

**Non-contact measurements of nanoscale phonon and electron
transport with ultrafast, coherent short-wavelength light**

by

Brendan Gray McBennett

B.S., University of North Carolina at Chapel Hill, 2016

M.S., University of Colorado Boulder, 2021

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

Mahmoud Hussein

Cindy Regal

Longji Cui

McBennett, Brendan Gray (Ph.D., Physics)

Non-contact measurements of nanoscale phonon and electron transport with ultrafast, coherent short-wavelength light

Thesis directed by Professor Margaret Murnane

The miniaturization of technology over the past few decades has generated enormous demand for experimental techniques capable of resolving material dynamics at ever shorter length and time scales. Ultrafast, coherent ultraviolet light provides a window into microscopic dynamics, combining femtosecond pulse durations, short wavelengths and tunable sensitivity to electronic and thermal processes. This thesis furthers the development of novel ultraviolet metrology tools and their application to nanoscale phonon and electron transport, where confinement and nonequilibrium conditions give rise to surprising new phenomena. Using extreme ultraviolet scatterometry, we extract the thermal and elastic properties of a 3D phononic crystal metalattice, nanostructured on $\ll 100$ nm length scales, and combine the results with previous experiments and atomistic simulations to propose a new effective description of highly-confined heat flow in nanostructured semiconductors. An interlude discusses the challenges associated with contact- and fabrication-based approaches to studying transport phenomena in the context of the wider array of materials appearing in modern nanotechnology. The final chapter introduces a new non-contact deep-ultraviolet (6.3 eV) transient grating experiment, capable of investigating nanoscale phonon and electron transport in ultrawide-bandgap materials at femtosecond timescales. This new technique bridges previous laboratory and facility-scale capabilities and provides new opportunities for studying emergent nanoscale transport phenomena of relevance to next-generation energy and semiconductor technologies.

Dedication

I would like to dedicate this thesis to my soon-to-be wife Emma Raszmann and to the mentors, teachers and role models who first sparked my love of learning and interest in math and science, including Daniel Harris, Irene Bouw, Alan Letarte and Nabeel Kandah, and, from the very earliest age, my parents Amy and Larry McBennett.

Acknowledgements

From start to finish in graduate school, I have been fortunate to be surrounded by outstanding mentors and friends. This begins with Josh, who first showed me the lab and who I've had the great privilege to work with all the way through to the present day. I am also very thankful for the opportunity to work with Albert, whose enthusiasm and insights have vastly enhanced our connection to theory. I am grateful to Begoña, Travis, Kevin, Nathan and Quynh for getting me started in lab, and to Nico, Mahmoud and Hossein for detailed discussions on phonon transport theory. Curtis, who helped me with several scary chemicals, and Tomoko, Jeff and Weilun who taught me lithography, were indispensable to the nanofabrication results, while fruitful collaborations with Yuka, Na, Nick, Tika, Anya, Sadegh, Justin and Sam provided crucial insight into a variety of complex samples. I would also like to thank Longji for letting me teach my first lecture, Alex for lively discussions at my first conference, and Sinéad and Drew for delightful Skull King lunches.

Margaret and Henry have been outstanding advisors, both because of their vast technical expertise and their welcoming and collaborative group culture. The ideas and results in this thesis originate from the phenomenal research environment which they and all the JILA staff have created. Emma and Teddy, I am excited to see where your research leads and look forward to the remainder of our time together at JILA. I would also like to thank SPIE and Siemens EDA for financial support. Visits with my parents Amy and Larry and brother Chris, calls with Melanie and trips with Will, Patrick, Kevin, Sadie and Clement provided many things to look forward to these past several years. Finally, the excitement of building a future together with Emma Raszmann has made this a joyous chapter in life. Thank you all for your love and support.

Contents

Chapter

1	Introduction	1
2	Ultrafast, coherent short-wavelength light sources	8
2.1	Ultrafast lasers across the electromagnetic spectrum	9
2.1.1	Kerr lens mode-locking	9
2.1.2	Chirped pulse amplification	10
2.1.3	Facility-scale ultrafast light sources	12
2.1.4	Nonlinear optics in the perturbative regime	12
2.1.5	High harmonic generation	15
2.2	Pump-probe experiments with ultrafast lasers	17
2.3	Material response of electrons, phonons and spins	17
2.4	Extreme ultraviolet dynamic scatterometry experiment	20
3	Nanoscale phonon transport	24
3.1	Phonon dispersion	27
3.2	Heat conduction in bulk materials from <i>ab initio</i>	28
3.3	Nanoscale thermal effects	31
3.3.1	Broken translation symmetry and boundary effects	32
3.3.2	Nonlocality, Fourier's law and the relaxation time approximation	34
3.3.3	Phonon hydrodynamics	37

3.3.4	Phonon coherence and wave effects	39
3.4	Phonon thermal transport from nanoscale transducer arrays	39
4	Acoustic and thermal behaviors of <100 nm 3D silicon phononic crystal metalattices	44
4.1	Overview of phononic crystals	44
4.2	3D silicon metalattice sample	49
4.3	Extracting metalattice pore geometry and porosity via electron tomography	50
4.4	Nondestructive EUV measurements of metalattice porosity, elastic properties and thermal conductivity	54
4.4.1	EUV scatterometry dataset	56
4.4.2	Elastic theory and surface acoustic waves	61
4.4.3	Finite element model of metalattice unit cell	64
4.4.4	Finite element mode of full sample	66
4.4.5	EUV diffraction modeling	75
4.4.6	Frequency-domain analysis of porosity and elastic properties	79
4.4.7	Time-domain analysis of thermal conductivity	82
4.4.8	Error bar calculations	85
4.5	Phonon hydrodynamic model of heat flow in highly-confined nanosystems	89
5	Nanofabricated metallic gratings in the EUV scatterometry experiment	99
5.1	Electron-beam lithography process	100
5.2	Amorphous carbon at metal-diamond interfaces	108
6	Tabletop femtosecond deep ultraviolet transient grating experiment	111
6.1	Overview of the transient grating technique	111
6.2	DUV transient grating beamline	118
6.2.1	Pump beamline	119
6.2.2	Probe beamline	127

6.2.3	Electronics and data acquisition	132
6.2.4	Beamline diagnostics	135
6.2.5	Carbon deposition and parasitic heterodyning	136
6.3	Validation using surface acoustic waves launched in gold films	139
6.4	Probing carrier diffusion in diamond at <300 nm TG period	141
6.5	Probing nanoscale phonon transport in silicon	147
7	Conclusions and outlook	152
	Bibliography	155
	Appendix	
A	Code repositories	173
B	EUV scatterometry and DUV transient grating experimental documentation	174
C	B1B35 Wyvern alignment procedure	175
D	Electron-beam lithography recipe	182
E	Electron-beam lithography chemical and instrument list	185

Tables

Table

4.1	Experimental EUV dataset for the silicon metalattice sample	58
4.2	Interpolation expressions from unit cell models relating porosity (ϕ), Young's modulus and Poisson's ratio.	66
4.3	Boundary conditions for thermal finite element simulations	74
4.4	Boundary conditions for acoustic finite element simulations	74
4.5	Finite element model material parameters	76
4.6	Fitted thermal conductivity of the silicon metalattice	83
E.1	Data on chemicals used in electron beam lithography process	185
E.2	Information on instruments used in the electron beam lithography process.	185

Figures

Figure

1.1	Manipulating energy flow from the macroscale to the nanoscale	3
1.2	Examples of nanostructuring for tailored material properties	6
2.1	Ultrafast lasers measure microscopic dynamics	11
2.2	Sum frequency generation and high harmonic generation	14
2.3	Material dynamics after excitation by an ultrafast laser	19
2.4	Extreme ultraviolet dynamic scatterometry experiment	22
3.1	Phonon dispersion and bulk <i>ab initio</i> thermal conductivity calculations	29
3.2	Phonon thermal transport in a thin film	33
3.3	Nanoscale thermal transport frameworks	35
3.4	Close-packing effect and thermal transport channeling	41
4.1	Overview of phononic crystals	46
4.2	Coherent and incoherent effects in phononic crystals	47
4.3	3D silicon phononic crystal “metalattices”	51
4.4	Electron tomography reconstruction of the 3D silicon metalattice geometry	55
4.5	EUV scatterometry measurements of the 3D silicon metalattice	57
4.6	Dynamic EUV scatterometry dataset for the silicon metalattice sample	60
4.7	Finite element model of metalattice unit cell	67
4.8	Full sample finite element model	69

4.9	SAW-likeness analysis	72
4.10	Metalattice porosity, thickness and elastic properties extracted via EUV scatterometry	81
4.11	Thermal conductivity vs. fit start time and error calculation	84
4.12	Fourier finite element model fit quality for silicon and metalattice samples	84
4.13	Universal thermal transport behavior in highly-confined nanostructured silicon systems	93
4.14	Atomistic and phonon hydrodynamic heat flux profiles in highly-confined geometries	97
5.1	Electron-beam lithography process overview	101
5.2	Characterization of nanoscale nickel gratings fabricated with EBL on silicon substrates	104
5.3	Amorphous carbon layer at metal-diamond interfaces after electron-beam lithography	109
6.1	Transient grating technique overview	113
6.2	Previous transient grating measurements of phonon thermal transport	117
6.3	Deep ultraviolet transient grating experiment pump beamline	121
6.4	Deep ultraviolet transient grating experiment probe beamline	128
6.5	Deep ultraviolet transient grating beamline data acquisition	134
6.6	Deep ultraviolet transient grating beamline diagnostics	137
6.7	Deep ultraviolet transient grating acoustic measurements and validation	140
6.8	Nanoscale carrier diffusion in single-crystal diamond	143
6.9	Deep ultraviolet transient grating measurements of silicon membranes	149

Chapter 1

Introduction

The flow of energy from power plant to transistor involves some of the most complicated machinery ever built. Yet as seamless as they appear, modern energy systems face the same age-old challenges of energy conversion, transport and storage as in the era of whale oil and steam power. Beginning with electricity generation and distribution, a large fraction of useful energy is lost as heat in route to its destination [1]. Five percent of electricity generated in the U.S. dissipates in the transmission lines, an often-cited application for high-temperature superconductivity [2]. Silicon solar cells only convert a quarter of the sun's energy to electricity, motivating massive investments in perovskite and multilayer solar technologies [3]. Decades of fusion research seek to provide an alternative to the environmental degradations and inefficiencies of coal, natural gas and nuclear fission, while ongoing improvements to electric vehicle batteries seek to link the transportation sector to the electrical grid as an alternative to oil. Our desire to harness energy at ever larger scales and higher efficiencies is likely the single largest driver behind scientific research in general.

One way to make an energy-consuming machine more efficient is to make it as small as possible. This is the approach of the semiconductor industry, driven by consumer demand for home and pocket-sized computation, and encapsulated by the famous Moore's law shown in Fig. 1.1(b). While the pace of miniaturization today seems as fast as ever, much has changed since the days of room-sized computers. Devices that could once be inspected with an optical microscope now require advanced metrology techniques. Two-dimensional chip designs are giving way to three-dimensional architectures, as the size of each transistor becomes increasingly limited by the size of the atom.

And while the number of transistors per chip continues to double every two years, the clock speed, a measure of the number of operations per second, has plateaued. If computers were run at clock speeds significantly faster than the several GHz used in many laptops today, they would catch on fire [4]. There simply is not enough time to evacuate the heat.

In electronics and electric power generation, heat is an unwanted byproduct. Once energy dissipates into heat, it is challenging to recover, and the accumulation of heat can impose constraints on device functionality. Unlike electricity, heat lacks charge, complicating its manipulation by electric and magnetic fields. Unlike light, heat prefers to spread out isotopically and inconveniently, rather than travel as a ray. Worst of all, heat moves slowly. While electrons can travel at a respectable one percent of the speed of light in most materials, thermal vibrations propagate about a thousand times slower, complicating heat evacuation in devices and the application of thermal energy towards long-distance communication or computing.

As thermal constraints increasingly limit both energy and computing technologies, new research into the atomic vibrations, or phonons, responsible for heat flow in crystalline semiconductors has shown that it is possible to do more with heat than was previously believed. New approaches to thermoelectrics, an old and relatively inefficient technology, wherein a steep thermal gradient drives a current, may lead to more widespread applications [5]. Thermal analogs of electronic devices, including thermal diodes, where a specialized material permits the flow of heat forwards, but not backwards, are sparking interest [6]. Many of these technologies exploit the fact that heat-carrying phonons propagate in the same wavelike manner as audible sound, with a wavelength on the order of a million times shorter. Just as audible sound can be cloaked and amplified in underground caverns and musical instruments, researchers have increasingly come to realize that heat can be manipulated, to a degree, at extremely small length and time scales [7].

Devices designed to manipulate heat are often smaller than those used to control photons and electrons. In crystalline solids, phonons have wavelengths of approximately one nanometer, or one billionth of a meter. This is roughly the length of ten atoms lined up in a row, an incredibly small distance even for a transistor. Another relevant length scale for a phonon is its mean free




Figure 1.1: Manipulating energy flow from the macroscale to the nanoscale: (a) Flow of energy to various sectors of the U.S. economy in 2022, illustrating the challenges of energy conversion, transport and storage. A quad is equal to approximately a billion billion joules. (b) Historically, increasingly compact computers have relied on the miniaturization and acceleration of the underlying transistors. While the number of transistors per chip continues to double about every two years, clock speeds leveled off in the early 2000s due to thermal limitations. Figure adapted from [1, 4].

path, the expected distance it will travel before colliding with another phonon, material defect or boundary. A phonon's mean free path is typically much longer than its wavelength, up to and above a micron for some thermal phonons. Structuring a material on this scale causes a significant fraction of phonon scattering to occur at boundaries and interfaces, changing the material thermal properties. Precisely structuring a material on even shorter scales can cause phonons to interfere coherently, granting even more delicate control over heat flow.

The study of matter at nanometer to micron scales is the domain of nanotechnology, a field that took off in the late twentieth century with the advent of new fabrication technologies. A crucial early insight was that objects consisting of tens to many millions of atoms are not just trivial aggregations of fundamental particles or miniaturized versions of bulk materials, but instead exhibit entirely new behaviors otherwise unattainable in nature [8]. Since technology is fundamentally limited by the materials available on earth, nanotechnology presented a new avenue of materials discovery alongside traditional chemical synthesis. Fig. 1.2(a) provides a recent example, where a nanoscale scaffolding of hollow metal glass is observed to crack under compression, until the walls are made sufficiently thin and it begins to deform like a sponge [9]. An even smaller system designed to modify thermal behavior is shown in Fig. 1.2(b), where heat is predicted to flow through nanometer-scale channels between silicon pores in a Poiseuille-like manner analogous to a fluid [10].

The brute force approach to nanotechnology is to create a new nanostructured material, test quickly whether it has the desired properties, and if not, throw it out, arbitrarily modify the recipe and start over. Fortunately, it is increasingly possible to circumvent this tedious process by measuring the dynamics of electrons, phonons and spins inside nanostructured materials, not only informing future designs, but giving fundamental insight into how these particles behave in very small groups, before the averaging of statistical mechanics fully applies. Resolving dynamics at nanoscale dimensions is no easy task, since even phonons cover the small distances very quickly. A phonon in silicon moves roughly as fast as the moon orbits the earth. As long as it does not collide with anything, it will cover the entirety of Fig. 1.2(b) from right to left in about 15 picoseconds, or trillionths of a second. This is much too fast for a camera, or even a fast, electronic instrument like

an oscilloscope. Instead, researchers often turn to ultrafast lasers, which provide the combination of spatial and temporal resolution necessary to resolve nanoscale dynamics.

Ultrafast lasers produce a series of light pulses of sub-picosecond, or even sub-femtosecond temporal duration. Like any laser, the photons in an ultrafast laser pulse are collimated, in that they all travel in the same direction, and coherent, in that they all have a synchronized phase. This enables the straightforward manipulation of ultrafast laser light using a variety of diffractive and refractive optics. When directed at a material, an ultrafast laser pulse provides a snapshot of its internal dynamics, averaged over the duration of the pulse. For resolving phonon dynamics, a femtosecond pulse is more than adequate. A common limitation of ultrafast lasers is their relatively long wavelength, since the energy levels in common laser gain media correspond to light with wavelengths in the vicinity of a micron. Such a long wavelength is inadequate to spatially resolve many nanostructured systems, including the example in Fig. 1.2(b). Until recently, accessing coherent light of significantly shorter wavelength required travel to an inflexible, time-constrained billion-dollar facility such as a synchrotron or X-ray free electron laser.

Over the past two decades, new techniques to generate laboratory-scale short-wavelength ultrafast light have emerged and matured. The most well known, high harmonic generation, uses a quantum nonlinear process to convert a number of lower energy photons into a single higher energy photon in the 1-100 nanometer wavelength range [13]. These wavelengths fall in the extreme ultraviolet (EUV) part of the electromagnetic spectrum, which has proven especially well suited for dynamic imaging, semiconductor nanolithography, identifying elemental composition and studying the dynamics of electrons, phonons and spins at the nanoscale. As with EUV light, deep ultraviolet (DUV) light of slightly longer, approximately 200 nm wavelength, holds many advantages over the visible spectrum, and can also be produced in a laboratory at high intensities.

This thesis is dedicated to the creation of ultrafast, coherent EUV and DUV light sources and their application to outstanding questions in nanoscale phonon transport. Chapters 2 and 3 discuss ultrafast short-wavelength light and nanoscale phonon transport, respectively, in more detail. In particular, Section 2.4 introduces JILA's dynamic EUV scatterometry experiment, used




Figure 1.2: Examples of nanostructuring for tailored material properties: (a) Mechanical properties often change at the nanoscale, affected by the high percentage of altered surface chemical bonds. Here, a normally brittle metallic glass nanolattice becomes deformable for sufficiently small wall thicknesses. (b) Crystalline silicon with nanoscale pores can have a thermal conductivity two orders of magnitude below bulk silicon, making it a potential thermoelectric material. (c) Thermoelectrics, (b) thermal diodes and (c) acoustic and thermal cloaking, together with many other emerging technologies, have the potential to revolutionize our ability to manipulate and harness thermal energy. Figure adapted from [6, 9, 10, 11, 12].

in Chapter 4 to investigate the nanostructured silicon “metalattice” sample shown in Fig. 1.2(b). Time-resolved observations of sound and heat propagation through the metalattice, combined with predictions from atomistic simulations, sparked the development of a new model of heat flow in the limit where phonons are confined in all dimensions to distances far less than their bulk mean free path. Chapter 5 discusses the challenges associated with applying nanofabrication-based metrology approaches such as EUV scatterometry to a wider variety of samples, in particular the formation of an anomalous low-density yet diamond-like amorphous carbon layer on diamond substrates during electron-beam lithography. Finally, Chapter 6 introduces a new DUV transient grating experiment constructed over the past three years at JILA. Using the interference of two high-power DUV beams, it is possible to excite a microscopic charge and temperature grating on a bulk sample, providing a non-contact, fabrication-free window into the nanoscale dynamics of wide-bandgap materials, which are transparent and inaccessible to longer laser wavelengths. The ultrafast, short-wavelength metrology tools advanced in this thesis provide new insight into heat flow under highly-confined conditions and have the potential to further the development of effective theories for phonon and electron transport at the nanoscale.

Chapter 2

Ultrafast, coherent short-wavelength light sources

By human standards, microscopic events happen very quickly, regardless of the exact speed at which the atoms and molecules move. This is because the distances involved are so small. Even if an electron were to travel at the speed your fingernails grow, about three millimeters per month, it would only require a quarter second to orbit a hydrogen atom. In reality, microscopic particles zip along much faster than their macroscopic counterparts, in the case of electrons, up to a significant fraction of the speed of light. As illustrated in Fig. 2.1(a), observing most micron-scale processes requires nanosecond temporal resolution, observing a nanoscale process requires picosecond resolution and so on down to the attosecond transitions of electrons between atomic energy levels.

Techniques such as electron and atomic force microscopy have exquisite spatial resolution, capable of resolving individual atoms, but generally require seconds of scanning to produce an image, resulting in a static average of the microscopic movement. Coherent ultrafast lasers, with their femtosecond and even attosecond pulses, have temporal resolution almost a million times shorter than the fastest electronic instruments; however, the most widely available laser gain media emit in the visible or near infrared, around 1 μm wavelength, limiting spatial resolution. This chapter will discuss approaches to produce short-wavelength ultrafast light sources, which with their combined high spatial and temporal resolution can give unique insight into microscopic processes from biology to hard condensed matter.

2.1 Ultrafast lasers across the electromagnetic spectrum

The electromagnetic spectrum shown in Fig. 2.1(b) ranges from long wavelength infrared to short wavelength X-ray radiation, with each region of the spectrum possessing certain advantages from an experimental and technological standpoint. Long-wavelength infrared photons have energies corresponding to many material vibrational modes, leading to their application in Raman spectroscopy, terahertz spectroscopy and Brillouin light scattering experiments [14, 15]. Short-wavelength X-ray radiation can diffract from an atomic lattice, enabling studies of crystal structure and chemical and elemental composition [16, 17]. The intermediate visible spectrum provides a happy medium. Easily and inexpensively manipulated using refractive optics, light with wavelengths ranging from 400 – 700 nm provides the sub-micron spatial resolution of modern optical microscopes. Visible and near-infrared photon energies also correspond to the electron energy level transitions in laser gain media, such as the ubiquitous Nd:YAG and Ti:Sapphire crystals. This makes femtosecond lasers with visible and near-infrared wavelengths the easiest to build and operate at tabletop scale.

2.1.1 Kerr lens mode-locking

The design of an ultrafast laser begins with the requirement that a short laser pulse have a broad frequency bandwidth due to the energy-time uncertainty relation. For this reason, laser gain media with many closely-spaced energy levels, such as titanium-doped sapphire (Ti:Sapphire) are attractive, as they are capable of fluorescing across a broad range of photon energies, in the case of Ti:Sapphire covering wavelengths from 600 to 1050 nm [18]. Experiments in the early 1990s revealed that Ti:Sapphire lasers can operate in a self-mode-locked configuration, capable of producing pulses <10 fs in temporal duration [19, 20, 21]. Self-mode-locked lasers, or oscillators, rely on the nonlinear optical Kerr effect, wherein the electric field, E , associated with an intense Gaussian beam mode causes the gain medium refractive index, n , to vary as $n = n_0 + n_1 E^2$. This in turn causes the self-focusing of pulsed cavity modes in the gain medium due to their high peak intensities [22]. Because the same self-focusing effect is not experienced by continuous cavity modes, it is possible

to design the oscillator to favor pulsed operation by compensating for the dispersion of the pulsed modes using a pair of prisms [23].

2.1.2 Chirped pulse amplification

The nanojoule pulse energies of a Ti:Sapphire oscillator are too weak to drive many of the nonlinear optical processes discussed in the following sections, but can be boosted in intensity by a factor of millions using a technology known as chirped pulse amplification. The amplification process consists of three main steps [24]. First the femtosecond pulse is stretched in time to lower its peak intensity. It then enters a laser cavity, making multiple passes through a Ti:Sapphire crystal excited by a powerful 532 nm pump laser. On each pass, the stretched pulse extracts energy from the crystal, before exiting the cavity and being compressed back to femtosecond pulsed duration.

The pulse stretcher and compressor employed in chirped pulse amplification utilize customized diffraction gratings which take advantage of the large bandwidth of the short oscillator pulses. Because each wavelength in the oscillator spectrum will diffract from a grating at a different angle, it is possible through precise alignment of the stretcher to retroreflect a broad bandwidth pulse back on itself in such a way that the wavelengths arrive at different times. This condition is known as temporal chirp. After passing through the stretcher, the oscillator pulse duration exceeds 100 picoseconds [25]. A long pulse duration is necessary in the amplifier cavity to prevent damage to the Ti:Sapphire laser gain medium and cavity optics.

Once inside the amplifier cavity, the stretched oscillator pulse, or seed, passes multiple times through a Ti:Sapphire crystal excited by a high power, 532 nm, ~ 100 ns pump laser [26]. The crystal must be cooled to cryogenic temperatures to prevent accumulated thermal energy from the pump laser from spatially modifying its refractive index, which would act as a dangerous thermal lens for the amplified beam. An electro-optical device known as a Pockels cell restricts the amplification process to a subset of the oscillator pulses, which lowers the repetition rate of the amplified beam and boosts the pulse energy. By rapidly switching on and off a high voltage applied across a birefringent crystal, the Pockels cell can selectively rotate the polarization of pulses spaced nanoseconds apart,




Figure 2.1: Ultrafast lasers measure microscopic dynamics: (a) As the size of an object or particle shrinks, its dynamics occur faster, creating a need for metrology tools with both high spatial and temporal resolution. (b) The ultraviolet (UV) and extreme ultraviolet (EUV) regions of the electromagnetic spectrum offer unique insights into the fast dynamics of phonons, electrons and spins. (c) The availability of coherent pulsed light sources with femtosecond temporal resolution has increased dramatically over the past three decades. However, aside from facility-scale X-ray free electron lasers (XFELs) and synchrotrons, there are few sources offering coherent, femtosecond light substantially below the visible diffraction limit. Perturbative nonlinear optical processes such as second and third harmonic generation (SHG, THG) and optical parametric amplification (OPA) allow limited conversion of the near-infrared wavelengths produced by common laser gain media, while high harmonic generation (HHG) offers a laboratory-scale route to much shorter wavelengths. Figure images from KMLabs, SPIE, Encyclopedia Britannica, CDC, and Christophe Dang Ngoc Chan and Horst Frank under CC BY-SA 3.0.

allowing a particular pulse to pass through a polarizer optic while neighboring pulses are reflected. In regenerative amplifiers like the one in B1B35, the Pockels cell allows a seed pulse to be amplified over a certain number of passes through the crystal, before rotating its polarization and releasing it from the cavity.

The compressor reverses the temporal chirp imposed by the stretcher, producing a femtosecond pulse train of millijoule pulse energy. When focused, the peak power of such a laser pulse is enough to ionize atoms, including the noble gases used in the high harmonic generation process described below. As the technology of ultrafast tabletop lasers advances, designs based on optical fibers and alternative gain media, with reduced alignment complexity, are becoming increasingly commercially available, furthering applications in materials science and industry [27].

2.1.3 Facility-scale ultrafast light sources

Alongside tabletop visible and near-infrared femtosecond lasers, fast, coherent light of EUV and X-ray wavelengths is available at a handful of specialized facilities worldwide. These light sources come in two types, X-ray free electron lasers (XFELs) and synchrotrons, both of which harness the radiation produced by electrons accelerated to relativistic speeds. The process by which XFELs and synchrotrons turn this radiation into coherent pulse trains is covered in a variety of sources [28, 29, 30]. While synchrotrons are generally limited to picosecond pulse durations, XFELs can achieve pulse durations comparable to ultrafast chirped pulse amplifiers [31]. As an example, the FERMI XFEL in Trieste, Italy can generate 20 – 80 nm radiation in $<30 \mu\text{J}$ pulses ~ 100 fs in duration [32]. While both synchrotrons and XFELs have the pulse energy and tunability necessary to investigate many of the dynamics presented in Fig. 2.1(a), their extremely high cost limits their availability to only a tiny fraction of the projects to which they could be applied.

2.1.4 Nonlinear optics in the perturbative regime

A more practical route to short wavelength ultrafast lasers involves the use of nonlinear optical processes to change the wavelength of light produced by the chirped pulse amplification process

described in Section 2.1.2. Nonlinear optics exploits the nonlinear response of the dipole moment, or polarization (P) induced in a material in response to an applied electric field,

$$P(t) = \epsilon_0[\chi^{(1)}E(t) + \chi^{(2)}E^2(t) + \chi^{(3)}E^3(t)], \quad (2.1)$$

where ϵ_0 is the vacuum permittivity, $\chi^{(1)}$ is the linear susceptibility and $\chi^{(2)}$ and $\chi^{(3)}$ are the second and third order nonlinear susceptibilities, respectively. From Eq. 2.1, it is evident that for a weak applied electric field, the relationship between $E(t)$ and $P(t)$ is approximately linear. However, for strong electric fields in select materials, the nonlinear terms become relevant. In the case of electromagnetic radiation, where the electric field oscillates with a fundamental frequency ω as $E(t) = E_0e^{-i\omega t}$, the polarization can exhibit higher order oscillations with integer multiples of the fundamental frequency. These oscillations in the induced dipole moment emit radiation of the corresponding harmonic frequency.

The process of second harmonic generation, or frequency doubling was demonstrated shortly after the invention of the laser itself, as shown in Fig. 2.2(a) [36, 33]. At the time, crystalline quartz served as the nonlinear medium for the 700 nm fundamental wavelength of a ruby laser. In the photon picture illustrated in Fig 2.2(b), second harmonic generation involves the combination of two 700 nm photons into a single 350 nm photon such that energy is conserved. A modern experiment would more typically use a beta barium borate (BBO) crystal and the 800 nm output of a Ti:Sapphire amplifier to produce ultrafast 400 nm light. Not long after the first demonstration of second harmonic generation in solid quartz, gases were also shown to exhibit nonlinear optical phenomena [37]. Since gases are spherically symmetric, the even terms in Eq. 2.1 are zero, making gases only capable of generating the third and other odd harmonics of the driving laser.

Second harmonic generation is a special case of sum frequency generation, where two photons of differing energy are added to produce a third. This is a perturbative nonlinear optical process, which can be treated quantum mechanically by considering the nonlinearities as small corrections to the overall Hamiltonian. Along with similar processes, such as difference frequency generation,




Figure 2.2: Sum frequency generation and high harmonic generation: (a) Second harmonic generation, a special case of sum frequency generation, was demonstrated shortly after the invention of the laser. In this image, a small portion of the fundamental 700 nm beam is converted into 350 nm light due to the second-order nonlinear susceptibility of the conversion medium. (b) Sum frequency generation involves the annihilation of two photons to produce a third photon equal to their combined energy. (c) A pulsed laser of high peak intensity focused in a neutral gas generates a plateau of higher harmonics, whose wavelengths can vary from deep ultraviolet to soft X-ray depending on the gas ionization energy and fundamental laser wavelength. (d) The classical three-step model of high harmonic generation begins when the high intensity driving laser distorts the Coulomb potential of the gas nuclei, allowing electrons to tunnel into the continuum. Once in the continuum, the electrons gain kinetic energy over a half oscillation of the electric field of the driving laser. Subsequent electron recombination events are accompanied by the release of high energy photons. Figure adapted from [13, 33, 34, 35].

optical parametric amplification and third order processes, it is possible to produce a variety of photon energies within a few multiples of the fundamental, as shown in Fig. 2.1(c) [34]. However, the efficiency of perturbative nonlinear optical processes decreases rapidly with harmonic order. Along with the ultraviolet and deep-ultraviolet absorption edges of most nonlinear crystals, this renders such approaches incapable of producing light of very short wavelength [38].

2.1.5 High harmonic generation

Experiments in the late 1980s revealed that by focusing an intense driving laser into a gas, it is possible to obtain a plateau of higher harmonics, illustrated in Fig. 2.2(c), with conversion efficiencies far exceeding the perturbative predictions [39, 40]. High harmonic generation (HHG), as it became known, is a highly nonlinear process outside the perturbative framework of Eq. 2.1 [41]. HHG can be explained classically by a three-step model, where high intensity electromagnetic radiation first distorts the Coulomb potential of the neutral gas atoms, allowing electrons to tunnel into the continuum [42]. Once free from the atoms, the electrons accelerate away from their parent ions in the electric field of the driving laser. When the oscillating electric field reverses sign, the electrons reverse course and some recombine with their parent ions. These electrons have a high kinetic energy due to their acceleration in the laser field, and their recombination is accompanied by the release of high energy photons. The three-step model is illustrated in Fig. 2.2(d). More rigorous quantum mechanical treatments of HHG provide additional insights and corrections [43, 44].

The maximum or cutoff energy of the high harmonics relates to the driving laser wavelength as

$$E_{\text{cutoff}} \approx I_p + 3.17U_p \quad (2.2)$$

where U_p is the ponderomotive, or average energy of an electron in the driving laser field and I_p is the atomic ionization potential. Because the ponderomotive energy can be shown to increase as the square of driving laser wavelength, the generation of soft X-ray light of <10 nm wavelength

via HHG requires high intensity coherent driving lasers of $>1\ \mu\text{m}$ wavelength, itself an ongoing field of research [27, 45]. An ultrafast, coherent soft X-ray light source from HHG could enable the imaging of biological processes in the high contrast “water window” between 290 and 540 eV, where carbon absorbs electromagnetic radiation but water does not [46].

Practically, the most important consideration in both HHG and traditional perturbative nonlinear optics is that of phase matching [47]. All nonlinear optical processes occur in a medium of finite thickness. If the phases of the higher harmonics produced at each point along the medium interfere destructively, the overall efficiency of the process will be highly limited. Because the refractive index of the nonlinear medium is different for the fundamental light and its harmonics, the problem of phase matching amounts to engineering the phase velocities of the two wavelengths to obtain the desired constructive interference. For sum frequency generation in solid media such as BBO crystals, this can be accomplished using birefringence, which often sets a practical limit on the nonlinear processes achievable in a given crystal [38]. For HHG, the harmonics of interest are usually in the extreme ultraviolet, where their refractive index in the gas is approximately one. Meanwhile, the refractive index experienced by the infrared driving laser increases with the atomic density of the gas and decreases with the density of ionized electrons. It further depends on the experimental geometry. For example, confining the HHG process to a gas-filled waveguide generates an additional negative contribution to group velocity [48]. By tuning the gas pressure and geometry, the positive and negative dispersion effects can be made to cancel, allowing for phase-matched generation of extreme ultraviolet and soft X-ray light [49].

The results in Chapters 4 and 6 of this thesis utilize both HHG and perturbative nonlinear optics. The 30 nm light used to probe nanoscale phonon transport in silicon in Chapter 4 originates from an HHG process in argon, while the 200 nm light to excite charge and energy distributions in wide-bandgap semiconductors in Chapter 6 originates from a multi-stage perturbative process in BBO crystals. While established technologies such as excimer lasers can produce coherent light at 200 nm, they are limited to nanosecond pulse durations. Generating intense ultrafast light in the deep ultraviolet region of the electromagnetic spectrum is a particular challenge addressed in this

thesis.

2.2 Pump-probe experiments with ultrafast lasers

An ultrafast laser will deposit energy into a sample only during its femtosecond pulse duration, generating a highly nonequilibrium material state. By monitoring the subsequent material evolution, it is possible to gain important insights into the fast dynamics of electrons, phonons and spins, three key constituent components of condensed matter systems. In practice, this involves taking snapshots of the dynamics at various time delays after the initial excitation. However, because the dynamics of interest occur over a few femtoseconds to nanoseconds, each frame must be captured over an extremely short integration time to avoid temporally averaging and blurring the dynamics. Most scientific cameras require millisecond integration times, making them far too slow for this purpose. Instead, in what is known as a pump-probe experiment, a portion of the ultrafast laser pulse is separated and delayed relative to the main beam [50]. At the speed of light, the introduction of a nanosecond delay requires a roughly one foot increase of the optical path length, making this a relatively straightforward mechanical process. The weaker “probe” pulse arrives some time after the main “pump” beam excites the sample and integrates the dynamics only over its femtosecond pulse duration. By repeating this process at a series of pump-probe delay times, it is possible to construct a movie of the fast dynamics. The material property of interest, such as temperature, must be inferred from changes to the probe’s reflectivity, phase or diffraction efficiency as a function of delay time. As discussed in the following section, this can pose particular modeling challenges depending on the probe wavelength.

2.3 Material response of electrons, phonons and spins

While mid-infrared photons can directly excite vibrational modes of correspondingly lower energy, the shorter-wavelength light used in most ultrafast laser experiments is absorbed by electrons, which rapidly transition to higher energy states. As shown in Fig. 2.3(a), this results in a highly nonequilibrium electron distribution, which because of its complexity has a very low entropy.

Collisions between electrons, occurring over tens of femtoseconds in most materials, rapidly relax this distribution to the well-known Fermi-Dirac distribution of statistical mechanics

$$n(\varepsilon) = [e^{(\varepsilon-\mu)/k_B T} + 1]^{-1} \quad (2.3)$$

where ε is the electron energy, μ is the chemical potential, T is the electron temperature and k_B is the Boltzmann constant. The temperature determines the likelihood of finding electrons above the Fermi level and can reach thousands of Kelvin after an ultrafast laser excitation, due to the low electron heat capacity. The heated electrons immediately begin to exchange energy with the lattice and spin systems, as show in Fig. 2.3(b), a process often modeled phenomenologically with a so-called three-temperature model, which incorporates the heat capacities and coupling constants of the three systems [51]. Energy exchange with the lattice is a relatively slow process, as shown in Fig. 2.3(c). However, because the heat capacity of the lattice is large, after it equilibrates with the electrons and spins over hundreds of femtoseconds to picoseconds, the common temperature is significantly lower.

The probe laser wavelength makes a large difference in terms of an experiment's relative sensitivity to the electron, phonon and spin dynamics described above. For example, while lattice thermal expansion causes a density reduction, which decreases the refractive index in most materials, refractive index changes due to electrons excited near the Fermi level can vary by material and wavelength according to the band structure [54]. An important advantage in using EUV light to study phonon transport is that it can be tuned away from core level electronic resonances to be relatively insensitive to electron dynamics. As illustrated in Fig. 2.3(a), after approximately 100 fs, the excited electrons and holes relax to a Fermi-Dirac distribution of elevated temperature. If the electron temperature is approximately 5000 Kelvin, the effects of this distribution only extend by approximately 0.5 eV on either side of the Fermi energy, below which almost all states are filled and above which few electrons are excited. The energies of visible photons fall near this range (green arrow), making a visible probe very sensitive to changes in the electron population near the Fermi




Figure 2.3: Material dynamics after excitation by an ultrafast laser. (a) Initial excitation by an ultrafast laser pulse creates a highly nonequilibrium electron distribution. Through electron-electron scattering over tens of femtoseconds, the electrons rapidly thermalize, relaxing to a Fermi-Dirac distribution with an elevated temperature, before equilibrating with the lattice on a picosecond timescale. Visible light (green arrow) is very sensitive to these electron dynamics, since its photon energy often corresponds to the distances between valence electron states near the Fermi level. Higher energy extreme ultraviolet photons (purple arrow) interact primarily with core electrons, and are mainly sensitive to material density. (b) An incident ultrafast visible or infrared laser pulse first excites electrons to higher energy levels, which then exchange energy with the lattice and spin systems. (c) A three-temperature model is commonly used to model the exchange of energy between electrons, phonons and spins. While electrons and spins equilibrate on femtosecond timescales, their equilibration with the lattice can require several picoseconds. Figure adapted from [51, 52, 53].

energy before and after an ultrafast excitation. EUV photons have energies corresponding to core level electron transitions (purple arrow). Since core levels remain occupied, at least after the electron system equilibrates, EUV probes are only sensitive to electron dynamics near discrete absorption edges corresponding to energy differences between a filled core level and levels near the Fermi energy. When tuned away from absorption edges, EUV probes are instead primarily sensitive to density-induced changes in the refractive index caused by thermal expansion [55, 56]. The following section discusses JILA’s EUV scatterometry experiment, which takes advantage of EUV light’s ability selectively probe phonon dynamics and is the basis for the results on nanostructured silicon presented in Chapter 4.

2.4 Extreme ultraviolet dynamic scatterometry experiment

The extreme ultraviolet scatterometry experiment discussed below has enabled a series of investigations of acoustic and thermal properties of nanostructured semiconductors over the past decade [10, 55, 57, 58, 59, 60, 61, 62, 63, 64, 65, 66, 67, 68]. Since it features prominently in the dissertations of several preceding graduate students, I will cover it in less detail here. Appendix B points to documentation describing the step by step operation of this experiment, indicating the best resources for troubleshooting particular aspects of the experiment, and providing technical detail on points not covered elsewhere.

The block diagram in Fig. 2.4(a) shows the EUV scatterometry experiment, which begins with the 35 fs, 5 kHz pulse train of a Ti:Sapphire amplifier. A discussion of ultrafast lasers and the chirped pulse amplification technology behind the amplifier appears in Section 2.1.2. Shortly after the amplifier, a beamsplitter separates its output into pump and probe components. The pump first passes through a half-wave plate and polarizer, the combination of which allows the pump power to be adjusted between measurements, from only a few μJ to approximately 200 μJ per pulse. It then proceeds to a series of motorized delay stages, with adjust the temporal interval between pump and probe pulse arrival times at the sample with femtosecond precision. This is an extremely small interval of time, the control of which is made possible by the extremely large speed

of light, which travels approximately one foot per nanosecond, or three microns per ten femtoseconds. These delay stages are one unique element of the EUV scatterometry experiment as compared with other tabletop EUV experiments. Since phonon transport occurs on the nanosecond timescale, as opposed to the femtosecond or picosecond timescales of electronic and magnetic phenomena, the EUV scatterometry experiment features delay stages capable of eight nanoseconds of adjustment. The first, or “short” delay stage allows only 600 ps of delay, at higher ~ 10 fs precision, while the second or “long” delay stage allows the remaining delay, at roughly 1 ps precision. Alignment of these delay stages can be challenging but is crucial, since a misaligned delay stage can cause the pump and probe beam overlap on the sample to change over the course of a measurement. The pump beam next proceeds through collimating lenses meant to reduce its spot size and onwards to a set of two chirped mirrors. The chirped mirrors compress the pump pulse to its final 35 fs duration. This could be accomplished by optimizing the compressor grating alignment in the amplifier. However, a highly compressed pulse accumulates nonlinearities when traveling through air due to its high peak intensity, which is a problem in the EUV scatterometry experiment due to the length of the delay stages. By locating the final pulse compression stage shortly before the vacuum system, these nonlinearities, or B-Integral, are minimized. The pump finally passes through an optical chopper, which is a rotating wheel designed to block half of the pump pulses, enabling a differential measurement of the experimental signal with and without the pump light present. The laser repetition rate is sufficiently high that the optical chopper cannot spin fast enough to block every other pump pulse. It instead must block a number of pump pulses, subsequently allowing an equal number of pulses to pass through. The first pulse passing through the optical chopper in each cycle activates a timing circuit, which in turn instructs the charge-coupled device (CCD) camera to begin an exposure, such that alternating CCD camera exposures cover time periods when the pump is either blocked or unblocked, respectively. In the data processing, the subtraction of every other camera frame yields the desired differential measurement.

