Title
Thermal Conductivity Measurements in Challenging Regimes

Author
Lubner, Sean

Publication Date
2016

Peer reviewed|Thesis/dissertation

Thermal Conductivity Measurements in Challenging Regimes

By

Sean Daniel Lubner

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Mechanical Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Christopher Dames, Chair
Professor Holger Müller
Professor David Steigmann

Fall 2016

Thermal Conductivity Measurements in Challenging Regimes

by

Sean Daniel Lubner

Abstract

Thermal Conductivity Measurements in Challenging Regimes

by

Sean Daniel Lubner

Doctor of Philosophy in Engineering – Mechanical Engineering

University of California, Berkeley

Professor Christopher Dames, Chair

Creating technologies to address increasingly diverse challenges ranging from biomedical devices to carbon-free energy solutions requires measuring material properties and system behaviors in increasingly challenging regimes. In the biomedical field, accurate determination of the thermal conductivity (k) of biological tissues is important for cryopreservation, thermal ablation, and cryosurgery, but is hampered by the delicate nature and often-small sizes of tissues. In the electronics and clean energy fields, it is increasingly necessary to reliably model the dissipation of heat from micro and nanoelectronics for thermal management, and the transport of heat through nanostructured materials for energy control and conversion technologies such as batteries and thermoelectrics. However, the classical equations of heat transfer break down at these short length scales, calling into question the validity of various formulations of heat transfer theory and the very concept of thermal conductivity itself. Here, too, is a need for challenging thermal conductivity measurements at micron and nanometer scales. In this thesis, we describe and demonstrate two techniques that combined are capable of measuring the key thermal transport properties in all of these regimes.

We adapt the 3ω method—widely used for rigid, inorganic solids—as a reusable sensor to measure k of soft biological samples, two orders of magnitude thinner than conventional tissue characterization methods. Analytical and numerical studies quantify the error of the commonly used "boundary mismatch approximation" of the bi-directional 3ω geometry, confirm that the generalized slope method is exact in the low-frequency limit, and bound its error for finite frequencies. The bi-directional 3ω measurement device is validated using control experiments to within ± 2% (liquid water, std. dev.) and ± 5% (ice). Measurements of mouse liver cover a temperature range from -69 ºC to + 33 ºC. The liver results are independent of sample thicknesses from 3 mm down to 100 μm, and agree with available literature for non-mouse liver to within the measurement scatter.

Next, we focus the laser spot $1/e^2$ radius in TDTR measurements down to single micron length scales to measure quasi-ballistic thermal transport at length scales where Fourier's law breaks down. We present an in-depth discussion of the instrumentation and provide comprehensive analyses of system sensitivities to all experimental parameters. The

system is first validated on sapphire and single crystal silicon control samples. We then measure two nano-grained Si samples (550 nm and 76 nm average grain size) and two SiGe alloys (1% and 9.9% Ge concentration), representing two classes of silicon-based materials with qualitatively different phonon scattering physics. All samples are 5 mm x 5 mm x 0.5 mm or larger. Sub-diffusion measurements are performed on all samples using $1/e^2$ laser spot radii down to 1.6 μm. Apparent thermal conductivity suppressions ranging from 18% to 76% are observed at room temperature, indicating that while most of the heat in sapphire, Si, and nano-grained Si is carried by phonons with mean free paths of a couple microns or less, much of the heat in SiGe alloys is still carried by phonons with mean free paths up to a few tens of microns at room temperature. We present a discussion of the microscale origins of this suppressed thermal conductivity and its physical interpretation, addressing some common misconceptions. Our results show that alloying and nanostructuring shift the spectral phonon mean free path distributions in opposite directions. Alloying skews the phonon distribution toward long mean free paths, increasing k suppression at small length scales, while nanostructuring skews the distribution toward short mean free paths, reducing k suppression.

Table Of Contents

List Of Figures

List Of Tables

Acknowledgements

I am grateful to UC Berkeley and the National Science Foundation for their generous financial support. At the end of the day, good and consistent science is impossible without adequate funding, and it is important to remember those people and institutions that enabled the research and all resulting benefits to exist in the first place by providing the necessary resources.

