earlier version of this system was used to explore the issue of optical parametric generation (OPG, often called superfluorescence) and its subsequent amplification [160]. A common test of OPG levels is to block the seed beam while leaving the pumps on the crystals. The authors found that the (unseeded) OPG signal from the three stages was 780 μJ , or 85% amplified signal of 920 μJ seeded by an estimated 4 nJ. However, the authors estimated that when seeded, the actual quantity of amplified OPG in the amplifier output was reduced to 184 μJ (20%).

The second sub-3 μm PPLN OPCPA of note is at ETH Zurich, which utilized PPLN for all

four OPA stages to produce 126 μJ , 14.4 fs, 100kHz pulses centered at 2.5 μm with an average power of 12.6 W [161]. The system was pumped by a 280 W, 1.8 ps Innoslab-type amplifier. Elliptical spatial modes were formed on the last two PPLNs in order to make more complete use of the rectangular clear aperture typical of PPLN crystals. We also note that the third and fourth stages were pumped with high peak intensities of 33 and 24 GW/cm^2 and high average powers of 50 and 135 W (0.5 and 1.35 mJ). The system was later upgraded and used for soft X-ray HHG [84].

In the spectral region above 3 μm , the system at ICFO utilized five consecutive PPLN stages followed by two bulk potassium niobate stages to amplify 3.25 μm pulses [73]. An earlier version of this system reported pumping the fourth stage PPLN with 0.62 mJ, 9.5 ps pulses at 100 W average power (1064 nm), but at peak intensities above 4.5 GW/cm^2 , a strong beam reshaping process was observed in the PPLN [162, 163]. ETH Zurich reported a 3.4 μm wavelength OPCPA with 44.2 fs, 21.8 μJ pulses that utilized a chain of three PPLN OPAs [164]. The authors achieved high gain bandwidth by: (1) using aperiodically poled crystals in the first two stages and (2) using the final stage PPLN in a noncolinear fashion. Another three stage PPLN OPCPA with a very low stretch ratio (pumped by 400 fs pulses) was reported by CNRS that notably compressed both the signal (20 μJ at 1.55 μm) and the idler (10 μJ at 3.1 μm) by using two bulk compressors: 110 mm of SiO_2 for the signal and 12 mm of Si for the idler [165]. The Fastlite system described in the Introduction utilized two PPLN stages followed by two bulk lithium niobate stages for power amplification, with the early stage PPLNs taking no more than 125 μJ of pump energy [77].

Of the systems reviewed here, the highest pump pulse energy and average power was reported by Bigler et al. at 1.35 mJ, 135 W with a peak intensity of 24 GW/cm^2 , corresponding to a peak fluence of 0.043 J/cm^2 . We compare this to the ~ 10 mJ pulse energy used to pump the final state of the OPCPA in this work. Using the empirical square root scaling of damage fluence from their 1.8 ps pulses to the 275 ps pulses used in this work, the corresponding peak fluence is 0.5 J/cm^2 - about half the peak fluence used in this work. We note that the PPLN reported by Bigler et al. (operated at room temperature) was mounted in indium foil to improve thermal and electrical contact in the hopes of reducing pyroelectrically induced photorefractive damage [166], and that




Figure 4.1: **Schematic of the cryogenically cooled Yb:YAG regenerative amplifier.** CLM: water cooled collimator, HR: high reflector, DM: dichroic mirror, TFP: thin film polarizer, PC: Pockels cell, $\lambda/2$: half-wave plate. $R = 15$ m refers to the radius of curvature of the mirrors. In the final configuration, the crystal was pumped from the right. Not shown are the optics to collect the unabsorbed pump and direct it to a water cooled beam dump. From Reference [152].

the crystal was pumped at a non-colinear angle of 3.5 deg. The work in this thesis demonstrates the operation of PPLN in a new regime of high pulse energy and high fluence, while at the same time, we acknowledge that we are operating near the limitations of PPLN and are actively seeking alternative crystals.

4.3 Cryogenically cooled Yb:YAG regenerative amplifier

The performance of an OPCPA is highly dependent on the pump laser. Instabilities in the pump of all kinds can be dramatically imprinted on the amplified signal and idler. As an example, the parametric gain experienced by an OPCPA (see Chapter 2 scales approximately as the exponential of the square root of the pump intensity; as a result, even small fluctuations in pump intensity across mode or in time can have outsized effects on the amplified beam. The OPCPA reported here is pumped by a single-stage cryogenically cooled Yb:YAG regenerative amplifier (RA), seeded by the YDFA chain described in Chapter 3. In this section, we will provide an overview of the system, a review of similar systems in the literature, discuss cryogenic Yb:YAG as a gain media, describe the issue of bifurcation, and finally provide some performance specifications. The




Figure 4.2: Photographs of the regenerative amplifier showing, from left to right, the custom enclosure and auxiliary equipment for the cryocell, water cooled external beam dump and collimator for thermal stability, and the layout of the optics in the amplifier.

regenerative amplifier (RA) described in this section is reported in References [153, 152].

The RA delivers a typical pulse energy of 18-19 mJ at a 1 kHz repetition rate, resulting in an average power of 18-19 W. The amplified pulses are ~ 275 ps FWHM in duration with a FWHM spectral bandwidth of ~ 0.19 nm. The laser delivers near TEM_{00} modes in the far and near field. The system is also relatively compact, housed inside an enclosure with a footprint of 55×125 cm². The spectrum and modes are presented in Figure 4.3. A schematic depicting the layout of the RA is given in Figure 4.1, and photographs of the system are presented in Figure 4.2. The RA is a four mirror bow-tie ring resonator, a simple design that yields good stability. The cavity consists of two dichroic mirrors reflective to the signal and transmissive to the pump, and two high reflectivity spherical mirrors with a radius of curvature of 15 m. Pulses make about 20 passes in the RA before being ejected from the cavity by a two-crystal BBO Pockels cell followed by a half-wave plate and a thin film polarizer. As an extra level of attenuation to ensure a high pulse contrast between the selected 1 kHz pulse train and the rejected pulses, a second Pockels cell is used between the front-end and the RA. The mode size in the gain medium is calculated to have a $1/e^2$ diameter of 2.1 mm. The gain medium is a 5 mm long, 8 mm diameter, 10% doped Yb:YAG crystal mounted

in a copper sleeve and maintained at ~ -220 C by a 90 W capacity Gifford-McMahon cryogenic cooler (PT90 from Cryomech). The closed cycle cooler circulates helium gas, and enables hands-off maintenance of sub-liquid nitrogen temperatures. The cold finger holding the crystal is maintained in a small vacuum chamber which is kept in the low 10^{-8} mbar range using an ion pump and the cryo-pumping effect of the cold finger. The crystal is pumped by a fiber coupled 940 nm wavelength diode bar array capable of delivering up to 200 W CW power and delivered to the crystal via a 400 μm diameter fiber and relay imaged onto the crystal with a single lens and a magnification of 5. The amplifier is seeded with ~ 2.5 nJ pulses at 100 MHz from the front-end (Chapter 3). The seed pulses have a linear chirp of 1.5 ns/nm from two chirped fiber Bragg gratings in the front end - applied to the 1.2 nm of bandwidth centered about 1029.5 nm, this results in 1.8 ns duration seed pulses.

A number of mid-IR OPCPAs have been reported over the past decade with pump lasers featuring either MOPA architecture with Nd:YVO₄ gain medium [167], flash-lamp pumped Nd:YAG [71], or CPA architecture using diode pumped Yb:YAG and Ho:YLF both at room and cryogenic




Figure 4.3: **Output modes and spectrum of the regenerative amplifier.** The near field mode (top image) is taken 2 m from the amplifier output. The far field mode (bottom image) is taken in the focus of a 50 cm lens. Both modes are near TEM₀₀. The spectrum was taken using a single mode fiber coupled to an Ando optical spectrum analyzer set to 0.05 nm resolution in double monochromator mode.

temperatures [168, 82, 77, 169, 170]. Here, we will describe a few previously reported Yb pump lasers to provide context for our work. A comparable cryo-cooled Yb:YAG RA generating 13.2 mJ pulses at 1 kHz was reported by an MIT group in 2010 [168], with comparable energy stability ($\sim 0.3\%$) to the system reported here (0.43%). The system was seeded with a low pulse energy of 0.12 nJ, and a two-pass booster amplifier brought the pulses to 40 mJ (at 2 kHz repetition rate). The MIT system utilized liquid nitrogen dewars to cool the crystals, requiring much more intervention than the closed-cycle helium refrigerator used in the system reported here. Since the amount of GDD needed to stretch a pulse to a given duration increases with decreasing bandwidth (see Chapter 1), narrowband pulse compressors tend to be inconveniently large - the authors overcame this challenge with a unique folded compressor design to bring the pulses to 15 ps. A four pass cryo-cooled Yb:YAG amplifier capable of an impressive 250 W of average power was developed by Hemmer et al.[171]. The 2.5 mJ pulses allowed for relatively short, sub-20 ps duration stretched pulses in the cavity. Another highly relevant system is a cryo-cooled Yb:YLF bow-tie RA that delivers 20 mJ pulses at up to 3.5 kHz repetition rate for an average power of up to 70 W [172]. Yb:YLF exhibits much broader emission at cryogenic temperatures than Yb:YAG, which allowed the authors to amplify a FWHM bandwidth of 2.1 nm centered at 1018.6 nm - approximately ten times the bandwidth generated by our comparable cryo-cooled Yb:YAG RA. As noted, the added bandwidth allows for a much more compact and practical pulse compressor. A 120 pass Yb:YAG thin-disk RA was developed with comparable specifications to the system reported here, delivering up to 17 mJ pulses at 3 kHz. We note that the front-end followed a similar architecture to our own.

Yb:YAG exhibits improved thermal and optical properties when cooled to cryogenic temperatures, but with cryogenic coolings comes some unique challenges [154, 173]. A primary concern of high average power gain media is the extraction of heat, which inevitably originates from sources such as unextracted pump and thermal relaxation pathways associated with the quantum defect. Yb:YAG's thermal conductivity increases from 17 to 47 W/mK as the crystal is cooled from 300 to 77 K. Despite the favorable scaling, even at cryogenic temperatures the thermal conductivity

is relatively low - for instance, it is about an order of magnitude lower than that of sapphire at cryogenic temperatures. Thermal gradients in the crystal over the beam profile can result in problematic thermal lensing effects due to the temperature dependence of the refractive index. Here again, cryo-cooling helps, with a decrease by approximately a factor of 9 when Yb:YAG is cooled from 300 K to 77 K. Over the same temperature decrease, the coefficient of thermal expansion reduces by about a factor of 3. Finally, we note that the high gain is enabled by a stimulated emission cross section that increases by about a factor of four, to the relatively high value of 11×10^{-20} cm^2 , when Yb:YAG is cooled from 300 to 77 K [154, 174].

Perhaps the most important change Yb:YAG undergoes as it is cooled to cryogenic temperatures is a transition from a quasi-three level to a four level gain medium, which is depicted in Figure 4.4. The $^2F_{7/2}$ lower manifold in ytterbium's active laser transition is split into closely spaced levels. The terminal level is separated from the active laser lower level by only 0.075 eV [154]. Due to this low energy spacing, at room temperature the lower laser level experiences a 5.5% fractional thermal population. Such a significant population of the lower laser level makes Yb:YAG effectively three level at room temperature, enabling absorption of the ~ 1030 nm laser light and leading to lower amplification efficiency and, potentially, runaway thermal issues. However, when Yb:YAG is cooled to 77 K, the fractional thermal population of the lower laser level reduces to 0.001% [154]. Due to the short lifetime of the lower laser level, the laser transition sees an effectively unpopulated lower level, making it effectively four level.

While cryogenic Yb:YAG is attractive in many ways, its low gain bandwidth and cryo-vacuum requirements present some challenges. The emission cross section (which peaks around 1029.3 nm) is reduced to around 0.25 nm FWHM bandwidth. This narrow emission cross section strictly limits the bandwidth of amplified pulses, forcing the use of large and unwieldy pulse compressors if shorter pulses are desired. We also note that operating the gain medium at cryogenic temperatures requires a carefully designed vacuum and cryogenic system. The most consistent limitation of the up-time of the OPCPA described here is the crystal cryocell. Due to the low temperature of the crystal, ambient molecules in the vacuum freeze onto the optical surface of the




Figure 4.4: **Cryogenically cooling Yb:YAG depopulates the lower laser level, causing the gain medium to transition from quasi-three level to four level.** At room temperature, the lower laser level is partially populated due to the thermal distribution of electron energies. Cryogenic cooling forces nearly all of these electrons to the terminal level. Without any significant population in the lower laser level, absorption of the laser (signal) radiation is vastly reduced, increasing the efficiency of amplification.

crystal. We have found that over a month or two timescale, this build-up can cause a decrease in output power of the amplifier from, for instance, around 18 W to around 15 W. Generally, the power can be regained by warming the crystal to room temperature while pumping out the cryocell with a turbomolecular pump for ~ 1 day, thus allowing sublimation and removal of any adhered material. We note that more careful design of the cryocell and vacuum system may be able to relieve this constraint. For instance, by introducing a cold tube or funnel near the crystal facet, the probability of molecules freezing on to the crystal surface could be reduced. Eliminating O-rings in the gate valve and the window flanges, more carefully cleaning and potentially baking all components (including the cold finger) in the cryocell, and improving the pumping speed of the always-running ion pump may be avenues towards continuous operation of the RA without periodic warming of the crystal.

Yb:YAG has an excited state lifetime of around 1 ms, which also corresponds to the time between pulses at our 1 kHz repetition rate. When these two time scales are similar, the amplifier becomes particularly vulnerable to an effect known as bifurcation [149]. To understand bifurcation, it is helpful to consider the Yb:YAG crystal after it has achieved a state of significant population inversion. An arriving seed pulse will see strong gain due to the strong population inversion, and

will significantly deplete the excited state population. Due to this depletion, the next pulse may not encounter a gain medium with as strong of population inversion as the first pulse, resulting in lower gain and lower energy extraction from the crystal. With the pump delivering constant power and the second pulse extracting less energy, the third pulse may again experience a higher gain. This oscillatory behavior in pulse energy is referred to as bifurcation, and in addition to being a cause of pulse-to-pulse instability, can cause catastrophic damage to an amplifier which may be operating somewhere near the the damage threshold of its optics. It has been demonstrated that an increase in seed pulse energy can help to mitigate bifurcation, which was one of the design considerations of the front-end seed laser described in Chapter 3.




Figure 4.5: **Output power of the regenerative amplifier over multiple days plotted with lab relative humidity and temperature.** Regen output power is clearly correlated with temperature and/or humidity, which track in the lab due to a coupled control system - see, for instance, the change in periodicity occurring near hour 45.

Due to the nonlinear dependence of the amplified signal and idler beams on the pump, and the even more highly nonlinear dependence of the HHG beam on the amplified idler, stability of the RA is of paramount importance. We measured the pulse to pulse instabilities over 10 minutes to be $<0.47\%$ RMS and average power instabilities to be $<0.43\%$ RMS over 15 hours. We find

that the average power instabilities are strongly correlated to environmental fluctuations in the lab due to the climate control system (see Figure 4.5). While temperature fluctuations are seen as the most likely culprit, we note that humidity and temperature control are coupled in the lab, so it is difficult to identify the relative contribution of these two environmental factors. Nevertheless, the periodicity of power fluctuations track with the periodicity of humidity and temperature. We also measured pointing stability by tracking the displacement of the beam focused by a 50 cm lens placed ~ 2 m from the output of the amplifier using a DataRay camera and associated software. We measured displacement instabilities of the centroid of the focus of $1.7 \mu\text{m}$ RMS over 50 hours. We also experimented with a method to improve amplifier stability by equipping an intra-cavity mirror mount with piezo-actuators coupled via an electronic feedback loop to a quadrant detector placed outside of the cavity. This reduced the aforementioned displacement instabilities to $1.06 \mu\text{m}$, or 0.3% of the $1/e^2$ diameter at the focus. We ultimately decided not to use active feedback, largely because the pointing stability was sufficient for our applications without it. To further improve stability, we constructed a robust enclosure for the RA, shown in Figure 4.2. The walls were constructed from $1/4$ in thick aluminum, and the lid was constructed with lightweight AluPOLY aluminum composite material (ACM), composed of a polyethylene core sandwiched between two bonded aluminum sheets. The material can easily be sheared and punched, making it highly workable compared to acrylics or equivalent thickness sheet metal. To prevent air currents and the introduction of dust, we used rubber weather-stripping type seals where the lid joined the enclosure walls, and rotary shaft oil seals (which work particularly well for this purpose) to mate the enclosure with beam transport tubes. Early in the development of the RA, we observed longer term drifts in system performance that were attributed to gradual warming of amplifier components. To reduce this instability, we carefully tracked down heat sources in the RA and used building process chilled water to remove heat where necessary. In particular, the unabsorbed pump light was directed to a water cooled beam dump which was partially outside of the enclosure, and the pump collimator was wrapped in water-cooled, flexible copper tubing. Both of these modifications are shown in the inset of Figure 4.2. It is important to note that, while the single stage, uncompressed RA introduces

challenges for OPCPA pumping, the simplicity of the OPCPA pump laser likely contributes strongly to its stability and robustness. As a point of comparison, we consider the Innoslab amplifier used for OPCPA pumping in Reference [84], which utilized active beam stabilization at the output, followed by spatial filtering, compression, and a second active beam stabilisation at the compressor output.

A temporal pump envelope with a flat top results in more uniform gain across the spectral seed components in an OPCPA. Although this was not pursued in this work, the introduction of a frequency-doubling nonlinear crystal between the RA and the OPCPA may have the effect of converting some of the most intense part of the pump envelope into second harmonic, while leaving the less intense parts of the envelope largely unaffected. While offering potential benefits in conversion efficiency and bandwidth, such a scheme would require more pump energy than is currently available - a problem potentially solved with a booster amplifier. Looking towards the future of OPCPA pump lasers, hybrid architectures may offer promise for the development of stable, high pulse energy and high average power systems. For instance, one of the main drawbacks of the thin-disk architecture is that it often requires >100 passes through the crystal. On the other hand, these systems are often seeded by pJ or nJ level pulses. While large mode area Yb-doped fiber (or rod) type amplifiers are currently limited to about 1 mJ of pulse energy, and are therefore less useful for directly pumping an OPCPA in our energy regime, such amplifiers could act as a preamplifier for a thin-disk or other amplifiers, dramatically reducing the required number of passes, thereby simplifying the system and improving stability. Secondly, it is worth mentioning that the optimum pulse duration to pump an OPCPA is ultimately an optimization problem which considers the damage threshold of the nonlinear crystal, available aperture sizes of the crystal and target bandwidth and pulse energy for the amplified signal or idler. While short pump pulses enable higher peak intensity, and therefore the use of shorter and lower d_{eff} nonlinear crystals, they also decrease the total pulse energy available before damage - the latter, of course, being the premise of chirped pulse amplification. Tuning the pump pulse duration could be readily achieved by adjusting the spacing between grating pairs in a pulse compressor. This optimization will be considered in

greater detail in Appendix A.

4.4 Dispersion control in the OPCPA

The physics of spectral phase was covered in detail in Chapter 2. In this section, we will describe the technology that allows us to precisely control the spectral phase of the signal and idler to enable chirped pulse amplification, resulting not only in transform limited 3 μm pulses, but also pulses of arbitrary, programmable spectral phase.

In Section 1.5, we introduced chirped pulse amplification, but did not go into detail in how spectral phase is actually applied to stretch and compress pulses. Ideally, in a well balanced chirped pulse amplifier, the various orders of spectral phase applied by the stretcher exactly cancel the spectral phase applied by the compressor: for instance, a Martinez-style grating stretcher can be aligned to apply equal and opposite GDD, TOD and FOD of a grating pair compressor acting on the same pulse. A summary of a few important dispersion control devices is presented in Table 4.1:

	GDD	TOD
Grating pair compressor	-	+
Martinez-style grating stretcher	+	-
Grism	-	-
SLM pulse shaper	+/-	+/-

Table 4.1: **Sign of the GDD and TOD for a grating pair compressor, Martinez-style grating stretcher and grism.** Note that -/- dispersion of the grism is the configuration used in this experiment, and may be tuneable with careful design.

There are three dispersion controlling elements used in the OPCPA after the front end. The first optical assembly seen by the output of the 1.5 μm arm of the front-end is a pulse shaper that uses a 2D spatial light modulator (SLM). Immediately following the pulse shaper, the 1.5 μm seed is stretched using a grism stretcher. The shaped and stretched 1.5 μm pulse then seeds the chain of

PPLN OPAs. After amplification, a double pass Treacy-style reflective grating compressor is used to compress the 3 μm pulse for immediate use in HHG or other experiments.

4.4.1 Programmable SLM-based pulse shaper

While the stretcher and compressor can generally be adjusted to remove GDD and TOD in the compressed pulse, higher order dispersion is mostly not accounted for - nor is spectral phase from nonlinear processes or propagation through material. Furthermore, fine-tuning the compressor to remove the final few radians of spectral phase can be time consuming. Additionally, not all nonlinear processes necessarily prefer a transform limited pulse. For instance, HHG was optimized using a deformable mirror pulse shaper, resulting in a temporal envelope of the driving pulse with a slight prepulse [175]. The use of a pulse shaper can be used to address each of these issues. While the 2D SLM can in principle shape a pulse in both amplitude and phase, we generally use it exclusively as a phase shaper. On a more fundamental level, it is worth commenting on the relative importance of phase and amplitude in describing optical pulses. When a real-valued electric field waveform is Fourier transformed, the resultant complex valued function contains both real and imaginary components, corresponding to the spectral amplitude and spectral phase, respectively. The SLM operates on the pulse in this spectral (Fourier) domain, with the ability to manipulate both spectral phase and amplitude. When a real-valued function, such as the electric field amplitude in the time domain, is reconstructed from a Fourier-domain phase and amplitude, the phase is in some sense the more important quantity. Among other demonstrations, this can be seen in few-cycle pulse experiments, in which the spectral amplitude might be highly structured, but the temporal envelope is smooth so long as the spectral phase is relatively flat [105]. Nevertheless, amplitude shaping can play an important role in the seeding of high gain amplifiers, as will be discussed below.

The pulse shaper used in this system is described in Reference [150]. An annotated photograph of the pulse shaper is shown in Figure 4.6. Light from the EDFA is sent through a polarizer beam-splitting (PBS) cube, a half-wave plate set to apply a polarization rotation of 45° and a Faraday rotator. In this configuration, light exiting the fiber collimator and propagating through the




Figure 4.6: **An annotated photograph of the SLM-based 1.5 μm pulse shaper.** The 940 grooves/mm grating is pictured above. The shaper was later upgraded to include a 600 groove/mm grating, decreasing the angular dispersion and enabling transmission of the full EDFA spectrum.

three optics retains a polarization parallel to the table, whereas light traveling the opposite direction experiences a 90° polarization rotation and is ejected by the PBS. Although it is possible to use slight angle offsets to eject shaped pulses out of the shaper, such designs typically introduce spatial chirp. After the rotator, the beam is diffracted by a 600 grooves/mm transmission diffraction grating (part 600/1550-25.4 from Wasatch Photonics). The grating is placed in the back focal plane of a 2", $f=10$ cm lens, which acts to make the various spectral components colinear and causes them to focus on the SLM, which is in the front focal plane of the lens. The SLM (part SLM-200 from Santec) applies a programmable phase mask and retroreflects the beam back through the system. Power before and after the pulse shaper was measured at 218 and 102 mW, respectively, for a transmission of 47%. Note that the photograph in Figure 4.6 depicts a slightly out of date iteration, which used a 940 grooves/mm grating and therefore had a lower transmission bandwidth of ~ 100 nm. In order to increase the pulse shaper bandwidth without changing the grating line density, we experimented with the use of a shorter focal length $f=7.5$ cm lens, but that resulted in as much as a few tens of radians of added spectral phase after the pulse shaper - likely originating

from spherical aberration.




Figure 4.7: **Algorithmic phase optimization of the compressed 3 μm pulse with the pulse shaper.** A second harmonic signal was used as a fitness metric, resulting in a near transform limited pulse. (Left) FROG temporal reconstructions of the pulse before optimization (blue) and after optimization (orange), compared to the Fourier-transform limit of the spectrum (shaded). (Right) FROG reconstructions of the spectral intensity (blue) and spectral phase (orange) of the optimized pulse, compared with the measured spectrum (shaded). We published this figure in Reference [150].

Code was written by Will Hettel to algorithmically optimize the spectral phase imparted by the SLM by using any feedback as a fitness metric. Two algorithms were implemented (freezing and stochastic parallel gradient descent, or SPGD), with no discernible difference in performance other than SPGD converging faster. An example of algorithmic phase optimization is shown in Figure 4.7. In this example, second harmonic generation was performed in a 400 μm thick AgGaS_2 crystal and measured with a fast photodiode. The photodiode signal was used as a fitness metric for the phase optimization algorithm. After optimization, the peak intensity of the 3 μm pulse increased by 63%, primarily by bringing amplitude in from side pulses. An experiment was also performed using the brightness of the HHG beam as a fitness metric, by integrating a region of interest on the X-ray CCD camera. While the fitness metric only improved by about 10%, we note that the FROG reconstruction of the resultant temporal profile of the HHG optimized driving pulse almost exactly matched the temporal profile observed by Bartels et al. in a similar experiment [175].

Owing to the horizontal dispersion of the transmission grating used in the pulse shaper, a variation in the applied phase in the horizontal dimension of the SLM results in the application of

spectral phase to the pulse. Phase mask variations across the vertical dimension of the 2D SLM are not employed for the purpose of phase shaping, but they can be used for amplitude shaping through the application of a periodic phase structure across the vertical dimension. This periodic structure diffracts light out of the plane of incidence, thereby removing spectral amplitude at a user-defineable spectral location. The modulation depth of the applied vertical phase grating can be used to tune the diffraction efficiency and therefore change the rejection ratio of a particular spectral component. Due to the large amount of GDD applied by the stretcher, the temporal envelope of the stretched pulses almost exactly matches the spectral envelope (see, for instance, Figure 4.9). In the literature, amplitude shaping the seed of an OPCPA was used to improve the overall gain, conversion efficiency and spectral bandwidth [176]. Although amplitude shaping is not currently employed in our day to day operations, preliminary results have been positive. For instance, the SPGD algorithm was applied to optimize the bandwidth of the amplified signal pulse by shaping the amplitude of the signal seed pulse. This resulted in a “U-shaped” signal seed spectrum, and a 70% increase in the amplified signal bandwidth, but the configuration was not usable at the time due to excessive attenuation of the seed amplitude. We propose an alternative approach to amplitude pulse shaping that is additive, rather than subtractive, and to our knowledge has never been demonstrated in theory or experiment. In the proposed approach, the seed pulse is spectrally dispersed and collimated (much the same as in the pulse shaper described here). However, instead of retroreflecting from an SLM, the spectrally dispersed seed interacts with a fan-out type PPLN. The fan-out PPLN and spectrally dispersed seed are aligned such that each spectral component of the seed encounters a phase-matched poling period. This scheme has been demonstrated and is sometimes referred to as frequency-domain OPA (or FOPA). A pump pulse is then dispersed using a cylindrical lens onto a SLM. The SLM can be used to amplitude shape the spatial profile of the pump, and in that way, selectively amplify user-defined spectral components by user-defined amounts. Since the gain in an early stage PPLN OPA can easily be $>10^4$, an extremely high dynamic range may be possible in this amplitude shaping scheme, while avoiding the problematic low-seed amplitude traditionally associated with subtractive amplitude shaping. This technique

has not been realized in the lab, but may find usefulness in future systems.

4.4.2 Grism stretcher

When optical parametric amplification is pumped by a narrowband, quasi-unchirped pump, and seeded by a highly chirped seed (as is the case in the system presented here), the phases of the signal and idler are necessarily inverted, mirror images of one another: while the magnitude of each term of the spectral phase stays the same, the even terms of the spectral phase (GDD, FOD, etc) switch sign, while the odd terms (TOD, 5th order, etc) keep the same sign. This can intuitively be understood by considering the photon energy picture of OPA. If a signal pulse is stretched with negative GDD (as is the case here), the blue frequencies are in the leading part of the pulse and the red frequencies in the trailing part of the pulse. If the pump is considered monochromatic, then, by energy conservation, the bluer signal will necessarily result in a redder idler, and vice versa. Therefore, the generated idler will lead with the red frequencies and trail with the blue frequencies: the opposite GDD of the signal. In contrast, TOD yields a frequency chirp with inversion symmetry about the center of the pulse, and therefore TOD is unchanged between the signal and idler.

The phase conjugation poses a particular problem for OPCPAs that seed with the signal beam and compress the idler: in these systems, the commonly used stretcher/compressor pairs in standard CPA systems will not properly compensate the phase. An earlier version of the system reported here was only able to achieve ~ 420 fs pulses as a result of this challenge [156]. The solution we have chosen is to stretch the signal with a double-pass transmission grating grism pair and to compress the idler with a standard Treacy double-pass reflective grating pair compressor, an arrangement analyzed in detail in Reference [177]. Grism stretchers were originally designed by Tournois as a way to apply negative GDD with zero TOD [178]. The scheme was later modified to yield both negative GDD and TOD [179] - of particular interest in order to compensate material dispersion (from, for instance, many meters of fiber), which in the sub-1100 nm spectral range generally yields positive GDD and TOD.

The grism stretcher used in this work is illustrated in Figure 4.8. Of the 102 mW at 100 MHz

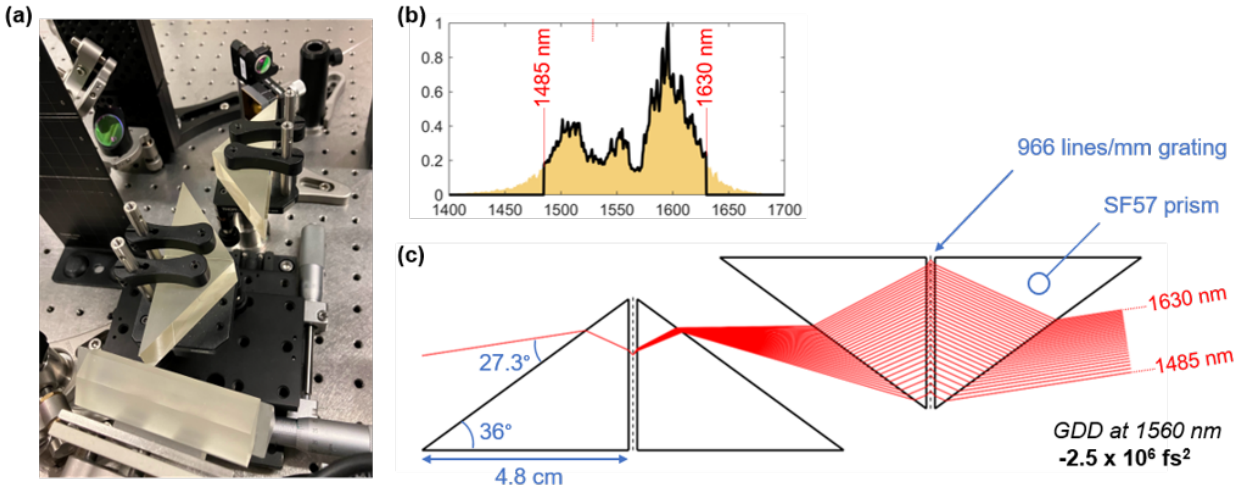


Figure 4.8: **The 1.5 μm grism stretcher.** (a) A photograph of the stretcher configured to amplify 100 fs compressed pulses. The photographed configuration matches the ray tracing simulation in (c), resulting in predicted temporal pulses depicted in Figure 4.9, and matches the compressor configuration described in the following section. Increasing the distance between the subunits increases the GDD but reduces the bandwidth. The incident angle is 1.7 deg off Brewster. Improvements to the bandwidth of the seed EDFA resulted in clipping of the spectrum in the grism - in this case between 1485 and 1630 nm - which we are actively addressing with larger prisms and gratings. This results in a tradeoff between seed bandwidth and GDD, which influence both the compressed pulse duration and the conversion efficiency in the OPAS by changing the temporal overlap. The GDD predicted by the ray tracing code includes both forward and backwards passes through the stretcher. The code was developed by Michaël Hemmer.

that enters the stretcher from the pulse shaper, about 46 mW can be measured at the output, for a transmission of 45%. The stretcher is composed of two identical subunits. Each subunit contains two identical right angle uncoated SF57 prisms (from Tower Optical) that are 20 mm tall. The two side lengths adjacent to the right angle are 48.5 and 35.2 mm, forming a $54^\circ\text{-}90^\circ\text{-}36^\circ$ triangle. Sandwiched between the two short-leg faces is a 966.18 lines/mm, 0.675 mm thick transmission grating (part T-966C-35X20-94 from LightSmyth/Finisar/II-VI/Coherent). The grating is held in place purely by mechanical force. The 1.5 μm beam leaving the pulse shaper enters the stretcher slightly off Brewster angle, for reduced reflection losses. The beam enters high, propagates through the two subunits, is brought down in height by a roof mirror, propagates back at a lower height, and is ejected by a gold mirror towards OPA1. By adjusting the relative position of the two subunits, the applied GDD and TOD can be tuned: increasing the distance between the two subunits increases

the applied GDD, but it also decreases the maximum bandwidth allowed before clipping. A ray tracing code was used to simulate and optimize the dispersion. Figure 4.8(c) shows a ray tracing simulation for the OPCPA configuration analyzed in this chapter, in which a GDD of $\sim -2.5 \times 10^6 \text{ fs}^2$ is applied to the seed. As can be seen, the second subunit is insufficiently large to accommodate the full bandwidth while still applying sufficient chirp to the seed. We have also experimentally found that some components near the clipped edges of the transmitted grism spectrum do not amplify in the OPAs. The ray tracing code was used to verify that the maximum GDD achievable for a given bandwidth scales with the linear dimension of the prisms. With this understanding, we are currently in the process of replacing the second grism subunit with prisms and gratings that are twice as large. With this new, large bandwidth grism, we will be able to apply an appropriately large stretch factor without clipping the seed spectrum from the EDFA. For an illustration of the current balance between bandwidth and stretch, see Figure 4.9.

4.4.3 Grating pair compressor

The grism stretcher was chosen because it ultimately results in a stretched idler pulse with the correct phase (positive GDD and negative TOD) to be correctly compensated by a Treacy-style grating compressor. Although a variety of pulse compression schemes exist, the combination of femtosecond, millijoule pulses, mid-IR wavelength and relatively large amount of GDD (few 10^6 fs^2) make a grating compressor a natural choice. While bulk material compressors have been used with high-efficiency, the large amount of GDD would require, for instance, 5-10 m of propagation through silicon. As an interesting aside, we note that the actual origin of spectral phase in a grating compressor is often misunderstood to be a simple matter of red frequencies traveling a different path length than blue frequencies, where the relative delay due to the path length difference accounts for the entirety of the pulse compression. In fact, this simple ray tracing approach fails to accurately predict the GDD, something that was recognized in Treacy's original 1969 paper [180], where the applied phase is described in terms of a "geometric term" and "correction term," the latter originating from the nature of grating diffraction. When light diffracts from a grating into the first




Figure 4.9: **Calculations of the OPCPA pump and signal seed in the time domain.** Few hundred ps pulses are somewhat inconvenient to measure in time, as they push the limits of fast photodiode technology and are generally too long for FROG techniques. Instead, we apply known spectral phase to measured spectra to estimate the time domain envelope. The dispersion on the pump spectrum (black line) is 1.5 ns/nm and arises from the front-end CFBGs; the dispersion on the 1.5 μm seed spectrum primarily comes from the grism stretcher. The geometric configuration of both the stretcher and the compressor verify that the applied GDD is $\sim -2.5 \times 10^6 \text{ fs}^2$. We also note that the stretch factor has evolved as the system bandwidth has been improved, and is a balance between desired gain bandwidth and conversion efficiency - therefore, this configuration may not be representative of all results presented in this thesis.

order, the angle of diffraction is set by the condition that neighboring lines emit radiation into that direction that is exactly 2π out of phase, which creates constructive interference and bright radiation in that particular direction. Since many grooves are illuminated, it is necessary to add 2π times the number of illuminated grating grooves to the geometrically calculated phase. Although rather interesting, we don't generally need to employ this observation as we calculate the spectral phase from a grating compressor.

It is very useful to be able to quickly and accurately calculate the GDD and TOD from a grating compressor for arbitrary conditions. Although some amount of dispersion originates from material and nonlinear phase between the stretcher and compressor, these values are generally much smaller than the dispersion applied by the stretcher and compressor. For this reason, and due to

the large and rather simple geometry of a grating compressor, simply measuring the optimized compressor is arguably the best way to accurately characterize the amount of stretch on the pulses in the OPCPA - a critical parameter for optimization. The equations describing GDD and TOD of a double-pass grating compressor from Reference [181] have been used by the author and are considered accurate. To use them, simply make note of the angle of incidence, θ , of the collimated beam onto the first grating, off-normal and in radians; the line spacing, d , in meters (by inverting the known groove density); the *perpendicular* distance, L , between the gratings, depicted in Figure 4.10 and not to be confused with the center to center distance, and apply these values to the following formulas:

$$\text{GDD} = -\frac{\lambda^3 L}{\pi c^2 d^2} \left[1 - \left(\frac{\lambda}{d} - \sin \theta \right)^2 \right]^{-3/2} \quad (4.1)$$

$$\text{TOD} = -\text{GDD} \frac{3\lambda}{2\pi c} \left[1 + \frac{\lambda}{d} \frac{\lambda/d - \sin \theta}{1 - (\lambda/d - \sin \theta)^2} \right], \quad (4.2)$$

where λ is the central wavelength, c is the speed of light, and the resultant GDD and TOD are in units of sec^2 and sec^3 , respectively. It is apparent from inspection of Equations 4.1 and 4.2 that the magnitude of the GDD can be adjusted by changing the perpendicular spacing between the gratings L , and the TOD/GDD ratio can be set by changing the angle of incidence θ . Despite this tuneability, it is often difficult to produce a perfectly transform-limited pulse from a compressor; the previously described pulse shaper is sufficient to remove any residual phase.

A photograph of the compressor used here is shown in Figure 4.10. The compressor is composed of two 300 lines/mm reflective gratings (part 53006CU06-780R from Newport/Richardson), with a blaze angle of 26.75° , blaze wavelength of $3.12 \mu\text{m}$, on a $58 \times 58 \times 10 \text{ mm}$ copper substrate with a protected silver coating, and a metallic roof mirror. The primary issue with the compressor is the low transmission, which is typically around 65%. The losses stem from sub-optimal diffraction efficiency, reflectivity losses due to the grating coating (gold would likely be better), low reflectivity of the roof mirror (dielectrics would likely be better), and potentially even losses associated with a collimating lens. Higher efficiency diffraction gratings can also be achieved using multilayer dielectric reflective gratings or dielectric transmission gratings, but both of these options




Figure 4.10: A photograph of the 3 μm grating compressor with approximate representation of beams overlaid. The perpendicular distance, L , used in Equation 4.1, is depicted in yellow ($L = 25.75$ cm for the compressor shown). This compressor configuration compensates the dispersion from the grism depicted in Figure 4.8.

would require very expensive, fully custom solutions. The specific geometrical configuration of the compressor changes depending on the configuration of the front-end (bandwidth) and the grism (stretch factor). The configuration depicted in Figure 4.10 has a perpendicular distance $L=25.75$ cm and an angle of incidence of $\sim 37^\circ$. Applying these values to Equations 4.1 and 4.2 results in a predicted $\text{GDD} = -2.5 \times 10^6 \text{ fs}^2$ and $\text{TOD} = 1.5 \times 10^7 \text{ fs}^3$. Figure 4.9 shows a time domain plot of the 1.5 μm signal spectrum with this applied GDD.² We also plot the predicted time domain pump pulse, which is calculated by applying a stretch factor of 1.5 ns/nm to the measured spectrum, as well as the bandwidth limitation of the current configuration of the grism stretcher.

Various post compression schemes have been proposed and demonstrated to increase the peak power of pulses in the mid-infrared. Carlson et al. numerically investigated nonlinear mid-IR pulse compression using a multipass cell containing thin dielectric plates [182]. Sub-two cycle 3.2 μm pulses have been generating using two thin silicon plates for post-compression [183]. The use of a

² The TOD is not applied in the calculation because it has a negligible contribution.

gas-filled anti-resonant hollow core fiber has been demonstrated with 75 μJ pulses at 3.1 μm [73]. The use of thick dielectric materials has also proven effective. Hemmer et al. demonstrated self compression at 3.1 μm from 70 to 30 fs using a 3 mm thick YAG plate [184]. In the high pulse energy regime, a 2 mm YAG plate was used to self compress 21 mJ, 3.9 μm pulses from 90 to 30 fs with a transmission of 95% [185]. While not implemented in this work, nonlinear post compression remains an intriguing option to increase the peak intensity for HHG.

4.5 Four stage PPLN optical parametric amplifier

The analytical nonlinear optical theory describing optical parametric amplification was developed in Chapter 2. Later, in Appendix A, we will simulate some practical elements of OPA in both periodic and bulk crystals. In this section, we aim to describe the OPA chain as it exists on the optical table, and give some representative experimental results.

4.5.1 Overview

The OPA chain is constantly being improved in order to optimize for amplified pulse bandwidth, compressability, and pulse energy. Therefore, the best that we can do is to provide a snapshot of the system, with explanations of how it got to its current state and where it is going. Over the course of our investigation, some elements were held constant: the regenerative amplifier afforded roughly the same, ~ 18 mJ, 275 ps FWHM pulses, and the only crystal used in the system was PPLN. The seed spectrum, on the other hand, has evolved several times - forcing modifications to allow for the larger bandwidth, as described in the dispersion management section, but also a rethinking of the OPAs. In general, it is not so difficult to get some parametric amplification out of PPLN; on the other hand, optimization of the OPAs can be an extremely meticulous process. At every stage of the amplifier, one must select an appropriate pump spot size by carefully balance the damage threshold of the crystal, the desired peak intensity and the available intensity for seeding. Also at every stage, one must select a seed spot size that optimizes conversion efficiency, noting that the seed mode evolves dramatically as it is amplified. One must find the optimum relative

pulse duration of the seed and the pump, ensuring the wings of the pulse see sufficiently high gain to maintain bandwidth, while also noting that too short of a seed leaves pump energy unused, decreases conversion efficiency and can even amplify OPG. Even if all of these factors are perfectly implemented, there are still the practical matters of ensuring spatial overlap and colinearity of the pump and the seed, adjusting delays for temporal overlap, tuning the crystal temperature for phase matching, and ensuring cleanliness of the seed pulse at each stage by filtering out either the signal or idler between stages.

Figure 4.11 illustrates a recent “snapshot” of the system. At the point in the project depicted here, harmonics had already been demonstrated with ~ 130 fs, $680 \mu\text{J}$ $3 \mu\text{m}$ pulses. After generating harmonics, motivated by analytical calculations and numerical simulations that suggested that the PPLN crystals could amplify significantly more bandwidth (see Appendices), we pursued a careful study and realignment of the system. The seed EDFA pump current was brought up in order to increase the seed bandwidth. Figure 4.11, and the associated results, is a snapshot of the system at this time. As we attempted to compress the pulses originating from this configuration, it became clear that the spectral phase was preventing total compression: while the bandwidth from OPA4 should in principle have supported sub-80 fs pulses, the best compression achievable was closer to 110 fs - with a heavily wrapped phase mask on the SLM pulse shaper. It was soon discovered that the EDFA seeding the OPCPA was overdriven - in fact, a FROG of the front-end seed pulse was practically unachievable, suggesting unexpected spectral phase. Shortly after, the length of PM1550 fiber after the EDFA was reduced to its shortest practical length, and the currents were further tuned - this time, with a careful eye to the resultant FROG. A 35 fs, nearly transform limited pulse was generated in this new configuration. Although some of the seed bandwidth is now clipped in the grism stretcher, a rough alignment of the system was done (illustrated in Figure 4.14), resulting in compressed ~ 100 fs pulses (illustrated in Figure 4.16) which have driven HHG. Through careful optimization of the OPAs, as well as a stretcher upgrade, we are now targeting compressed pulses significantly below 100 fs.