The probe portion of the amplified beam requires more pulse energy than the pump in order to drive the high harmonic generation process, typically $\gtrsim 500$ μJ . Probe pulses first pass through a




Figure 2.4: Extreme ultraviolet dynamic scatterometry experiment: (a) Block diagram of the experiment. The fundamental 800 nm Ti:Sapphire output is split into pump and probe components, whose relative arrival time at the sample can be adjusted by up to 8 ns via two optical delay stages. The probe is partially converted to EUV via high harmonic generation, with photon energies ranging from 35 – 50 eV. After the pump heats the sample, the EUV probe diffracts from its surface and onto a CCD camera. Adjustment of the delay stages allows for sample dynamics to be monitored with sub-picosecond time resolution and spatial resolution below the visible diffraction limit. (b) EUV diffraction pattern on the CCD camera. Subtraction of exposures with the pump incident on the sample surface (red dashed) and blocked by the optical chopper (black) yields a differential signal (blue), which isolates the effect of the pump from various noise sources.

half-wave plate, which, in combination with the polarizing Brewster-cut vacuum entrance window allows for tunable probe power. Chirped mirrors similar to those used in the pump beamline compress the probe directly before its entrance into vacuum. A lens then focuses the probe through the vacuum entrance window to a focal spot at the entrance to the argon-filled waveguide, or fiber, in which the high harmonic generation process occurs. See Section 2.1.5 for details on high harmonic generation. In the EUV scatterometry experiment, typical parameters for the high harmonic generation process are ~ 30 Torr argon pressure, 5 cm fiber length, 150 μm fiber diameter, and ~ 90 μm probe spot size at the fiber entrance. Exiting the fiber, the probe encounters two aluminum filters of 200 nm thickness, which block the fundamental 800 nm probe wavelength but transmit the EUV harmonics. The transmitted harmonics reflect off a focusing toroidal mirror, then diffract from the nanostructured sample surface and onto the CCD camera. A further set of aluminum filters protects the CCD camera chip from stray pump light.

Fig. 2.4(b) shows an example of the EUV probe diffraction pattern on the CCD camera at a fixed temporal delay between pump and probe. The blue line shows the differential measurement between the pump on (red dashed) and pump off (gray) diffraction patterns, which is only a 0.1 – 1% perturbation. Since drift in the HHG spectrum or other experimental factors can cause shifts exceeding this magnitude over the course of a full data scan, it is important to ensure that each millisecond-scale pump on camera exposure is immediately followed by, and subtracted from, the corresponding pump off camera exposure. In this way, the differential nature of the EUV scatterometry measurements eliminates a significant portion of the experimental noise. Finally, as discussed in Section 2.3, it is important to note that the differential EUV diffraction pattern primarily records the thermal expansion of the nickel nanostructures after excitation by the infrared pump [55]. In contrast, visible lasers, with 1.8 – 3.3 eV photon energies, interact strongly with valence electrons excited by the pump beam and therefore exhibit a complex and material-specific dependence on the index of refraction. This enables the EUV probe to measure thermal dynamics at short timescales, which are obscured by electronic effects in visible measurements.

Chapter 3

Nanoscale phonon transport

Unlike the solar radiation on earth's surface or the convection currents at its core, heat propagates through a solid without the apparent net movement of its constituent particles. In metals, this occurs predominantly via the flow of electrons, and for this reason the electrical and thermal conductivities of metals often increase in proportion. In insulators and semiconductors, electrons are bound to the nuclei, yet thermal conductivity can vary by orders of magnitude between materials. Instead, heat propagates using the same mechanism as sound, through tiny vibrations of atoms about their equilibrium positions.

In a 3D crystal, there exist three normal modes of vibration for each of the periodically spaced atoms [69]. For example, each atom may oscillate sinusoidally about its equilibrium position with the exact opposite phase of its immediate neighbors. Alternatively, neighboring atoms may oscillate in phase, as a wavelike disturbance that propagates through the lattice, analogous to an audience waving their arms in unison at a stadium. Any arbitrary lattice distortion is expressible as the superposition of the many normal modes, where each contributes with a specific amplitude. Quantum mechanics dictates that the energy, and therefore amplitude, of a normal mode take on only discrete values, or quanta. In analogy to photons, these individual quanta of vibrational energy are called phonons; each phonon carries the energy $\hbar\omega$, where ω is its frequency.

While audible sound occurs at frequencies of tens to thousands of hertz, heat-carrying thermal phonons have frequencies exceeding a terahertz. Consequently, while the wavelength of audible sound is on the order of a meter, thermal phonons can have wavelengths less than a nanometer.

Given their short wavelength, thermal phonons, like electrons, can be approximated in many contexts as localized particles. Phonons, however, are far more elusive than electrons. They carry no charge or internal spin, making them unresponsive to electronic and magnetic fields. They have no mass, and their number is not conserved, allowing two phonons to merge into one, or a single phonon to split into many. The associated experimental challenges of studying phonon behavior mean that they are less well understood than electrons or photons [7].

As slippery as they may seem, phonons are constrained by the fundamental requirements of thermodynamics and statistical mechanics [70]. A nonequilibrium distribution of phonons, expressible as the number of phonons $n_q(\mathbf{x}, t)$ associated with mode q in the vicinity of position $\mathbf{x}$ and time t , will generate entropy as it relaxes towards the equilibrium Bose-Einstein distribution function,

$$n_q^{\text{eq}} = [e^{\hbar\omega_q/k_B T} - 1]^{-1}, \quad (3.1)$$

where T is the local equilibrium temperature and k_B is the Boltzmann constant. Collisions between phonons must conserve energy, limiting the phase space of their interactions. Phonons also inherit various properties from the normal mode they quantize, including wavevector and group velocity. These considerations facilitated the application of kinetic theory to phonons, a framework originally designed to model the collisions of discrete particles in a liquid or gas. Kinetic theory endows phonons with further, statistical, properties, such as a spatial mean free path, or temporal lifetime, indicating the average distance or time lapse between collision events.

At the macroscopic scale, thermal conduction in solids has long been known to follow Fourier's law,

$$\mathbf{q} = -\kappa \nabla T, \quad (3.2)$$

which states that the heat flux $\mathbf{q}$ and temperature gradient ∇T are related proportionally by the spatially-invariant thermal conductivity κ . For an anisotropic material, the thermal conductivity

appears in Eq. 3.2 as a tensor. Fourier's law often appears alongside a statement of energy conservation

$$\nabla \cdot \mathbf{q} = c_V \frac{\partial T}{\partial t}, \quad (3.3)$$

where c_V is the volumetric specific heat capacity in $\text{J}/(\text{m}^3 \cdot \text{K})$. Combining Fourier's law with Eq. 3.3 results in the heat equation,

$$\frac{\partial T}{\partial t} = \alpha \nabla^2 T, \quad (3.4)$$

where $\alpha = \kappa/c_V$ is the thermal diffusivity. The heat equation models the evolution of the temperature distribution as a diffusive process. The great advance of the twentieth century in understanding heat conduction in solids was the connection established between macroscopic Fourier's law and the microscopic phonon picture. For example, kinetic theory made it possible to express the heat flux inside a crystalline solid as a sum over the different phonon modes of the nonequilibrium distribution,

$$\mathbf{q}(\mathbf{x}, t) = \sum_q \hbar \omega_q \mathbf{v}_q n_q(\mathbf{x}, t), \quad (3.5)$$

where ω_q and $\mathbf{v}_q$ are the frequency and velocity associated with mode q [71]. The thermal conductivity can also be expressed as the sum over microscopic quantities, but its definition is more complex than the heat flux in that it also requires the calculation of statistical phonon properties such as lifetimes and mean free paths. Recently, *ab initio* methods have enabled the calculation of such properties, which has allowed the theoretical prediction of thermal conductivity in new materials, which could previously only be determined empirically.

As discussed in the remainder of this chapter, an exact kinetic theory approach to thermal transport is impractical for systems dominated by nanoscale boundaries, interfaces and highly nonequilibrium statistical phonon distributions. In an infinite bulk system, the perfect spatial

symmetry of the lattice greatly simplifies the calculation of phonon properties. Absent this symmetry, it becomes necessary to make other approximations, often to the phonon collision processes or distribution functions. Alternatively, one can turn to atomistic simulations on increasingly powerful computers, bypassing the notion of a phonon entirely, or return to the fundamental wave nature of phonons, with the coherent effects that wavelike interactions entail.

Regardless of the theoretical approach taken, investigations of nanoscale phonon transport reveal behaviors unrealizable macroscopic bulk systems. At the nanoscale, heat can flow as a fluid [72], temperature can oscillate as a wave [73, 74] and a diffusive Fourier’s law, even with a fitted thermal conductivity, proves insufficient to model many systems’ thermal evolution [67]. Experiments are crucial to understanding non-diffusive thermal transport behaviors, both because they establish when it is advantageous to apply a particular theoretical framework and because they sometimes reveal entirely new and unpredicted phenomena, as discussed in Section 3.4 for the case of closely spaced nanoscale heaters. As both semiconductor technology and energy engineering delve deeper into the nanoscale, the degree to which these novel thermal effects will provide opportunities, or impose constraints, remains an open question.

3.1 Phonon dispersion

The normal vibrational modes of atoms arranged in a lattice greatly increase in complexity with the crystal structure. In the simple case of a 1D chain of N atoms with a harmonic potential, the total energy can be expressed as

$$E = \sum_{n=1}^N [u_n^2 + \omega_0^2(u_n - u_{n-1})^2], \quad (3.6)$$

leading to the dispersion relation

$$\omega_q = 2\omega_0 |\sin(\frac{q}{2})|, \quad (3.7)$$

where q is the wave number, which takes on the values $2\pi/N \times (0, \pm 1, \pm 2, \dots)$. In the case of

a 1D lattice with basis vector $\mathbf{a}$, the wavevector $\mathbf{q}$ is related to wave number as $\mathbf{q} = q\hat{\mathbf{a}}/|\mathbf{a}|$. The group velocity $\mathbf{v}_q$ of a particular mode is the slope of the dispersion, $\partial\omega_q/\partial\mathbf{q}$, and the wavelength is $\lambda = 2\pi/|\mathbf{q}|$. The quantity $\hbar\mathbf{q}$ is often referred to as crystal momentum. It should not be confused with ordinary momentum, which is zero in a solid because the ions experience no net displacement. From Eq. 3.7, it is evident that at low frequencies the dispersion is approximately linear, with modes propagating at a constant sound velocity, while at higher frequencies things become more complicated. Just how much more complicated is illustrated by the phonon dispersion for diamond shown in Fig. 3.1(a) [75]. The multiple branches extending from the origin correspond to the longitudinal and transverse polarizations of phonons propagating in a 3D lattice. At frequencies below approximately one terahertz, the acoustic modes (blue) carry little thermal energy because of their low heat capacity. Phonons of several terahertz frequency (yellow) provide the dominant contribution to thermal transport due to their combination of higher heat capacity and large group velocity. Finally, high frequency optical phonons (green) consist of out-of-phase atomic vibrations whose low group velocity limits their contribution to thermal conduction.

3.2 Heat conduction in bulk materials from *ab initio*

While it is straightforward to solve for the dispersion of a one-dimensional atomic chain subject to a harmonic interatomic potential, this toy model is far removed from the calculation needed to describe heat conduction in a real crystal. For the harmonic Hamiltonian used in Eq. 3.6, the normal modes of oscillation are energy eigenstates, whose amplitudes are fixed in time. Since energy is not redistributed between modes, phonons propagate with infinite mean free paths, which leads to an infinite thermal conductivity. To advance further, is necessary to consider anharmonic corrections to the potential as shown in Fig. 3.1(b). The first anharmonic term, cubic in the atomic displacements, leads to transition rates between modes caused by three-phonon interactions. In a so-called *normal* three-phonon process, wavevector is conserved as $\mathbf{q} + \mathbf{q}' = \mathbf{q}''$, which means that the total crystal momentum,




Figure 3.1: Phonon dispersion and bulk *ab initio* thermal conductivity calculations: (a) Experimental (black diamonds) and theoretical (black curves) phonon dispersion for diamond, defining the spectrum of phonon modes as a function of their wavevector and frequency. The slope of the dispersion corresponds to the phonon group velocity. Low-frequency acoustic phonons (blue) have an approximately linear dispersion and large group velocity, but low heat capacities, limiting their contribution to thermal conductivity. The multiple acoustic branches correspond to longitudinal and transverse polarizations. Phonons of several THz frequency (yellow) generally make the largest contribution to thermal conductivity, while even higher frequency optical phonons (green) generally have a low group velocity, limiting their contribution. (b) Starting from a crystalline unit cell, the interatomic potential energy can be expanded as a series in the atomic displacements (u_n). The coefficients in this expansion, called force constants, are calculated using computationally expensive *ab initio* methods. The second- and third-order force constants determine harmonic and anharmonic phonon properties, respectively. By inserting these phonon properties into the linearized Boltzmann transport equation, it is possible to solve for the thermal conductivity of a bulk material, which exhibits good agreement with experiment. Figure adapted from [75, 76].

$$\mathbf{P}(\mathbf{x}, t) = \sum_{\mathbf{q}} \hbar \mathbf{q} n_{\mathbf{q}}(\mathbf{x}, t), \quad (3.8)$$

where $n_{\mathbf{q}}(\mathbf{x}, t)$ is the phonon distribution function, is also conserved. A constant, nonzero total crystal momentum means a net flux of phonons and heat, even in the absence of a temperature gradient, which again implies an infinite thermal conductivity [71]. However, three-phonon processes sometimes produce a phonon with wavelength smaller than the interatomic spacing. Such phonons are physically equivalent to aliased versions of themselves, with wavevectors differing by a reciprocal lattice vector $\mathbf{R}$. A three-phonon interaction in which crystal momentum is conserved up to a nonzero reciprocal lattice vector is known as an Umklapp processes and modifies the crystal momentum as $\mathbf{q} + \mathbf{q}' = \mathbf{q}'' + \mathbf{R}$, accounting for the dissipation of heat flux necessary to produce a finite thermal conductivity [77].

To quantify thermal conductivity, we would need to know the coefficients, or force constants, for both the harmonic and anharmonic terms in the interatomic potential, as shown in Fig. 3.1(b). The force constants are material-specific quantities and their calculation requires computationally expensive *ab initio* methods such as density functional theory [76]. The harmonic force constants are coefficients in an eigenvalue problem whose solution gives the phonon dispersion. A phonon's frequency, velocity and heat capacity are therefore referred to as harmonic properties. The anharmonic force constants can be used in combination with Fermi's golden rule to determine the three-phonon scattering rates shown in Fig. 3.1(b), where $P_{\mathbf{q}, \mathbf{q}'}^{q''}$ accounts for processes where two phonons combine into one and $P_{\mathbf{q}}^{q', q''}$ accounts for processes where one phonon splits into two [78]. The scattering rates alone are insufficient to determine phonon lifetimes and mean free paths, since they do not indicate whether a particular normal mode involved in a three-phonon process is actually populated with phonons. This information comes from the phonon distribution functions $n_{\mathbf{q}}(\mathbf{x}, t)$, which indicate the number of phonons in each mode as a function of space and time. To determine the phonon distribution functions, it is necessary to solve the Boltzmann transport equation (BTE), the key equation of kinetic theory,

$$\frac{\partial n_q}{\partial t} + \mathbf{v}_q \cdot \nabla n_q = C(n_q). \quad (3.9)$$

In the Boltzmann transport equation, the time evolution of the distribution function associated with mode q is determined by the balance of phonon drift, $\mathbf{v}_q \cdot \nabla n_q$, and collisions, $C(n_q)$. The complexity of the Boltzmann transport equation is in the collision term, which integrates over the scattering rates and space- and time-dependent populations of each phonon mode. To make progress, the collision term is linearized by expressing $n_q(\mathbf{x}, t)$ as the equilibrium Bose-Einstein distribution function plus a small perturbation $\Delta n_q(\mathbf{x}, t)$ and then dropping terms of second order in the perturbation, leading to

$$C(n_q) \approx \sum_{q', q''} \left[P_{q, q'}^{q''} (\Psi_{q''} - \Psi_{q'} - \Psi_q) + \frac{1}{2} P_q^{q', q''} (\Psi_{q''} + \Psi_{q'} - \Psi_q) \right], \quad (3.10)$$

where $\Psi_q = \Delta n_q / [n_q^{\text{eq}}(n_q^{\text{eq}} + 1)]$ [76]. The drift term can also be linearized by neglecting the contribution of $\Delta n_q(\mathbf{x}, t)$ and assuming a constant, uniform thermal gradient, resulting in

$$\mathbf{v}_q \cdot \nabla n_q(\mathbf{x}, t) = \mathbf{v}_q \cdot \frac{\partial n_q(\mathbf{x}, t)}{\partial T} \nabla T(\mathbf{x}, t) \approx \mathbf{v}_q \cdot \frac{\partial n_q^{\text{eq}}}{\partial T} \nabla T. \quad (3.11)$$

The distribution functions obtained by iteratively solving the BTE under these simplifications can then be used to determine phonon lifetimes and mean free paths and calculate a thermal conductivity. In bulk materials, this procedure works quite well, as shown for germanium and silicon in Fig. 3.1(b) [76]. In a macroscopic geometry, the thermal conductivity, determined experimentally or through the method described above, need only be plugged into Fourier's law to solve for the temperature field.

3.3 Nanoscale thermal effects

The framework presented in the previous section relies on three key assumptions. First, the *ab initio* force constant calculations require that all lattice sites be physically identical, or, in other

words, have translational symmetry. Second, the reduction to Fourier’s law requires there be a perfectly local connection between the microscopic distribution functions and macroscopic variables such as temperature and heat flux. Finally, the treatment of phonons as particles requires that their phases be sufficiently randomized that they do not exhibit coherent effects.

On length and time scales roughly corresponding to the phonon mean free paths and lifetimes, some or all of these assumptions can fail. Because most thermal phonon mean free paths are sub-micron, the behaviors which emerge upon the breakdown of the macroscopic picture are collectively referred to as nanoscale thermal transport. The cartoons in Fig. 3.2 illustrate a classic example, where heat flow is confined to a thin film, whose top and bottom are held at different temperatures [79]. In Fig. 3.2(a), the film is still sufficiently large that the vast majority of phonons scatter internally. Translational invariance between lattice sites approximately holds and the macroscopic framework is a good approximation. In the extreme case of Fig. 3.2(b), the film is so thin that no phonons scatter internally. Instead they radiate between the two boundaries. The lack of internal scattering means that the distribution functions are equal at all points inside the film, and the local temperature is ill-defined, an example of nonlocality. The remainder of this section discusses the intermediate case in Fig. 3.2(c), where the onset of broken translational symmetry, nonlocality and wave effects can lead to the emergence of new, non-diffusive transport behaviors. Due to the broad phonon mean free path spectrum in most materials, the complex situation in Fig. 3.2(c) applies over a wide range of length, time and temperature scales, including those of most modern semiconductors and quantum devices [80].

3.3.1 Broken translation symmetry and boundary effects

If the extent of a crystal lattice is very large, such that the vast majority of atoms are far away from its edges, then it is safe to assume that the interatomic force constants depend only on the relative locations of atoms in the unit cell, rather than the absolute locations of the atoms in the crystal. However, if a significant fraction of atoms resides near boundaries, as is the case in a nanostructured system, the force constants will vary spatially, introducing a spatial variation of the




Figure 3.2: Phonon thermal transport in a thin film: (a) For macroscopic film thicknesses, the vast majority of phonons scatter internally with other phonons or impurities rather than boundaries. Thermal transport is well described by the diffusive Fourier's law, with a bulk thermal conductivity calculated using the *ab initio* methods discussed above. (b) In the thin film limit where no internal scattering occurs, phonons propagate between boundaries in analogy to the blackbody radiation of photons. (c) The intermediate case, where phonon mean free paths are comparable to system dimensions, applies to most room temperature nanoscale devices. Both internal and boundary scattering contribute significantly here, making it much more difficult to simplify the microscopic dynamics.

phonon properties.

Because the calculation of spatially-varying force constants is so computationally difficult, the degree to which phonons near a boundary feel its effect remains an open question [81]. Recent experiments have highlighted a behavior known as elastic softening, wherein the thermal phonon velocity is observed to decrease in the vicinity of a boundary, an effect which has also been reported for longer-wavelength acoustic oscillations due to the weaker chemical bonding environment near surfaces [82, 65]. If harmonic phonon properties change in the vicinity of boundaries, it will be interesting to see if future research reveals that this is also the case for anharmonic properties such as phonon lifetimes.

Alongside the spatial variation of phonon properties near surfaces, there is also the question of what happens when a phonon actually collides with a boundary. This question has spawned a vast area of research concerned with the stochastic modeling of phonon distributions constrained to percolate through nanostructured geometries [83, 84]. At each boundary encounter, phonon phase and momentum are either preserved or destroyed depending on the boundary's roughness. The main constraint on thermal transport in these models is the neck size or smallest aperture of the nanostructured geometry, a key consideration also in the theory developed in Chapter 4 [85]. Many of these stochastic models make the major assumption that phonons experience perfectly bulk conditions when propagating between boundaries, such that when boundaries are spaced much closer than the bulk mean free path, phonons experience no internal scattering. This is the ballistic framework of nanoscale thermal transport, where the characteristic distance between boundaries truncates the bulk mean free path. One consequence of the ballistic framework is that sufficiently small geometries support ray-like phonon transport effects, as illustrated in Fig. 3.3(b) [86].

3.3.2 Nonlocality, Fourier's law and the relaxation time approximation

The assumption of locality resides in the series of approximations that reduce the Boltzmann transport equation to the macroscopic Fourier law, the most notable of which is the relaxation time approximation (RTA) [88]. In the RTA, the distribution function of each phonon mode is assumed




Figure 3.3: Nanoscale thermal transport frameworks: The effects of nanoscale heat sources and boundaries are often investigated through the lens of (a) a wave mechanical, or phononics, framework centered around changes to the phonon dispersion, (b) a ballistic framework, which uses the phonon gas model with a simplified collision operator and (c) a hydrodynamics framework, which again uses the phonon gas model but relies on thermodynamically-motivated simplifications to the distribution function. The ballistic framework predicts that nanoscale structuring will enable the ray-like manipulation of phonons, in analogy to light, while hydrodynamics anticipates fluid-like heat flux in the vicinity of nanostructures, even at length scales far below the bulk phonon mean free paths. Panels (a) and (b) adapted from [86, 87].

to evolve independently, allowing the collision term to be approximated as

$$C(n_q) \approx -\frac{n_q - n_q^{\text{eq}}}{\tau_q}. \quad (3.12)$$

In Eq. 3.12, τ_q is a characteristic phonon lifetime for mode q calculated by assuming all other modes are in equilibrium. While the RTA greatly reduces computational requirements while preserving microscopic details, it does so by violating fundamental aspects of the underlying statistical mechanics. In particular, energy is not inherently conserved in RTA calculations, although energy conservation can be imposed using certain formulations, and the RTA is known to fare especially poorly in situations where normal collisions dominate over Umklapp [89].

Another common approximation to the BTE is the linearization of the drift term as in Eq. 3.11, which assumes that the temperature gradient is constant over the system dimensions. Using the drift term linearization and the RTA, the Boltzmann transport equation can be reduced to Fourier's law, with the thermal conductivity defined microscopically as the sum over phonon mean free paths (Λ), velocities (v) and specific heats (C),

$$\kappa = \frac{1}{3} \sum_q C_q \Lambda_q v_q, \quad (3.13)$$

associated with each mode q .

Fourier's law is perfectly local, meaning that the heat flux in Eq. 3.5 depends only on the temperature gradient at the exact location $\mathbf{x}$ and time t . Macroscopically, this is a sound approximation, but at the nanoscale it is unclear whether it is meaningful to even define a temperature or heat flux localized to an area smaller than a phonon mean free path [90]. For example, if Fourier's law and the bulk thermal conductivity were applied to the film in Fig. 3.2(b), the implicit assumption would be that internal phonon scattering establishes a local equilibrium and local temperature at every point. In the limit of negligible film thickness, this would require a negligibly small phonon mean free path. In order to maintain the bulk thermal conductivity in the face of a decreasing phonon mean free path, the phonon group velocity would need to increase, which would in turn

cause the heat flux predicted by Eq. 3.5 to trend to infinity. In practice, experiments measure an effective cross-plane thermal conductivity for thin films, which decreases with film thickness [91, 92]. For the case of a 1D film with a single length scale, this can be explained via the truncation of the phonon mean free path spectrum to include only those phonons with bulk mean free paths shorter than the film thickness [93]. The extension of this ballistic approach to geometries with additional length scales leads to the stochastic simulations described in the previous section.

3.3.3 Phonon hydrodynamics

When solving the full BTE with the collision term linearized, one finds that local conditions depend on the average properties of the distribution function within a spatial and temporal radius approximately defined by the phonon mean free path and lifetime [94]. Ideally, one could solve the linearized BTE for any arbitrary initial condition and geometry. However, the impossibility of deducing the exact initial microscopic distribution function and the complexities associated with boundary conditions preclude this, motivating the approximations described above [95]. An alternative to simplifying the collision term in the BTE is to solve the BTE using an ansatz, or assumed form for the distribution function. This is the approach taken by the phonon hydrodynamics framework of nanoscale thermal transport [96, 97]. The tradeoff is that while this approach obscures the microscopic details of the distribution function, it maintains the nonlocality and conservation laws of the full kinetic theory. A typical ansatz for the distribution function in the hydrodynamic theory is

$$n_q = n_q^{\text{eq}} + \beta_q \cdot \mathbf{q} + \gamma_q \cdot \frac{\partial \mathbf{q}}{\partial t} + \mathbf{G}_q : \nabla \mathbf{q} \quad (3.14)$$

where β_q and γ_q are vectors and $\mathbf{G}_q$ is a rank two tensor [98]. At first glance, this is a surprising formula, since one typically sums over the distribution functions to arrive at the macroscopic variables, as in Eq. 3.5, and not the other way around. In nonequilibrium thermodynamics, this practice of expanding the distribution function in terms of macroscopic variables is known as the moment

method; depending on the situation it may be necessary to involve more or less macroscopic variables, or moments, in order to capture the physics [99, 100]. In using the ansatz of Eq. 3.14, the assumption is that higher order moments decay rapidly, because their decay produces a large amount of entropy, and are negligible on intermediate length and time scales [101]. When inserted into the BTE, this leads to a macroscopic transport equation of the form [102]

$$\tau \frac{\partial \mathbf{q}}{\partial t} + \mathbf{q} = -\kappa \nabla T + \ell^2 [\nabla^2 \mathbf{q} + \alpha \nabla (\nabla \cdot \mathbf{q})] \quad (3.15)$$

where the coefficients τ , κ and ℓ can be calculated using *ab initio* methods, while α is set to 1/3 for silicon [67]. Eq. 3.15 resembles the Navier-Stokes equation for fluid flow, giving the hydrodynamic framework its name. Because the simpler ansatz $n_q = n_q^{\text{eq}} + \beta_q \cdot \mathbf{q}$ reduces the BTE to Fourier's law [103], Eq. 3.15 with *ab initio* coefficients can be considered a first order refinement which captures the onset of nanoscale effects. Fig. 3.3(c) shows heat flux profiles in nanostructured geometries predicted by phonon hydrodynamics using Eq. 3.15, which can be solved, with boundary conditions, using finite element methods. Instead of the ray-like behavior characteristic of the ballistic framework, the heat flux adopts a parabolic Poiseuille profile within each channel, and curls to accommodate obstacles [10, 104, 105, 106, 107]. This is a consequence of the nonlocality imposed by the $\nabla^2 \mathbf{q}$ and $\partial \mathbf{q} / \partial t$ terms in Eq. 3.15.

A great deal of confusion has arisen over the name phonon hydrodynamics. This is because Eq. 3.15 can also be derived from the BTE without the distribution function ansatz but by assuming that normal (crystal-momentum-conserving) three-phonon collisions dominate over Umklapp [108, 109, 110, 111]. However, this does not imply that hydrodynamic-like behaviors associated with Eq. 3.15 only appear in the normal collision limit [74]. For example, when the $\partial \mathbf{q} / \partial t$ term dominates over $\nabla^2 \mathbf{q}$, Eq. 3.15 predicts wavelike oscillations in the temperature. This has been observed in graphite, a material with strong normal scattering, but also in germanium, where Umklapp processes are more prevalent [112, 113].

3.3.4 Phonon coherence and wave effects

In certain instances, the entire picture of phonons as particles falls apart, and it is necessary to go back to viewing them as semi-localized wave packets. For example, if the phases of individual phonons are not randomized through collisions with boundaries, impurities or other phonons, they may exhibit Bragg-like coherent interference. Practically, phonon coherence requires materials with very few defects and either extremely smooth surfaces or low temperatures, since at low temperatures the wavelengths of heat-carrying phonons are larger [114, 115]. Under these conditions, the phonon wavelength may exceed the boundary roughness, allowing phonons to reflect specularly, preserving their phase.

However, the utility of the wave perspective extends beyond this low-temperature limit. Atomistic simulations, which use empirical formulas for the interatomic potential, implicitly consider anharmonic effects beyond the third-order approximation to the interatomic potential. Using these computational approaches, it is possible to calculate the phonon dispersion and thermal conductivity independently of the kinetic framework [116]. Fig. 3.3(a) shows dispersions obtained using lattice dynamics for a thin silicon film (left) compared to a silicon nanophononic metamaterial (right), which consists of the thin film with a pillared array on its surface [87]. In the latter case, resonant modes in the pillars hybridize with film modes, reducing group velocity and thermal conductivity [117]. This perspective also relates to the case of highly anharmonic and disordered materials, where broad, closely-spaced bands facilitate the types of vibrational interactions that account for heat transport in amorphous solids [118, 119, 120].

3.4 Phonon thermal transport from nanoscale transducer arrays

A recent body of work at JILA has used the extreme ultraviolet light sources discussed in Chapter 2 to investigate the phonon transport dynamics of heated nanoscale transducer arrays [58, 62, 64, 67]. The typical system is shown in the inset to Fig. 3.4(a) and consists of a periodic 1D nickel grating of linewidth (L) and period (P) on a silicon substrate. An infrared pump pulse

first heats electrons in the nickel structures, which pass energy to the phonons on a sub-picosecond timescale. Phonon scattering across the nickel-silicon interface then gradually transfers energy to the substrate below. A time-delayed EUV probe pulse measures the change in nickel grating diffraction efficiency as a function of pump-probe delay time. The EUV probe is not sensitive to the nickel or silicon electron dynamics because its photon energy is too high to measure transitions between valence levels and tuned away from transitions involving core levels. Instead, due to nickel's large thermal expansion coefficient and higher temperature rise relative to the silicon substrate, the EUV probe primarily measures the nickel grating's thermal expansion and subsequent relaxation.

Fig. 3.4(a) shows the change in EUV diffraction efficiency for a nickel grating with linewidth $L = 20$ nm and period $P = 400$ nm. At 20 nm, the nickel structures are much smaller than the average bulk phonon mean free path in silicon. As a result, and in line with theoretical predictions, the nickel structures cool more slowly than predicted by bulk Fourier's law [58]. The reduction in apparent thermal conductivity occurs for the same reasons of nonlocality discussed in Section 3.3.2 in the context of a thin film, since there are very few phonon-phonon collisions near each nickel nanostructure, and no local thermal equilibrium in its vicinity. It is important to note that regardless of the nickel grating linewidth, there is likely a nanoscale region adjacent to each structure which experiences nonlocality; however, the size of this region does not scale with structure size and is therefore negligible for sufficiently large geometries.

Quite surprisingly, subsequent experiments revealed that for two nanoscale nickel gratings of equal linewidth, the grating with the shorter period will cool faster [62, 64]. This behavior is clearly visible in the comparison between the $P = 400$ nm and $P = 80$ nm geometries in Fig. 3.4(a) and Fig. 3.4(b) respectively. This absolute increase in thermal dissipation efficiency cannot be a coherent interference effect, as the nickel-silicon interface randomizes the phonon phases. It might be related to a change in interatomic force constants near the silicon surface due to mass loading from the nickel structures, or be caused by the nonlocal correlation of heat flux from nearby heat sources. Due to the difficulty of solving the BTE in such a complex geometry, it has thus far defied a rigorous explanation.




Figure 3.4: Close-packing effect and thermal transport channeling: (a), (b), (d) Dynamic extreme ultraviolet scatterometry experiments measure the thermal relaxation of 1D nanoscale nickel gratings of linewidth (L) and period (P) on a silicon substrate. Surprisingly, nickel structures of equivalent linewidth are observed to cool faster when spaced more closely together. (c) Nonequilibrium atomistic simulations of periodic silicon heaters on a silicon thin film substrate predict a higher cross-plane thermal conduction when the heaters are closely spaced (orange) vs. effectively isolated (blue). (e) Spectral energy density analysis of the atomistic simulations reveals that in-plane phonon lifetimes decrease as heaters are packed more closely, suggesting a channeling mechanism whereby in-plane phonons can more readily scatter downwards, increasing cross-plane thermal conduction. Figure adapted from [64, 121].

Recent atomistic nonequilibrium molecular dynamics simulations have provided an exciting clue as to the origin of this so-called close-packing effect [121]. The simulations are limited in scale due to the intensive computational requirements and so only consider silicon structures of varying linewidth and period on top of a thin silicon membrane, which limits the phonon mean free path to approximately the membrane thickness. The silicon structures and a thin heat sink layer on the bottom of the membrane are held at constant temperatures T_2 and T_1 respectively, in order to calculate the steady-state temperature profile in the remainder of the membrane. From the temperature profile and heat flux, it is possible to calculate an apparent thermal conduction value, as described in Ref. [121]. The apparent thermal conduction calculated as a function of L and P for the molecular dynamics simulations exhibits similar behavior as observed in the EUV experiment, although the former measures cross-plane thermal conduction while the latter fits an effective nickel-silicon thermal boundary resistivity. As shown in Fig. 3.4(c) and Fig. 3.4(d), for geometries with a fixed grating period (blue), thermal conduction varies monotonically with linewidth, while for geometries with a fixed ratio between linewidth and period (orange), conduction decreases then increases as the grating size shrinks.

The molecular dynamics simulations solve for the atomic displacements and velocities as a function of time. From these outputs, a spectral energy density calculation, essentially a massive Fourier transform, reveals the phonon dispersion, as illustrated in Fig. 3.4(e) [122]. Because the phonon dispersion calculation includes anharmonicity, the phonon lifetimes can be extracted by fitting the band linewidth, where a broader linewidth indicates a shorter lifetime. Fig. 3.4(e) shows phonon lifetimes calculated in the in-plane direction (parallel to the sample surface) for two geometries of equal linewidth and periods $P = 69.5$ nm (blue) and 17.4 nm (orange). Collisions between phonons originating from closely-spaced adjacent heaters result in a much shorter in-plane phonon lifetime for the $P = 17.4$ nm case. When phonons from adjacent heaters collide, their in-plane crystal momentum cancels, while their cross-plane crystal momentum adds, resulting an energy current preferentially directed in the cross-plane direction. This “thermal transport channeling” phenomenon may account for the close-packing effect observed in the EUV experiments and inform

heat management considerations in future nanoscale device design.

Chapter 4

Acoustic and thermal behaviors of <100 nm 3D silicon phononic crystal metalattices

Having made a brief overview of ultrafast short-wavelength laser-based metrology techniques and nanoscale phonon transport in the previous chapters, we now turn to a specific example at the intersection of the two fields. While it has long been known that boundaries and nanostructuring can modify phonon thermal transport [123], it has only recently become possible to create materials in which the nanostructuring is tailored to produce desirable thermal properties, which would otherwise be unrealizable in a bulk material. One such class of nanostructured thermal materials are phononic crystals. Named in analogy to photonic crystals, which manipulate photons via wavelength-scale periodic structuring, phononic crystals use periodic nanoscale voids or inclusions to modify microscopic thermal phonon behavior, thereby altering the thermal properties at the macroscale [85]. Because the thermal phonon wavelength in common semiconductors at room temperature is on the order of a single nanometer, almost three orders of magnitude less than visible light, the fabrication requirements for phononic crystals are far more demanding, which explains the immaturity of phononic crystals as compared to their optical counterparts.

4.1 Overview of phononic crystals

Fig. 4.1 provides an overview of phononic crystals and their development over the last two decades. The top panel shows the phononic analogy of the electromagnetic spectrum, with heat-carrying phonons occupying the terahertz frequency range, far above the lower-frequency and

longer-wavelength audible and ultrasonic acoustic vibrations more easily manipulated at longer length scales. Fig. 4.1(b) shows examples of phononic crystals in one to three dimensions. A one-dimensional phononic crystal, like the superlattice shown here, is far easier to scale down to the smallest length scales than a two- or three-dimensional sample, since it requires alternating a thin film deposition process, without the lithographic or self-assembly steps necessary to realize the nanoscale periodicity along additional axes. There has been extensive work on two-dimensional phononic crystals in the last decade, which usually consist of periodic holes in a film structure, as discussed below. Three-dimensional phononic crystals, like those shown in Fig. 4.1(b) and investigated in the present chapter, are only very recently becoming available at dimensions below 100nm [124].

The nanoscale inclusions or voids in phononic crystals affect phonon transport either by presenting obstacles with which phonons incoherently scatter, or by using periodic structuring to induce coherent Bragg-like interference effects. While both incoherent and coherent effects can change the phonon lifetimes and mean free paths, coherent effects can restrict the ability of certain modes to propagate, resulting in band gaps in the phonon dispersion. While coherent effects can be readily observed in atomistic simulations, experimentally they require materials with extremely smooth surfaces and few defects, since both impurities and boundaries with roughness on the order of the single-nanometer phonon wavelength scatter phonons diffusely, randomizing their phases. Thus far, the observation of coherent effects in phononic crystals is limited to 1D superlattices, for which a large body of literature exists [125, 127], and 2D phononic crystals at very low temperatures, where wavelengths associated with heat-carrying phonons increase dramatically, reducing the stringent surface roughness and impurity requirements [115]. Fig. 4.2 shows calculated band structures for a thin SiN film and two 2D phononic crystals consisting of holey SiN films with different periods. Signatures of the band gap predicted for the smaller hole period were observed at <1 K [114].

In 2D and 3D phononic crystals at higher temperatures, incoherent scattering introduced by the periodic structures can result in a reduction of thermal conductivity to 1% or less of the bulk value. It is important to stress that this is an incoherent effect in which the phonons are treated as




Figure 4.1: Overview of phononic crystals: (a) The phononic spectrum, in analogy to the electromagnetic spectrum, classifies vibrations according to their frequency, from long wavelength acoustic waves to the terahertz thermal phonons that serve as the primary heat carriers in insulators. (b) In phononic crystals, in analogy to photonic crystals, nanoscale structures or voids introduced into a material modify its bulk phonon transport environment to generate desirable thermal transport properties. One-dimensional phononic crystals, including superlattices, can be fabricated on much smaller length scales than 2D and more recently available 3D phononic crystals. Figure adapted from [7, 85, 124, 125, 126].




Figure 4.2: Coherent and incoherent effects in phononic crystals: (a) Phononic crystals with length scales on the order of the broadband phonon mean free path spectrum (nanometers to microns in many semiconductors at room temperature) can greatly reduce thermal conductivity from its bulk value by introducing additional incoherent phonon scattering. If electrical conductivity remains intact, this can enhance a material's potential as a thermoelectric. The neck size, or smallest channel through which heat can flow between voids, is found to roughly correlate with effective thermal conductivity across a wide variety of phononic crystal geometries. (b) At single-nanometer phonon wavelength scales, or at cryogenic temperatures, phononic crystals can coherently modify the phonon band structure via the introduction of band gaps, an effect of interest for applications such as thermal diodes and cloaking. Figure adapted from [114, 85].

particles, rather than waves, but it is still a microscopic effect. In a macroscopic Fourier model, the removal of a certain volume of material results in a classical “volume reduction” modification to the thermal conductivity, which can be represented by the models of Eucken or Russell developed in the early part of the 20th century [128, 129] as

$$\kappa_{\text{Eucken}} = \kappa_{\text{bulk}}(1 - \phi)/(1 + \phi/2), \quad (4.1)$$

$$\kappa_{\text{Russell}} = \kappa_{\text{bulk}}(1 - \phi^{2/3})/(1 - \phi^{2/3} + \phi), \quad (4.2)$$

where κ_{bulk} is the bulk thermal conductivity and ϕ is the porosity. This volume reduction effect is, to first order, a scaling by one minus the porosity, otherwise known as the filling fraction, with additional corrections accounting for the simple geometric shadowing effects of the pores. Since it does not account for the microscopic details, classical volume reduction models are accurate when pore size greatly exceeds the average phonon mean free paths, but spectacularly overpredict apparent thermal conductivity at the nanoscale. Instead, the incoherent scattering in a phononic crystal can, to first order, be compared to obstacles in mini-golf, which prevent a certain fraction of randomly-putted particle-like golf balls from reaching their target. Crucially, in this first order approximation, it is assumed that in their time between collisions with obstacles, the phonons behave exactly as they would in a bulk sample, which allows their properties to be extracted from *ab initio* calculations and plugged into a simulation of the geometry, often realized in practice using a stochastic Monte Carlo approach [83, 84]. Due to the complexity of the boundaries, Monte Carlo solvers generally invoke the relaxation time approximation discussed in Section 3.3.2, which assumes that phonon modes evolve independently in their propagation through the system. These assumptions constitute the ballistic picture of incoherent phonon propagation in a phononic crystal. Intuitively, in ballistic picture, if the phonons are launched with a random distribution of initial wavevectors, the probability that any one makes it through the obstacles most closely depends on the smallest constriction through which they must pass, or neck size. Fig. 4.2(a) shows a compilation

of 2D phononic crystal neck sizes vs. experimentally-measured effective thermal conductivities, whose increasing trend illustrates this intuitive effect. However, the ballistic picture is only an approximation; in reality, phonon properties in the vicinity of nanostructures may deviate from their bulk values. The most well-known manifestation of this effect is elastic softening, wherein the phonon group velocity decreases due to the softening induced by dangling atomic bonds at the material edge, reducing the thermal conductivity [82, 130]. Phonon velocity is a harmonic property, in that it does not arise from phonon-phonon scattering. As we shall see, anharmonic properties, such as the mean free path between internal phonon-phonon scattering events, also deviate from their bulk values in highly-confined phononic crystals. The ~ 20 nm pore diameter and 36 nm period of the 3D silicon phononic crystal metalattice studied here places it in this highly-confined category. The experimentally-characterized phonon transport dynamics of this sample and molecular dynamics simulations of similar materials can be modeled using a mesoscale phonon hydrodynamic approach with modified *ab initio* parameters. The remainder of this chapter first introduces the metalattice sample, then discusses the non-trivial extraction of the metalattice porosity and elastic properties, necessary to model its thermal behavior, before summarizing and generalizing the nanoscale thermal transport findings.