I would like to offer tremendous thanks to all of my mentors, collaborators, and colleagues for offering invaluable guidance and insight, and to all friends, family, and loved ones for their unwavering support both in academics and in life. No person can accomplish alone anything close to what they can achieve with the help of a caring and diverse support network, and I have been incredibly fortunate to have a support network far beyond everything I had hoped for.

In addition, I give specific thanks to my advisor, Professor Chris Dames. Chris is one of those rare mentors who is a role model in every way. He genuinely cares about his students and sets their best interests and professional development as his top priority, even above his own. I have learned so much from Chris regarding diligence, humility, teaching, communication, dedication, organization, and of course thermal physics and all aspects of good scientific research.

Most importantly, I thank my family. My parents and my brother have supported and encouraged me from the day I was born. The older I get, the more I realize just how lucky I am, and the less I take for granted. This body of work is just a drop in the bucket compared to the hard work and sacrifices my family has happily made without hesitation to help me achieve my goals and pursue my life's passion in science, and I will be forever grateful.

Chapter 1 - Introduction

1.1 Setting The (Thermal) Stage

Any technology that performs work, processes information, or generally transfers or converts energy from one form into another—really all conceivable technology—will at some point have to deal with the dissipation of heat. So says the second law of thermodynamics. The macroscopic dissipation of heat in solid matter is described by the heat conduction equation, which requires knowledge of the system geometry and a few physical properties including the material's thermal conductivity, k. At steady-state operation, the thermal conductivity is the only material property that must be known in order to calculate the temperature distributions and thermal energy fluxes in a material. For many systems, the thermal conductivity is easily measured and Fourier's law of heat conduction is valid. For these systems, thermal engineering has flourished for over a century and we are all the fortunate beneficiaries of the technologies that have been produced as a result.

However, as we have created more advanced technologies, we have encountered more challenging regimes where the thermal conductivity cannot be measured so easily, and the classical equations of heat conduction are not always valid. The challenges of such regimes include samples that are wet, soft, and chemically sensitive, or systems with important characteristic lengths at micron or nanometer scales where the diffusive Fourier law paradigm of heat conduction no longer holds true.

The potential benefits of measuring thermal conductivities and characterizing thermal transport for these nontrivial systems are many. A swath of thermal biomedical procedures rely on the ability to accurately model heat transport in tissues in order to use heat and cold to kill diseased or dangerous tissues or to keep good tissues and organs alive. The data storage and transistor densities in microelectronics, the thermal management of batteries large and small, and the efficiency of LED lighting and solid state energy conversion devices are all becoming increasingly dependent on knowledge of heat conduction on nanometer and micron length scales. There exists a growing opportunity for broader impacts across many areas with the ability to more accurately, efficiently, and reliably measure thermal conductivities and heat transfer physics in challenging regimes.

1.2 Organization Of This Dissertation

This thesis can be divided into two contiguous parts, each with its own more personalized introduction and conclusions sections. The first part presents the bi-directional 3ω method and how it can be applied to measure k of liquids, and frozen and non-frozen hydrated biological tissues. The second part presents the variable spot size TDTR technique and how it can be used as a phonon mean free path spectroscopy tool.

Chapter 2 - Bi-Directional 3-Omega introduces the technique, deriving and fully explaining both the mathematical theory and experimental implementation. A rigorous analysis of error quantification is presented alongside discussions of sensitivities and our own observed experimental uncertainties. The technique is demonstrated on control samples of water, ice, and

agar gel, and then applied to thermal conductivity measurements of both frozen and non-frozen mouse tissue samples ranging in thicknesses from 3 mm down to 100 μm. Recommended avenues for future research are presented.

Chapter 3 - Variable Spot Size TDTR introduces the technique and provides broader context in which to understand the utility and motivation for developing such a technique. It generally discusses all of the practical details of actually implementing the technique experimentally from start to finish, including key pieces of hardware and their purpose, robust chromatic aberration-corrected laser spot size measurements, sample preparation and characterization, mathematical modeling and data analysis, and uncertainty and sensitivity analyses.