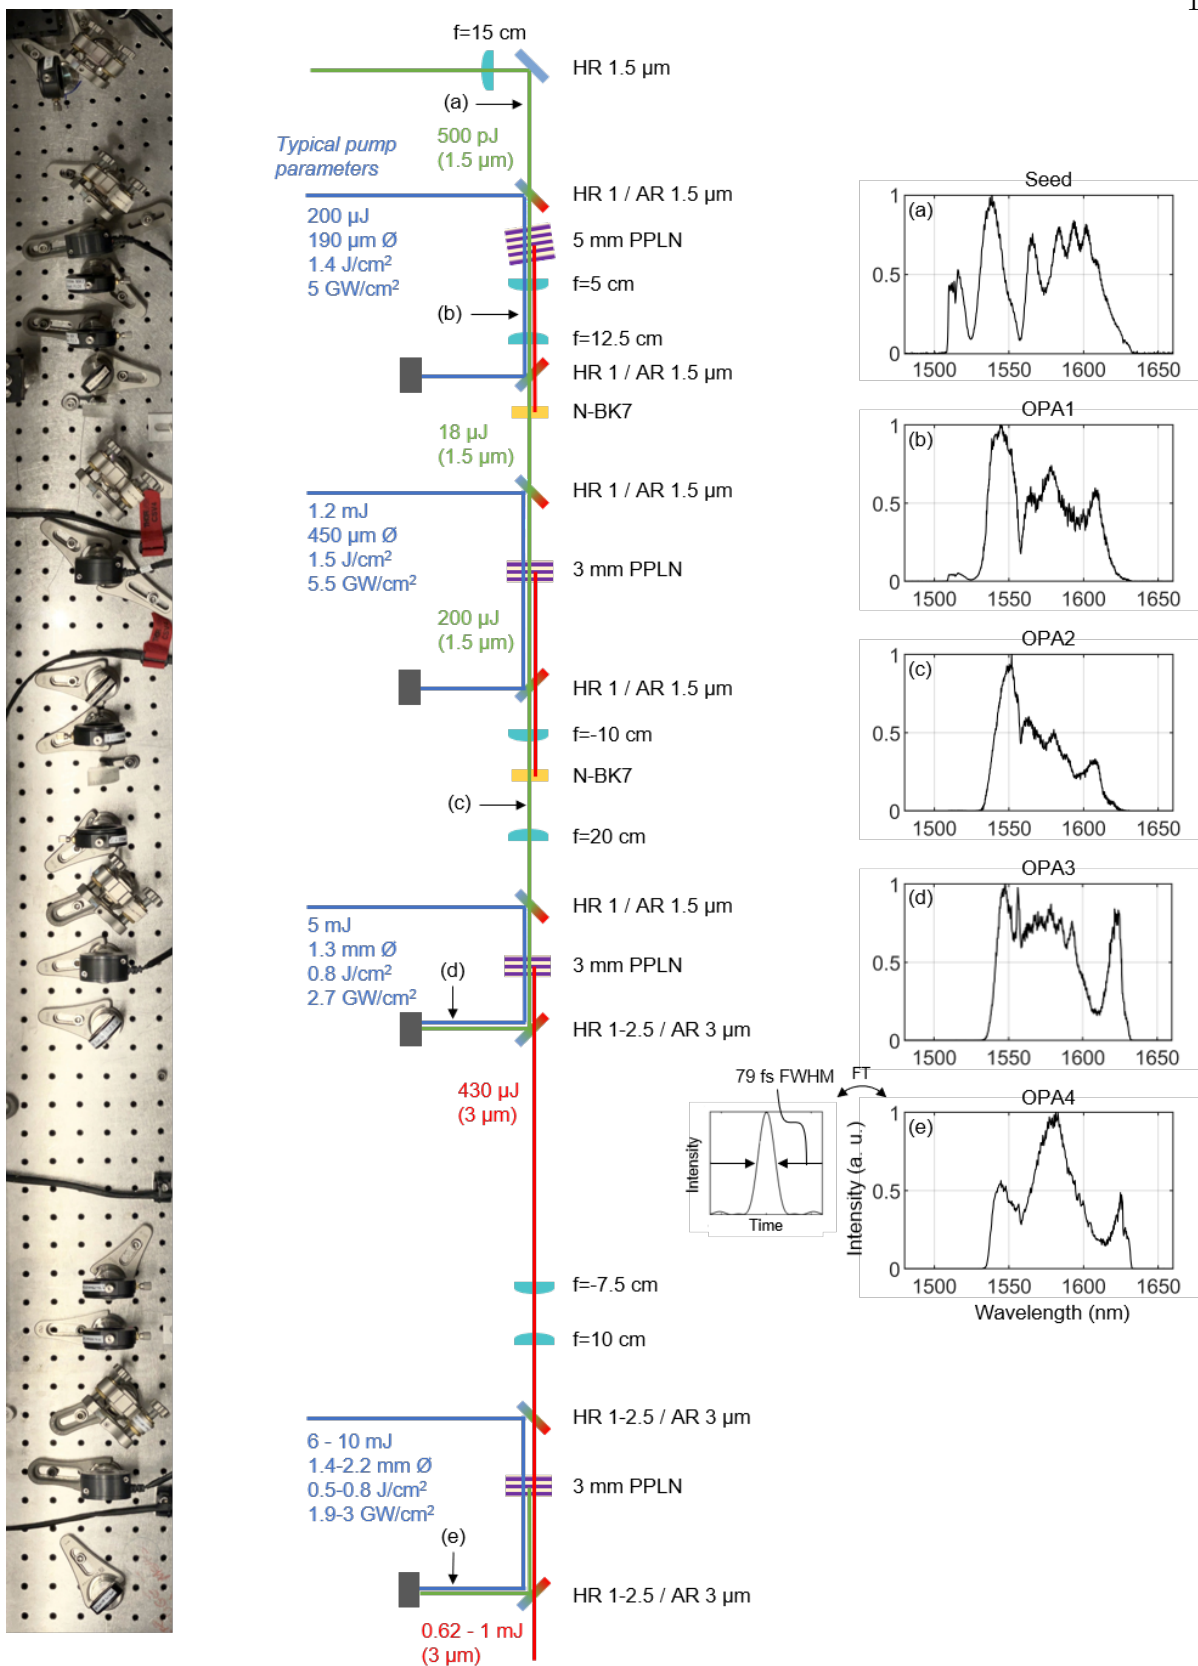


Figure 4.11: A photograph and annotated schematic of the OPA chain, with a “snapshot” of input and output parameters. Note that the seed spectrum has since been modified.

4.5.2 Stage by stage performance

We will now walk through the OPA chain, highlighting important results and considerations at each stage. All four stages utilize anti-reflective coated 5% magnesium oxide doped periodically poled lithium niobate (MgO:PPLN) crystals which are all typically heated in the range of 100 - 130 C, with the exact temperature dependent on the spectrum being amplified. All four crystals are mounted in ovens from Covesion (United Kingdom). Each oven contains a heating element and a thermometer, and is controlled by an external unit allowing for precise temperature control. The ovens are sealed aside from entrance and exit apertures for the laser. While it is possible to run the PPLN ovens without their lid, we have found that, in addition to improving temperature stability and uniformity, the downward force applied by the lid is important for consistent thermal contact. The first OPA stage is typically seeded by 350 - 500 pJ of pulse energy at 1.5 μm , corresponding to 35-50 mW at 100 MHz. It is colinearly combined with the pump beam using a dichroic optic that is anti-reflective coated at 1.5 μm and highly reflective at 1 μm . The OPA1 pump is typically focused to a $1/e^2$ diameter of 190 μm and has 200 μJ of energy. This results in a peak fluence of ~ 1.4 J/cm² and a peak intensity of ~ 5 GW/cm². We note that this fluence is apparently tolerable to the PPLN only due to the small beam size/low average power - peak fluences of 1-1.2 J/cm² have caused damage in the third and fourth OPAs, which see much more pulse energy. Nevertheless, OPA1 has withstood this fluence for multiple years. The author tends to target a seed size ~ 1.4 times larger than the pump size in all the OPAs, striking a balance between maintaining sufficient seed intensity across the pump mode for efficient and uniform pump extraction and not wasting seed energy on unpumped regions of the crystal. The first PPLN is 5 mm long and contains 5 different poling periods (from Covesion). We operate on a region of the crystal with a nominal poling period of 29.98 μm , and heat the crystal to around 130 C - typically about 30 C warmer than the 30.1 μm poling periods used in the later stages. OPA1 is by far the highest gain OPA of the system, with a total gain of around 3.6×10^4 or 46 dB. One consequence of this extremely high gain is that backreflections at the exit surface of the crystal can reseed the OPA, leading to satellite

pulses that steal energy from the main pulse. This is particularly problematic with long pump pulses, which occupy a length in free space of around 8 cm - so that backreflected light has ample time to be amplified. In our normal operation, we tilt the OPA1 crystal by around 8° off-normal. This eliminated the issue of backreflection, but it also appears to significantly increase the gain bandwidth of the OPA - which would otherwise be reduced by the long crystal. The theory of bandwidth enhancement by tilting PPLNs is described in Reference [186].




Figure 4.12: **Optical parametric generation (or OPG) spectra can be used to determine the optimally phase matched signal wavelengths. Here, we see the phase matching shift to longer signal wavelengths as a 5 mm long, normal incidence PPLN with a 30.49 μm poling period is heated.** Spectra were generated by pumping an unseeded PPLN and measured with an OSA. Pumping conditions were identical to the OPA1 pump parameters shown in Figure 4.11.

Optical parametric generation (OPG), also commonly called superfluorescence or parametric fluorescence, in OPA1 was measured by pumping the unseeded crystal and measuring the power in the $1.5 \mu\text{m}$ band. We observed significantly reduced OPG with the tilted crystal compared to a

normal incidence crystal, which could possibly be explained by the reduced influence of backreflections. With OPA1 tilted, we measured 400 nJ of OPG in an unseeded crystal, compared to 18 μ J of amplified pulse energy when seeded. We also note that 400 nJ does not necessarily reflect the total amount of OPG present in the seeded pulse, as described in Reference [160]. OPG originates from a broad spectrum of vacuum fluctuations. Only those spectral components which are phase matched see significant gain in a pumped crystal. Therefore, OPG can be an excellent way to measure the phase matching conditions of a nonlinear crystal. In Figure 4.12, we present a series of OPG spectra from an unseeded, 5 mm long, normal incidence PPLN with a poling period of 30.49 μ m. The center of the OPG spectrum reflects the phase-matched signal wavelength for that temperature. By heating the crystal, we see a red shift in the OPG spectrum, which reflects the increase of the poling period due to the thermal expansion of the crystal.

Depending on the seed spectrum and relative pulse durations, OPA1 typically produces around 10 - 20 μ J of 1.5 μ m pulse energy. Over the course of \sim 1 hour, the mean total output power of OPA1 was measured at 72 mW with a standard deviation of 60 μ W. Assuming that all fluctuations come from the amplified pulse, and not the seed, this corresponds to an RMS instability of 0.33%. We also verified a high mode quality after amplification. This is particularly important in OPAs, because backconversion can result in dramatic modal reshaping. For an example of these measurements of OPA1, see Chapter ???. There is a general lack of small-pixel, affordable cameras in the 1.5 μ m spectral range, which are necessary to analyze the mode structure of few hundred μ m diameter beams. We solved this problem by carefully employing the two-photon response of a standard silicon CCD camera (Mightex), which is insensitive to 1.5 μ m light in the single photon absorption regime. This very useful tool is described in Appendix B.

After OPA1, the rapidly expanding, amplified signal is refocused using a composite system of two lenses onto OPA2. Between OPA1 and OPA2, and again between OPA2 and OPA3, we insert a normal incidence, C-coated (AR coating for 1050 - 1700 nm), 12 mm thick window segment made of N-BK7 (from Thorlabs) to absorb the idler. Due to absorption in the NBK7, 3 μ m transmission of this window is less than 1%, while the signal transmission was measured at 98%. Early iterations




Figure 4.13: **A 12 mm thick piece of AR coated NBK7 removes the idler light while transmitting the signal light after OPAs 1 and 2.** Two amplified spectra from OPA2 are shown above: one with the NBK7 filter between OPA1 and OPA2, and one without the filter. The spectral modulations are reduced with the idler filter. The dramatic hole in the middle of the spectrum was caused by a deep hole in the middle of the pump spectrum/temporal envelope, which precluded gain in that region. The pump problem, which originated from a circulator failure in the front end, has since been resolved!

of the OPA chain without these filters produced spectra with deep modulations, likely due to the OPA process being seeded by both the signal and the idler. Due to material dispersion between OPA stages, the signal and idler become phase shifted relative to one another. Figure 4.13 shows how the addition of the NBK7 filter mitigated this effect, and decreased the modulations in the amplified spectra.

The remaining three OPA stages use 30.1 μm poling period (single period) PPLNs from HC Photonics. The crystals are each 3 mm tall, 3 mm long and 10 mm wide and are operated at normal incidence. Two examples of spectral evolution of the signal are shown in Figures 4.11 and 4.14. Typical operating parameters are shown in Figure 4.11. OPA2 and OPA3 are both seeded with 1.5 μm light (with the idler removed), while OPA4 is idler seeded. The signal is removed after OPA3 and OPA4 by dichroic optics that are highly reflective in the 1-2.5 μm range, and transmissive to the 3 μm idler. D- or E-coated CaF_2 lenses from Thorlabs are used for the 3 μm beam, and exhibit highly variable transmission, ranging from 93% for both D- and E-coatings (the same as uncoated CaF_2) to as high as 98% for a D-coated lens. A primary issue with OPA3 and OPA4 is sub-optimal conversion efficiency. In general, we target $\sim 10\%$ pump to idler energy conversion and

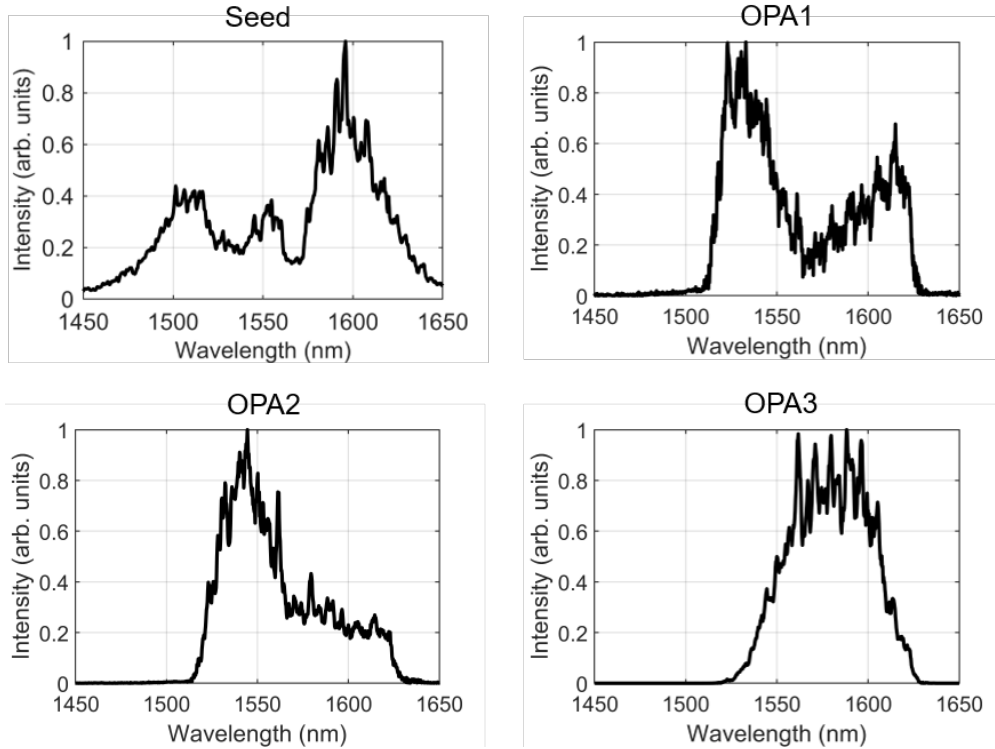


Figure 4.14: **A series of signal spectra from the OPCPA using an improved configuration of the seed EDFA.** These spectra are representative of OPA chain configuration that delivered the 99 fs pulse presented in Figure 4.16. We partially attribute the reduction in bandwidth between the seed and the OPA1 output to clipping in the grism stretcher - however, only future experimental measurements can determine the full gain bandwidth of the OPAs.

$\sim 20\%$ pump to signal conversion. In the first two stages, conversion efficiency is not a concern, so long as sufficient bandwidth and a quality beam profile can be amplified. However, the efficiency of the final two stages determines the output power of the system. Using a fixed ~ 1.3 mm $1/e^2$ diameter pump spot size on OPA3, we varied the pump energy from 2.2 to 8 W and characterized the pump to idler conversion efficiency. The conversion efficiencies ranged from 7.2% at 2.2 W to 8.6% for all pump energies in the range of 3.5 to 5 W and then decreased again to 7.5% for the 8 W pump. For this reason, we settled on ~ 5 mJ of pump at 2.7 GW/cm² for OPA3. The greatest conversion efficiency deficit is in the final stage, which is seeded by about 300 μ J of idler energy. We note that lossy optics reduce the seed by more than 30% between OPA3 and OPA4. In this stage, the conversion efficiency was measured to monotonically decrease with pump fluence, from 6.2%

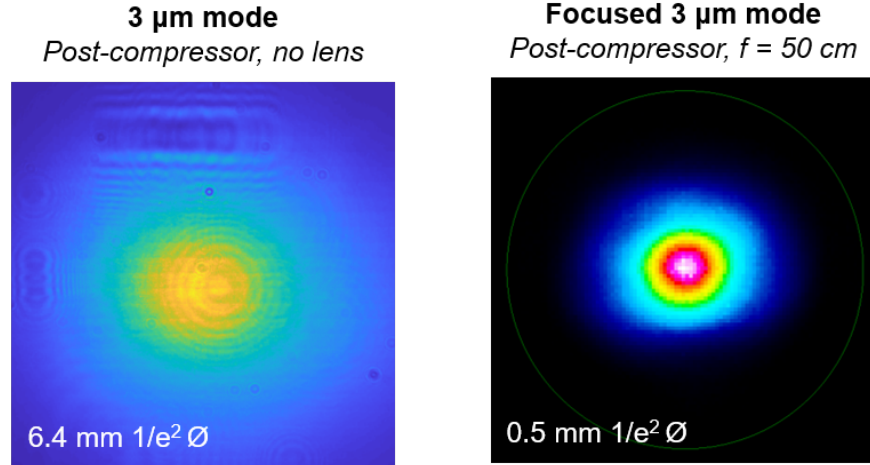


Figure 4.15: **Unfocused and focused modes of the compressed 3 μm beam.** Both images were captured with a VOx microbolometric beam profiling camera (WinCamD-IR-BB from DataRay). The structure at the top of the unfocused mode was later determined to be caused by damage on a post-compressor steering mirror, and was removed for the measurement of the focused beam. The concentric rings on the right side of the unfocused mode are attributed to artifacts from the attenuating filters.

when pumped with 1 mJ to 5% when pumped with 6.4 mJ, while maintaining a constant spot size of ~ 1.4 mm $1/e^2$ diameter. The low conversion efficiency of this final stage ultimately limits our pulse energy for HHG experiments, and is likely responsible for our inability to generate harmonics with helium. Simulations predict that thinner, 1-2 mm crystals would be more appropriate for the low gain OPA3 and OPA4 stages, and that excessively long crystals can result in sub-optimal conversion efficiency. These shorter crystals are soon to be incorporated into the system in an effort to improve both the conversion efficiency and gain bandwidth.

4.5.3 Characterization of the compressed 3 μm pulses

After the dataset for Figure 4.11 was recorded, the front-end seed EDFA was modified to deliver pulses with cleaner phase and more bandwidth. A rough alignment of the OPCPA was done with the improved seed. The evolution of the signal spectrum in this configuration is presented in Figure 4.14, where the dramatic spectral narrowing between the seed and the output of OPA1 is partially attributed to clipping in the under-sized grism stretcher. A higher bandwidth stretcher is

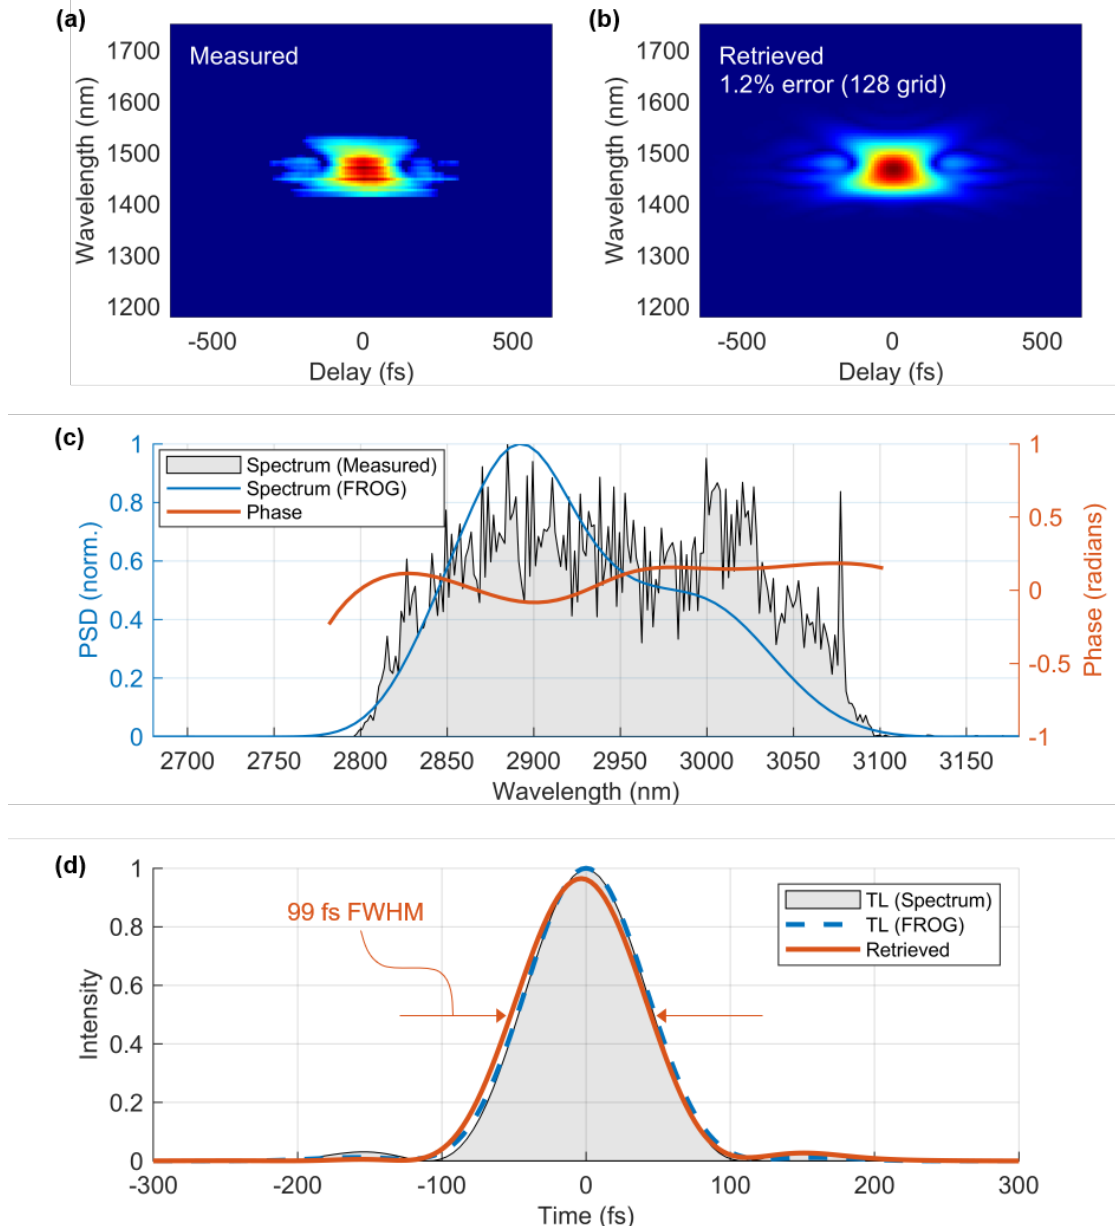


Figure 4.16: **SHG FROG pulse characterization of the compressed 3 μm pulse, showing a near-transform limited 99 fs FWHM duration.** (a) Experimentally measured SHG FROG trace using a home-built FROG apparatus with a 0.4 mm thick AGS crystal. (b) The retrieved FROG trace using a 128x128 grid exhibited a 1.2% retrieval error. Note that both (a) and (b) plot intensity to the 1/3 power in order to highlight low amplitude features. (c) The reconstructed spectral intensity (blue) and spectral phase (orange) is presented alongside the measured spectrum (grey). The spectral phase varies by less than 0.5 radian across the spectrum. (d) The reconstructed temporal intensity envelope (orange) has a FWHM duration of 99 fs.

actively being constructed. The 3 μm pulse ultimately generated from this new configuration was compressed with the help of an iterative phase optimization algorithm on the pulse shaper that used an SHG signal from an AgGaS_2 (AGS) crystal as a fitness metric. A home-built SHG FROG utilizing a 0.4 mm thick AGS crystal for frequency doubling was used to characterize the compressed pulse. The FROG data and retrieval is presented in Figure 4.16. The compressed 3 μm pulse exhibited less than 0.5 rad of spectral phase across the spectrum, resulting in a 99 fs FWHM pulse. The spectrum was measured using a mid-IR spectrometer (S2 from Miriad Technologies combined with home-written readout code) and found to have a transform limit of 94 fs. The reconstructed spectrum (which has a transform limit of 95 fs) is in good agreement with the measured spectrum, and the retrieval error using a 128x128 grid was 1.2%, suggesting a trustworthy reconstruction.

The spatial mode of the compressed 3 μm beam is excellent. In Figure 4.15 we present an unfocused (near-field) profile and a profile in the focus of an $f=50$ cm lens (far-field). The far-field profile exhibits a 95% fit with a Gaussian profile in both the x and y coordinates, calculated by summing the difference between the measured beam profile and a Gaussian and dividing the sum by the area under the curve. The power after the compressor was measured at 620 μJ and the peak power was calculated to be 5.9 GW.

4.6 Distributed parametric amplification for higher pulse energies: numerical simulations of the strong seed regime, and the role of crystal length in conversion efficiency

One approach to amplify energetic signal and idler pulses in the face of a low damage threshold nonlinear crystal is to distribute amplification over a number of crystals. For instance, if the goal is to amplify pulses to 3 mJ idler pulses, and 30 mJ of pump is available, a distributed amplification scheme might involve three crystals in series, each pumped with 10 mJ and generated 1 mJ and seeding the next. A couple immediate concerns come to mind with this approach. The inter-stage transmissive optics, including lenses for beam resizing and dichroic optics for pump and signal filtering, are generally lossy in the 3 μm spectral region. Nevertheless, the 10-30% loss between

stages³ may be an acceptable cost. There is also a significant growth in size and system complexity, as each OPA stage requires its own delay line, as well as material cost, in particular for the large aperture PPLNs. Nevertheless, the requirement of high d_{eff} crystals imposed by the long, ~ 275 ps pump pulses makes this option highly appealing. Ultimately, we have begun pursuing this approach in the final two stages of the OPA chain. While the final OPA stage could in principle be seeded with the signal from OPA3, that would result in wasting ~ 300 μJ of idler pulse energy.

A critical question remains: how does a strong seed influence the onset of back-conversion? It is known from our review of OPA theory in Chapter 2 that the change in amplitude of the pump is dependent on both the amplitudes of the signal and idler, and we further note that the concepts of signal and idler are interchangeable in OPA. We also know, from theory and experiment, that when the pump becomes significantly depleted, back-conversion occurs, in which signal and idler photons combine through the process of sum frequency generation to form pump photons. It is therefore reasonable to ask if a strong seed brings the OPA process that much closer to back-conversion, ultimately limiting the achievable conversion efficiency.

To investigate this, we used *sisyfos* to simulate the fourth OPA stage under typical experimental conditions. In the simulations, the energy of the idler seed is varied from 40 μJ to 16 mJ - or in terms of the pump energy, $0.005E_{\text{pump}}$ to $2E_{\text{pump}}$. Pulse energy data was saved during the simulation at 29 evenly spaced slices along the crystal length. The parameters of the simulation are summarized in Table 4.2.

Crystal	2 mm PPLN ($d_{\text{eff}} = 14.8$ pm/V)
Pump	300 ps FWHM, 8 mJ, 1.4 mm $1/e^2$ diam.
Pump spectrum	T.L. Gaussian at 1.03 μm
Seed	150 ps FWHM, 0.04 - 16 mJ, 1.82 mm $1/e^2$ diam.
Seed spectrum	Gaussian at 3 μm supporting 80 fs T.L. (idler)

Table 4.2: **Simulation parameters used for strong seed simulations.** T.L.: transform limit.

³ This is comparable to what we observe between OPA3 and OPA4, which is idler seeded.




Figure 4.17: **Numerical simulations of OPA in a 2 mm PPLN with variable idler seed pulse energy.** All parameters are the same for each of the four simulations presented above (see Table 4.2), except for a variable seed strength. For each simulation, the 8 mJ pump results in a peak intensity of 3.5 GW/cm^2 and a peak fluence of 1.04 J/cm^2 . The simulations indicate that distributing amplification across multiple crystals is a feasible - if somewhat cumbersome - approach to dealing with the low-damage threshold of PPLN.

Figure 4.17 illustrates four representative simulation results. We plot the pulse energy of the three waves as a function of position in the crystal. We further analyzed the results by measuring the amount of generated idler energy at the peak of the gain - at whatever crystal plane this occurs. An illustration of this analysis is presented in Figure 4.18. We calculated the percent of pump energy that was converted to the idler pulse in the single crystal that was studied, and plotted that percentage of various seed strengths.




Figure 4.18: **Analysis of the strong seed simulation set.** (Left) A representative dataset (1.6 mJ idler seed) is chosen to illustrate how the analysis was conducted. (Right) The analysis was done for each of the 9 seed energies simulated. For each seed energy, the percent of pump energy converted to the idler at the gain maximum is plotted. We see an increasing trend in conversion efficiency as a stronger seed is used and the gain peak occurs over a shorter crystal length. We caution that to fully utilize the improved conversion efficiency, the correct length of crystal must be used.

The result presented in this section suggest that a strong seed does not cause back-conversion to quench efficient pump to idler conversion. In fact, the opposite is observed: strong seeding can result in improved optimal conversion efficiency. This is in line with a trend generally observed in the simulations: when the peak of the gain can be brought forward, to a shallower depth in the crystal - whether through a higher intensity pump or a stronger seed - the optimal conversion efficiency increases. However, this result is not so simple to implement in the lab: the improved conversion efficiency can only be realized if the peak of the gain is placed near the back-surface of the crystal - otherwise back-conversion quickly strips away any improvements in conversion efficiency. Large aperture (3x10.5 mm), thin (1-2 mm) PPLN crystals are currently \$5000-6000 each with a generally long lead time, making a fine-tuning optimization procedure in the lab nearly unfeasible. Nevertheless, the general trends are important to take note of: there are strong benefits to conversion efficiency and gain bandwidth if amplification can be achieved in a shorter crystal, and similarly, the use of an excessively long crystal can significantly reduce conversion efficiency -

even if the seed is reduced to push the peak of the gain curve to the back-facet of the crystal.

It is also interesting to consider this result in the context of published work on the enhancement of OPA conversion efficiency through the attenuation of either the signal or idler waves [66, 67, 68]. The numerical results presented here clearly show that *both* the signal and idler fields need to reach a threshold intensity before back-conversion turns on and begins limiting conversion efficiency - therefore, by attenuation one field, the other could be allowed to grow to intensities that would otherwise trigger back-conversion.

4.6.1 Numerical simulation of conversion efficiency and crystal length in OPA4

A primary bottleneck of the OPCPA is low-pulse energy from OPA4, ultimately stemming from low conversion efficiency. This has prevented us from driving harmonics in higher ionization potential gases such as neon and argon - a critical goal of the project. To understand this low conversion efficiency, we modified the simulation parameters in Table 4.2 slightly to more accurately represent the current state of OPA4: the crystal length was set to 3 mm, the seed energy was set to 300 μJ , and the pump energy was varied from 2 - 8 mJ (0.26 - 1.04 J/cm^2). In Table 4.3, we summarize the results of this simulation:

Pump energy	Maximum C.E.	C.E. at 3 mm	Experimental OPA4 C.E.
2 mJ (0.26 J/cm^2)	15.1% at 2 mm	9.8%	5.8%
4 mJ (0.52 J/cm^2)	15.4% at 1.7 mm	8.2%	5.3%
6 mJ (0.78 J/cm^2)	15.4% at 1.6 mm	7.4%	5%
8 mJ (1.04 J/cm^2)	15.3% at 1.5 mm	7.3%	-

Table 4.3: **Simulation and experimental data on OPA4 conversion efficiency.** C.E.: conversion efficiency, measured by dividing the generated 3 μm pulse energy by the pump energy as shown in Figure 4.18.

The simulation results presented in Table 4.3 suggest the the effect of back-conversion in the overly long 3 mm crystal of OPA4 could easily account for the low conversion efficiency observed

in OPA4. It is important to note that numerical simulations generally offer a best-case scenario for generated pulse energy. In practice, a reasonable goal for pump to 3 μm conversion efficiency is 10%. The simulation results suggest that the conversion efficiency from a 3 mm crystal may be less than half of what is optimally possible. For this reason, the experimentally measured conversion efficiency of $\sim 5\%$ appears to be explainable by the use of an overly long crystal. 1 and 2 mm crystals from HC Photonics will soon be integrated into the system in an effort to address this.

4.7 Stability of the OPA chain

Stability is of paramount importance for lasers designed to drive HHG for applications. Longer wavelength HHG driving lasers in particular benefit from high pointing stability. In part, this is because longer wavelength drivers typically require longer interaction lengths for optimal HHG brightness - a condition that practically only be obtained through the use of a hollow core, gas filled waveguide. Due to the extremely high intensity ($\sim 10^{14}$ W/cm²) near the entrance facet of the fiber, any wander of the focused beam commonly causes damage to the capillary, resulting in poor coupling and decreased HHG brightness. For titanium:sapphire lasers, this is often addressed through the use of active beam stabilization. Good pointing stability is also necessary for high pressure gas cells. The gas cell used in our HHG experiments has an aperture of 150 μm - the small diameter is necessary to reduce conductance of the high pressure gas into the vacuum chamber. Small misalignments can easily cause damage to the entrance aperture of the gas cell. The OPCPA described here delivers excellent pointing stability (shown in Figure 4.19), making active beam stabilization unnecessary. A strong advantage of the OPA architecture is the possibility of very high single pass gain - the entire four-stage parametric amplifier is composed of a single pass that is ~ 1.5 m long. This improves pointing stability compared to multi-pass systems. Pointing instabilities of the compressed 3 μm beam were measured by focusing the beam onto a beam profiling camera (DataRay WinCamD-IR-BB) and measuring the wander of the calculated beam center over the course of 24 hours. The standard deviation of the beam center in the x (y) direction was 1.8 μm (2.1 μm), representing 0.35% (0.43%) of the mean $1/e^2$ diameter of 525 μm (475 μm).

The few micron wander is insufficient to cause loss of coupling or damage to typical HHG setups. As an anecdotal example, we have generated harmonics with the gas cell on Friday, blocked the beam, then unblocked the beam on Monday to find a comparably bright HHG beam. Aside from mitigating damage to the HHG setup, this type of stability is a pre-requisite for a usable HHG source for applications. The relatively low average flux of long-wavelength driven HHG may require spectroscopic or imaging experiments to average over long periods of time - demanding high source stability.

Power stability of the driving laser is also critical for an applications-oriented HHG source. The tunnel ionization rate predicted by the ADK model [187] has a highly nonlinear dependence on the driving laser intensity - therefore, small fluctuations in driving laser power can dramatically impact HHG power. Furthermore, many applications demand good HHG flux stability. For instance, the pump-probe signal from a magneto-optical Kerr effect (MOKE) spectroscopy experiment can have a very low signal to noise ratio - driving some researchers to go so far as to install active noise cancellation in their HHG spectroscopy beamlines [188]. We measured an RMS power instability of 0.97% in the compressed 3 μm beam (shown in Figure 4.19).

Environmental, thermal management and mechanical considerations are all critical to the stability of the OPCPA. Air currents are known to cause instabilities in high power beams, and are reduced through the use of aluminum composite material (ACM) enclosures around all components of the system and beam tubes to connect separate enclosures. Whenever possible, vibration causing and heat producing heavy equipment, such as the cryogenic compressor and roughing pumps, are kept in a service corridor that is separated from the lab. The lab temperature and humidity is carefully monitored (using a TSP01 USB probe from Thorlabs), and fluctuations are minimized to the best of our ability. An effort is made to water cool any significant source of heat that is in thermal contact with the optical table (such as high power beam dumps) to avoid warping. Finally, the beam height is held at 2.5 in (6.35 cm) above the optical table in all areas except the dispersion controlling modules (pulse shaper, stretcher and compressor). Maintaining the beam as low as convenience allows results in increased mechanical stability of optics.

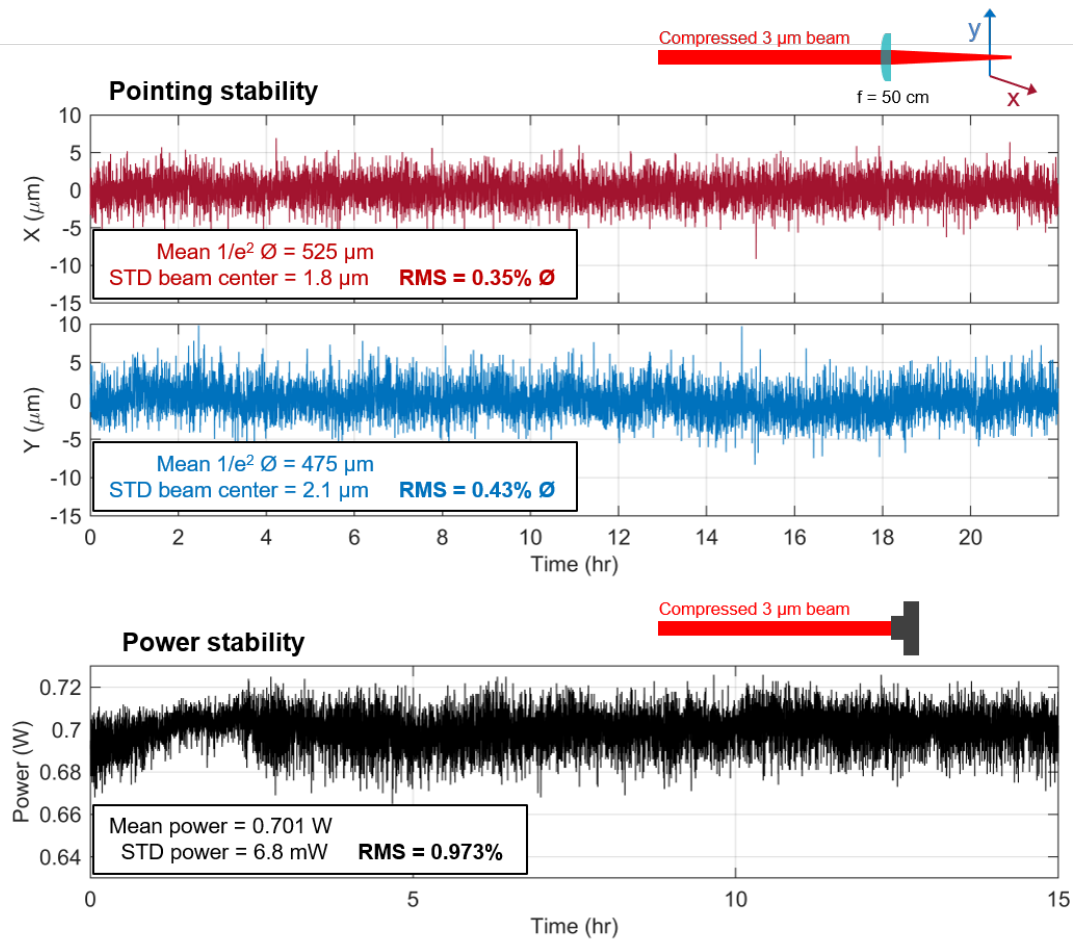


Figure 4.19: **Pointing and power stability of the compressed 3 μm beam.** (Top) Pointing instabilities were measured by focusing the beam onto a beam profiling camera (WinCamD-IR-BB from DataRay) with a 50 cm focal length lens and recording x and y coordinates of the calculated beam center. For a profile of the focused beam, see Figure 4.15. Importantly, the few micron deviation of the focused beam is sufficiently low to maintain good coupling into a hollow core fiber without damage. (Bottom) Power stability was measured with a thermopile power meter (3A from Ophir), with less than 1% instability in the average power of the compressed beam.

Many factors contribute to the stability of the 3 μm OPCPA - the predominant factor being the energy and pointing stability of the regenerative amplifier pump laser, which was reviewed in Section 4.3, and to a lesser degree, the power stability of the 1.5 μm front-end fiber laser. Laser amplifiers (such as the regenerative amplifier) can operate in a saturation regime, in which the output power is more determined by the amount of available energy stored in the gain medium than the seed energy. This benefits power stability significantly. In contrast, there is no saturation

in an OPA in the true sense of the word. Rather, the evolution of the amplitudes of the three waves in an OPA approximately follow the coupled wave equations presented in Chapter 2 (Equations 2.55), where it can be seen that the amplitude of each wave depends on the amplitudes of the other two waves. When the signal or idler grow to a sufficiently strong field compared to the pump, such that the pump is significantly depleted, sum-frequency generation occurs, in which signal and idler photons combine to form pump photons. Thus, unlike in a saturated laser amplifier, increasing the pump intensity on an OPA that is already operating at peak conversion efficiency will actually decrease the signal and idler powers in a process known as back-conversion.

4.7.1 Numerical simulations of OPA stability: output energy sensitivity to pump and seed fluctuations

The literature suggests that the process of back-conversion has important consequences for the stability of OPAs [189, 190]. In particular, Guardalben et al. used numerical simulations to argue that when an OPA gain curve is adjusted such that back conversion begins shortly before the back surface of the crystal, the signal and idler output power is less sensitive to fluctuations in pump intensity. This is because a small variation of the pump intensity causes the peak of the gain to move forward or backward in the crystal. When the pump intensity is increased, back conversion occurs earlier in the depth of the crystal, such that the increased gain due to the higher pump is cancelled by the increased amount of back conversion. The opposite process occurs when the pump intensity is decreased. Here, we use Sisyfos to simulate the fourth OPA stage to understand the effect of crystal length on output power instabilities induced by both pump and seed energy fluctuations. The simulation parameters are identical to those summarized in Table 4.2, except that an idler seed energy of 0.16 mJ is chosen as a central value and the crystal length is 3 mm.