4.2 3D silicon metalattice sample

The silicon metalattice sample investigated in the remainder of this chapter was created by collaborators in John Badding’s group at Penn State University. Its fabrication begins with the self-assembly of a monodisperse colloid of silica nanospheres using a seed regrowth process, wherein nanospheres are nucleated and then grown to their desired size in two separate stages [131, 132, 133]. Fig. 4.3(b) illustrates an example silica nanosphere template. A vertical evaporation technique next transfers the nanospheres from the colloid to a silicon substrate, where they assemble as a close-packed crystalline film [134, 135]. Because the vertical evaporation requires aligning the substrate a precise angle to optimize the deposition, it results in a gradient in film thickness across the sample surface. The remainder of the fabrication consists of inverting the silica nanosphere

template to achieve the metalattice structure shown in Fig. 4.3. A high-pressure confined chemical vapor deposition step first infiltrates amorphous silicon into the voids in the silica nanosphere template, terminating when a silicon layer begins to form on top of the sample [136]. The silicon overlayer is then removed with reactive ion etching, followed by an annealing step to crystallize the silicon, wherein the sample is heated to 800C for 30 minutes. Etching in a hydrofluoric acid vapor removes the silica template, leaving only the silicon metalattice structure. The porosity at this point is the approximately 75% packing fraction of the face-centered cubic (FCC) silica nanosphere template. For the sample studied here, additional silicon infiltration, reactive ion etching and annealing steps reduce the porosity to an unknown lower value. A lower porosity is desirable because it was hypothesized to support increased electron transport while leaving the thermal transport relatively impaired, making this particular metalattice a candidate thermoelectric material [124].

It is necessary to fabricate nickel gratings on the surface of the metalattice sample for the EUV scatterometry experiment, both to act as absorbers for the infrared pump, and as a diffraction grating for the EUV probe. This is accomplished with an electron-beam lithography processes similar to that presented in Chapter 5, and detailed in Ref. [66]. The nickel gratings are ~ 12 nm in height, range in period from 120 nm to 4 μm , and have linewidths generally equal to one-quarter of the period. Fig. 4.3(c) shows the full sample geometry, with the metalattice as a ~ 500 nm film on a silicon substrate and the nickel gratings on the surface. There is also a thin ~ 15 nm silicon layer between then nickel structures and the beginning of the metalattice, which is left over from the reactive ion etching step.

4.3 Extracting metalattice pore geometry and porosity via electron tomography

A 3D electron tomography reconstruction of the metalattice sample structure reveals its FCC pore distribution, ~ 36 nm pore period, 10-20 nm average pore size and a $40 \pm 2\%$ porosity. Electron tomography is a destructive technique, requiring focused ion beam milling of the sample. The procedure is described extensively in Ref. [68] and summarized briefly here. Since the measurement




Figure 4.3: 3D silicon phononic crystal “metalattices”: (a) Metalattice fabrication procedure, wherein a monodisperse silica nanosphere template is first infiltrated with silicon via high pressure chemical vapor deposition, then etched away, yielding a 3D face-centered cubic void structure with porosity near the $\sim 75\%$ filling fraction. The porosity can then be reduced through a further silicon infiltration step. The nickel gratings on the sample surface are fabricated using an electron-beam lithography process. (b) Example monodisperse silica nanosphere template. (c) Full sample schematic (left) showing the ~ 500 nm thick metalattice film, ~ 10 nm silicon overlayer and nickel grating on the silicon substrate. The metalattice unit cell (right) shows the ~ 20 nm pore diameter and ~ 36 nm pore period. Figure adapted from [68, 133].

requires transmission of the electron beam through a cross-section of the sample, it is first necessary to obtain a thin slice, or lamella, of the metalattice film. This is accomplished using the FEI Nova 600i Dual Beam tool in the COSINC-CHR facility on East Campus. This is a two-component system designed for transmission electron microscopy sample preparation, consisting of a Ga focused ion beam (FIB) and a field emission scanning electron microscope. We obtain lamellae at three locations on the sample surface, corresponding to the nickel gratings with period 4 μm , 400 nm, and 120 nm (see Fig. 4.6(b)). The lamellae were subsequently polished using a low energy Ga beam to mitigate any damage from the FIB process. The electron tomography uses a FEI Titan Themis system, also in COSINC, to obtain a high-angle annular dark-field scanning transmission electron microscopy (HAADF-STEM) single-axis tilt series on one of the lamellae, which consists of an array of 1024 x 1024 pixel 2D images (0.7 nm pixel size) projections of the sample at angles varying from -69° to $+65^\circ$. The HAADF acronym in electron microscopy refers to the use of an annular detector, which excludes the primary electron beam and collects only electrons scattered at certain angles; for this reason, dark areas in the images correspond to void, while the lighter areas are silicon. Reconstruction of the 2D slices via the GENFIRE algorithm ultimately produces a 3D 512 x 512 x 512 image of a section of the metalattice sample [137]. The number of voxels along each axis is half the number of pixels in the 2D images due to computational constraints, resulting in a 1.4 nm voxel size. To verify that the voxel resolution is the same as the voxel size, it is necessary to perform a Fourier shell correlation analysis, as discussed in the Supplementary Information of Ref. [68], along with further details on the pre-processing of the 2D slices. The 3D reconstruction is ultimately cropped to a volume of 251 x 219 x 125 voxels to account for geometric effects and artifacts in the reconstruction process and to exclude the nickel grating and silicon overlayer film from the subsequent determination of the metalattice porosity.

Fig. 4.4 illustrates the process of analyzing the 3D reconstruction to determine porosity. While there are many techniques which one could use to determine a sample's porosity, they are generally ill-suited to the thin film geometry and unknown pore interconnectivity of the metalattice sample. Infiltration-based approaches like adsorption porosimetry cannot detect isolated pores,

which may be expected in this sample given the silicon reinfiltration step in its fabrication [138]. Many optical and X-ray techniques, such as ellipsometry or reflectometry, require effective medium modeling or assumptions about the sample geometry [139, 140, 141, 142]. This leaves electron tomography, while destructive and time consuming, as one of the only viable metrologies [143]. The difficult step in extracting porosity from the 3D reconstruction is to assign each of the grayscale voxels to be either silicon or void, a process known as segmentation. Two-dimensional slices of the 3D tomographic reconstruction are first smoothed with an anisotropic diffusion algorithm, and sharpened with an unsharp mask sharpening algorithm, as shown in Fig. 4.4(a) [144, 145]. This enhances the contrast at the pore boundaries in preparation for the segmentation step. Three segmentation algorithms are tested, global thresholding [146], local thresholding [147] and k-means clustering [148], to establish that the segmentation does not significantly depend on the choice of algorithm. The resulting segmentation is shown in the upper right corner of Fig. 4.4(a), where now the white pixels correspond to void and the black pixels to silicon. Fig. 4.4(b) shows the 3D result of the segmentation over the cropped portion of the reconstruction. The extracted porosity is $40 \pm 2\%$, with the error arising from differences between the segmentation algorithms.

In addition to porosity, the tomography dataset provides important geometric details that aid in developing finite element models for the EUV scatterometry analysis. A Fourier transform of the 3D reconstruction shows a body-centered cubic distribution, which implies that the metalattice pores are packed in the inverse face-centered cubic geometry. The 2D HAADF-STEM images include the nickel structures and silicon substrate on either side of the metalattice film, which allows us to precisely measure its thickness at the locations where lamellae were extracted. The thickness varies between 398 ± 4 nm and 612 ± 3 nm for the 4 μ m and 120 nm period gratings, respectively. The 400 nm period grating, which is located in between the other two, has an intermediate thickness. This gradient in metalattice thickness across the sample surface is expected, and caused by the vertical deposition step in its fabrication. Finally, by fitting ellipses to contiguous areas of void in both the 2D projections and slices of the 3D reconstruction, it is possible to estimate the metalattice pore size and period, which are 10-20 nm and 35-36 nm, respectively. The high error in pore size is

due to misidentifications of pores in the fitting algorithm due to pore interconnectivity and image artifacts. Further details on this process appear in the Supplementary Information of Ref. [68].

4.4 Nondestructive EUV measurements of metalattice porosity, elastic properties and thermal conductivity

The dynamic EUV scatterometry measurements described and analyzed in the remainder of this chapter provide a nondestructive means to extract the metalattice porosity, Young's modulus and apparent thermal conductivity, which are determined to be $38.5 \pm 2\%$, 66 ± 3 GPa and 1.6 ± 0.4 W/(m·K), respectively. The porosity value agrees well with the electron tomography measurements in Section 4.3. The ability of the EUV scatterometry experiment to simultaneously measure so many material properties lies in its sensitivity to both thermal and acoustic dynamics, as shown in Fig. 4.5. The thermal signal corresponds to the gradual reduction in average nickel structure height as phonons scatter incoherently across the nickel-silicon interface and deposit thermal energy in the metalattice film. The acoustic signal consists of the coherent oscillations written on top of the thermal data, which correspond to a standing surface acoustic wave (SAW) mode launched by the impulsive expansion of the nickel gratings upon excitation by the infrared pump pulse. The surface acoustic wave propagates in the metalattice film and silicon substrate with penetration depth $\delta \sim \lambda_{\text{SAW}}/\pi$, where λ_{SAW} is the SAW wavelength. Since λ_{SAW} corresponds to the grating period, it is possible through measurements of multiple nickel gratings and finite element modeling to determine the SAW velocity and elastic properties of only the metalattice. The following sections first give an overview of the EUV dataset and finite element models used in the analysis, before summarizing the results. It is found that the metalattice elastic properties can be well described by classical effective medium models, while the apparent thermal conductivity cannot. This is because the SAW wavelength is much larger than, and insensitive to, the exact pore geometry, while the pores significantly affect the propagation of thermal phonons. Section 4.4.8 discusses the error bar calculation for the various metalattice material properties, based on variation between the experimental scans and uncertainties in certain finite element model parameters. The error bar




Figure 4.4: Electron tomography reconstruction of the 3D silicon metalattice geometry: (a) 2D slices of the 3D tomographic reconstruction, which are smoothed with an anisotropic diffusion algorithm, sharpened with an unsharp mask sharpening algorithm and finally segmented in order to assign each pixel to silicon (black) or void (white). Comparison of three segmentation algorithms, global thresholding, local thresholding and k-means clustering, yields an estimated porosity of $40 \pm 2\%$. (b) 3D visualization of the tomographic reconstruction in the tomviz software (www.tomviz.org) with resolution equal to the 1.4 nm voxel size, as verified with a Fourier shell correlation analysis. Analysis of the electron tomography data confirms the FCC pore arrangement and 36 nm pore period, and provides a pore size estimate of 10-20 nm, which is complicated by the variability, roughness and interconnectedness of the pores. Figure adapted from [68].

approach given here could be generalized to the extraction of elastic and thermal properties of other samples in the dynamic EUV scatterometry experiment. The understanding of the metalattice sample gained here makes possible the development of a new theoretical description of thermal transport in highly-confined semiconductor nanosystems discussed in Section 4.5.

4.4.1 EUV scatterometry dataset

The dynamic EUV scatterometry dataset used for the acoustic and thermal analysis consists of the 47 experimental scans summarized in Table 4.1. Each scan measures the change in EUV diffraction efficiency as a function of delay time between the infrared pump and EUV probe, as shown in Fig. 4.5. The raw data for each scan are arrays of size $\mathcal{T} \times 1024 \times \mathcal{L}$, where $\mathcal{T}$ is the number of time steps, or unique temporal delays between the pump and probe, 1024 is the number of pixels on the CCD camera and $\mathcal{L}$ is the number of loops, or repetitions, of the data acquisition process. A typical scan consists of three loops to verify data repeatability, with additional loops as necessary to improve the signal to noise ratio through averaging.

There are two raw data arrays associated with each scan, which measure the EUV probe diffraction pattern with and without the infrared pump. Subtracting these two arrays gives the difference signal shown in Fig. 4.5(b). The difference signal primarily tracks the effect of the thermally-induced change in nickel grating height on the EUV diffraction pattern. Secondary effects include changes in the grating and substrate reflectivities. The difference signal is aggregated by summing across the 1024 camera pixels at each time step. Because a change in nickel structure height will conserve the overall diffracted intensity, the diffracted orders (AC) and specular beam (DC) in the differential EUV diffraction generally have opposite sign. When aggregating across the camera chip, signal on pixels associated with the AC orders is inverted, as shown by the red shading in Fig. 4.5(b). Areas of the camera chip clearly unimpacted by the EUV diffraction can be excluded to improve the signal to noise ratio. The result of this data processing is a “consolidated” two column array, where the first column gives the time delay between pump and probe and the second column gives the difference in aggregated EUV diffraction efficiency with and without the




Figure 4.5: EUV scatterometry measurements of the 3D silicon metalattice: (a) The nickel gratings absorb the infrared pump pulse and expand, launching surface acoustic waves (yellow) and injecting heat (red) into the metalattice layer below. (b) A CCD camera monitors the difference in the diffraction of the EUV probe from the nickel grating with and without the infrared pump excitation. The differential diffraction pattern (top) is aggregated at each time delay by multiplying the positive signal by one and the negative signal by negative one to produce an experimental scan (bottom). The experimental scans feature an acoustic oscillation with wavelength equal to the grating period superimposed on a thermal decay corresponding to the thermal relaxation of the nickel structures. Analysis of the acoustic signal reveals the metalattice porosity and elastic properties, while analysis of the thermal signal reveals the apparent metalattice thermal conductivity and transport behavior. Figure adapted from [68].

Linewidth [nm]	Period [nm]	Max delay [ps]	Time step [ps]	Scans
50	200	8000	40	6
50	200	750	5	4
50	200	1500	10	1
100	400	8000	40	6
100	400	1500	10	6
100	800	8000	40	4
100	800	3000	15	1
100	800	3500	20	2
100	800	4000	20	1
500	2000	8000	40	8
1000	4000	8000	40	8

Table 4.1: Experimental EUV dataset for the silicon metalattice sample. The 47 scans measure different time delays to capture the thermal and acoustic dynamics across five nickel grating geometries.

infrared pump, $\Delta I(t)$, where

$$\Delta I(t) = \sum_{p \in \mathcal{P}_{\text{DC}}} [I_p^{\text{on}}(t) - I_p^{\text{off}}(t)] - \sum_{p \in \mathcal{P}_{\text{AC}}} [I_p^{\text{on}}(t) - I_p^{\text{off}}(t)]. \quad (4.3)$$

In Eq. 4.3, $I_p^{\text{on}}(t)$ and $I_p^{\text{off}}(t)$ are the counts at delay time t on CCD pixel p with the infrared pump on (unblocked) and off (blocked), respectively. The sums are over the subsets of pixels $\mathcal{P}_{\text{DC}}$ and $\mathcal{P}_{\text{AC}}$ capturing the DC beam and AC diffracted orders. Eq. 4.3 is often normalized to the static diffraction profile, revealing a change in diffraction efficiency due to the infrared pump of 10^{-3} or 10^{-4} for most samples. It is expected that $\Delta I(t)$ will be zero for $t < 0$, meaning that the EUV probe arrives before the infrared pump. In practice, small amounts of pump light can bypass the aluminum filters in front of the CCD cameras, resulting a small, time-independent signal before time zero which is averaged and subtracted from each scan.

Each scan measures a particular nickel grating geometry of linewidth (L) and period (P).

Hereafter, grating geometries are referenced using their linewidth and period in nanometers, such that 100-400 means the grating with 100 nm linewidth and 400 nm period. The nickel structure height (h) is nominally the same between gratings. Fig. 4.6(b) shows a top-down optical microscope image of the metalattice sample, where each grating is repeated as a column of 150 μm squares with centers separated by 400 μm . Each row corresponds to a different dose in the electron-beam lithography fabrication procedure. Because the infrared pump beam spot size is $\sim 400 \mu\text{m}$ on the sample, it heats a particular nickel grating square relatively uniformly. The penetration depth of 800 nm light in silicon is $\sim 8 \mu\text{m}$, versus $\sim 13 \text{ nm}$ for nickel, making it safe to assume that initially only the nickel structures are meaningfully heated. The $\sim 70 \mu\text{m}$ EUV probe spot size fits neatly into a square grating array, so that all of the EUV light incident on the sample participates in the diffracted signal. The number of EUV diffracted orders captured on the camera chip depends on the grating period. At $P = 200 \text{ nm}$, the diffraction angle is so large that only one diffracted order makes it on the chip. For the larger gratings, the diffracted orders may overlap, especially since their profile on the CCD camera is broadened by the high surface roughness of the metalattice sample. The time step of each scan must be selected for the dynamics of interest. Due to the low apparent metalattice thermal conductivity, which bottlenecks the thermal relaxation of the nickel gratings, the full 8 ns temporal delay available in the EUV scatterometry experiment is necessary to capture the thermal decay. Because the nickel grating period sets the surface acoustic wave wavelength, smaller gratings have higher SAW frequencies. For the 50-200, 100-400 and 100-800 gratings, the 40 ps time steps used in the 8 ns scans alias the SAW frequencies. Additional, “acoustic” scans for these grating sizes, shown in Fig. 4.6 and Table 4.1, have shorter time steps and reveal the SAW oscillations.

Finally, while EUV data was acquired for the 30-120 grating as indicated by Fig. 4.6(b), this data was not used in the thermal and acoustic analysis. The thermal data had a noise level too high for it to be reliably fit. The acoustic data, taken with shorter time steps, showed a SAW oscillation, however, this SAW was too closely confined to the surface, including the $\sim 15 \text{ nm}$ silicon overlayer and first layer of metalattice pores, to give useful information about the average porosity and elastic




Figure 4.6: Dynamic EUV scatterometry dataset for the silicon metalattice sample: (a) Average of the acoustic and thermal scans collected at each grating size. The 1000-4000 and 500-2000 scans launch SAWs of sufficiently long wavelength that an 8 ns scan with 40 ps step size captures both the acoustic and thermal dynamics. (b) Top-down map of the sample surface taken with an optical microscope. The nickel gratings are visible as an array of squares of $150\text{ }\mu\text{m}$ side length spaced $400\text{ }\mu\text{m}$ apart, where each column is a different grating geometry and each row is different dose in the electron-beam lithography fabrication process. The insets show SEM images of the individual nickel lines for a 1000-4000 and 100-400 grating. These inset images are from a different fabrication on a bulk silicon substrate. Blue dots correspond to the gratings and doses measured with dynamic EUV scatterometry, while red crosses indicate the locations at which lamellae were extracted for the electron microscopy. Figure adapted from [10, 68].

properties of the ~ 500 nm metalattice film. Gratings of $L = 30$ nm, $P = 400$ nm and $L = 50$ nm, $P = 400$ nm were also fabricated on the metalattice but never reliably measured, due to challenges associated with their small size and low duty cycle.

4.4.2 Elastic theory and surface acoustic waves

In order to relate the metalattice porosity and elastic properties to the experimentally-measured SAW frequencies, we construct finite element models of the full metalattice sample and a single unit cell of the FCC pore geometry. These finite element models are necessary due to the sample complexity, which prevents the use of analytic expressions. However, a brief discussion of the underlying elastic theory based on the presentation in Refs. [25, 149, 150] illustrates the logic behind the finite element modeling of the following two sections.

Linear elastic theory begins with a generalized linear relationship between stress, the force acting on a body and strain, its infinitesimal deformation, written as

$$\sigma_{ij} = C_{ijkl}\epsilon_{kl}, \quad (4.4)$$

where σ_{ij} and ϵ_{ij} are elements of the stress and strain tensors, $\boldsymbol{\sigma}$ and $\boldsymbol{\epsilon}$, respectively. In Eq. 4.4, the indices donate the three real space coordinate axes and repeated indices are summed. Elements of the strain tensor are defined as

$$\epsilon_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right), \quad (4.5)$$

where $u_i(\mathbf{x})$ is the displacement of the material from equilibrium along axis i at location $\mathbf{x}$. The stress tensor can be related back to the displacement through an equation of motion

$$\nabla \cdot \boldsymbol{\sigma} = \rho \frac{\partial^2 \mathbf{u}}{\partial t^2}, \quad (4.6)$$

where ρ is the material density. In the context of thermal expansion, Eq. 4.4 can be linearly

coupled to the temperature field as

$$\sigma_{ij} = C_{ijkl}(\epsilon_{kl} - \alpha(T - T_0)\delta_{kl}), \quad (4.7)$$

where α is the material-dependent coefficient of thermal expansion, T_0 is the equilibrium temperature, and δ_{kl} is the Kronecker delta. Eq. 4.7 accounts for thermoelastic coupling in the finite element model and results in the impulsive excitation of acoustic waves as the nickel structure experiences a rapid temperature increase due to heat deposited by the simulated infrared pump pulse.

The fourth-order tensor $\mathbf{C}$ relating stress and strain in Eq. 4.4 is often referred to as the elasticity tensor, and in general has 81 components C_{ijkl} . For example, the component C_{2211} describes the effect of the gradient along the the x direction of displacement in the x direction on the stress in the y direction on the plane perpendicular to the y axis. It can be shown that only 21 components of $\mathbf{C}$ are independent even in an anisotropic material. For this reason, elements of the elasticity tensor are often represented using the more compact Voigt notation as C_{nm} , wherein the indices n and m abbreviate ij and kl , respectively.

For a perfectly isotropic material, there are only two independent degrees of freedom in C_{nm} , often selected as C_{11} and C_{44} . All other values in $\mathbf{C}$ can be written as linear combinations of these two elements. There are several alternative notations with which to express these two degrees of freedom, of which three are relevant to the discussion below. First, the Lamé parameters λ and μ relate to the elasticity tensor components as

$$C_{11} = \lambda + 2\mu, \quad C_{44} = \mu. \quad (4.8)$$

Second, Young's modulus (E) and Poisson's ratio (ν), which measure the stiffness and compressibility of the isotropic material, respectively, can be expressed as

$$\lambda = \frac{E\nu}{(1+\nu)(1-2\nu)}, \quad \mu = \frac{E}{2(1+\nu)}. \quad (4.9)$$

Finally, the longitudinal and transverse acoustic wave velocities, c_L and c_T , relate to the elasticity tensor elements as

$$c_L = \sqrt{\frac{C_{11}}{\rho}}, \quad c_T = \sqrt{\frac{C_{44}}{\rho}}, \quad (4.10)$$

where ρ is the density. A longitudinal acoustic wave corresponds to a displacement along the direction of propagation, while for a transverse wave the displacement is perpendicular the propagation direction. In an anisotropic material, there will be additional acoustic velocities corresponding to the different directions of propagation and polarization. The longitudinal and transverse wave velocities are the slopes, at low frequencies, of the two branches in the phonon dispersion shown in Fig. 4.7(a).

The relationships above assume an infinite bulk medium, while the EUV experiment measures surface acoustic waves, whose velocity does not depend so simply on the elastic constants. The simplest type of surface acoustic wave is a Rayleigh wave, which propagates at the interface between a semi-infinite medium and vacuum. The velocity c_R of a Rayleigh wave is the numerical solution to the transcendental equation

$$\left(\frac{c_R}{c_T}\right)^6 - 8\left(\frac{c_R}{c_T}\right)^4 + 8\left[3 - 2\left(\frac{c_T}{c_L}\right)^2\right]\left(\frac{c_R}{c_T}\right)^2 - 16\left[1 - \left(\frac{c_T}{c_L}\right)^2\right] = 0, \quad (4.11)$$

whose solution generally lies in the range $0.87c_T < c_R < 0.95c_T$ [150]. There are two principle additional complications in the present experiment, which render even Eq. 4.11 insufficient and necessitate the use of finite element modelling. First, the substrate consists of three layers rather

than a single semi-infinite medium, introducing the possibility of additional acoustic modes such as Sezawa waves, which propagate inside supported films [151]. Second, the mass loading of the nickel structures modifies the SAW dispersion, reducing the acoustic velocity and scattering a portion of the acoustic energy, which reduces the SAW lifetime and causes it to spread out in frequency space [59]. While perturbative approaches can extend the Rayleigh wave framework to cases of thin films or short structures, they fail when the structure height or film thickness reaches a significant fraction of the SAW wavelength [152]. This is the case for the EUV scatterometry experiment, such that the SAWs observed here are sometimes referred to as pseudo-SAWs [153].

4.4.3 Finite element model of metalattice unit cell

The purpose of the finite element unit cell model is to define functions relating the metalattice Young's modulus and Poisson's ratio to its geometry and porosity. This is necessary because the EUV experiment measures only a single SAW velocity, making it impossible to fit Young's modulus, Poisson's ratio and porosity separately without constraining the relationship between the three parameters. The unit cell model consists of a single period of the metalattice geometry implemented in the COMSOL Multiphysics software and incorporates the FCC pore geometry and 36 nm pore period evident in the electron tomography. The silicon is assumed to be polycrystalline and therefore isotropic, with a Young's modulus and Poisson's ratio of 160 GPa and 0.22, respectively, while the elastic constants of the void are effectively zero [154]. By imposing Bloch boundary conditions,

$$\mathbf{u}_1 = \mathbf{u}_2 e^{-i\mathbf{k} \cdot (\mathbf{r}_1 - \mathbf{r}_2)} \quad (4.12)$$

where $\mathbf{k}$ is the wavevector and $\mathbf{u}_1$ and $\mathbf{u}_2$ are the displacements at boundaries located at $\mathbf{r}_1$ and $\mathbf{r}_2$, respectively, it is possible to solve for the acoustic dispersion shown in Fig. 4.7(a). The slopes of the first two dispersion branches give the longitudinal and transverse acoustic velocities, which can be translated into Young's modulus and Poisson's ratio via Eqs. 4.8, 4.9 and 4.10.

The main difficulty with the unit cell model is that the exact metalattice geometry is

unknown. Specifically, the electron tomography does not precisely indicate the pore diameter or interconnectivity, which may vary from pore to pore depending on the intricacies of the silicon reinfiltration step in the fabrication process. To address this, we construct three unit cell models representing different possible geometries. Summarizing briefly from Ref. [68], the isolated model consists of spherical pores arranged in the FCC geometry expected from the fabrication and confirmed by the electron tomography, whereas the interconnected model is identical to the isolated model except with cylindrical channels of void connecting the centers of the nearest neighbor pores. If the diameter of the pores in the isolated or interconnected model is increased sufficiently, the pores will overlap (eventually the connecting hole diameter will be larger than the channel diameter), creating a network of overlapping voids which we denote as the overlapped model. Thus, the overlapped geometry is a limiting case of the isolated and interconnected models at large pore diameters.

In the case of the isolated model, porosity relates simply to pore diameter D and unit cell side length a as

$$\phi_{\text{isolated}} = \frac{16\pi}{3a^3} \left(\frac{D}{2}\right)^3. \quad (4.13)$$

An analogous relationship can be derived analytically for the overlapping model, resulting in the more complex expression given in Ref. [68], while the porosity in the interconnected model is calculated numerically in the finite element software. Three variants of the interconnected model consider interconnection channels of radius 2 nm, 3 nm and 5 nm. The relationships between porosity, Young's modulus and Poisson's ratio calculated for each of these models appear as discrete points in Fig. 4.7(b), and are approximated as continuous functions using the second order polynomial fits in Table 4.2. For geometrical reasons, it is impossible to realize a porosity below a certain threshold in the 2 nm, 3 nm and 5 nm interconnected and overlapping model, while the isolated model is compatible with arbitrarily small porosity values. When fitting the porosity to the experimental data, this limitation informs the types of unit cell modes with which the metalattice is compatible.

The results of the unit cell model are robust to slight changes in the parameters assumed from

the electron tomography, such as pore period, as discussed in the Supplementary Information of Ref [68]. However, by assigning all volume elements to be either silicon or void, the unit cell model does not consider oxidation or surface elastic effects near the pore boundaries. Oxidation is not expected, since the silicon reinfiltration step in the fabrication continues until a surface layer accumulates, indicating that the pore network is sealed off from the external environment. Surface effects could play a role due to the altered chemical bonding environment near surfaces and the high surface to volume ratio in the metalattice. Accounting for surface effects would mean reducing the average polysilicon elastic properties in the unit cell model, which would result in a smaller fitted porosity. However, because the porosity extracted using the EUV technique agrees well with the electron tomography value, and the Young's modulus obeys an effective medium description, it is unlikely that surface effects greatly affect the elastic behavior.

Unit cell geometry	Young's modulus [GPa]	Poisson's ratio
Isolated	$172.0(1 - \phi)^2 - 35.7(1 - \phi) + 23.2$	$-0.082(1 - \phi)^2 + 0.208(1 - \phi) + 0.094$
Interconnects (2 nm)	$109.7(1 - \phi)^2 + 34.5(1 - \phi) - 1.2$	$-0.101(1 - \phi)^2 + 0.221(1 - \phi) + 0.080$
Interconnects (3 nm)	$80.8(1 - \phi)^2 + 48.9(1 - \phi) - 4.2$	$-0.09 - (1 - \phi)^2 + 0.139(1 - \phi) + 0.103$
Interconnects (5 nm)	$2.6(1 - \phi)^2 + 60.5(1 - \phi) - 4.7$	$1.255(1 - \phi)^2 - 1.246(1 - \phi) + 0.346$

Table 4.2: Interpolation expressions from unit cell models relating porosity (ϕ), Young's modulus and Poisson's ratio.

4.4.4 Finite element mode of full sample

By solving the full sample finite element model, it is possible to determine which values of the metalattice porosity, Young's modulus and thermal conductivity are compatible with the SAW frequency and thermal relaxation timescale observed in the EUV experiment. While a simulation involving the full set of thermoelastic equations (Eq. 3.2, Eq. 3.3, Eq. 4.6 and Eq. 4.7) would capture both the thermal relaxation and SAW oscillations, the inertial coupling between the thermal and elastic equations makes this a computationally expensive proposition. Instead, the SAW frequency can be estimated using an eigenfrequency calculation, while the thermal relaxation can be simulated




Figure 4.7: Finite element model of metalattice unit cell: (a) The unit cell model represents a single period of the metalattice geometry and calculates its acoustic dispersion, whose slope provides the longitudinal (V_l or c_L) and transverse or shear (V_s or c_T) acoustic velocities. (b) Because the average interconnectedness of the metalattice pores is unknown, the simulations consider five geometries: fully isolated spherical pores (dark blue), overlapping spherical pores (light blue) and spherical pores connected by cylindrical channels of varying radius (green, yellow, red). By varying porosity and using relations from continuum mechanics to convert V_l and V_s into Young's modulus and Poisson's ratio, it is possible to relate the metalattice elastic properties with porosity for each geometry. Using these relations, porosity, Young's modulus and Poisson's ratio can be treated as a single fit parameter. Figure adapted from [68].

in the absence of acoustic oscillations using the quasi-static approximation discussed below.¹ While the eigenfrequency and quasi-static thermal simulations share the same geometry and material parameters, slight tweaks customize the finite element model for each case.

The full sample model represents a single period of the nickel grating and underlying substrate and treats the metalattice as an effective medium with a uniform Young's modulus, Poisson's ratio and thermal conductivity. Fig. 4.8 shows the geometry, which involves a stack of the nickel structure, polysilicon layer, metalattice and crystalline <100> silicon substrate. Technically, both the polysilicon and nickel will have ~1 nm native oxide layers. The thermal simulations are only impacted by the native oxide between the nickel and silicon, which can be modeled as part of the thermal boundary resistance discussed below, while the silicon native oxide layer slightly impacts the SAW frequency in the acoustic simulations for the smaller period gratings and is therefore included as shown in Fig. 4.8. The material parameters, including density, elastic constants, thermal expansion coefficient and heat capacity used in these simulations appear in Table 4.5.

The defining characteristic of a finite element model is its mesh, which subdivides the simulation domain into a number of smaller areas. A variety of mesh designs may work equally well for a particular finite element simulation, as can be tested by varying the many mesh parameters in COMSOL. A general rule of thumb is that when initially building a finite element model, the mesh should be made simpler, with fewer elements, until the simulation results begin to deviate from those achieved with a more complex mesh.² The required mesh complexity depends on the geometry and equations to be solved. In the case of the eigenfrequency simulations, this is the eigenvalue problem obtained by assuming a harmonically oscillating solution to Eq. 4.6, written as

$$-\rho\omega^2\mathbf{u} = \nabla \cdot \boldsymbol{\sigma}, \quad (4.14)$$

¹In COMSOL, the eigenfrequency calculation uses the solid mechanics module, while the thermal simulation uses both the solid mechanics and the heat transfer in solids modules. In COMSOL versions 4.3 and earlier, thermal simulations only require the thermal stress module, reducing license requirements.

²Finite element models of the metalattice sample, along with the meshes used for the simulations in Refs. [10] and [68] are in `home/brendan/Empty_Metalattices/comsol-models/` on Tesseract and backed up on the Nanothermal Dropbox.




Figure 4.8: Full sample finite element model: (a) Geometry and mesh used in the eigenfrequency simulations of one period of the $L = 100$ nm, $P = 400$ nm nickel grating. (b) Zoomed-in geometry and mesh showing the nickel structure, silicon overlayer and thin native oxide layer. The native oxide layer is only included in the eigenfrequency simulations. (c) Simulated temperature profile for a quasi-static thermal simulation some time after the heat pulse given by Eq. 4.16 is injected into the nickel structure. The pink line represents the thermal boundary resistance between the nickel and silicon. (d) Simulated displacement profile from an eigenfrequency calculation showing the mode identified by the SAW-likeness analysis. Figure panel (d) adapted from [68].

where ρ is the material density, $\mathbf{u}$ is displacement and $\boldsymbol{\sigma}$ is the stress tensor, which relates back to strain and displacement through Eq. 4.4 and Eq. 4.5. Eq. 4.14 has as its solution all eigenmodes, indexed with m , with frequencies ω_m . This calculation is quick and generally requires less mesh complexity, although the substrate depth sets the spacing between modes in frequency space. The presence of nickel structures complicates identification of the SAW frequency by broadening its bandwidth in frequency space. The resulting eigenfrequency spectrum can feature multiple modes adjacent in frequency space which all appear SAW-like. Fortunately, the SAW-likeness analysis developed in Ref. [153] can help identify the one true SAW, by calculating the fraction of energy localized near the surface in each mode. If $\mathbf{u}(\mathbf{r})$ is the local displacement associated with mode m and $\rho(\mathbf{r})$ is the local density, then the SAW-likeness coefficient is the ratio of a mode's mechanical energy near the surface or in the nanostructure to its total mechanical energy,

$$\alpha_{\text{SAW}}(\omega_m) = \frac{\int_{\text{nanostructure}} \rho(\mathbf{r}) \mathbf{u}_m^2(\mathbf{r}) d^3\mathbf{r} + \int_{\text{surface}} \rho(\mathbf{r}) \mathbf{u}_m^2(\mathbf{r}) d^3\mathbf{r}}{\int_{\text{total}} \rho(\mathbf{r}) \mathbf{u}_m^2(\mathbf{r}) d^3\mathbf{r}}. \quad (4.15)$$

By launching an eigenfrequency simulation, calculating SAW-likeness coefficients, identifying the frequency of the mode with the highest SAW-likeness coefficient, and comparing to the experimentally-measured frequency, the simulated elastic constants can be fit to those that best match the experiment. While the eigenfrequency simulations and SAW-likeness analysis are computationally very efficient, they can be tricked by the presence of other surface-confined modes, such as Sezawa waves, or longitudinal modes in thin films [151]. These modes can appear in the simulations but not be activated or observed in the experiment. The first check when identifying a particular mode in an eigenfrequency simulation as the experimentally-relevant SAW is to use the COMSOL software to plot the associated displacement. A SAW will have the distinctive displacement profile shown in Fig. 4.8(d). Additional verification is possible via the brute force method of launching a full inertial simulation in the time domain, and then extracting the SAW frequency using the same chirp Z-transform procedure as in the experimental analysis. This requires a long simulation of

multiple oscillation periods to get high frequency resolution, which can be avoided by using the projection method discussed in Ref. [59]. Fig. 4.9(c) shows a validation of the SAW-likeness analysis for the metalattice finite element model using both the full inertial and projection approaches.

The thermal simulations solve the thermoelastic equations in the time domain. In the quasi-static approximation, the inertial term $\partial^2 \mathbf{u} / \partial t^2$ in Eq. 4.6 is set to zero, which means that the medium instantly responds to a change in temperature through linear thermal expansion. This approximation greatly simplifies the calculation, and reduces the necessary mesh complexity. However, comparison of a quasi-static simulation result to experimental data in which acoustic oscillations are present can be a complicated endeavor. A combination of two approaches may be used to deal with this issue. First, oscillations in the experimental data may be isolated and removed using an analysis technique known as the matrix pencil method (MPM), discussed in the Supplementary Information of Refs. [10] and [67] and summarized briefly here. The MPM algorithm begins by constructing a time-lagged covariance matrix from the experimental diffraction efficiency signal. It then performs a process known as singular value decomposition, wherein the covariance matrix is decomposed onto a basis that allows for separation of the time-varying components of the experimental signal from the time-invariant background noise. By fitting a complex exponential $e^{i\omega t}$ to each time-varying component and subtracting the oscillatory components, where $\text{Im}(\omega) = 0$, the MPM filters acoustics from the experimental signal, leaving only the monotonic quasi-static decay. The MPM algorithm can outperform more traditional methods like a fast Fourier transform in situations where oscillations are highly damped or partially obscured by noise. However, even with acoustic oscillations removed, the experimental and simulated scans may not be directly proportional at short timescales due the initial inertial expansion of the nickel structure not captured by the simulation. This can be accounted for by starting any fits after a certain pump-probe delay time, generally tens of picoseconds, or by fitting an amplitude alongside the thermal parameter of interest. In cases where the experimental data has high noise, the amplitude fit may be unavoidable, since it does not make sense to normalize the simulation to the noise peak in the experimental scan.

The time domain thermal simulations calculate the evolution of an initial nonequilibrium




Figure 4.9: SAW-likeness analysis: (a) Sample EUV scatterometry data at different nickel grating periods in the time domain (top) and as a Fourier transform (bottom), showing the fundamental and second harmonic surface acoustic wave (SAW) frequencies. The grating period sets the SAW wavelength, enabling a sampling of its dispersion through the measurement of gratings of varying period. (b) The metalattice porosity and elastic properties are fit to the EUV scatterometry experiment by first identifying which eigenmode in the full sample FEA simulation corresponds to the experimentally-observed SAW using a SAW-likeness analysis, and then determining which parameter values reproduce the experimentally-observed SAW frequencies. (c) Validation of the SAW-likeness analysis at two porosity values using the projection method from Ref. [59] and by taking the chirp Z-transform of the surface oscillations predicted by a time-domain full inertial simulation. Figure adapted from [68].

temperature profile meant to represent the injection of energy by the ultrafast pump pulse into the nickel structures. Different circumstances may require differing levels of complexity in modeling this ultrafast excitation. If the short timescale (sub-picosecond) dynamics are of interest, it is advisable to simulate both the electron and phonon temperatures of the nickel using a two-temperature model of the type discussed in Section 2.3. Details on the two-temperature model used in past simulations of the EUV scatterometry experiment, including the temperature-dependent electron-phonon coupling coefficients, appear in Ref. [25]. Because we are interested in the metalattice thermal behavior, which occurs on nanosecond timescales, these complexities are unnecessary here. Instead, the thermal metalattice simulations only include a phonon temperature and phenomenologically model the heat injected into the nickel as

$$\mathcal{Q}(t) = \mathcal{Q}_0 e^{-t/\tau_{\text{ep}}}, \quad (4.16)$$

where $\tau_{\text{ep}} = 450$ fs is an experimentally-measured equilibration timescale for electrons and phonons in a nickel film excited by an ultrafast infrared laser [155]. While in a two-temperature simulation, the electrons and phonons have different heat capacities, the single nickel heat capacity here accounts for both (see Table 4.5).

The evolution of the temperature profile after heat is injected into the nickel is influenced by the inclusion of a thermal boundary resistance (TBR) between the nickel and silicon overlayer, as indicated by the pink line in Fig. 4.8(c). The concept of thermal boundary resistance originates from the so-called Kapitza resistance observed at the interface between liquid helium and solids [156], but now refers the temperature drop associated with imperfect thermal coupling at the boundary between any two materials

$$\Delta T = qR, \quad (4.17)$$

where q is the magnitude of the heat flux across the interface and R is the TBR in units of $\text{nK}\cdot\text{m}^2/\text{W}$. The term resistance mirrors the analogous voltage drop in an electrical system.

In the context of interfaces between crystalline solids, the various TBR models in the literature generally struggle to predict the TBR values observed in real systems, where interface roughness and imperfections can play a major role [157, 158]. This makes it difficult to determine the appropriate TBR between the nickel and silicon for the present model. However, the low metalattice thermal conductivity serves as the main bottleneck to the thermal relaxation of the nickel grating, limiting the effect of uncertainty in the TBR within reasonable bounds.

Boundary	Thermal condition	Thermal equation
Top	Insulating	$\mathbf{n} \cdot \mathbf{q} = 0$
Bottom	Insulating	$\mathbf{n} \cdot \mathbf{q} = 0$
Sides	Periodic	$T_1 = T_2, \mathbf{n}_1 \cdot \mathbf{q}_1 = -\mathbf{n}_2 \cdot \mathbf{q}_2$

Table 4.3: Boundary conditions for thermal finite element simulations. The variables $\mathbf{q}$, T and $\mathbf{n}$ refer to the heat flux, temperature and the boundary normal vector, respectively. The subscripts refer to the two boundaries involved in the periodic boundary condition.

Boundary	Elastic condition	Elastic equation
Top	Free	
Bottom	Fixed or Low Reflecting	$\mathbf{u} = 0$ or $\boldsymbol{\sigma} \cdot \mathbf{n} = -d_{\text{im}} \frac{\partial \mathbf{u}}{\partial t}$
Sides	Periodic	$\mathbf{u}_1 = \mathbf{u}_2$

Table 4.4: Boundary conditions for acoustic finite element simulations. The variables $\mathbf{u}$, $\boldsymbol{\sigma}$ and $\mathbf{n}$ refer to the displacement, stress tensor and the boundary normal vector, respectively. A low reflecting boundary condition is only necessary in a time domain simulation and tailors the mechanical impedance d_{im} to minimize the amplitude of reflected waves, which could otherwise lead to nonphysical interference effects.

Because the 2D finite element model captures only one period of the nickel grating geometry and a finite substrate depth, it is necessary to impose boundary conditions on all four edges as summarized in Tables 4.3 and 4.4. The top of the simulation domain is an interface with vacuum, assumed to be free in the acoustics and insulating in the thermal simulations, neglecting the effects of thermal radiation on the short experimental timescales. The left and right side of the simulation domain are subjected periodic boundary conditions, equivalent to the constraint that the finite

element solution be identical across every period of the nickel grating. The main difficulty arises with the bottom of the simulation domain. It is impractical to simulate the entire silicon substrate in COMSOL, since this would require a huge disparity in mesh element size between the bottom of the substrate and the area of interest near the nickel grating. Instead, the model boundary is located 20 and 1 μm below the sample surface, and subjected to either fixed or low reflecting and insulating boundary conditions for the acoustic and thermal simulations, respectively. The adequacy of the simulated substrate size should be verified by increasing it and checking that there is no significant impact on the simulation results.