Chapter 4 - Effects of Alloying and Nanostructuring on Spectral Phonon Mean Free Path Distributions explains the microscopic origins of thermal conductivity and the correct physical way to interpret the results of Fourier law-based measurements of quasi-ballistic thermal transport experiments. Measurements of sapphire, Si, nanograined-Si, and SiGe samples are presented and the results are carefully discussed. Differing trends in thermal conductivity suppression due to grain boundary scattering versus alloy scattering are explained in the context of spectral phonon mean free path distributions. Recommended avenues for future research are presented.

Dissertation Summary and Conclusions offers a summary of all experimental results from the thesis and their implications.

Chapter 2 - Bi-Directional 3-Omega

This chapter is very closely based on our recent publication [1]:

S. D. Lubner, J. Choi, G. Wehmeyer, B. Waag, V. Mishra, H. Natesan, J. C. Bischof, and C. Dames, "Reusable bi-directional 3ω sensor to measure thermal conductivity of 100-μm thick biological tissues," *Rev. Sci. Instrum.*, vol. 86, no. 1, p. 014905, Jan. 2015.

I would like to thank Shannon Yee for helpful suggestions regarding the dielectric layer, Yasuhiro Hasegawa and Anthony Fong for help with instrumentation, the UC Berkeley Office of Laboratory Animal Care for mouse liver tissue samples, and the UC Berkeley Biological Imaging Facility for use of their cryotome.

2.1 Introduction

Thermal bioengineering uses controlled heating and cooling to preserve healthy tissues as well as treat conditions such as prostate, breast, and renal cancer, atrial fibrillation, peripheral artery disease, and renal hypertension [2]–[7]. The efficacy and safety of these treatments requires precisely controlled heating and/or cooling rates, which is predicted using the thermal conductivity (k) combined with other thermal and geometric properties. However, traditional techniques to measure k of biological tissues, such as the cut-bar, guarded hot plate, and embedded thermistor methods, require samples at least ~10 mm thick [8]–[16]. Thus, numerous important thinner tissues cannot be characterized, including heart valves (1 - 2 mm), pulmonary vein (1 - 3 mm), esophagus (1 - 3 mm), small diameter arteries (1 mm), phrenic nerve (0.5 - 1 mm), cornea (0.5 mm), and fascia (0.1 mm). Furthermore, existing macroscopic techniques are more susceptible to errors due to parasitic heat losses to the environment and thermal contact resistances [9]–[11].

To address these shortcomings, we adapt the "3ω" method [17]–[19] for use with biological tissues. This frequency-domain electrothermal method is well established for measuring the thermal conductivity of dry, rigid, inorganic, solids down to sub-micron thicknesses [19], [20], such as semiconductor wafers and thin films. However, the traditional 3ω method cannot be applied directly to biological samples, because deposition of the 3ω heater line requires harsh microfabrication, which is incompatible with hydrated, soft, or chemically sensitive materials.

To overcome this severe limitation, here we separate the functions of sensor and sample by using a 2-sided "supported" or "bi-directional" 3ω method (Figure 2-1), which has also been used recently for some liquid and solid samples [21]–[24]. It has even been extended to measure individual cells [25], although this introduces additional challenges in microfabrication, isolating the appropriate cells and ensuring their vitality, and restricting the thermally-probed volume to be within the single cell. It also is not yet clear how to determine tissue-level thermal properties from corresponding measurements at the single-cell level. Most of the supported 3ω studies invoke a "boundary mismatch approximation" (BMA) [21]–[24] which forbids heat transfer through the interface between upper and lower subdomains everywhere except at the heater line

itself. This approximation greatly simplifies the analysis and thus is commonly assumed for bi-directional 3ω measurements, but has not been rigorously evaluated.

Here, we first analytically and numerically quantify the error that results from applying this BMA and the closely related generalized slope method. We then describe our experimental apparatus and control experiments on ice and liquid water. Finally, we apply the device to measure k of bulk biological tissues (mouse liver) down to 100 μm thick, over two orders of magnitude thinner than previous tissue measurements [8]–[16].