In the simulation shown in Figure 4.20(a), we maintain a constant idler seed energy of 0.160 mJ, and plot the gain curves for pump energies of 7.6, 8 and 8.4 mJ. While the simulations of Guardalben et al. [189] predict a crystal depth shortly past the gain maximum which is almost insensitive to 5% fluctuations in pump energy, we find no such obvious plane, with only the possi-

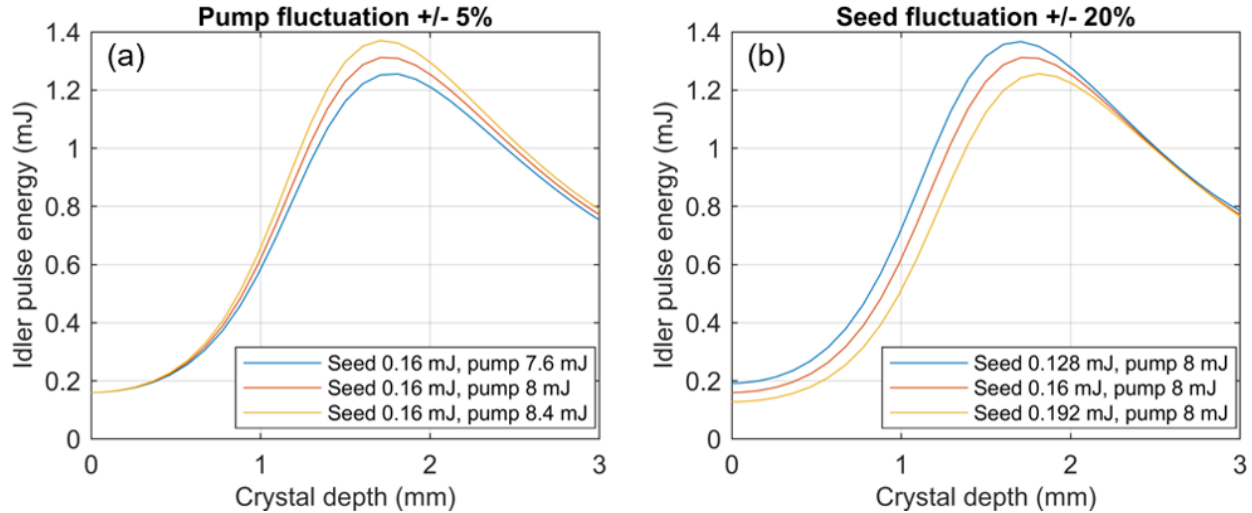


Figure 4.20: **Numerical simulations of the sensitivity of OPA4 to fluctuations in pump and seed energy.** See Table 4.2 for simulation parameters (noting that the idler seed energy has been changed to a central value of 0.16 mJ and the crystal length set to 3 mm). (a) 3 μm pulse energy vs depth for +/- 5% pump pulse energy. In different OPA conditions, other groups have found that locating the back surface of the crystal just beyond the peak of the gain curve makes the generated pulse energy less sensitive to pump fluctuations. In this parameter space, we do not see the same effect. (b) 3 μm pulse energy vs depth for +/- 20% seed pulse energy. Here, the simulation suggest that placing the peak of the gain curve just in front of the crystal back facet may decrease the sensitivity of the output power to fluctuations in the seed energy. This suggests that amplified pulse energy stability from the final OPA stage is mostly dependent on pump pulse energy stability.

bility of very modest stability improvements in the post gain-peak region. We also note that the total gain of the OPA simulated here is dramatically lower than what was studied by Guardalben et al., who studied an LBO OPA with an ~ 25 mm crystal. The result is not particularly problematic for our OPCPA, which exhibits high stability of pump energy. What is perhaps more interesting is the influence of seed energy on the output power. In Figure 4.20(b), we maintain a constant pump energy of 8 mJ and plot the gain curves for seed energies of 0.128, 0.160, and 0.192 mJ - corresponding to fluctuation of +/- 20%. Here, it is immediately apparent that operating the OPA in the post-back-conversion regime significantly reduces the sensitivity of the output energy to seed fluctuations. Interestingly, we find that the generated idler energy converges to the same value for all three seed energies at a particular crystal depth - however, operating so deep in the back-conversion regime would likely result in significant modal reshaping and $\sim 30\%$ sub-optimal

output power. This result agrees with the work of Guardalben et al., who found that $\pm 5\%$ seed fluctuations resulted in output power fluctuations of much less than 1% when operating in the post-back-conversion regime. In our experiment, insensitivity to seed fluctuations is likely more valuable than insensitivity to pump fluctuations, as seed fluctuations can be amplified from one stage to the next without careful design. These simulation results suggest that with a highly stable pump source and a crystal length chosen to allow for a small amount of back-conversion, OPA4 can be optimized to deliver very stable output pulse energies. On the other hand, with an unstable pump, such stability will be difficult to achieve.

The actual gain curves of the parametric amplifiers in the OPCPA presented in this chapter have not yet been optimized in this way for stability, largely because the high cost of large-area PPLN crystals has made length optimization a challenge, and because this level of control over the OPA process is very challenging - nevertheless, it is a future direction worth pursuing.

Chapter 5

3 μm driven harmonic generation

5.1 Introduction

The 3 μm wavelength OPCPA described in previous chapters was used to drive high-harmonic generation in both argon and nitrogen. This represents the second ever report of high-harmonic generation in gases driven by a mid-infrared laser.¹ Using nitrogen, HHG photons with energies up to the carbon K-edge near 284 eV were observed, representing nearly the 700th harmonic of the driving laser. In this chapter, we first describe the theory of HHG in two parts: we begin with the microscopic theory from the single-atom perspective, followed by the macroscopic theory that includes the physics of phase matching and reabsorption. Next, we describe the design, fabrication and testing of a novel, highly modular high harmonic source, that has enabled optimization based on experimental feedback. We then present beam profiles and spectra characterizing the harmonic beam. Finally, we describe preliminary observations of below threshold harmonics and their context in computational nonlinear optics.

5.2 Theory of high-harmonic generation

5.2.1 The microscopic, single atom picture of HHG

In the introduction, we described the semiclassical, three-step model to describe HHG at the level of a single atom. We presented a simple equation for the single atom cut-off energy, which

¹ According to ISO standard 20473:2007, the mid-infrared spectral region spans from 3 to 50 μm .

motivated the development of the long-wavelength OPCPA in this thesis:

$$E_{\text{SA cut-off}} = I_p + 3.17U_p \propto I_L \lambda_L^2. \quad (5.1)$$

Here, we expand this microscopic picture to describe a number of key aspects of HHG, including the characteristic odd-ordered harmonic comb spectrum, the intensity-dependent *intrinsic* or *atomic dipole phase*, and the decrease in single-atom conversion efficiency associated with long-wavelength drivers. While the single-atom approach does not capture the full picture of HHG, it provides necessary context for the subsequent discussion of macroscopic phase matching effects.

A qualitative single atom HHG spectrum is illustrated in Figure 5.1(a). The occurrence of only odd-ordered harmonics from pseudo-monochromatic drivers can be described in several ways. Perhaps most intuitively, this result can be understood by simple Fourier analysis: the HHG process in a single-atom can occur twice per laser cycle, as the symmetry of the medium does not discern between tunnel ionization in one direction or the other. This results in a pulse train of attosecond HHG bursts every half-period in time. When Fourier transformed to the frequency domain, the spectrum is a comb separated by 2ω , where ω is the frequency of the driving laser. We further note that the centrosymmetry of the gaseous medium allows exclusively for the generation of odd-ordered harmonics (see Chapter 2 for why the symmetry of a medium dictates the harmonics that it can produce). Lastly, odd-ordered harmonics are enforced by conservation of spin angular momentum, as described in Reference [191]. We note that two-color driven HHG can produce spectra with both odd- and even-ordered harmonics [192, 193].

The specific time that tunnel ionization occurs after the peak of the electric field is referred to as the ionization phase, and has important consequences for the spatial coherence of the generated beam and phase matching. An electron ionized at precisely the peak of the electric field will have an ionization phase of zero and will return to the parent ion with no kinetic energy, thereby not contributing to HHG. A plot of the kinetic energy of the returning electron as a function of ionization phase is shown in Figure 5.1(b). The maximum kinetic energy of the returning electron, corresponding to the single atom cut-off energy, occurs when the ionization phase is approximately

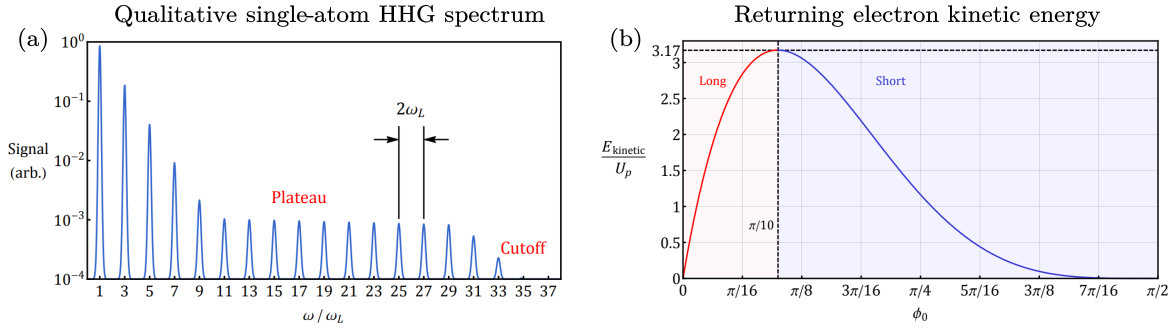


Figure 5.1: **(a) A qualitative single-atom high harmonic spectrum. (b) Kinetic energy of returning electron as a function of ionization phase.** (a) Low-order harmonics are generated in the perturbative regime, followed by a plateau and eventually a cut-off. The width of each individual harmonic relates to the temporal window over which the HHG process occurs. Single color HHG produces odd-ordered harmonics only, and the comb spacing is twice the frequency of the driving laser. (b) The kinetic energy of a returning electron depends on the ionization phase ϕ_0 . The maximum kinetic energy occurs with an ionization phase of $\phi_0 \approx \pi/10$. Ionization at lesser and greater phases causes electrons to undergo short and long trajectories, respectively. Figure adapted from Reference [158].

$\phi \approx \pi/10$ radians after the peak. Ionization phases less than this maximum result in a so-called *long trajectory*, while phases greater than that result in a *short trajectory*, referring to the amount of time the electron spends traveling in the continuum. The trajectories can have important implications for the temporal and spatial coherence properties of the resultant harmonic beam. The phase front curvature of beams generated through the short trajectory is lower than that of the long trajectory, causing reduced divergence and a higher degree of spatial coherence [194, 195].

The ionization phase determines the relative phase of the driving laser and the generated harmonics. This relative phase is often referred to as the intrinsic or atomic dipole phase, and is intensity dependent [196, 197]. The observation that the relative phase of the harmonics and the driving laser depends on intensity has significant consequences for phase matching: in a tight focusing geometry, the variation in intensity through the focus can reduce the interaction region over which phase matching can occur. On the other hand, the relative invariance of intensity over the length of a waveguide mitigates this effect.

Finally, we briefly mention the quantum mechanical nature of high-harmonic generation.




Figure 5.2: **Visualizations of an electron wavefunction driven by a three cycle pulse of different wavelengths.** As the wavelength of a driving laser is increased, the electron wavepacket spends more time in the continuum and undergoes more diffusion, reducing the overlap with the bound state during recombination, and in turn reducing the efficiency of HHG. Reference [158] with source material found in Reference [36].

When the electron is ionized, it transitions from a bound state into the continuum. While the electron is in the continuum, its phase evolves according to the Schrödinger equation. When the electron recombines with the parent ion, the bound state overlaps with the continuum state and causes spatial and temporal interference which modulates the wavefunction. The rapid (high frequency) temporal fluctuations of the electron wavefunction radiates as a dipole, causing the high energy radiation associated with high-harmonic generation. A visualization of the spatially varying

wavefunction for three different driving laser wavelengths is shown in Figure 5.2. This model has significant implications for the efficiency of long wavelength driven HHG. When driven with longer wavelengths, the free electron wavepacket spends more time in the continuum and experiences a greater degree of diffusion. This quantum diffusion reduces the free electron wavefunction amplitude that overlaps with the bound state, reducing the strength of the quantum interference which gives rise to HHG by a factor which scales as λ^{-3} . Furthermore, longer wavelength drivers increase the cut-off energy, effectively decreasing the generated harmonic power in a given photon energy interval by a factor of λ^{-2} . All told, a variety of experimental and theoretical methods have suggested that the combined scaling of single-atom HHG efficiency with wavelength is in the range of λ^{-5} to λ^{-9} [32, 40, 37, 41]. We will soon see in the macroscopic picture that phase matching with longer wavelength drivers requires a lower ionization fraction that scales as $\eta_{crit} \propto \lambda^{-2}$, which would further decrease the number of available emitters in long wavelength driven HHG if it were not for the observation that gas pressure for phase matching scales as $p \propto \lambda^2$, so that the line density of ionized atoms (and therefore potential emitters) stays roughly constant with wavelength. Further, we can recoup some of the efficiency losses caused by longer wavelength driving lasers by noting that helium exhibits dramatically lower absorption of higher energy harmonics - further increasing the accessible pressure-length product, and therefore the total number of emitters that can be harnessed, before the maximum brightness plateaus.

5.2.2 The macroscopic picture of HHG: phase matching and efficiency

5.2.2.1 Phase matching

In order to generate a bright, well-collimated HHG beam, the emission of individual emitters must be phase matched over a macroscopic volume. Just as we saw in our study of second harmonic generation in a crystal, the generated light must travel at nearly the same phase velocity as the driving light in the interaction medium to achieve the phase matching condition. If the harmonics remain in phase over the interaction volume, their fields will add constructively and grow over

the length of the medium, as illustrated in Figure 5.3. At the single-atom level, HHG emission follows a dipole radiation pattern, but only the emitted light which satisfies the phase-matching conditions with the driving laser will grow to significant brightness. Therefore, phase-matching is responsible for the high-degree of spatial coherence generally present in an HHG beam. The macroscopic efficiency of HHG is further determined by reabsorption of the emitted harmonics by the gas. This feature depends strongly on both the photon energy of the emitted harmonics and the species of gas, and is illustrated in Figures 5.7 and 5.8.




Figure 5.3: **The generation of a bright, spatially coherent HHG beam requires the emitted harmonics and the driving laser to be well phase matched.** In this illustration HHG in a gas-filled waveguide, the emitted harmonics (shown in purple) are able to add constructively. Figure from Reference [37].

For the generation of the q -th harmonic to be perfectly phase matched, the wavevector of that harmonic, \mathbf{k}_q , must be equal to the sum of the wavevectors of the contributing laser photons, which each have a wavevector k_L . In general, these are vector quantities, and for tight focusing geometries, it is important to consider the vectorial nature. Nevertheless, for the purpose of this analysis, we restrict ourselves to the 1D case, where all beams propagate in the z -direction. In this case, the wavevector mismatch can be written simply as:

$$\Delta k = k_q - qk_L \quad (5.2)$$

$$= \frac{\omega_q n_q}{c} - \frac{q\omega_L n_L}{c}, \quad (5.3)$$




Figure 5.4: **An illustration of phase matching in HHG.** The phase mismatch between the q -th harmonic and the driving laser is calculated by determining the effective wavevector in the medium for the harmonic (k_q), and subtracting from that q times the effective wavevector of the driving laser ($\Delta k = k_q - qk_L$). The dramatically different responses to soft X-ray and mid-IR light mean that contributions to k_L and k_q by material and waveguide dispersion can be quite different. It's also important to remember that the units of Δk are [radians/meter] - to determine the accrued phase mismatch, one must calculate the product ΔkL .

where in the second line we have expanded the wavevector according to $k = \omega n/c$, where ω is the angular frequency and n is the refractive index. This wavevector mismatch is illustrated in Figure 5.4. We note that the units of Δk are [radians/meter], but the relevant parameter to describe the efficiency of HHG is the actual phase slip between the harmonic waves emitted throughout the medium. When emitted harmonics become π radians out of phase with harmonics emitted earlier, they will destructively interfere and decrease the total emitted brightness. We therefore introduce the *coherence length*, which is the length of the medium over which the phase slip reaches π radians. This can be easily calculated according to:

$$\Delta k L_{coh} = \pi, \quad (5.4)$$

and will be an important parameter in determining the maximum harmonic yield in the presence of reabsorption [38].

From Equation 5.3, it is apparent that the case of perfect phase matching ($\Delta k = 0$) occurs when the effective refractive indices experienced by the harmonic and the driving laser are equal. Earlier, in our discussion of phase matching in optical parametric amplification, we described two common techniques to phase match different frequencies in media: birefringent phase matching and quasi-phase matching (QPM). In QPM, harmonic emission in the spatial regions which would otherwise destructively interfere is reduced through modulation in the gas pressure or through

the introduction of additional counterpropagating light, in principle allowing harmonic emission beyond the phase matching cutoff energy [198, 199]. Unlike in crystal phase matching, the phase mismatch in HHG has many contributions, some of which are oppositely signed and can be adjusted with experimentally accessible parameters such as gas pressure and laser intensity. This allows for phase matching to occur even without birefringence or QPM techniques, but it requires careful accounting of all contributions to the wavevectors of the driving laser and the emitted harmonics. The following overview of this wavevector accounting follows References [36] and [200].

We first summarize the various terms that we will consider in our expression for the phase mismatch:

$$\Delta k = \Delta k_{neutrals} + \Delta k_{plasma} + \Delta k_{geometric} + \Delta k_{dipole} \quad (5.5)$$

where $\Delta k_{neutrals}$ and Δk_{plasma} are the phase mismatches caused by the difference in refractive index for the laser and the harmonics due to the non-ionized, neutral atoms and the plasma, respectively. In general, $\Delta k_{neutrals} < 0$, while $\Delta k_{plasma} > 0$. The remaining two terms depend on the particular geometry chosen for HHG. $\Delta k_{geometric}$ is in general positive. For waveguide harmonics, it stems from the dispersion of the waveguide experienced by the driving laser, while for harmonics generated in a free focus, it is primarily due to the Gouy phase. Lastly, Δk_{dipole} is caused by the intrinsic or atomic dipole phase introduced with the microscopic picture of HHG. We will now quantitatively discuss these contributions to the phase mismatch, with an eye towards how each term can be balanced to generate bright, phase matched harmonics.

The phase velocity of the driving laser is generally slowed in the presence of a neutral gas, as the refractive indices of gases in the visible through mid-IR is typically greater than one. This causes k_L to be larger in a neutral gas than in vacuum. On the other hand, for EUV and soft X-ray frequencies, the refractive index can generally be approximated to be one. Therefore, using the convention $\Delta k = k_q - qk_L$, the $\Delta k_{neutrals}$ contribution is negative. In order to quantify this mismatch, we first consider a pressure dependent refractive index of the driving laser, given by

$$n_L(P) = P(n_L^{(0)} + n_{L,2}^{(0)}I_L), \quad (5.6)$$

where P is the pressure and $n_L^{(0)} + n_{L,2}^{(0)}I_L$ is the refractive index of the gas at $P = 1$ atm, including the intensity-dependent nonlinear index due to the $\chi^{(3)}$ nonlinearity. We also introduce η , which is the fraction of the gas atoms that are ionized. Making the approximation that $n_q = 1$, we can summarize the phase mismatch due to the neutral gas:

$$\Delta k_{neutrals} = -q(1 - \eta)P \frac{2\pi}{\lambda_L} (n_L^{(0)} + n_{L,2}^{(0)}I_L - 1). \quad (5.7)$$

In HHG, some fraction of the gas is ionized to create a plasma, which has a frequency dependent refractive index that is less than one. Balancing the effect of the neutral atoms, the plasma generally introduces a positive phase mismatch. The critical parameter defining the interaction between the beams and the plasma is the plasma resonance frequency,

$$\omega_p = \sqrt{\frac{q_e^2 n_e}{\epsilon_0 m_e}}, \quad (5.8)$$

where $n_e = \eta N_a P$ is the free electron number density and N_a is the number density of atoms in a gas at standard pressure and temperature. From the plasma frequency ω_p , we can deduce the plasma refractive index for radiation at frequency ω :

$$n_{plasma}(\omega) = \sqrt{1 - \left(\frac{\omega_p}{\omega}\right)^2}. \quad (5.9)$$

We note that the plasma refractive index is less than one and increases towards unity as the optical frequency ω is increased. The effect of the plasma is to speed up the phase velocity of the driving laser relative to the phase velocity of the harmonics, such that the phase mismatch $\Delta k = k_q - qk_L$ is positive. The result of some fraction η of the gas being ionized into a plasma is a phase mismatch given by:

$$\Delta k_{plasma} = \eta P N_a r_e \frac{q^2 - 1}{q} \lambda_L \quad (5.10)$$

where $r_e \approx 2.82 \times 10^{-15}$ m is the classical electron radius.

The combined effect of $\Delta k_{neutrals}$ and Δk_{plasma} determines the phase mismatch due to the partially ionized medium. By increasing the ionization fraction η , we can increase the contribution of the plasma term to make the phase mismatch more negative; alternatively, by decreasing η we

can make it more positive. The ionization fraction that causes the neutral and plasma dispersion to exactly cancel out is the *critical ionization*:

$$\eta_c = \left[\frac{\lambda_L^2 N_a r_e}{2\pi(n_L^{(0)} + n_{L,2}^{(0)} I_L - 1)} + 1 \right]^{-1} \quad (5.11)$$

Above the critical ionization ($\eta > \eta_c$), the phase mismatch from the medium is positive ($\Delta k_{medium} > 0$) with a magnitude proportional to the gas pressure P . Below the critical ionization ($\eta < \eta_c$), the phase mismatch is negative ($\Delta k_{medium} < 0$), also with a magnitude proportional to the gas pressure P . We will soon see that the geometric contributions to the wave-vector mismatch from waveguide dispersion or the Gouy phase are generally positive. As a consequence, when the medium is overdriven and the ionization is above critical, tuning the pressure cannot balance the geometric dispersion, and there is generally no way to reach perfect phase matching ($\Delta k = 0$). On the other hand, if the ionization fraction is slightly below critical, the gas pressure be tuned such that the phase mismatch of the medium exactly cancels the geometric contributions, allowing for ideal phase matching. The critical ionization places an upper limit on the laser intensity that can be used to drive phase matched HHG, limiting the maximum photon energies that can be generated. Experiments have determined that the phase-matched cut-off energy scales as $E_{PMcut-off} \propto \lambda^{1.6-1.7}$ [37].

Gas species	Refractive index at $\lambda_L = 3 \mu\text{m}$	Critical ionization η_c at $\lambda_L = 3 \mu\text{m}$
Helium	1.000034682	0.0343 %
Argon	1.00026312	0.259 %

Table 5.1: **Calculated refractive indices and critical ionization levels at a driving wavelength of 3 μm .** Refractive index data for argon is from [201]; helium data is from [202]. Data was measured from 0.48 to 2.06 μm and extrapolated using Sellmeier equations. Critical ionization was calculated using Equation 5.11, and the nonlinear refractive index was ignored.

The refractive indices of argon [201] and helium [202] have been experimentally determined up to a wavelength of about 2 μm . We can use the Sellmeier equations (dispersion formulas) to extrapolate refractive indices to a wavelength of 3 μm and calculate the critical ionization (Table

5.1). Ammosov, Delone and Krainov (ADK) derived an expression for the rate of tunnel ionization in a laser field [187], which is commonly used to determine the laser intensity needed to reach a particular ionization fraction. A clear tutorial on this calculation can be found in [158]. The critical ionization and the laser intensity needed to reach it is graphically illustrated in Figure 5.5 for the wavelength range of 0.8 to 10 μm . When harmonics are driven in a waveguide, the effect of the plasma dispersion on the phase velocity of the driving laser can be reduced due to averaging of the mode over the waveguide area [203]. This means that in a waveguide, it is possible to use marginally higher laser intensities, and achieve a higher ionization fraction, while still maintaining phase matching. Compared to unguided geometries, the higher ionization fraction due to modal averaging can result in brighter emission due to more emitters and higher phase matching cut-off energies (Figure 5.5, dashed lines).

We conclude our discussion of the macroscopic phase matching effects by considering the geometric contributions to the phase-mismatch. The phase velocity of the driving laser can be increased relative to the harmonic beam ($\Delta k > 0$) due to contributions from the Gouy phase, which is most prominent in a free-focused geometry, or waveguide dispersion. The atomic dipole phase, on the other hand, can either increase or decrease the relative phase of the harmonics depending on the intensity gradient.

We first consider the case of harmonics generated in a free focus, which are influenced by both the Gouy phase and the atomic dipole phase. The Gouy phase refers to the phase that a wavefront accrues as it propagates through a focus (Figure 5.6(b)). A Gaussian beam will acquire an extra phase of π radians going through a focus compared to the same beam comprised entirely of plane waves. The accrued Gouy phase ϕ_{Gouy} as a function of longitudinal position z (where the beam waist is at $z = 0$) is given by:

$$\phi_{Gouy} = -\arctan\frac{z}{z_R}, \quad (5.12)$$

where $z_R = \pi\omega_0^2/\lambda_L$ is the Rayleigh length (the length away from the waist where the beam radius grows by a factor of $\sqrt{2}$). By inspection of Equation 5.12, we can see that the rate at which the

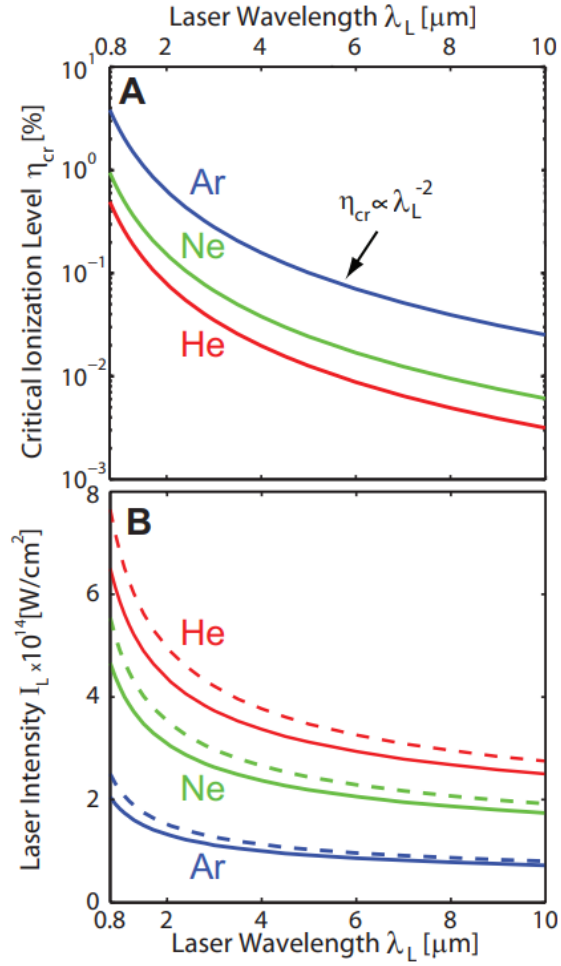


Figure 5.5: (a) Critical ionization levels and (b) requisite laser intensities for driving laser wavelengths between 0.8 and 10 μm . Critical ionization can be easily calculated from Equation 5.11, and scales as $\eta_c \propto \lambda^{-2}$. Laser intensities are peak intensities on axis. The dashed lines represent the slightly higher laser intensities that can be used in a waveguide while still maintaining phase matching - an effect related to modal averaging [203]. Figure from Reference [41].

phase changes near the focus is greater for tighter focuses (smaller Rayleigh lengths). We also see that the phase of the driving laser decreases through the focus relative to the harmonic beam (Equation 5.12 is a monotonically decreasing function with z), speeding up the phase velocity of the driving laser relative to the harmonics. The influence of the Gouy phase depends strongly on where the HHG occurs relative to the waist: the phase mismatch due to the Gouy phase will be greater if the phase matched region is closer to the focus than if it is farther away.

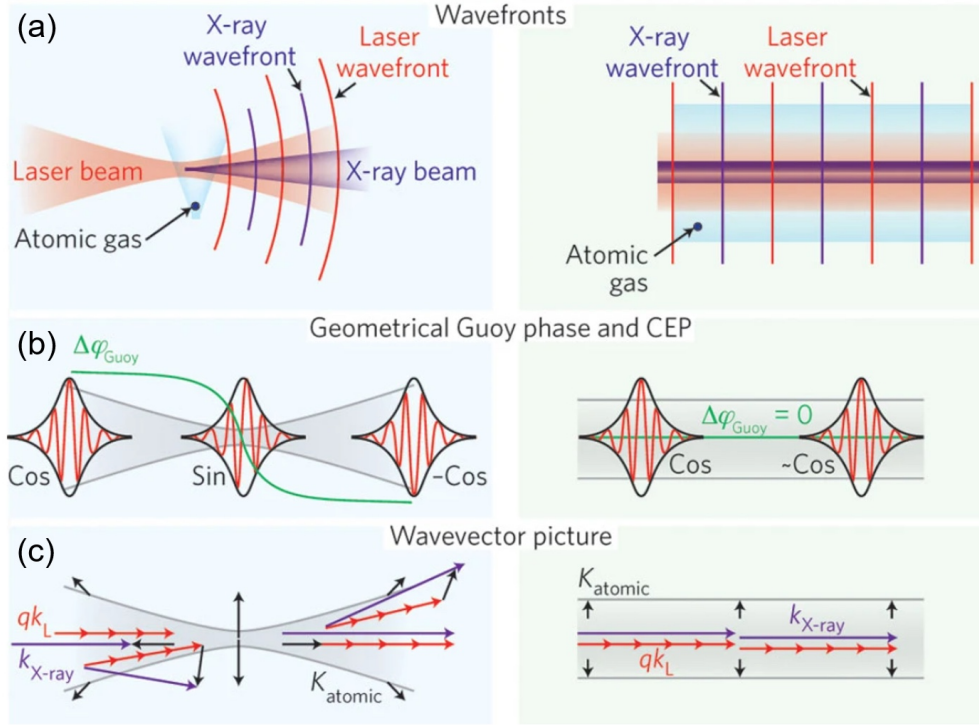


Figure 5.6: **Phase matching harmonics in a free focus and a waveguide.** The Guoy phase results in a phase shift of the driving laser through a free focus. The intensity gradient through the focus results in a spatially varying atomic dipole phase. In the region immediately after the focus, these two effects can balance to enable phase matching over a fraction of the confocal parameter. Meanwhile, in a gas filled waveguide, the interaction length over which phase matching is possible can in principle be longer than the confocal parameter. Figure adapted from Reference [36].

The second contribution to the geometric phase in the free focus geometry is the atomic dipole phase or intrinsic phase. The phase of the driving field at which tunnel ionization occurs depends on the intensity of the field. This ionization phase also determines the amount of time the ionized electron spends in the continuum, and therefore, it determines the phase of the emitted harmonics relative to the driving field. The contribution of the atomic dipole phase to the total phase mismatch is [204]:

$$\Delta k_{dipole} = \frac{-8zb^2}{[b^2 + 4z^2]^2} \alpha_q I_0 \quad (5.13)$$

where α_q is a parameter that depends primarily on travel time during the trajectory, and is smaller for the short trajectories and larger for the long trajectories, $b = 2z_R$ is the confocal parameter, z is the on-axis position relative to the waist, and I_0 is the peak intensity. The odd-ordered dependence

of the atomic dipole phase on z causes it to change both sign and magnitude, preventing ideal phase matching over the full focus. The model from Balcou et al. [205] treats the atomic dipole phase as a separate wavevector in the phase matching calculation. They demonstrate how, in particular in front of the focus, the atomic dipole phase can enable off-axis phase matching and therefore degrade the modal quality of the generated beam (Figure 5.6(c)). The atomic dipole phase is generally negative after the focus, as the laser intensity decreases, which can be used to compensate the positive contribution from the Gouy phase. This optimizes phase matching when the laser is focused just before the region of gas density in a gas jet or a gas cell. Indeed, we observed the brightest HHG beam when the focus was moved just in front of the gas cell.

When intense laser light is guided in a hollow waveguide, the longitudinal gradient in intensity is vastly reduced compared to the case of a free-focus, reducing the influence of the atomic dipole phase on the cumulative phase mismatch. Further, the plane wavefronts that propagate in a waveguide do not exhibit a Gouy phase and maintain high intensity over a larger interaction volume. These factors contribute to the possibility of brighter harmonics from the waveguide geometry. The phase mismatch accounting in a waveguide is considerably simpler: in addition to the terms due to the medium, we need only consider how the waveguide changes the phase velocity of the driving laser. This contribution is generally positive (increasing the phase velocity), and is generally given by:

$$\Delta k_{\text{waveguide}} \approx \frac{qu^2 \lambda_L}{4\pi a^2} \quad (5.14)$$

where q is the harmonic order, u is a parameter related to the waveguide mode, and a is the radius of the hollow section of the waveguide. In the waveguide geometry, ideal phase matching is achievable by pressure tuning as long as the ionization level is below critical ionization.

5.2.2.2 Efficiency limits due to reabsorption

The noble gases which are frequently used to generate harmonics also absorb the emitted harmonic light. The absorption length, or attenuation length, is defined as the length of material through which the transmission is $1/e$ (36.8%). Absorption lengths for helium, neon and argon


are presented in Figure 5.7(a) using data from the Center for X-ray Optics (CXRO) database [9]. The reabsorption of harmonics by the gases that emit them has two significant consequences in the macroscopic picture of HHG: (1) the total harmonic flux from a source will not increase indefinitely with the length of the medium (number of emitters), even for the case of perfect phase matching [38]; and (2), one must take care in designing an HHG source to pump out the gas as quickly as possible after the interaction region to prevent reabsorption. This second consideration informed the design of the HHG source presented in the next section.

In the case of perfect phase matching in a perfectly transparent medium, the harmonic intensity would grow quadratically with the number of emitters - for the case of uniform pressure, the quadratic growth would be over the length of the medium. This is simple enough to understand: if all the emitted waves are travelling perfectly in phase, the generated fields are superposed linearly, and the resultant sum is squared to deduce the intensity. This result is shown in the dotted line in Figure 5.8. In reality, the harmonic yield does not grow indefinitely over the length of the medium, but is limited by reabsorption. Constant et al. developed a model to predict the evolution of the harmonic flux in an absorbing medium, and found that even in the case of perfect phase matching, the generated harmonic flux saturates after around 10 absorption lengths (L_{abs}), and reaches around 80% of the maximum yield after about 5 absorption lengths [38]. When the coherence length, defined by $L_{coh} = \pi/\Delta k$, is reduced due to sub-optimal phase matching, the maximum harmonic yield decreases and is reached after just a few absorption lengths (Figure 5.8).


The behavior in Figure 5.8 is nicely captured by the following formula (from Reference [206]) for harmonic intensity as a function of depth in the medium L , phase mismatch Δk , and absorption length:

$$|E_q|^2 \propto e^{\frac{-L}{2L_{abs}}} \frac{\cosh \frac{L}{2L_{abs}} - \cos \Delta k L}{\Delta k^2 + \frac{1}{4L_{abs}^2}}. \quad (5.15)$$

This result has critical implications for high-energy harmonics driven by mid-IR lasers. In general, the efficiency of harmonic generation at the single atom level scales poorly with wavelength (between λ^{-5} to λ^{-9} [32, 40, 37, 41]). Beyond the single atom picture, the critical ionization, which determines



(a) The attenuation length is defined as the length of gas which reduces the intensity to $1/e$ of the incident value. The densities used are listed in the legend and come from the CXRO database. The values are for 1 atm of pressure; at higher pressures (generally required for phase matching), the attenuation length may be shorter.



(b) Even short lengths of residual gas can dramatically reduce the brightness of an HHG source due to reabsorption. Argon, for instance, is all but opaque at 250 eV.

Figure 5.7: **The soft X-ray attenuation lengths and transmission of gases relevant for HHG.** Due to vastly superior transparency and high ionization potential, helium is the preferred gas for soft X-ray HHG. Data was compiled from the Center for X-ray Optics database [9].




Figure 5.8: **The influence of reabsorption by gas on the optimum HHG flux for various phase matching conditions.** Even in the case of perfect phase matching ($L_{coh} \gg L_{abs}$), the generated HHG flux stops increasing with interaction length after around 10 absorption lengths. As phase matching degrades, the coherence length decreases, while the absorption length remains a fixed property of the gas. Figure from Reference [38].

the fraction of gas atoms which can actually emit light, scales as $\eta_c \propto \lambda_L^{-2}$. Nevertheless, some of these losses can in principle be regained for phase matched harmonics due to the higher transparency of noble gases to high energy soft X-ray photons. For example, the attenuation length in helium of 500 eV photons is two orders of magnitude larger than that of 100 eV photons (see Figure 5.7). The result is that, in the case of perfect phase matching, the harmonic yield will saturate after the driving laser has interacted with (not necessarily ionized) two orders of magnitude more gas atoms. For the case of 1 keV photons, the improvement is three orders of magnitude. Consider 3 μm driven HHG in helium, where an optimal phase matching pressure is predicted to be around 13 atmospheres [37]. The absorption length for atmospheric pressure helium at 750 eV is 38 cm - at 13 atmospheres of pressure, this length is reduced to 2.9 cm. A well phase matched source for these harmonics may require around five absorption lengths for optimal brightness, meaning the interaction length may be up to around 15 cm in length. In order to sustain the high peak intensity needed for HHG over these lengths, low loss waveguides are required. One approach to overcome

the generally higher waveguide losses in the mid-infrared is the use of anti-resonant hollow-core fibers.

5.3 A versatile, modular HHG source

In order to use our laser to drive high-harmonic generation in gas, we designed and fabricated a novel source. In this section, we will describe important factors that influenced the design, how the source actually works, and lastly present some data on gas and vacuum handling.

5.3.1 Introduction: context and design criteria

Earlier in this Chapter, we described the physics of phase matching in HHG. One result was that in order to increase the brightness of HHG - especially in the mid-IR - it is necessary to phase match over a length longer than the typical confocal parameters associated with the intensities needed to drive HHG, motivating the use of gas filled waveguides. The only other reported HHG result driven by mid-IR lasers utilized a 200 μm diameter, 5 cm long, unstructured gas-filled hollow waveguide [3], driven by 80 fs, 10 mJ pulses at 3.9 μm wavelength. The peak power accessible by that OPCPA was more than an order of magnitude higher than what is accessible by the current build of the 3 μm OPCPA presented here, and higher than what has been used to drive HHG in the 2-3 μm window. The lower peak powers used by many long-wavelength HHG experiments makes the use of gas-filled fibers challenging due to the scaling of the attenuation, which, in decibels, goes as λ/a , where a is the waveguide diameter. Reducing the waveguide diameter to account for lower peak power therefore increases hollow-core fiber transmission losses, which are already more severe in the mid-infrared due to detrimental wavelength scaling. In order to overcome this scaling, we continue to explore the use of anti-resonant hollow core fibers.

In the Introduction, we reviewed four important laser sources which have driven HHG with driving wavelengths above 2 μm . Here, we briefly mention their chosen HHG geometries as context for our own design. The 2.6 mJ, 40 fs source from MIT and the 2.7 mJ, 30 fs source from the MBI Berlin both utilized a 6 mm interaction length, differentially pumped gas cell to drive harmonics

in argon and (for the MIT system) nitrogen, with the MBI noting that higher flux could likely be achieved with improved gas target design [82, 83]. The 250 μJ , 16.5 fs ETH Zurich source utilized a pressurized stainless-steel needle with a 1 mm inner diameter and 100 μm thick walls, which were drilled by the laser itself, effectively forming a 1.2 mm long gas cell. The shorter interaction length is well justified by the tighter focus required to reach sufficient intensity. They also considered a gas jet but found it unsuitable due to the high pressures required for long-wavelength drivers [84]. The 2.5 μm source at The Ohio State University is the longest wavelength of four sources considered in this section and delivers 7 mJ, 100 fs pulses. They utilized a gas jet with a 0.5 mm inner diameter connected to a pulsed piezo-valve, which was sufficient to generate 0.6 keV HHG when backed by 2.5 bar of argon - however, the authors note that for future studies using higher pressure helium, they intend to switch to a gas cell [86]. In conclusion, three distinct HHG geometries have been used at wavelengths longer than 2 μm : waveguides, gas cells and gas jets. The physics of phase matching suggests that waveguides should be the brightest source, but they have only been demonstrated once before in this wavelength regime, with significantly more peak power than we have access to in the current build of the system. Further, given our lower peak power, we require sophisticated dispersion engineered fibers. Gas cells are the most widely utilized in the field, and there is at least one demonstration using a gas jet - the latter being the simplest and easiest to implement. It is clear that a significant amount of source optimization work has yet to be done in this field.

In consideration of the macroscopic phase matching and reabsorption effects presented earlier, and the literature just reviewed, we can pinpoint a few key design considerations in building our HHG source:

- (1) The high pressures required for phase matching in the mid-IR results in significantly more gas flow into the vacuum system. This gas must be removed rapidly to prevent reabsorption of the harmonic flux using differential pumping.
- (2) The optimization of the HHG source will require significant experimentation, in particular with regard to type of gas confinement (waveguide, gas cell, or gas jet), interaction lengths,

aperture sizes and pressure profiles.

- (3) Experience has proven that focusing high intensity lasers on solid targets frequently results in damage to the target; the source must return quickly to service after damage.
- (4) Having unfettered access (without clipping) to the driving laser in the vicinity of, and downstream of, the HHG target is necessary for the purpose of alignment and coupling diagnostics.

5.3.2 Design and fabrication of the source




Figure 5.9: **The full HHG source assembly is highly modular, enabling rapid optimization as well as reconfiguration for different driving laser parameters and gas types.** Cross sectional renderings are shown along a longitudinal plane of symmetry. The rendering at the bottom shows the source configured for a gas cell driven by a high NA beam, which is the configuration that generated the HHG spectra shown in the following section. The beam propagates from left to right. The three renderings at the top illustrate examples of targets from the three categories that were tested: gas cells, gas filled capillaries, and gas jets. The target can be easily swapped out without change to alignment or the differential pumping mechanism. Credit to Jeremy Thurston for his assistance in artistic renderings.

In order to satisfy the design criteria, the HHG source is highly modular (Figure 5.9). A rendering depicting the propagation of the driving laser, generation of a harmonic beam and gas flow is shown in Figure 5.11. The source is comprised of six distinct modular components: the differential pumping chassis, the target, the two differential pumping inserts, the laser entrance port insert and the lid. The differential pumping chassis is the main body (100 mm long x 60 mm wide x 30 mm tall), onto which the remaining four components attach. The chassis is the most complex and difficult to fabricate component, and is designed to be the only fixed component of the system - working with any other possible configuration, and free of any optically active components which may be damaged by the laser. The chassis is the main body onto which the remaining four components are attached. It is comprised of two vacuum chambers - the larger of which is the main chamber, and the smaller is the differential pumping chamber. During operation, gas spews from either side of the target into the main chamber. Both sides of the main chamber are pumped directly by ports on the floor (connected to the same high speed scroll pump). The differential chamber is also pumped out through the floor, and is connected to a separate roughing pump. High pressure gas is delivered to the target via a port in the bottom of the main chamber. The bottom of the main chamber contains two parallel rows of precisely positioned 2 mm holes used for target registration. The targets have identical holes in their mating surface. 10 mm long, 2 mm diameter precision stainless steel dowel pins are used to register the target in the plane of the chamber floor to within $\sim\pm 10\ \mu\text{m}$ tolerances. Parallel to each row of dowel pin holes are a row of threaded holes, used to bolt the target to the chamber floor. In this way, targets of varying lengths up to about 30 mm can be implemented, and the target can easily be exchanged without loss of alignment - indeed, we have retained good coupling through a hollow core fiber before and after removing the target. The use of a fixed target position is well suited to the good pointing stability of the OPCPA.

The chassis is secured to a pedestal (from KMLabs), depicted in Figure 5.10, which is in turn bolted to the optical table. In this way, the chassis is highly robust to misalignment or mechanical instabilities, despite being attached to two roughing pumps. The pedestal sets the

height of the driving laser at 4.5 in above the optical table. The pedestal contains internal passages for the delivery of gas and connection to the two vacuum pumps. The pedestal also incorporates a pressure relief valve in case of accidental pressurization of the vacuum system and water cooling, although the low average powers in this experiment did not necessitate water cooling. The chassis is sealed to the pedestal with four separate O-rings.

The main and differential chambers are separated by a low vacuum conductance insert. The standard insert has a 0.75 mm diameter bore and is ~ 12 mm long. The differential chamber and the soft X-ray beamline are separated by another low conductance insert which has a 1.5 mm diameter bore and is 38.45 mm long. The driving laser enters the main chamber through the laser entrance port insert, which is identical to the exit side insert except for its bore. In our experiments, we epoxied a 2 mm thick, broadband AR (2-5 μm) AR coated sapphire window to the laser entrance insert. A CaF_2 window is also used. For the latter, the calculated B-integral is only 0.07 radians.²



Figure 5.10: **The full HHG source assembly on the pedestal, as it appears on the optical table.** The large ports on the bottom of the pedestal connect to two roughing-level vacuum pumps. The lid is not pictured. Position registration of the chassis is ensured by a lip on the near side of the pedestal. In the configuration which generated harmonics, a window was epoxied to the brass laser entrance port insert; the backside insert is connected to the beamline with an ultraTorr fitting.