4.4.5 EUV diffraction modeling

The thermal finite element model calculates the displacement from equilibrium of each spatial mesh point at each simulated time due to the thermal expansion of the surrounding system. Selecting only those points on the sample surface, this gives a profile of thermal surface deformation due to the simulated pump excitation. Although thermal expansion varies linearly with temperature through Eq. 4.7, the relationship between surface deformation and the experimentally-measured EUV diffraction signal is in general nonlinear. A home-built diffraction model therefore simulates diffraction from a deformed surface profile, facilitating a quantitative comparison between the finite element model and experiment.

Given the ~ 6 cm distance between the sample and CCD camera, the scalar diffraction amplitude, $U(x)$, of the EUV probe some distance z after the 1D nickel grating is well described by the Fresnel approximation,

$$U(x) = \frac{e^{ikz}}{i\lambda z} e^{i\frac{k}{2z}x^2} \int_{-\infty}^{\infty} \left\{ U(\xi) e^{i\frac{k}{2z}\xi^2} \right\} e^{-i\frac{2\pi}{\lambda z}x\xi} d\xi \quad (4.18)$$

where λ is the probe wavelength, k is the probe wavevector, ξ is the spatial coordinate on the object plane and x is the spatial coordinate on the image plane [162]. For a static grating of linewidth L , period P and height h , the scalar amplitude at the object plane is

Parameter	Value	Source
ρ_{Ni}	8.9 g/cc	[159]
$c_p \text{ Ni}$	444 J/kg/K	[159]
κ_{Ni}	90.9 W/m/K	[159]
α_{Ni}	12.77e-6 K^{-1}	[160]
E_{Ni}	219 GPa	[59]
ν_{Ni}	0.31	[59]
ϕ	0.385 ± 0.02	[68]
ρ_{ML}	$2.33(1-\phi) = 1.43 \text{ g/cc}$	[159]
$c_p \text{ ML}$	712 J/kg/K	[159]
α_{ML}	$2.6\text{e-}6 \text{ } K^{-1}$	[161]
E_{ML}	66.4 GPa	[68]
ν_{ML}	0.19	[68]
ρ_{Si}	2.33 g/cc	[159]
$c_p \text{ Si}$	712 J/kg/K	[159]
κ_{Si}	149 W/m/K	[159]
α_{Si}	$2.6\text{e-}6 \text{ } K^{-1}$	[161]
E_{Si}	160 GPa	[154]
ν_{Si}	0.22	[154]

Table 4.5: Finite element model material parameters. Values for nickel (Ni), the metalattice (ML) and silicon (Si), including porosity (ϕ), density (ρ), specific heat capacity (c_p), thermal conductivity (κ), coefficient of thermal expansion (α), Young's modulus (E) and Poisson's ratio (ν).

$$U_0(\xi) = r_S + (r_G e^{-i\phi_0} - r_S) \text{comb}_P(\xi) \star \text{rect}\left(\frac{\xi}{L}\right), \quad (4.19)$$

where $r_S = |r_S|e^{-i\phi_S}$ and $r_G = |r_G|e^{-i\phi_G}$ are the complex substrate and grating reflectivities, respectively, and the symbol $\star$ denotes a convolution. The geometric phase $\phi_0 = 4\pi h \cos \theta / \lambda$ accounts for the change in optical path length due to the static grating height for light incident at an angle θ from normal. To determine the values of r_S and r_G , we use the IMD code written by Yuka Esashi, which uses the Parratt formalism to compute the complex EUV reflectivity of a

multilayer stack [163]. The diffraction intensity profile at the image (camera) plane for the grating profile in Eq. 4.19 is

$$I(x) = |U(x)|^2 = \mathcal{R}_0 \delta^2(x) + \mathcal{R}_1 \sum_{n \neq 0} \delta^2 \left(x - \frac{n\lambda z}{P} \right) \text{sinc}^2 \left(\frac{Ln}{P} \right), \quad (4.20)$$

where the sum is over both negative and positive $n \neq 0$ and the factors $\mathcal{R}_0$ and $\mathcal{R}_1$ are

$$\begin{aligned} \mathcal{R}_0 &= |r_S|^2 \left(1 - \frac{L}{P} \right)^2 + \frac{L^2}{P^2} |r_G|^2 + 2 \frac{L}{P} |r_G| |r_S| \cos(\phi_S - \phi_G - \phi_0) \left(1 - \frac{L}{P} \right), \\ \mathcal{R}_1 &= \frac{L^2}{P^2} \left(|r_G|^2 + |r_S|^2 - 2 |r_G| |r_S| \cos(\phi_S - \phi_G - \phi_0) \right). \end{aligned} \quad (4.21)$$

Eq. 4.20 has the same specular (DC) and diffracted (AC) peaks as the experimentally-observed diffraction profiles in Fig. 4.5, and corresponds to diffraction from the static grating unperturbed by the infrared pump. In the pump on case, an arbitrary, small perturbation in the sample height $\Delta h(\xi, t)$ can be expanded in a Fourier series as

$$\Delta h(\xi, t) = \frac{a_0(t)}{2} + \sum_{n=1}^{\infty} a_n(t) \cos \left(\frac{2\pi n \xi}{P} \right). \quad (4.22)$$

This results in the scalar field

$$U(\xi, t) = r_S + (r_G e^{-i\phi_0} - r_S) \text{comb}_P(\xi) \star \text{rect} \left(\frac{\xi}{L} \right) \times \exp \left[-4\pi i \frac{\cos \theta}{\lambda} \sum_{n=1}^{\infty} a_n(t) \cos \left(\frac{2\pi n \xi}{P} \right) \right], \quad (4.23)$$

at the object plane, where the factor in brackets is the small correction to the static expression appearing in Eq. 4.19. The term involving $a_0(t)$ has been neglected, since it does not contribute towards the diffracted intensity. Plugging both Eq. 4.23 and Eq. 4.19 into Eq. 4.18, and keeping only terms of linear order in $a_n(t)/\lambda$, the normalized change in diffraction efficiency due to the infrared pump is

$$\Delta I(x, t) = \frac{4\pi \cos \theta}{\lambda} \mathcal{R} \sin(\phi_S - \phi_G - \phi_0) \sum_{n \neq 0} \left\{ a_{|n|} \text{sinc} \left(\frac{Ln}{P} \right) \left[\delta^2(x) - \delta^2 \left(x - \frac{\lambda z n}{P} \right) \right] \right\}, \quad (4.24)$$

where the sum is over negative and positive $n \neq 0$ and the factor $\mathcal{R}$ is

$$\mathcal{R} = |r_S| |r_G| \left[|r_G|^2 + |r_S|^2 \left(\frac{P}{L} - 1 \right) \right]^{-1}. \quad (4.25)$$

The first term inside the square brackets in Eq. 4.24 is the DC signal and the second term represents the AC diffracted orders. Details on this rather lengthy derivation and an additional derivation for the more complex case of a time-dependent grating reflectivity appear in the Analysis/Spatial Frequency folder on Dropbox.³ Eq. 4.24 reveals that the EUV diffraction patterns contain important information about the spatial temperature profile, which is lost during the typical procedure of summing over the diffracted orders to aggregate the experimental signal. To linear order, the n th diffracted order depends only on a_n , which accounts for the n th spatial frequency of the surface deformation profile. By capturing sufficient diffracted orders on the CCD camera, it should therefore be possible to spatially reconstruct the sample surface deformation and temperature profile in the EUV experiment. Extracting the dynamic evolution of a nanoscale temperature profile with high spatial resolution would add a new dimension to the EUV scatterometry experiment, which could provide insight into theoretically-predicted temperature dips adjacent to nanoscale heat sources, or even reveal the wavelike propagation of heat pulses.

Rather than rely on these linearized analytical expressions, the present experiment uses a numerical Python code to translate the COMSOL surface displacements into a simulated EUV diffraction efficiency (see Appendix A). This code relies on specific input file formats from COMSOL and IMD and is capable of quickly calculating the diffraction efficiencies for large COMSOL parameter sweeps. Options in the code allow the user to choose the number of AC orders included in the

³The expressions in Eqs. 4.20 and 4.24 have similar functional forms in both the Fresnel and Fraunhofer approximations, with the Fresnel solution involving the convolution of each delta function with a quadratic phase, which is omitted here for simplicity.

calculation to match the number captured in the experiment, and include or exclude the change in reflectivity of the nickel as it thermally expands. The importance of considering a change in nickel reflectivity can be checked by summing the experimental diffraction profile across the camera chip without flipping the sign of the AC orders relative to the DC as in Fig. 4.5(b). If the signal mostly disappears in this case, this is an indication that the infrared pump does not induce a significant change in the nickel reflectivity. In most experiments, the nickel temperature rise is much higher than that of the substrate, as heat is bottlenecked behind the thermal boundary resistance at the interface. For this reason, it is generally safe to assume that any change in substrate reflectivity is a secondary effect.

The EUV diffraction efficiency can depend greatly on the complex reflectivities, nickel structure height, angle of incidence and wavelength, which combine to form the factor $\phi_S - \phi_G - \phi_0$ in Eq. 4.24. If this factor is approximately an integer multiple of 2π , such that $\sin(\phi_S - \phi_G - \phi_0)$ sweeps through a zero between the different wavelengths in the harmonic comb, the result is a very messy and nonlinear relationship between the EUV diffraction efficiency and underlying temperature profile. For example, given a diamond substrate and harmonic comb extending from approximately 25 to 32 nm and a 45 degree probe incidence angle, the typical 12 nm nickel grating works well, while a ~ 10 nm height severely complicates the experiment, as was observed for the diamond sample 5963. The effects of an inconvenient nickel grating height can be fixed by slightly adjusting the incidence angle of the EUV probe.

4.4.6 Frequency-domain analysis of porosity and elastic properties

Using the full sample finite element model, SAW-likeness analysis and the isolated unit cell model, the metalattice porosity and Young's modulus can be fit to $38.5 \pm 2\%$ and 66 ± 3 GPa. Because only one acoustic frequency is measured, the fit quality is relatively insensitive to the Poisson's ratio, since for an isotropic material E and ν are functions of both c_L and c_T as discussed in Section 4.4.2. The fitted values, shown as the blue bars in Fig. 4.10, use the experimental data from the 50-200 and 100-400 nickel gratings, which are most sensitive to the metalattice properties

due to their small periods. The fitted porosity agrees well with the value of $40 \pm 2\%$ extracted from electron tomography. Fig. 4.10 also includes results based on the unit cell model with 2 nm interconnections (green bars), which predicts a lower porosity. The remaining unit cell models, with 3 nm and 5 nm interconnections and overlapping pores, fit a porosity so low as to be incompatible with the assumed geometry and so are discarded.

The larger error bars for the 50-200 grating are partially attributable to the uncertainty in the metalattice thickness, which is known more precisely for the 100-400 grating as it is the location of one of the lamellae extracted for the electron microscopy (see Fig. 4.6). Due to a coupling effect between the longitudinal acoustic modes propagating in the metalattice film and SAW modes, the SAW mode frequency depends on the metalattice thickness, especially for the 50-200 grating. The ability of the SAW-likeness analysis to reproduce the correct magnitude of this coupling effect is confirmed with a full inertial simulation.

The intermediate grating size, with 800 nm period, launches a surface acoustic wave with a penetration depth sufficiently large to be sensitive to the metalattice thickness, allowing it to be fit as shown in Fig. 4.10(b). The fitted value of 515 ± 87 nm falls between the thicknesses measured using electron microscopy for the 30-120 and 100-400 gratings located on either side, consistent with a monotonic variation in metalattice thickness across the sample surface due to the vertical deposition process. Finally, SAWs launched by the larger 500-2000 and 1000-4000 gratings are sensitive primarily to the silicon substrate properties, confirming the $\langle 100 \rangle$ orientation.

Unlike porosity, the fitted Young's modulus is relatively insensitive to the use of the 2 nm interconnected vs. isolated unit cell model, consistently fitting to a value ~ 65 GPa. This is true even if porosity and density are fixed using the 40% electron tomography value, Poisson's ratio is set to 0.19 based on the isolated pore unit cell calculation in Table 4.2 and Young's modulus is fit on its own. The measured Young's modulus is consistent with or perhaps slightly higher than previous experimental measurements of mesoporous silicon at 40% porosity, which report Young's moduli in the 50-60 GPa range [164, 165]. Mesoporous silicon, which consists of <50 nm irregular, interconnected pores, is well described by the effective medium model developed by Gibson and




Figure 4.10: Metalattice porosity, thickness and elastic properties extracted via EUV scatterometry: (a) Metalattice porosity extracted from the electron tomography reconstruction and the EUV scatterometry experiment. The EUV scatterometry results are provided for both the isolated pores model and the model with 2 nm interconnections. The other interconnection models fit a porosity which is too low to satisfy their geometrical requirements and can therefore be ruled out. The EUV scatterometry results agree for the two nickel grating periods, 200 nm and 400 nm, for which the SAW is mainly confined to the metalattice film and therefore most sensitive to the metalattice properties. (b) Metalattice thickness extracted using analysis of the EUV scatterometry results for the 800 nm period nickel grating. The fitted thickness compares well to a linear extrapolation of the thickness measured at two points via the electron microscopy. A linear variation in thickness across the metalattice surface is expected from the vertical deposition step in the fabrication process [135]. (c) and (d) Young's modulus and Poisson's ratio of the metalattice extracted using EUV scatterometry as compared to the bulk polycrystalline silicon values given in gray. Figure adapted from [68].

Ashby,

$$\frac{E_{\text{eff}}}{E_0} = C\bar{\rho}^2, \quad (4.26)$$

where E_0 is the Young's modulus of the constituent material, $\bar{\rho} = 1 - \phi$ is the relative density, or fraction of the bulk density, and C is an empirical constant close to unity [166]. At 40% porosity, using $C = 1$ and $E_0 = 160$ GPa for bulk polysilicon, the Gibson-Ashby model predicts a Young's modulus of 57 GPa. The measured value of 66 ± 3 GPa for the metalattice is somewhat higher and closer to the Hashin-Shtrikman upper bound on the effective Young's modulus of an isotropic cellular solid,

$$\frac{E_{\text{eff}}}{E_0} = \frac{2\bar{\rho}(7 - 5\nu)}{15(\bar{\rho} - 1)\nu^2 + 2(\bar{\rho} - 6)\nu - 13\bar{\rho} + 27}, \quad (4.27)$$

which is 69 GPa for porous silicon at 40% porosity [167, 168]. This suggests that the regular FCC pore arrangement slightly boosts the metalattice stiffness. Poisson's ratio is set to 0.19 for this calculation, although despite its appearance, Eq. 4.27 is not very sensitive to the exact value of ν . The finding that the metalattice elastic properties can be described using classical effective medium theories supports the conclusions of a previous investigation of a silica-filled metalattice sample, and can be explained by the long acoustic wavelength (≥ 200 nm) relative to the microscopic geometrical detail (< 40 nm) [66]. In this limit, acoustic waves experience Rayleigh scattering with the pores, a process which is insensitive to their exact geometry.

4.4.7 Time-domain analysis of thermal conductivity

The apparent metalattice thermal conductivity can be fit to $\kappa = 1.6 \pm 0.4$ W/(m·K) by searching for the value at which the results of thermal finite element and EUV diffraction simulations match the experimental data. Prior to the fits, the acoustics are removed from the experimental data using the MPM algorithm described above. The experimental amplitude is normalized to one and the fitted amplitude is floated to account for noise spikes, as discussed in the Supplementary

Information of Ref. [10]. Three fitting algorithms, MATLAB’s GlobalSearch, lsqnonlin and nlinfit are tested and deliver the same results. The code used for these fits is versatile in that it allows for the fitting of multiple parameters simultaneously (*e.g.*, thermal conductivity and thermal boundary resistance) and the selective imposition of constraints across groupings of experimental scans in the dataset (see Appendix A). The fitted thermal conductivity is two orders of magnitude lower than bulk silicon and far below the predictions of the classical volume reduction models of Eucken and Russell introduced in Section 4.1.

Linewidth [nm]	Period [nm]	Fitted Conductivity W/(m·K)
50	200	1.0 ± 0.3
100	400	1.5 ± 0.5
100	800	1.5 ± 0.6
500	2000	2.2 ± 1.1
1000	4000	4.2 ± 1.9
Aggregate		1.6 ± 0.4

Table 4.6: Fitted thermal conductivity of the silicon metalattice. Each experimental scan is fit separately, and results are aggregated by nickel grating and over the entire dataset. The final row excludes the 12 experimental scans for which fitted conductivity declines with fit start time.

The fitted thermal conductivity is independent of the nickel grating geometry within the error bars, with the exception of the 1000-4000 grating, as summarized in Table 4.6. However, within the error bars, there is a noticeable decrease in the central fit value with decreasing grating linewidth. This may be the same effect observed for nanoscale nickel heaters on a bulk silicon substrate in Ref. [64]. However, the effect of the metalattice layer and the size of the error bars complicates the interpretation in the present experiment. No indication of a close-packing effect is observed, even for the 50-200 geometry, likely because the nanostructured pores scatter thermal phonons from adjacent heaters before they can interact.

The larger fitted conductivity for the 1000-4000 grating also depends on the fit start time, decreasing when the fit is started later, as shown in Fig. 4.11. This dependence of fitted thermal




Figure 4.11: Thermal conductivity vs. fit start time and error calculation for the (a) 1000-4000 and (b) 500-2000 nickel gratings. Error comes from uncertainty in the finite element model parameters (red vs. blue curves), variation between experimental scans (shading), and changes in the fitted conductivity with fit start time. The three sources of error are added in quadrature to determine the overall error. Diffraction error was found to be relatively insignificant and excluded. The fitted conductivity for the 1000-4000 grating and some of the 100-800 data depends on fit start time, while the fitted conductivity is approximately constant with fit start time for the remaining nickel gratings. Figure adapted from [10].




Figure 4.12: Fourier finite element model fit quality for silicon and metalattice samples. (a) The exponential plus power law form predicted by Fourier's law (green) for a heated nickel nanostructure on a silicon substrate cannot capture the slow decay observed at large delay times in the experiment. Instead, the experimental data exhibits a double-exponential decay predicted by a phonon hydrodynamic model (red). (b) The thermal decay dynamics in the metalattice sample are well explained by a Fourier finite element model with a low apparent thermal conductivity. Both panels show the raw experimental data (light gray) and data with acoustics removed using the MPM algorithm (black). Figure adapted from [67].

conductivity on fit start time also appears for the four 8 ns 100-800 traces. For this reason, we hypothesize these measurements may be affected by inertial effects, variations in the silicon overlayer, or interface quality, and exclude them from the aggregate thermal conductivity fit quoted above. The fitted thermal conductivity of the remaining 35 of 47 experimental measurements does not depend on fit start time, as shown in Fig. 4.11. This differs from the case of nanoscale heaters on bulk silicon, discussed in Ref. [67]. In this case, the functional form of the nickel temperature predicted by a 2D Fourier model for a single heater is

$$T_{\text{Ni}}(t) \sim A_1 e^{-t/(c_{\text{Ni}} h R)} + \frac{A_2}{\kappa_{\text{Si}} t}, \quad (4.28)$$

where c_{Ni} is the nickel specific heat, h is the nickel structure height, R is the thermal boundary resistance at the nickel-silicon interface, κ_{Si} is the thermal conductivity of the silicon substrate and A_1 and A_2 are constants. The functional form of Eq. 4.28 does not match the experimental data for nanoscale nickel gratings on a bulk silicon substrate, but works well for similar nickel gratings on the metalattice sample, as illustrated in Fig. 4.12. The observation that phonon transport dynamics in the metalattice obey a Fourier model motivates the theoretical model presented in Section 4.5.

4.4.8 Error bar calculations

Determination of the error bars quoted in the previous two sections can be a tedious process, involving additional finite element and diffraction simulations to determine the effect of various uncertainties on the fit results. The present section briefly outlines the procedure used to determine error bars on the metalattice porosity, elastic properties and thermal conductivity in Refs. [10] and [68], with the aim of presenting a general procedure for future studies. Details specific to each analysis appear in the relevant Supplementary Information documents.

The error calculation begins by taking the standard deviation of the individual fits to each experimental measurement. In the case where all experimental measurements can be said to be drawn from the same statistical distribution, the standard deviation σ can be replaced by the

standard error of the mean

$$\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{N}}, \quad (4.29)$$

where N is the number of measurements. However, EUV scatterometry measurements are generally made over multiple days under slightly different experimental conditions. If, for example, several measurements are made on a day with a different harmonic spectrum due to shift in gas pressure, the variation in fits to this subset of measurements comes from a different statistical distribution than that of the overall dataset. For this reason, use of the standard deviation is generally more appropriate. When N is small, it is also important to remember to use the sample rather than population standard deviation.

In cases where some measurements in the experimental dataset have significantly higher noise than others, such that fits to these measurements appear less trustworthy, it may be appropriate to weight according to noise when calculating the mean and error of a fitted parameter across the dataset. A means of quantifying relative experimental noise is to normalize each scan and then take the standard deviation of the signal before time zero. A weighting method is given by Eqs. 4.17 and 4.19 in Ref. [169] as,

$$\bar{x} = \frac{\sum_i (x_i / \sigma_i^2)}{\sum_i (1 / \sigma_i^2)}, \quad (4.30)$$

$$\sigma_{\bar{x}} = \frac{1}{\sum_i (1 / \sigma_i^2)}, \quad (4.31)$$

where $\bar{x}$ and $\sigma_{\bar{x}}$ are the weighted mean and standard error of the mean, respectively, and x_i and σ_i are the fit values and standard deviations associated with the individual measurements. Eq. 4.31 should be approached with caution, since the σ_i refer to the error in the fit parameter, rather than some relative uncertainty metric like the pre-time zero experimental noise. While the fit parameter error can be estimated from the covariance matrix in the fitting algorithm, Eq. 4.31

assumes that all measurements are drawn from the same statistical distribution, which may lead to an underestimation of the error.

The finite element model and diffraction simulations introduce additional error due to uncertainties in their material and geometrical parameters. Simulation parameters fall into four categories, depending on their certainty and sensitivity. An example of a certain, yet sensitive parameter is the silicon substrate density. Changing it would dramatically alter the metalattice Young's modulus fit, but the density of crystalline silicon is well known and therefore does not contribute significantly to the modeling error. An example of an uncertain yet insensitive parameter is the native oxide thickness on the sample surface. Regardless of whether it is 1 nm or 2 nm, it has a relatively minor impact on the results. Parameters included in the error bars are generally both uncertain and sensitive, such as the thermal boundary resistance between nickel and silicon. In the metalattice thermal conductivity fits, the TBR is allowed to vary between 2 and 8 nK·m²/W based on various experimental measurements [64, 67, 170], and contributes significantly to the error bars. In determining which model parameters are sufficiently uncertain and sensitive to contribute to the error bars, it is usually possible to draw the line at a cluster of parameters whose variation is far more significant than all the others. Even if other parameters vary slightly, any error they introduce will be dwarfed by the main contributors. Once the significant model parameters have been identified, the uncertainty of each must be determined to quantify its contribution to the error. This is often done using AFM or SEM for geometrical dimensions, and literature values for the more obscure material parameters. The contributions of the significant model parameters to the error can then be aggregated in a number of ways. The simplest approach is to set all parameters to the extreme values within their uncertainties necessary to produce the highest or lowest fit value. This is a conservative approach which likely overestimates the error in the case of many parameters. An alternative approach when the simulations are not too computationally expensive, as in the case of eigenfrequency calculations, is to stochastically simulate a large number of parameter combinations, obtain the fit value for each and then draw up a confidence interval.

Inclusion of error from diffraction simulations poses a particular challenge due to the fact

that the EUV probe consists of multiple wavelengths whose relative magnitudes depend on the high harmonic generation process. There is no straightforward way to measure the relative magnitudes of the various harmonics, since they arrive at the CCD camera after reflecting from the sample and passing through additional aluminum filters. The typical procedure to account for diffraction error involves repeating the diffraction simulations for the 25th, 27th, 29th and 31st harmonics of the ~ 800 nm infrared pump and observing the change in the fits. If this change is large, then the standard deviation can be taken between the various wavelengths or experimentally-realistic combinations thereof. It is also likely that the probe incidence angle may deviate by <2 degrees from its nominal value of 45 degrees, making it important to check the sensitivity of the fits to this angle. The error introduced by uncertainties in the diffraction simulations is generally minor, unless the diffraction occurs near a zero in the function $\sin(\phi_S - \phi_G - \phi_0)$ in Eq. 4.24, in which case it can be large. In this case, it can be helpful to examine the magnitudes of the simulated diffraction signal for various angle and wavelength combinations. Simulations of wavelengths close to the zero in Eq. 4.24 generally lead to large changes in fit values, but also have low diffraction efficiency, making it unlikely that they are strong contributors to the experimental signal.

The time window of fits to thermal simulations may also contribute error. This can be systematic, as discussed in Section 4.4.7, or due to the inclusion or exclusion of a particular noise spike. Starting and ending the fits at various times should be tested as a consistency check and may be included in the error bars. The error of individual fit values calculated using the MATLAB algorithms `lsqnonlin` and `nlinfit` can be calculated from the covariance matrices supplied with the solution. This error should be checked and incorporated as necessary. In cases where it is significantly less than some other source of error it can be safely neglected.

The main consideration when combining error from different sources, such as the experimental, model, diffraction and fit error discussed above is whether the various sources of error are independent. In most cases they are, allowing the total error σ_{tot} to be calculated through their addition in quadrature as

$$\sigma_{\text{tot}} = \sqrt{\sigma_{\text{E}}^2 + \sigma_{\text{M}}^2 + \sigma_{\text{D}}^2 + \sigma_{\text{F}}^2}, \quad (4.32)$$

which is the method used in the metalattice analysis. Fig. 4.11 illustrates these various sources of error in the context of the metalattice thermal conductivity fits.

4.5 Phonon hydrodynamic model of heat flow in highly-confined nanosystems

The following section is an adapted version of the content of Ref. [10], which places the low apparent thermal conductivity and Fourier-like transport behavior observed in the silicon metalattice in the broader context of highly-confined nanostructured systems. Highly-confined in this context refers to systems in which boundaries are spaced much more closely than the mean free paths (MFPs) of thermal phonons in the bulk constituent material. For silicon, bulk phonon MFPs have a broad spectrum centered around 300 nm. Very few phonons with MFP below ~ 30 nm are predicted to contribute to thermal conduction in bulk silicon [80].

If, as in the ballistic framework presented in Chapter 3, phonons in a highly-confined system experienced perfectly bulk conditions except when colliding with a boundary, then virtually no internal phonon-phonon collisions would occur inside the metalattice, given its FCC arrangement of pores at 36 nm period. Instead, evidence from atomistic simulations, the EUV experiment described above and a number of other previous investigations suggests that, in the limit of extreme nanoscale confinement, phonons experience a dramatically modified phase space, which affects both their harmonic and anharmonic properties, including their mean free paths [106, 107].

To develop a mesoscale theory of phonon transport in these systems, we turn to the phonon hydrodynamics framework introduced in Chapter 3. Eq. 3.15,

$$\tau \frac{\partial \mathbf{q}}{\partial t} + \mathbf{q} = -\kappa \nabla T + \ell^2 [\nabla^2 \mathbf{q} + \alpha \nabla (\nabla \cdot \mathbf{q})], \quad (4.33)$$

known as the Guyer-Krumhansl equation (GKE), is derived under the assumption that the phonon distribution function relaxes quickly to a function of the heat flux and its immediate

derivatives. Previous work has demonstrated that nanoscale size and frequency effects on heat transport in silicon can be modeled in terms of non-local and memory effects, as described by the GKE, with *ab initio* geometry-independent coefficients τ , κ_{GK} and ℓ , with $\alpha = 1/3$ [67, 98, 171, 172]. However, the predictive capabilities of this approach are restricted to systems with dimensions sufficiently large compared to the average bulk phonon MFP [173], since nanoscale confinement modifies the *ab initio* phonon properties used to determine the GKE coefficients [107, 105]. This would seem to exclude the silicon metalattice sample.

We overcome this limitation by using the similarities between the GKE and the Navier-Stokes equation for mass transport in fluids. In analogy to Stokes flow in incompressible fluids at low Reynolds numbers, the GKE may be simplified by assuming that, in highly-confined nanosystems, the thermal gradient is almost fully compensated by the viscous resistance described by the Laplacian term [174]. This is equivalent to assuming that the rapid spatial fluctuations in heat flux necessary to accommodate the pores cause $\nabla^2 \mathbf{q}$ to dominate over the memory $\partial \mathbf{q} / \partial t$ and heat flux $\mathbf{q}$ terms, with $\nabla \cdot \mathbf{q} = 0$ due to the lack of heat sources inside the medium, reducing the GKE to the Stokes equation

$$\nabla T = \frac{\ell^2}{\kappa_{\text{GK}}} \nabla^2 \mathbf{q}. \quad (4.34)$$

In Eq. 4.34, the quantity $\mu \equiv \ell^2 / \kappa_{\text{GK}}$ with units of $\text{K} \cdot \text{m}^3 / \text{W}$ may be identified as the dynamic phonon gas viscosity, which models the strength of the interaction between the pore walls and the phonon population. Inside a porous medium such as the metalattice, Eq. 4.34 can be further simplified in analogy to Darcy's law for fluid flow in porous media as

$$\nabla T = -\frac{\mu}{K} \mathbf{q}, \quad (4.35)$$

where K is the metalattice permeability, which encompasses its geometrical properties. The derivation of Eq. 4.35 from Eq. 4.34 is analogous to the Kozeny-Carman approach for fluids [175] and appears in the Supplementary Information of Ref. [10]. It begins by solving the Stokes equation, Eq.

4.34, for the heat flux inside a cylindrical pipe of diameter D subject to a homogeneous temperature gradient, which represents a silicon channel surrounded by pores. The no-slip, $q(D/2) = 0$, boundary condition used when solving the Stokes equation for a macroscopic fluid is inappropriate here, since phonons colliding with the boundary will have originated a non-negligible distance inside the silicon and therefore have a non-negligible crystal momentum. Instead, we use the slip boundary condition commonly applied in the case of a rare gas

$$q(r = D/2) = -C\ell_{\text{slip}}\frac{\partial q(r = D/2)}{\partial r}, \quad (4.36)$$

where ℓ_{slip} is the characteristic distance from the boundary at which phonons originate and the parameter C is a specular parameter which varies between 1 and ∞ depending on the degree to which boundary collisions conserve crystal momentum [176]. At room temperature, boundary collisions are diffuse, or momentum-destroying, due to the small phonon wavelength compared to the surface roughness of experimentally-realizable phononic crystals, and in this context $C \approx 1$. After having calculated the heat flux in an individual channel, the derivation imagines the metalattice as a system of interconnected channels surrounding the pores, such that the pore and channel surface areas may be set equal. Accounting for various geometrical factors, the metalattice permeability is given as

$$K = \frac{d^2\Psi^2(1-\phi)^3}{24\Gamma^2\phi^2}, \quad (4.37)$$

where d is the average empty pore diameter, Ψ is the sphericity of the pores (*i.e.*, the surface area ratio of the ideal spherical pore to the actual pore), and Γ is the tortuosity (*i.e.*, the ratio of the average length of a heat flux streamline to the distance between its ends). In the derivation of Eq. 4.37, the slip length ℓ_{slip} is assumed to be $D/4$ based on the observation in atomistic simulations that it scales linearly with D in highly-confined geometries (details in Ref. [10]).

As observed in the EUV experiment, Eq. 4.35 predicts Fourier-like thermal transport in the metalattice, with an apparent thermal conductivity of

$$\kappa = \frac{K}{\mu}. \quad (4.38)$$

The apparent thermal conductivity can be interpreted as a ratio between a permeability K , which only depends on geometry, and the viscosity μ , which is determined by the average phonon properties in the highly-confined phase space.

The expression for the permeability in Eq. 4.37 is general to any highly-confined 3D system consisting of nanostructured material channels separated by voids. The 2D case is treated separately in the Supplementary Information of Ref. [10]. This facilitates the comparison of a broad variety of silicon nanosystems with varying pore distributions, sizes, shapes and periods. The framework presented here is only valid for sufficiently constrained systems, with characteristic sizes of a few tens of nanometers for silicon. It cannot be extended to larger-scale structured materials (*e.g.*, [177, 178]) at room temperature due to the breakdown of the Stokes regime assumption.

Fig. 4.13 shows the apparent conductivity κ , permeability K and viscosity μ as functions of porosity for new and previous MD simulations in metalattices [179] and 3D nanowire networks [180], along with present and past experimental results in metalattices [124], nanomeshes [126], and porous nanowires [130]. In Fig. 4.13(a), the conductivity comes directly from experimental measurements or the results of the atomistic simulations. The calculation of the permeability in Fig. 4.13(b) is straightforward for the atomistic simulations, where the input geometry is precisely defined, while in the case of experiments, parameters such as the sphericity Ψ , tortuosity Γ , pore size d and porosity are estimated from available electron microscopy images (details in Ref. [10]). The viscosity in Fig. 4.13(c) is then calculated by taking the ratio of K and κ as in Eq. 4.38. While the thermal conductivity and permeability at a given porosity depend strongly on the system-specific geometrical details, the viscosity depends solely on porosity, ϕ , across all nanostructured systems in both experiments and MD simulations. It scales as $(1 - \phi)/\phi$, the ratio of the available to non-available volume for heat transport. Combining this empirical relationship with Eqs. 4.37 and 4.38, one obtains an analytical expression for the apparent thermal conductivity in nanostructured




Figure 4.13: Universal thermal transport behavior in highly-confined nanostructured silicon systems: (a) Porosity vs. apparent thermal conductivity for a variety of experiments, including a high-porosity metalattice (green squares) [124], 2D nanomeshes (brown diamonds) [126], porous nanowires (pink triangles) [130] and the present experiment (red square). Also included are atomistic simulations of square-pore metalattices (black circles), circular-pore metalattices (blue circles) [179] and nanowire networks (yellow circles) [180]. While a higher porosity correlates with lower apparent conductivity, it fails to predict the precise value by over an order of magnitude. (b) Porosity vs. permeability, calculated using Eq. 4.37, where the geometrical parameters are inferred from electron microscopy images in the case of the experiments. (c) Porosity vs. viscosity, which scales empirically as $\mu \sim (1 - \phi)/\phi$ across all geometries. This behavior is reminiscent of a non-Newtonian fluid and accounts for intrinsic changes to harmonic and anharmonic phonon properties such as the reduction in velocities due to elastic softening and modifications to the intrinsic bulk scattering rates. Figure adapted from [10]. Insets from [124, 126, 130, 179, 180].

silicon systems with feature sizes far below the average bulk phonon mean free path.

Since viscosity depends only on porosity, the variations in thermal conductivity at a given porosity observed in Fig. 4.13(a) can be attributed to the permeability K . Permeability not only accounts for the classical volume reduction effect described by the Eucken and Russell models in Eqs. 4.1 and 4.2, but also captures nanoscale geometrical effects, and can be used to reinterpret correlations between thermal conductivities and geometric descriptors observed in the literature [181, 182, 183, 184]. For example, the neck size, or limiting dimension, of a system has been shown to be a good indicator of the effects of nanostructuring on thermal conductivity [85, 185]. The influence of the neck size on the permeability can be investigated by identifying the neck size with the channel diameter D and inserting the 3D geometrical relation, $D = 2(1 - \phi)\Psi d/3\phi$ into Eq. 4.37. The present model agrees with previous findings that thermal conductivity varies as the square of neck size under highly-confined transport conditions [184]. Other purely geometrical effects commonly invoked in phononic crystals like pore anisotropy [186], backscattering [177], or surface disorder [187] can be related to the permeability. Pore anisotropy might be described by the tortuosity Γ , and surface disorder may modify the slip flux boundary condition, which determines the numerical factor in Eq. 4.37. However, it is worth noting that it was not necessary to characterize the influence of these effects when calculating each sample's permeability in order to uncover the scaling of μ , suggesting that these are not the key mechanisms distinguishing the different cases considered. Finally, the resistive effect associated with backscattering in ballistic interpretations can be quantified using the geometrical parameters appearing in K , which depend on the size and frequency of channels orthogonal to the thermal gradient.

The dependency of viscosity on porosity observed here is reminiscent of a non-Newtonian fluid [188] and provides insight into the intrinsic mechanisms influencing phonon transport in highly-confined nanosystems.⁴ Specifically, it points to a modified dispersion and phase space for phonon-phonon interactions, in which phonon properties deviate from their bulk values. Previous

⁴Blood, for example, exhibits a lower viscosity in small capillaries than in larger channels, enabling it to move effectively throughout the circulatory system.

works have observed elastic softening effects in nanostructured systems, wherein phonon velocities decrease with increasing porosity, regardless of the specific geometric details [82, 130]. This effect would influence μ but not K in Fig. 4.13. Anharmonic effects such as phonon-phonon collision rates and phonon mean free paths may also be impacted by the proximity of boundaries. Beyond the simple truncation of phonon MFPs due to boundary scattering, this can occur through the modification of force constants in the interatomic potential due to surface proximity or through the inability of nonequilibrium distribution functions to quickly relax in the confined geometry, modifying phonon populations. For example, kinetic theory predicts that the thermal conductivity κ_{GK} and non-local length ℓ in a bulk system are related to the phonon MFP Λ as $\kappa_{\text{GK}} \propto \Lambda$ and $\ell^2 \propto \Lambda^2$ respectively, which implies that the viscosity μ should be proportional to Λ [98]. This would imply a relationship between viscosity and neck size, which is the primary geometric factor influencing the MFP. However, in Fig. 4.13(c), systems with the same porosities but different neck sizes are shown to have the same viscosity. The likely consequence of this apparent contradiction is that in a highly-confined nanosystem, MFPs depend on additional factors beyond those in the geometric permeability, limiting approaches which combine boundary effects and bulk *ab initio* calculations using a Matthiessen rule [189, 190]. In the limit where porosity approaches 100%, μ declines to zero. However, we observe that even at high porosities above 80%, MD simulation geometries with dramatically different values of κ collapse to the same value of μ , indicating that a hydrodynamic-like description of the phonon flow can be extended to extremely constrained systems. These results suggest that the phonon distribution evolves collectively under nanoscale confinement, and should be characterized by averaged phonon properties including ℓ , κ_{GK} and μ , rather than multiscale description where each phonon mode interacts independently with the geometry. This perspective is consistent with recent electron transport experiments, where current profiles exhibit hydrodynamic-like vortices at low temperatures [191]. Since momentum-conserving electron-electron collisions are scarce under these conditions, the observed behavior is instead induced by nanoscale confinement, in analogy to the emergence of phonon hydrodynamics in silicon nanosystems.

Because the framework of permeability and viscosity is derived within the phonon gas model,

it cannot account for coherent effects on the phonon dispersion, such as changes in group velocity due to elastic softening, which originate from the wavelike nature of phonons. The breakdown of the current framework in the presence of coherent effects would likely manifest as the emergence of a more complex relationship between viscosity and various geometrical parameters, in particular a dependence on the degree of system disorder [115]. In the experimental studies considered here, we observe that viscosity is independent of system disorder, *e.g.*, the periodically arranged 3D metalattices display the same viscosity as the randomly distributed pores in the nanowires. This supports the assumption that phonons undergo diffuse boundary scattering in these systems, which randomizes phase and destroys coherence, in contrast to coherent effects previously reported in periodically structured silicon systems at very low temperatures [114, 115]. In contrast, the simulations examined here all consider periodic pores. It would be interesting to check whether a molecular dynamics simulation of a disordered metalattice exhibits a viscosity which deviates from the $(1 - \phi)/\phi$ scaling.

To gain local information about heat transport in highly-confined systems, we consider a 3D metalattice geometry with square pores in a simple cubic arrangement, and use molecular dynamics (MD) simulations to calculate the steady-state the heat flux profile shown in Fig. 4.14(a) (details in Ref. [10]). The cross-sectional heat flux profiles in the channels between pores are Poiseuille-like, and are incompatible not only with the exponential profiles predicted by the Fuchs-Sondheimer approximation to the Boltzmann transport equation [107], but also with more general ballistic descriptions using the bulk MFP spectrum [86]. Instead, the flux profiles are reminiscent of those predicted by the steady-state GKE,

$$\mathbf{q} = -\kappa_{\text{GK}} \nabla T + \ell^2 \nabla^2 \mathbf{q} \quad (4.39)$$

which is fit to the MD results as shown in Fig. 4.14(b), with ℓ , κ_{GK} and ℓ_{slip} as the fit parameters. Remarkably, the viscosity value calculated from the fit parameters as $\mu = \ell^2/\kappa_{\text{GK}}$ follows the same scaling with porosity as is obtained previously using κ and K , as shown in the inset




Figure 4.14: Atomistic and phonon hydrodynamic heat flux profiles in highly-confined geometries: Heat flux profiles in a square silicon metalattice of 6.5 nm pore diameter subject to an external temperature gradient, (a) calculated from molecular dynamics (MD) simulations and (b) fitted to the MD results using the Guyer-Krumhansl equation (GKE) with effective non-local length (ℓ) and thermal conductivity (κ_{GK}) parameter values. (c) Cross-sectional heat flux profile in a channel between pores from the MD simulation (red), GKE (black) and Stokes approximation to the GKE (blue). The Stokes approximation is the starting point for the derivation of the Darcy law analog. The inset shows that the viscosity parameter, $\mu = \ell^2 / \kappa_{\text{GK}}$, resulting from the GKE fits to ℓ and κ_{GK} exhibits the same $(1 - \phi) / \phi$ dependence on porosity as the results in Fig. 4.13. Figure from [10].

to Fig. 4.14(c). The observed decrease in μ with increasing porosity is due to a reduction in the size of strongly correlated heat flux regions surrounding the boundaries, whose extent is described by ℓ . However, the conductivity parameter κ_{GK} is also slightly reduced. At fixed porosity, both ℓ and κ_{GK} increase when increasing the pore spacing, but the resulting viscosity μ remains the same. Finally, Fig. 4.14(c) illustrates that the Stokes approximation in Eq. 4.34 can fit the local heat flux profile in the passages between contiguous pores, as required by the Darcy's law derivation. However, the Stokes approximation is not sufficient to predict the complete heat flux profile in MD simulations at the areas of interconnection between multiple channels.

Chapter 5

Nanofabricated metallic gratings in the EUV scatterometry experiment

The experiments of the previous chapter rely on the fabrication of nanoscale metallic gratings on the sample surface, to act as absorbers for the infrared pump and a diffraction grating for the extreme ultraviolet probe. The gratings are fabricated using the electron-beam lithography (EBL) process described below. In previous EUV scatterometry experiments, collaborators performed the EBL process, enabling investigation of the nanoscale elastic and thermal behaviors of silicon and a variety of insulators, including sapphire, fused silica and dielectric thin films [58, 65, 66]. New samples of interest, including diamond, lithium lanthanum zirconium oxide (LLZO), 2D metal halide perovskites, and hexagonal boron nitride, have either higher surface roughness or more fragile surfaces, which have complicated the typical EBL procedure. To ease with troubleshooting, lead times and repeatability, we developed a standardized EBL procedure at the University of Colorado, using the shared instrumentation facilities in JILA and COSINC. This chapter details the new EBL process, which was derived from the recipes used by our collaborator, Weilun Chao, at LBNL, with input from Justin Shaw at NIST and students Terrell Curley and Alex Motler from Jeff Jessing's group at Fort Lewis College. When applied to diamond substrates, both the new and older LBNL electron-beam lithography recipes triggered the formation of an anomalous, low-density, diamond-like amorphous carbon layer between the diamond substrate and nickel structures. The collaborative effort to characterize this layer's novel material properties, using EUV scatterometry, EUV imaging reflectometry, X-ray reflectivity and several modalities of electron microscopy will appear in an upcoming publication. The challenges encountered during this foray

into nanofabrication demonstrate the value of advanced metrology techniques for increasingly complex problems in nanofabrication, and the attractiveness of fabrication-free approaches, such as the deep ultraviolet transient grating experiment introduced in Chapter 6.