The outer diameter of the externally facing inserts is 0.5 in, allowing for easy coupling to the beamline using standard ultraTorr connections. In the case of shorter wavelength or higher peak power drivers, it can be useful to connect the source directly to a vacuum chamber containing beam delivery optics to avoid issues associated with B-integral; for this reason, the entrance side insert can also easily be connected to a vacuum fitting. All three of the inserts are fitted with O-rings at the neck of the bolt to improve vacuum sealing. Finger tightening is sufficient for the externally facing inserts, while the smaller, internal insert requires a hexhead driver for removal. All three inserts can quickly be exchanged without influence to system alignment, which is useful in a variety of cases. In particular, during alignment of the beam through a target, it is necessary to quantify the transmission efficiency and measure the transmitted mode. This is not possible with the differential pumping inserts in place, due to the restrictive apertures and the rapidly diverging, high NA beam. Therefore, we align the assembly with the differential pumping inserts removed, and replace them for high-harmonic generation. Swapping the inserts is also useful for vacuum diagnostics or optimization of the differential pumping.

A single lid with seals both chambers with two O-rings. In normal operation, we use a clear acrylic lid to give visual access to the often colorful process (see, for instance, Figures 5.12 and 5.23). In diagnostic operation, we use a modified aluminum lid fitted with a vacuum flange that allows measurement of the pressure in the main chamber.

With all differential pumping considerations handled by other components, the targets are designed to be effectively consumable, allowing for rapid prototyping and replacement of laser damaged units. An example of each type of target tested is rendered in Figure 5.9. All targets, regardless of type, connect to the chassis with the same arrangement of dowel pins and bolts. The bottom of each target contains an O-ring to isolate the high pressure gas from the main chamber. 15 atm of gas was applied to a blanked-off target to test the effectiveness of the isolation between the high pressure gas and the vacuum system; no measurable change of vacuum levels was detected. The gas cell used for successful HHG is illustrated in detail in Figure 5.16, and involves a central

² Assuming a window thickness of 3 mm, and a 1 mJ, 120 fs pulse in a 2.4 mm radius beam.




Figure 5.11: **A cross-sectional rendering of the modular HHG source driven by a high NA beam with a capillary cartridge.** Red beam: mid-IR driving laser focused with a 75 mm lens to a 35 μm radius spot size; pink beam: high harmonic beam; teal: applied gas (typically argon, neon, nitrogen or helium). Beams are not drawn to scale.

high pressure chamber connected to the main vacuum chamber via two low vacuum conductance drilled holes. The driving laser and harmonic beams pass through these holes. The gas jet utilized a 1.2 mm outer diameter, 150 μm inner diameter quartz capillary angled vertically and epoxied to a flange to deliver gas to the main chamber. Finally, the capillary targets operate by holding a hollow core fiber colinear to the laser propagation on both ends using UV curable epoxy. Gas inlet holes are drilled on the side of the capillary using a laser, and gas exits the fiber through either end facet into the main chamber. Figure 5.11 depicts the operating HHG source configured for a capillary.

The chassis and the variety of targets were all machined from 6061 aluminum using a 3-axis Haas CNC mill at the JILA Instrument Shop. The three inserts were machined from brass precision shoulder bolts using a Hardinge lathe at the JILA Staff Shop. Shoulder bolts were chosen as starting stock due to their ready availability and tight tolerances on the outer surface of the cylindrical section, enabling high concentricity and repeatable registration with the chassis. A photograph showing three completed chassis and a variety of targets and inserts is presented in Figure 5.12. For more detailed specifications of the HHG source, including technical drawings and user notes, see Appendix D.




Figure 5.12: A photograph showing a selection of chassis, inserts and targets, taken shortly after the units were machined.

5.3.3 Gas handling and differential pumping characterization

The multi-atmosphere pressures required for phase matching in the mid-IR result in significant gas flow into the main chamber from all three interaction geometries considered here. We first characterized the amount of gas that flowed from the target into the main chamber by isolating the main chamber with two solid (blank) inserts and connecting a hose to the lid. Gas was applied to both a gas cell and a 12 mm long, 100 μm inner diameter quartz capillary with a single hole drilled in the middle with a CO_2 laser (Figure 5.14(a)). The measured flow rates are presented in Figure 5.14(b). We note that the measured gas cell consumes gas at a $\sim 5.5\times$ higher rate than the capillary. Beyond being a useful data point in developing gas flow models, gas consumption is useful to understand in the context of the high cost of neon gas and non-renewability of helium gas. A further experiment was done demonstrating that the exhaust of the main chamber roughing pump is comprised almost entirely of the feed gas, enabling a route towards recapture of precious gases.




Figure 5.13: A photograph of the HHG source configured with an argon filled gas cell being driven by the 3 μm driver, indicating jet-like behavior of the argon gas as well as density modulations thought to be shock diamonds. The glowing stream is thought to be fluorescing argon, and is angled slightly down with respect to the driving beam (towards the vacuum port on the floor of the chamber), suggesting that fluorescence excitation happens near the focus, with the glowing gas travelling towards the port.

Figure 5.15 presents pressure measurements in the main chamber, differential chamber and the beamline for a range of applied pressures of nitrogen and helium. Details of this measurement, as well as the specific vacuum pumps used in the beamline, are given in the caption. The differential pumping scheme was found to be sufficient for even high backing pressures of helium, which is considered the most difficult gas to pump: at 15 bar of applied gas, the beamline pressure stayed below 2×10^{-5} Torr. During HHG in argon, we applied a backing pressure of ~ 2.2 bar to a gas cell. In the case of the capillary backed by ~ 2.2 bar of nitrogen, the main chamber pressure was ~ 0.6 Torr. According to the CXRO database [9], an 8 mm path through 0.6 Torr of argon results in transmission of greater than 92% for the entire soft X-ray region, including at the absorption peak near 250 eV. Vacuum measurements were also made with the gas cell in place using argon gas. With applied argon pressures in the typical range used for HHG, the main chamber pressure was ~ 0.85 Torr and the differential chamber pressure was ~ 60 mTorr. At these pressures, argon is more than 99% transmissive at 240 eV. While these pressure readings would suggest that reabsorption is

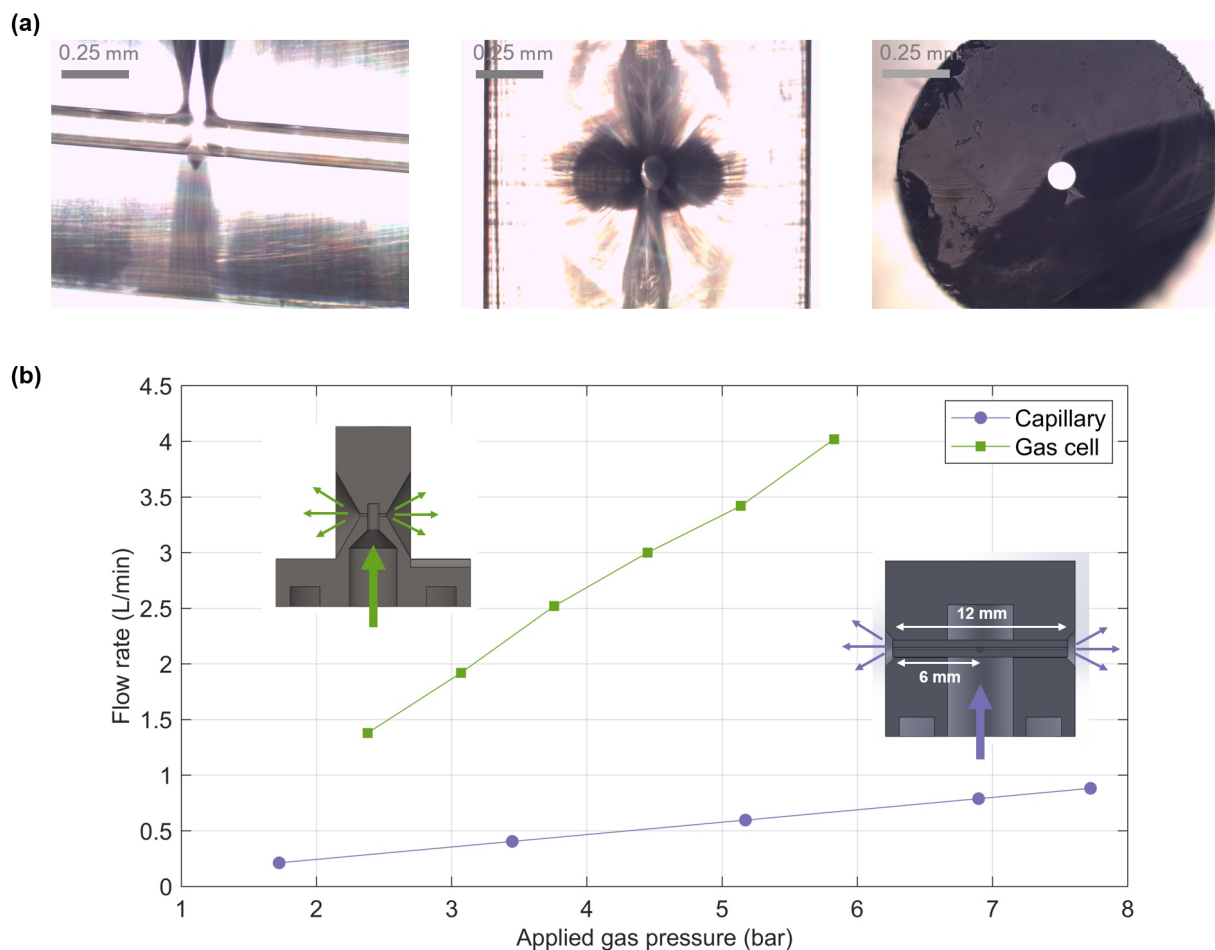


Figure 5.14: (a) Microscope images of a CO_2 laser drilled capillary and (b) measured gas flow rates through a laser drilled hollow core capillary and a gas cell. A hole was drilled in the side halfway along the length of a 12 mm long, 1.2 mm outer diameter, 100 μm inner diameter quartz capillary using a CO_2 laser. The capillary was epoxied into the fiber holder with UV curable epoxy. The gas inlet port was pressurized with nitrogen gas, and the flow rate out the ends of the capillary was measured. The measurement conducted by measuring the amount of time the discharged gas took to displace one liter of water in a graduated cylinder. A similar measurement was taken using a gas cell (for dimensions, see Figure 5.16) with argon gas. For the capillary, pressures were measured on the bottle regulator; for the gas cell, pressures were measured on the inline pressure gauge.

not problematic in the HHG source, it is likely that they do not necessarily capture the actual gas density seen by the HHG beam. The pressure in the main chamber is thought to be inhomogeneous, with streams of high pressure gas forming higher density regions on either side of the target orifices. Figure 5.13 illustrates this with a photograph showing a stream of glowing argon gas moving from

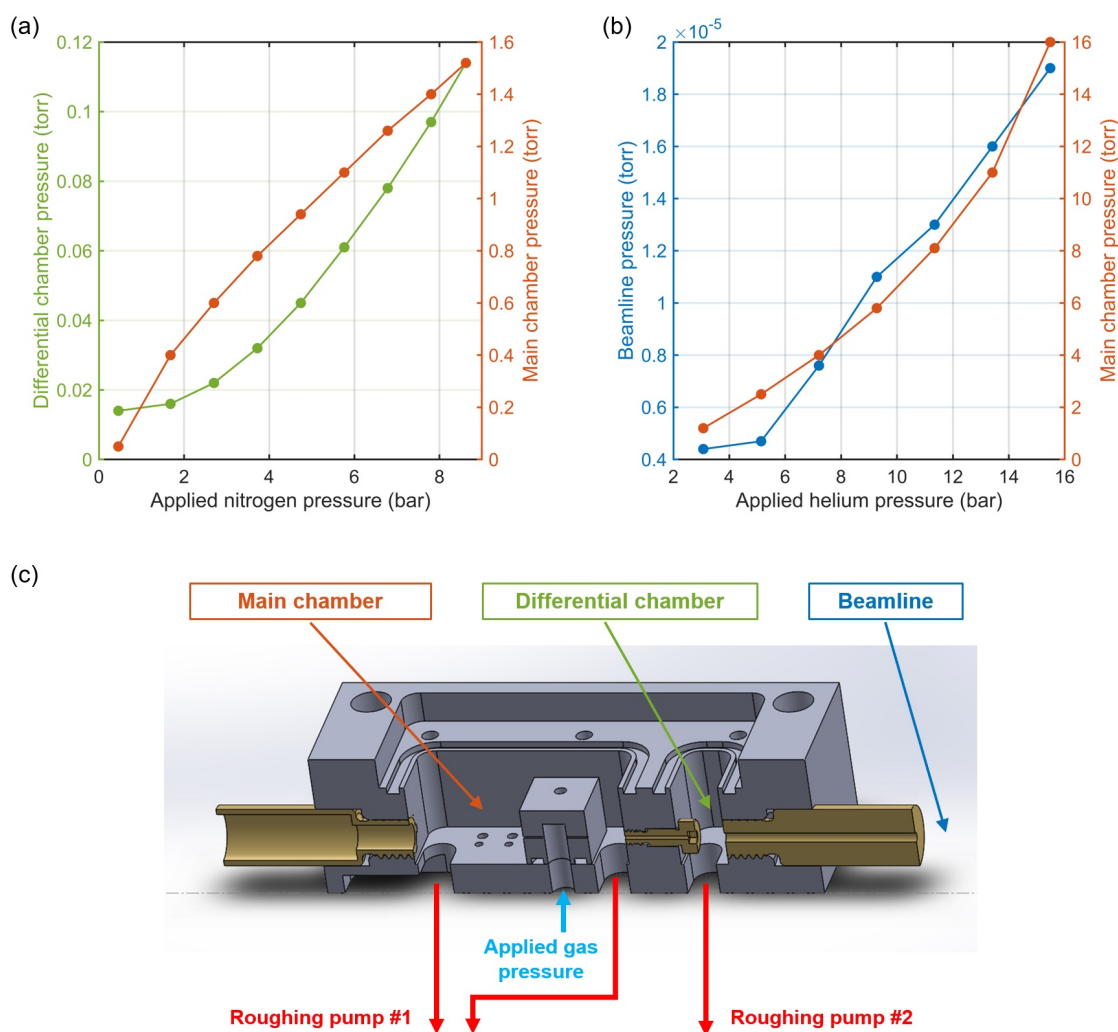


Figure 5.15: **Differential pumping pressure measurements for multiple atmospheres applied to a laser drilled, 100 μm diameter hollow core capillary.** In both measurements, a hole was drilled in the side halfway along the length of a 12 mm long, 1.2 mm outer diameter, 100 μm inner diameter quartz capillary using a CO_2 laser. In both (a) and (b), the input-side insert was sealed to the environment with an epoxied laser window. (a) The effect of the small differential pumping insert with applied nitrogen. The exit-side insert was replaced with a high gas conductance, large bore diagnostic insert, and a vacuum gauge was placed on the output. The 0.75 mm diameter inner differential pumping insert remained. (b) The effect of both differential pumping inserts with applied helium. Even in the case of 15 bar of continuously applied helium, the beamline pressure remained in the low 10^{-5} Torr level. The second insert is 1.5 mm inner diameter. Helium is more difficult for the pumps to handle, causing the $\sim 3\text{x}$ increase in main chamber pressure for same applied pressure of helium compared to nitrogen. (c) In all measurements, the pumps are: Roughing pump #1: “Busch Fossa 0035 A” 35 m^3/hr scroll pump; Roughing pump #2: “Adixen ACP28” 27 m^3/hr roots pump. The beamline is pumped on by a Pfeiffer HiPace 300 turbomolecular pump.




Figure 5.16: **The gas cell used to generate 3 μm driven argon harmonics.** The high pressure region does not neck down until a few mm from the interaction region, suggesting that the backing pressure should be a relatively accurate representation of the pressure in the interaction region.

the gas cell to the chamber floor vacuum port. The glowing gas is thought to be excited argon atoms undergoing fluorescence during their journey from the gas cell to the vacuum port. We note that the fluorescence, which is excited at the gas cell orifice by the driving laser, is localized in a jet. We further note the presence of shock diamonds, indicating supersonic expansion of the gas. Similar shock diamonds were observed in an HHG gas cell driven by a 1.85 μm laser and used to model pressure gradients [207]. Recent work has used the fluorescence signal from an argon gas jet to characterize the density length product of the medium [208].

5.4 HHG results

HHG was observed using the argon backed gas cell backed by an estimated 2.3 bar of pressure. Exact determination of the backing pressure was challenging with the pressure gauge utilized,

which gave measurements in psig and had questionable accuracy below 1 bar of pressure; pressure estimates were made by adding 1 bar of pressure to the measured psig pressure. A dimensioned drawing of the gas cell used is presented in Figure 5.16. Approximately 630 μJ (as measured in front of the focusing lens) was focused into the gas cell using a 75 mm focal length CaF_2 lens. The $1/e^2$ beam diameter at the focus was estimated to be about 70 μm , and the pulse duration was measured using SHG-FROG to be ~ 130 fs FWHM.³ The intensity on gas was estimated to be $\sim 2.5 \times 10^{14}$ W/cm². This estimate was validated using the below-threshold harmonic plateau, as described in the following section.

Careful alignment of the driving beam to the gas cell and the beamline was critical to the observation of the HHG beam. The 7th harmonic (below-threshold harmonic, see Figure 5.23) near 422 nm is colinear with the HHG and is dramatically less divergent and easier to observe than the 3 μm driving beam. The 7th harmonic was successfully used to align the beamline in the absence of an HHG beam. This ensured that the camera was in the correct location during parameter scanning. The height of the 7th harmonic beam increased by only ~ 3 mm over ~ 1 m of propagation after the HHG source. In this configuration, the 3 μm transmission of the gas cell (including the window and lens) was measured to be 73%, but some losses may be accounted for by difficulty positioning the power meter.

5.4.1 Beamline

The soft X-ray beamline is depicted in Figure 5.17. The source is connected to a gate valve with a flexible bellows. When the gate valve is closed, beamline pressure typically stabilizes in the low 10^{-7} Torr range. The gate valve allows the beamline to remain under vacuum, and the X-ray camera to remain cold, during any modifications to the HHG source. The beamline is pumped on by a Pfeiffer HiPace 300 turbomolecular pump with a 260 l/s pumping speed for nitrogen gas. The filter wheel assembly (Artemis from KMLabs) contains four separate filter wheels, each with a capacity of 7 filters. A 2f-2f single lens imaging system uses a 50 cm focal length, 5 deg

³ Later improvements in the OPCPA brought this pulse duration down to around 100 fs




Figure 5.17: **The soft X-ray beamline (drawn partially to scale).** Labelled distances and angles are drawn to scale. Positions of the gate valve, 260 l/s turbopump, iris and filter wheel are approximate. The nickel toroid images ($2f$ - $2f$) the source onto the entrance slit plane of the grating.

glancing incidence nickel coated toroid (ARW Optical Corporation) to image the HHG target to the entrance slit plane of the grating. No physical slit was used for the spectrometer. The argon and nitrogen harmonic spectra in the following section were taken using a Hitachi aberration corrected concave grating (part 001-0437) with 1200 grooves/mm, a 5649 mm radius of curvature, 10 nm blaze wavelength and 3.2 deg blaze angle. The grating is optimized for the spectral range of 5-20 nm (248-62 eV). This grating (as well as the 2400 grooves/mm grating intended for higher photon energies) is well characterized in Reference [209]. The camera used for spectroscopy and beam characterization was a TEC-cooled Andor Newton 940 CCD (part number DO940P-BN) with a 2048 x 512 array of 13.5 x 13.5 μm pixels.

5.4.2 Modes and parameter scans

During beam characterization measurements, the toroidal focusing mirror (Figure 5.17) was removed and replaced with the X-ray camera. The camera sensor was about 89 cm past the HHG target. Beam images taken in this configuration are presented in Figure 5.18. In Figure 5.18(a), we show that the harmonic beam is fully blocked by a 100 μm glass coverslip, which was attached to a filter wheel.⁴ Figure 5.18(b) highlights the sensitivity to the longitudinal position of the 7.5 cm focusing lens: moving the lens 50 μm (representing 4% of the Rayleigh length) in either direction dramatically decreased the brightness of the harmonics.⁵ We also display sensitivity to the applied gas pressure in Figure 5.18(c). The narrow range of pressures resulting in bright HHG is distinctly different than in the case of below threshold harmonics, for which brightness monotonically increased with gas pressure. The narrow range of pressures generating harmonics is likely due to phase matching, but we also note that increased gas pressure in the region in front of the gas cell can result in both Kerr lensing and, depending on intensity, plasma defocusing, which could account for decreases in brightness at higher pressures.

An enlarged image of the beam is presented in Figure 5.18(d). We note that the cross-

⁴ The use of a glass plate is an excellent way to distinguish between HHG and below threshold, low order harmonics, which can penetrate thin metal foils with measurable flux and confuse the observer.

⁵ We note that the Rayleigh length, which is the distance from the waist to the plane where the cross-sectional area is doubled, is estimated to be 1.3 mm for the driving laser.

like shape imprinted in the beam profile is likely due to damage to the gas cell aperture: during alignment, even 10s of mW can ablate the aluminum cell, and this particular cell exhibited cross-like damage from scanning the lens in the X and Y directions to optimize alignment. A cross-sectional lineout of the beam suggests a Lorentzian like shape, and exact characterization of the beam width is difficult. Nonetheless, depending on how the fitting is done, the beam radius is on the order of 1.25 mm. Using the formula for the evolution of the beam width of a Gaussian beam,

$$\omega(z) = \omega_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2}, \quad (5.16)$$

where $\omega(z)$ and ω_0 are the $1/e^2$ radii of the beam at some position z and the waist ($z = 0$), respectively, and the Rayleigh range is given by $z_R = \frac{\pi\omega_0^2}{\lambda}$, we can estimate the size of the beam at the source, 89 cm before the camera. Using a photon energy of 150 eV (8.3 nm), which is approximately the brightest part of the spectrum, the estimated beam radius at source is 4.8 μm , compared to an estimated driving laser radius of 35 μm . While this effective source size is not unexpected, we also note that in a tight focusing geometry, noncolinear phase matching of the type predicted by Balcou et al. (Reference [205]) may also be influencing the divergence of the beam. Finally, we note that the driving laser power was increased from 630 to 690 μJ by increasing the pump power on the final OPCPA crystal, but this did not result in an observable increase in HHG brightness, again likely due to phase matching but also possibly influenced by nonlinearities immediately before the gas cell. HHG optimization is often achieved by irisng the driving laser. An adjustable iris was placed before the focusing lens and reduced slightly, bringing the pulse energy incident on the lens from 630 μJ to 600 μJ . This resulted in a nearly complete reduction in the harmonic flux.

We now provide some context for the observed optimal backing pressure of 2.3 bar. A 1.85 μm OPA (TOPAS from LightConversion) at ICFO was used to drive HHG using a similar gas cell, composed of a 0.5 mm ID, 1.5 mm OD tube with 0.3 mm entrance and exit holes for the laser, with an optimum argon backing pressure of 1.2 bar [207]. The only other system to date that has to date driven HHG in the mid-infrared used a 3.9 μm laser in a capillary. A phase-matched peak

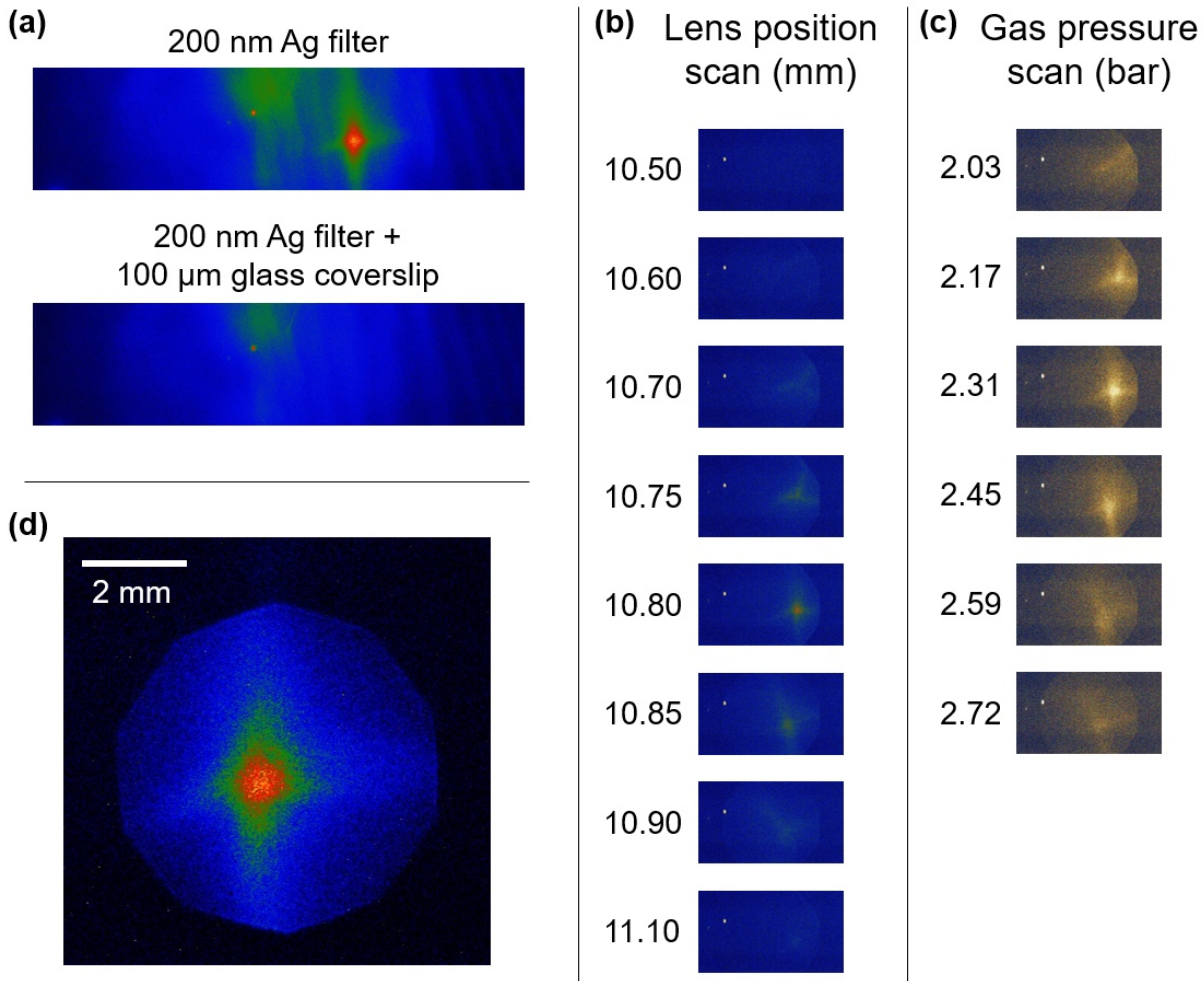


Figure 5.18: **Raw X-ray CCD images of the HHG beam, including pressure and lens position scans.** (a) The HHG beam is blocked by a 100 μm glass coverslip, effectively ensuring that the beam is not comprised of below threshold harmonics. (b) and (c) Scans taken with a 100 nm + 200 nm Ti filter and 8 sec exposure time per image, at (b) 2.3 bar and (c) 10.8 mm lens position. (d) The harmonic beam through a 200 nm Ag filter with an 8 sec exposure time. The windowing is from the in-vacuum iris. Note: the CCD has a region of hot pixels to the upper left of the beam.

was reported for argon at a pressure of 3.5 bar [3].

Although our observed optimal backing pressure fits well within the aforementioned data points, the question of optimal backing pressure is considerably more complicated. Both the MBI and MIT used a 6 mm long gas cell for their 2.1 μm systems. The MBI reported an optimum gas backing pressure of 0.24 bar using a 75 cm (loose) focusing lens [83], while the MIT group

reported an optimum 0.2 bar backing pressure [82]. To complicate things further, the ETH Zurich group reported an optimum backing pressure of 45 bar of helium for a 1.2 mm long gas cell [84] - although argon would be expected to have a considerably lower backing pressure than helium. While it is very difficult to quantify the exact pressure in the interaction region, these seemingly disparate results may be somewhat clarified if one assumes absorption limited HHG, which results in a constant pressure-length product for optimized brightness. Under this condition, there is a hyperbolic relationship between pressure and length. Reference [206] suggests that at lower pressures, phase matching can be reached at a lower ionization level, while absorption pegs the optimum pressure-length product at a fixed value. Those authors concede that there is a dearth of experimental measurements of HHG conversion efficiency as a function of pressure and length to verify these models. In this context, we note that the modular source presented in the previous section may be an appropriate platform to provide rigorous measurements in this parameter space.

5.4.3 Spectra

High harmonic spectra were recorded using the beamline depicted in Figure 5.17. The camera was positioned to capture the $m=-1$ diffraction from the grating. The spectra were calibrated by fitting the observed location of absorption features of boron (at 188 eV) and argon (at 248.4 eV) [9] to a simple geometric model derived from the grating equation. Figure 5.19 displays four measured spectra all taken in the same measurement run, driven by 680 μJ in the argon filled gas cell. Details of the measurement are in the caption. We note that the shape of the measured spectrum is largely determined by the combined transmission of the filters and argon gas. Importantly, the spectral amplitude consistently drops to zero at the argon absorption edge of 248 eV, which is below the predicted phase matching cutoff of around 330 eV [37]. There are two related explanations for this behavior based on reabsorption, and both are likely to be contributing. Firstly, if the harmonic generating region is more than a few absorption lengths away from the backside gas cell aperture, most of the flux will be reabsorbed. We note that the attenuation length for 2.3 bar of argon drops from 524 μm before the edge to 33 μm after the edge [9]. The high pressure region of the gas cell,

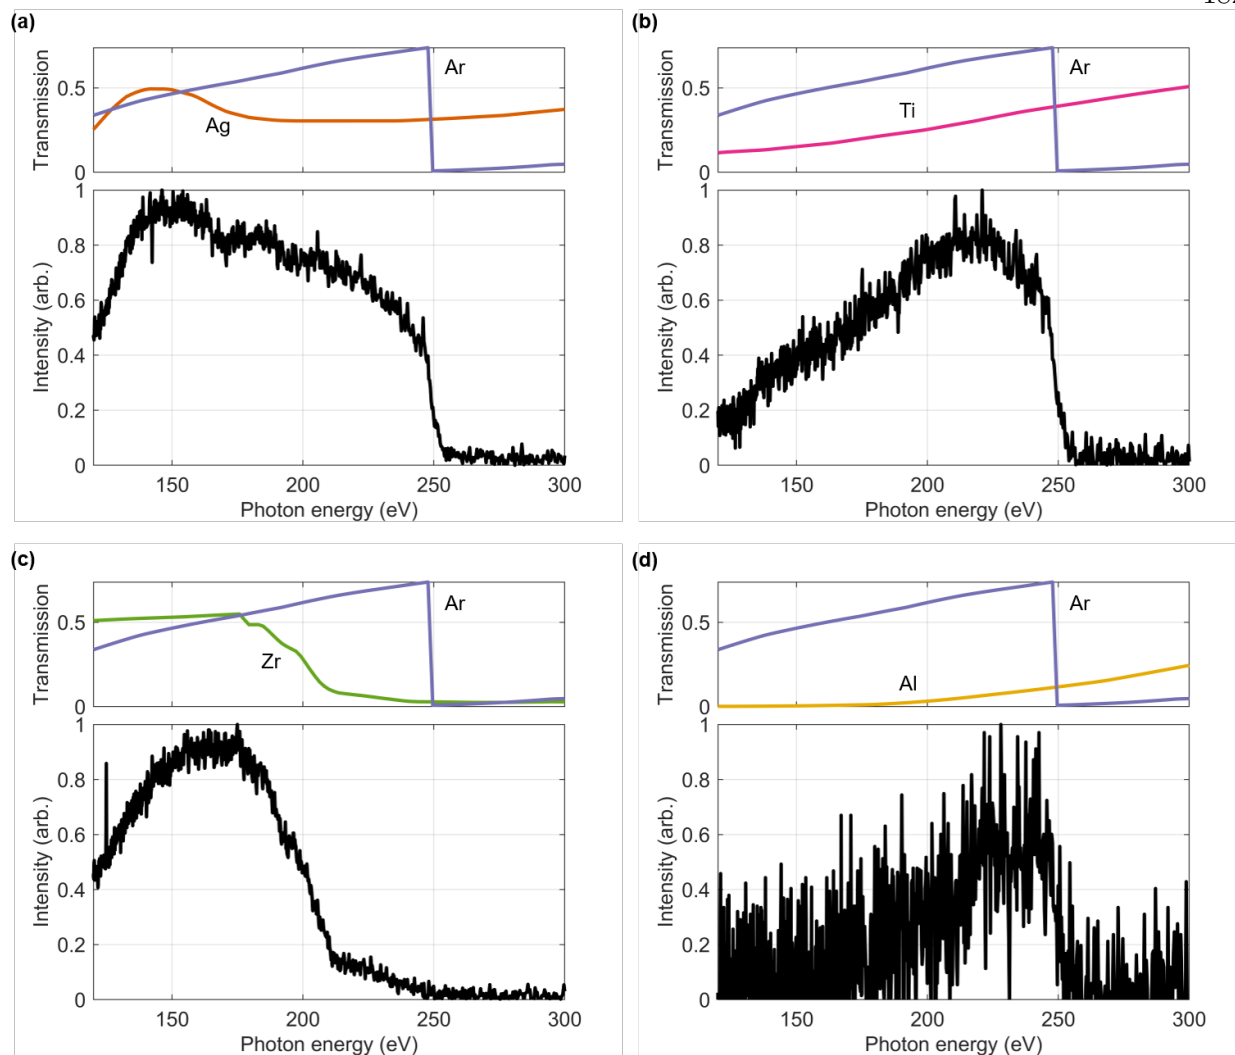


Figure 5.19: **Four measured HHG spectra from the argon filled gas cell using different filters.** All measurements were taken with the Andor CCD camera cooled to -10 C and a 2 min exposure time. Full vertical binning was performed over a 150×2048 pixel region of the sensor. The displayed region corresponds to ~ 880 pixels, or about 43% of the sensor. A background scan was taken using the $100 \mu\text{m}$ thick glass coverslip and subtracted from the measured signal. The top panels display transmission data for the metallic filter used for that measurement: (a) 150 nm silver; (b) 200 nm titanium; (c) 200 nm zirconium; (d) 2×100 nm aluminum. Also displayed is the transmission of 1 cm of argon gas at 30 Torr. Transmission data is from CXRO [9].

as defined from aperture to aperture, is 1.75 mm long (Figure 5.16). Therefore, a photon energies below the argon edge, reabsorption is less severe because the total length of the cell is only around 3 absorption lengths. At photon energies above the edge, any harmonics generated more than around $100 \mu\text{m}$ from the backside aperture will be more than 95% absorbed. This is particularly




Figure 5.20: **Low noise HHG spectrum from the argon filled gas cell with a silver filter, highlighting the stability of the harmonic spectrum over 10 minutes.** Measurement parameters were the same as in Figure 5.19, except that the camera was cooled to -80 C and the integration time was increased to 10 min. Individual harmonics are resolved in the zoomed in inset.

problematic when we consider that the optimal position of the focus is just before the region of high gas density, as described in the macroscopic phase matching section. Secondly, we note that even in the case of ideal phase matching, emission can only out compete reabsorption for a maximum of around 10 attenuation lengths (far fewer in our case of sub-optimal phase matching). This means that the effective number of emitters contributing to the harmonic beam is vastly lower at photon energies above the edge, further decreasing brightness. Clearly, absorption is playing a dominant role in argon HHG. For comparison, we note that the attenuation length of helium at the expected phase matching pressure of 13 bar [37] is 2.9 cm at 750 eV, without any absorption edges across

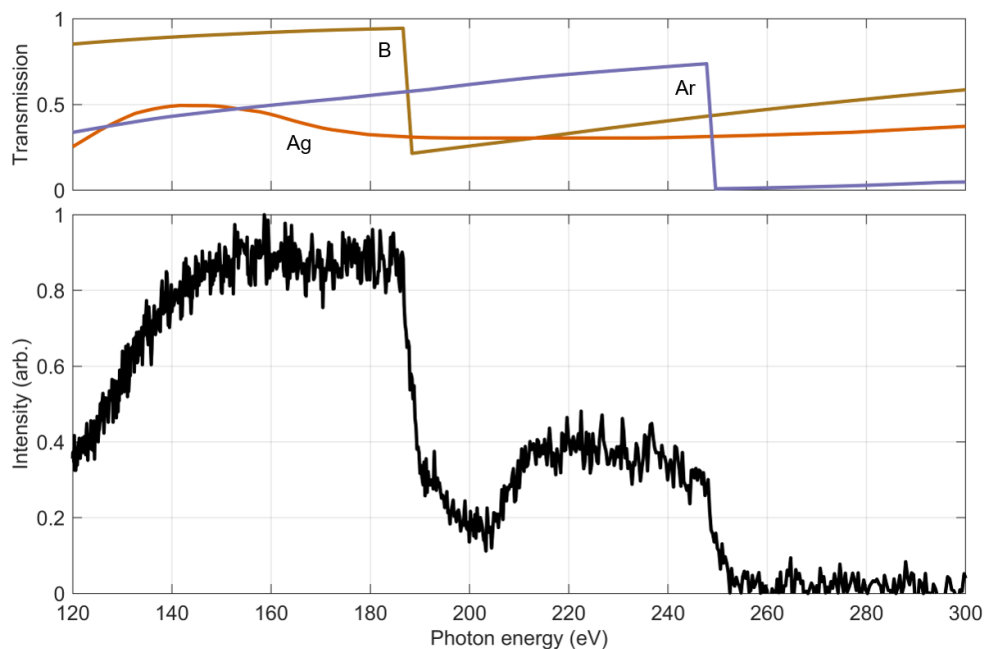


Figure 5.21: **HHG transmission spectroscopy of a 75 nm boron film mounted on a 500 nm parylene substrate.** Boron exhibits a strong absorption feature, which can be seen at 188 eV. A 150 nm silver filter was used to block residual driving light. The measurement was conducted in the same way as those presented in Figure 5.19 (2 min exposure, camera at -10 C).

the soft X-ray region.

Figure 5.20 highlights the stability of the harmonic beam and suggests that this source is useful for low signal-to-noise spectroscopy applications, such as near edge X-ray absorption fine structure (NEXAFS) or magneto-optical Kerr effect spectroscopies. In this measurement of an argon HHG spectrum through a 200 nm silver filter, we reduced noise by cooling the camera to -80 C and increasing the exposure time to 10 min. Beam pointing and spectral fluctuations are common in HHG sources, and over longer integration times, these effects can be problematic for high resolution spectroscopy. We attribute the high fidelity of the spectrum in Figure 5.20 to the high stability of the OPCPA and the HHG source. We further note the observation of individual harmonics, separated by around twice the photon energy of the driving laser.

In Figure 5.21, we present transmission spectroscopy of a 75 nm boron film mounted on a 500 nm parylene substrate. Boron exhibits a strong and sharp absorption feature at 188 eV that




Figure 5.22: **HHG spectra generated in nitrogen and argon under identical measurement conditions.** The relative intensity of the two spectra is accurate, with the spectral peak at around 140 eV around 29% weaker in nitrogen than argon. The measurements were taken with a 200 nm silver filter and a 2 min acquisition time. The nitrogen spectrum just reaches the carbon K-edge (284 eV), corresponding to approximately the 689th harmonic of the driving laser.

is clearly resolved.

HHG was also demonstrated in unaligned nitrogen gas (N₂). A harmonic spectrum generated with nitrogen is presented in Figure 5.22 alongside an argon spectrum that was taken under identical acquisition settings with the same filter (2 min exposure with a 200 nm silver filter). The relative amplitude of the two spectra reflects the relative brightness of the two beams. At the brightest part of the spectra near 140 eV, the nitrogen harmonics are about 29% as bright as the argon harmonics. Unlike monatomic gases, the harmonic emission from diatomic molecular nitrogen depends on the relative orientation of the internuclear axis and the recombining electron wave vector. A separate, lower intensity laser beam is sometimes used to pre-align the nitrogen molecules to improve harmonic yield [210]. Misalignment between the molecular axis and the laser polarization may partially explain the difference in brightness. We further note that the ionization potentials of

argon (15.8 eV) and nitrogen (15.6 eV) are quite similar, so it is unsurprising that harmonics are optimally generated in a similar spectral region. The optimal backing pressure was found to be 2.3 bar - nearly the same as with argon. Unlike argon, however, nitrogen does not exhibit any distinct absorption edges in the 100-300 eV range. We therefore attribute the photon cut off energy of around 285 eV to the phase matching cutoff, approximately corresponding to a harmonic order of 689. The highest energy component of the nitrogen HHG spectrum just reaches the carbon K-edge of 284 eV.

5.5 Below-threshold harmonics

5.5.1 Introduction

Below-threshold harmonics (BTH) refer to the generation of harmonics with photon energies that are below the ionization potential of the generating gas. BTH are occasionally referred to as low-order harmonics to distinguish from high-harmonic generation. BTH generation has received considerable attention as a possible source of coherent VUV radiation for spectroscopy and metrology [211]. Interestingly, BTHs can be used to alter the waveform that drives HHG, resulting in flux (and potentially cut-off energy) enhancements. This has been demonstrated at 800 nm [212] and later at 2.1 μm [213], but never in the mid-infrared. In addition to potentially useful applications, the mechanism behind below-threshold harmonics has been hotly debated by in the literature by theorists and those employing advanced computational methods. In particular, the relative role of perturbative interactions (higher-order Kerr effects, or HOKE) and non-perturbative interactions is controversial [214, 215]. See [216] for a review of this topic.

Furthermore, recent work published by Zuffi et al. [217] introduced and experimentally validated a model to explain a peculiar effect of below-threshold harmonic generation: as the driving laser intensity is increased, the generated harmonic intensity reaches a plateau, and ceases to increase with driving laser intensity. The authors attribute this effect to the perfect balancing of depletion of electrons due to laser-induced ionization and the increasing harmonic yield with

excitation power. The authors also observed that the driving laser intensity to reach the plateau is independent of the driving laser wavelength and the harmonic order, instead depending only on the ionization threshold of the gas, the laser beam confocal parameter, and the interaction length of the medium. Besides providing us an opportunity to test this newly presented hypothesis in a new parameter space, the model also provides a mechanism for us to validate our own experimental gas delivery geometry and laser intensity on target. In practice, both of these parameters can be difficult to precisely measure experimentally. The gas pressure profile that the laser encounters - far from being a simple step function between vacuum and high pressure gas - arises from complicated vacuum and gas flow dynamics, and it can be difficult to quantitatively determine the actual length of the interaction region. Secondly, the intensity on target can also be difficult to verify experimentally. While FROG measurements combined with focal spot size measurements and average power measurements can in principle allow for the calculation of peak laser intensity on target, effects such as temporal pulse pedestals, nonlinear Kerr lensing, plasma defocusing, modal distortions that might not be easy to resolve in the mid-infrared, and even misbehaving optics can all contribute to the actual peak intensity on target being different than predicted.

Up to now, the use of mid-infrared drivers to explore below-threshold harmonics has been primarily explored theoretically [218, 219], with the exception of some work by the Baltuška group which explored third and fifth harmonic generation using their 3.9 μm OPCPA [220]. Nevertheless, mid-infrared lasers may provide fascinating new insights into below-threshold harmonic generation. The ionization potentials of argon and helium are 15.8 eV and 24.6 eV respectively. Most lasers that have been used for exploring the HOKE debate and investigating BTH are Ti:sapphire systems [221, 222], delivering photon energies of around 1.55 eV (800 nm wavelength). For the case of argon gas, the highest order harmonic of a Ti:sapphire laser which can be generated while staying below the ionization potential is the 9th. In contrast, the 3 μm laser presented here delivers photons with energies around 0.42 eV. With a 3 μm laser, up to the 37th harmonic can be generated in argon while still staying below the ionization potential. To our knowledge, this is a parameter space which has not been explored until now.




Figure 5.23: A photograph of 3 μm driven low-order harmonic generation in an argon filled gas cell. See the text for a complete description of the various effects seen here. The output side differential pumping inserts were removed and replaced with a UVFS lens, which held vacuum.

5.5.2 Experiment

Prior to our observation of HHG, we observed below-threshold harmonics in argon up to the 13th order ($\lambda = 227 \text{ nm}$). This BTH signal exhibited many of the hallmarks of HHG - for instance, the light penetrated thin metal filters, and vanished when gas was removed from the interaction region, or when a quarter-wave plate was put in the driving beam. All measurements reported here were taken by focusing about 600 μJ of pulse energy at 130 fs with a 7.5 cm focal length lens to a $1/e^2$ radius of around 35 μm into argon gas. BTH was observed in both an unstructured, 150 μm core hollow-core waveguide and using the gas cell. The gas cell produced a brighter BTH beam. The brightness of the BTH generation from the gas cell increased monotonically with pressure, with a signal beginning from around 2.5 bar and increasing up to around 18 bar of




Figure 5.24: **A representative below-threshold harmonic spectrum measured from argon driven in a gas cell with no filtering optics.** The spectral components are labelled, and can be compared to the photograph in Figure 5.23.

backing pressure. Although the vacuum and gas handling system can in principle sustain higher pressures, we stopped at 18 bar because the system had not been pressure tested beyond that. We note that these pressures are considerably higher than the backing pressure required for optimum argon high-harmonic generation in the gas cell, which was around 2.3 bar. A photograph of the gas cell with a clear acrylic lid generating BTHs is shown in Figure 5.23, illustrating a number of interesting effects. The red light illuminating the entrance window occurs around 625 nm and is the sum-frequency of the pump (1029.5 nm) and the signal (near 1590 nm), generated in the PPLN crystals. This light co-propagates with the 3 μm idler beam even through the compressor. The pink glow of on the entrance side of the gas cell can be attributed to fluorescence of the argon gas. On the output side of the gas cell, the yellow emission corresponds to the 5th harmonic (near 590 nm) and the blue emission, visible through a 500 nm shortpass filter, is the 7th harmonic near 422 nm.

A spectrometer (HR4000 from OceanOptics) sensitive in the range of 200 - 1100 nm was placed behind the gas cell. A UVFS lens was placed on the back side of the chassis using an O-ring to provide a vacuum seal. This helped focus the BTH into the spectrometer. The resulting




Figure 5.25: **Low-order harmonic generation spatial modes.**(a) Low-order harmonic beam as measured through a 200 nm silver filter by an X-ray silicon CCD camera located 154 cm behind an argon-filled gas cell backed by 14.8 bar of pressure. (b) Low-order harmonic beam as measured on a silicon CCD beam profiling camera. An uncoated CaF_2 window was placed on the back side of the differential pumping chassis to allow for measurement of the beam in air. The beam was generated from an argon filled gas cell backed by 5.4 bar, and measured 14 cm from the back of the cell. The $1/e^2$ diameter of the beam is 4 mm.

spectrum is shown in Figure 5.24.

Two spatial modes from the BTH beam are shown in Figure 5.25. A DataRay beam profiling CCD camera was placed 14 cm after the gas cell, with a CaF_2 window at the back side of the chassis and both differential pumping ports removed to avoid clipping the beam. The $1/e^2$ diameter of the beam was measured to be ~ 4 mm, with an excellent mode. We note that the relative brightness of the third and fifth harmonics has not been well established experimentally or theoretically. Nevertheless, initial indications are that the third harmonic beam diverges rapidly. Using Equation 5.16, and assuming a wavelength of 590 nm (fifth harmonic) and noting the near-Gaussian mode, we can predict the size of the beam upon emission at the gas cell to have a $1/e^2$ radius of around $\omega_0 = 13 \mu\text{m}$, smaller than the predicted $1/e^2$ radius of $3 \mu\text{m}$ driving laser ($35 \mu\text{m}$). We further note that silver, titanium and zirconium foils in the range of 100-200 nm thickness all transmitted a BTH signal which was brighter than the eventual HHG beam. Figure 5.25(a) shows the BTH beam transmitted through a 200 nm silver foil on the X-ray camera; a dimmer beam was even observed on the X-ray camera through a pair of two 200 nm silver foils. In the visible spectrum, thin metallic foils are more transmissive to shorter wavelengths [223]. Furthermore, surface plasmon enhanced transmission of light through thin silver films has been reported [224]. Although the precise mechanism for BTH transmission in the foils has not been established, we caution researchers

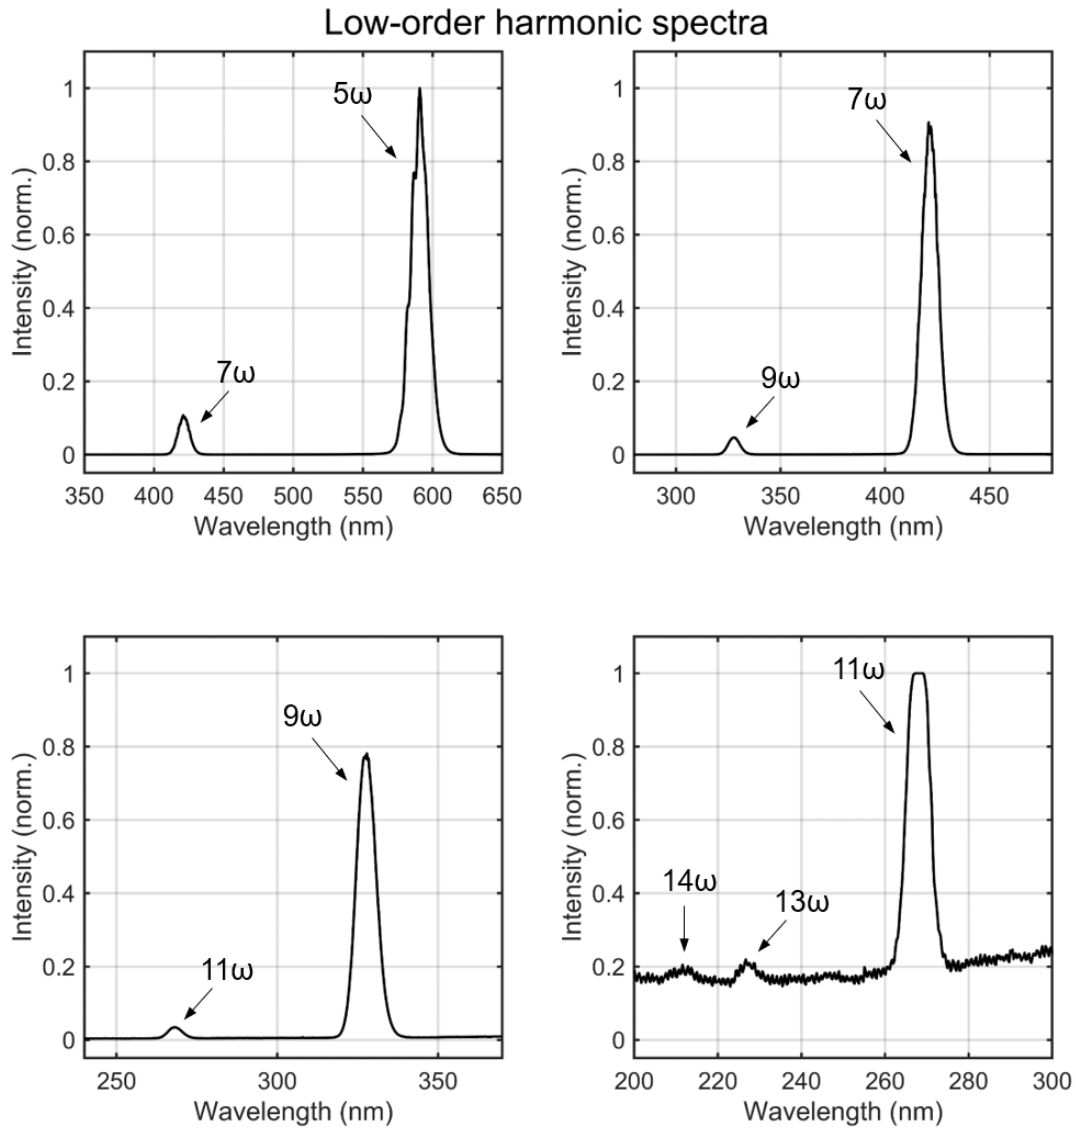


Figure 5.26: **Measured below-threshold harmonic spectra.** Below-threshold harmonics span more than 5 orders of magnitude in brightness between the 5th and 13th harmonics. To estimate the relative brightness of each harmonic, we took a series of spectra with different integration times and, for the case of the 5th harmonic, optical filters to prevent saturating the spectrometer. We then adjusted the counts for exposure time and attenuation, and integrated the area under each peak to create a “high dynamic range” spectrum, which is presented in Figure 5.27. Unexpectedly, we observed a very small spectral peak near 211 nm, the 14th harmonic of the driving laser. However, we caution that this is near the edge of the spectral response window of the spectrometer, and may be accounted for by measurement artifacts such as higher diffractive orders.

in mid-IR driven HHG that low-order harmonics may have unexpectedly high transmission through metallic filters.

5.5.3 Perturbative effects in below-threshold harmonics

Two recent papers by Spott et al. have presented calculations for below-threshold harmonic generation using both *ab initio* numerical methods as well as perturbative power series expansions of the third and fifth order susceptibilities. The authors conclude that the transition from perturbative to nonperturbative interactions occurs with driving laser intensities at about 10^{13} W/cm². Calculations were done assuming an 800 nm driving laser [214] and a 1600 nm driving laser [215]. Both papers arrive at nearly the same conclusion regardless of driving wavelength: near intensities of 1×10^{13} to 2×10^{13} W/cm², a particular transition is noted. For lower intensities, the contributions of the higher-order terms is much smaller than the lowest-order term, such that the perturbative power series expansion converges. For intensities above the threshold, the sum of the higher order terms reaches about 10% of the contribution of the lowest order term. Furthermore, we note that in the perturbative regime, the intensity of the *n*th harmonic should scale to the *n*th power of the laser intensity [101].

Here, we present preliminary results in the context of these numerical studies, but caution that more rigorous experimental techniques and analysis are needed to draw conclusions. A BTH beam was generated in an argon gas cell backed by 15 bar using the focusing conditions described in the introduction, with driving laser power varied using a waveplate in front of the grating compressor. Spectroscopy was performed on the BTH beam using an OceanOptics HR4000 spectrometer with sensitivity to down to 200 nm. In order to make a semi-quantitative measurement of harmonic intensities over such a large dynamic range, we optimized each spectral measurement for a single harmonic peak by adjusting the integration time and, for the case of measuring the very bright 5th harmonic beam, through the addition of calibrated neutral density filters before the spectrometer. Examples of spectra showing the decreasing amplitude of subsequent harmonics are shown in Figure 5.26. This technique allowed us to stitch together a high dynamic-range spectrum. Two




Figure 5.27: **High dynamic range reconstructed spectra of below-threshold harmonics.** We took separate spectra of the type shown in Figure 5.26 for each harmonic. The area under the relevant harmonic was integrated and is represented by a single data point. Data was used directly from the HR4000 spectrometer without any spectral responsivity post-processing, introducing an unknown spectrally dependent error in integrated counts. We show data for two driving laser intensities: 9.6×10^{13} W/cm² and 1.8×10^{14} W/cm². References [214, 215] predict that a transition between perturbative to nonperturbative interactions should occur around 10^{13} W/cm². Nonperturbative effects are not immediately clear from the data.

examples of these high-dynamic range harmonic spectra are presented in Figure 5.27, for driving laser intensities of 9.6×10^{13} and 1.8×10^{14} W/cm², corresponding to incident pulse energies of 239 and 445 μ J respectively. We note that analysis using the BTH plateau is presented in the next section to verify that intensities on target agree with those predicted by pulse energy, focal spot size and pulse duration measurements. Both of the intensities shown in Figure 5.27 are around an order of magnitude greater than the predicted intensity for the onset of nonperturbative effects by Spott et al. Nevertheless, nonperturbative effects for the first 13 harmonics are not immediately apparent from the raw data, although we note that more careful harmonic analysis and comparison with theory should be made before concluding that nonperturbative effects are not present in some capacity. In particular, the USB4000 spectrometer is characterized by an instrument response func-

tion (IRF) that depends on the detector response, diffraction grating efficiency curve, and other spectrally dependent effects. The IRF is not provided by the manufacturer. In the data presented here, the IRF is not corrected for, introducing error in comparing the brightness of different harmonics. In future work, a dispersive CaF₂ prism may be used for a more quantitative comparison of harmonic intensities.

5.5.4 Below-threshold harmonic plateau: verifying intensity and interaction length

Zuffi et al. [217] presented an experimental study of BTH from noble gases driven by a 550 μJ pulse energy Ti:sapphire with pulse durations in the range of 25 to 200 fs. They observed that harmonic yield for the third, fifth and seventh harmonic increased for driving intensities below some threshold I_{p-th} , and plateaued above that threshold. The third harmonic intensity as a function of driving laser intensity is reproduced from Zuffi et al. in Figure 5.28, where the intensity plateaus at the predicted value of 1.3×10^{14} W/cm². Below I_{p-th} , the authors found that the third harmonic intensity can be fit to a simple power law (dashed line in Figure 5.28):

$$I_{3\omega} = AI_L^p \quad (5.17)$$

where I_L is the laser intensity and p is the harmonic order.

We approximately measured the power the seventh harmonic as a function of laser intensity in the following way. We focused our 3 μm laser into an argon filled gas cell. We adjusted the laser power with a half wave plate in front of the compressor (which acts as a polarizer). We placed a notch filter (part FESH0500 from Thorlabs) after the gas cell, which has OD >5 for wavelengths >500 nm, ~90% transmission for wavelengths between 390 and 500 nm, and around OD 2 for wavelengths <390 nm. The wavelengths of the 5th, 7th and 9th harmonics are around 590 nm, 421 nm and 328 nm, respectively. We also note that the 9th harmonic is typically greater than an order of magnitude weaker than the 7th harmonic. In this way, we isolate the 7th harmonic with high contrast. The isolated 7th harmonic is focused onto a silicon photodiode using a UVFS lens, which also acted as the backside vacuum window of the chassis. The amplitude of the photodiode

signal was measured using an oscilloscope, and is plotted in mV in Figure 5.29, with the RMS fluctuations plotted as the error bars. This signal was ensured to be in the linear response regime of the photodiode.

Using a model which considers the depletion of the bound valence electrons by the driving laser, and considers only the first ionization of the gas, Zuffi et al. develop a simple formula for I_{p-th} , which is independent of harmonic order and depends only on the ionization threshold of the gas I_{ion-th} , the interaction length l and the confocal parameter $b = 2z_R$:

$$I_{p-th} = \frac{I_{ion-th}}{2} \left(1 + \frac{l^2}{b^2} \right) \quad (5.18)$$

where the Rayleigh range is $z_R = \pi\omega_0^2/\lambda$. The first ionization threshold for argon is 2.4×10^{14} W/cm² [217]. In our experiment, the 3 μm central wavelength beam is focused to a $1/e^2$ radius of $\omega_0 = 35 \mu\text{m}$. We assume an interaction length $l = 1.75$ mm, which is the distance between the entrance and exit apertures of the gas cell (illustrated in Figure 5.29 inset). By applying these parameters to Equation 5.18, we calculate an expected $I_{p-th} = 1.76 \times 10^{14}$ W/cm², corresponding to an applied pulse energy of around 430 μJ . This intensity is illustrated with a vertical line in Figure 5.29, and agrees well with the onset of the measured plateau. We note that the laser intensity




Figure 5.28: **Third-harmonic intensity versus driving laser intensity for below-threshold harmonics in argon gas driven by a Ti:sapphire laser.** Note that the plateau onset intensity I_{p-th} is well described by Equation 5.18. Figure from Reference [217].




Figure 5.29: **Seventh-harmonic intensity versus driving laser intensity for below-threshold harmonics in argon gas driven by the 3 μm wavelength mid-IR OPCPA.** Intensity was measured using a silicon photodiode read out by an oscilloscope, with voltage displayed on the vertical axis in mV. The error bars are the RMS fluctuations. While the data does not fit the expected power law (possibly due to the photodiode not sampling the full beamwidth) the plateau threshold intensity is well predicted by Equation 5.18, suggesting that our laser intensity on target is well predicted by secondary measurements of power, spot size and pulse duration, and that the interaction length can be considered the full length of the gas cell from entrance to exit aperture. Inset: the gas cell used in this experiment, with Gaussian beam propagation overlaid in red.

plotted on the horizontal axis in Figure 5.29 was calculated using our independently measured beam parameters: a 130 fs pulse measured with SHG-FROG, a radius of 35 μm estimated with a VOx microbolometric beam profiling camera and a pulse energy deduced from average power measurements using a thermal power meter. We therefore find good agreement between our own estimates of laser intensity and interaction length with theory.

Chapter 6

Concluding remarks

The field of near- and mid-infrared ultrafast laser physics is in an exciting stage of rapid development. As these laser technologies mature, and new technologies are developed, they will be well positioned to find numerous applications outside of the optics laboratory. In this thesis, we have highlighted one such application: the generation of coherent soft X-ray radiation through high-harmonic generation, but many others abound - from molecular spectroscopy in the so-called fingerprint region [225] to driving ion accelerators for cancer therapy [53].

In our own laboratory, there is no shortage of promising future directions to take our work. The greatest promise of mid-infrared driven harmonics has yet to be realized on our optical table: the generation of keV-level harmonic photons in helium gas. Nevertheless, the ingredients are there: previous work with a lower repetition rate system has demonstrated that it is possible [3], and our own HHG demonstration in argon and nitrogen has shown us that we are tantalizingly close to the needed peak intensity. On the month-scale horizon, the OPCPA will undergo two upgrades which may get us to this target. Firstly, an upgrade to the grism stretcher will increase the transmission bandwidth of the device by incorporating larger gratings and prisms, allowing us to seed the OPCPA with the full bandwidth of the front-end. The influence of this upgrade on the amplified pulse duration has yet to be seen, but numerical simulations suggest that it holds promise to decrease the amplified pulse duration and increase peak intensity on target. Secondly, custom, large aperture PPLN crystals with lengths of 1 and 2 mm will soon be incorporated into the later stages of the OPA. Simulations presented in Chapter 4 suggest that this may significantly

improve the conversion efficiency of pump energy to amplified 3 μm energy.

On the slightly longer term horizon, there are several avenues towards scaling the pulse energy of the OPCPA. In the appendix that follows, we numerically simulate the performance of alternative nonlinear crystals, and find that incorporating either PPLT or PPKTP in the later amplification stages of the OPA chain may present an immediate solution to scaling pulse energy, depending on their damage threshold characteristics. Before 3 μm energy increases of more than 2x can be realized, the pump laser system must be upgraded to deliver more pulse energy. While adding to system complexity, the use of few-pass booster amplifiers is well established in the literature [168]. Should periodically-poled crystals be found to be insufficient for power scaling, the simulations in Appendix A indicate that bulk potassium niobate may be a viable candidate in the power OPA stages without modification to the pump pulse duration - however, at present, no vendor for potassium niobate has been located. The readily available bulk crystals MgO-doped lithium niobate and potassium titanyl arsenate may hold promise for power scaling, but will likely require the use of a higher peak intensity/lower fluence pump that is only attainable through compression of the pump pulses. While compression of cryo-Yb:YAG amplifier pulses is challenging due to the narrow bandwidth, it has been demonstrated in the literature using a folded compressor design [168]. If compression is pursued, a pulse duration optimization may be done to balance beam sizes, pulse energy, bandwidth and other considerations.

A recent trend in mid-IR OPCPA technology is the use of 2 μm pump lasers based on holmium-doped gain media [80, 81]. These lasers can be used to pump bulk non-oxide crystals, such as ZGP and BGSe [226], which can exhibit impressive nonlinearities and phase-matching properties. Our group is actively developing such a system, with plans to generate a signal near 3 μm and an idler near 6.5 μm . Bringing the pump wavelength closer to the amplified wavelength of interest increases the energy conversion efficiency of OPA. This system may be an interesting candidate to explore the use of idler absorption to prevent back-conversion and enhance conversion efficiency [66, 67, 68] - for instance, CSP naturally starts to absorb at around 6-7 μm [227]. Another active area of research in our group is the use of large mode area ytterbium rod amplifiers as compact, robust

high-average power pumps for nonlinear processes such as OPA and VUV generation. The high stability, high single pass gain, and good thermal management properties of large mode area fiber amplifiers may also make them promising pre-amplifiers in high pulse energy hybrid laser amplifier architectures. Direct laser amplification of broadband, long wavelength pulses using chromium- and iron-doped chalcogenides, as well as holmium- and thulium-doped gain media, is another exciting avenue for soft X-ray high-harmonic generation and other applications.

High-harmonic generation also presents a wellspring of exciting future work. We have already seen in Chapter 5 that the combination of low critical ionization fraction, high phase-matching gas pressure, and high SXR transparency will make bright mid-IR driven HHG in helium challenging without the large interaction volumes afforded by the wave-guide geometry. On the other hand, small-diameter unstructured hollow-core waveguides are particularly lossy in the mid-IR: the attenuation coefficient is proportional to $\lambda^2/\text{radius}^3$ [228, 158]. To solve this dilemma, we are actively exploring the use of gas-filled anti-resonant hollow core (AR-HC) fiber waveguides for HHG, which can be engineered to exhibit low transmission losses and low field amplitude in the glass fiber structure even in the mid-IR [229, 230, 231, 232]. In a separate experiment, building on the clear observation of below-threshold harmonics from the 3 μm OPCPA, interesting work may be considered in collaboration with our computational colleagues to better understand the transition between perturbative and non-perturbative harmonic generation, as well as multi-color waveform control for enhanced HHG. We also note that recently developed X-ray CMOS cameras have been used by others to characterize the photon energy of soft X-ray radiation incident on the detector based on the number of generated electrons [233] - potentially paving the way for high frame rate, SXR hyperspectral imaging and other applications. Lastly, we consider the macroscopic physics of HHG itself, which is in general challenging to optimize and study experimentally. The versatile, modular HHG source presented in Chapter 5 offers a unique combination of repeatability and ease of modification - which we plan to leverage to experimentally study important effects such as phase matching and reabsorption.

Bibliography

- [1] Paul Gorenstein and Wallace H. Tucker. Soft x-ray sources. Annual Review of Astronomy and Astrophysics, 14(1):373–416, 1976.
- [2] Ariel Gordon and Franz X Kärtner. Scaling of keV HHG photon yield with drive wavelength. Optics Express, 13(8):2941–2947, 2005.
- [3] Tenio Popmintchev, Ming-Chang Chen, Dimitar Popmintchev, Paul Arpin, Susannah Brown, Skirmantas Ališauskas, Giedrius Andriukaitis, Tadas Balčiunas, Oliver D Mücke, Audrius Pugzlys, et al. Bright coherent ultrahigh harmonics in the keV x-ray regime from mid-infrared femtosecond lasers. Science, 336(6086):1287–1291, 2012.
- [4] Christine Grienberger, Andrea Giovannucci, William Zeiger, and Carlos Portera-Cailliau. Two-photon calcium imaging of neuronal activity. Nature Reviews Methods Primers, 2(1):67, 2022.
- [5] Benjamin P Abbott, Richard Abbott, TDe Abbott, MR Abernathy, Fausto Acernese, Kendall Ackley, Carl Adams, Thomas Adams, Paolo Addesso, RX Adhikari, et al. Observation of gravitational waves from a binary black hole merger. Physical Review Letters, 116(6):061102, 2016.
- [6] T.H. Maiman. Stimulated optical radiation in ruby. Nature, 187(4736):493–494, 1960.
- [7] Ernst Abbe. Beiträge zur theorie des mikroskops und der mikroskopischen wahrnehmung. Archiv für Mikroskopische Anatomie, 9(1):413–468, 1873.
- [8] P.B. Corkum and Ferenc Krausz. Attosecond science. Nature Physics, 3(6):381–387, 2007.
- [9] Burton L Henke, Eric M Gullikson, and John C Davis. X-ray interactions: photoabsorption, scattering, transmission, and reflection at $e= 50\text{-}30,000$ ev, $z= 1\text{-}92$. Atomic data and nuclear data tables, 54(2):181–342, 1993.
- [10] Wilhelm Conrad Röntgen. Ueber eine neue Art von Strahlen. Phys. Med. Gesellschaft, 1895.
- [11] FK Goward and David Errington Barnes. Experimental 8 mev. synchrotron for electron acceleration. Nature, 158(4012):413–413, 1946.
- [12] Eberhard J Jaeschke, Shaukat Khan, Jochen R Schneider, and Jerome B Hastings. Synchrotron light sources and free-electron lasers: accelerator physics, instrumentation and science applications. 2016.

- [13] Shaukat Khan. Ultrashort pulses from synchrotron radiation sources. Synchrotron Light Sources and Free-Electron Lasers: Accelerator Physics, Instrumentation and Science Applications, pages 51–81, 2020.
- [14] A Madsen, J Hallmann, T Roth, and G Ansaldo. Scientific instrument materials imaging and dynamics (mid). Technical Design Report TR-2013-005. European XFEL, Hamburg, Germany, 2013.
- [15] Daniel E Rivas, Svitozar Serkez, Thomas M Baumann, Rebecca Boll, Marie Kristin Czwalinna, Simon Dold, Alberto de Fanis, Natalia Gerasimova, Patrik Grychtol, Björn Lautenschlager, et al. High-temporal-resolution x-ray spectroscopy with free-electron and optical lasers. Optica, 9(4):429–430, 2022.
- [16] Felicie Albert, AGR Thomas, SPD Mangles, S Banerjee, Sébastien Corde, Alessandro Flacco, Michael Litos, D Neely, J Vieira, Zufikar Najmudin, et al. Laser wakefield accelerator based light sources: potential applications and requirements. Plasma Physics and Controlled Fusion, 56(8):084015, 2014.
- [17] Wentao Wang, Ke Feng, Lintong Ke, Changhai Yu, Yi Xu, Rong Qi, Yu Chen, Zhiyong Qin, Zhijun Zhang, Ming Fang, et al. Free-electron lasing at 27 nanometres based on a laser wakefield accelerator. Nature, 595(7868):516–520, 2021.
- [18] Thomas Gaumnitz, Arohi Jain, Yoann Pertot, Martin Huppert, Inga Jordan, Fernando Ardana-Lamas, and Hans Jakob Wörner. Streaking of 43-attosecond soft-x-ray pulses generated by a passively cep-stable mid-infrared driver. Optics Express, 25(22):27506–27518, 2017.
- [19] Gerry McDermott, Mark A Le Gros, Christian G Knoechel, Maho Uchida, and Carolyn A Larabell. Soft x-ray tomography and cryogenic light microscopy: the cool combination in cellular imaging. Trends in Cell Biology, 19(11):587–595, 2009.
- [20] Weiwei Gu, Laurence D Etkin, Mark A Le Gros, and Carolyn A Larabell. X-ray tomography of schizosaccharomyces pombe. Differentiation, 75(6):529–535, 2007.
- [21] Gerd Schneider. Cryo x-ray microscopy with high spatial resolution in amplitude and phase contrast. Ultramicroscopy, 75(2):85–104, 1998.
- [22] Werner Meyer-Ilse, Donna Hamamoto, Ajit Nair, SA Lelievre, Gregory Denbeaux, L Johnson, Angelic Lucero Pearson, Deborah Yager, Mark A Legros, and Carolyn A Larabell. High resolution protein localization using soft x-ray microscopy. Journal of Microscopy, 201(3):395–403, 2001.
- [23] Jianwei Miao, Pambos Charalambous, Janos Kirz, and David Sayre. Extending the methodology of x-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens. Nature, 400(6742):342–344, 1999.
- [24] Jianwei Miao, Tetsuya Ishikawa, Ian K Robinson, and Margaret M Murnane. Beyond crystallography: Diffractive imaging using coherent x-ray light sources. Science, 348(6234):530–535, 2015.

- [25] Bin Wang, Nathan J Brooks, Peter Johnsen, Nicholas W Jenkins, Yuka Esashi, Iona Binnie, Michael Tanksalvala, Henry C Kapteyn, and Margaret M Murnane. High-fidelity ptychographic imaging of highly periodic structures enabled by vortex high harmonic beams. Optica, 10(9):1245–1252, 2023.
- [26] Vivek M Prabhu, Sharadha Sambasivan, Daniel Fischer, Linda K Sundberg, and Robert D Allen. Quantitative depth profiling of photoacid generators in photoresist materials by near-edge x-ray absorption fine structure spectroscopy. Applied Surface Science, 253(2):1010–1014, 2006.
- [27] F Willems, S Sharma, C v Korff Schmising, JK Dewhurst, Leandro Salemi, D Schick, P Hensing, C Strüber, WD Engel, and S Eisebitt. Magneto-optical functions at the 3 p resonances of fe, co, and ni: Ab initio description and experiment. Physical Review Letters, 122(21):217202, 2019.
- [28] A McPherson, G Gibson, H Jara, U Johann, Ting S Luk, IA McIntyre, Keith Boyer, and Charles K Rhodes. Studies of multiphoton production of vacuum-ultraviolet radiation in the rare gases. Journal of the Optical Society of America B, 4(4):595–601, 1987.
- [29] M Ferray, Anne L’Huillier, XF Li, LA Lompre, G Mainfray, and C Manus. Multiple-harmonic conversion of 1064 nm radiation in rare gases. Journal of Physics B: Atomic, Molecular and Optical Physics, 21(3):L31, 1988.
- [30] K. C. Kulander and B. W. Shore. Calculations of multiple-harmonic conversion of 1064-nm radiation in xe. Physical Review Letters, 62:524–526, Jan 1989.
- [31] Jeffrey L Krause, Kenneth J Schafer, and Kenneth C Kulander. High-order harmonic generation from atoms and ions in the high intensity regime. Physical Review Letters, 68(24):3535, 1992.
- [32] P. B. Corkum. Plasma perspective on strong field multiphoton ionization. Physical Review Letters, 71:1994–1997, Sep 1993.
- [33] Kjeld SE Eikema. Twisted high-harmonic generation. Nature Photonics, 9(11):710–712, 2015.
- [34] Zenghu Chang, Andy Rundquist, Haiwen Wang, Margaret M Murnane, and Henry C Kapteyn. Generation of coherent soft x rays at 2.7 nm using high harmonics. Physical Review Letters, 79(16):2967, 1997.
- [35] Ch Spielmann, NH Burnett, S Sartania, R Koppitsch, M Schnurer, Ch Kan, M Lenzner, P Wobrauschek, and F Krausz. Generation of coherent x-rays in the water window using 5-femtosecond laser pulses. Science, 278(5338):661–664, 1997.
- [36] Tenio Popmintchev, Ming-Chang Chen, Paul Arpin, Margaret M Murnane, and Henry C Kapteyn. The attosecond nonlinear optics of bright coherent x-ray generation. Nature Photonics, 4(12):822–832, 2010.
- [37] 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.

- [38] E. Constant, D. Garzella, P. Breger, E. Mével, Ch. Dorrer, C. Le Blanc, F. Salin, and P. Agostini. Optimizing high harmonic generation in absorbing gases: Model and experiment. Physical Review Letters, 82:1668–1671, Feb 1999.
- [39] 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.
- [40] J. Tate, T. Augustine, H. G. Muller, P. Salières, P. Agostini, and L. F. DiMauro. Scaling of wave-packet dynamics in an intense midinfrared field. Physical Review Letters, 98:013901, Jan 2007.
- [41] Tenio Popmintchev. Tunable Ultrafast Coherent Light in the Soft and Hard X-ray Regions of the Spectrum: Phase Matching of Extreme High-Order Harmonic Generation. PhD thesis, University of Colorado at Boulder, 2009.
- [42] Elias Snitzer. Cylindrical dielectric waveguide modes. Journal of the Optical Society of America A, 51(5):491–498, 1961.
- [43] Advanced scientific background on the nobel prize in physics 2009. Available at <https://www.nobelprize.org/prizes/physics/2009/advanced-information/>.
- [44] Elias Snitzer. Optical maser action of Nd^{3+} in a barium crown glass. Physical Review Letters, 7(12):444, 1961.
- [45] Biography of Elias Snitzer. Available at https://www.optica.org/History/Biographies/bios/Elias_Snitzer.
- [46] Charles J Koester and Elias Snitzer. Amplification in a fiber laser. Applied Optics, 3(10):1182–1186, 1964.
- [47] Stewart Ash. The history of the repeater. Available at <https://atlantic-cable.com/Article/SA/65/index.htm>.
- [48] Stuart Abbott. Review of 20 years of undersea optical fiber transmission system development and deployment since tat-8. In 2008 34th European Conference on Optical Communication, pages 1–4. IEEE, 2008.
- [49] Tarun Kumar Gangopadhyay, Pathik Kumbhakar, and Mrinal Kanti Mandal. Photonics and fiber optics: foundations and applications. CRC Press, 2019.
- [50] RJ Mears, L Reekie, SB Poole, and DN Payne. Low-threshold tunable cw and q-switched fibre laser operating at $1.55 \mu\text{m}$. Electronics Letters, 3(22):159–160, 1986.
- [51] Emmanuel Desurvire, Jay R Simpson, and PC Becker. High-gain erbium-doped traveling-wave fiber amplifier. Optics Letters, 12(11):888–890, 1987.
- [52] Jan Rothhardt, Steffen Hädrich, Yariv Shamir, Maxim Tschernajew, Robert Klas, Armin Hoffmann, Getnet K Tadesse, Arno Klenke, Thomas Gottschall, Tino Eidam, et al. High-repetition-rate and high-photon-flux 70 eV high-harmonic source for coincidence ion imaging of gas-phase molecules. Optics Express, 24(16):18133–18147, 2016.

- [53] Leily Kiani, Tong Zhou, Seung-Whan Bahk, Jake Bromage, David Bruhwiler, E Michael Campbell, Zenghu Chang, Enam Chowdhury, Michael Downer, Qiang Du, et al. High average power ultrafast laser technologies for driving future advanced accelerators. Journal of Instrumentation, 18(08):T08006, 2023.
- [54] Karel Veselský, Jan Sulc, Helena Jelínková, and Karel Nejezchleb. Yb concentration influence of on thermal lensing in Yb:LuAG and Yb:YAG lasers at cryogenic temperatures: modeling and experimental study. In W. Andrew Clarkson and Ramesh K. Shori, editors, Solid State Lasers XXIX: Technology and Devices, volume 11259, page 112590Z. International Society for Optics and Photonics, SPIE, 2020.
- [55] Why use cryogenic cooling? Available at <https://www.kmlabs.com/cryo-app-note-short>.
- [56] Adolf Giesen, H Hügel, A Voss, K Wittig, U Brauch, and H Opower. Scalable concept for diode-pumped high-power solid-state lasers. Applied Physics B, 58:365–372, 1994.
- [57] Jonathan Fischer, Alexander-Cornelius Heinrich, Simon Maier, Julian Jungwirth, Daniele Brida, and Alfred Leitenstorfer. 615 fs pulses with 17 mj energy generated by an yb:thin-disk amplifier at 3 khz repetition rate. Optics Letters, 41(2):246–249, Jan 2016.
- [58] Ihsan Fsaifes, Louis Daniault, Severine Bellanger, Matthieu Veinhard, Jerome Bourderionnet, Christian Larat, Eric Lallier, Eric Durand, Arnaud Brignon, and Jean-Christophe Chanteloup. Coherent beam combining of 61 femtosecond fiber amplifiers. Optics Express, 28(14):20152–20161, Jul 2020.
- [59] Weijian Sha, Jean-Christophe Chanteloup, and Gérard Mourou. Ultrafast fiber technologies for compact laser wake field in medical application. Photonics, 9(6), 2022.
- [60] NKT Photonics. aeroGAIN ROD: high power PM ytterbium rod fiber gain module, 2023.
- [61] Report of workshop on laser technology for k-bella and beyond. May 2017.
- [62] Advanced scientific background on the nobel prize in physics 2018. Available at <https://www.nobelprize.org/prizes/physics/2018/advanced-information/>.
- [63] Donna Strickland and Gerard Mourou. Compression of amplified chirped optical pulses. Optics Communications, 55(6):447–449, 1985.
- [64] Simon Hooker and Colin Webb. Laser physics, volume 9. Oxford University Press, 2010.
- [65] Shian Zhou, Frank W Wise, and Dimitre G Ouzounov. Divided-pulse amplification of ultra-short pulses. Optics Letters, 32(7):871–873, 2007.
- [66] TH Jeys. Multipass optical parametric amplifier. Optics Letters, 21(16):1229–1231, 1996.
- [67] Jingui Ma, Jing Wang, Bingjie Zhou, Peng Yuan, Guoqiang Xie, Kainan Xiong, Yanqing Zheng, Heyuan Zhu, and Liejia Qian. Broadband, efficient, and robust quasi-parametric chirped-pulse amplification. Optics Express, 25(21):25149–25164, 2017.
- [68] Noah Flemens, Nicolas Swenson, and Jeffrey Moses. Efficient parametric amplification via simultaneous second harmonic generation. Optics Express, 29(19):30590–30609, 2021.

- [69] A Dubietis, G Jonušauskas, and A Piskarskas. Powerful femtosecond pulse generation by chirped and stretched pulse parametric amplification in bbo crystal. Optics Communications, 88(4-6):437–440, 1992.
- [70] Audrius Dubietis and Aidas Matijošius. Table-top optical parametric chirped pulse amplifiers: past and present. Opto-Electron. Adv, 6(3):220046, 2023.
- [71] Giedrius Andriukaitis, Tadas Balčiūnas, Skirmantas Ališauskas, Audrius Pugžlys, Andrius Baltuška, Tenio Popmintchev, Ming-Chang Chen, Margaret M Murnane, and Henry C Kapteyn. 90 gw peak power few-cycle mid-infrared pulses from an optical parametric amplifier. Optics Letters, 36(15):2755–2757, 2011.
- [72] Dimitar Popmintchev, Benjamin R Galloway, Ming-Chang Chen, Franklin Dollar, Christopher A Mancuso, Amelia Hankla, Luis Miaja-Avila, Galen O’Neil, Justin M Shaw, Guangyu Fan, et al. Near-and extended-edge x-ray-absorption fine-structure spectroscopy using ultrafast coherent high-order harmonic supercontinua. Physical Review Letters, 120(9):093002, 2018.
- [73] Ugaitz Elu, Matthias Baudisch, Hugo Pires, Francesco Tani, Michael H Frosz, Felix Köttig, Alexey Ermolov, Philip St J Russell, and Jens Biegert. High average power and single-cycle pulses from a mid-ir optical parametric chirped pulse amplifier. Optica, 4(9):1024–1029, 2017.
- [74] Ugaitz Elu, Luke Maidment, Lenard Vamos, Francesco Tani, David Novoa, Michael H Frosz, Valeriy Badikov, Dmitrii Badikov, Valentin Petrov, Philip St. J. Russell, et al. Seven-octave high-brightness and carrier-envelope-phase-stable light source. Nature Photonics, 15(4):277–280, 2021.
- [75] Jordi Alcalà, Utso Bhattacharya, Jens Biegert, Marcelo Ciappina, Ugaitz Elu, Tobias Graß, Piotr T Grochowski, Maciej Lewenstein, Anna Palau, Themistoklis PH Sidiropoulos, et al. High-harmonic spectroscopy of quantum phase transitions in a high-*tc* superconductor. Proceedings of the National Academy of Sciences, 119(40):e2207766119, 2022.
- [76] U Elu, T Steinle, D Sánchez, L Maidment, K Zawilski, P Schunemann, UD Zeitner, C Simon-Boisson, and J Biegert. Table-top high-energy 7 μm opcpa and 260 mj ho: Ylf pump laser. Optics Letters, 44(13):3194–3197, 2019.
- [77] Nicolas Thiré, Raman Maksimenka, Bálint Kiss, Clément Ferchaud, Grégory Gitzinger, Thomas Pinoteau, Hervé Jousset, Sebastian Jarosch, Pierre Bizouard, Vittorio Di Pietro, et al. Highly stable, 15 w, few-cycle, 65 mrad cep-noise mid-ir opcpa for statistical physics. Optics Express, 26(21):26907–26915, 2018.
- [78] Matthias Kübel, Sebastian Hell, Apurba Manna, Matthias Awad, Bo Ying, Levente Abrók, Zsolt Diveki, Eric Cormier, Balint Kiss, Jan Böhmer, et al. Few-cycle laser pulse characterization on-target using high-harmonic generation from nano-scale solids (pre-print). 2023.
- [79] Fangjie Zhou, Yi Wu, Alphonse Marra, and Zenghu Chang. Efficient generation of femtosecond millijoule pulses at 3.1 μm . Optics Letters, 47(23):6057–6060, 2022.
- [80] Lorenz von Grafenstein, Martin Bock, Dennis Ueberschaer, Esmerando Escoto, Azize Koç, Kevin Zawilski, Peter Schunemann, Uwe Griebner, and Thomas Elsaesser. Multi-millijoule, few-cycle 5 μm opcpa at 1 khz repetition rate. Optics Letters, 45(21):5998–6001, 2020.

- [81] Pia Fuertjes, Martin Bock, Lorenz von Grafenstein, Dennis Ueberschaer, Uwe Griebner, and Thomas Elsaesser. Few-cycle 65- μ j pulses at 11.4 μ m for ultrafast nonlinear longwave-infrared spectroscopy. *Optica*, 9(11):1303–1306, 2022.
- [82] Kyung-Han Hong, Chien-Jen Lai, Jonathas P Siqueira, Peter Krogen, Jeffrey Moses, Chun-Lin Chang, Gregory J Stein, Luis E Zapata, and Franz X Kärtner. Multi-mj, khz, 2.1 μ m optical parametric chirped-pulse amplifier and high-flux soft x-ray high-harmonic generation. *Optics Letters*, 39(11):3145–3148, 2014.
- [83] Tianli Feng, Anke Heilmann, Martin Bock, Lutz Ehrentraut, Tobias Witting, Haohai Yu, Holger Stiel, Stefan Eisebitt, and Matthias Schnürer. 27 w 2.1 μ m opcpa system for coherent soft x-ray generation operating at 10 khz. *Optics Express*, 28(6):8724–8733, 2020.
- [84] Justinas Pupeikis, P-A Chevreuil, Nicolas Bigler, Lukas Gallmann, Christopher R Phillips, and Ursula Keller. Water window soft x-ray source enabled by a 25 w few-cycle 2.2 μ m opcpa at 100 khz. *Optica*, 7(2):168–171, 2020.
- [85] Yi Wu, Fangjie Zhou, Esben W Larsen, Fengjiang Zhuang, Yanchun Yin, and Zenghu Chang. Generation of few-cycle multi-millijoule 2.5 μ m pulses from a single-stage cr2+: Znse amplifier. *Scientific Reports*, 10(1):7775, 2020.
- [86] Vyacheslav E Leshchenko, Bradford K Talbert, Yu Hang Lai, Sha Li, Yaguo Tang, Stephen J Hageman, Greg Smith, Pierre Agostini, Louis F DiMauro, and Cosmin I Blaga. High-power few-cycle cr: Znse mid-infrared source for attosecond soft x-ray physics. *Optica*, 7(8):981–988, 2020.
- [87] Weichao Yao, Yicheng Wang, Shahwar Ahmed, Martin Hoffmann, Marcel van Delden, Thomas Musch, and Clara J Saraceno. Low-noise, 2-w average power, 112-fs kerr-lens mode-locked ho: Calgo laser at 2.1 μ m. *Optics Letters*, 48(11):2801–2804, 2023.
- [88] Jakub Cajzl, Pavel Peterka, Maciej Kowalczyk, Jan Tarka, Grzegorz Sobon, Jaroslaw Sotor, Jan Aubrecht, Pavel Honzátko, and Ivan Kašík. Thulium-doped silica fibers with enhanced fluorescence lifetime and their application in ultrafast fiber lasers. *Fibers*, 6(3):66, 2018.
- [89] Justin Cook, Alex Sincore, Nicholas Vail, Ria G Krämer, Thorsten A Goebel, Patrick Roumayah, Joshua Bryan, Nathan Bodnar, Stefan Nolte, and Martin Richardson. 100 w, tunable in-band thulium fiber amplifier pumped by incoherently combined 1.9 μ m fiber lasers. *Optics Express*, 31(18):29245–29254, 2023.
- [90] C Gaida, M Gebhardt, T Heuermann, F Stutzki, C Jauregui, and J Limpert. Ultrafast thulium fiber laser system emitting more than 1 kw of average power. *Optics Letters*, 43(23):5853–5856, 2018.
- [91] C Gaida, M Gebhardt, F Stutzki, C Jauregui, J Limpert, and A Tünnermann. Thulium-doped fiber chirped-pulse amplification system with 2 gw of peak power. *Optics Letters*, 41(17):4130–4133, 2016.
- [92] Martin Gebhardt, Tobias Heuermann, Ziyao Wang, Mathias Lenski, Christian Gaida, Robert Klas, Alexander Kirsche, Steffen Hädrich, Jan Rothhardt, and Jens Limpert. Soft x-ray high order harmonic generation driven by high repetition rate ultrafast thulium-doped fiber lasers. In *Fiber Lasers XVII: Technology and Systems*, volume 11260, pages 84–89. SPIE, 2020.