5.1 Electron-beam lithography process

Electron-beam lithography uses a nanoscale focused electron beam and a nanometer-precision motorized stage to write the desired lithographic pattern in a resist layer as the sum of many point exposures. The very short electron wavelength allows the focal spot size itself to set the resolution, eliminating the need for a photomask and making EBL in principle a higher resolution technique than photolithography. However, EBL is much slower than photolithography, making it unsuitable for mass production in the semiconductor industry.

Fig. 5.1 outlines the basic process steps of the EBL process developed in JILA and COSINC. It is important to start with a clean substrate. We verified using X-ray reflectivity that sequential one-minute sonications in IPA, methanol and ethanol baths, followed by a one-minute rinse under deionized water and drying with compressed nitrogen will remove standard hydrocarbon contamination from diamond. This recipe comes from Samuel Marks from Mike Toney's group. It is also possible to clean with acetone; however, care must be taken to follow up with IPA and methanol before the acetone dries, since acetone alone will leave a thin residue.

The cleaned sample should be very dry before spin coating, which can be accomplished by leaving the sample on a hot plate set to $\sim 100^\circ\text{C}$ for ten minutes. The hot plates in the JILA clean room are most likely very dirty. If the cleanliness of the back side of the sample should be preserved, a clean, expendable silicon wafer can be placed underneath. The goal of the spin coating step is to deposit a uniform 50 nm PMMA layer on the sample surface to serve as a resist for the electron beam. Of the two spin coaters in the JILA clean room, we use the enclosed, programmable Laurell Technologies model on the left. The spin coating recipe uses an A2 950K PMMA solution from Kayaku Advanced Materials (formerly MicroChem), where A2 refers to a 2% concentration in anisole and 950K to the molecular weight. Concentration, along with spin speed affects the final




Figure 5.1: Electron-beam lithography process overview: The substrate is first cleaned in solvents, then dried on a hot plate. A spin-coated 50 nm PMMA layer serves as the resist. If the substrate is insulating, it is necessary to deposit a conductive discharge layer on top of the PMMA, typically 15 nm of Al. The electron beam exposes the PMMA with which it comes into contact, altering it chemically and recording the desired nanoscale pattern. Removal of the discharge layer and development of PMMA leaves only the unexposed PMMA. A uniform ~ 12 nm nickel layer is evaporated on and between the PMMA structures. Finally, a liftoff step removes the remaining, unexposed PMMA and the nickel on top of it, leaving the nickel pattern in the areas where the PMMA was exposed.

film thickness, as shown in the spin curves on the Kayaku website. Each spin coater can have a slightly different spin curve. Program 13 in the JILA spin coater menu contains our recipe, which spins at 6000 rpm (2000 rpm/s acceleration) for 45 seconds. The recipe also includes a 5 second, 500 rpm (250 rpm/s acceleration) step at the beginning. When depositing the PMMA solution on the sample before spinning, it is very important to keep the pipette close to the sample surface, cover the majority of the sample with solution, and start the program immediately thereafter. If the anisole solution splashes or has time to evaporate, it will result in an uneven PMMA layer. However, if too much solution is applied, especially at low spin speeds and for small samples, it can coat the underside of the sample, making it difficult to remove from the chuck. If the sample is very valuable, it can be worthwhile to test that the vacuum chuck is working with a less precious sample, since if the chuck is dirty it can fail to properly seal, causing the sample to fly off when the spinning starts. A dirty chuck and O-ring can be cleaned with standard solvents.

Immediately after the sample stops rotating, it should be transferred to a hot plate set to 180C and baked for 90 seconds. Keeping track of the bake time is important. A 30 second deviation in bake time will noticeably modify the PMMA's optical properties in the JILA ellipsometer. It is also useful to visually inspect the PMMA layer. Radial streak marks can mean that the substrate was not sufficiently clean, while discolorations can indicate thickness variations. When spin coating a small or non-circular substrate, it is important to remember that the PMMA thickness will vary within approximately one millimeter of the edges, especially near corners. The PMMA thickness can be verified with the ellipsometer.

Unlike photolithography, electron-beam lithography requires a conductive substrate. Otherwise charge will accumulate along the pattern, deflecting the electron beam and worsening the pattern resolution. Doped silicon counts as a conductive substrate, however, intrinsic silicon and insulators such as sapphire and diamond require the deposition of an additional discharge layer before the writing step. This discharge layer can be a thin metal film, or one of a variety of conductive polymers. Conductive polymers are spin coated on the sample after the PMMA and have the advantage of being water-soluble and therefore easy to remove after the writing step. Our tests of two such

products, aquaSAVE from Mitsubishi Chemical Corporation and DisCharge H₂O from DisChem indicate that both are clearly inferior to a 15 nm aluminum layer at preventing charging effects during EBL. An aluminum layer can be deposited on top of the PMMA using the CVC thermal evaporator, currently located in the COSINC-FAB clean room on East Campus. Instructions for the thermal evaporator are straightforward, provided at the machine and covered in a mandatory training process. It is important to deposit the aluminum relatively quickly, so that the PMMA does not exceed its 160°C melting point. A 140 Ampere current should be sufficient to evaporate aluminum using the CVC, or slightly higher when melting a new pellet for the first time. The current can be safely ramped at one Amp per second on the way up, and faster on the way down. Use of alumina-coated boats is advisable to prevent reactions with the tungsten and molybdenum in the boat itself.

The EBL step itself is the same regardless of whether a discharge layer is present. The electron beam source in our process is the Hitachi SU3500 SEM in COSINC-CHR on East Campus. Its maximum 30 keV energy, while typical for an SEM, is sub-optimal for an EBL process, since lower energy beams tend to generate more secondary electrons when passing through the PMMA layer. These secondary electrons travel laterally in the PMMA film, broadening the area of exposure and reducing pattern resolution. Dedicated EBL systems, such as the JEOL JBX 6300-FS at NIST, use 100 keV beams and proximity correction algorithms to reduce proximity effects. The lithography sample should be mounted in the SEM along with a reference sample, used for alignment, and a Faraday cup, used to calibrate the electron beam current. When mounting the sample, use copper tape on the corners in addition to any insulating tape on the backside. This provides an electrical connection to prevent charge buildup. It can also be helpful to make a ~1 mm scratch in the PMMA using a scribe, to provide a reference for alignment on the otherwise featureless sample surface.

The electron beam writing process on the Hitachi SEM relies on the Nanometer Pattern Generation System (NPGS) developed by J.C. Nabity, a product which includes the necessary hardware and software to enable the SEM to rapidly and precisely alternate between overlapping point-like exposures and stage movements to draw the desired lithographic pattern. The details




Figure 5.2: Characterization of nanoscale nickel gratings fabricated with EBL on silicon substrates: (a) Electron beam contamination spot generated on PMMA using the method described in the text, with good (left) and poor (right) astigmatism alignment. The diameter of the contamination spot depends on charging effects and exposure time and does not necessarily correspond to the real beam size. (b) SEM image of the COSINC Hitachi Au standard at 100,000x magnification after alignment. The fuzziness in the image arises from contamination on the Au standard itself and is not indicative of poor alignment. (c) – (f) AFM (left) and SEM (right) images of nickel gratings of varying periodicities on silicon substrates. The gratings were fabricated in COSINC and JILA using the recipe in Appendix D. Periodicities $<100\text{ nm}$ are likely beyond the capabilities of the converted Hitachi SU3500 SEM, and would require a dedicated EBL system.

on NPGS, including the file preparation, SEM alignment and diagnostics appear in detail in the Quick Tutorial section of the NPGS manual, available online or in print next to the Hitachi SEM. The NPGS file preparation can be a time-consuming exercise in CAD. However, there are test files and common grating patterns located in the “KM” folder on the Hitachi SEM computer, which can be used to troubleshoot. The most challenging aspect of EBL on a converted SEM is the alignment, which includes focusing, aperture alignment and astigmatism steps and can take some practice. Fig. 5.2(a) shows two indicators of SEM alignment quality, produced by allowing the electron beam to rest at a single location on the PMMA for ~ 10 seconds. This is a very long time for the electron beam to be in any one place, and as a result it forms a contamination spot consisting of hydrocarbon molecules “cracked” by the high electron energies and deposited on the sample surface. By subsequently imaging the contamination spot, it is possible to determine the beam mode shape. The image on the left is typical, while the image on the right exhibits obvious astigmatism and should be realigned. The size of the contamination spot depends on the exposure time and therefore does not indicate the precise electron beam diameter. Fig. 5.2(b) shows a typical image of the reference sample, or “gold standard” after SEM alignment. This sample consists of gold nanoparticles embedded in a graphite substrate. The nanoparticle edges are blurry in Fig. 5.2(b) because the Hitachi SEM gold standard is covered in carbon contamination. Because it is possible to see the small nanoparticles, of <50 nm diameter, the image indicates the SEM is well-aligned.

Several points are worth mentioning that caused us trouble when first developing the EBL process on the Hitachi SEM. First, positive and negative stage movements in the NPGS and SEM software are defined differently. Verify that the NPGS software will move the stage in the correct direction by executing a movement-only run file near a reference feature and checking the microscope before and after. When switching back and forth between the NPGS and SEM software, make sure to consider whether or not a particular action will unblock the electron beam, potentially exposing unwanted areas of the PMMA and destroying the lithographic pattern. To avoid this, always set up the picoammeter early in the process, since it will only register a current if the beam is unblocked. After completing the lithography, make a large manual stage movement away from the patterned

area before beginning the shutdown procedure. When setting the lithography parameters, make sure that the center-to-center and line spacing parameters in the NPGS run files are equal and set to values <20 nm. This distance will ensure that the point exposures overlap. Finally, while the NPGS manual suggests that writing with a smaller current will improve pattern resolution at the cost of slowing down the lithography, we have observed that the Hitachi's resolution does not improve below 50 pA.

After electron beam writing, the aluminum discharge layer, if present, can be removed using a 2.5% NaOH solution prepared by Curtis Beimborn in 2021 and stored in the X110 corrosives cabinet. The shelf life of this solution should be years if stored properly. It is only necessary to submerge the sample for approximately 30 seconds in this solution. At 2.5% concentration, NaOH will still burn exposed skin, but is far less dangerous than higher-concentration mixtures. Afterwards, the sample should be immediately rinsed in deionized (DI) water and blown dry.

To develop the PMMA, we vertically submerge the sample for 90 seconds in a 3:1 IPA:H₂O mixture while gently swirling it back and forth. Reverse action tweezers are very helpful in this step, to prevent the sample from being accidentally dropped into the bath. After 90 seconds, the sample should be immediately rinsed in DI water and blown dry, to terminate the process. While many development recipes use a mixture of methyl isobutyl ketone (MIBK) and IPA, we observed no evidence that this extra chemical improves the process [192]. EBL pattern resolution has been shown to improve when development occurs at low temperatures [193]. In high-resolution fabrications, we therefore situate the glassware containing the development chemicals inside a larger bath of ice water. The ice can be obtained from the JILA ice machine and transported to clean room in a secondary container. This process will cool the development bath to $\sim 5^{\circ}\text{C}$ in a matter of minutes.

When developing an EBL process for the first time, it is extremely important to inspect the sample with atomic force microscopy (AFM) after the development. The AFM will indicate the quality of the unexposed PMMA structures shown in the bottom left panel of Fig. 5.1, which will in turn affect the outcome of the liftoff process described below. If the PMMA structures are rounded, due to limitations in the SEM resolution, a poor SEM alignment, or charging effects, the

liftoff process may go poorly, since metal will cover the positively-sloped PMMA sidewalls during deposition and shield the PMMA from the liftoff chemicals. If the PMMA structure height deviates from nominal, it may be necessary to adjust the spin coating recipe in future iterations. For more established recipes, or lower resolution fabrications, it is usually sufficient to inspect the sample with an optical microscope, where features of $>1\text{ }\mu\text{m}$ size should be easily visible.

The desired metal layer, in our case nickel, can now be deposited on top of the PMMA using the CVC thermal evaporator mentioned above. It is especially important to deposit the nickel relatively quickly so that the PMMA does not exceed its 160°C melting point, which would destroy the lithographic pattern. A 220 Ampere current should be sufficient to evaporate nickel using the CVC, or slightly higher when melting a new nickel pellet for the first time. The current can be safely ramped at one Amp per second on the way up, and faster on the way down. Use alumina-coated boats for nickel to prevent it reacting with the tungsten and molybdenum in the boat itself. The target metal layer thickness should be no more than one-quarter of the PMMA thickness, in this case $\sim 12\text{ nm}$ for a 50 nm PMMA film. The Edwards electron-beam evaporator in the JILA clean room can also deposit the nickel layer. However, we observe that after passing through the Edwards, the PMMA layer is no longer soluble in liftoff chemicals such as acetone or N-methylpyrrolidone (NMP). This likely occurs because the electron beam, or byproducts thereof, deliver such a large exposure dose to the PMMA as to crosslink it and make it behave as a negative resist [194].

A process known as liftoff removes the remaining PMMA and the metal lying on top of it, leaving metal on the sample surface only in the areas of PMMA exposed to the electron beam. Liftoff processes involve a solvent for the unexposed PMMA, heat, sonication and repetition. A template recipe from which we often deviate appears in Appendix D. Higher-resolution lithographic patterns may require longer liftoff processes. The Remover PG chemical from Kayaku Advanced Materials serves as our solvent. It contains the active ingredient NMP, which is a reproductive toxin and should be used only in small quantities and only when working in the clean room fume hood with gown and gloves. Acetone can also serve as a liftoff solvent. However, it should not be heated beyond room temperature due to its low flash point. We abandoned acetone after observing

it to leave tall spikes, or “bunny ears” on the sides of nickel structures after liftoff, which NMP will remove. After liftoff, the sample should be immediately rinsed in DI water and blown dry.

The bottom four panels in Fig. 5.2 show AFM and SEM images of nickel gratings of varying linewidth and periodicity on silicon substrates fabricated using the EBL process described above. As shown in Fig. 5.2(f), a 50 nm grating pitch appears to be the resolution limit for the Hitachi SEM, and for such small features it is desirable to use a 100 keV EBL system.

5.2 Amorphous carbon at metal-diamond interfaces

The exceptional material properties of diamond make it an attractive candidate for a variety of power electronics, radiofrequency, quantum sensing and quantum communication applications and an exciting material to study at the nanoscale using the dynamic EUV scatterometry experiment [195, 196, 197, 198]. Diamond has the highest thermal conductivity amongst naturally occurring materials, with a long but relatively narrowband thermal phonon mean free path spectrum, suggesting that it will exhibit nanoscale thermal effects at relatively large length scales [199]. However, diamond is difficult to study with visible ultrafast lasers due to its 5.5 eV bandgap, which necessitates either the use of a deep ultraviolet pump, as discussed in Chapter 6, or the fabrication of metallic transducers, discussed here.

The fabrication of nanoscale nickel gratings on diamond proceeds broadly along the same lines as the recipe described above, with special attention paid to the aluminum discharge layer, necessary to prevent charge buildup on the insulating substrate from adversely affecting the pattern resolution. Two different recipes, detailed in Appendix D, succeeded in fabricating nickel gratings with periods down to 200 nm on diamond substrates. The unanticipated result of both recipes is an amorphous carbon layer on the diamond surface beneath and the between the nickel structures, as shown in Fig. 5.3. It is unclear which fabrication step produces the amorphous layer. The layer is revealed by electron microscopy. High-angle annular dark-field scanning transmission electron microscopy (HAADF-STEM) in Fig. 5.3(a) shows the contrast of the dark, lower-density amorphous layer with the crystalline diamond underneath, while energy-dispersive X-ray spectroscopy (EDS)




Figure 5.3: Amorphous carbon layer at metal-diamond interfaces after electron-beam lithography: (a) HAADF-STEM image of the amorphous carbon layer between a nickel structure and the diamond substrate, illustrating its lack of crystallinity. (b) Schematic of the 1D nickel grating of linewidth (L), period (P) and height (h), with labelled approximate amorphous carbon dimensions. (c) EDS image showing elemental composition of a single nickel nanostructure and the area underneath. The amorphous carbon appears as a dark red layer of lower density, and is thinner beneath the nickel structures. The platinum and aluminum layers are deposited for the electron microscopy.

in Fig. 5.3(c) reveals the elemental composition. The lack of oxygen in the EDS establishes that the amorphous layer is not PMMA or generic hydrocarbon contamination. Surprisingly, the amorphous layer is ~ 5 nm underneath and ~ 20 nm between the nickel structures. Its presence complicates the study of the sample using the EUV scatterometry experiment, since its material and optical properties are uncertain, in particular its absorption of the infrared pump beam. A collaborative characterization of this amorphous layer using both EUV and electron microscopy techniques will appear in an upcoming publication.

Chapter 6

Tabletop femtosecond deep ultraviolet transient grating experiment

The present chapter is dedicated to a new beamline constructed at JILA over the past three years, which uses fast, deep ultraviolet (DUV) light at 196 nm to overcome the fabrication challenges discussed in the previous chapter. The spatial and temporal interference of two DUV pulses on a sample creates a sinusoidal transient grating (TG) excitation. A time-delayed probe pulse of adjustable wavelength diffracts from the TG, monitoring the propagation of excited carriers, thermal phonons and acoustic waves between interference fringes. The shorter DUV wavelength enables fringe spacing down to ~ 230 nm using standard two-inch optics, half the previous tabletop record, while its 6.3 eV photon energy permits the investigation of wide-bandgap samples inaccessible to non-contact visible techniques. The following sections provide a brief overview of transient grating techniques and their recent relevant applications, a detailed overview of the new DUV-TG beamline, and a discussion of ongoing investigations of nanoscale electron and phonon transport in diamond and silicon.

6.1 Overview of the transient grating technique

Transient grating experiments utilize the sinusoidal interference pattern created by crossing two coherent light sources, as shown in Fig. 6.1(a). For two plane waves of wavelength λ with wavevectors $\mathbf{k}_{1,2} = k_y \hat{\mathbf{y}} \pm k_x \hat{\mathbf{x}}$, intersecting at an angle θ such that $k_x = |\mathbf{k}| \sin(\theta/2)$, the total electric field amplitude along the x direction parallel to the plane of intersection is

$$E(x) = E_0 e^{ik_x x} + E_0 e^{-ik_x x} = 2E_0 \cos(k_x x) \quad (6.1)$$

The resulting intensity profile is a sinusoidal grating of wavevector $\mathbf{q} = \mathbf{k}_2 - \mathbf{k}_1 = 2k_x \hat{\mathbf{x}}$, and period

$$\Lambda = \frac{2\pi}{|\mathbf{q}|} = \frac{\lambda}{2 \sin(\theta/2)}, \quad (6.2)$$

which can be reduced by decreasing the wavelength or increasing the crossing angle [200]. The same basic relationships hold for ultrafast pulsed lasers, making it possible to illuminate a material of interest with a femtosecond, microscopic, contactless grating and monitor its material response.

For visible and ultraviolet excitation wavelengths, the incident photons typically trigger a fast, electronic response, as electrons absorb the photons and change energy levels, modifying the material's refractive index. Electron-electron scattering occurs within tens of femtoseconds, equilibrating the electron population at an elevated temperature, which, due to the low heat capacity of the electron system, can be several thousand Kelvin. The heated electrons immediately begin to propagate in the material. At the same time, they begin to scatter with the ions, generating lattice vibrations and injecting energy into the phonon system. While the excited electrons can also relax through various radiative recombination processes, these generally occur on longer timescales or require high carrier densities, making them secondary to carrier diffusion and electron-phonon scattering as energy dissipation pathways.

As phonons are excited, the lattice expands, inducing a further, thermal change in refractive index and a sinusoidal sample surface deformation. The thermal grating decays as phonons propagate and scatter outwards from the excited region, gradually bringing the material into thermal equilibrium on the nanosecond to microsecond timescale, before the next excitation pulse arrives. The rapid lattice expansion also impulsively launches lower energy, longer wavelength, acoustic vibrations, including a resonant, coherent surface acoustic wave (SAW) whose wavelength is set by the transient grating period.




Figure 6.1: Transient grating technique overview: (a) The interference of two monochromatic plane waves produces a sinusoidal interference pattern, with fringe spacing determined by the angle between the two interfering beams and their wavelength. (b) In a transient grating experiment involving ultrafast lasers, the two pump beams must be spatially and temporally overlapped. A time-delayed probe beam diffracts from the pump excitation and a photodetector collects only the diffracted portion of the probe beam, which is absent before the pump arrives and after the transient grating relaxes. (c) The number of diffracted probe beams and their diffraction efficiencies depend on the probe incidence angle and whether the transient grating is thin (left) or thick (right) relative to the probe wavelength. Figure adapted from [200].

A key challenge when probing a transient grating is to isolate the behavior of interest, whether electronic, thermal or acoustic, from the variety of often simultaneous ongoing material responses. The typical approach involves diffracting an additional probe laser, often of a different wavelength, from the transient grating. The probe laser diffraction can occur in reflection, or in transmission for a sample sufficiently transparent to the probe wavelength λ_P . For transmission measurements, as shown in Fig. 6.1(c), the diffraction will be more or less sensitive to the probe incidence angle depending on whether the excitation depth is thick (right) or thin (left), as discussed further in Section 6.2.2. Approximating the transient grating as a thin sinusoidal variation of wavevector magnitude q , thickness d , and time-dependent real refractive index $\Delta n(t)$, absorption coefficient $\Delta k(t)$ and sample surface height $\Delta h(t)$, and using scalar diffraction theory [162], the amplitude transmittance function $t_A(\xi, t)$ imparted on a normally-incident probe as it passes through the sample appears quite complicated,

$$t_A(\xi, t) = \left[1 + (e^{-2\pi\Delta k(t)d/\lambda_P} - 1) \cos(q\xi) \right] \exp \left[i \frac{2\pi\Delta n(t)d}{\lambda_P} \cos(q\xi) \right] \exp \left[i \frac{4\pi n\Delta h(t)}{\lambda_P} \cos(q\xi) \right], \quad (6.3)$$

where ξ is the spatial coordinate along the sample plane. Assuming the terms involving $\Delta n(t)$, $\Delta k(t)$ and $\Delta h(t)$ are small and expanding them to linear order, Eq. 6.3 can be simplified as

$$t_A(\xi) \approx 1 - \frac{2\pi}{\lambda_P} \cos(q\xi) [\Delta k(t)d - i(\Delta n(t)d + 2\Delta h(t))]. \quad (6.4)$$

The amplitude, $E_p(t)$ and intensity, $I_p(t)$, of the first diffracted order of a plane wave incident on this transmittance function can be obtained in the Fraunhofer approximation by taking the Fourier transform of Eq. 6.4 with respect to ξ ,

$$E_p(t) \approx \Delta k(t)d - i(\Delta n(t)d + 2\Delta h(t)), \quad (6.5)$$

$$I_p(t) = |E_p(t)|^2 \approx [\Delta k(t)d]^2 + [\Delta n(t)d]^2 + 4[\Delta h(t)]^2 + 2\Delta h(t)\Delta n(t)d. \quad (6.6)$$

The many terms in Eq. 6.6 make it difficult to separate the thermal and electronic contributions to the experimental signal. In order to simplify the experimental signal, the diffracted probe is often interfered with a static reference beam of constant amplitude E_s and relative phase ϕ , resulting in a signal of the form

$$I_p^{\text{het}}(t) = |E_p(t) + E_s e^{i\phi}|^2 = |E_s|^2 + |E_p(t)|^2 + 2E_s E_p(t) \cos(\phi) \quad (6.7)$$

This approach is known as heterodyning [151, 201]. If the reference beam is significantly more powerful than the diffracted probe, $|E_p(t)|^2$ can be neglected, resulting in a linear time-dependent diffracted signal of the form

$$I_p^{\text{het}}(t) \approx 2E_s \{\cos(\phi)\Delta k(t)d - \sin(\phi)[\Delta n(t)d + 2\Delta h(t)]\}, \quad (6.8)$$

where the relative phase ϕ can be adjusted to more effectively isolate the various signal components [202]. It is also possible for scattered light with a relatively stable phase to interfere with the diffracted probe, in a process known as parasitic heterodyning and discussed further in Section 6.2.5. In the case of parasitic heterodyning, the reference and diffracted light can be of similar intensity, complicating interpretation of the probe signal.

The diffracted probe beam is measured on a photodetector, as illustrated in Fig. 6.1(b). This can pose a practical challenge, since the changes in refractive index and sample height associated with a transient grating are often quite small, making the diffraction efficiency quite weak. The desired time resolution also influences the choice of probe. A continuous probe laser is adequate to observe microsecond or nanosecond dynamics, since its temporal resolution is limited by the photodetector

bandwidth, while a pulsed probe laser is necessary to resolve sub-nanosecond dynamics.

Fig. 6.2 shows results from several recent transient grating experiments to illustrate some of the concepts discussed above. The first panel shows transient grating measurements of thermal diffusion in 400 nm silicon membranes. The pump wavelength is 515 nm, which has a sufficient absorption length that the sinusoidal excitation is approximately uniform across the membrane depth, resulting in a one-dimensional transport problem. The 532 nm probe laser is continuous, and so the experiment has approximately nanosecond time resolution, sufficient for the micron-scale TG periods measured. The inset in Fig. 6.2(a) shows a negative spike at early times. This is carrier diffusion, which happens much faster than thermal diffusion in silicon. The sign of the electronic decay is flipped because the probe is heterodyned and therefore its intensity is linear with the change in refractive index, which changes in opposite directions for the electronic and thermal excitations. For a non-heterodyned experiment, the signal would be squared, as in Eq. 6.6, and entirely positive.

Fig. 6.2(b) plots the thermal decay timescale vs. the square of the transient grating wavevector. The heat equation predicts a linear relationship between these two variables, which breaks down at small TG periods. This is an example of non-diffusive thermal transport, where heat conduction becomes less efficient than macroscale predictions when confined, in this case by the small TG period, to dimensions on the order of the mean free paths of the heat carrying phonons. Since the average mean free path of heat carrying phonons in silicon is ~ 300 nm [80] and the smallest period measured in Fig. 6.2(b) is $2.4 \mu\text{m}$, this is merely the onset of non-diffusive transport. Even more interesting behaviors may occur at the smaller length scales accessible to the deep ultraviolet transient grating experiment introduced in the following sections.

Fig. 6.2(c) shows another interesting thermal behavior observed using a transient grating experiment, second sound in graphite at 200 K. Second sound refers to the propagation of heat as a wave and is discussed extensively in the literature [74, 109, 112, 113]. In the present experiment, the temperature wave causes the local refractive index to oscillate. Because of the heterodyning, this results in a negative signal at a delay time sufficient for the original temperature maxima to become minima, and vice versa.




Figure 6.2: Previous transient grating measurements of phonon thermal transport: (a) Transient grating measurements of a 400 nm silicon membrane at varying periods using a 515 nm pulsed pump and a 532 nm continuous probe. The main figure shows the phonon-mediated thermal relaxation, while the negative spike in the inset shows the much faster carrier diffusion, which is flipped in sign due to the heterodyning. (b) Thermal relaxation timescales for the data in (a), which deviates from bulk diffusive predictions at small TG periods. (c) Transient grating observation of second sound at 150 K in graphite at micron length scales. (d) EUV transient grating on a Si_3N_4 membrane implemented at a free electron laser, and measured with a visible (blue) and EUV (black) probe beam. While the visible probe is very sensitive to electron dynamics, the off-resonance EUV probe measures only surface deformation induced by lattice heating. Figure adapted from Refs. [56, 112, 203].

Finally, Fig. 6.2(d) illustrates the sensitivity of the transient grating signal to the probe wavelength. The demand for short transient grating periods has been partially met by the coherent, extreme ultraviolet (EUV) light available at X-ray free electron lasers, whose femtosecond pulse duration and short wavelength enables excitation periods down to tens of nanometers [56, 204]. EUV light has the added benefit of being insensitive to most valence electron dynamics due to its high photon energy [55], while its short wavelength makes it especially sensitive to the phase differences caused by sample surface deformation. The blue curve in Fig. 6.2(d) shows an EUV transient grating on a Si_3N_4 membrane probed by a 400 nm pulsed laser, which mainly measures the carrier diffusion. The black curve shows the dynamics of a similar transient grating probed with EUV. The electron dynamics are nearly absent, but the signal from the picosecond-scale sample surface deformation is enhanced.

While visible transient grating experiments excel at producing micron-scale excitations, they can only measure metals and low-bandgap semiconductors. EUV transient gratings can excite almost any material, at much shorter length scales, but thus far require a facility-scale free electron laser to generate sufficient pump fluence. This complicates the implementation of heterodyning, and due to the vacuum and optics requirements of EUV, makes it difficult to reconfigure the experiment to measure a variety of TG periods. EUV light also easily damages soft matter, limiting its ability to reliability measure resists and polymers. A transient grating experiment pumped by deep ultraviolet (DUV) light is therefore attractive, since it would combine the ability to investigate wide-bandgap materials at short length scales with the throughput and flexibility of a tabletop experiment. Chapter 2 discusses the advantages and disadvantages of a few approaches to producing femtosecond, coherent, DUV light, including the Ti:Sapphire upconversion used in the JILA experiment described below.

6.2 DUV transient grating beamline

The present section breaks down the DUV transient grating experiment into its pump and probe components, and includes a subsection describing the various electronic devices and software necessary to adjust the optical delay stages and acquire data. The following treatment is meant to

be relatively exhaustive, both in terms of explaining how the experiment works, and documenting its day-to-day operational procedure. Additional detail can be found in an upcoming publication and in the alignment and troubleshooting documentation referenced in Appendix B.

6.2.1 Pump beamline

The process of producing a ~ 1 mW DUV transient grating on a sample begins with kilowatts of electrical power, approximately enough to drive an electric vehicle at 60 mph down a highway. The high input power is required to drive the many stages necessary to produce light that is not only femtosecond pulsed, but also collimated, coherent and of the correct wavelength. Fortunately, the output of a Ti:Sapphire chirped pulsed amplification system already meets three of these stringent requirements. The Wyvern system in B1B35 outputs pulses of ~ 50 fs duration at 1 kHz repetition rate. The light is collimated, coherent and at 7 W average power contains 7 mJ of energy per pulse. The alignment of this laser is covered in Appendix C, while general details on chirped pulse amplification and oscillators appear in Section 2.1. The task of the pump beamline is to change the amplifier wavelength and focus it into a TG pattern on the sample. Fig. 6.3 depicts the post-amplifier DUV pump beamline. The majority of optics, shown in Fig. 6.3(b) are dedicated to the production and separation of the 200 nm DUV light. Subsystems dedicated to the transient grating and the generation of the third harmonic are shown in Fig. 6.3(a) and Fig. 6.3(c), respectively. See Section 2.1.4 for an overview of the various nonlinear optical concepts discussed below.

The two primary considerations in the experimental design are ease of day-to-day use and the maximization of the DUV pulse energy incident on the sample. Pulse duration and mode quality are secondary considerations. The experimental signal increases with the DUV fluence, in mJ/cm^2 , on the sample, which depends on the DUV pulse energy and the pump spot size. Practically, the pump spot size can only be made so small, leaving DUV pulse energy as a key experimental parameter. While the high harmonic generation process discussed in Section 2.1.5 is excellent at producing coherent short wavelength light, it has a low conversion efficiency. It is most efficient when the harmonics are produced using a short-wavelength input beam, which for the B1B35 Wyvern would

mean driving with the second harmonic, at 400 nm, in krypton to produce the 7th harmonic at 57 nm. However, tests from February 2020 indicated a flux of $\sim 10^{12}$ photons per second at 57 nm at the fiber exit, which translates to a pulse energy of only ~ 3.5 nJ, approximately two orders of magnitude lower than the DUV pulse energies used in the present experiment. Because pump pulse energies < 10 nJ are likely too low to yield detectable transient grating signal, the current pump beamline instead relies on three solid beta barium borate (β -BaB₂O₄), or BBO, nonlinear crystals. While this approach is more cumbersome, in that it involves multiple stages and additional optics, it ultimately produces ~ 20 μ J of DUV light at 200 nm wavelength. This combination of wavelength and pulse energy enables the excitation and detection of transient gratings in most wide-bandgap materials.

The beamline begins with a series of mirrors directly after the amplifier exit (not pictured) used to navigate the beam to the correct area on the optical table. There are magnetic mirrors in this area which can be easily inserted and removed to direct the beam towards the various experiments in B1B35. It is important to check, especially after aligning the amplifier, that the beam is still passing correctly through this “switchyard”. Clipping can generally be fixed by adjusting the first magnetic mirror after the amplifier. At this point, the beam should be incident on the beamsplitter optic pictured in Fig. 6.3(b). Around 90% of the horizontally polarized incident light reflects from the beamsplitter to supply the pump beamline. Alignment through the first and second irises (labelled I1 and I2) in Fig. 6.3(b), using the beamsplitter and first mirror (labelled M1) directs the pump beam onto the first nonlinear optics. There is a plano-convex lens of 2 meter focal length which should be removed for the alignment, then reinserted and realigned to the beam path. This lens focuses the beam throughout the nonlinear upconversion process, such that the focal plane occurs shortly after the final BBO.

The first two BBO crystals are packaged into a series of optics shown in Fig. 6.3(c) and sold commercially by Eksma Optics as a Femtokit 800-50-10. They are used to produce the third harmonic of the Ti:Sapphire amplifier at 267 nm. The first BBO receives only the fundamental Ti:Sapphire wavelength, and converts a portion of it to the second harmonic. The second harmonic pulse exits




Figure 6.3: Deep ultraviolet transient grating experiment pump beamline: (a) Transient grating optics. The DUV beam passes through a phase mask of 2 μm pitch, which is optimized for 196 nm wavelength to direct power into the +1 and -1 diffracted orders. The +1 and -1 diffracted orders are focused onto the sample surface using a 4f lens system. The two-inch lenses can be exchanged to alter the system's magnification, thereby changing the TG period on the sample surface. (b) Starting with 800 nm light from a Ti:Sapphire amplifier, a series of BBO crystals and separator optics (S1-S7) produces the fourth harmonic at 200 nm. Steering mirrors (M1-M2) and irises (I1-I2) facilitate the alignment (details in text). (c) Detailed view of optics used to produce the third harmonic (267 nm) and polarizations of the various harmonics during the process.

the BBO slightly after the fundamental pulse due to the higher refractive index it experiences in the crystal. It is also polarized orthogonally relative to the fundamental. To maximize the efficiency of the third harmonic generation process, it is necessary to overlap the two pulses temporally and make their polarizations equivalent. Taking advantage of the orthogonal polarizations, an adjustable birefringent calcite plate first aligns the pulses temporally. A dichroic wave plate, constructed to act as a half-wave plate for the second harmonic and a full-wave plate for the fundamental, subsequently rotates the polarization of only the second harmonic. Finally, a second BBO outputs the third harmonic via sum frequency generation. This BBO has a different cut angle, due to the different phase matching requirements of the desired nonlinear process. The range of possible cut angles restricts which nonlinear processes a BBO can support; for example, it is impossible to phase match the conversion of two second harmonic (400 nm) photons into the fourth harmonic, which prevents the use of a second frequency doubling step to produce the DUV. While the abovementioned optics all nominally have anti-reflective coatings, they still have significant back reflections which are blocked using foil around the irises.

Returning to Fig. 6.3(b), the fundamental amplifier output now propagates collinearly with its first two harmonics. Ideally, we could just fix temporal overlap, rotate polarizations and use another BBO to create the fourth harmonic from the first and third. Unfortunately, the 267 nm would likely burn the calcite, and it is difficult to create a suitable dichroic waveplate. Instead, following the approaches of Refs. [205, 206], we separate the third harmonic from the fundamental, introduce an adjustable delay and recombine them in the final BBO. This process begins with two 45-degree separating mirrors (S1 and S2), whose specialized dielectric coatings reflect the third and second harmonics, respectively, while transmitting the remaining wavelengths. The second harmonic travels to a beam block, as it is no longer needed. The third harmonic reflects from two turning mirrors and off an additional separator (S3) and into the final BBO. The fundamental beam passes through the separators, through a half-wave plate, which rotates its polarization to match the third harmonic, and off two 45 degree mirrors on an adjustable micrometer stage. The micrometer screw on this stage can be adjusted to temporally overlap the fundamental and third harmonic

pulses on the final BBO. Because the pulse durations are on the order of 100 femtoseconds, the spatial region of overlap is approximately 30 microns, necessitating very precise movements. Finally, the fundamental beam reflects from two additional steering mirrors before passing through the back side of the separator (S3). It is important to remember that significant power remains in the fundamental beam at this point. We generally begin with 3-4 mJ at the first BBO, of which roughly 1 mJ is converted to the second harmonic and 200 μ J to the third, leaving roughly 2 mJ of 800 nm light. This power is more than sufficient to create a plasma if allowed to focus.

After sum frequency generation in the final BBO, the third and fourth harmonics propagate approximately colinearly with the fundamental. It is important to separate them quickly, because the beam is approaching its focus. Because the fourth harmonic is horizontally polarized at this point, we use two periscopes, each with two separators (S4-S7), to achieve a geometry where the fourth harmonic is s-polarized on the separators, maximizing reflectivity, while maximizing transmissivity for the remaining p-polarized wavelengths. Directly after the first separator S4, a razor blade beam block catches most of the very powerful first and third harmonics before they can focus. After the fourth separator, there is generally ~ 20 μ J of the fourth harmonic and < 1 μ J each of the other wavelengths. However, because CMOS cameras are so much more sensitive in the infrared, care must be taken when trying to image the DUV mode.

High fluence, approaching the damage threshold, is important when maximizing nonlinear conversion efficiency in a BBO. This often requires a small laser spot size on the crystal. A disadvantage of the present setup is that, when using a single lens, it is impossible to place all of the BBOs in their optimal locations due to the space required for the intermediate optics. As a result, the beam is larger than optimal on the first two BBOs. This costs not only conversion efficiency, but also worsens the quality of the DUV mode, since it is derived from lower-quality out-of-focus mode profiles, in contrast to high harmonic generation, where the fiber entrance is placed at the focal plane. The DUV mode can be improved by placing a pinhole at its focus. This is because coherent light incident on a lens is Fourier transformed at the lens's back focal plane, such that the higher spatial frequency content is located furthest from the optical axis. A pinhole at the focal plane

spatially filters this high frequency content, resulting in a more Gaussian mode, without many of the fine defects amplified during the nonlinear processes. The high fluence of the DUV beam at its focus will burn a standard optical pinhole. Instead, we use a wire die from Fort Wayne Wire Die, placed in a translatable optical mount. The wire die should be replaced every few months, since the DUV light will gradually ablate material, enlarging the aperture. Fig. 6.6(a) shows the fourth harmonic, third harmonic and fundamental modes near the focus without a pinhole, while Fig. 6.6(b) shows the improvement in the fourth harmonic mode after insertion of the wire die. Conveniently, because the fundamental and third harmonic beams are not perfectly colinear as they pass through the last BBO, and because momentum conservation dictates that the fourth harmonic beam therefore appear between them, the wire die also acts as the final wavelength separator. Because smaller apertures offer improved spatial filtering, but also block more of the DUV light, the aperture size determines the tradeoff between DUV power and mode quality. We find that a 100 μm aperture diameter is optimal, reducing DUV power by approximately two-thirds while greatly improving the mode.

The DUV next passes through an optical chopper, a spinning wheel of alternating closed and open apertures. If set to rotate at the correct frequency, the chopper is capable of blocking every other DUV pulse, reducing the repetition rate from 1 kHz to 500 Hz. The timing system controlling the chopper is discussed in Section 6.2.3, which covers the experiment's electronics and software. It is important to locate the chopper near the focus, since a large beam could be clipped on the apertures, resulting in a train of pulses with uneven intensities. It is also important that the chopper come after the separators and pinhole, since it would otherwise chop the 800 nm light. Since most photodetectors are very sensitive to this wavelength, this can create an unwanted signal on the photodetector, which at 500 Hz is not filtered by the lock-in amplifier.

The DUV beam now reflects from a turning mirror (M2) towards the transient grating phase mask and lenses shown in Fig. 6.3(a). The phase mask (Ibsen Photonics) is a transmissive diffraction grating of 2 μm pitch designed to maximize the diffraction efficiency of 196 nm light into the +1 and -1 orders. This creates two beams which are refocused onto the sample by the lenses to create

the transient grating. Practically, the phase mask is far less than 100% efficient at diffracting the DUV light into the desired orders, suffering from three main inefficiencies. First, fabrication limitations cause a small amount of light to leak into the remaining diffraction orders. Second, despite an anti-reflective coating on its back side, some light reflects backwards from the phase mask. This light is also diffracted. Combined with the forward diffracted orders, there are many weak beams exiting the phase mask at various angles. Anodized foil placed around the phase mask blocks many of these beams, however, it is important to keep them in mind when working with the phase mask. Finally, the phase mask absorbs some of the DUV light. It is made from UV fused silica designed for 193 nm ArF excimer lasers, whose ~ 9 eV bandgap would ideally make it completely transparent to the fourth harmonic photons at 6.3 eV. In practice, even high-grade fused silica contains some impurities and defects. These defects, along with two-photon absorption, can result in significant absorption at high illumination intensities. Because the DUV light used in the present experiment has a pulse duration orders of magnitude smaller than most excimer lasers, its peak intensity is relatively high, causing the phase mask to glow upon illumination. In a future iteration of the experiment, the phase mask would ideally be made thinner than its current 2 mm thickness. The phase mask absorption makes the pinhole especially important, because an uneven beam mode with localized hot spots would result in areas of very high absorption, and even damage to the phase mask. In practice, the losses from the pinhole are largely compensated by savings at the phase mask. The pinhole and phase mask reduce the DUV pulse energy from ~ 20 μJ exiting the final BBO to approximately 1.5 μJ in each of the +1 and -1 diffracted orders.