- [93] M Gebhardt, T Heuermann, R Klas, C Liu, A Kirsche, M Lenski, Z Wang, C Gaida, JE Antonio-Lopez, A Schülzgen, et al. Bright, high-repetition-rate water window soft x-ray source enabled by nonlinear pulse self-compression in an antiresonant hollow-core fibre. Light: Science & Applications, 10(1):36, 2021.
- [94] Rüdiger Paschotta. Peak power. Available at https://www.rp-photonics.com/peak_power.html.
- [95] P. A. Franken, A. E. Hill, C. W. Peters, and G. Weinreich. Generation of optical harmonics. Physical Review Letters, 7(4):118, 1961.
- [96] M Bass, PA Franken, AE Hill, CW Peters, and G Weinreich. Optical mixing. Physical Review Letters, 8(1):18, 1962.
- [97] Charles C Wang and George W Racette. Measurement of parametric gain accompanying optical difference frequency generation. Applied Physics Letters, 6(8):169–171, 1965.
- [98] George I Stegeman and Robert A Stegeman. Nonlinear optics: phenomena, materials and devices. John Wiley & Sons, 2012.
- [99] G. P. Agrawal. Nonlinear fiber optics. Academic Press, 2001.
- [100] Gunnar Arisholm. General numerical methods for simulating second-order nonlinear interactions in birefringent media. Journal of the Optical Society of America B, 14(10):2543–2549, Oct 1997.
- [101] R. W. Boyd. Nonlinear optics. Academic Press, 2020.
- [102] A. M. Weiner. Ultrafast optics. John Wiley & Sons, Inc., 2009.
- [103] J. A. Valdmanis, G. Mourou, and C. W. Gabel. Picosecond electro-optic sampling system. Applied Physics Letters, 41(3):211–212, 08 1982.
- [104] Q. Wu and X.-C. Zhang. Free-space electro-optic sampling of terahertz beams. Applied Physics Letters, 67(24):3523–3525, 12 1995.
- [105] Rick Trebino. Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses. Springer Science & Business Media, 2000.
- [106] Harald R Telle, Guenter Steinmeyer, Amy E Dunlop, Joern Stenger, Dirk H Sutter, and Ursula Keller. Carrier-envelope offset phase control: A novel concept for absolute optical frequency measurement and ultrashort pulse generation. Applied Physics B, 69:327–332, 1999.
- [107] 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. Nature Communications, 5(1):3331, 2014.
- [108] Joseph W Goodman. Introduction to Fourier optics. Roberts and Company publishers, 2005.
- [109] Guan Gui, Amitava Adak, Manika Dandapat, Daniel Carlson, Drew Morrill, Alexander Guggenmos, Henry Kapteyn, Margaret Murnane, Vladimir Pervak, and Chen-Ting Liao. Measurement and control of optical nonlinearities in dispersive dielectric multilayers. Optics Express, 29(4):4947–4957, 2021.

- [110] John David Jackson. Classical electrodynamics. John Wiley and Sons, Inc., 1999.
- [111] Ayman Hoblos. Electro-optical photonic crystal devices for the modulation and detection of weak electric fields. PhD thesis, Université Bourgogne Franche-Comté, 02 2021.
- [112] Daniele Faccio, John Dudley, and Matteo Clerici. Frontiers in modern optics, volume 190. IOS Press, 2016.
- [113] JA Armstrong, N Bloembergen, J Ducuing, and Peter S Pershan. Interactions between light waves in a nonlinear dielectric. Physical Review, 127(6):1918, 1962.
- [114] Cristian Manzoni and Giulio Cerullo. Design criteria for ultrafast optical parametric amplifiers. Journal of Optics, 18(10):103501, 2016.
- [115] Zhensong Cao, Xiaoming Gao, Weidong Chen, Huan Wang, Weijun Zhang, and Zhiben Gong. Study of quasi-phase matching wavelength acceptance bandwidth for periodically poled linbo3 crystal-based difference-frequency generation. Optics and Lasers in Engineering, 47(5):589–593, 2009.
- [116] M. Nieto-Vesperinas and G. Lera. Solution to non-linear optical mixing equations with depletion and diffraction: difference-frequency generation. Optics Communications, 69(3):329–333, 1989.
- [117] O Gayer, Z Sacks, E Galun, and A Arie. Temperature and wavelength dependent refractive index equations for mgo-doped congruent and stoichiometric linbo 3. Applied Physics B, 91:343–348, 2008.
- [118] Arlee V Smith. Crystal nonlinear optics: with SNLO examples. AS-Photonics Albuquerque, NM, USA, 2018.
- [119] Paul J Wegner, Jerome M Auerbach, Charles E Barker, Scott C Burkhart, Scott A Couture, Jim J DeYoreo, RL Hibbard, Lisa W Liou, Mary A Norton, Pamela K Whitman, et al. Frequency converter development for the national ignition facility. In Third International Conference on Solid State Lasers for Application to Inertial Confinement Fusion, volume 3492, pages 392–405. SPIE, 1999.
- [120] Rogers H Stolen and Chinlon Lin. Self-phase-modulation in silica optical fibers. Physical Review A, 17(4):1448, 1978.
- [121] Rüdiger Paschotta. Self-phase modulation causes spectral broadening – does it really? The Photonics Spotlight, July 2015. Available at https://www.rp-photonics.com/spotlight_2015_07_01.html.
- [122] D Morrill, D Carlson, D Lesko, T-H Wu, S Diddams, M Murnane, H Kapteyn, and M Hemmer. All-fiber mid-ir opcpa front-end and cryogenic yb: Yag pump laser for soft x-ray generation. In Advanced Solid State Lasers, pages AW2A–7. Optica Publishing Group, 2021.
- [123] Visual inspection and cleaning of connectors. Available at <https://www.thefoa.org/tech/ref/testing/test/scope.html>.
- [124] 980nm band panda fiber. Available at https://www.fujikura.co.jp/eng/products/optical/opticalfibers/04/2050105_12895.html.

- [125] Core pumping and cladding pumping. Available at <https://www.fiberlabs.com/glossary/core-pumping-and-cladding-pumping/>.
- [126] E. Fred Schubert. *Optical Communication*, page 367–381. Cambridge University Press, 2 edition, 2006.
- [127] Nils Raabe, Tianli Feng, Mark Mero, Haochen Tian, Youjian Song, Wolfgang Hänsel, Ronald Holzwarth, Alexander Sell, Armin Zach, and Günter Steinmeyer. Excess carrier-envelope phase noise generation in saturable absorbers. *Optics Letters*, 42(6):1068–1071, 2017.
- [128] Menlo Systems. *Why 9 Cuts a Good Figure: Technology Primer*.
- [129] Wolfgang Hänsel, Heinar Hoogland, Michele Giunta, Sebastian Schmid, Tilo Steinmetz, Ralf Doubek, Peter Mayer, Sven Dobner, Carsten Cleff, Marc Fischer, et al. All polarization-maintaining fiber laser architecture for robust femtosecond pulse generation. *Exploring the World with the Laser: Dedicated to Theodor Hänsch on his 75th birthday*, pages 331–340, 2018.
- [130] Andy Chong, Joel Buckley, Will Renninger, and Frank Wise. All-normal-dispersion femtosecond fiber laser. *Optics Express*, 14(21):10095–10100, 2006.
- [131] Florian Tauser, Alfred Leitenstorfer, and Wolfgang Zinth. Amplified femtosecond pulses from an er: fiber system: Nonlinear pulse shortening and self-referencing detection of the carrier-envelope phase evolution. *Optics Express*, 11(6):594–600, Mar 2003.
- [132] JW Nicholson and MF Yan. Cross-coherence measurements of supercontinua generated in highly-nonlinear, dispersion shifted fiber at 1550 nm. *Optics Express*, 12(4):679–688, 2004.
- [133] Kazunori Suzuki, Yasuo Kimura, and Masataka Nakazawa. Subpicosecond soliton amplification and transmission using er³⁺-doped fibers pumped by ingaasp laser diodes. *Optics Letters*, 14(16):865–867, 1989.
- [134] BJ Ainslie, KJ Blow, AS Gouveia-Neto, PGJ Wigley, ASB Sombra, and JR Taylor. Femtosecond soliton amplification in erbium doped silica fibre. *Electronics Letters*, 3(26):186–188, 1990.
- [135] I Yu Khrushchev, AB Grudinin, EM Dianov, DV Korobkin, VA Semenov, and AM Prokhorov. Amplification of femtosecond pulses in Er/sup³⁺-doped single-mode optical fibres. *Electronics Letters*, 7(26):456–458, 1990.
- [136] J. W. Nicholson, A. D. Yablon, P. S. Westbrook, K. S. Feder, and M. F. Yan. High power, single mode, all-fiber source of femtosecond pulses at 1550 nm and its use in supercontinuum generation. *Optics Express*, 12(13):3025–3034, Jun 2004.
- [137] Günther Krauss, Sebastian Lohss, Tobias Hanke, Alexander Sell, Stefan Eggert, Rupert Huber, and Alfred Leitenstorfer. Synthesis of a single cycle of light with compact erbium-doped fibre technology. *Nature Photonics*, 4(1):33–36, 2010.
- [138] Daniel MB Lesko, Henry Timmers, Sida Xing, Abijith Kowligy, Alexander J Lind, and Scott A Diddams. A six-octave optical frequency comb from a scalable few-cycle erbium fibre laser. *Nature Photonics*, 15(4):281–286, 2021.

- [139] M. E. Fermann, V. I. Kruglov, B. C. Thomsen, J. M. Dudley, and J. D. Harvey. Self-similar propagation and amplification of parabolic pulses in optical fibers. Physical Review Letters, 84:6010–6013, Jun 2000.
- [140] Ruoyu Sun, Dongchen Jin, Fangzhou Tan, Shouyu Wei, Chang Hong, Jia Xu, Jiang Liu, and Pu Wang. High-power all-fiber femtosecond chirped pulse amplification based on dispersive wave and chirped-volume Bragg grating. Optics Express, 24(20):22806–22812, 2016.
- [141] JW Nicholson, MF Yan, P Wisk, J Fleming, F DiMarcello, E Monberg, A Yablon, C Jørgensen, and T Veng. All-fiber, octave-spanning supercontinuum. Optics Letters, 28(8):643–645, 2003.
- [142] M. Erkintalo, Y. Q. Xu, S. G. Murdoch, J. M. Dudley, and G. Genty. Cascaded phase matching and nonlinear symmetry breaking in fiber frequency combs. Physical Review Letters, 109:223904, Nov 2012.
- [143] Dane R. Austin, C. Martijn de Sterke, Benjamin J. Eggleton, and Thomas G. Brown. Dispersive wave blue-shift in supercontinuum generation. Optics Express, 14(25):11997–12007, Dec 2006.
- [144] K. E. Webb, Y. Q. Xu, M. Erkintalo, and S. G. Murdoch. Generalized dispersive wave emission in nonlinear fiber optics. Optics Letters, 38(2):151–153, Jan 2013.
- [145] Eli Yablonovitch, C. Flytzanis, and N. Bloembergen. Anisotropic interference of three-wave and double two-wave frequency mixing in GaAs. Physical Review Letters, 29:865–868, Sep 1972.
- [146] Govind P. Agrawal. Chapter 12 - novel nonlinear phenomena. In Govind P. Agrawal, editor, Nonlinear Fiber Optics (Sixth Edition), pages 503–556. Academic Press, sixth edition edition, 2019.
- [147] L. Misoguti, S. Backus, C. G. Durfee, R. Bartels, M. M. Murnane, and H. C. Kapteyn. Generation of broadband vuv light using third-order cascaded processes. Physical Review Letters, 87:013601, Jun 2001.
- [148] pynlo: Nonlinear optics modeling for python. Available at <https://github.com/pyNLO/PyNLO>.
- [149] Lorenz von Grafenstein, Martin Bock, and Uwe Griebner. Bifurcation analysis in high repetition rate regenerative amplifiers. IEEE Journal of Selected Topics in Quantum Electronics, 24(5):1–13, 2018.
- [150] W Hettel, D Morrill, D Carlson, D Lesko, T-H Wu, S Diddams, H Kapteyn, M Murnane, and M Hemmer. Dispersion compensation in a 3 μm wavelength opcpa system by shaping the 1.5 μm signal input. In Mid-Infrared Coherent Sources, pages MW3C–5. Optica Publishing Group, 2022.
- [151] D Carlson, W Hettel, D Morrill, D Lesko, T-H Wu, F Yu, S Diddams, H Kapteyn, M Murnane, and M Hemmer. 3 μm wavelength opcpa for soft x-ray generation. In 2022 Conference on Lasers and Electro-Optics (CLEO), pages 1–2. IEEE, 2022.

- [152] Drew Morrill, Daniel Carlson, Will Hettel, Daniel Lesko, Tsung-Han Wu, Scott Diddams, Margaret Murnane, Henry Kapteyn, and Michael Hemmer. A cryogenic yb: Yag regenerative amplifier pumped, mj-class, 1 khz, 3 μm wavelength ultrafast opcpa. In Advanced Solid State Lasers, pages AM5A–5. Optica Publishing Group, 2022.
- [153] Drew Morrill, Daniel Carlson, Will Hettel, Daniel Lesko, Tsung-Han Wu, Scott Diddams, Margaret Murnane, Henry Kapteyn, and Michael Hemmer. 20-mj class, 1 khz cryogenic yb: Yag regenerative amplifier for mid-ir opcpa pumping. In Frontiers in Optics, pages JTU4A–18. Optica Publishing Group, 2022.
- [154] Michael Gerrity. Development of a High Energy, kHz, Mid-Infrared OPCPA Laser for keV High Harmonic Generation. PhD thesis, University of Colorado, 2015.
- [155] Susannah Rachel Wang. High Peak and Average Power Mid-Infrared Laser for High Harmonic Generation of Soft X-rays. PhD thesis, University of Colorado at Boulder, 2016.
- [156] Susannah Wang, Michael Gerrity, Sterling Backus, Margaret M. Murnane, Henry C. Kapteyn, and Seth L. Cousin. Multi-mj, 1khz, 3.1 μm opcpa. In Conference on Lasers and Electro-Optics, page STh1L.5. Optica Publishing Group, 2017.
- [157] Ming-Chang Chen, Christopher Mancuso, Carlos Hernández-García, Franklin Dollar, Ben Galloway, Dimitar Popmintchev, Pei-Chi Huang, Barry Walker, Luis Plaja, Agnieszka A Jaroń-Becker, et al. 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.
- [158] Benjamin Galloway. High-Order Harmonic Generation Driven by Mid-Infrared Laser Light. PhD thesis, University of Colorado at Boulder, 2016.
- [159] Yunpei Deng, Alexander Schwarz, Hanieh Fattahi, Moritz Ueffing, Xun Gu, Marcus Ossian-der, Thomas Metzger, Volodymyr Pervak, Hideki Ishizuki, Takunori Taira, et al. Carrier-envelope-phase-stable, 1.2 mj, 1.5 cycle laser pulses at 2.1 μm . Optics Letters, 37(23):4973–4975, 2012.
- [160] Xun Gu, Gilad Marcus, Yunpei Deng, Thomas Metzger, Catherine Teisset, Nobuhisa Ishii, Takao Fuji, Andrius Baltuska, Rytis Butkus, Volodymyr Pervak, et al. Generation of carrier-envelope-phase-stable 2-cycle 740- μj pulses at 2.1- μm carrier wavelength. Optics Express, 17(1):62–69, 2009.
- [161] Nicolas Bigler, Justinas Pupeikis, Stefan Hrisafov, Lukas Gallmann, Christopher R Phillips, and Ursula Keller. High-power opcpa generating 1.7 cycle pulses at 2.5 μm . Optics Express, 26(20):26750–26757, 2018.
- [162] Matthias Baudisch, H Pires, H Ishizuki, T Taira, M Hemmer, and Jens Biegert. Sub-4-optical-cycle, 340 mw peak power, high stability mid-ir source at 160 khz. Journal of Optics, 17(9):094002, 2015.
- [163] M Baudisch, M Hemmer, H Pires, and Jens Biegert. Performance of mgo: Ppln, kta, and knbo 3 for mid-wave infrared broadband parametric amplification at high average power. Optics Letters, 39(20):5802–5805, 2014.

- [164] Benedikt W Mayer, Christopher R Phillips, Lukas Gallmann, and Ursula Keller. Mid-infrared pulse generation via achromatic quasi-phase-matched opcpa. Optics Express, 22(17):20798–20808, 2014.
- [165] Philippe Rigaud, Aymeric Van de Walle, Marc Hanna, Nicolas Forget, Florent Guichard, Yoann Zaouter, Khmaies Guesmi, Frédéric Druon, and Patrick Georges. Supercontinuum-seeded few-cycle mid-infrared opcpa system. Optics Express, 24(23):26494–26502, 2016.
- [166] Judith R Schwesyg, Matthias Falk, Chris R Phillips, Dieter H Jundt, Karsten Buse, and Martin M Fejer. Pyroelectrically induced photorefractive damage in magnesium-doped lithium niobate crystals. Journal of the Optical Society of America B, 28(8):1973–1987, 2011.
- [167] M Baudisch, B Wolter, M Pullen, M Hemmer, and Jens Biegert. High power multi-color opcpa source with simultaneous femtosecond deep-uv to mid-ir outputs. Optics Letters, 41(15):3583–3586, 2016.
- [168] 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. Optics Letters, 35(11):1752–1754, 2010.
- [169] M Hemmer, D Sanchez, M Jelínek, H Jelínková, V Kubeček, and J Biegert. Fiber-seeded, 10-ps, 2050-nm, multi-mj, cryogenic ho: Ylf cpa. In 2014 Conference on Lasers and Electro-Optics (CLEO)-Laser Science to Photonic Applications, pages 1–2. IEEE, 2014.
- [170] Lorenz von Grafenstein, Martin Bock, Uwe Griebner, and Thomas Elsaesser. High-energy multi-kilohertz ho-doped regenerative amplifiers around 2 μm . Optics Express, 23(11):14744–14752, 2015.
- [171] Luis E Zapata, F Reichert, M Hemmer, and FX Kärtner. 250 w average power, 100 khz repetition rate cryogenic yb: Yag amplifier for opcpa pumping. Optics Letters, 41(3):492–495, 2016.
- [172] Umit Demirbas, Huseyin Cankaya, Yi Hua, Jelto Thesinga, Mikhail Pergament, and Franz X Kärtner. 20-mj, sub-ps pulses at up to 70 w average power from a cryogenic yb: Ylf regenerative amplifier. Optics Express, 28(2):2466–2479, 2020.
- [173] Theodor Seuss Geisel. The Cat in the Hat. Random House, 1957.
- [174] Tso Yee Fan, Daniel J Ripin, Roshan L Aggarwal, Juan R Ochoa, Bien Chann, Michael Tilleman, and Joshua Spitzberg. Cryogenic yb³⁺-doped solid-state lasers. IEEE Journal of selected topics in Quantum Electronics, 13(3):448–459, 2007.
- [175] Randy Bartels, S Backus, E Zeek, L Misoguti, G Vdovin, IP Christov, MM Murnane, and HC Kapteyn. Shaped-pulse optimization of coherent emission of high-harmonic soft x-rays. Nature, 406(6792):164–166, 2000.
- [176] Peter Fischer, Alexander Muschet, Tino Lang, Roushdey Salh, and Laszlo Veisz. Saturation control of an optical parametric chirped-pulse amplifier. Optics Express, 29(3):4210–4218, 2021.

- [177] S Bucht, D Haberberger, J Bromage, and DH Froula. Methodology for designing grism stretchers for idler-based optical parametric chirped-pulse-amplification systems. Journal of the Optical Society of America B, 36(8):2325–2337, 2019.
- [178] P Tournois. New diffraction grating pair with very linear dispersion for laser pulse compression. Electronics Letters, 16(29):1414–1415, 1993.
- [179] Steve Kane and Jeff Squier. Grating compensation of third-order material dispersion in the normal dispersion regime: sub-100-fs chirped-pulse amplification using a fiber stretcher and grating-pair compressor. IEEE Journal of Quantum Electronics, 31(11):2052–2057, 1995.
- [180] Edmond Treacy. Optical pulse compression with diffraction gratings. IEEE Journal of Quantum Electronics, 5(9):454–458, 1969.
- [181] Frédéric Druon, Marc Hanna, Gaëlle Lucas-Leclin, Yoann Zaouter, Dimitris Papadopoulos, and Patrick Georges. Simple and general method to calculate the dispersion properties of complex and aberrated stretchers-compressors. Journal of the Optical Society of America B, 25(5):754–762, 2008.
- [182] Daniel Carlson, Michael Tanksalvala, Drew Morrill, Julio San Roman, E Conejero Jarque, Henry C Kapteyn, Margaret M Murnane, and Michaël Hemmer. Nonlinear post-compression in multi-pass cells in the mid-ir region using bulk materials. Optics Letters, 47(20):5289–5292, 2022.
- [183] Mate Kurucz, Roland Flender, Ludovit Haizer, Roland S Nagymihaly, Wosik Cho, Kyung T Kim, Szabolcs Toth, Eric Cormier, and Balint Kiss. 2.3-cycle mid-infrared pulses from hybrid thin-plate post-compression at 7 w average power. Optics Communications, 472:126035, 2020.
- [184] Michaël Hemmer, Matthias Baudisch, Alexandre Thai, Arnaud Couairon, and Jens Biegert. Self-compression to sub-3-cycle duration of mid-infrared optical pulses in dielectrics. Optics Express, 21(23):28095–28102, 2013.
- [185] Valentina Shumakova, Pavel Malevich, S Ališauskas, Alexander Voronin, AM Zheltikov, D Faccio, Daniil Kartashov, A Baltuška, and A Pugžlys. Multi-millijoule few-cycle mid-infrared pulses through nonlinear self-compression in bulk. Nature Communications, 7(1):12877, 2016.
- [186] Minyuan Hu, Xiaoyan Liang, Baozhen Zhao, Ruxin Li, and Zhizhan Xu. Broad-bandwidth semi-noncollinear optical parametric amplification in periodically poled linbo3 based on tilted quasi-phase-matched gratings. Japanese Journal of Applied Physics, 46(8R):5148, 2007.
- [187] Maxim V Ammosov, Nikolai B Delone, and Vladimir P Krainov. Tunnel ionization of complex atoms and of atomic ions in an alternating electromagnetic field. Soviet Journal of Experimental and Theoretical Physics, 64(6):1191, 1986.
- [188] Peter C Johnsen, Sinéad A Ryan, Christian Gentry, Anya Grafov, Henry Kapteyn, and Margaret Murnane. A beamline for ultrafast extreme ultraviolet magneto-optical spectroscopy in reflection near the shot noise limit. Review of Scientific Instruments, 94(3), 2023.
- [189] MJ Guardalben, J Keegan, LJ Waxer, V Bagnoud, IA Begishev, J Puth, and JD Zuegel. Design of a highly stable, high-conversion-efficiency, optical parametric chirped-pulse amplification system with good beam quality. Optics Express, 11(20):2511–2524, 2003.

- [190] SK Zhang, M Fujita, M Yamanaka, M Nakatsuka, Y Izawa, and C Yamanaka. Study of the stability of optical parametric amplification. Optics Communications, 184(5-6):451–455, 2000.
- [191] Nathan Brooks. Structured High-Harmonic Light Sources for Enhanced Extreme Ultraviolet Microscopy. PhD thesis, University of Colorado at Boulder, 2022.
- [192] J. Mauritsson, P. Johnsson, E. Gustafsson, A. L’Huillier, K. J. Schafer, and M. B. Gaarde. Attosecond pulse trains generated using two color laser fields. Physical Review Letters, 97:013001, Jul 2006.
- [193] C. Chen, C. Hernández-García, Z. Tao, W. You, Y. Zhang, D. Zusin, C. Gentry, P. Tengdin, A. Becker, A. Jaron-Becker, H. Kapteyn, and M. Murnane. Influence of microscopic and macroscopic effects on attosecond pulse generation using two-color laser fields. Optics Express, 25(23):28684–28696, Nov 2017.
- [194] L. Le Déroff, P. Salières, and B. Carré. Beam-quality measurement of a focused high-order harmonic beam. Optics Letters, 23(19):1544–1546, Oct 1998.
- [195] T. Auguste, P. Salières, A. S. Wyatt, A. Monmayrant, I. A. Walmsley, E. Cormier, A. Zaïr, M. Holler, A. Guandalini, F. Schapper, J. Biegert, L. Gallmann, and U. Keller. Theoretical and experimental analysis of quantum path interferences in high-order harmonic generation. Physical Review A, 80:033817, Sep 2009.
- [196] Maciej Lewenstein, Ph Balcou, M Yu Ivanov, Anne L’huillier, and Paul B Corkum. Theory of high-harmonic generation by low-frequency laser fields. Physical Review A, 49(3):2117, 1994.
- [197] A Willner, F Tavella, Mark Yeung, Thomas Dzelzainis, C Kamperidis, M Bakarezos, D Adams, R Riedel, M Schulz, MC Hoffmann, et al. Efficient control of quantum paths via dual-gas high harmonic generation. New Journal of Physics, 13(11):113001, 2011.
- [198] Xiaoshi Zhang, Amy L Lytle, Tenio Popmintchev, Xibin Zhou, Henry C Kapteyn, Margaret M Murnane, and Oren Cohen. Quasi-phase-matching and quantum-path control of high-harmonic generation using counterpropagating light. Nature Physics, 3(4):270–275, 2007.
- [199] Oren Cohen, Xiaoshi Zhang, Amy L Lytle, Tenio Popmintchev, Margaret M Murnane, and Henry C Kapteyn. Grating-assisted phase matching in extreme nonlinear optics. Physical Review Letters, 99(5):053902, 2007.
- [200] Ba Khuong Dinh. Phase-Matched High Order Harmonic Generation and Applications. PhD thesis, Swinburne University of Technology, 2012.
- [201] Edson R Peck and Donald J Fisher. Dispersion of argon. Journal of the Optical Society of America, 54(11):1362–1364, 1964.
- [202] Charles R Mansfield and Edson R Peck. Dispersion of helium. Journal of the Optical Society of America, 59(2):199–204, 1969.

- [203] Charles G. Durfee, Andy R. Rundquist, Sterling Backus, Catherine Herne, Margaret M. Murnane, and Henry C. Kapteyn. Phase matching of high-order harmonics in hollow waveguides. Physical Review Letters, 83:2187–2190, Sep 1999.
- [204] T. Auguste, B. Carré, and P. Salières. Quasi-phase-matching of high-order harmonics using a modulated atomic density. Physical Review A, 76:011802, Jul 2007.
- [205] Philippe Balcou, Pascal Salieres, Anne L’Huillier, and Maciej Lewenstein. Generalized phase-matching conditions for high harmonics: The role of field-gradient forces. Physical Review A, 55(4):3204, 1997.
- [206] R Weissenbilder, S Carlström, L Rego, C Guo, CM Heyl, P Smorenburg, Eric Constant, CL Arnold, and A L’huillier. How to optimize high-order harmonic generation in gases. Nature Reviews Physics, 4(11):713–722, 2022.
- [207] Seth Lucien Cousin. Towards the generation of isolated attosecond pulses in the water window. PhD thesis, Universitat Politècnica de Catalunya, 2016.
- [208] Antoine Comby, S Beaulieu, Eric Constant, D Descamps, S Petit, and Y Mairesse. Absolute gas density profiling in high-order harmonic generation. Optics Express, 26(5):6001–6009, 2018.
- [209] F Frassetto, S Coraggia, L Poletto, N Guerassimova, S Dziarzhytski, E Ploenjes, and H Weigelt. Compact spectrometer for the analysis of high harmonics content of extreme-ultraviolet free-electron-laser radiation. In Advances in X-Ray/EUV Optics and Components V, volume 7802, pages 66–73. SPIE, 2010.
- [210] Brian K McFarland, JP Farrell, PH Bucksbaum, and M Gühr. High-order harmonic phase in molecular nitrogen. Physical Review A, 80(3):033412, 2009.
- [211] Dylan C Yost, Thomas R Schibli, Jun Ye, Jennifer L Tate, James Hostetter, Mette B Gaarde, and Kenneth J Schafer. Vacuum-ultraviolet frequency combs from below-threshold harmonics. Nature Physics, 5(11):815–820, 2009.
- [212] Fernando Brizuela, CM Heyl, Piotr Rudawski, David Kroon, Linnea Rading, J Marcus Dahlström, Johan Mauritsson, Per Johnsson, CL Arnold, and Anne L’Huillier. Efficient high-order harmonic generation boosted by below-threshold harmonics. Scientific Reports, 3(1):1410, 2013.
- [213] Tobias Kroh, Cheng Jin, Peter Krogen, Philip D Keathley, Anne-Laure Calendron, Jonathas P Siqueira, Houkun Liang, Edilson L Falcão-Filho, CD Lin, Franz X Kärtner, et al. Enhanced high-harmonic generation up to the soft x-ray region driven by mid-infrared pulses mixed with their third harmonic. Optics Express, 26(13):16955–16969, 2018.
- [214] Andrew Spott, Agnieszka Jaroń Becker, and Andreas Becker. Ab initio and perturbative calculations of the electric susceptibility of atomic hydrogen. Physical Review A, 90:013426, Jul 2014.
- [215] A Spott, A Becker, and A Jaroń-Becker. Transition from perturbative to nonperturbative interaction in low-order-harmonic generation. Physical Review A, 91(2):023402, 2015.

- [216] Miroslav Kolesik and Jerome V Moloney. Modeling and simulation techniques in extreme nonlinear optics of gaseous and condensed media. Reports on Progress in Physics, 77(1):016401, 2013.
- [217] Armando Valter Felicio Zuffi, Nilson Dias Vieira Junior, and Ricardo Elgul Samad. Below-threshold-harmonics-generation limitation due to laser-induced ionization in noble gases. Physical Review A, 105(2):023112, 2022.
- [218] Peng-Cheng Li, Yae-Lin Sheu, Cecil Laughlin, and Shih-I Chu. Dynamical origin of near- and below-threshold harmonic generation of cs in an intense mid-infrared laser field. Nature Communications, 6(1):7178, 2015.
- [219] Zhi-Bin Wang, Yang-Yang Chen, Tong-Gang Jia, Zhi-Hong Jiao, and Peng-Cheng Li. Carrier-envelope-phase-dependent below-threshold harmonic generation in few-cycle mid-infrared laser fields. Optics Express, 31(2):1567–1582, 2023.
- [220] D Kartashov, S Ališauskas, Audrius Pugzlys, Aleksandr A Voronin, Aleksei M Zheltikov, and A Baltuška. Third- and fifth-harmonic generation by mid-infrared ultrashort pulses: beyond the fifth-order nonlinearity. Optics Letters, 37(12):2268–2270, 2012.
- [221] V. Loriot, E. Hertz, O. Faucher, and B. Lavorel. Measurement of high order kerr refractive index of major air components. Optics Express, 17(16):13429–13434, Aug 2009.
- [222] P. B ejot, J. Kasparian, S. Henin, V. Loriot, T. Vieillard, E. Hertz, O. Faucher, B. Lavorel, and J.-P. Wolf. Higher-order kerr terms allow ionization-free filamentation in gases. Physical Review Letters, 104:103903, Mar 2010.
- [223] L.E. Murr and W.R. Bitler. Investigation of optical absorption and the structure of thin metal foils. Materials Research Bulletin, 2(8):787–798, 1967.
- [224] Armando Giannattasio, Ian R. Hooper, and William L. Barnes. Transmission of light through thin silver films via surface plasmon-polaritons. Optics Express, 12(24):5881–5886, Nov 2004.
- [225] Christian Rosenberg Petersen, Uffe M oller, Irnis Kubat, Binbin Zhou, Sune Dupont, Jacob Ramsay, Trevor Benson, Slawomir Sujecki, Nabil Abdel-Moneim, Zhuoqi Tang, et al. Mid-infrared supercontinuum covering the 1.4–13.3 μm molecular fingerprint region using ultra-high na chalcogenide step-index fibre. Nature Photonics, 8(11):830–834, 2014.
- [226] Jinwei Zhang, Qing Wang, Jingjie Hao, Heyan Liu, Jiyong Yao, Zhuang Li, Jie Liu, and Ka Fai Mak. Broadband, few-cycle mid-infrared continuum based on the intra-pulse difference frequency generation with bgse crystals. Optics Express, 28(25):37903–37909, 2020.
- [227] Andre Peremans, Dan Lis, Francesca Cecchet, Peter G Schunemann, Kevin T Zawilski, and Valentin Petrov. Noncritical singly resonant synchronously pumped opo for generation of picosecond pulses in the mid-infrared near 6.4 μm . Optics Letters, 34(20):3053–3055, 2009.
- [228] Enrique AJ Marcatili and Robert A Schmeltzer. Hollow metallic and dielectric waveguides for long distance optical transmission and lasers. Bell System Technical Journal, 43(4):1783–1809, 1964.
- [229] NM Litchinitser, AK Abeeluck, C Headley, and BJ Eggleton. Antiresonant reflecting photonic crystal optical waveguides. Optics Letters, 27(18):1592–1594, 2002.

- [230] OH Heckl, CRE Baer, Christian Kränkel, SV Marchese, F Schapper, M Holler, T Südmeyer, JS Robinson, JWG Tisch, F Couny, et al. High harmonic generation in a gas-filled hollow-core photonic crystal fiber. Applied Physics B, 97:369–373, 2009.
- [231] Piotr Jaworski, Paweł Koziół, Karol Krzempek, Dakun Wu, Fei Yu, Piotr Bojes, Grzegorz Dudzik, Meisong Liao, Krzysztof Abramski, and Jonathan Knight. Antiresonant hollow-core fiber-based dual gas sensor for detection of methane and carbon dioxide in the near-and mid-infrared regions. Sensors, 20(14):3813, 2020.
- [232] Young-Gyun Jeong, Riccardo Piccoli, Andrea Rovere, Luca Zanotto, Gabriel Tempea, Derrek Wilson, Maksym Ivanov, Alicia Ramirez, Roberto Morandotti, François Légaré, et al. Guiding of laser pulses at the theoretical limit—97% throughput hollow-core fibers. In Advanced Solid State Lasers, pages AW2A–5. Optica Publishing Group, 2021.
- [233] Tetsuo Harada, Nobukazu Teranishi, Takeo Watanabe, Quan Zhou, Jan Bogaerts, and Xinyang Wang. High-exposure-durability, high-quantum-efficiency ($> 90\%$) backside-illuminated soft-x-ray cmos sensor. Applied Physics Express, 13(1):016502, 2019.
- [234] H Turcicova, O Novak, J Muzik, D Stepankova, M Smrz, and T Mocek. Laser induced damage threshold (lidt) of β -barium borate (bbo) and cesium lithium borate (clbo)—overview. Optics & Laser Technology, 149:107876, 2022.
- [235] Richard DeSalvo, Ali A Said, David J Hagan, Eric W Van Stryland, and Mansoor Sheik-Bahae. Infrared to ultraviolet measurements of two-photon absorption and $n/2$ wide bandgap solids. IEEE Journal of Quantum Electronics, 32(8):1324–1333, 1996.
- [236] Daniele Brida, Günther Krauss, Alexander Sell, and Alfred Leitenstorfer. Ultrabroadband er: fiber lasers. Laser & Photonics Reviews, 8(3):409–428, 2014.
- [237] Yuan Hung Lo, Jihan Zhou, Arjun Rana, Drew Morrill, Christian Gentry, Bjoern Enders, Young-Sang Yu, Chang-Yu Sun, David A Shapiro, Roger W Falcone, et al. X-ray linear dichroic ptychography. Proceedings of the National Academy of Sciences, 118(3):e2019068118, 2021.
- [238] Nist standard reference database 69: Nist chemistry webbook.
- [239] Cheng-Tien Chiang, Michael Huth, Andreas Trützschler, Mario Kiel, Frank O Schumann, Jürgen Kirschner, and Wolf Widdra. Boosting laboratory photoelectron spectroscopy by megahertz high-order harmonics. New Journal of Physics, 17(1):013035, 2015.
- [240] Christian Erny, Clemens Heese, M Haag, Lukas Gallmann, and Ursula Keller. High-repetition-rate optical parametric chirped-pulse amplifier producing $1\text{-}\mu\text{j}$, sub-100-fs pulses in the mid-infrared. Optics Express, 17(3):1340–1345, 2009.
- [241] Shima Gholam-Mirzaei, John Beetar, and Michael Chini. High harmonic generation in zno with a high-power mid-ir opa. Applied Physics Letters, 110(6), 2017.
- [242] Ian N Ross, Pavel Matousek, Geoffrey HC New, and Karoly Osvay. Analysis and optimization of optical parametric chirped pulse amplification. Journal of the Optical Society of America B, 19(12):2945–2956, 2002.

- [243] Xiao Zou, Wenkai Li, Houkun Liang, Kun Liu, Shizhen Qu, Qi Jie Wang, and Ying Zhang. 300 μJ , 3 ns, few-cycle, 3 μm opcpa based on periodically poled stoichiometric lithium tantalate crystals. Optics Letters, 44(11):2791–2794, 2019.
- [244] Yuxi Fu, Bing Xue, Katsumi Midorikawa, and Eiji J Takahashi. Tw-scale mid-infrared pulses near 3.3 μm directly generated by dual-chirped optical parametric amplification. Applied Physics Letters, 112(24), 2018.
- [245] Kun Zhao, Haizhe Zhong, Peng Yuan, Guoqiang Xie, Jing Wang, Jingui Ma, and Liejia Qian. Generation of 120 gw mid-infrared pulses from a widely tunable noncollinear optical parametric amplifier. Optics Letters, 38(13):2159–2161, 2013.
- [246] Mark Mero, Zsuzsanna Heiner, Valentin Petrov, Horst Rottke, Federico Branchi, Gabrielle M Thomas, and Marc JJ Vrakking. 43 ns, 1.55 μm and 12.5 ns, 3.1 μm dual-beam, sub-10 cycle, 100 khz optical parametric chirped pulse amplifier. Optics Letters, 43(21):5246–5249, 2018.
- [247] Manuel Bridger, OA Naranjo-Montoya, Alexander Tarasevitch, and Uwe Bovensiepen. Towards high power broad-band opcpa at 3000 nm. Optics Express, 27(22):31330–31337, 2019.
- [248] Xiao Zou, Wenkai Li, Shizhen Qu, Kun Liu, Hao Li, Qi Jie Wang, Ying Zhang, and Houkun Liang. Flat-top pumped multi-millijoule mid-infrared parametric chirped-pulse amplifier at 10 khz repetition rate. Laser & Photonics Reviews, 15(6):2000292, 2021.
- [249] Junji Hirohashi. Characterization of domain switching and optical damage properties in ferroelectrics. PhD thesis, Royal Institute of Technology, Stockholm, Sweden, 2006.
- [250] Mauro Nisoli, Sandro De Silvestri, and Orazio Svelto. Generation of high energy 10 fs pulses by a new pulse compression technique. Applied Physics Letters, 68(20):2793–2795, 1996.
- [251] An-Chun Tien, Sterling Backus, Henry Kapteyn, Margaret Murnane, and Gérard Mourou. Short-pulse laser damage in transparent materials as a function of pulse duration. Physical Review Letters, 82(19):3883, 1999.
- [252] I Dolev, A Ganany-Padowicz, O Gayer, A Arie, J Mangin, and Grégory Gadret. Linear and nonlinear optical properties of mgo: Litao 3. Applied Physics B, 96:423–432, 2009.
- [253] Kiyoshi Kato and Eiko Takaoka. Sellmeier and thermo-optic dispersion formulas for ktp. Applied Optics, 41(24):5040–5044, 2002.
- [254] M Katz, D Eger, MB Oron, and A Hardy. Refractive dispersion curve measurement of ktiopo 4 using periodically segmented waveguides and periodically poled crystals. Journal of Applied Physics, 90(1):53–58, 2001.
- [255] Joana Alves, Hugo Pires, Celso P João, and Gonçalo Figueira. Multi-mj scaling of 5-optical cycle, 3 μm opcpa. In Photonics, volume 8, page 503. MDPI, 2021.
- [256] BC Stuart, MD Feit, AM Rubenchik, BW Shore, and MD Perry. Laser-induced damage in dielectrics with nanosecond to subpicosecond pulses. Physical Review Letters, 74(12):2248, 1995.
- [257] S Manjooran, H Zhao, IT Lima, and A Major. Phase-matching properties of ppktp, mgo: Ppslt and mgo: Ppcnl for ultrafast optical parametric oscillation in the visible and near-infrared ranges with green pump. Laser Physics, 22:1325–1330, 2012.

- [258] Qiang Liu, Jianhui Liu, Zilong Zhang, and Mali Gong. A high energy 3.75 μm kta optical parametric oscillator at a critical angle. Laser Physics Letters, 10(7):075407, 2013.
- [259] Josef Feldhaus, John Arthur, and JB Hastings. X-ray free-electron lasers. Journal of Physics B: Atomic, Molecular and Optical Physics, 38(9):S799, 2005.
- [260] Beam profiling in the SWIR range: What you need to know. Available at <https://www.ophiropt.com/en/a/beam-profiling-in-swir>.
- [261] Fei Xia, David Sinefeld, Bo Li, and Chris Xu. Two-photon Shack-Hartmann wavefront sensor. Optics Letters, 42(6):1141–1144, 2017.

Appendix A

Numerical simulations of alternative $\chi^{(2)}$ crystals for 1 μm -pumped OPCPA

A.1 Introduction

Optical parametric amplification relies on the existence of a suitable $\chi^{(2)}$ nonlinear medium. The two requirements for nonlinear crystals laid out by Franken et. al in their 1961 paper [95] reporting the first demonstration of second harmonic generation are much the same today as in the early days of nonlinear optics: a material must have a sufficiently large nonlinear dielectric coefficient for the process of interest, and it must be sufficiently transparent to all wavelengths involved in the process. However, with the vast array of nonlinear optical media developed and reported over the last 6 decades, selecting an appropriate crystal and its relevant parameters for a particular application can be challenging.

In this Appendix, we will address one of the foremost concerns limiting the performance of the OPCPA as an effective HHG driver: the limitations of PPLN (periodically poled lithium niobate) and the suitability of alternative crystals. We will describe some alternative crystals, and use the `sisyfos` code package to run numerical simulations to determine the viability of those alternatives - including quasi-phase matched PPLT (lithium tantalate) and PPKTP (potassium titanyl phosphate), as well as birefringently phase-matched KTA (potassium titanyl arsenate), LN (lithium niobate), and KN (potassium niobate). The basic analytical theory of OPA, which is presented in Chapter 2, is extremely informative in selecting crystals. The scaling of amplification - both in terms of energy and bandwidth - is strongly informed by the small-signal gain in the undepleted pump approximation, which is readily understood without the use of simulation. Among

other things, this analytical theory tells us that the phase-matching bandwidth is insufficient to determine the actual amplification bandwidth of an OPA, which is strongly determined by gain. This illustrates a general point about analyzing OPA crystals: many factors are highly coupled. For instance, the Sellmeier equations alone cannot singularly determine the gain bandwidth of an OPA - for that, it is also necessary to know the effective nonlinearity and how much peak intensity the crystal can handle before damaging. Another example of this coupling is in damage threshold, which is in general dependent on wavelength, pulse duration, spot size and repetition rate.

While the analytical treatment of small signal gain is informative, many applications require the use of simulations. The basic foundation of the *sisyfos* code package [100] used for numerical simulations is presented in Chapter 2. Numerical simulations are important for answering the often expensive question of crystal length. We have made a point of including a number of curves showing the evolution of pulse energy as a function of crystal depth with this in mind. We also note that absorption was not considered in the simulations presented in this Appendix. While the crystals considered generally exhibit high transparency to the wavelengths of interest, we note that oxide crystals (in particular, KN) generally exhibit an absorption spike near 2.8-2.9 μm associated with an oxygen bond, and also that KTP begins to gradually absorb as wavelengths exceed 3 μm . The latter absorption was analyzed, and it was found that a 10 mm long KTP crystal would likely absorb 4% of the total idler energy. This was considered to be within acceptable levels, especially considering the crystals would likely be 1-2 mm long.