The two lenses pictured in Fig. 6.3(a) direct the +1 and -1 diffracted orders from the phase mask onto the sample to form the transient grating [207]. They constitute a 4f imaging system, with magnification equal to the ratio in the lens's focal lengths f_2 and f_1 . This results in a transient grating period on the sample of

$$\Lambda = \frac{\Lambda_{\text{PM}}}{2} \times \frac{f_2}{f_1}, \quad (6.9)$$

where Λ_{PM} is the phase mask period of 2 μm . Eq. 6.9 can be reconciled with Eq. 6.2 by noting that a lower value of f_2 will result in the beams crossing at a larger angle θ at the sample plane. The available lenses have nominal focal lengths of 6 cm, 7.5 cm, 10 cm, 15 cm and 20 cm, allowing for a minimum TG period of 300 nm when using the 2 μm phase mask. Chromatic aberrations due to the high refractive index experienced by the DUV light in the glass cause the real focal lengths to be approximately 17% shorter than nominal. This reduction in focal length was measured empirically and is fairly consistent across the various lenses. The lenses are made from varying grades of UV fused silica and absorb some of the DUV light. However, since the DUV power in each diffracted order is lower and the beam size larger, absorption is less of an issue as compared to the phase mask. However, substantial back reflections from the lenses, even those with anti-reflective coatings (20 cm, 15 cm and 10 cm), reduce the total DUV pulse energy at the sample to ~ 1 μJ , split between both beams. It necessary to block the zeroth diffracted order, transmitted without deflection through the phase mask, in order to produce a transient grating on the sample. If the two lenses are allowed to refocus all light exiting the phase mask, they will produce an image of the beam on the sample exactly corresponding to the beam as it exits the phase mask. Since the phase mask introduces a phase rather than amplitude modulation in the electric field, its replication on the sample will yield no variation in intensity. It is only by filtering the zeroth order and higher diffracted orders that an amplitude grating is produced. Blocking the higher diffracted orders also ensures that the transient grating excitation is a pure sinusoid, rather than the superposition of several spatial frequencies.

It is important that the DUV beams be well-aligned through the two lenses. A poor alignment can alter the TG period or rotate the TG fringes on the sample. The lenses are mounted in an optical cage, which comes with a “t-shirt” alignment tool designed to indicate the optical axis passing through the lens centers. While it is very important to place the phase mask and sample at the focal planes of the lenses, the experiment will still work if the two lenses are separated by a distance slightly different than the sum of their focal lengths. This is evident from a Fourier optics calculation in the thin lens approximation, using the operator notation introduction in Section 5.4 of Ref. [162],

$$\begin{aligned}
\mathcal{S} &= \mathcal{R}[f_2] \mathcal{Q} \left[-\frac{1}{f_2} \right] \mathcal{R}[f_2] \mathcal{R}[d - f_1 - f_2] \mathcal{R}[f_1] \mathcal{Q} \left[-\frac{1}{f_1} \right] \mathcal{R}[f_1] \\
&= \mathcal{V} \left[\frac{1}{\lambda f_2} \right] \mathcal{F} \mathcal{R}[d - f_1 - f_2] \mathcal{V} \left[\frac{1}{\lambda f_1} \right] \mathcal{F} \\
&= \mathcal{Q} \left[\frac{-(d - f_1 - f_2)}{f_2} \right] \mathcal{V} \left[\frac{-f_1}{f_2} \right], \tag{6.10}
\end{aligned}$$

where the system operator $\mathcal{S}$ gives the aggregate effect of the TG lens system on light incident on the phase mask. The operators $\mathcal{F}$, $\mathcal{R}[d]$, $\mathcal{V}[t]$ and $\mathcal{Q}[c]$ represent a Fourier transform, free space propagation over a distance d , multiplication by a factor t and multiplication by a quadratic phase $e^{ik(cx^2/2)}$, respectively. Eq. 6.10 shows that separation of the two lenses by a distance $d \neq f_1 + f_2$ only introduces a quadratic phase at the sample plane, in addition to the magnification factor $\mathcal{V}[-f_1/f_2]$. The quadratic phase will not affect the intensity of the sinusoidal TG pattern on the sample. However, ensuring that the distance between the two lenses is not significantly different than $f_1 + f_2$ does ensure that the beams cross at their respective focuses. In the case of a heterodyned probe as in Ref. [202], it also eliminates a potentially significant phase factor.

6.2.2 Probe beamline

To probe the transient grating some time after its excitation, it is necessary to divert a piece of the Ti:Sapphire amplifier output into the separate beamline shown in Fig. 6.4. The 1 kHz probe pulse train can then be delayed to arrive at the sample slightly after the pump. The probe beamline begins with the portion of the Ti:Sapphire amplifier output that transmits through the beamsplitter shown in Fig. 6.4, which subsequently reflects from two 45-degree mirrors (M1 and M2) and into a retroreflector mounted on a Newport XMS160-S delay stage (Delay 1). This delay stage can travel 16 cm, which translates to a temporal delay of approximately one nanosecond. As a “closed-loop” stage, it is also capable of extremely small, accurate motions, with an advertised repeatability of ~ 300 nm or < 1 fs in light travel time.

After the first delay stage, the beam reflects from two more 45-degree mirrors (M3 and M4)




Figure 6.4: Deep ultraviolet transient grating experiment probe beamline: Steering mirrors (M1-M8) and irises (I1-I4) facilitate the alignment. After separation from the pump via a beamsplitter, the probe beam passes through two delay stages. The first delay stage provides a 1 ns delay at femtosecond precision, while the second delay stage provides an additional 4 ns delay at picosecond precision. A BBO crystal converts the 800 nm to its second harmonic at 400 nm, to allow it to diffract from smaller TG periods. Alignment into a UV-Al mirror positioned on a rotation stage (M8), allows for measurements at the Bragg angle. A small fraction of the probe beam diffracts from the transient grating and onto a photodiode. See text for detailed alignment procedure.

before passing through a half-wave plate, polarizer and irises I1 and I2. It is not necessary to exactly align through these irises, instead they serve as a reference to indicate the degree to which the probe beamline is misaligned. If the probe beam is severely misaligned on the irises, it is often possible to recover its alignment by testing whether the movement of a single mirror, such as M2 or the magnetic mirror after the amplifier, will improve the alignment on both irises simultaneously. The combination of the half-wave plate and polarizer allows for the independent adjustment of the probe beamline power. The polarizer is rotated such that its back reflection travels to a beam block. The transmitted portion of the probe beam proceeds to the second delay stage (Applied Motion). This “open-loop” delay stage is far less precise than the Newport, having a repeatability that we estimate to be ~ 1 ps, but is longer, capable of introducing delays up to four nanoseconds. It is ideal for looking at long-timescale dynamics with lower time resolution. The retroreflectors mounted on the optical delay stages are triangular optics which have the property of reflecting light parallel to, but displaced from, its incident path. In the present context, this allows the delay stages, when properly aligned, to modify the probe beam path length without changing its position or pointing. Details on aligning the delay stages appear in the experimental documentation linked in [Appendix B](#).

In order to diffract from the smaller transient grating periods, it is necessary to use a shorter wavelength than the 800 nm amplifier output, because, while the diffraction limit would theoretically allow measurement of a ~ 400 nm grating, the very large diffraction angles involved make this impractical. The most straightforward option to decrease the probe wavelength is second harmonic generation in a BBO crystal, the same process occurring first in the pump beamline upconversion. Because it is important to minimize spot size on the BBO in order to improve conversion efficiency and mode quality, the combination of a plano-convex ($f = 15$ cm) and plano-concave ($f = -2.5$ - 3 cm) lens after the final delay stage form a telescope, which decreases the beam diameter by a magnification equal to the ratio of the focal lengths, while leaving it approximately collimated. The separation between the two telescope lenses must be very precise in order to produce a collimated beam, and care should be taken when aligning them to not accidentally produce a mode which focuses on, and burns, one of the subsequent optics. It is important that this telescope be located after the

delay stages, since if the probe beam were to be slightly uncollimated entering the delay stages, then stage movements would change the probe spot size, and therefore fluence, at the sample. The BBO used is identical to the first BBO in the pump beamline. In the current configuration, and with the half-wave plate maximized, the BBO can produce up to $\sim 60 \mu\text{J}$ of 400 nm light, far more than needed to probe a sample.

After the BBO, the probe beam immediately encounters a separator optic, which reflects the second harmonic while transmitting most of the fundamental wavelength, which travels to a beam block. The relative polarizations of the fundamental and second harmonic greatly impact the efficiency of the separator and subsequent dielectric mirrors in filtering the fundamental. The situation is complicated by the presence of the two retroreflectors, which change the polarization of incident light, generally resulting in an elliptically polarized beam [208]. Because the experiment is not greatly affected by the probe polarization at the sample, we optimize the position of the polarizer before the second delay stage until only 400 nm, but no 800 nm light is visible on a CMOS camera placed at the sample plane. Given the sensitivity of CMOS cameras to infrared, this is sufficient to ensure that the impact of any 800 nm light on the experiment is minimal, since due to dispersion in the BBO, 800 nm light arrives hundreds of femtoseconds before the main probe pulse and could excite electrons in an opaque sample, distorting the measurement. After the first separator, two dielectric mirrors direct the probe beam to a plano-convex lens of 50 cm focal length, which focuses the probe beam to a $\sim 200 \mu\text{m}$ spot. The long focal length is necessary to produce a sufficient depth of focus such that the probe beam is small, both at the sample and on the 1 mm detector. It is possible to lengthen the depth of focus, if needed, by tuning the telescope to be very slightly divergent, and adjusting the lens position.

The remaining optics in the probe beamline enable the relatively precise adjustment of the probe incidence angle on the sample. Depending on the penetration depth of the transient grating excitation, the probe incidence angle can greatly affect the observed diffraction signal. For thick gratings, often defined as those extending to a depth

$$d > \frac{\Lambda^2 n}{\lambda_P}, \quad (6.11)$$

where λ_P is the probe wavelength, Λ is the grating period and n is the sample refractive index, it is desirable to probe at the Bragg angle, θ_B ,

$$\sin(\theta_B) = \frac{\lambda_P}{2\Lambda}, \quad (6.12)$$

such that the diffracted signal adds constructively as the probe beam passes through the area of excitation (see Section 9.7 in Ref. [162]). Alignment using mirrors M5 and M7 through irises I3 and I4 sets the probe beam to be orthogonal to the optical axis of the pump beam as it passes through the two TG lenses. The probe beam then reflects from mirror M8, which is mounted on a tower of stages, allowing it to be both rotated and translated along two axes. Translation along the z direction indicated in Fig. 6.4 allows M8 to be moved out of the way during the alignment through I4. By first aligning M8 to retroreflect the probe beam along its incident path, and then rotating it by an angle

$$\theta_{M8} = \frac{\pi}{4} - \frac{\theta_B}{2}, \quad (6.13)$$

the probe beam will be incident on the sample at the Bragg angle. For thin transient gratings, or for measurements of highly absorbent samples such as metals, where only a reflected diffracted order is present, the exact probe angle is less important. It is now necessary to overlap the probe beam with the pump at the sample plane through adjustments to the vertical knob on M5 and translation of M8 in the x direction indicated in Fig. 6.4. In the current overlap procedure, a CMOS camera images a sample of similar thickness and adjacent to the target sample. The adjacent sample is covered in a thin layer of phosphor powder (LDP LLC VUV blue phosphor), which fluoresces at both the pump and probe wavelengths such that both beams appear as bright, easily visible spots.

When switching transient grating periods, it is often necessary to replace one or both of the

imaging lenses shown in Fig. 6.3(a). Since the second lens's focal length sets the working distance between the sample and imaging system, this often requires both M8 and the sample mount to be moved significantly. When this movement is outside the range of the translation stages, or would cause M8 to crash into the imaging system, some tedious trial-and-error rebuilding is required in order to maintain the Bragg angle alignment. A future iteration of this experiment may involve a cleverer approach to set the probe beam at the Bragg angle, such as the heterodyning scheme discussed in Section 6.1 and [202]. However, since chromatic aberrations cause the fused silica lenses to have very different focal lengths for a 200 nm pump and a 400 nm, or even 267 nm, probe, this future design will likely involve only reflective optics, such as off-axis parabolic mirrors.

6.2.3 Electronics and data acquisition

After the pump and probe beams are overlapped on the sample, the alignment portion of the experiment is finished. The following section covers the process by which a photodiode and lock-in amplifier measure the average intensity of the very weak diffracted probe beam at each delay stage position to produce a time-resolved measurement of the type shown in Fig. 6.6(f).

Fig. 6.5 illustrates the data collection process. A photoreceiver (New Focus 2107) with an adjustable gain of up to 3×10^4 measures the diffracted probe beam. This high gain is critical when measuring the very low intensity light. However, the photoreceiver also amplifies any incident background scattered light, which is inevitably generated from the surface roughness of the various optics, their back reflections and the sample itself. The intensity of the background scatter is generally much higher than the diffracted probe beam. A lock-in amplifier (SRS SR830 DSP) is used to filter the desired signal, which requires the diffracted probe beam and noise to have different frequencies. This is accomplished by using the optical chopper (Newport 3502) discussed in Section 6.2.1, which blocks every other DUV pump pulse, reducing the frequency of the DUV pump from 1 kHz to 500 Hz. The chopper triggers off a 1 kHz laser reference frequency supplied from a BNC port on the amplifier electronics rack. The same reference frequency is supplied to a digital delay generator (DDG) (SRS DG535), which is programmed to divide it in half and supply this 500 Hz

frequency as a reference to the lock-in amplifier. The chopper is also capable of outputting a 500 Hz frequency to trigger the lock-in, however it is much less stable than the DDG output. The 500 Hz DUV pump and 1 kHz probe allows the lock-in amplifier to filter the 500 Hz diffracted probe from the 1 kHz background scatter. While background scatter from the DUV pump at 500 Hz could interfere with the diffracted probe, the photoreceiver is relatively insensitive at 200 nm and is further protected by a BG39 bandpass filter. The analog output from the lock-in amplifier is in the range of -10 V to +10 V. It is digitalized by a DAQ card (National Instruments BNC 2110), so that it can be read by Bruno Mars, the B1B35 lab computer. The DAQ card handles roughly half the voltage range of the lock-in amplifier, so it is important to set the lock-in sensitivity to a sufficiently low value that the DAQ will not saturate.

The data collection process described above generates a single voltage measurement corresponding to the diffraction efficiency of the probe from the transient grating at a given pump-probe delay time. Recording this voltage measurement as a function of delay stage position yields the time-resolved data shown in Fig. 6.6(f). The first delay stage (Newport XMS 160) communicates with the lab computer via its motion controller (Newport XPS-Q). The second delay stage consists of a motor (Applied Motion HT23-400) attached to a long screw, on which the retroreflector is mounted and a driver (Applied Motion Si3540), which also communicates with the lab computer. The Python code used to iteratively move the delay stages and acquire data is saved on Bruno and in a GitHub repository (see Appendix A).

Finally, it can be very difficult to place the photoreceiver at the location of the diffracted probe beam when setting up the experiment, since this beam is too weak to see on a notecard or with a camera. By setting the photoreceiver height and finding time zero using the transmitted probe beam, it should only be necessary to sweep the detector in one dimension. If the signal still cannot be found, it is often due to a vertical offset between the diffracted and transmitted probe, which can be caused by a rotated TG pattern, and can be fixed by re-aligning the DUV pump beam through the two lenses so that the two pump beams are equidistant from the optical axis. Alternatively, if there is parasitic heterodyning, a signal can sometimes be found by slightly




Figure 6.5: Deep ultraviolet transient grating beamline data acquisition: A reference signal at the 1 kHz amplifier repetition rate triggers the optical chopper (i), which spins to block every other DUV pump pulse. It also triggers a digital delay generator (ii), which is programmed to output a reduced reference signal at 500 Hz, corresponding to the repetition rate of the chopped DUV pump. The photoreceiver (iii) consists of a photodiode and high-gain amplifier, which converts the weak 500 Hz diffracted probe beam and stronger 1 kHz background scatter into an electronic signal of MHz bandwidth. This signal travels to the lock-in amplifier (iv), which also receives the 500 Hz signal from the digital delay generator as a reference, allowing it to filter the diffracted probe from the noise. The lock-in amplifier passes the filtered signal to a DAQ card (v), which digitalizes it and transmits it to the lab computer. Depending on the scan type, the computer moves either the long delay stage (vi) via its stepper (vii) or the short delay stage (viii) via its controller (ix) to adjust the pump-probe time delay, before repeating the data collection process. Images from the vendor websites.

translating the sample or phase mask.

6.2.4 Beamline diagnostics

The DUV-TG beamline introduced in the previous sections is unique among tabletop experiments in that it combines both high spatial and temporal resolution with a high pump photon energy capable of investigating dynamics in wide-bandgap samples. Fig. 6.6 provides several important diagnostics acquired during the development of the beamline, which allow for a comparison to previous TG implementations. Fig. 6.6(a) shows the third and fourth harmonic and fundamental beam modes, measured by a CMOS camera, near their focus after separator S7. For the reasons discussed in Section 6.2.1, the DUV mode is quite poor. However, it can be improved by the insertion of a 100 μm Fort Wayne wire die as a pinhole at the focal plane, as shown in Fig. 6.6(b), which has the added benefit of spatially filtering the remaining light from the other wavelengths. Fig. 6.6(d) shows the DUV spectrum, centered around 197 nm, as measured by an Ocean Optics spectrometer. The DUV spectrum has a low shoulder extending to nearly 205 nm wavelength. Since the Ti:Sapphire spectrum extends from roughly 755 nm to 820 nm (tail to tail), this shoulder may be caused by poor phase matching of the longer Ti:Sapphire wavelengths during the upconversion process. The bandwidth cutoff around 194 nm could be caused by a variety of factors, including poor phase matching and cutoffs in the dielectric separator reflectivities, whose coatings are difficult to optimize for broad bandwidths.

Fig. 6.6(e) shows a measured upper bound on the DUV pulse duration. This is determined by placing a copy of BBO3 directly after M2 in Fig. 6.3(b). By rotating the 800 nm probe beam's polarization and directing it through the back side of M2 such that it overlaps with the DUV, this BBO can be made to generate the third harmonic through difference frequency generation. The third harmonic can be spatially separated by making the incident beams slightly non-collinear, and will only appear when the pump and probe are temporally overlapped in the BBO. By measuring the intensity of the third harmonic as a function of delay stage position, we arrive at a correlation with a Gaussian-fitted 285 fs FWHM, which is the convolution of the pump and probe pulse durations,

plus some broadening due to dispersion in the BBO itself. This value increases by approximately 50 fs when the compressor grating in the amplifier is detuned by a half turn, which also causes a noticeable decline in DUV power. It also increases by approximately 50 fs when one of the imaging lenses is placed in the DUV beam before the BBO, suggesting that the glass in the imaging system only moderately increases the DUV pulse duration at the sample. Fig. 6.6(c) shows an atomic force microscopy image of a TG interference pattern etched into a 50 nm spin-coated PMMA film. The period is 287 nm, achieved using the 2 μm phase mask period, reduced to 1 μm on the sample via the factor of two in Eq. 6.9, and then de-magnified by $M \sim 0.3$ using lenses of 20 cm and 6 cm nominal focal length. It is possible to generate a smaller TG period by increasing the ratio of focal lengths or tilting the phase mask as discussed in Ref. [209]. However, it is difficult to probe a smaller grating with 400 nm light. The smallest TG period we have etched in a PMMA film using the setup described above is ~ 230 nm.

The example data in Fig. 6.6(f) illustrates the importance of short pump and probe pulse durations. It is from the dataset discussed in Section 6.4 and measures carrier diffusion in diamond at 287 nm period. The non-negligible signal rise time is attributable to the pulse duration, since electrons thermalize on a < 100 fs timescale before diffusing outwards from the interference fringes. A picosecond-scale pulse duration would be too long to accurately measure the diffusion dynamics at this length scale.

6.2.5 Carbon deposition and parasitic heterodyning

Carbon deposition and parasitic heterodyning can affect the scan-to-scan repeatability of the transient grating signal. Carbon deposition occurs when the high energy pump and probe photons “crack” atmospheric hydrocarbons in the vicinity of the sample, causing a contamination layer to accumulate on its surface [210, 211]. On diamond, carbon deposition causes the TG signal to increase approximately linearly on a minutes-long timescale, even when the delay stages do not move. The linear scaling of an otherwise decaying signal complicates analysis of the electron and phonon dynamics. Because we have only regularly observed carbon deposition on diamond, it is




Figure 6.6: Deep ultraviolet transient grating beamline diagnostics: (a) CMOS camera image of the fundamental, third and fourth harmonics after the last BBO crystal near the focus. The modes of the higher harmonics are poor, because the upconversion process amplifies the nonlinearities and imperfections in the out-of-focus fundamental mode. (b) A 100 μm pinhole (Fort Wayne Wire Die) placed near the focus both improves the fourth harmonic mode by filtering its higher spatial frequency components, and spatially filters the remaining harmonics. (c) Atomic force microscopy image of a 287 nm transient grating period etched into a PMMA film. (d) Spectrum and (e) pulse duration of the fourth harmonic. The pulse duration measurement is a correlation of the fourth harmonic and fundamental obtained via difference frequency generation in a BBO. (f) Example data on single-crystal diamond at 287 nm TG period, showing the probe diffraction signal as a function of pump-probe delay time. The signal rise time lasts hundreds of femtoseconds due to the pulse duration.

unknown whether it always results in an increasing signal.

Carbon deposition is a multiphoton process, since the photon energies used in the DUV-TG experiment are lower than the chemical bonding energies. It is therefore possible to eliminate significant carbon deposition by reducing the pump and probe fluence. On diamond, we observe no significant pump-induced carbon deposition, perhaps because the higher-energy pump photons also create ozone, which can etch away carbon contamination. Probe-induced carbon deposition on diamond is virtually absent below 1 μJ pulse energy, at a 200 μm spot size, but very significant above 2-3 μJ .

Heterodyning in transient grating experiments is discussed in Section 6.1. In the context of the present experiment, we sometimes observe “parasitic” heterodyning caused by probe light scattering from sample surface roughness or defects at the same angle as the diffracted probe beam [151]. Probe light which scatters off anything besides the sample and onto the detector cannot cause heterodyning, since it will not be temporally overlapped with the diffracted probe. Pump light also cannot cause heterodyning, since it has a different wavelength and will not interfere with the probe. Since the reference light must come from the sample, the degree of parasitic heterodyning depends on the magnitude and spatial frequency of its surface roughness. On diamond, we observe that heterodyning increases at larger, micron-scale TG periods, suggesting that micron-scale surface roughness either dominates, or is better at producing reference light with a stable phase. We also observe increased heterodyning after significant carbon deposition, or on diamond samples with known surface contamination. Contamination-induced heterodyning might be explained by the DUV pump’s etching of a permanent grating into the contamination layer, which diffracts reference probe light along the same optical path as the transient grating signal.

The degree to which a TG signal is heterodyned can be checked by observing how it responds when the sample or phase mask is translated horizontally by fractions of a TG period. Either adjustment will shift the TG pattern on the sample, which will change the phase of the diffracted probe light, but not the reference light, modifying the linear, heterodyned term in Eq. 6.7. If parasitic heterodyning is contributing significantly to the TG signal, it will cause the signal on the

lock-in amplifier to oscillate periodically as the sample or phase mask translation sweeps through multiple TG periods. If the signal can be made fully negative, the heterodyned term likely dominates. Except in this extreme case, it is difficult to quantify the relative magnitudes of the heterodyned and non-heterodyned contributions due to phase and amplitude instabilities in the scattered reference light. However, even if the relative amplitudes of the two terms in Eq. 6.7 are changing from scan to scan, the decay time constants should remain consistent, as was observed for the diamond carrier diffusion data discussed in Section 6.4.

6.3 Validation using surface acoustic waves launched in gold films

As a validation of the new experiment, we measure the velocity of surface acoustic waves (SAWs) and longitudinal acoustic waves (LAWs) launched by the transient grating excitation in gold films purchased from Angstrom Engineering.¹ The gold films are 50 nm thick on 10 mm square <100> p-doped silicon substrates, with an approximately 5 nm chromium adhesion layer underneath the gold. The transient grating period and film thickness set the SAW and LAW wavelengths, respectively. The experimental data, shown in Fig. 6.7, provides the frequency, and the product of the wavelength and frequency is the acoustic velocity. The frequency can be extracted from the experimental data using a simple Fourier transform. Alternatively, a chirp Z-transform gives more precise sampling of the frequency space [212]. In cases where only a few oscillations can be captured within the range of the optical delay stage or when the oscillations are highly damped, the matrix pencil method discussed in Section 4.4.4 performs best, and is used in the analysis presented here.

Because the longitudinal acoustic wave reflects from the gold-chromium interface and is therefore confined to the gold film, the 3125 m/s velocity measured by the DUV-TG experiment can be directly compared to the 3240 m/s tabulated longitudinal acoustic velocity in gold [159]. The slight deviation could be attributable to strain or a nanometer-scale difference in the film thickness from its nominal value.

The penetration depth δ of a surface acoustic wave depends on its wavelength as $\delta \sim \lambda_{\text{SAW}}/\pi$,

¹<https://substratasolutions.com/gold-coated-silicon/>




Figure 6.7: Deep ultraviolet transient grating acoustic measurements and validation: Experimental measurements of (a) longitudinal acoustic waves and (b) surface acoustic waves propagating in a 50 nm gold film on a $\langle 100 \rangle$ silicon substrate. (c) Finite element model to calculate the SAW dispersion. The top, bottom and sides of the simulation are subject to free, fixed and periodic boundary conditions, respectively. The color map shows displacement as a function of position for the SAW mode identified by the eigenfrequency calculation. (d) SAW velocity as a function of wavelength, which is equal to the TG period, for the experimental data and finite element simulations.

as discussed in Section 4.4. As with the silicon metalattice, we construct a finite element model of the full sample, shown in Fig. 6.7(c), using nominal elastic properties and densities for the three layers. Due to the lack of nanostructures, the SAW mode can be identified manually from an eigenfrequency calculation, without the SAW-likeness analysis used for the metalattice. Fig. 6.7(d) shows good agreement between the measured and predicted SAW dispersion, confirming that the experimental time delay is calibrated correctly. For each experimental point, the SAW wavelength is measured by exposing a 50 nm PMMA film to the DUV pump beam for approximately 30 seconds, which etches the sinusoidal pattern into the soft polymer, allowing its period to be measured using an optical microscope.

6.4 Probing carrier diffusion in diamond at <300 nm TG period

Due to the difficulty of studying diamond with visible metrology techniques, little is known about its transport properties in device-relevant nanoscale geometries. This is particularly relevant for thermal transport in nanostructured diamond, because the mean free paths of diamond's heat carrying phonons mostly fall between 0.5-1 μm , far larger than most modern semiconductor devices [60]. Transient grating measurements have been used previously as a non-contact means to observe both electron and thermal diffusion in diamond, with previous investigations falling into three broad categories. Early experiments used an infrared pump, either near 4 μm to take advantage of multi-phonon absorption or tuned to an absorption band associated with some defect [213, 214]. These early experiments succeeded in measuring diamond's thermal diffusivity, but were limited to TG periods $>10 \mu\text{m}$ due to the long excitation wavelength. A decade later, experiments began using the fifth harmonic of a 1064 nm laser, which at 213 nm is just below diamond's 225 nm optical absorption edge [215, 216]. However, pulse durations were limited to ~ 25 ps, preventing measurements at sub-micron length scales where electron and phonon dynamics happen too quickly. Additional experiments used femtosecond lasers to excite transient gratings in diamond via two-photon absorption at slightly longer wavelengths [217]. Finally, a recent free electron laser TG implementation using 98 eV extreme ultraviolet light produced a 277 nm transient grating on a

diamond substrate, which provides a useful comparison to the present experiment [204].

Fig. 6.8 illustrates the relaxation, through carrier diffusion, of DUV transient gratings in single-crystal diamond. The CVD diamond sample is purchased from Applied Diamond and has ~ 500 μm thickness, <1 ppm nitrogen impurities and is double-side polished to <3 nm Ra surface roughness to enable measurements in the transmission configuration. After polishing, the diamond is dry-etched in O_2 to remove any subsurface damage and then cleaned in a boiling nitric and sulfuric acid mixture. X-ray reflectivity confirms minimal surface contamination prior to measurement. As shown in Fig. 6.8(a), the achievable TG period in the current experiment extends down to 287 nm, limited only by the 400 nm probe wavelength. It is possible to produce a ~ 230 nm grating using standard two-inch optics, verified by etching it into a PMMA film and then taking an AFM image. Investigation of electron and phonon transport at these lengths scales was previously only possible via nanofabrication or through free electron laser (FEL) TG implementations [64, 204].

Modeling of the DUV-TG signal proceeds from the assumption that the non-equilibrium electron and hole distributions excited by the DUV pump thermalize locally on a <100 fs timescale, shorter than the smallest experimental time step. Afterwards, the local excited carrier density is assumed to induce a proportional change in refractive index, Δn , which serves as an initial condition for the subsequent electron and hole diffusion,

$$\Delta n(x, z, t = 0) = \Delta n_0 [1 + \cos(qx)e^{-\beta z}], \quad (6.14)$$

where x and z are coordinates in the in-plane and cross-plane directions respectively, $1/\beta$ is the pump penetration depth, and $q = 2\pi/\Lambda$ is the TG wavevector magnitude. The penetration depth of 200 nm light in diamond is approximately 500 nm [219], comparable to Λ , such that carrier diffusion occurs in both the in-plane and cross-plane dimensions on similar timescales. The 2D diffusion equation in a semi-infinite medium with Eq. 6.14 as its initial condition can be solved analytically, yielding




Figure 6.8: Nanoscale carrier diffusion in single-crystal diamond: (a) Measured decays due to in-plane carrier diffusion at various TG periods. The decays have a double-exponential form because probe light scattered from the sample interferes with the diffracted signal to introduce a heterodyned term. (b) Fitted time constants for the non-heterodyned (red) and heterodyned (blue) components of the decays. As expected for carrier diffusion, the time constant scales as the square of the TG period and is twice as large for the linear heterodyned term. (c) Fitted ambipolar diffusion coefficient as a function of estimated carrier concentration. In the present experiment (black), fits to the heterodyned and non-heterodyned decay components are aggregated at each TG period, corresponding to different excited carrier concentrations, because the magnification system used to change TG period also changes the pump spot size. Past DUV [218] (red, green) transient grating experiments observe a similar decrease in ambipolar diffusion coefficient with carrier concentration. (d) Fitted amplitude of the heterodyned and non-heterodyned decay components as a function of TG period, illustrating the lower heterodyned contribution at small sizes.

$$\Delta n(x, z, t) = \Delta n_0 [1 + \cos(qx)] e^{-Dq^2 t} e^{-\beta(z - D\beta t)} \left[2 - \text{Erfc} \left(\frac{z - 2D\beta t}{2(Dt)^{1/2}} \right) + e^{\beta z} \text{Erfc} \left(\frac{z + 2D\beta t}{2(Dt)^{1/2}} \right) \right], \quad (6.15)$$

where D is the isotropic ambipolar diffusion coefficient and $\text{Erfc}[x] = 2/\sqrt{\pi} \int_x^\infty e^{-u^2} du$ is the complementary error function. The probe beam diffracts from this 2D refractive index profile as it passes through the sample at some time delay. It can be demonstrated using the analytical methods in Section 9.7 of Ref. [162] or a numerical approach such as rigorous coupled-wave analysis that a probe beam of wavelength λ_p incident on the transient grating excitation at the Bragg angle $\sin \theta_B = \lambda_p/2\Lambda$, will result in a diffraction intensity $I_p(t)$ on the detector sensitive only to the in-plane component of the diffusion equation solution in Eq. 6.15, such that

$$I_p(t) \sim |E_p(t)|^2 \sim e^{-2Dq^2 t}, \quad (6.16)$$

where $E_p(t)$ is the electric field of the diffracted probe (see Section 6.1). This simplification at the Bragg angle occurs due to the constructive interference of the probe beam as it diffracts from the excitation profile across the sample depth. As the probe beam incidence angle deviates from Bragg, there is an increasing sensitivity to the more complex cross-plane dynamics.

As illustrated in Fig. 6.8(a), the experimental TG decays are observed to have a double-exponential form,

$$I_p^{\text{exp}}(t) = A_1 e^{-t/\tau_1} + A_2 e^{-t/\tau_2}, \quad (6.17)$$

with time constants τ_1 and τ_2 . Both time constants scale as $\tau \sim q^{-2} \sim \Lambda^2$, as expected from the in-plane diffusive process outlined in Eq. 6.16. Because τ_2 is observed to be approximately twice τ_1 across all TG periods, the second exponential decay can be attributed to parasitic heterodyning [215]. This occurs when the time-varying diffraction signal of amplitude $E_p(t)$ interferes with scattered probe light of constant amplitude E_s and relative phase ϕ on the detector to produce a signal whose

intensity is proportional to $E_s E_p(t) \cos(\phi)$, as given by Eq. 6.7 in Section 6.1. Comparison with Eqs. 6.16 and 6.17 allows the identification $\tau_2 = 2\tau_1 = D/q^2$. The resulting ambipolar diffusivity fit is $13.8 \pm 3.0 \text{ cm}^2/\text{s}$ across the full dataset presented in Fig. 6.8. The attribution of τ_2 to parasitic heterodyning is further supported by the observation that the fitted amplitude A_2 of the second exponential decay decreases with decreasing TG period, as shown in Fig. 6.8(d), because at larger diffraction angles there is less static background scatter from the sample surface. The parasitic heterodyning observed here is due to random scatter, in contrast with the heterodyning approach discussed in Section 6.1 and Ref. [202], where a portion of the probe beam itself is deliberately overlapped spatially, and if necessary temporally, with the diffracted signal on the detector. However, even for parasitic heterodyning it is possible to adjust the relative phase ϕ in Eq. 6.7 by translating the phase mask laterally by a fraction of its $2 \text{ }\mu\text{m}$ period using a sensitive micrometer screw. This translates the TG pattern on the sample by some fraction of its period, which introduces a phase shift into the diffracted light. For the data presented in Fig. 6.8, we crudely optimize the phase mask position to maximize the interference, such that the observed A_2 is positive.

We do not observe thermal signal in the present experiment. Previous studies, both those using infrared and ultraviolet pump lasers, observed thermal signal on diamond only at TG periods above $10 \text{ }\mu\text{m}$. For ultraviolet excitations, this is likely because the microsecond decay timescales of these large transient gratings exceed typical electron-hole recombination timescales, allowing energy from the electron system to be transferred to the phonons before the in-plane diffusion begins. In the present experiment, the in-plane carrier diffusion occurs before significant recombination, so that electrons, rather than phonons, are the main energy carriers. Recombination timescales in diamond are on the order of a nanosecond, but can vary greatly with sample quality and impurities. The free electron laser TG implementation did observe a thermal phonon decay in diamond at 277 nm TG period [204]. This is partially attributable to the high signal to noise ratio of the experiment, since the thermal signal was $>100\times$ weaker than the electronic signal. However, the FEL experiment also uses extreme ultraviolet light at 98 eV , compared to 6.3 eV in the present experiment. Given diamond's 5.5 eV bandgap, the FEL experiment excites electrons far above the

conduction band minimum, allowing them to exchange most of their energy with the lattice through fast electron-phonon scattering even before recombination. This additional energy in the phonon system at short timescales likely enhances the magnitude of the thermal signal as compared to the present experiment.

In the present experiment, the diffusion coefficient increases slightly with TG period when fitted separately for each value of Λ . This occurs because the magnification system used to vary the TG period results in an increasing pump beam spot size, and therefore decreasing excited carrier concentration N (in electrons or holes per cm^3) as Λ increases. As shown in Fig. 6.8(c), a decrease in the diffusion coefficient with increasing carrier concentration is consistent with other DUV transient grating experiments on diamond [216, 218, 220, 221]. Time of flight experiments at much lower carrier concentrations, $N \sim 10^{13}/\text{cm}^3$, measure $D > 50 \text{ cm}^2/\text{s}$, although this is generally a conversion from a mobility measurement via the Einstein relation [222, 223]. The decrease in D with carrier concentration is generally attributed to a combination of carrier-carrier scattering [218, 224] and many-body effects, which create a gradient in the bandgap along the spatial variation of N in opposition to the carrier diffusion [225]. Although the many-body effect is modest in many other semiconductors, it is expected to be larger in diamond due to diamond's large effective carrier mass [226]. We estimate the present experiment's excited carrier density to be $4 \cdot 10^{17} < N < 4 \cdot 10^{19}/\text{cm}^3$ by assuming a DUV penetration depth of 500 nm and assuming that each photon excites one electron-hole pair. We then estimate a pump fluence for each value of Λ by measuring the width of the TG pattern etched into a PMMA film placed at the sample plane.

Other mechanisms have been proposed to modify the diffusion coefficient at high carrier concentration, including exciton diffusion and bandgap shifts due to lattice temperature gradients [227]. The effects of exciton diffusion are thought to be significant only at lower temperatures [228, 229, 230]. The thermal bandgap effect however is quite interesting in the context of comparisons between the present experiment and EUV transient grating implementations. In an EUV transient grating, electrons are excited very high above the conduction band minimum and can scatter rapidly to establish a large lattice temperature gradient relative to a DUV excitation of equal fluence.

A large lattice temperature gradient can then induce a bandgap gradient similar to that of the many-body effect. A controlled comparison of carrier diffusion induced by EUV and DUV transient gratings could elucidate the relative magnitudes of these various effects.

6.5 Probing nanoscale phonon transport in silicon

Previous transient grating measurements of silicon have shown signal from both electronic and thermal processes on differing time scales, and the onset of non-diffusive phonon transport at TG periods over $1\text{ }\mu\text{m}$ [54, 203]. Using the new deep ultraviolet transient grating experiment at JILA, it should be possible to measure much smaller TG periods and inform ballistic and hydrodynamic interpretations of nanoscale phonon transport at shorter length scales. The current section briefly discusses these ongoing measurements.

After initially measuring semi-infinite bulk silicon substrates in a reflection geometry, where the photodetector collects the reflected diffracted probe beam, we observed higher quality signal in thin silicon membranes measured in transmission. This is consistent with the general observation that transmission measurements give higher signal but may also relate the $\sim 50\text{ nm}$ membrane thickness, which traps hot electrons near the surface, keeping the full excitation in view of the probe. After receiving a batch of membranes from Silson, which were far thicker than expected, we began verifying membrane thickness using the spectrophotometer and ellipsometer in the JILA Keck Lab. The former measures membrane optical transmissivity as a function of wavelength while the latter measures the change in polarization upon reflection. The spectrophotometer and ellipsometer data can be fit to optical models, using rigorous coupled-wave analysis (RCWA) and custom software on the ellipsometer computer, respectively. The membrane thickness fits obtained using these two methods generally agree within $\sim 2\text{ nm}$. Fig. 6.9(a) shows spectrophotometer data and RCWA fits for two 4 mm square silicon membranes in 10 mm square frames manufactured by Norcada. While both membranes are nominally 50 nm thick, their real thicknesses can be quite different, an important consideration given that the apparent thermal conductivity of a thin film varies as a function of its thickness. In conversation, Norcada revealed to us that they measure membrane thickness when the

silicon layer is still a film on a dielectric substrate, which is subsequently etched away. Membranes from the same batch with very different thicknesses likely come from different locations on the wafer. Neither the RCWA nor ellipsometer fits improved with the inclusion of a significant oxide layer, while the inclusion of a ~ 1 nm native oxide layer had a negligible impact. Finally, the Norcada membranes are boron doped, likely to the 10^{15} or 10^{16} atoms/cm³ level necessary to produce a standard $1 - 30 \text{ } \Omega \cdot \text{cm}$ resistivity, which may be an important factor to consider in the thermal modeling.

The silicon membrane TG signal depends greatly on the probe wavelength, as shown in Fig. 6.9(b) and (c) for a $1 \text{ } \mu\text{m}$ TG period. When using an 800 nm probe, there is a fast decay on early timescales which resembles the diamond data shown in Fig. 6.8(a) and Fig. 6.6(f), and which is also attributable to in-plane carrier diffusion. The subsequent, longer-timescale dynamics are attributable to phonon thermal transport. It can be inferred from Fig. 6.9(c) that thermal expansion and electronic excitation have opposite effects on the refractive index experienced by the 800 nm probe. Assuming two contributions to the $\Delta n(t)$ term in Eq. 6.5 of opposite sign, which decay exponentially with time constants τ_e representing the electrons and τ_p the phonons, the diffracted probe amplitude would have the form

$$E_p(t) \sim -A_e e^{-t/\tau_e} + A_p e^{-t/\tau_p}, \quad (6.18)$$

Eq. 6.18 would result in a sign flip between the electronic and thermal signals if the measurement were heterodyned, as observed in Ref. [203]. A negative signal corresponds to destructive interference between the probe and reference light, such that the overall intensity on the photodetector decreases in the presence of the diffracted probe. However, since the signal in Fig. 6.9 only drops to zero around 50 ps, it is likely that the signal is not appreciably heterodyned in the present experiment and therefore varies as $|E_p(t)|^2$. The reduction in parasitic heterodyning as compared to the diamond measurements in Section 6.4 is likely due to the low silicon membrane surface roughness, $<1 \text{ nm Ra}$, as quoted by Norcada, versus $\sim 3 \text{ nm}$ for diamond. Because of the short probe pulse duration,




Figure 6.9: Deep ultraviolet transient grating measurements of silicon membranes: (a) Spectrophotometer measurements of silicon membrane transmission reveal that membranes of 50 nm nominal thickness from Norcada can vary in thickness by over 25 nm. The solid dark and light blue curves are transmission data for membranes of 23 nm and 35 nm thickness, respectively, and the dashed red curves are fits calculated using rigorous coupled-wave analysis. Comparison of short (b) and long (c) timescale dynamics of a 1 μm TG period in a membrane of 50 nm nominal thickness measured with an 800 nm (red) and 400 nm (blue) probe. The 800 nm probe has a far stronger electronic signal at early times. Because the 800 nm signal fully decays at ~ 50 ps before rising again, it can be inferred that the electronic and thermal signals have opposite phase, which is not the case for the 400 nm probe. The 800 nm scans exhibit additional peaks at approximately 300 ps and 600 ps due to a multiple pulsing issue with the amplifier discussed in the text.

scattered probe light must follow a very similar optical path as the diffracted probe beam in order to temporally interfere on the detector, which practically means it must scatter from the sample surface, from surface roughness containing the same spatial frequency as the transient grating excitation.

The additional peaks at approximately 300 and 600 ps in the 800 nm signal in Fig. 6.9(c) are due to smaller pulses which escape the Ti:Sapphire amplifier laser cavity before the main pulse. This issue with the B1B35 amplifier went undetected before the DUV-TG experiment, since unless one is measuring phonon transport it is generally unnecessary to delay the pump and probe pulses by hundreds of picoseconds. However, the multiple pulses are observable at the amplifier exit using a fast photodiode and the 3 GHz oscilloscope from the electronics shop. The amplitudes of the side pulses are not affected by the Pockels cell alignment but can be reduced by adjusting the pump timing and seed overlap on the amplifier crystal (see Appendix C). However, it is difficult to reduce the side pulse amplitude below a factor of 1% of the main pulse. At <1 ns, the pulse spacing is too short to correspond to a full cavity round trip, and instead may originate from an amplified back reflection or an additional pulse from the ~ 80 MHz oscillator pulse train which enters the amplifier cavity and makes a different number of passes relative to the main pulse before exiting.

Because the conversion efficiency of the 800 nm light varies nonlinearly with intensity, the side pulses do not generate appreciable DUV light after the three BBOs in Fig. 6.3, meaning that the sample is not pumped multiple times in succession. Instead, the multiple spikes evident in Fig. 6.9(c) correspond to two smaller 800 nm probe pulses arriving before the main one. Even a single BBO is sufficient to greatly reduce the amplitude of the side pulses, as illustrated by their disappearance for the 400 nm probe measurement in Fig. 6.9(c). At 400 nm probe wavelength, the electronic and thermal signals have the same sign, in agreement with Ref. [54], and the thermal signal is relatively much stronger. The 50 nm membrane thickness is particularly important here, since even at 200 nm, transmission measurements would be severely attenuated due to the short probe penetration depth.

Given the strong thermal signal and the potential to measure very short TG periods, this

configuration may reveal interesting nanoscale thermal transport behaviors. For example, the ballistic framework predicts a ~ 50 nm phonon mean free path limited by the membrane thickness, which would imply that the apparent in-plane thermal conductivity for TG periods above a few hundred nanometers should not vary with the grating size, as it does in thicker films [203]. In the hydrodynamic picture, the $\nabla^2 \mathbf{q}$ term discussed in Section 3.3.3 will likely be attenuated due to the reduced value of its coefficient ℓ^2 inside the film, which could cause the dynamics to appear Fourier-like even at small TG periods. Finally, the 200 nm pump beam has a very short penetration depth in silicon and likely does not heat the membrane evenly in the cross-plane direction. In the case of uniform and nonuniform cross-plane heating, phonon transport occurs in one and two dimensions, respectively. Since a 400 nm pump will heat a 50 nm membrane relatively uniformly, while a 267 nm pump will have an even shorter penetration depth than at 200 nm, it may be possible to explore the behavior of 1D vs. 2D phonon transport in thin films by varying the pump wavelength.

Chapter 7

Conclusions and outlook

The rise and fall of experiments in B1B35 continues with the intense deep ultraviolet light source developed in this thesis, which is poised to produce a variety of new and exciting results. Immediate improvements to the transient grating experiment will likely include a thinner phase mask to boost pump power and perhaps an all-reflective design using off-axis parabolic mirrors to enable heterodyned measurements. Temperature control is always on the horizon, due to the interesting phonon dynamics accessible in both bulk samples and phononic crystals as thermal phonon mean free paths and wavelengths increase, and quantum effects emerge. To access even shorter length scales, the 400 nm probe could drive an HHG process to produce 57 nm light, which would have the added benefit of isolating the surface deformation from the electronic signal. However, this would be a significant experimental challenge due to the fluence and stability requirements for both the pump and probe arms. As an imaging adaptation, the development of a DUV interference microscope would allow a visible probe to spatially resolve dynamics in previously inaccessible wide-bandgap samples [231]. There is also the possibility of moving the DUV-TG experiment to the PHAROS in the X-wing, whose higher repetition rate would improve signal quality. Ongoing experiments with the PHAROS aim to produce bright vacuum ultraviolet light at 147 nm [232].

A strength of the DUV-TG experiment is the wide variety of samples which it can measure. This not only includes non-contact investigations of wide-bandgap semiconductors such as diamond, gallium oxide and gallium nitride, but also materials for which the electron-beam lithography process outlined in Chapter 5 would be ineffective or destructive. This includes energy materials such as

vanadium oxide and LLZO, with their high surface roughness, and delicate 2D materials including TMDs and metal halide perovskites. The ultimate challenge and opportunity would be to access thermal and acoustic dynamics in polymers and soft materials. While the DUV pump clearly etches away PMMA, and the 4 ns delay stage may prove inadequate, if these challenges can be overcome it opens up a relatively unexplored area as compared to semiconductors, starting with the complex interaction between the DUV photons and the organic molecules. Collaborations will be key in identifying promising samples to pursue.

While the DUV experiment is just reaching adolescence, the EUV scatterometry experiment in B1B19 continues to deliver surprises in its old age. New transient EUV reflectivity measurements of VO_2 in collaboration with the magnetism team will feature in an upcoming publication. Preliminary measurements of nanostructured diamond, albeit in the presence of the amorphous carbon layer discussed in Chapter 5 showed surprising nanoscale effects. A relatively straightforward experiment involving nickel gratings on supported silicon films of varying thickness could elucidate the 2D nature of the close-packing effect suggested by the atomistic simulations of Ref. [121]. When the EUV scatterometry experiment is eventually rebuilt, it will likely feature many of the modern advancements soon to be ubiquitous in the KM group, including noise cancellation via an EUV beamsplitter and improved differential pumping for HHG stability. The rebuild can also include a new twist, whether it is temperature control, an adjustable camera capable of capturing many diffracted orders, enabling nanoscale spatial reconstruction of the surface profile using the methodology of Section 4.4.5, or the long-awaited realization of tabletop EUV transient gratings.

From the theoretical side, the same debates on nonlocality, collective effects and emergent hydrodynamics discussed in the context of highly-nonequilibrium phonon transport in Chapter 3 are playing out for electrons and other more exotic excitations across condensed matter physics [233, 234]. The development of effective theories to describe these phenomena is entering a golden age, poised to provide rich physical insight and complement the brute force approaches enabled by modern computation. Harnessing the time resolution of the DUV and EUV sources in this thesis, it may be possible to investigate transient collective effects in systems beyond phonons, such

as for electrons, where most prior research has concentrated on vortices and other phenomena in the spatial domain [235]. A further challenge is to formalize the rather fuzzy distinction between boundary and nonlocal or “intrinsic” effects in highly-nonequilibrium situations presented in Chapter 3. Boundaries disrupt the elegance of *ab initio* approaches, and there is the temptation to graft on a simplistic boundary correction at the end of a calculation otherwise performed in the bulk. As shown by the analysis of nanostructured silicon systems in Section 4.5, the reality is more complicated, in that the phase space adapts thermodynamically to the constrained geometry at every point in the system. Connecting the microscopic and mesoscopic pictures in such situations is a general challenge in condensed matter physics and an area where experiments offering an ever-higher combination of spatial and temporal resolution will continue to be indispensable.

Bibliography

- [1] Lawrence Livermore National Laboratory and U.S. Department of Energy, “Energy flow charts,” 2022. Accessed on Feb 20, 2024.
- [2] U.S. Energy Information Administration, “U.S. electricity flow, 2022,” 2022. Accessed on Feb 20, 2024.
- [3] National Renewable Energy Laboratory, “Best research-cell efficiencies,” 2024. Accessed on Feb 20, 2024.
- [4] M. M. Waldrop, “The chips are down for Moore’s law,” Nature News, vol. 530, no. 7589, p. 144, 2016.
- [5] S. K. Bux, R. G. Blair, P. K. Gogna, H. Lee, G. Chen, M. S. Dresselhaus, R. B. Kaner, and J.-P. Fleurial, “Nanostructured bulk silicon as an effective thermoelectric material,” Advanced Functional Materials, vol. 19, no. 15, pp. 2445–2452, 2009.
- [6] N. Li, J. Ren, L. Wang, G. Zhang, P. Hänggi, and B. Li, “Colloquium: Phononics: Manipulating heat flow with electronic analogs and beyond,” Reviews of Modern Physics, vol. 84, no. 3, pp. 1045–1066, 2012.
- [7] M. Maldovan, “Sound and heat revolutions in phononics,” Nature, vol. 503, no. 7475, pp. 209–217, 2013.
- [8] P. W. Anderson, “More is different,” Science, vol. 177, no. 4047, pp. 393–396, 1972.
- [9] R. Lontas and J. R. Greer, “3D nano-architected metallic glass: Size effect suppresses catastrophic failure,” Acta Materialia, vol. 133, pp. 393–407, 2017.
- [10] B. McBennett, A. Beardo, E. E. Nelson, B. Abad, T. D. Frazer, A. Adak, Y. Esashi, B. Li, H. C. Kapteyn, M. M. Murnane, and J. L. Knobloch, “Universal behavior of highly confined heat flow in semiconductor nanosystems: From nanomeshes to metalattices,” Nano Letters, vol. 23, no. 6, pp. 2129–2136, 2023.
- [11] D. Torrent and J. Sánchez-Dehesa, “Acoustic cloaking in two dimensions: a feasible approach,” New Journal of Physics, vol. 10, no. 6, p. 063015, 2008.
- [12] A. J. Minnich, M. S. Dresselhaus, Z. Ren, and G. Chen, “Bulk nanostructured thermoelectric materials: current research and future prospects,” Energy & Environmental Science, vol. 2, no. 5, pp. 466–479, 2009.

- [13] T. Popmintchev, M.-C. Chen, P. Arpin, M. M. Murnane, and H. C. Kapteyn, “The attosecond nonlinear optics of bright coherent X-ray generation,” Nature Photonics, vol. 4, no. 12, pp. 822–832, 2010.
- [14] A. C. Jones, S. Berweger, J. Wei, D. Cobden, and M. B. Raschke, “Nano-optical investigations of the metal- insulator phase behavior of individual VO₂ microcrystals,” Nano letters, vol. 10, no. 5, pp. 1574–1581, 2010.
- [15] H. Hirori, A. Doi, F. Blanchard, and K. Tanaka, “Single-cycle terahertz pulses with amplitudes exceeding 1 MV/cm generated by optical rectification in LiNbO₃,” Applied Physics Letters, vol. 98, no. 9, p. 091106, 2011.
- [16] J. Als-Nielsen and D. McMorrow, Elements of Modern X-ray Physics. West Sussex, United Kingdom: John Wiley & Sons Ltd., 2 ed., 2011.
- [17] D. Attwood and A. Sakdinawat, X-rays and Extreme Ultraviolet Radiation: Principles and Applications. Cambridge, United Kingdom: Cambridge University Press, 2 ed., 2016.
- [18] P. F. Moulton, “Spectroscopic and laser characteristics of Ti:Al₂O₃,” J. Opt. Soc. Am. B, vol. 3, no. 1, pp. 125–133, 1986.
- [19] D. E. Spence, P. N. Kean, and W. Sibbett, “60-fsec pulse generation from a self-mode-locked Ti:sapphire laser,” Opt. Lett., vol. 16, no. 1, pp. 42–44, 1991.
- [20] M. T. Asaki, C.-P. Huang, D. Garvey, J. Zhou, H. C. Kapteyn, and M. M. Murnane, “Generation of 11-fs pulses from a self-mode-locked Ti:sapphire laser,” Opt. Lett., vol. 18, no. 12, pp. 977–979, 1993.
- [21] J. Zhou, G. Taft, C.-P. Huang, M. M. Murnane, H. C. Kapteyn, and I. P. Christov, “Pulse evolution in a broad-bandwidth Ti:sapphire laser,” Opt. Lett., vol. 19, no. 15, pp. 1149–1151, 1994.
- [22] E. G. Lariontsev and V. N. Serkin, “Possibility of using self-focusing for increasing contrast and narrowing of ultrashort light pulses,” Soviet Journal of Quantum Electronics, vol. 5, no. 7, p. 796, 1975.
- [23] F. Salin, J. Squier, and M. Piché, “Mode locking of Ti:Al₂O₃ lasers and self-focusing: a Gaussian approximation,” Opt. Lett., vol. 16, no. 21, pp. 1674–1676, 1991.
- [24] D. Strickland and G. Mourou, “Compression of amplified chirped optical pulses,” Optics Communications, vol. 55, no. 6, pp. 447–449, 1985.
- [25] J. L. Knobloch, New Understanding of Nanoscale Thermal Transport and Mechanical Properties Uncovered Using Coherent Extreme Ultraviolet Light. PhD thesis, University of Colorado at Boulder, 2020.
- [26] S. Backus, R. Bartels, S. Thompson, R. Dollinger, H. C. Kapteyn, and M. M. Murnane, “High-efficiency, single-stage 7-kHz high-average-power ultrafast laser system,” Opt. Lett., vol. 26, no. 7, pp. 465–467, 2001.
- [27] W. Hettel, G. Golba, D. Morrill, D. Carlson, P. Chang, T.-H. Wu, S. Diddams, H. Kapteyn, M. Murnane, and M. Hemmer, “Compact, ultrastable, high repetition-rate 2 μ m and 3 μ m fiber laser for seeding mid-IR OPCPA,” Opt. Express, vol. 32, no. 3, pp. 4072–4080, 2024.

- [28] S. Mobilio, F. Boscherini, and C. Meneghini, Synchrotron Radiation: Basics, Methods, and Applications. Berlin, Germany: Springer, 2015.
- [29] E. A. Seddon, J. A. Clarke, D. J. Dunning, C. Masciovecchio, C. J. Milne, F. Parmigiani, D. Rugg, J. C. H. Spence, N. R. Thompson, K. Ueda, S. M. Vinko, J. S. Wark, and W. Wurth, “Short-wavelength free-electron laser sources and science: a review,” Reports on Progress in Physics, vol. 80, no. 11, p. 115901, 2017.
- [30] C. L. Porter, Complex Extreme Ultraviolet Imaging Reflectometry: Quantitative Lensless Imaging with Short-Wavelength Light in Reflection Geometries. PhD thesis, University of Colorado at Boulder, 2018.
- [31] S. Khan, Ultrashort Pulses from Synchrotron Radiation Sources, pp. 51–81. New York: Springer International Publishing, 2016.
- [32] E. Allaria, R. Appio, L. Badano, W. Barletta, S. Bassanese, S. Biedron, A. Borga, E. Busetto, D. Castronovo, P. Cinquegrana, et al., “Highly coherent and stable pulses from the FERMI seeded free-electron laser in the extreme ultraviolet,” Nature Photonics, vol. 6, no. 10, pp. 699–704, 2012.
- [33] P. A. Franken, A. E. Hill, C. W. Peters, and G. Weinreich, “Generation of optical harmonics,” Physical Review Letters, vol. 7, no. 4, pp. 118–119, 1961.
- [34] R. W. Boyd, Nonlinear Optics. Burlington: Academic Press, 3 ed., 2008.
- [35] C. Winterfeldt, C. Spielmann, and G. Gerber, “Colloquium: Optimal control of high-harmonic generation,” Reviews of Modern Physics, vol. 80, no. 1, pp. 117–140, 2008.
- [36] T. H. Maiman, “Stimulated optical radiation in ruby,” Nature, vol. 187, no. 4736, pp. 493–494, 1960.
- [37] G. H. C. New and J. F. Ward, “Optical third-harmonic generation in gases,” Physical Review Letters, vol. 19, no. 10, pp. 556–559, 1967.
- [38] X. Zhou, S. He, G. Liu, L. Zhao, L. Yu, and W. Zhang, “New developments in laser-based photoemission spectroscopy and its scientific applications: a key issues review,” Reports on Progress in Physics, vol. 81, no. 6, p. 062101, 2018.
- [39] A. McPherson, G. Gibson, H. Jara, U. Johann, T. S. Luk, I. McIntyre, K. Boyer, and C. K. Rhodes, “Studies of multiphoton production of vacuum-ultraviolet radiation in the rare gases,” JOSA B, vol. 4, no. 4, pp. 595–601, 1987.
- [40] M. Ferray, A. L’Huillier, X. F. Li, L. A. Lompré, G. Mainfray, and C. Manus, “Multiple-harmonic conversion of 1064 nm radiation in rare gases,” Journal of Physics B: Atomic, Molecular and Optical Physics, vol. 21, no. 3, p. L31, 1988.
- [41] B. K. Dinh, Phase-Matched High Order Harmonic Generation and Applications. PhD thesis, Swinburne University of Technology, 2012.
- [42] P. B. Corkum, “Plasma perspective on strong field multiphoton ionization,” Physical Review Letters, vol. 71, no. 13, pp. 1994–1997, 1993.

- [43] J. L. Krause, K. J. Schafer, and K. C. Kulander, “High-order harmonic generation from atoms and ions in the high intensity regime,” Physical Review Letters, vol. 68, no. 24, pp. 3535–3538, 1992.
- [44] M. Lewenstein, P. Balcou, M. Y. Ivanov, A. L’Huillier, and P. B. Corkum, “Theory of high-harmonic generation by low-frequency laser fields,” Physical Review A, vol. 49, no. 3, pp. 2117–2132, 1994.
- [45] 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 Laser Congress 2021 (ASSL,LAC), p. AW2A.7, Optica Publishing Group, 2021.
- [46] T. Popmintchev, M.-C. Chen, D. Popmintchev, P. Arpin, S. Brown, S. Ališauskas, G. Andriukaitis, T. Balčiunas, O. D. Mücke, A. Pugzlys, A. Baltuška, B. Shim, S. E. Schrauth, A. Gaeta, C. Hernández-García, L. Plaja, A. Becker, A. Jaron-Becker, M. M. Murnane, and H. C. Kapteyn, “Bright coherent ultrahigh harmonics in the keV X-ray regime from mid-infrared femtosecond lasers,” Science, vol. 336, no. 6086, pp. 1287–1291, 2012.
- [47] J. A. Armstrong, N. Bloembergen, J. Ducuing, and P. S. Pershan, “Interactions between light waves in a nonlinear dielectric,” Physical Review, vol. 127, no. 6, pp. 1918–1939, 1962.
- [48] C. G. Durfee, A. R. Rundquist, S. Backus, C. Herne, M. M. Murnane, and H. C. Kapteyn, “Phase matching of high-order harmonics in hollow waveguides,” Physical Review Letters, vol. 83, no. 11, pp. 2187–2190, 1999.
- [49] A. Rundquist, C. G. Durfee, Z. Chang, C. Herne, S. Backus, M. M. Murnane, and H. C. Kapteyn, “Phase-matched generation of coherent soft X-rays,” Science, vol. 280, no. 5368, pp. 1412–1415, 1998.
- [50] S. L. Shapiro, ed., Ultrashort Light Pulses: Picosecond Techniques and Applications, vol. 18. Berlin, Germany: Springer-Verlag, 1977.
- [51] G. M. Müller, J. Walowski, M. Djordjevic, G.-X. Miao, A. Gupta, A. V. Ramos, K. Gehrke, V. Moshnyaga, K. Samwer, J. Schmalhorst, A. Thomas, A. Hütten, G. Reiss, J. S. Moodera, and M. Münzenberg, “Spin polarization in half-metals probed by femtosecond spin excitation,” Nature Materials, vol. 8, no. 1, pp. 56–61, 2009.
- [52] J. Hohlfeld, S.-S. Wellershoff, J. Güdde, U. Conrad, V. Jähnke, and E. Matthias, “Electron and lattice dynamics following optical excitation of metals,” Chemical Physics, vol. 251, no. 1, pp. 237–258, 2000.
- [53] M. E. Siemens, Nanoscale Thermal, Acoustic, and Magnetic Dynamics Probed with Soft X-ray Light. PhD thesis, University of Colorado at Boulder, 2010.
- [54] H. J. Eichler, F. Massmann, E. Biselli, K. Richter, M. Glotz, L. Konetzke, and X. Yang, “Laser-induced free-carrier and temperature gratings in silicon,” Physical Review B, vol. 36, no. 6, pp. 3247–3253, 1987.
- [55] R. I. Tobey, M. E. Siemens, O. Cohen, M. M. Murnane, H. C. Kapteyn, and K. A. Nelson, “Ultrafast extreme ultraviolet holography: dynamic monitoring of surface deformation,” Optics letters, vol. 32, no. 3, pp. 286–288, 2007.

- [56] F. Bencivenga, R. Mincigrucci, F. Capotondi, L. Foglia, D. Naumenko, A. A. Maznev, E. Pedersoli, A. Simoncig, F. Caporaletti, V. Chiloyan, R. Cucini, F. Dallari, R. A. Duncan, T. D. Frazer, G. Gaio, A. Gessini, L. Giannessi, S. Huberman, H. Kapteyn, J. Knobloch, G. Kurdi, N. Mahne, M. Manfredda, A. Martinelli, M. Murnane, E. Principi, L. Raimondi, S. Spampinati, C. Spezzani, M. Trovò, M. Zangrando, G. Chen, G. Monaco, K. A. Nelson, and C. Masciovecchio, “Nanoscale transient gratings excited and probed by extreme ultraviolet femtosecond pulses,” *Science Advances*, vol. 5, no. 7, p. eaaw5805, 2019.
- [57] M. E. Siemens, Q. Li, M. M. Murnane, H. C. Kapteyn, R. Yang, E. H. Anderson, and K. A. Nelson, “High-frequency surface acoustic wave propagation in nanostructures characterized by coherent extreme ultraviolet beams,” *Applied Physics Letters*, vol. 94, no. 9, p. 093103, 2009.
- [58] M. E. Siemens, Q. Li, R. Yang, K. A. Nelson, E. H. Anderson, M. M. Murnane, and H. C. Kapteyn, “Quasi-ballistic thermal transport from nanoscale interfaces observed using ultrafast coherent soft X-ray beams,” *Nature Materials*, vol. 9, no. 1, pp. 26–30, 2010.
- [59] D. Nardi, M. Travaglini, M. E. Siemens, Q. Li, M. M. Murnane, H. C. Kapteyn, G. Ferrini, F. Parmigiani, and F. Banfi, “Probing thermomechanics at the nanoscale: impulsively excited pseudosurface acoustic waves in hypersonic phononic crystals,” *Nano Letters*, vol. 11, no. 10, pp. 4126–4133, 2011.
- [60] Q. Li, K. Hoogetboom-Pot, D. Nardi, M. M. Murnane, H. C. Kapteyn, M. E. Siemens, E. H. Anderson, O. Hellwig, E. Dobisz, B. Gurney, R. Yang, and K. A. Nelson, “Generation and control of ultrashort-wavelength two-dimensional surface acoustic waves at nanoscale interfaces,” *Physical Review B*, vol. 85, no. 19, p. 195431, 2012.
- [61] K. M. Hoogetboom-Pot, J. N. Hernández-Charpak, X. Gu, T. D. Frazer, E. H. Anderson, W. Chao, R. W. Falcone, R. Yang, M. M. Murnane, H. C. Kapteyn, and D. Nardi, “A new regime of nanoscale thermal transport: collective diffusion increases dissipation efficiency,” *Proceedings of the National Academy of Sciences*, vol. 112, no. 16, pp. 4846–4851, 2015.
- [62] K. M. Hoogetboom-Pot, E. Turgut, J. N. Hernández-Charpak, J. M. Shaw, H. C. Kapteyn, M. M. Murnane, and D. Nardi, “Nondestructive measurement of the evolution of layer-specific mechanical properties in sub-10 nm bilayer films,” *Nano Letters*, vol. 16, no. 8, pp. 4773–4778, 2016.
- [63] J. N. Hernández-Charpak, K. M. Hoogetboom-Pot, Q. Li, T. D. Frazer, J. L. Knobloch, M. Tripp, S. W. King, E. H. Anderson, W. Chao, M. M. Murnane, H. C. Kapteyn, and D. Nardi, “Full characterization of the mechanical properties of 11-50 nm ultrathin films: influence of network connectivity on the Poisson’s ratio,” *Nano Letters*, vol. 17, no. 4, pp. 2178–2183, 2017.
- [64] T. D. Frazer, J. L. Knobloch, K. M. Hoogetboom-Pot, D. Nardi, W. Chao, R. W. Falcone, M. M. Murnane, H. C. Kapteyn, and J. N. Hernández-Charpak, “Engineering nanoscale thermal transport: size- and spacing-dependent cooling of nanostructures,” *Physical Review Applied*, vol. 11, no. 2, p. 024042, 2019.
- [65] T. D. Frazer, J. L. Knobloch, J. N. Hernández-Charpak, K. M. Hoogetboom-Pot, D. Nardi, S. Yazdi, W. Chao, E. H. Anderson, M. K. Tripp, S. W. King, H. C. Kapteyn, M. M. Murnane, and B. Abad, “Full characterization of ultrathin 5-nm low- k dielectric bilayers: influence of dopants and surfaces on the mechanical properties,” *Physical Review Materials*, vol. 4, no. 7, p. 073603, 2020.

- [66] B. Abad, J. L. Knobloch, T. D. Frazer, J. N. Hernández-Charpak, H. Y. Cheng, A. J. Grede, N. C. Giebink, T. E. Mallouk, P. Mahale, N. N. Nova, A. A. Tomaschke, V. L. Ferguson, V. H. Crespi, V. Gopalan, H. C. Kapteyn, J. V. Badding, and M. M. Murnane, “Nondestructive measurements of the mechanical and structural properties of nanostructured metalattices,” *Nano Letters*, vol. 20, no. 5, pp. 3306–3312, 2020.
- [67] A. Beardo, J. L. Knobloch, L. Sendra, J. Bafaluy, T. D. Frazer, W. Chao, J. N. Hernández-Charpak, H. C. Kapteyn, B. Abad, M. M. Murnane, F. X. Alvarez, and J. Camacho, “A general and predictive understanding of thermal transport from 1D-and 2D-confined nanostructures: theory and experiment,” *ACS Nano*, vol. 15, no. 8, pp. 13019–13030, 2021.
- [68] J. L. Knobloch, B. McBennett, C. S. Bevis, S. Yazdi, T. D. Frazer, A. Adak, E. E. Nelson, J. N. Hernández-Charpak, H. Y. Cheng, A. J. Grede, P. Mahale, N. N. Nova, N. C. Giebink, T. E. Mallouk, J. V. Badding, H. C. Kapteyn, B. Abad, and M. M. Murnane, “Structural and elastic properties of empty-pore metalattices extracted via nondestructive coherent extreme UV scatterometry and electron tomography,” *ACS Applied Materials & Interfaces*, vol. 14, no. 36, p. 41316–41327, 2022.
- [69] N. W. Ashcroft and N. D. Mermin, *Solid State Physics*. Holt, Rinehart, and Winston, 1976.
- [70] J. M. Ziman, *Electrons and Phonons: The Theory of Transport Phenomena in Solids*. Oxford, United Kingdom: Oxford University Press, 1962.
- [71] R. E. Peierls, *Quantum Theory of Solids*. Oxford, United Kingdom: Oxford University Press, 1955.
- [72] Y. Machida, A. Subedi, K. Akiba, A. Miyake, M. Tokunaga, Y. Akahama, K. Izawa, and K. Behnia, “Observation of Poiseuille flow of phonons in black phosphorus,” *Science Advances*, vol. 4, no. 6, p. eaat3374, 2018.
- [73] J. Sussmann and A. Thellung, “Thermal conductivity of perfect dielectric crystals in the absence of umklapp processes,” *Proceedings of the Physical Society*, vol. 81, no. 6, p. 1122, 1963.
- [74] R. J. Hardy, “Phonon Boltzmann equation and second sound in solids,” *Physical Review B*, vol. 2, no. 4, p. 1193, 1970.
- [75] A. Ward, D. A. Broido, D. A. Stewart, and G. Deinzer, “Ab initio theory of the lattice thermal conductivity in diamond,” *Physical Review B*, vol. 80, no. 12, p. 125203, 2009.
- [76] D. A. Broido, M. Malorny, G. Birner, N. Mingo, and D. A. Stewart, “Intrinsic lattice thermal conductivity of semiconductors from first principles,” *Applied Physics Letters*, vol. 91, no. 23, p. 231922, 2007.
- [77] R. Peierls, “Zur kinetischen Theorie der Wärmeleitung in Kristallen,” *Annalen der Physik*, vol. 395, no. 8, pp. 1055–1101, 1929.
- [78] G. Fugallo, M. Lazzeri, L. Paulatto, and F. Mauri, “Ab initio variational approach for evaluating lattice thermal conductivity,” *Physical Review B*, vol. 88, no. 4, p. 045430, 2013.
- [79] A. Majumdar, “Microscale heat conduction in dielectric thin films,” *Journal of Heat Transfer*, vol. 115, no. 1, pp. 7–16, 1993.

- [80] K. Esfarjani, G. Chen, and H. T. Stokes, “Heat transport in silicon from first-principles calculations,” Physical Review B, vol. 84, no. 8, p. 085204, 2011.
- [81] A. Jain, Y.-J. Yu, and A. J. H. McGaughey, “Phonon transport in periodic silicon nanoporous films with feature sizes greater than 100 nm,” Physical Review B, vol. 87, no. 19, p. 195301, 2013.
- [82] K. Takahashi, M. Fujikane, Y. Liao, M. Kashiwagi, T. Kawasaki, N. Tambo, S. Ju, Y. Naito, and J. Shiomi, “Elastic inhomogeneity and anomalous thermal transport in ultrafine Si phononic crystals,” Nano Energy, vol. 71, p. 104581, 2020.
- [83] Q. Hao, G. Chen, and M.-S. Jeng, “Frequency-dependent Monte Carlo simulations of phonon transport in two-dimensional porous silicon with aligned pores,” Journal of Applied Physics, vol. 106, no. 11, p. 114321, 2009.
- [84] J.-P. M. Péraud and N. G. Hadjiconstantinou, “Efficient simulation of multidimensional phonon transport using energy-based variance-reduced Monte Carlo formulations,” Physical Review B, vol. 84, no. 20, p. 205331, 2011.
- [85] M. Nomura, R. Anufriev, Z. Zhang, J. Maire, Y. Guo, R. Yanagisawa, and S. Volz, “Review of thermal transport in phononic crystals,” Materials Today Physics, vol. 22, p. 100613, 2022.
- [86] R. Anufriev, A. Ramiere, J. Maire, and M. Nomura, “Heat guiding and focusing using ballistic phonon transport in phononic nanostructures,” Nature Communications, vol. 8, no. 1, p. 15505, 2017.
- [87] B. L. Davis and M. I. Hussein, “Nanophononic metamaterial: thermal conductivity reduction by local resonance,” Physical Review Letters, vol. 112, no. 5, p. 055505, 2014.
- [88] A. J. H. McGaughey and M. Kaviani, “Quantitative validation of the Boltzmann transport equation phonon thermal conductivity model under the single-mode relaxation time approximation,” Physical Review B, vol. 69, no. 9, p. 094303, 2004.
- [89] G. Fugallo, A. Cepellotti, L. Paulatto, M. Lazzeri, N. Marzari, and F. Mauri, “Thermal conductivity of graphene and graphite: collective excitations and mean free paths,” Nano Letters, vol. 14, no. 11, pp. 6109–6114, 2014.
- [90] G. Chen, “Nonlocal and nonequilibrium heat conduction in the vicinity of nanoparticles,” Journal of Heat Transfer, vol. 118, no. 3, pp. 539–545, 1996.
- [91] Y. S. Ju and K. E. Goodson, “Phonon scattering in silicon films with thickness of order 100 nm,” Applied Physics Letters, vol. 74, no. 20, pp. 3005–3007, 1999.
- [92] J. Cuffe, J. K. Eliason, A. A. Maznev, K. C. Collins, J. A. Johnson, A. Shchepetov, M. Prunnila, J. Ahopelto, C. M. Sotomayor Torres, G. Chen, and K. A. Nelson, “Reconstructing phonon mean-free-path contributions to thermal conductivity using nanoscale membranes,” Physical Review B, vol. 91, no. 24, p. 245423, 2015.
- [93] A. J. Minnich, J. A. Johnson, A. J. Schmidt, K. Esfarjani, M. S. Dresselhaus, K. A. Nelson, and G. Chen, “Thermal conductivity spectroscopy technique to measure phonon mean free paths,” Physical Review Letters, vol. 107, no. 9, p. 095901, 2011.

- [94] C. Hua and L. Lindsay, “Space-time dependent thermal conductivity in nonlocal thermal transport,” Physical Review B, vol. 102, no. 10, p. 104310, 2020.
- [95] V. Chiloyan, L. Zeng, S. Huberman, A. A. Maznev, K. A. Nelson, and G. Chen, “Variational approach to solving the spectral Boltzmann transport equation in transient thermal grating for thin films,” Journal of Applied Physics, vol. 120, no. 2, p. 025103, 2016.
- [96] C. de Tomas, A. Cantarero, A. F. Lopeandia, and F. X. Alvarez, “From kinetic to collective behavior in thermal transport on semiconductors and semiconductor nanostructures,” Journal of Applied Physics, vol. 115, no. 16, p. 164314, 2014.
- [97] Y. Guo and M. Wang, “Phonon hydrodynamics and its applications in nanoscale heat transport,” Physics Reports, vol. 595, pp. 1–44, 2015.
- [98] L. Sendra, A. Beardo, P. Torres, J. Bafaluy, F. X. Alvarez, and J. Camacho, “Derivation of a hydrodynamic heat equation from the phonon Boltzmann equation for general semiconductors,” Physical Review B, vol. 103, no. 14, p. L140301, 2021.
- [99] H. Grad, “On the kinetic theory of rarefied gases,” Communications on Pure and Applied Mathematics, vol. 2, no. 4, pp. 331–407, 1949.
- [100] H. Struchtrup and M. Torrilhon, “Regularization of Grad’s 13 moment equations: derivation and linear analysis,” Physics of Fluids, vol. 15, no. 9, pp. 2668–2680, 2003.
- [101] Y. Guo, D. Jou, and M. Wang, “Nonequilibrium thermodynamics of phonon hydrodynamic model for nanoscale heat transport,” Physical Review B, vol. 98, no. 10, p. 104304, 2018.
- [102] Y. Guo and M. Wang, “Phonon hydrodynamics for nanoscale heat transport at ordinary temperatures,” Physical Review B, vol. 97, no. 3, p. 035421, 2018.
- [103] S. R. D. Groot and P. Mazur, Non-Equilibrium Thermodynamics. New York: Dover, 1962.
- [104] Y. Zhou, X. Zhang, and M. Hu, “Nonmonotonic diameter dependence of thermal conductivity of extremely thin Si nanowires: competition between hydrodynamic phonon flow and boundary scattering,” Nano Letters, vol. 2, no. 17, pp. 1269–1276, 2017.
- [105] C. Melis, R. Rurali, X. Cartoixà, and F. X. Alvarez, “Indications of phonon hydrodynamics in telescopic silicon nanowires,” Physical Review Applied, vol. 11, no. 5, p. 054059, 2019.
- [106] M. Verdier, Y. Han, D. Lacroix, P.-O. Chapuis, and K. Termentzidis, “Radial dependence of thermal transport in silicon nanowires,” Journal of Physics: Materials, vol. 2, no. 1, p. 015002, 2018.
- [107] P. Desmarchelier, A. Beardo, F. X. Alvarez, A. Tanguy, and K. Termentzidis, “Atomistic evidence of hydrodynamic heat transfer in nanowires,” International Journal of Heat and Mass Transfer, vol. 194, p. 123003, 2022.
- [108] R. A. Guyer and J. A. Krumhansl, “Solution of the linearized phonon Boltzmann equation,” Physical Review, vol. 148, no. 2, pp. 766–778, 1966.
- [109] R. A. Guyer and J. A. Krumhansl, “Thermal conductivity, second sound, and phonon hydrodynamic phenomena in nonmetallic crystals,” Physical Review, vol. 148, no. 2, pp. 778–788, 1966.

- [110] A. Cepellotti, G. Fugallo, L. Paulatto, M. Lazzeri, F. Mauri, and N. Marzari, “Phonon hydrodynamics in two-dimensional materials,” Nature Communications, vol. 6, no. 1, p. 6400, 2015.
- [111] S. Lee, D. Broido, K. Esfarjani, and G. Chen, “Hydrodynamic phonon transport in suspended graphene,” Nature Communications, vol. 6, no. 1, p. 6290, 2015.
- [112] S. Huberman, R. A. Duncan, K. Chen, B. Song, V. Chiloyan, Z. Ding, A. A. Maznev, G. Chen, and K. A. Nelson, “Observation of second sound in graphite at temperatures above 100 K,” Science, vol. 364, no. 6438, pp. 375–379, 2019.
- [113] A. Beardo, M. López-Suárez, L. A. Pérez, L. Sendra, M. I. Alonso, C. Melis, J. Bafaluy, J. Camacho, L. Colombo, R. Rurali, F. X. Alvarez, and J. S. Reparaz, “Observation of second sound in a rapidly varying temperature field in Ge,” Science Advances, vol. 7, no. 27, p. eabg4677, 2021.
- [114] N. Zen, T. A. Puurtinen, T. J. Isotalo, S. Chaudhuri, and I. J. Maasilta, “Engineering thermal conductance using a two-dimensional phononic crystal,” Nature Communications, vol. 5, no. 1, p. 3435, 2014.
- [115] J. Maire, R. Anufriev, R. Yanagisawa, A. Ramiere, S. Volz, and M. Nomura, “Heat conduction tuning by wave nature of phonons,” Science Advances, vol. 3, no. 8, p. e1700027, 2017.
- [116] P. K. Schelling, S. R. Phillpot, and P. Keblinski, “Comparison of atomic-level simulation methods for computing thermal conductivity,” Physical Review B, vol. 65, no. 14, p. 144306, 2002.
- [117] M. I. Hussein, C.-N. Tsai, and H. Honarvar, “Nanophononics: thermal conductivity reduction in a nanophononic metamaterial versus a nanophononic crystal: a review and comparative analysis,” Advanced Functional Materials, vol. 30, no. 8, p. 2070047, 2020.
- [118] P. B. Allen and J. L. Feldman, “Thermal conductivity of glasses: Theory and application to amorphous Si,” Physical Review Letters, vol. 62, no. 6, pp. 645–648, 1989.
- [119] P. B. Allen, J. L. Feldman, J. Fabian, and F. Wooten, “Diffusons, locons and propagons: character of atomic vibrations in amorphous Si,” Philosophical Magazine B, vol. 79, no. 11-12, pp. 1715–1731, 1999.
- [120] M. Simoncelli, N. Marzari, and F. Mauri, “Unified theory of thermal transport in crystals and glasses,” Nature Physics, vol. 15, no. 8, pp. 809–813, 2019.
- [121] H. Honarvar, J. L. Knobloch, T. D. Frazer, B. Abad, B. McBennett, M. I. Hussein, H. C. Kapteyn, M. M. Murnane, and J. N. Hernández-Charpak, “Directional thermal channeling: a phenomenon triggered by tight packing of heat sources,” Proceedings of the National Academy of Sciences, vol. 118, no. 40, p. e2109056118, 2021.
- [122] J. A. Thomas, J. E. Turney, R. M. Iutzi, C. H. Amon, and A. J. H. McGaughey, “Predicting phonon dispersion relations and lifetimes from the spectral energy density,” Physical Review B, vol. 81, no. 8, p. 081411, 2010.
- [123] H. Casimir, “Note on the conduction of heat in crystals,” Physica, vol. 5, no. 6, pp. 495–500, 1938.

- [124] W. Chen, D. Talreja, D. Eichfeld, P. Mahale, N. N. Nova, H. Y. Cheng, J. L. Russell, S.-Y. Yu, N. Poilvert, G. Mahan, S. E. Mohney, V. H. Crespi, T. E. Mallouk, J. V. Badding, B. Foley, V. Gopalan, and I. Dabo, "Achieving minimal heat conductivity by ballistic confinement in phononic metalattices," ACS Nano, vol. 14, no. 4, pp. 4235–4243, 2020.
- [125] J. Ravichandran, A. K. Yadav, R. Cheaito, P. B. Rossen, A. Soukiassian, S. J. Suresha, J. C. Duda, B. M. Foley, C.-H. Lee, Y. Zhu, A. W. Lichtenberger, J. E. Moore, D. A. Muller, D. G. Schlom, P. E. Hopkins, A. Majumdar, R. Ramesh, and M. A. Zurbuchen, "Crossover from incoherent to coherent phonon scattering in epitaxial oxide superlattices," Nature Materials, vol. 13, no. 2, pp. 168–172, 2014.
- [126] J. K. Yu, S. Mitrovic, D. Tham, J. Varghese, and J. R. Heath, "Reduction of thermal conductivity in phononic nanomesh structures," Nature Nanotechnology, vol. 5, no. 10, pp. 718–721, 2010.
- [127] M. N. Luckyanova, J. Garg, K. Esfarjani, A. Jandl, M. T. Bultsara, A. J. Schmidt, A. J. Minnich, S. Chen, M. S. Dresselhaus, Z. Ren, E. A. Fitzgerald, and G. Chen, "Coherent phonon heat conduction in superlattices," Science, vol. 338, no. 6109, pp. 936–939, 2012.
- [128] J. Wang, J. K. Carson, M. F. North, and D. J. Cleland, "A new approach to modelling the effective thermal conductivity of heterogeneous materials," International Journal of Heat and Mass Transfer, vol. 49, no. 17, pp. 3075–3083, 2006.
- [129] H. W. Russell, "Principles of heat flow in porous insulators," Journal of the American Ceramic Society, vol. 18, no. 1-12, pp. 1–5, 1935.
- [130] Y. Zhao, L. Yang, L. Kong, M. H. Nai, D. Liu, J. Wu, Y. Liu, S. Y. Chiam, W. K. Chim, C. T. Lim, B. Li, J. T. L. Thong, and K. Hippalgaonkar, "Ultralow thermal conductivity of single-crystalline porous silicon nanowires," Advanced Functional Materials, vol. 27, no. 40, p. 1702824, 2017.
- [131] A. Van Blaaderen, J. Van Geest, and A. Vrij, "Monodisperse colloidal silica spheres from tetraalkoxysilanes: particle formation and growth mechanism," Journal of Colloid and Interface Science, vol. 154, no. 2, pp. 481–501, 1992.
- [132] K. D. Hartlen, A. P. Athanasopoulos, and V. Kitaev, "Facile preparation of highly monodisperse small silica spheres (15 to >200 nm) suitable for colloidal templating and formation of ordered arrays," Langmuir, vol. 24, no. 5, pp. 1714–1720, 2008.
- [133] R. Watanabe, T. Yokoi, E. Kobayashi, Y. Otsuka, A. Shimojima, T. Okubo, and T. Tatsumi, "Extension of size of monodisperse silica nanospheres and their well-ordered assembly," Journal of Colloid and Interface Science, vol. 360, no. 1, pp. 1–7, 2011.
- [134] S.-L. Kuai, X.-F. Hu, A. Haché, and V.-V. Truong, "High-quality colloidal photonic crystals obtained by optimizing growth parameters in a vertical deposition technique," Journal of Crystal Growth, vol. 267, no. 1-2, pp. 317–324, 2004.
- [135] J. L. Russell, G. H. Noel, J. M. Warren, N.-L. L. Tran, and T. E. Mallouk, "Binary colloidal crystal films grown by vertical evaporation of silica nanoparticle suspensions," Langmuir, vol. 33, no. 39, pp. 10366–10373, 2017.