A.1.1 Damage and pulse duration scaling

Throughout this thesis, we have cited a number of OPCPAs using a wide array of crystals. For many systems, PPLN is utilized in early stages due to its very high d_{eff} of 14.8 - 16.1 (from SNLO and HC Photonics) - but is rarely used for high pulse energy applications due to a low damage threshold. The system presented in this thesis appears to be unique in this regard, with the next highest pump pulse energy found in the literature being 1.35 mJ [161]. The use of a high d_{eff} crystal such as PPLN is motivated by the long, ~ 300 ps pump pulses from our uncompressed




Figure A.1: **Expected scaling of PPLN’s damage fluence with pulse duration for a 1030 nm laser, and the corresponding peak intensity.** The use of shorter pulses allows for higher peak intensity before damage. For a crystal with a known damage threshold and/or aperture size, and a target pulse energy, the duration of a pump pulse is in principle a solvable optimization problem.

cryo-Yb:YAG regenerative amplifier.

Arguably, more limiting than any other factor in the OPCPA presented in this thesis is occurrence of laser damage to the nonlinear crystals. Rather unfortunately, it is very difficult to model and predict than the occurrence of laser induced damage for alternative crystals. Crystal properties, coatings, surface quality - as well as laser parameters such as beam size, pulse duration, and repetition rate/average power - all contribute to the damage threshold. Nevertheless, some rough scaling can be applied within a small range of parameters. In the 10 to few-hundred ps regime, damage fluence typically scales with pulse duration as [256, 251]:

$$LIDT_{\tau_2} \approx LIDT_{\tau_1} \sqrt{\frac{\tau_2}{\tau_1}}. \quad (\text{A.1})$$

In Figure A.1, we plot a reasonable estimate of the damage fluence of PPLN, scaled from a measured value of 10 J/cm² for 10 ns pulses. It is important to take note of the trade-offs in pump pulse duration scaling: shorter pulses enable higher peak intensity while still staying below the damage threshold. This allows for the use of lower d_{eff} bulk crystals, which - in comparison to

PPLN - often have a higher damage threshold, are less expensive, and can be made with a larger clear aperture. On the other hand, the use of shorter pump pulses leads to a lower damage fluence. For the same sized beam, this results in a lower total pulse energy. This can be addressed by scaling up the beam diameters - but only within the limits of the available aperture sizes. Furthermore, damage threshold is also known to decrease with larger beams. These competing factors result in an important optimization problem that balances desired pulse energies, pulse duration and the available crystals. Here, using numerical simulations we present some of the first steps towards that optimization problem, but note that much more can be done in this space.

A.1.2 Summary of results

Of the alternative crystals analyzed, and in consideration of practical availability concerns, PPKTP is thought to be the immediate term most promising alternative for peak power scaling of the 3 μm OPCPA *if* the pump pulse duration is not shortened from its current value of 275 ps FWHM. PPLT and PPKTP are both expected to exhibit higher damage thresholds than PPLN [257, 163]. For instance, a 300 μJ , 3W, few-cycle, 3 μm OPCPA based on PPSLT was reported in Reference [243]. They observed a damage threshold of PPSLT that was 8-18 times higher than PPLN. PPLT offers slightly better performance than PPKTP in terms of amplification bandwidth and conversion efficiency (though still not as impressive as PPLN), but depending on damage fluence, may require cylindrical optics for power scaling. PPKTP is currently available with a 4x4 mm clear aperture (from SLF, Svenska LaserFabriken AB). PPLT is currently available with a maximum clear aperture of 2x10 mm, with the possibility of extension to a 2.8 mm height with some development work (from Oxide). An elliptical beam profile, as was demonstrated with PPLN in Reference [161], may allow for more power handling with the rectangular aperture of PPLT, with the added complication of cylindrical focusing optics. The colinear geometry of periodically poled crystals delivers favorable beam profiles when compared to the noncolinear geometry needed for many bulk crystals, but this effect is likely mitigated if higher peak intensity pulses/shorter crystals are used.

Of the bulk crystals studied, the simulations suggest that potassium niobate is the only option which may potentially be implemented without a dramatic reduction in pulse duration (damage fluence depending). Unfortunately, at the present time, we must discount KN because no vendors has been located. The simulations suggest that low conversion efficiencies exclude both LN and KTA without reduction in pulse duration/increase in peak intensity. Further simulation is needed to understand the behavior of these two crystals with shortened pulses, but it is likely that they will both be able to offer high conversion efficiencies at the appropriate pump pulse duration. Of the available bulk crystals, it appears that lithium niobate offers dramatically improved amplification bandwidth compared to KTP in either colinear or noncolinear geometries. On the other hand, the colinear phase matching geometry of KTA produces significantly better beam profiles. The simulations suggest that amplification bandwidth and modal distortions may need to be weighed in deciding between LN and KTA. Due to the generally large apertures available for these bulk crystals, power scaling may be highly promising if a reduction in pump pulse duration is possible. It is worth noting that in our case, if a hybrid periodically-poled/bulk crystal OPCPA is built, the early QPM stages need not - and probably should not - be pumped with the shortened pulses, to employ their full capabilities and deliver a powerful seed for the later bulk stages. Such an arrangement could be realized by simply picking off some component of the regenerative amplifier beam prior to compression.

In the future, if PPLT becomes available with a larger aperture, or if potassium niobate becomes available at all, based on these simulation results, those crystals should likely replace PPKTP and MgO-doped lithium niobate as the most promising alternatives. Lastly, we note that while simulations are useful guides, experimentation is necessary to make a complete determination. In this regard, we point the reader to a particular useful reference in which MgO:PPLN, KTA and KN were compared using experimental methods for mid-IR OPCPA [163], causing that group to pursue the use of KN for the high power stage of the OPCPA presented in Reference [73].

A.2 Comparative numerical simulations of periodically-poled crystals: PPLN, PPLT and PPKTP

We used sisyfos to numerically simulate the performance of PPLN, PPSLT and PPKTP in the final stage of the OPCPA (OPA4). The bulk approximation, which is described in detail in Chapter 2, was used to simulate QPM. We note that sisyfos can be implemented to simulate the actual periodically poled gain medium without approximation, but such an implementation requires a separate program to be prepared by the FFI. Further, the provider of the code tested the code used for these simulations with full QPM and found no significant change in the generated spectra.

Each crystal was simulated for 3 different pump fluences, resulting in nine total simulations. The fluence is modified by increasing the pump energy with a fixed spot size. Each simulation took ~ 2.5 min on a desktop computer. The general parameters for each simulation are presented in Table A.1, and the parameters specific to each crystal are presented in Table A.2. The Sellmeier equations are for 5% MgO-doped PPLN, 0.5% MgO-doped PPSLT, and bulk KTP.¹ Importantly, we note that we have experimentally observed damage in PPLN above ~ 1.2 J/cm² with comparable spot sizes to this simulation - although the exact damage fluence has been higher or lower from crystal to crystal. This means that the PPLN data points at 2 and 3 J/cm² are for academic purposes only. Optimistically, PPKTP and/or PPLT may be able to handle these higher fluences, but such a determination is extremely difficult to make without actual measurements.

At an arbitrary z-position, the simulation can produce a 3D array of complex numbers representing the electric field in the two transverse coordinates (x and y) and in time. We extract this data at 39 slices at evenly spaced depths in the crystal, including at the front and back surfaces. At each slice, we integrate the total pulse energy of the pump, signal and idler. In Figures A.4 - A.6 (bottom), we plot the evolution of the 3 μm “idler” energy in the crystal for three different pump fluences. We maintain the convention of referring to the 3 μm pulse as the idler, even though it is the seed in this simulation. These curves can be used to identify best-guess optimum crystal thick-

¹ The Sellmeier equations for PPKTP from Katz, 2001 [254] and the associated erratum are considered preferable, but were not available for the simulation.

Crystal length	2.5 mm
QPM	Bulk approx.: $\Delta k=0$ at 99.9, 291.2, 191.3 THz
Seed	3 μm ctr. λ , stretched 160 ps FWHM Gaussian supporting 66 fs T.L.
Seed spot size	2.52 mm $1/e^2$ diam.
Seed energy	0.5 mJ
Pump	1029.5 nm ctr. λ , 275 ps FWHM T.L. Gaussian
Pump spot size	1.8 mm $1/e^2$ diam.
Pump energy	Variable: 12.7, 25.4, 38.1 mJ
Pump peak fluence	1, 2, 3 J/cm ²
Pump peak intensity	3.6, 7.3, 11 GW/cm ²

Table A.1: **Simulation parameters used for comparison of PPLN, PPSLT and PPKTP as candidates for the final OPA stage.** T.L. indicates the pump is at its transform limit. The Gaussian seed is stretched from its transform limited duration of 66 fs to 160 ps FWHM through the application of 2.08×10^6 fs² of GDD - comparable parameters to the actual OPCPA.

	PPLN	PPSLT	PPKTP
d_{eff}	14.8 pm/V	10.4 pm/V	8.5 pm/V
Polarization	fff	fff	sss
Sellmeier eq.	Gayer, 2008	Dolev, 2009	Kato, 2002

Table A.2: **Crystal parameters used in sisyfos simulations.** Gayer, 2008: [117]; Dolev, 2009: [252]; Kato, 2002: [253]. f and s stand for fast (low index) and slow (high index) and refer to the 3 μm , 1 μm and 1.5 μm waves, in that order. For all three crystals, the angle of propagation of the incident beams relative to the optic axis of the crystal is 90°. All three beams are colinear.

nesses. We also note that if the crystal aperture allows and the amplified bandwidth is sufficient, simply increasing the spot size without changing the fluence can result in more energetic amplified pulses - but in this case, the gain peak will shift due to the lower seed intensity. For a sense of how the gain peak shifts with a variable seed intensity, see Chapter 4. For each curve in Figures A.4 - A.6 (bottom), we identify the crystal plane where the idler has reached its maximal energy. At this




Figure A.2: The solid blue line shows the temporal envelope of the pulse used to seed the crystal in the simulation. The dashed yellow line shows the temporal envelope of the pump pulse before the crystal. A simulation was performed with PPLN pumped at a peak fluence of 1 J/cm^2 . We extracted the temporal envelopes of both the pump and the $3 \mu\text{m}$ pulse at the peak of the gain curve. At this plane, the pump pulse energy has been reduced from 12.7 to 6.7 mJ. Shortening the seed pulse relative to the pump pulse can improve conversion efficiency, but it also means that the spectral/temporal components near the wings of the spectrum see significantly less gain. Inspection of the dashed purple line shows the start of back-conversion - at the most intense part of the $3 \mu\text{m}$ pulse, corresponding to a region of depleted pump, signal and idler photons combine to form pump photons. In the spatial domain, this will eventually lead to a doughnut-shaped mode of the signal and idler.

plane, we extract the idler spectrum, and plot it in Figures A.4 - A.6 (top). Due to the different phase matching and gain characteristics of each crystal, the amplified idler spectrum varies slightly. In general, we find relatively little variance in the width of the amplified spectrum of a particular crystal by changing the pump fluence. On the other hand, we do observe some variability in gain bandwidth between the three crystals.

We further analyzed the simulation results to understand variations in gain bandwidth between the three crystals. The $3 \mu\text{m}$ idler spectra plotted in Figures A.4 - A.6 (top) is the full spectrum out of the crystal, and therefore contains components from the fairly sizeable 0.5 mJ seed that were not generated in the crystal. The transform limited pulse duration supported by




Figure A.3: **Transform-limit of the generated 1.5 μm signal pulse at the maximum of the gain curve.** The very high figure of merit (FOM) for PPLN enables high gain bandwidth, but that can only be realized for low fluences due to the low damage threshold. The marginally lower FOM of PPLT and PPKTP can be somewhat compensated by a higher pump fluence - if the crystals will tolerate it.

the amplified idler pulse is therefore artificially improved by the broad seed bandwidth, and is a sub-optimal metric of the actual spectral gain properties of the crystal. Instead, we considered the 1.5 μm signal spectra that was entirely generated in the crystal - extracted at the crystal plane of maximal gain. For each of the nine simulations, we computed the Fourier transform of this signal spectrum, and interpolated it to extract the supported FWHM transform limited pulse duration. This pulse duration is plotted against the pump fluence for each crystal in Figure A.3. PPLN supported the shortest amplified pulses, followed by PPLT and then PPKTP. We note that separate simulations comparing PPKTP and PPLN have indicated that the resultant amplified bandwidth depends strongly on the seed bandwidth - it was generally found that for both crystals, increasing the seed bandwidth will increase amplified bandwidth in this regime. Therefore, if amplified bandwidth is found to be problematically low, increasing the seed bandwidth may be an avenue forward, if possible. Nevertheless, with the 66 fs transform limited seed pulse used in this simulation, all

crystals supported near or sub-100 fs signal pulses - and even shorter idler pulses .

Considering the reasonable conversion efficiencies and only slightly reduced amplified bandwidth suggested by the simulations, both PPLT and PPKTP may exhibit superior performance to PPLN for the high power, final amplification stage of the OPCPA if they are able to withstand higher fluences. We note that the amplified bandwidth is expected to improve with higher pump fluence. If amplified bandwidth is problematic, improvements may be possible by tilting the crystal [186].

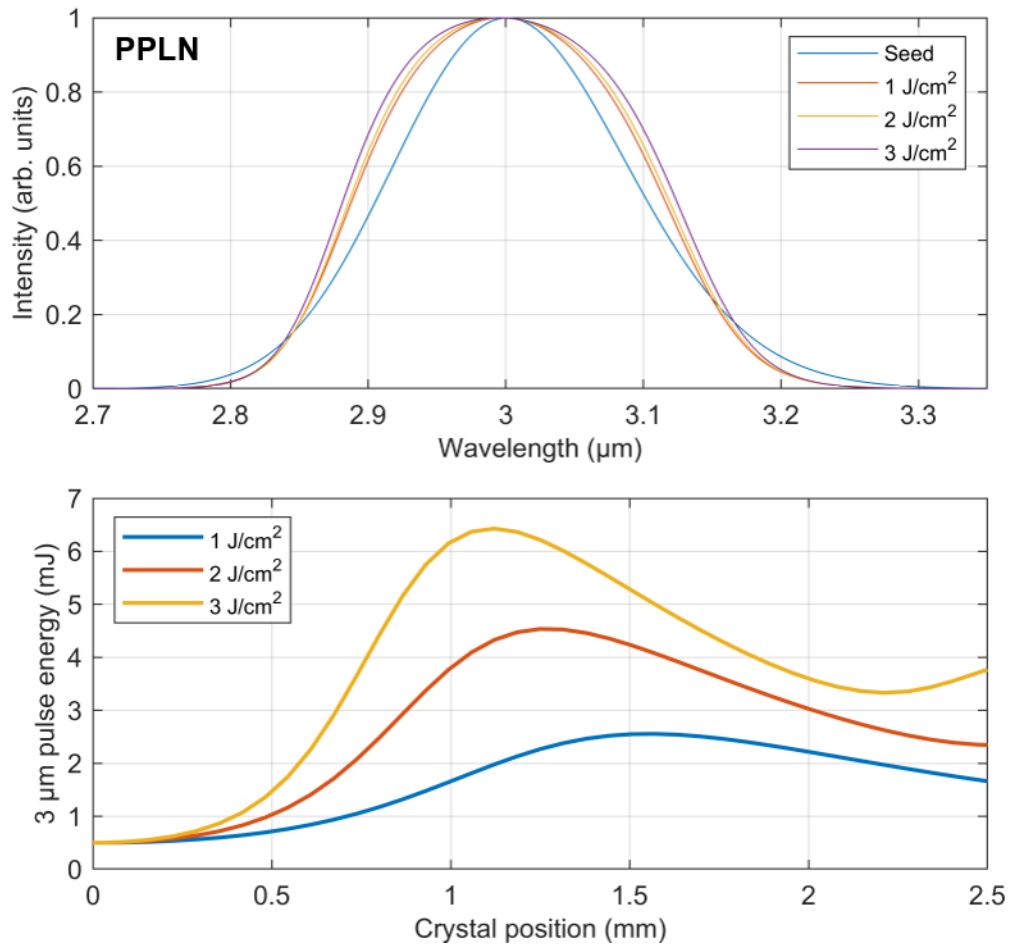


Figure A.4: **Numerical simulation of OPA in 5% MgO-doped PPLN.** Top: 3 μm idler spectra as they appear at the peak of the gain curve - not at the back facet of the crystal. Bottom: evolution of idler pulse energy through the crystal at various pump fluences. For parameters, see Tables A.1 and A.2.

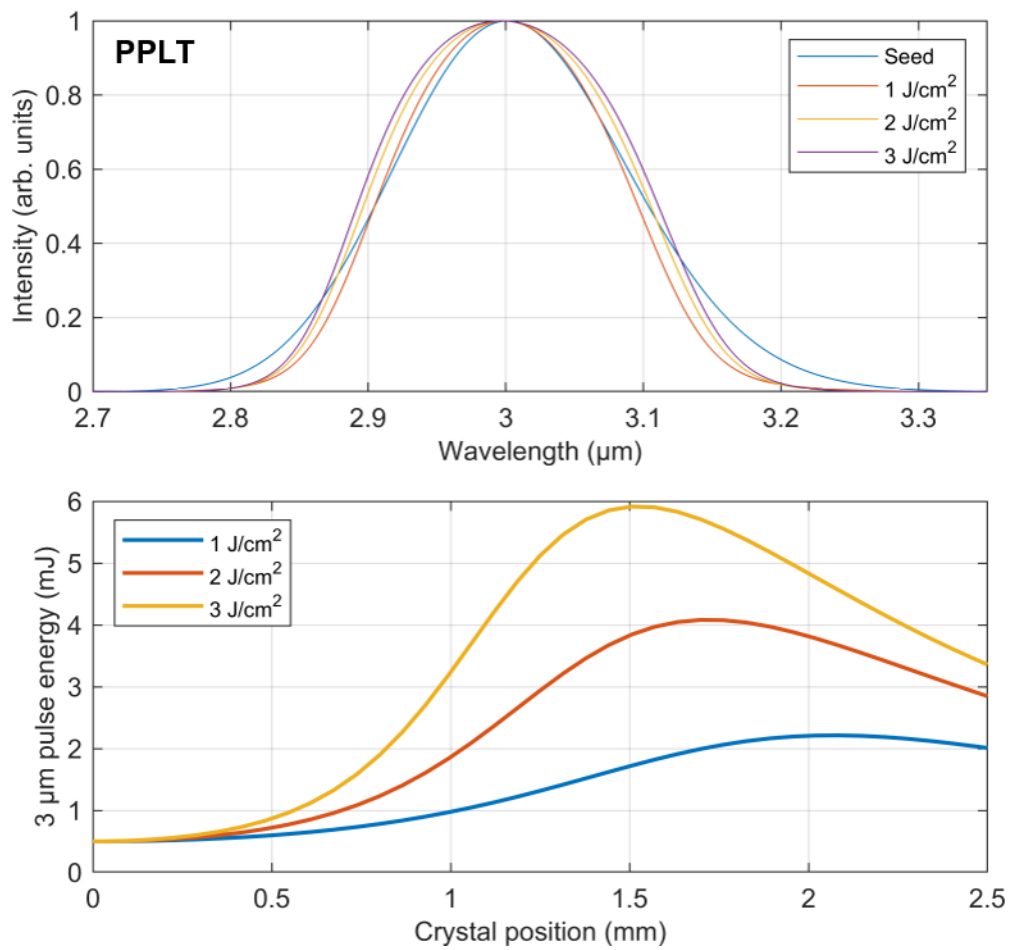


Figure A.5: **Numerical simulation of OPA in 0.5% MgO-doped PPLT.** Top: 3 μm idler spectra as they appear at the peak of the gain curve - not at the back facet of the crystal. Bottom: evolution of idler pulse energy through the crystal at various pump fluences. For parameters, see Tables A.1 and A.2.

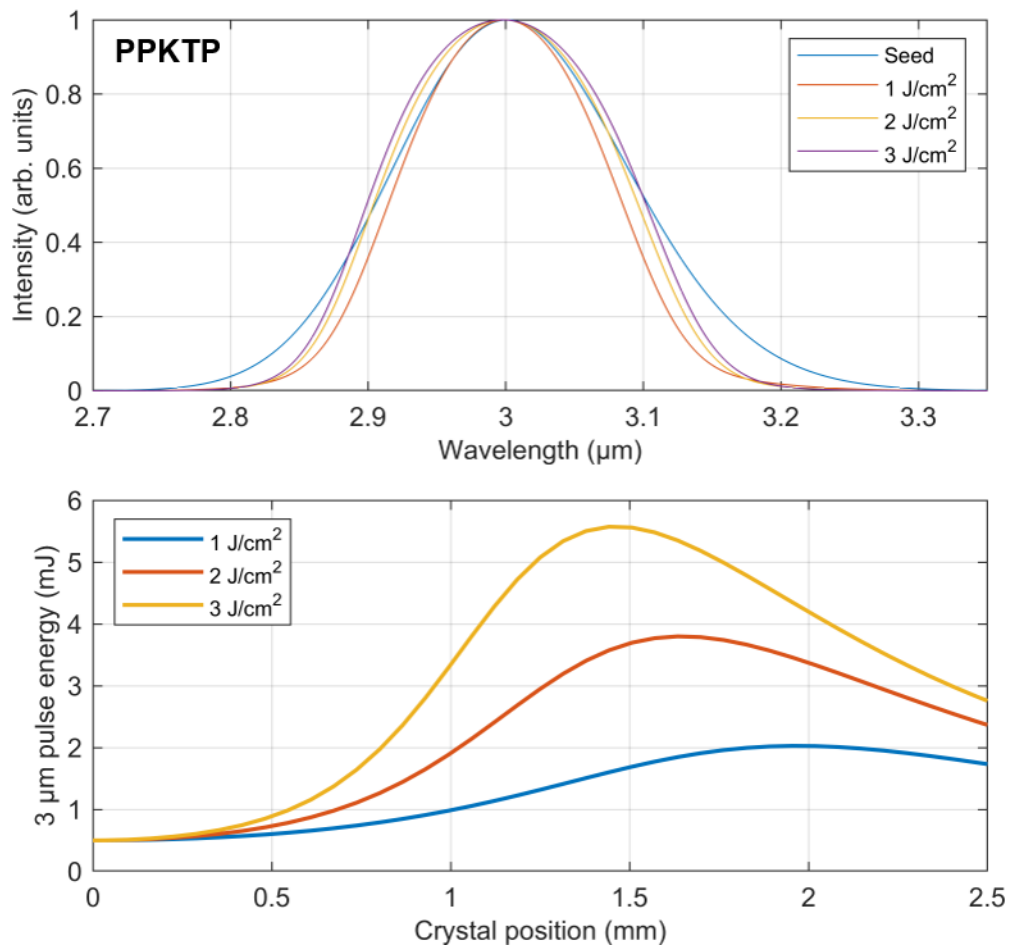


Figure A.6: **Numerical simulation of OPA in PPKTP.** Top: 3 μm idler spectra as they appear at the peak of the gain curve - not at the back facet of the crystal. Bottom: evolution of idler pulse energy through the crystal at various pump fluences. For parameters, see Tables A.1 and A.2.

A.3 Comparative numerical simulations of bulk crystals: KTA, LN and KN

We also used sisyfos to simulate the performance of KTA, 5% MgO-doped:LN and KN as candidates for the final stage of the OPCPA (OPA4). Generally, these bulk, birefringently phase matched crystals can be made with larger clear apertures than periodically poled (PP) crystals, whose aperture size is in principle limited by the generally high coercive field needed to apply poling. Furthermore, bulk crystals are generally less expensive and easier to fabricated than PP crystals. On the other hand, the d_{eff} of bulk crystals is generally much lower than for PP crystals, because the requirements of birefringent phase matching generally preclude accessing the highest valued coefficients of the dielectric tensor. As was shown in our analytical calculations using small-signal gain in Chapter 2, the small signal gain depends exponentially on the figure of merit determined by $\text{FOM} = d_{\text{eff}}^2/n_{\text{pump}}n_{\text{signal}}n_{\text{idler}}$. Accordingly, the gain is even more sensitive to changes in d_{eff} than it is to changes in peak intensity of the pump. Beyond hindering total energy conversion efficiency, the reduced gain can also deflate the amplification bandwidth.

A theoretical study of OPCPAs in the same parameter space as the system presented here, including a review of these crystals, can be found in Reference [255]. KTA is a positive biaxial crystal that can be used in both a colinear and noncolinear configuration, but we find it performs better in the colinear regime. KTA is widely available and has the lowest d_{eff} of the crystals we considered. LN is a negative uniaxial crystal that is generally used in a noncolinear configuration, with a moderate d_{eff} and relatively high availability. Bulk LN, like PPLN, can also be MgO-doped to reduce photorefractive effects and green induced infrared absorption. It is thought that bulk LN generally has a higher damage threshold than PPLN. Lastly, KN is a highly promising negative biaxial crystal known to be highly resistant to laser damage and exhibiting the highest d_{eff} of the bulk crystals considered here. However, at the time of writing, despite extensive searching, no supplier of KN is known to the authors. Other potentially promising crystals include the lithium chalcogenide crystals LGS and LGSe.

In general, bulk crystals are used when shorter pump pulse are available. This allows access

to higher peak intensities before reaching the damage fluence, enabling higher gain and shorter crystals. Among the longer pump pulses used with KTA are the 70 ps pulses used to pump the OPCPA that generated 1.6 keV HHG in 2012 [71]. Those pulses had flat-top, multimode profiles, energies above 250 mJ and resulted in peak intensities on the crystal of about 13 GW/cm² and peak fluences of around 0.9 J/cm². Interestingly, despite the use of 10 mm long KTA crystals and the use of colinear geometry, they were able to amplify bandwidth supporting 83 fs pulses. More common is the use of KTA in OPAs and OPCPA with few ps pump pulses. Another interesting result claims to have used bulk LN pumped with long, 430 ps pulses at a low peak intensity of 2 GW/cm² to amplify 111 fs pulses at 3.3 μm up to 30 mJ, with a gain of over 4000 [245]. The authors apparently used a 40 mm long crystal. This result is distinct from any others found in the literature.

The simulations were similar to the periodically poled simulations presented in the previous section. Symmetry considerations allowed the PP crystal simulations to be done in a quarter spatial grid (17x17 -j, 32x32); due to the spatial asymmetry from noncolinear phase matching, we perform these simulations on a full grid (32x32). For each simulation, the 3 μbeam seeds the OPA at normal incidence. Nevertheless, the amplified 3 μm beams can exhibit distortions, in part due to the shifting overlap between the pump and the seed through the length of the crystal. These crossing angle effects can be mitigated by using large beam diameters - as would likely occur if the pump were compressed. Modes of the pump before and after the 4 mm crystal are shown partly as an illustration of the relative pump angle. The asymmetry in the amplified profiles from colinear KTA may arise from spatial walk-off [258]. In noncolinear OPA, the unseeded beam (in this case, the 1.5 μm beam) generally experiences angular dispersion to complete the phase matching across a broad spectral range.

The parameters for the simulation and the crystals are presented in Tables A.3 and A.4. Compared to the periodically poled simulations, we have reduced the beam diameters by a factor of two. This was done to give more realistic pump energy values to reach the higher fluences at which bulk crystals may operate at. While 8 J/cm² may be a very high fluence for some of

these crystals, at a shorter pulse duration, 29 GW/cm² may not be. The data in Table A.4 was extracted from the SNLO software. We note that the listed Sellmeier equation sources refer to what was used for the sisyfos simulations, and may differ from the default SNLO sources. Figure A.7 attempts to summarize the gain bandwidth of each of the 16 simulations performed. As before, this is done by extracting the generated 1.5 μm spectrum at the crystal plane corresponding to the most energetic 3 μm pulse (these planes can be identified in the subsequent figures). We then Fourier transform this spectrum to the time domain, interpolate it, and characterize the temporal FWHM of the transform limited pulse - thereby reducing the influence of the seed spectrum on the characterization of amplification bandwidth. It can be seen that both KN and LN perform similarly well with regard to amplification bandwidth, even with the long pump pulses considered - but only KN delivers sufficient conversion efficiency.

Crystal length	4 mm
Phase matching	Birefringent: no bulk approximation
Seed	3 μm ctr. λ , stretched 160 ps FWHM Gaussian supporting 66 fs T.L.
Seed spot size	1.26 mm $1/e^2$ diam.
Seed energy	0.5 mJ
Pump	1029.5 nm ctr. λ , 275 ps FWHM T.L. Gaussian
Pump spot size	0.9 mm $1/e^2$ diam.
Pump energy	Variable: 6.35, 12.7, 19.05, 25.4 mJ
Pump peak fluence	2, 4, 6, 8 J/cm ²
Pump peak intensity	7.3, 14.5, 21.8, 29 GW/cm ²

Table A.3: **Simulation parameters used for comparison of KTA, LN and KN as candidates for the final OPA stage.** T.L. indicates the pump is at its transform limit. The Gaussian seed is stretched from its transform limited duration of 66 fs to 160 ps FWHM through the application of 2.08×10^6 fs² of GDD - comparable parameters to the actual OPCPA.

	KN	LN	KTA non-col.	KTA col.
d_{eff}	5.98 pm/V	-3.93 pm/V	-2.49 pm/V	-2.14 pm/V
Polarization	sfs	sfs	ffs	ffs
Sellmeier eq.	Ghosh, 1994	Gayer, 2008	Fenimore, 1995	Fenimore, 1995
k_{trans} (rad/m)	9.27×10^5	1.16×10^6	-1.49×10^6	0
$\theta_{3 \mu\text{m seed}}$	37.0°	41.7°	63.1°	42.1°
$\Delta\theta_{1 \mu\text{m pump}}$	4°	5°	4°	0°
P.M. type	Type I, XZ	Type I	Type II, XZ	Type II, XZ

Table A.4: **Crystal parameters used in sisyfos simulations of bulk crystals.** f and s stand for fast (low index) and slow (high index) and refer to the 3 μm , 1 μm and 1.5 μm waves, in that order. k_{trans} refers to the transverse component of the wave-vector for both the 1 and 1.5 μm beams, in rad/m. For noncolinear phase matching, these values must be the same for both waves. The transverse wave-vector for the 3 μm beam is zero. This is how noncolinear angles are specified in the sisyfos simulations. $\theta_{3 \mu\text{m seed}}$ refers to the angle of propagation of the seed beam relative to the optic axis of the crystal. SNLO would refer to $\theta_{3 \mu\text{m seed}}$ as the “red1 angle”, where red1 is the 3 μm beam. $\Delta\theta_{1 \mu\text{m pump}}$ is the angle added to $\theta_{3 \mu\text{m seed}}$ to determine the angle of the pump relative to the optic axis. SNLO would refer to $\Delta\theta_{1 \mu\text{m pump}}$ as “blue tilt.”




Figure A.7: **Transform-limit of the generated 1.5 μm signal pulse at the maximum of the gain curve, or the back facet of the crystal - whichever plane produces more energetic 3 μm pulses.** KTP in the colinear and noncolinear configurations has significantly narrowed amplification bandwidth than the seed spectrum, while KN and LN both exhibit very high amplification bandwidth. However, only KN exhibits reasonable conversion efficiencies when pumped below the likely damage threshold.

Potassium niobate (KN): noncollinear

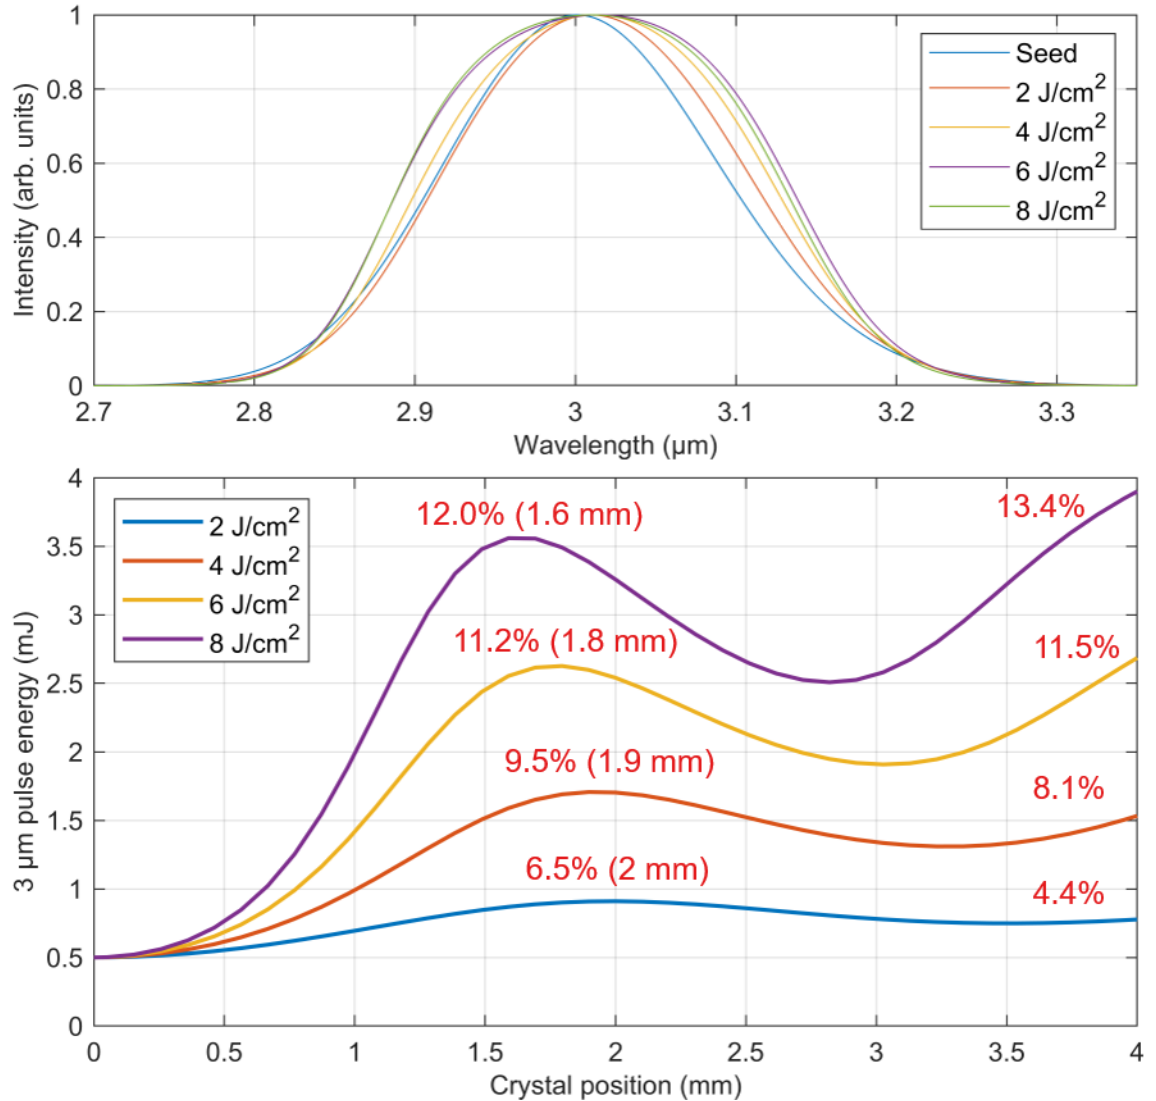


Figure A.8: **Numerical simulation of noncollinear OPA in bulk potassium niobate.** Top: spectra of the amplified 3 μm pulses at various fluences. For fluences which exhibit a defined gain peak in the bottom plot, the spectrum is extracted at the position of the gain peak. Otherwise, the spectrum is extracted at the back facet of the crystal ($z=4$ mm). Bottom: evolution of the total 3 μm pulse energy through the crystal at various pump fluences. Red text indicates the conversion efficiency of pump to 3 μm , calculated by subtracting the seed energy from the amplified energy and dividing by the incident pump energy. Percentages on the right were calculated at $z=4$ mm (crystal back facet). Simulation parameters are in the text.

Potassium niobate (KN): noncollinear

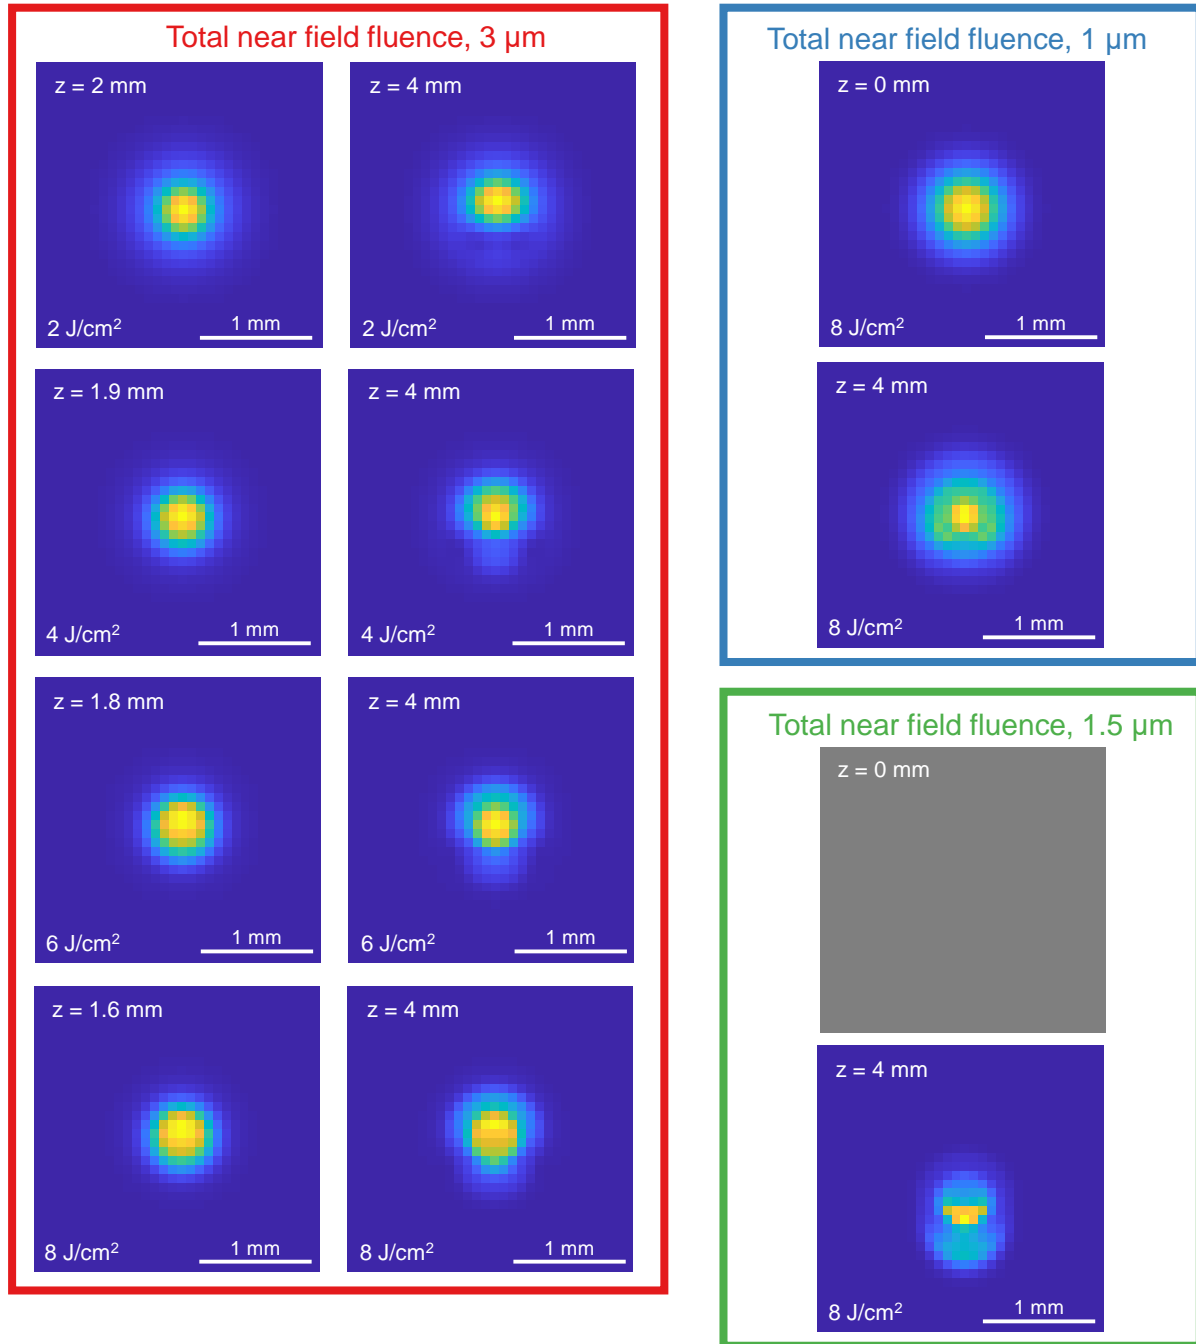


Figure A.9: Total near field fluence spatial modes (intensity integrated over the duration of the pulse) from the simulation of noncollinear OPA in potassium niobate. Selected crystal depths (left column) correspond to peaks in the gain curve. Simulation parameters are in the text.

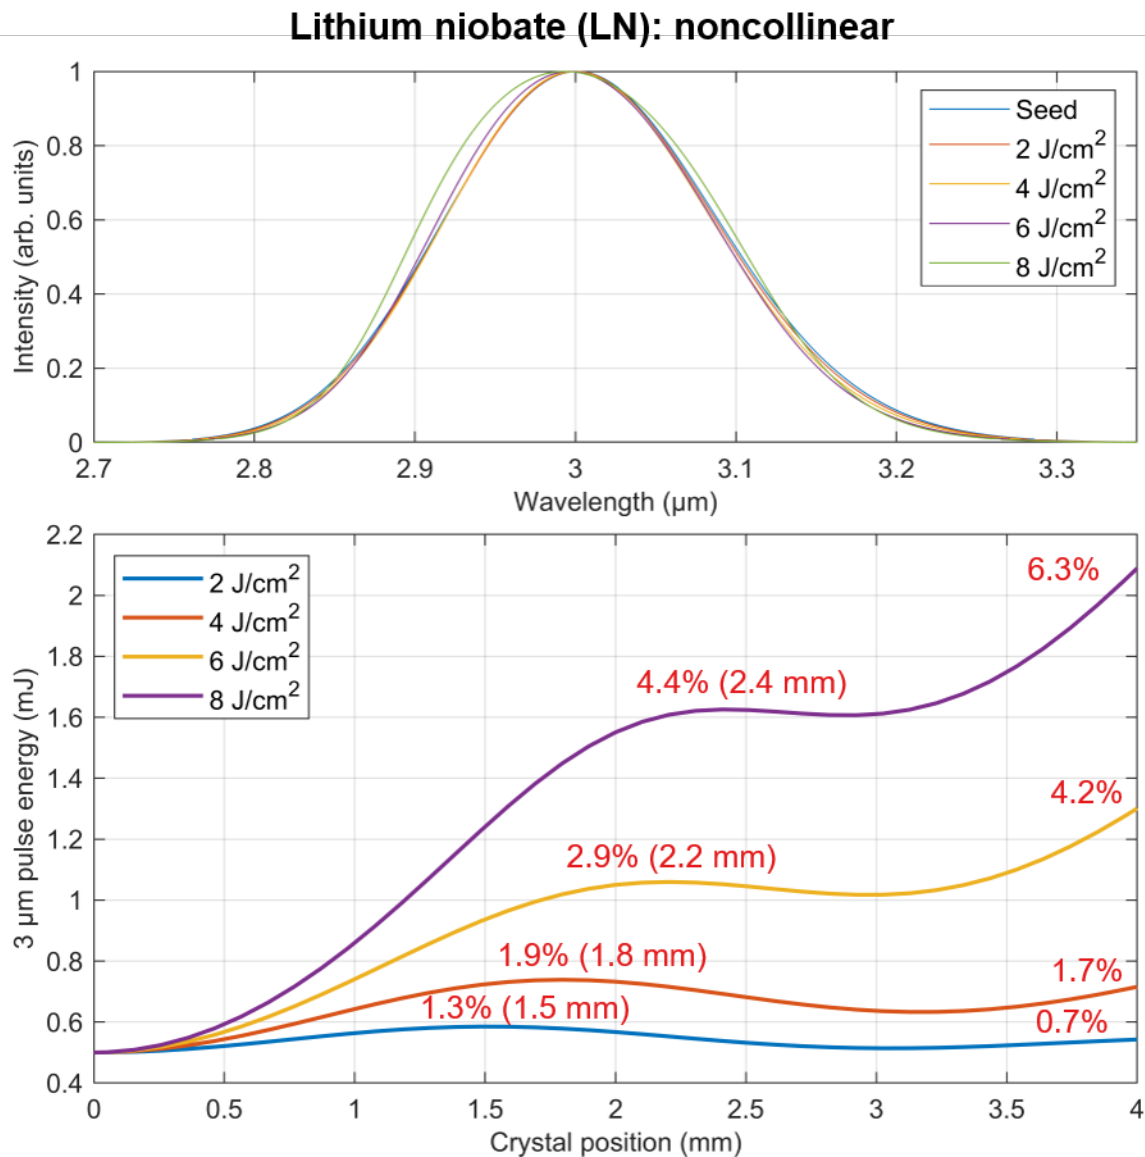


Figure A.10: **Numerical simulation of noncollinear OPA in bulk lithium niobate.** Top: spectra of the amplified 3 μm pulses at various fluences. For fluences which exhibit a defined gain peak in the bottom plot, the spectrum is extracted at the position of the gain peak. Otherwise, the spectrum is extracted at the back facet of the crystal ($z=4$ mm). Bottom: evolution of the total 3 μm pulse energy through the crystal at various pump fluences. Red text indicates the conversion efficiency of pump to 3 μm, calculated by subtracting the seed energy from the amplified energy and dividing by the incident pump energy. Percentages on the right were calculated at $z=4$ mm (crystal back facet). Simulation parameters are in the text.