- [136] N. F. Baril, B. Keshavarzi, J. R. Sparks, M. Krishnamurthi, I. Temnykh, P. J. Sazio, A. C. Peacock, A. Borhan, V. Gopalan, and J. V. Badding, "High-pressure chemical deposition for void-free filling of extreme aspect ratio templates," Advanced Materials, vol. 22, no. 41, pp. 4605–4611, 2010.
- [137] A. Pryor Jr., Y. Yang, A. Rana, M. Gallagher-Jones, J. Zhou, Y. H. Lo, G. Melinte, W. Chiu, J. A. Rodriguez, and J. Miao, "GENFIRE: a generalized Fourier iterative reconstruction algorithm for high-resolution 3D imaging," Scientific Reports, vol. 7, no. 1, p. 10409, 2017.
- [138] M. Thommes, K. Kaneko, A. V. Neimark, J. P. Olivier, F. Rodriguez-Reinoso, J. Rouquerol, and K. S. Sing, "Physisorption of gases, with special reference to the evaluation of surface area and pore size distribution (IUPAC Technical Report)," Pure and Applied Chemistry, vol. 87, no. 9-10, pp. 1051–1069, 2015.
- [139] D. J. Taylor, P. F. Fleig, and S. L. Hietala, "Technique for characterization of thin film porosity," Thin Solid Films, vol. 332, no. 1-2, pp. 257–261, 1998.
- [140] Y. K. Hwang, K. R. Patil, S. H. Jhung, J.-S. Chang, Y. J. Ko, and S.-E. Park, "Control of pore size and condensation rate of cubic mesoporous silica thin films using a swelling agent," Microporous and Mesoporous Materials, vol. 78, no. 2-3, pp. 245–253, 2005.
- [141] S. Dourdain, A. Mehdi, J. F. Bardeau, and A. Gibaud, "Determination of porosity of mesoporous silica thin films by quantitative X-ray reflectivity analysis and GISAXS," Thin Solid Films, vol. 495, no. 1-2, pp. 205–209, 2006.
- [142] E. Ortel, A. Hertwig, D. Berger, P. Esposito, A. M. Rossi, R. Kraehnert, and V.-D. Hodoroba, "New approach on quantification of porosity of thin films via electron-excited X-ray spectra," Analytical Chemistry, vol. 88, no. 14, pp. 7083–7090, 2016.
- [143] T. Galy, M. Marszewski, S. King, Y. Yan, S. H. Tolbert, and L. Pilon, "Comparing methods for measuring thickness, refractive index, and porosity of mesoporous thin films," Microporous and Mesoporous Materials, vol. 291, p. 109677, 2020.
- [144] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 12, no. 7, pp. 629–639, 1990.
- [145] A. P. Sheppard, R. M. Sok, and H. Averdunk, "Techniques for image enhancement and segmentation of tomographic images of porous materials," Physica A: Statistical Mechanics and its Applications, vol. 339, no. 1-2, pp. 145–151, 2004.
- [146] N. Otsu, "A threshold selection method from gray-level histograms," IEEE Transactions on Systems, Man, and Cybernetics, vol. 9, no. 1, pp. 62–66, 1979.
- [147] D. Bradley and G. Roth, "Adaptive thresholding using the integral image," Journal of Graphics Tools, vol. 12, no. 2, pp. 13–21, 2007.
- [148] D. Arthur and S. Vassilvitskii, "K-means++: the advantages of careful seeding," in SODA '07: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1027–1035, 2007.
- [149] B. A. Auld, Acoustic Fields and Waves in Solids: Volume 1. Malabar, Florida: Krieger Publishing Company, 2 ed., 1973.

- [150] B. A. Auld, Acoustic Fields and Waves in Solids: Volume 2. Malabar, Florida: Krieger Publishing Company, 2 ed., 1973.
- [151] J. A. Rogers, A. A. Maznev, M. J. Banet, and K. A. Nelson, “Optical generation and characterization of acoustic waves in thin films: fundamentals and applications,” Annual Review of Materials Science, vol. 30, no. 1, pp. 117–157, 2000.
- [152] A. A. Maznev and A. G. Every, “Surface acoustic waves in a periodically patterned layered structure,” Journal of Applied Physics, vol. 106, no. 11, p. 113531, 2009.
- [153] D. Nardi, F. Banfi, C. Giannetti, B. Revaz, G. Ferrini, and F. Parmigiani, “Pseudosurface acoustic waves in hypersonic surface phononic crystals,” Physical Review B, vol. 80, no. 10, p. 104119, 2009.
- [154] M. A. Hopcroft, W. D. Nix, and T. W. Kenny, “What is the Young’s modulus of silicon?,” Journal of Microelectromechanical Systems, vol. 19, no. 2, pp. 229–238, 2010.
- [155] F. Dalla Longa, J. T. Kohlhepp, W. J. M. de Jonge, and B. Koopmans, “Influence of photon angular momentum on ultrafast demagnetization in nickel,” Physical Review B, vol. 75, no. 22, p. 224431, 2007.
- [156] P. Kapitza, “The study of heat transfer in helium II,” Zh. Eksp. Teor. Fiz., vol. 4, p. 181, 1941.
- [157] E. T. Swartz and R. O. Pohl, “Thermal boundary resistance,” Reviews of Modern Physics, vol. 61, no. 3, pp. 605–668, 1989.
- [158] A. Giri and P. E. Hopkins, “A review of experimental and computational advances in thermal boundary conductance and nanoscale thermal transport across solid interfaces,” Advanced Functional Materials, vol. 30, no. 8, p. 1903857, 2020.
- [159] D. R. Lide, CRC Handbook of Chemistry and Physics. Boca Raton, Florida: CRC Press, 85 ed., 2005.
- [160] T. G. Kollie, “Measurement of the thermal-expansion coefficient of nickel from 300 to 1000 K and determination of the power-law constants near the Curie temperature,” Physical Review B, vol. 16, no. 11, pp. 4872–4881, 1977.
- [161] Y. Okada and Y. Tokumaru, “Precise determination of lattice parameter and thermal expansion coefficient of silicon between 300 and 1500 K,” Journal of Applied Physics, vol. 56, no. 2, pp. 314–320, 1984.
- [162] J. W. Goodman, Introduction to Fourier Optics. Englewood, Colorado: Roberts & Company, 3 ed., 2005.
- [163] L. G. Parratt, “Surface studies of solids by total reflection of X-rays,” Physical Review, vol. 95, no. 2, pp. 359–369, 1954.
- [164] D. Bellet, P. Lamagnère, A. Vincent, and Y. Bréchet, “Nanoindentation investigation of the Young’s modulus of porous silicon,” Journal of Applied Physics, vol. 80, no. 7, pp. 3772–3776, 1996.

- [165] S. Fakiri, A. Montagne, K. Rahmoun, A. Iost, and K. Ziouche, “Mechanical properties of porous silicon and oxidized porous silicon by nanoindentation technique,” Materials Science and Engineering: A, vol. 711, pp. 470–475, 2018.
- [166] L. J. Gibson and M. F. Ashby, “The mechanics of three-dimensional cellular materials,” Proceedings of the Royal Society of London. A, vol. 382, no. 1782, pp. 43–59, 1982.
- [167] Z. Hashin and S. Shtrikman, “A variational approach to the theory of the elastic behaviour of multiphase materials,” Journal of the Mechanics and Physics of Solids, vol. 11, no. 2, pp. 127–140, 1963.
- [168] C. Crook, J. Bauer, A. Guell Izard, C. Santos de Oliveira, J. Martins de Souza e Silva, J. B. Berger, and L. Valdevit, “Plate-nanolattices at the theoretical limit of stiffness and strength,” Nature Communications, vol. 11, no. 1, p. 1579, 2020.
- [169] P. R. Bevington and D. K. Robinson, Data Reduction and Error Analysis for the Physical Sciences. New York: McGraw-Hill, 3 ed., 2003.
- [170] R. Cheaito, J. T. Gaskins, M. E. Caplan, B. F. Donovan, B. M. Foley, A. Giri, J. C. Duda, C. J. Szejewski, C. Constantin, H. J. Brown-Shaklee, J. F. Ihlefeld, and P. E. Hopkins, “Thermal boundary conductance accumulation and interfacial phonon transmission: measurements and theory,” Physical Review B, vol. 91, no. 3, p. 035432, 2015.
- [171] A. Ruiz-Clavijo, O. Caballero-Calero, C. V. Manzano, X. Maeder, A. Beardo, X. Cartoixá, F. X. Álvarez, and M. Martín-González, “3D Bi₂Te₃ interconnected nanowire networks to increase thermoelectric efficiency,” ACS Applied Energy Materials, vol. 4, no. 12, p. 13556, 2021.
- [172] Z. Xiang, P. Jiang, and R. Yang, “Time-domain thermoreflectance (TDTR) data analysis using phonon hydrodynamic model,” Journal of Applied Physics, vol. 132, no. 20, p. 205104, 2022.
- [173] A. Beardo, M. Calvo-Schwarzwälder, J. Camacho, T. G. Myers, P. Torres, L. Sendra, F. X. Alvarez, and J. Bafaluy, “Hydrodynamic heat transport in compact and holey silicon thin films,” Physical Review Applied, vol. 11, no. 3, p. 034003, 2019.
- [174] F. X. Alvarez, D. Jou, and A. Sellitto, “Pore-size dependence of the thermal conductivity of porous silicon: a phonon hydrodynamic approach,” Applied Physics Letters, vol. 97, no. 3, p. 033103, 2010.
- [175] P. C. Carman, Flow of Gases through Porous Media. New York: Academic Press, 1956.
- [176] L. J. Klinkenberg, “The permeability of porous media to liquids and gases,” Drilling and Production Practice, vol. 41, p. 200, 1941.
- [177] J. Lee, W. Lee, G. Wehmeyer, S. Dhuey, D. L. Olynick, S. Cabrini, C. Dames, J. J. Urban, and P. Yang, “Investigation of phonon coherence and backscattering using silicon nanomeshes,” Nature Communications, vol. 8, no. 1, p. 14054, 2017.
- [178] P. E. Hopkins, C. M. Reinke, M. F. Su, R. H. Olsson, E. A. Shaner, Z. C. Leseman, J. R. Serrano, L. M. Phinney, and I. El-Kady, “Reduction in the thermal conductivity of single crystalline silicon by phononic crystal patterning,” Nano Letters, vol. 11, no. 1, pp. 107–112, 2011.

- [179] L. Yang, N. Yang, and B. Li, “Extreme low thermal conductivity in nanoscale 3D Si phononic crystal with spherical pores,” Nano Letters, vol. 14, no. 4, pp. 1734–1738, 2014.
- [180] M. Verdier, D. Lacroix, and K. Termentzidis, “Thermal transport in two- and three-dimensional nanowire networks,” Physical Review B, vol. 98, no. 15, p. 155434, 2018.
- [181] X. Zianni, V. Jean, K. Termentzidis, and D. Lacroix, “Scaling behavior of the thermal conductivity of width-modulated nanowires and nanofilms for heat transfer control at the nanoscale,” Nanotechnology, vol. 25, no. 46, p. 465402, 2014.
- [182] H. Wei, H. Bao, and X. Ruan, “Machine learning prediction of thermal transport in porous media with physics-based descriptors,” International Journal of Heat and Mass Transfer, vol. 160, p. 120176, 2020.
- [183] K. R. Hahn, C. Melis, and L. Colombo, “Thermal conduction and rectification phenomena in nanoporous silicon membranes,” Physical Chemistry Chemical Physics, vol. 24, no. 22, pp. 13625–13632, 2022.
- [184] J. Lim, H.-T. Wang, J. Tang, S. C. Andrews, H. So, J. Lee, D. H. Lee, T. P. Russell, and P. Yang, “Simultaneous thermoelectric property measurement and incoherent phonon transport in holey silicon,” ACS Nano, vol. 10, no. 1, pp. 124–132, 2016.
- [185] R. Anufriev, J. Maire, and M. Nomura, “Reduction of thermal conductivity by surface scattering of phonons in periodic silicon nanostructures,” Physical Review B, vol. 93, no. 4, p. 045411, 2016.
- [186] P. Ferrando-Villalba, L. D’Ortenzi, G. G. Dalkiranis, E. Cara, A. F. Lopeandia, Ll Abad, R. Rurali, X. Cartoixà, N. De Leo, Z. Saghi, M. Jacob, N. Gambacorti, L. Boarino, and J. Rodríguez-Viejo, “Impact of pore anisotropy on the thermal conductivity of porous Si nanowires,” Scientific Reports, vol. 8, no. 1, p. 12796, 2018.
- [187] P. Martin, Z. Aksamija, E. Pop, and U. Ravaioli, “Impact of phonon-surface roughness scattering on thermal conductivity of thin Si nanowires,” Physical Review Letters, vol. 102, no. 12, p. 125503, 2009.
- [188] R. P. Chhabra, Non-Newtonian Fluids: An Introduction, pp. 3–34. New York: Springer, 2010.
- [189] X. Cartoixà, R. Dettori, C. Melis, L. Colombo, and R. Rurali, “Thermal transport in porous Si nanowires from approach-to-equilibrium molecular dynamics calculations,” Applied Physics Letters, vol. 109, no. 1, p. 013107, 2016.
- [190] S. A. Hosseini, A. Davies, I. Dickey, N. Neophytou, P. A. Greaney, and L. de Sousa Oliveira, “Super-suppression of long phonon mean-free-paths in nano-engineered Si due to heat current anticorrelations,” Materials Today Physics, vol. 27, p. 100719, 2022.
- [191] A. Aharon-Steinberg, T. Völkl, A. Kaplan, A. K. Pariari, I. Roy, T. Holder, Y. Wolf, A. Y. Meltzer, Y. Myasoedov, M. E. Huber, B. Yan, G. Falkovich, L. S. Levitov, M. Hücker, and E. Zeldov, “Direct observation of vortices in an electron fluid,” Nature, vol. 607, no. 7917, p. 74–80, 2022.

- [192] S. Yasin, D. G. Hasko, and H. Ahmed, “Comparison of MIBK/IPA and water/IPA as PMMA developers for electron beam nanolithography,” Microelectronic Engineering, vol. 61-62, pp. 745–753, 2002.
- [193] W. Hu, K. Sarveswaran, M. Lieberman, and G. H. Bernstein, “Sub-10 nm electron beam lithography using cold development of poly (methymethacrylate),” Journal of Vacuum Science & Technology B, vol. 22, no. 4, pp. 1711–1716, 2004.
- [194] I. Zailer, J. E. F. Frost, V. Chabasseur-Molyneux, C. J. B. Ford, and M. Pepper, “Crosslinked PMMA as a high-resolution negative resist for electron beam lithography and applications for physics of low-dimensional structures,” Semiconductor Science and Technology, vol. 11, no. 8, p. 1235, 1996.
- [195] D. Araujo, M. Suzuki, F. Lloret, G. Alba, and P. Villar, “Diamond for electronics: materials, processing and devices,” Materials, vol. 14, no. 22, p. 7081, 2021.
- [196] M. Tordjman, “Diamond electronics with high carrier mobilities,” Nature Electronics, vol. 5, no. 1, pp. 21–22, 2022.
- [197] T. Schröder, S. L. Mouradian, J. Zheng, M. E. Trusheim, M. Walsh, E. H. Chen, L. Li, I. Bayn, and D. Englund, “Quantum nanophotonics in diamond,” Journal of the Optical Society of America B, vol. 33, no. 4, pp. B65–B83, 2016.
- [198] R. Schirhagl, K. Chang, M. Loretz, , and C. L. Degen, “Nitrogen-vacancy centers in diamond: nanoscale sensors for physics and biology,” Annual Review of Physical Chemistry, vol. 65, p. 83–105, 2013.
- [199] W. Li, N. Mingo, L. Lindsay, D. A. Broido, D. A. Stewart, and N. A. Katcho, “Thermal conductivity of diamond nanowires from first principles,” Physical Review B, vol. 85, no. 19, p. 195436, 2012.
- [200] H. J. Eichler, P. Guenter, and D. W. Pohl, Laser-Induced Dynamic Gratings. Berlin: Springer-Verlag, 1986.
- [201] A. A. Maznev, K. A. Nelson, and J. A. Rogers, “Optical heterodyne detection of laser-induced gratings,” Optics Letters, vol. 23, no. 16, pp. 1319–1321, 1998.
- [202] J. A. Johnson, A. A. Maznev, M. T. Bulsara, E. A. Fitzgerald, T. C. Harman, S. Calawa, C. J. Vineis, G. Turner, and K. A. Nelson, “Phase-controlled, heterodyne laser-induced transient grating measurements of thermal transport properties in opaque material,” Journal of Applied Physics, vol. 111, no. 2, p. 023503, 2012.
- [203] J. A. Johnson, A. A. Maznev, J. Cuffe, J. K. Eliason, A. J. Minnich, T. Kehoe, C. M. S. Torres, G. Chen, and K. A. Nelson, “Direct measurement of room-temperature nondiffusive thermal transport over micron distances in a silicon membrane,” Physical Review Letters, vol. 110, no. 2, p. 025901, 2013.
- [204] A. A. Maznev, F. Bencivenga, A. Cannizzo, F. Capotondi, R. Cucini, R. A. Duncan, T. Feurer, T. D. Frazer, L. Foglia, H.-M. Frey, H. Kapteyn, J. Knobloch, G. Knopp, C. Masciovecchio, R. Mincigrucci, G. Monaco, M. Murnane, I. Nikolov, E. Pedersoli, A. Simoncig, A. Vega-Flick, and K. A. Nelson, “Generation of coherent phonons by coherent extreme ultraviolet radiation in a transient grating experiment,” Applied Physics Letters, vol. 113, no. 22, p. 221905, 2018.

- [205] J. Ringling, O. Kittelmann, F. Noack, G. Korn, and J. Squier, “Tunable femtosecond pulses in the near vacuum ultraviolet generated by frequency conversion of amplified Ti:sapphire laser pulses,” Optics Letters, vol. 18, no. 23, pp. 2035–2037, 1993.
- [206] F. Rotermund and V. Petrov, “Generation of the fourth harmonic of a femtosecond Ti:sapphire laser,” Optics Letters, vol. 23, no. 13, pp. 1040–1042, 1998.
- [207] A. A. Maznev, T. F. Crimmins, and K. A. Nelson, “How to make femtosecond pulses overlap,” Optics Letters, vol. 23, no. 17, pp. 1378–1380, 1998.
- [208] J. Liu and R. M. A. Azzam, “Polarization properties of corner-cube retroreflectors: theory and experiment,” Applied Optics, vol. 36, no. 7, pp. 1553–1559, 1997.
- [209] A. Vega-Flick, J. K. Eliason, A. A. Maznev, A. Khanolkar, M. Abi Ghanem, N. Boechler, J. J. Alvarado-Gil, and K. A. Nelson, “Laser-induced transient grating setup with continuously tunable period,” Review of Scientific Instruments, vol. 86, no. 12, p. 123101, 2015.
- [210] R. R. Kunz, V. Liberman, and D. K. Downs, “Experimentation and modeling of organic photocontamination on lithographic optics,” Journal of Vacuum Science and Technology B, vol. 18, no. 3, pp. 1306–1313, 2000.
- [211] J. Hollenshead and L. Klebanoff, “Modeling radiation-induced carbon contamination of extreme ultraviolet optics,” Journal of Vacuum Science and Technology B, vol. 24, no. 1, pp. 64–82, 2006.
- [212] L. R. Rabiner, R. W. Schafer, and C. M. Rader, “The chirp z-transform algorithm,” IEEE Transactions on Audio and Electroacoustics, vol. 17, no. 2, pp. 86–92, 1969.
- [213] V. V. Gorbachev, V. M. Durasov, R. B. Zezin, E. V. Ivakin, A. S. Rubanov, and N. A. Tatyana, “Light-induced thermal gratings in natural diamond,” Physica Status Solidi B, vol. 150, no. 2, pp. 901–905, 1988.
- [214] A. Tokmakoff, W. F. Banholzer, and M. D. Fayer, “Thermal diffusivity measurements of natural and isotopically enriched diamond by picosecond infrared transient grating experiments,” Applied Physics A, vol. 56, pp. 87–90, 1993.
- [215] E. V. Ivakin, A. V. Sukhodolov, V. G. Ralchenko, A. V. Vlasov, and A. V. Khomich, “Measurement of thermal conductivity of polycrystalline CVD diamond by laser-induced transient grating technique,” Quantum Electronics, vol. 32, no. 4, p. 367, 2002.
- [216] T. Malinauskas, K. Jarašiūnas, E. Ivakin, N. Tranchant, and M. Nesladek, “Determination of carrier diffusion coefficient and lifetime in single crystalline CVD diamonds by light-induced transient grating technique,” Physica Status Solidi A, vol. 207, no. 9, pp. 2058–2063, 2010.
- [217] M. Kozák, F. Trojáněk, and P. Malý, “Optical study of carrier diffusion and recombination in CVD diamond,” Physica Status Solidi A, vol. 210, no. 10, pp. 2008–2015, 2013.
- [218] P. Ščajev, T. Malinauskas, L. Lubys, E. Ivakin, M. Nesladek, K. Haenen, and K. Jarašiūnas, “Optical monitoring of nonequilibrium carrier diffusion in single crystalline CVD and HPHT diamonds under high optical excitation,” Physica Status Solidi RRL, vol. 5, no. 5-6, pp. 193–195, 2011.

- [219] A. Taylor, P. Ashcheulov, P. Hubík, Z. Weiss, L. Klimša, J. Kopeček, J. Hrabovsky, M. Veis, J. Lorinčík, I. Elantýev, and V. Mortet, “Comparative determination of atomic boron and carrier concentration in highly boron doped nano-crystalline diamond,” Diamond and Related Materials, vol. 135, p. 109837, 2023.
- [220] M. Kozák, F. Trojánek, and P. Malý, “Temperature and density dependence of exciton dynamics in IIa diamond: experimental and theoretical study,” Physica Status Solidi A, vol. 211, no. 10, pp. 2244–2250, 2014.
- [221] N. Naka, H. Morimoto, and I. Akimoto, “Excitons and fundamental transport properties of diamond under photo-injection,” Physica Status Solidi A, vol. 213, no. 10, pp. 2551–2563, 2016.
- [222] J. Isberg, J. Hammersberg, E. Johansson, T. Wikström, D. J. Twitchen, A. J. Whitehead, S. E. Coe, and G. A. Scarsbrook, “High carrier mobility in single-crystal plasma-deposited diamond,” Science, vol. 297, no. 5587, pp. 1670–1672, 2002.
- [223] M. Nesladek, A. Bogdan, W. Deferme, N. Tranchant, and P. Bergonzo, “Charge transport in high mobility single crystal diamond,” Diamond and Related Materials, vol. 17, no. 7-10, pp. 1235–1240, 2008.
- [224] C.-M. Li, T. Sjodin, and H.-L. Dai, “Photoexcited carrier diffusion near a Si(111) surface: non-negligible consequence of carrier-carrier scattering,” Physical Review B, vol. 56, no. 23, pp. 15252–15255, 1997.
- [225] J. F. Young and H. M. van Driel, “Ambipolar diffusion of high-density electrons and holes in Ge, Si, and GaAs: many-body effects,” Physical Review B, vol. 26, no. 4, pp. 2147–2158, 1982.
- [226] N. Naka, K. Fukai, Y. Handa, and I. Akimoto, “Direct measurement via cyclotron resonance of the carrier effective masses in pristine diamond,” Physical Review B, vol. 88, no. 3, p. 035205, 2013.
- [227] H. M. van Driel, J. S. Preston, and M. I. Gallant, “Confinement of laser-generated carriers in semiconductors by induced lattice temperature gradients,” Applied Physics Letters, vol. 40, no. 5, pp. 385–387, 1982.
- [228] P. Ščajev, “Excitation and temperature dependent exciton-carrier transport in CVD diamond: diffusion coefficient, recombination lifetime and diffusion length,” Physica B: Condensed Matter, vol. 510, pp. 92–98, 2017.
- [229] M. Nagai, R. Shimano, K. Horiuchi, and M. Kuwata-Gonokami, “Phase diagram of the quantum degenerate electron-hole system in diamond,” Physica Status Solidi B, vol. 238, no. 3, pp. 509–512, 2003.
- [230] D. E. Kane and R. M. Swanson, “The effect of excitons on apparent band gap narrowing and transport in semiconductors,” Journal of Applied Physics, vol. 73, no. 3, pp. 1193–1197, 1993.
- [231] M. Delor, H. L. Weaver, Q. Yu, and N. S. Ginsberg, “Imaging material functionality through three-dimensional nanoscale tracking of energy flow,” Nature Materials, vol. 19, no. 1, pp. 56–62, 2020.

- [232] J. Thurston, T. Kaffle, L. Weiner, M. Murnane, and H. Kapteyn, “Bright tunable ultrafast deep- and vacuum-ultraviolet harmonic combs,” in Frontiers in Optics + Laser Science 2023 (FiO, LS), p. FTh3C.7, Optica Publishing Group, 2023.
- [233] A. Gromov, A. Lucas, and R. M. Nandkishore, “Fracton hydrodynamics,” Physical Review Research, vol. 2, no. 3, p. 033124, 2020.
- [234] A. Gupta, J. J. Heremans, G. Kataria, M. Chandra, S. Fallahi, G. C. Gardner, and M. J. Manfra, “Hydrodynamic and ballistic transport over large length scales in GaAs/AlGaAs,” Physical Review Letters, vol. 126, no. 7, p. 076803, 2021.
- [235] G. Varnavides, A. Yacoby, C. Felser, and P. Narang, “Charge transport and hydrodynamics in materials,” Nature Reviews Materials, vol. 8, no. 11, pp. 726–741, 2023.

Appendix A

Code repositories

The various codes discussed in the main text are located in the GitHub repositories at the locations below, which are currently private and can be shared upon request. The Readme files in these repositories contain documentation on their inputs, usage and capabilities.

Fresnel Diffraction Code - <https://github.com/bennettmcbrendan/Diffract>

Multiparameter Fitting Code - https://github.com/bennettmcbrendan/Fit_Multiparam

DUV-TG Experiment Code - <https://github.com/bennettmcbrendan/DUV-Code>

VO₂ Heater Code - <https://github.com/bennettmcbrendan/VO2-Code>

Appendix B

EUV scatterometry and DUV transient grating experimental documentation

Three documents in the Lab Setup and planning/Manuals folder on the Nanothermal Dropbox discuss alignment and troubleshooting procedures for the EUV scatterometry and DUV transient grating experiments described in this thesis.

B1B19 EUV Scatterometry Lab Guide 2019 Edition - By Begoña, discusses FROG, M2, alignment, HHG and gas system

B1B19 EUV Scatterometry Lab Guide 2021 Edition - Discusses the B1B19 RAEA amplifier, alignment, making and installing fibers, HHG, timing and Smaract stage

B1B35 DUV-TG Lab Guide 2024 Edition - Discusses alignment, finding TG signal and troubleshooting

Appendix C

B1B35 Wyvern alignment procedure

The B1B35 Wyvern amplifier alignment procedure is given in its manual in B1B35. However, we have found the following slightly modified and more detailed procedure to work well in bringing it back to lasing when for some reason it requires a total realignment. I would like to thank former graduate student Nathan Brooks for teaching me many of the tricks included in this procedure. This procedure presupposes that the oscillator is working well. A discussion of the oscillator is given in the “Amy’s Guide to the Oscillator” pamphlet stored behind the oscillator itself. In the procedure below, bolded text indicates steps that are very important to avoid damaging the laser.

- (1) **Make sure oscillator is working well and has sufficient bandwidth (at least 40nm FWHM, ideally more).**
- (2) Align seed beam from oscillator into stretcher using first two mirrors inside amplifier box and the two irises before stretcher.
- (3) Defeat first Pockels cell (PC) by turning half-wave plate (HWP) directly after it from 260° to $\sim 305^\circ$.
- (4) Align seed into cavity using the two mirrors after the first PC cell and the two irises along the long line before cavity. This is difficult to do well since beam is small on first iris and the mirrors are closely spaced, necessitating lots of iteration. Align by looking at red border on black iris when clamping down (with IR viewer if necessary) rather than red center on

card after iris. This synchronizes this procedure with the cavity alignment retroreflecting in steps 5 and 11.

- (5) Defeat second PC by turning quarter-wave plate (QWP) inside cavity to the position where white lines are vertical and angle reads $\sim 45^\circ$ (hereafter the ‘vertical’ position). This causes the seed to enter cavity by passing through the thin film polarizer (TFP), QWP and second PC, reflect off the cavity end mirror and then exit the cavity by traveling back through second PC, QWP and TFP, proceeding backwards along the line you aligned in step 4 before being picked off by another TFP to go to the compressor.
- (6) Adjust the cavity end mirror behind second PC to retroreflect the seed onto the second iris from the step 4 alignment. Use IR viewer to symmetrize its border on the back side of this iris, the mirror image of how you aligned in step 4.
- (7) You can optionally align the second PC at this point using the procedure in the manual and the Liam tune-up post, but there’s no need to do this every time.
- (8) Turn the cavity QWP so that the white lines are at 45° and angle reads $\sim 0^\circ$ (hereafter the ‘diagonal’ position). This will allow the seed to make exactly one pass in the cavity by first transmitting through the TFP, then reflecting off the TFP after its first two trips through the QWP and then passing through the TFP after its second two trips through the QWP. The seed will then travel backwards along the line aligned in step 4 before being picked off by another TFP to go to the compressor.
- (9) Next use the two irises stored in the right corner of the amplifier box to align the seed through the crystal. There are two posts for the two irises on either side of the crystal, the iris with the smaller collar goes in the post on the far side (pump laser input side) of the crystal. Use the 2” mirror and awkward 1” mirror on the closer side of the crystal input tube to do this. This step is annoying because the seed looks bad and is diverging a lot due to lack of thermal lensing, so you have to just pick a consistent reference and do your best.

(10) Turn on the pump laser at its lowest power (9.6 A current). It should be pretty well lined up with the seed through the irises. Tweak the seed alignment to get it to match the pump laser on either side of the crystal as best you can. Only mess with the pump steering if absolutely necessary. Next remove the first two ND filters from the crystal webcam. You should see the seed and pump beam on the crystal. The seed will be much larger. Tweak seed further if needed to make sure alignment is good here. Then turn pump down below threshold current so that it no longer lases.

(a) Note: The Brewster-cut crystal will refract the green pump light more than the red seed light. This means it is advantageous to have the seed slightly above the pump on the pump laser input side and slightly below the pump on the other side. However, this correction should be very slight, no more than 20% of the beam width, and isn't strictly necessary at all.

(11) Now let the seed go all the way through the cavity for one pass by reflecting it off the motorized cavity end mirror and back out the line aligned in step 4. Tweak the motorized cavity end mirror to overlap the retroreflected seed beam with itself. Note the retroreflected seed beam will be very large, so this step is difficult and often requires the IR viewer. I like to use a notecard on a post with a small hole cut in it that lets the incoming seed beam through but traps most of the large retroreflected seed beam on the back side of the card. I place the card-on-post near the second iris used in step 4, and observing the retroreflected beam with the lights off (no IR viewer required!) move the motorized mirror with large steps (size 100) to translate the retroreflected beam far to either side before splitting the difference. The card-on-post is stored in the right corner of the amplifier box along with the irises used in step 9. Note there is also another large back reflection, perhaps from the QWP that can confuse things. This back reflection however doesn't move with the cavity end mirror adjustment. There are many other ways to do this step:

(a) Retroreflect onto the same iris used to align the other cavity end mirror in step 5

- (b) Retroreflect onto a card placed just after the TFP used to kick out the post-cavity seed to the compressor and maximize intensity of this card
 - (c) Retroreflect onto the back side of iris placed before crystal in step 9 (closer so easier to see but less precise)
 - (d) Retroreflect onto a card that you wave in front of the beam on either side of the crystal, blocking half of the incoming beam to see half of retroreflected beam (helpful if you're way off, but not precise enough on its own)
 - (e) The manual says at this point to also check the retroreflection on the crystal camera, but I can never see the reflected beam on the camera well and the beams are huge anyways, so have been skipping this step
- (12) Now the cavity is aligned! Prepare for cavity lasing via following steps:
- (a) **Make sure all irises are open**
 - (b) Make sure all cards are removed, except one at amplifier exit
 - (c) Turn HWP after PC 1 to 260°
 - (d) **Turn on PC 1 and PC 2 if not on already.** Note you can also do step 13 with both PCs off, but we have observed empirically that it doesn't make a difference, since PC 1 only touches the seed and PC 2 is only on a tiny fraction of the time. Since they definitely need to be on when seeding, especially PC 2, it's safest to turn on here.
 - (e) Turn QWP in cavity to 'vertical' position. This will trap any lasing in the cavity and allow you to see it more easily (but be careful with this, see next step!)
 - (f) Replace all ND on crystal camera
 - (g) **BLOCK SEED MULTIPLE TIMES. If seed goes in with bad cavity mode and QWP in wrong position at high pump power, it can destroy laser.**
- (13) Slowly turn up pump laser, 0.5 - 1A at a time until you see any lasing at all in cavity or on card thereafter. Note you will see a little red scatter everywhere with green-blocking goggles

on as the pump laser ramps up. This is florescence, not lasing. Check with IR viewer on cavity optics not contacted by pump beam, where you should see a distinct beam if lasing. If you don't see any lasing by the time you get to about 85-95% of the normal pump laser operating current (a very likely scenario), **try tweaking the cavity end mirror behind PC 2**. Make slow, small adjustments, trying to center them about where you started to avoid messing up your alignment too much. If you see any red lasing, immediately reverse the move you made so that it disappears, don't try to optimize it yet. If you optimize the red with the QWP trapping the mode in the cavity, you can cause it to build up uncontrollably, which will burn the cavity optics. If can't see red lasing, try the other cavity end mirror or tiny adjustments to pump steering. If it still doesn't work, you will probably need to realign.

- (14) **Once you see the slightest bit of lasing, don't optimize it. Instead turn cavity end mirror just enough to make it go away.** Start rotating cavity QWP towards the 'diagonal' position. This will prevent uncontrolled buildup in cavity. Do 5-10 degrees at first. Tweak end mirrors to get a tiny amount of lasing back. This may be hard to do at first, since now it will be harder to build up in cavity and so you will have to have end mirrors in a better position to see something. Now iterate the following steps:
- (a) Turn up pump current slightly towards the normal operating value (currently 19.5A as of March 2024)
 - (b) Tweak manual cavity end mirror to make red mode look nicer. Don't worry if mode looks terrible at first, it likely will improve with more fishing around and iterating. Don't shoot for perfection or make any big moves on a single iteration, just make the smallest move necessary to improve the most awful part of the clipping
 - (c) Click around motorized cavity end mirror to make pump and cavity modes align on crystal camera. This occurs when the depletion is even across the pump beam mode. If you see a very bright, pretty, symmetric pump beam mode with camera exposure

set to $6\text{e-}6$, then there is likely very little depletion anywhere and you are way off.

- (d) Turn QWP in cavity another few degrees towards ‘diagonal’ position
 - (e) Keep doing these iterations till you are satisfied with ASE mode on card at amplifier exit, the pump current is all the way up and the QWP is approximately in its normal ‘diagonal’ position (reading $\sim 0^\circ$).
- (15) Align ASE beam into compressor using the two mirrors before compressor and irises before and after compressor. **Be careful not to close down on iris before compressor without beam blocked directly on card thereafter, since an airy pattern can burn the compressor gratings.**
- (16) Next measure power at amplifier exit. Make small (1-3 degree) movements on the cavity QWP to optimize power. **However, don’t depart more than a few degrees from the ‘diagonal’ position.** Make sure HWP before compressor is optimized for power. Ideally ASE power will be around 5-6 W. You can make small adjustments to pump steering to increase power more, but I typically don’t do this. **Note that moving the pump beam onto the crystal housing at full power can deposit metal on the crystal, necessitating its cleaning, however as long as you keep the pump beam visible on the crystal camera, you should be far enough away.**
- (17) Now you should be ready to seed! Close amplifier box lids and replace black foil, **double check that PCs are on**, get oscilloscope reading the photodiodes measuring cavity buildup and pulse train at amplifier exit. Seed amplifier and look at following things:
- (a) By how much did power increase when seeding? Want to maximize this.
 - (b) Does the cavity buildup look good on oscilloscope and are there large pre- and post-pulses after the main pulse?
 - (c) Does seeded spectrum look symmetric, and did it red shift to around 790nm central wavelength?

- (18) Tune seed-pump overlap with second seed steering mirror used in step 4 to align seed into cavity. Don't open lid all the way to do this, just reach hand under. Maximize size of pulses on oscilloscope and spectrum quality, not power. Ideally want to use this to remove any blue shoulder to the seeded spectrum, an indication of bad overlap. Make small adjustments and **be very careful not to block part of beam in stretcher as you reach over stretcher for seed mirror, since this will likely burn optics in cavity.**
- (19) Now adjust pump timing on timing screen. Want it to be at the value that maximizes the difference between seeded and unseeded power, not at the value that maximizes overall power. Usually, this is a pump timing value that is large enough for overall power to start decrease, but right before it falls off a cliff.
- (20) Next adjust PC 2 timings. Try adding or subtracting a pass by making 2-3 10 ns moves to PC off and see if you can dramatically increase power by adding a pass or keep the same power by subtracting a pass. If so, run with that new number of passes, if not, switch back. **Don't run with too many passes, especially when the last pass is no longer increasing cavity buildup on oscilloscope.** Next make smaller 1 ns adjustments within the number of passes you've selected. May need to iterate this a bit with pump timing if you add or subtract passes.
- (21) Finally, try adjusting PC 1 timings and PC 2 voltage (knob on rack by power switch) by small amounts to minimize pre- and post-pulsing on oscilloscope on the 10 ns timescale. Ideally want main pulse to be 50x higher than pre and post pulses. Laser is now ready to go!

Appendix D

Electron-beam lithography recipe

The following recipe provides the essential steps to fabricate nanoscale nickel structures on a substrate. Steps 4 and 6 are only necessary for insulating substrates. When using the sonicator, the power level should be set based on the sample size, to ensure that the sample vibrates slightly but doesn't flip over.

- (1) Clean substrates using sequential 1 minute sonications in IPA, methanol and ethanol, followed by a 1 minute rinse under flowing deionized (DI) water. Blow dry with compressed N₂, then place on clean hot plate for 10 minutes at 100C.
- (2) Spin coat A2 950K PMMA at 6000 rpm (2000 rpm/s acceleration) for 45 seconds. Also include a 5 second, 500 rpm (250 rpm/s acceleration) step at the beginning. This should result in an approximately 50 nm PMMA layer.
- (3) Immediately after spin coating, bake the sample on a hot plate at 180C for 90 seconds.
- (4) Deposit 15 nm of aluminum on top of the PMMA using the CVC thermal evaporator.
- (5) Align the Hitachi SEM and use the NPGS software, pattern and run files to write the desired pattern into the PMMA. Any current ≤ 50 pA will produce equally high-resolution patterns.
- (6) Remove the aluminum layer by submerging the sample in a 2.5% NaOH, 97.5% H₂O solution for approximately 30 seconds. Immediately rinse in DI water and blow dry with N₂.

- (7) Develop the PMMA in a 3:1 IPA:H₂O solution cooled to <5C. Submerge the sample for 90 seconds, gently swirling it around the beaker. Immediately rinse in DI water and blow dry with N₂.
- (8) Deposit 12 nm of nickel, or approximately 25% of the PMMA thickness, using the CVC thermal evaporator.
- (9) Submerge the sample in Remover PG on a hot plate set to 100C for 2-3 hours, then sonicate for 10 minutes, then switch NMP baths, repeating this cycle 2-3 times or until no PMMA remains. Immediately after the last bath, rinse in DI water and blow dry with N₂.

At several points in the recipe, it is possible to perform checks with various metrology techniques to verify that the process is on track. The PMMA thickness can be verified after step 2 using the JILA ellipsometer. After step 7, the fidelity of the PMMA structures can be checked with an optical microscope or AFM, depending on their size. Finally, AFM can verify the quality of the nickel structures after step 9.

The two recipes for diamond substrates discussed in Chapter 5 contain various deviations from the above process. The first recipe, performed at LBNL by our collaborator Weilun Chao, uses a diamond substrate purchased from READE Advanced Materials. The sample was polished, but there is no record of what cleaning READE performed prior to shipment. The fabrication process began with additional sample cleaning, first with acetone and IPA, and then UV ozone for 10 minutes. The sample was next spin-coated with PMMA using a solution of 1% PMMA dissolved in anisole at 1000 rpm, and then baked in an oven on a copper plate at 115C for 20 minutes. Electron-beam evaporation was used to evaporate 20 nm of aluminum on top of the PMMA to serve as a conduction layer during the lithography. The EBL writing step was performed using a Vistec VB300 system at 100 keV energy and ~1 nA current. After the EBL step, the aluminum conduction layer was removed by soaking the sample in the resist developer MF-26A, which contains the (dangerous) active ingredient TMAH. The development step was identical to the recipe above. After development and optical inspection, 12 nm of nickel was deposited on the sample, again using

the electron-beam evaporator. The liftoff step was very similar to that described above, with minor temperature and timing differences.

The second recipe was carried out in JILA and at COSINC. The company Applied Diamond supplied the substrates, which were polished, then dry-etched to remove any sub-surface damage, then cleaned in a boiling mixture of nitric and sulfuric acids prior to shipment. We performed no additional cleaning upon receipt, but the sample was baked at 180C for 10 minutes on a hot plate before spin coating. The remainder of the fabrication was very similar to the recipe above, with minor temperature and timing differences.

Appendix E

Electron-beam lithography chemical and instrument list

The following tables document the chemicals and instruments used in the electron beam lithography processes described in Chapter 5 and Appendix D.

Chemical	Supplier	Item No.	Location	Notes
PMMA	Kayaku	M230002 0500L1GL	JILA clean room	Dissolved in anisole
NaOH	Curtis	-	X110	2.5% NaOH in H ₂ O
Remover PG	Kayaku	G050200 4000L1PE	JILA clean room	Contains NMP

Table E.1: Data on chemicals used in electron beam lithography process. The NaOH is only necessary when washing off an aluminum discharge layer.

Instrument	Model	Location
Spin coater	WS-650Mz-23NPPB	JILA clean room
Hot plate	-	JILA clean room
Converted SEM	Hitachi SU3500 VP SEM	COSINC-CHR
Evaporator	CVC SC-300	COSINC-FAB clean room
Sonicator	-	JILA clean room

Table E.2: Information on instruments used in the electron beam lithography process.

The one-year expiration dates on the PMMA and Remover PG can be ignored for quite some time thereafter as long as the chemicals remain in their original bottles in a stable environment.

New bottles of the PMMA and Remover PG cost roughly \$500 and \$200, respectively. All chemicals should be used exclusively in fume hoods, especially the Remover PG, which contains NMP, a reproductive toxin that is dangerous to inhale. Acetone is a safer alternative to NMP for lift-off, however, it tends to produce inferior results. Acetone should not be heated due to its low flash point. The chemical list omits readily available solvents such as IPA, which are also used in the EBL process.