Lithium niobate (LN): noncollinear

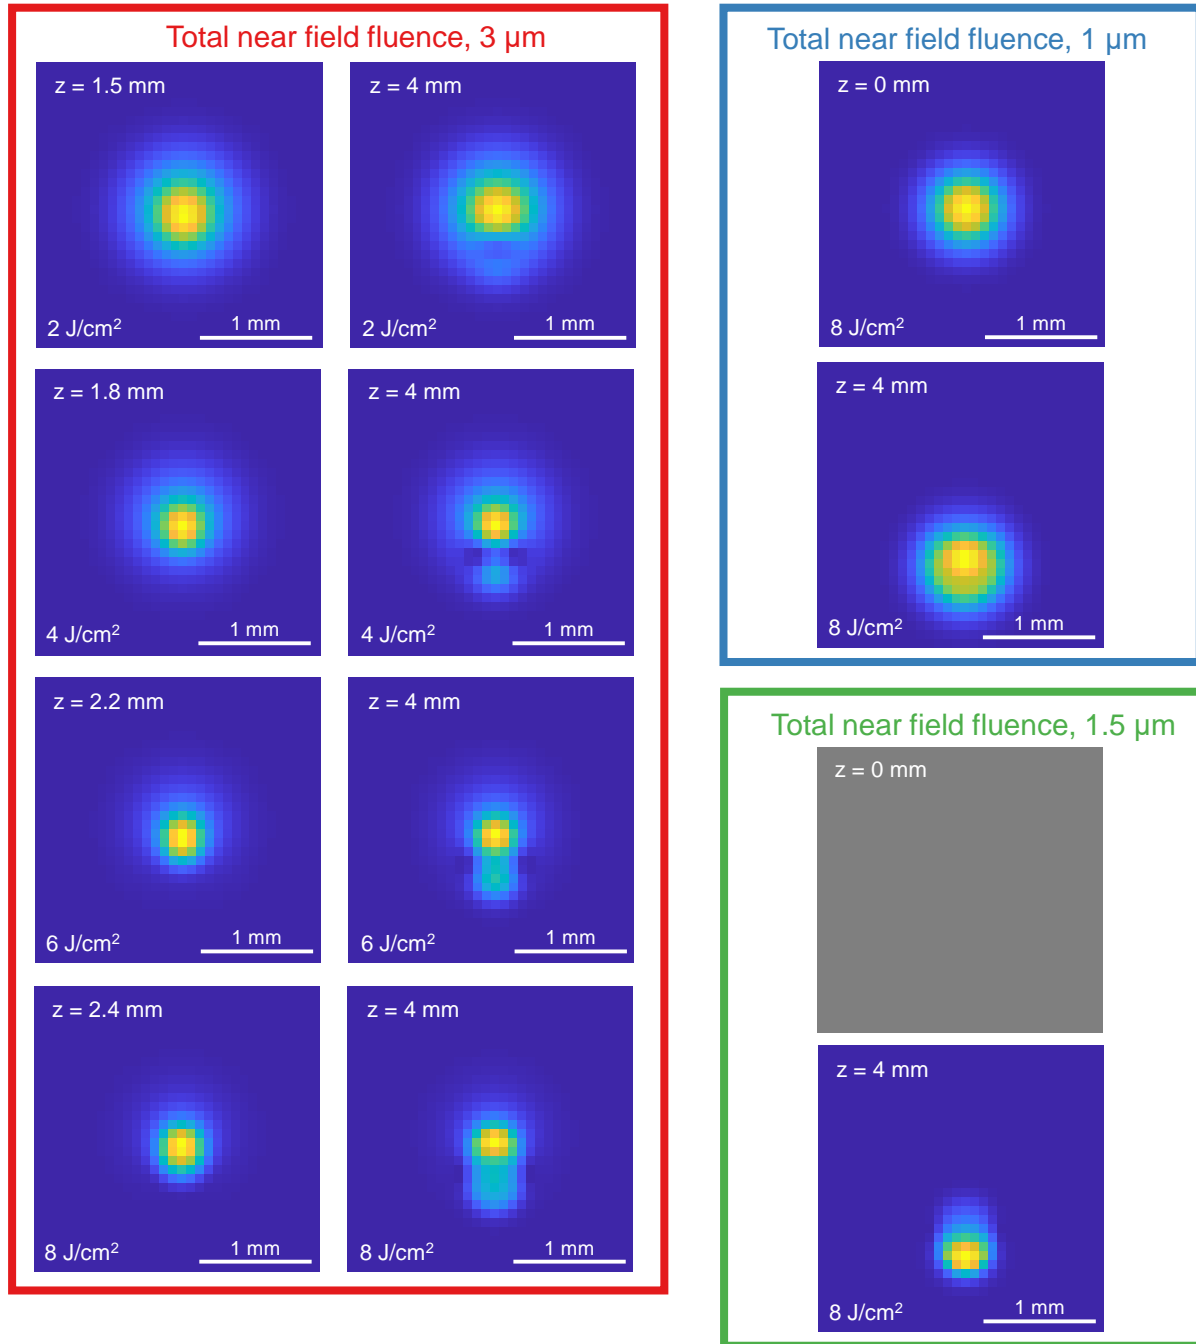


Figure A.11: Total near field fluence spatial modes (intensity integrated over the duration of the pulse) from the simulation of noncollinear OPA in lithium niobate. Selected crystal depths (left column) correspond to peaks in the gain curve. Simulation parameters are in the text.

Potassium titanyl arsenate (KTA): noncollinear

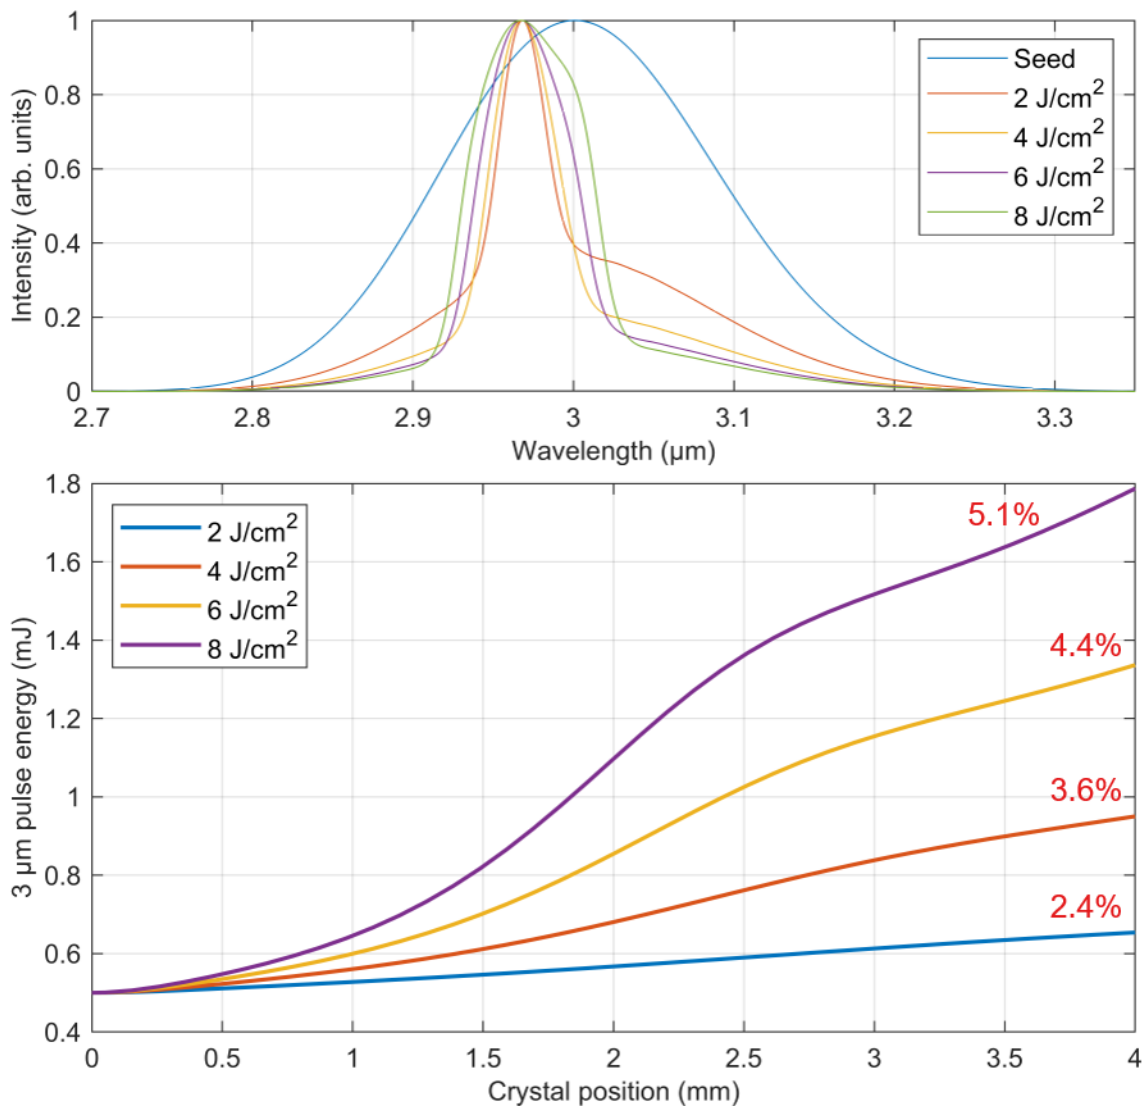


Figure A.12: **Numerical simulation of noncollinear OPA in bulk potassium titanyl arsenate.** Top: spectra of the amplified 3 μm pulses at various fluences. For fluences which exhibit a defined gain peak in the bottom plot, the spectrum is extracted at the position of the gain peak. Otherwise, the spectrum is extracted at the back facet of the crystal ($z=4$ mm). Bottom: evolution of the total 3 μm pulse energy through the crystal at various pump fluences. Red text indicates the conversion efficiency of pump to 3 μm , calculated by subtracting the seed energy from the amplified energy and dividing by the incident pump energy. Percentages on the right were calculated at $z=4$ mm (crystal back facet). Simulation parameters are in the text.

Potassium titanyl arsenate (KTA): noncollinear

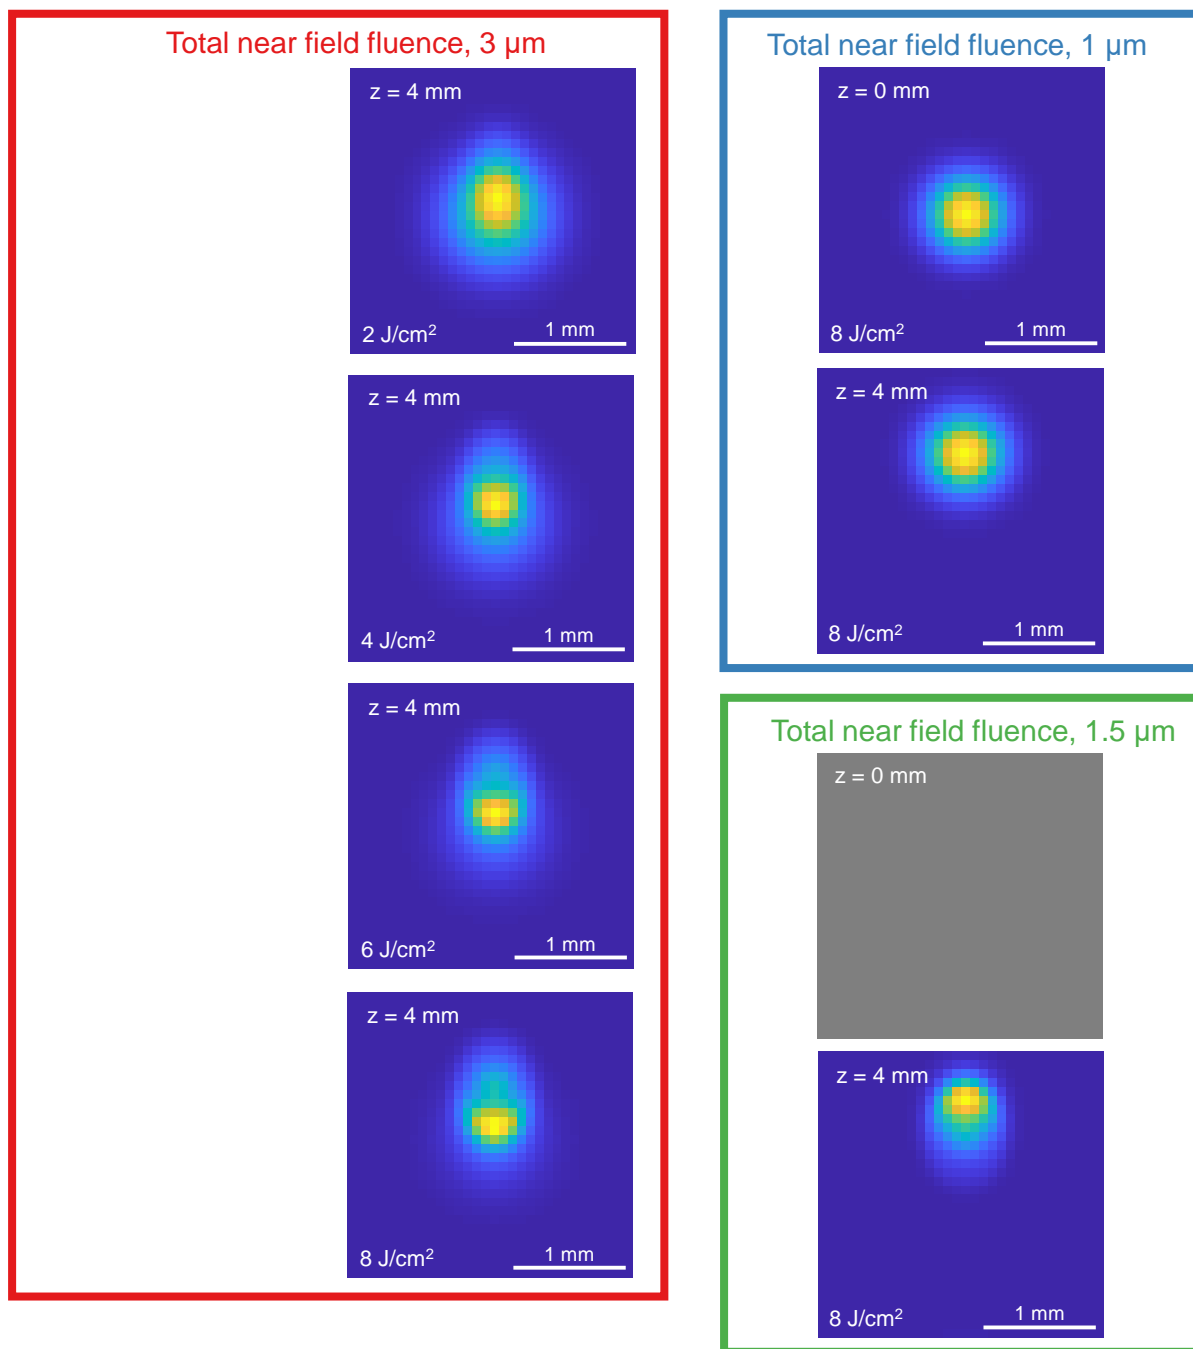


Figure A.13: Total near field fluence spatial modes (intensity integrated over the duration of the pulse) from the simulation of noncollinear OPA in potassium titanyl arsenate (KTA). Simulation parameters are in the text.

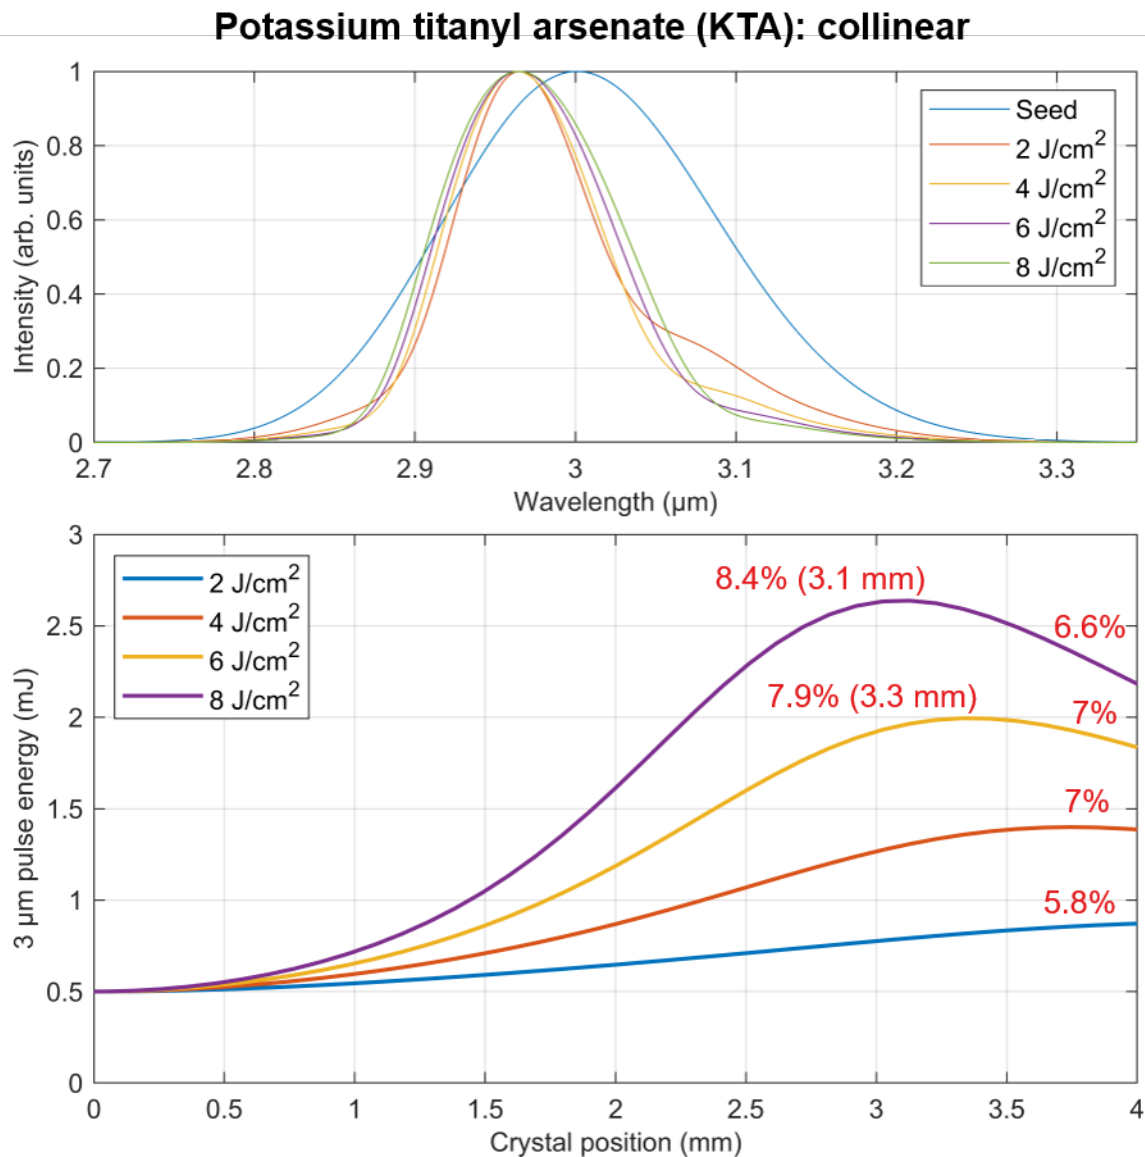


Figure A.14: **Numerical simulation of collinear OPA in bulk potassium titanyl arsenate.** Top: spectra of the amplified 3 μm pulses at various fluences. For fluences which exhibit a defined gain peak in the bottom plot, the spectrum is extracted at the position of the gain peak. Otherwise, the spectrum is extracted at the back facet of the crystal ($z=4$ mm). Bottom: evolution of the total 3 μm pulse energy through the crystal at various pump fluences. Red text indicates the conversion efficiency of pump to 3 μm , calculated by subtracting the seed energy from the amplified energy and dividing by the incident pump energy. Percentages on the right were calculated at $z=4$ mm (crystal back facet). Simulation parameters are in the text.

Potassium titanyl arsenate (KTA): collinear

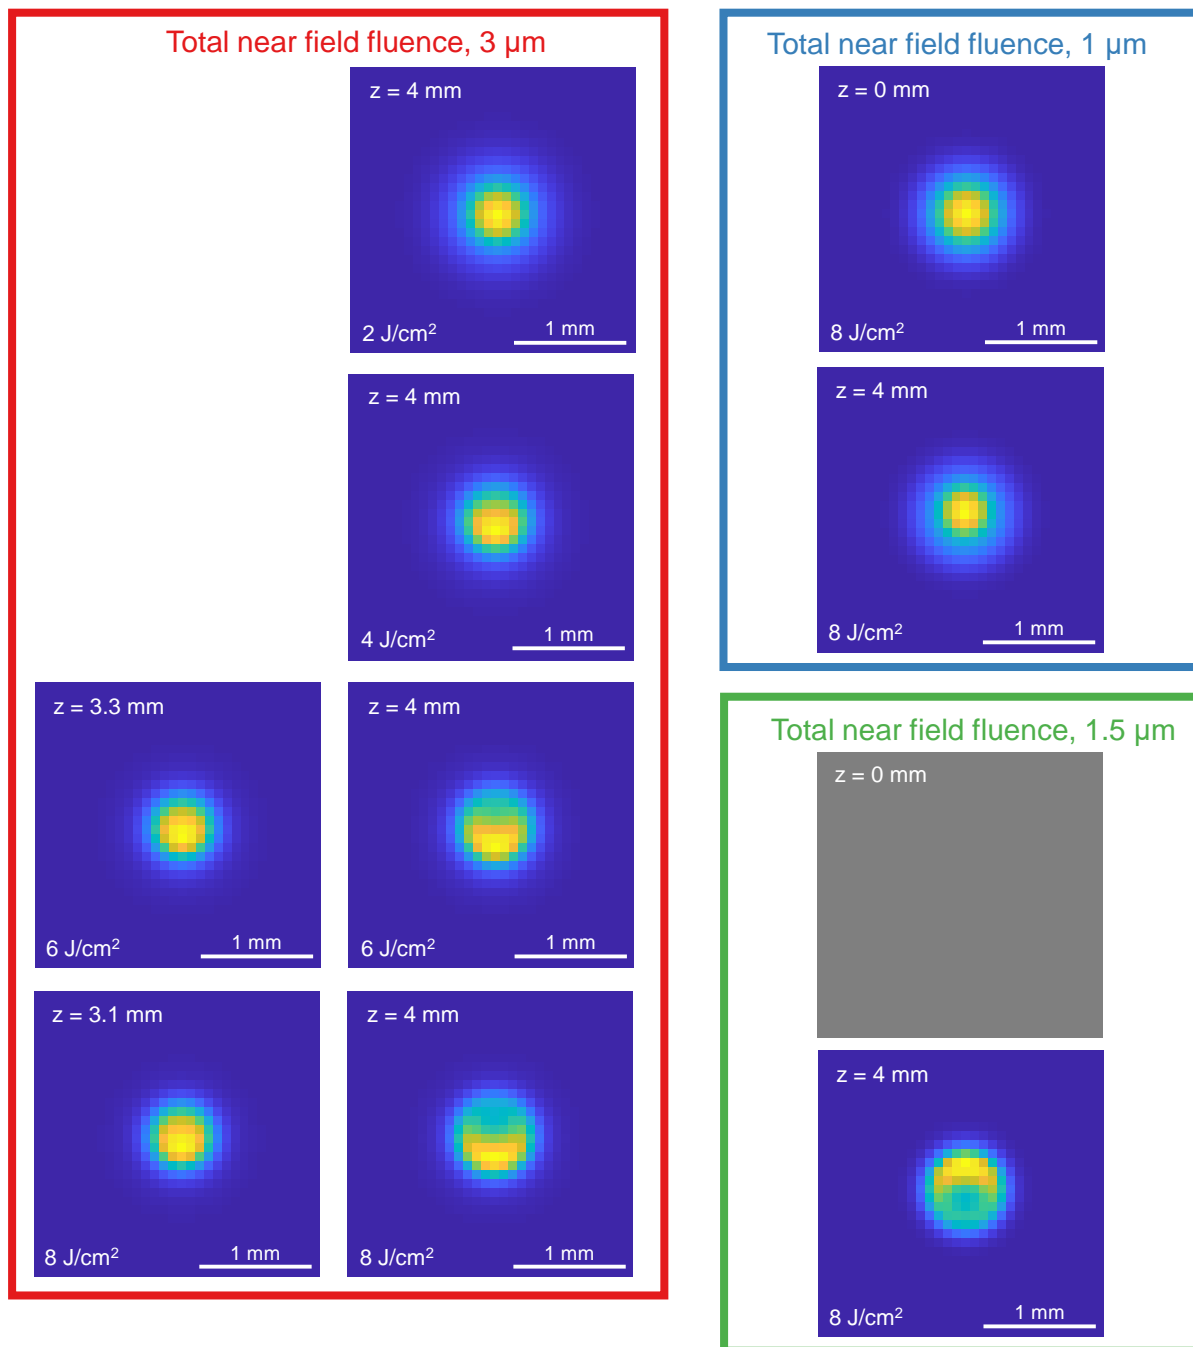


Figure A.15: Total near field fluence spatial modes (intensity integrated over the duration of the pulse) from the simulation of collinear OPA in potassium titanyl arsenate (KTA). Simulation parameters are in the text. Selected crystal depths (left column) correspond to peaks in the gain curve.

Appendix B

Two-photon beam profiling at telecom wavelengths with a standard silicon sensor

Silicon-based image sensors, found in virtually every consumer image digital camera, are extremely common and relatively affordable. The technology is highly mature, with pixel dimensions commonly as low as a few microns. Silicon-based cameras are the technology of choice for profiling visible and near-IR beams with wavelengths less than 1100 nm, and were used extensively in this thesis to characterize the beam of the 1030 nm wavelength OPCPA pump laser. However, light with wavelengths greater than 1100 nm does not carry enough photon energy to excite electrons across silicon's band gap of around 1.1 eV, rendering silicon cameras insensitive to those wavelengths in the single photon absorption regime.

InGaAs image sensors typically exhibit sensitivity in the 900-1700 nm (SWIR) range, with pixel sizes generally somewhat larger than with silicon detectors. These detectors are uncommon and costly, but are considered effective. An alternative technology involves the application of a wavelength up-converting phosphor to a silicon image sensor. While this avoids the use of costly InGaAs sensors, there are a number of considerable drawbacks: fine structure of the beam is lost due to the phosphorescence process, dynamic range is dramatically reduced, and the phosphor response is highly wavelength dependent even over a few 10's on nm [260]. Pyroelectric array cameras can also be sensitive to SWIR, but are again expensive, inconveniently large, and have very large pixel size - for instance, the Pyrocam in our lab contains 80 μm pixel pitch with a 320x320 sensor array - all but obscuring fine features of the beam. Despite these shortcomings in




Figure B.1: **Raw image of an estimated 12 μJ , 1 kHz, 1.5 μm beam from a silicon CCD.** A total 1.5 μm power of around 25 mW was incident on the CCD array, but about half of this was estimated to be 100 MHz, sub-nJ pulses. The nonlinear response of this technique enables imaging the kHz amplified pulses even when the 100 MHz background contains more average power. In the image above, we note a large hole on the left side which was later attributed to a damaged optics. The profile shown above represents a beam after it passed through the (unpumped) fourth PPLN crystal. The image before the crystal does not exhibit the vertical striation; we attribute those lines to diffraction from the rectangular aperture of the crystal. The full sensor area is shown.

technology, the OPCPA presented a strong need for detailed beam profiling in the SWIR region. As was shown in numerical simulations in Chapter 4 and Appendix A, optimization of the OPA process relies on accurate characterization of the gain process, which manifests in the spatial domain in, for instance, the formation of a doughnut-shaped beam after back-conversion. Due to the high-degree of coupling between spectral, temporal and spatial effects in OPA, detailed diagnostics are critical for optimization. Parametric amplification aside, beam profiling is a necessary part of the construction of any free space laser system: for instance, spatial profiling aids in diagnostic of optical damage, alignment, characterization of beam size, and so forth.

We addressed these shortcomings by using a standard silicon CCD camera to analyze the spatial profile of the 1.5 μm pulses in the OPA chain through two-photon absorption in a silicon sensor. While this technique may not be applicable to lower intensity optical characterization, we found excellent success using the μJ level, ~ 200 ps pulses in the OPCPA. Pulses in this regime of energy and duration delivered a clear two-photon response without causing damage to the detector. Briefly reviewing the literature, we found a similar use of a silicon sensor in the two-photon regime

applied to Shack-Hartmann wavefront sensing [261]. Those authors found slopes of 1.96-1.97 relating optical power and detector counts on a log-log plot over ~ 1 order of magnitude of optical power for a silicon CCD, suggesting that there is a region of optical power for which a true image can be calculated by simply taking the square root of the measured two-photon image. Further characterization of this nonlinear response, and the power range to which it applies, needs to be performed.

In the remainder of this appendix, we offer a few examples of images taken using this technique. We used a compact USB monochrome CCD camera (model CGE-B013-U from Mightex), which incorporates a near-IR enhanced sensor (Sony ICX445ALA). The sensor comprised 1280x960 pixels of $3.75 \times 3.75 \mu\text{m}$ size, resulting in an active imager size of $6.26 \times 5.01 \text{ mm}$. The retail price for this camera is 559 USD. Figure B.1 shows a raw camera image of a $1.5 \mu\text{m}$ beam. The beam was carefully spectrally filtered using a long-pass filter, and it was verified that no $1 \mu\text{m}$ pump light was detected by the sensor. The beam contained an average power of 25 mW. We estimate that half of this power was residual from the 100 MHz seed laser (sub-nJ), while the other half was amplified 1 kHz pulses with about $12 \mu\text{J}$ of energy. This technique enabled separation of the μJ pulses from




Figure B.2: **Raw CCD images (plotted using scaled colors) showing the simultaneous imaging of a $1.5 \mu\text{m}$ signal beam and a $1 \mu\text{m}$ pump beam.** Images were measured in the plane of OPA2, using the seed from OPA1 and the pump for OPA2. We caution future users that optical filters may cause the beams to shift and appear misaligned. The full sensor area is shown in each image.

the high-rep rate background. We had previously measured this beam using a Pyrocam, and were not able to resolve the detailed structure seen in Figure B.1: the hole on the left side of the profile allowed us to diagnose a burn spot on an optic, and the diffraction pattern allowed us to characterize the degree of modal distortion as a result of clipping on the rectangular aperture of the PPLN crystal.

Another useful application of this technique is shown in Figure B.2. Our OPA requires the careful alignment and relative sizing of a 1 μm pump and 1.5 μm signal. By adjusting the relative power of the pump and signal, it is straightforward to achieve comparable signal strengths from both beams on the image array - despite the fact that one beam employs the linear response and the other the two-photon response. Figure B.2 shows the 1.5 and 1 μm beams on their own, and then both impinging on the camera at the same time. As a word of caution, we note that the snell shift from optical filters can cause the beams to spuriously become offset relative to their position without the filters. Care must be taken to ensure that power attenuation does not displace the beams if this technique is to be used for alignment.

As a final application, we use the two-photon beam profiling technique to diagnose the gain properties of an OPA. Sisyfos simulations readily predict the modal evolution of a signal beam as it is amplified: in the early parts of the gain curve, the mode shrinks as a result of spatial gain narrowing. Then, as the gain curve approaches back-conversion, the simulations show the mode size increasing as the pump becomes more depleted in the central part of the beam. Finally, back-conversion sets in, causing the center of the signal mode to be converted back to pump, resulting in a doughnut-shaped mode. Figure B.3 shows the evolution of a beam profile, measured in the same plane after OPA1, under different amounts of amplification. In the top row, a single 5 mm PPLN crystal is used for three different pump powers. In the bottom row, a second, 3 mm crystal is placed after the 5 mm crystal, using residual pump light from the first crystal and increasing the total gain. The modal evolution reflects the pattern seen in simulation, demonstrating that this technique can provide useful diagnostic feedback in OPA optimization.

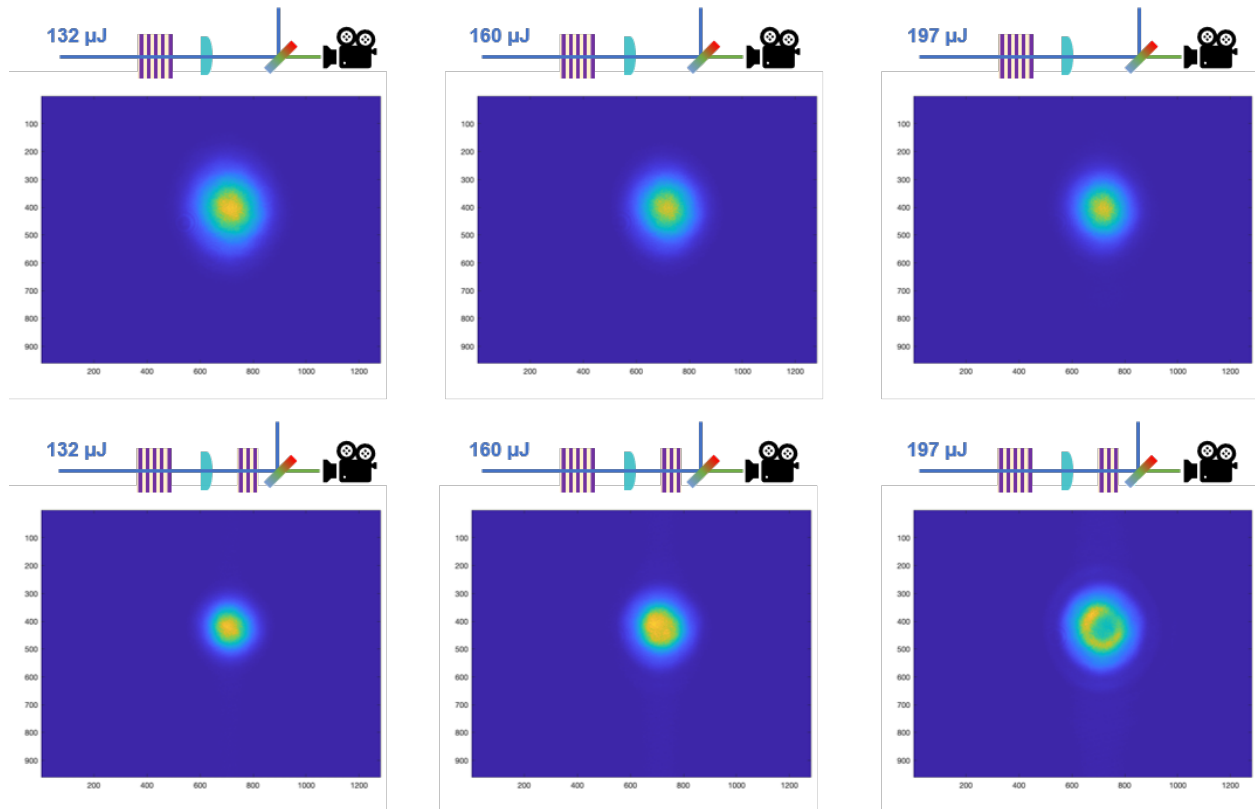


Figure B.3: **Raw CCD images (plotted using scaled colors) showing amplification of a $1.5 \mu\text{m}$ signal beam in OPA1.** Images were measured in the plane of OPA2. Observed reshaping of the signal due to amplification reflects trends seen in Sisyfos simulations. The full sensor area is shown in each image.

Appendix C

Analytical studies of OPA gain scaling with crystal length, pump intensity and effective nonlinearity

The time-dependent coupled wave equations that describe OPA are in general challenging to solve analytically, leading the application of numerical simulation methods such as the Sisyfos code package [100]. Nevertheless, we can derive the spectral gain properties of broadband optical parametric amplifiers in the small signal, undepleted pump regime from the monochromatic, time invariant coupled wave-equations presented in Chapter 2. This is possible by considering the Fourier decomposition of the pulses into individual frequency components. We first calculate the frequency dependent wave-vector mismatch $\Delta k(\omega)$, then plug it in to the Equation 2.64 for the small-signal parametric gain $G(L)$, which is derived in Chapter 2 and reproduced here:

$$G(L) = \frac{I_1(L)}{I_1(0)} = 1 + \left[\frac{\Gamma}{g} \sinh(gL) \right]^2. \quad (\text{C.1})$$

where $I_1(0)$ is the incident signal intensity at $z = 0$, and we have defined the small-signal gain g as:

$$g = \sqrt{\Gamma^2 - \frac{\Delta k^2}{4}}, \quad (\text{C.2})$$

where

$$\Gamma^2 = \frac{2d_{eff}^2 \omega_1 \omega_2}{c^3 \epsilon_0 n_1 n_2 n_3} I_3, \quad (\text{C.3})$$

and I_3 is the pump intensity.

For the calculations presented in this appendix, the wave-vector mismatch was calculated using the temperature dependent Sellmeier equations from the erratum of reference [117]. All three

beams have extraordinary polarization, which is necessary to access the highest nonlinear coefficient d_{33} .

We first plot the small signal parametric gain of PPLN and overlay the full spectrum from the front-end EDFA. This is particularly relevant for the first OPA stage, which has a typical gain of around 3.6×10^4 or 46 dB. We have addressed the problem of gain narrowing in this stage by tilting the crystal, which has allowed us to use an even longer 5 mm crystal.




Figure C.1: **Calculated small-signal parametric gain of a $30.1 \mu\text{m}$ poling period PPLN at 80 C pumped at $4 \text{ GW}/\text{cm}^2$, with front-end seed spectrum (arb. amplitude) overlaid.** We note that while the FWHM of the parametric gain broadens with a shorter crystal, the magnitude of the gain decreases dramatically - owing to the exponential nature of OPA. Using undepleted pump approximation (small-signal), we first calculate the frequency dependent wave-vector mismatch $\Delta k(\omega)$, then plug it into Equation 2.64 for parametric gain.

In the remaining three plots, we consider how small signal parametric gain scales with crystal length, d_{eff} and pump intensity. The top plot shows the maximum parametric gain, which occurs for the central frequency where Δk is minimized. The bottom plot shows the normalized spectral gain. We see that using a short crystal is an extremely effective way at broadening the amplification bandwidth, stemming from the fact that wave-vector mismatch is determined by multiplying Δk by length - although we caution that the maximum gain value for crystals shorter than 1 mm was in the few 10's.

Perhaps more interesting is the observation that, for the same crystal length, increasing either d_{eff} or the pump intensity will also increase the gain bandwidth - despite the fact that the wave-vector mismatch hasn't changed. This effect is analyzed in the main text.




Figure C.2: Dependence of the small signal parametric gain on crystal length for a 30.1 μm poling period PPLN at 100 C, with a $d_{\text{eff}} = 14.8 \text{ pm/V}$ and a pump intensity of 4 GW/cm^2 . The gain spectra (bottom) are normalized at each crystal length.




Figure C.3: Dependence of the small signal parametric gain on d_{eff} for a 3 mm long, 30.1 μm poling period PPLN at 100 C, with a pump intensity of 4 GW/cm^2 . The gain spectra (bottom) are normalized for each value of d_{eff} .






Figure C.4: Dependence of the small signal parametric gain on pump intensity for a 3 mm long 30.1 μm poling period PPLN at 100 C, with a $d_{\text{eff}} = 14.8$ pm/V. The gain spectra (bottom) are normalized at each pump intensity.

Appendix D

Technical drawings and user notes for the HHG source assembly





Label	Description	Qty.	Notes
(a)	Bolt, M4 x 40 mm	4	Requires M4 to M6 helicoil insert to mount to pedestal. Optimal performance is achieved by locking the helicoil insert in place with, for instance, Loctite.
(b)	Threaded 1/2" insert	2	Tight tolerance between insert and bore. Ensure collinearity of insert and bore while threading. Finger tight is enough for vacuum seal. Galling can occur; if so, use soft-jawed pliers, alcohol as lubricant when unscrewing, remove aluminum deposit on insert, ideally with an abrasive on the lathe, and, if necessary, ream bore. Customized from McMaster part: 95446A663.
(c)	O-ring, -022	1	Material not critical
(d)	Bolt, M3 x 12 mm	6	Bolts required to attach lid. Lid is 0.25" thick.
(e)	Threaded 5/32" insert	1	See notes for (b) to avoid galling. To insert/remove, first remove threaded 1/2" insert, then use a 3/32" straight hex (not ball end) driver to bring the bolt into the hole. Standard bore is 0.75 mm. McMaster part: 95446A638.
(f)	Bolt, M2 x 15 mm	2	Holds the cartridge to the chassis.
(g)	O-ring, -031	1	O-ring has a habit of popping out of place. Silicone holds in place better. More silicone O-rings are on order.
(h)	Threaded hole		Allows for the attachment of various lengths of fiber cartridges.
(i)	Dowel pin hole		2 mm x 10 mm long dowel pins are used to precisely locate the fiber cartridge to the chassis. Dowel pins can easily slip out of the fiber cartridge when installing/removing. Avoid allowing dowel pins to fall into vacuum port.



INCONSISTENT SCALES
Not for dimensions

SXR Cartridge Full Assembly

Author: Drew Morrill
6 Sept 2022




SXR Cartridge

Author: Drew Morrill
6 Sept 2022


- Break all edges
- Tolerance ± 0.003 unless noted
- minimum qty needed: 2

Groove for 1/16" nominal O-ring



DETAIL B
SCALE 6 : 1

Note: bottom surface should be smooth for O-ring mating




BOTTOM VIEW
SCALE 1:2


SXR Cartridge

Author: Drew Morrill
6 Sept 2022

- Break all edges
- Tolerance ± 0.003 unless noted
- minimum qty needed: 2




DETAIL A
SCALE 8 : 1




DETAIL C
SCALE 2 : 1


Note: groove for 1/16" CS ring, internal pressure



DETAIL C
SCALE 6 : 1




SECTION A-A



Note: chamfer to radius of hole

DETAIL E
SCALE 8 : 1





SECTION B-B

concentricity ±0.013

Ø 2.00 press fit dowel pin
ream √ 8


Ø 0.33 + 0.01 / -0.005



Fiber holder:
13 mm fiber length for unsupported 317 um diameter anti-resonant hollow core fiber

Author: Drew Morrill
7 Sept 2022

- Interior holes must be deburred for passage of glass fiber
- Break all edges
- Tolerance ±0.003" unless noted
- Minimum qty needed: 2




1-2 mm diameter is acceptable, preferring smaller diameter centered ± 0.075 mm $[\pm 0.003]$ "

SECTION A-A

Gas conductance:

- 1 mm = 3.4×10^{-4} L/sec
- 2 mm = 2.6×10^{-3} L/sec



Output large plug

Drew Morrill
15 Sept 2022