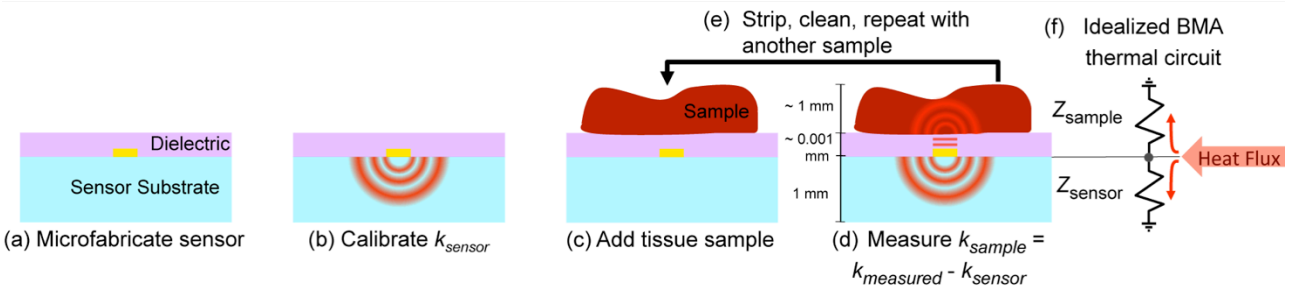

Figure 2-1: The bi-directional 3-omega technique for measuring soft tissues (not to scale). (a) The sensor base is a traditional 3ω geometry fabricated using standard photolithography, with an electrically insulating dielectric layer over the heater line. (b) Calibration of the sensor using Eq. (2-3) (introduced below). (c) Sample placed on top of the heater line. (d) Sample is measured. (e) Within 10 min, a new sample can replace the old one and be ready for measurement. (f) Idealized thermal circuit of impedances in parallel, corresponding to the common boundary mismatch approximation. Z denotes thermal transfer function.

2.2 Measurement Principle and Theory

2.2.1 Traditional 3ω

The traditional 3ω technique has been extensively described in the literature [17], [19]. A metal heater line deposited directly on the sample is driven by an electrical current of angular frequency ω, causing Joule heating P at frequency 2ω. The sample's thermal properties determine the resulting 2ω temperature response of the heater line, which cause a 2ω oscillation in the electrical resistance of the heater. This, combined with the 1ω driving electrical current, results in a voltage drop across the heater with a 3ω component, $V_{3\omega}$, whose magnitude can be related back to the thermal properties of the sample. The temperature oscillation amplitude decays into the sample over a characteristic penetration depth $\lambda = \sqrt{\alpha/\omega}$, where α is the thermal diffusivity. This tunable temperature localization reduces parasitic heat losses and enables measurements of thinner samples than are accessible with macroscopically-assembled steady-state methods. The amplitude of the AC temperature rise averaged over the heater line, ΔT, can be conveniently represented as a complex thermal transfer function $Z = \Delta T/P$ [18]. Throughout this paper we assume that the heater line length L is much larger than λ and the heater half-width b, and thus we use two-dimensional thermal models for Z.

For traditional one-sided 3ω (Figure 2-1(a)-(b)), Cahill [17] showed that the thermal transfer function for an isotropic semi-infinite sample is

$$\mathrm{Z}_{1-sided} = \frac{1}{\pi L}\int_0^\infty \frac{1}{\Upsilon_1}\frac{\sin^2(\xi b)}{(\xi b)^2}\,d\xi, \tag{2-1}$$

where $\Upsilon_j = k_j\sqrt{\xi^2 + i\frac{2\omega}{\alpha_j}}$, subscript j indexes the material, and ξ is the Fourier transform variable of the in-plane spatial coordinate x. When $\lambda \gg b$ (referred to as the "low-frequency" limit), the solution approaches [26]

$$\mathrm{Z}_{1-sided,low\ \omega} = \frac{1}{\pi L k_1}\left[-\frac{1}{2}\ln(\omega) - \frac{1}{2}\ln\left(\frac{b^2}{\alpha_1}\right) + \frac{3}{2} - \gamma - i\frac{\pi}{4}\right], \tag{2-2}$$

where γ is the Euler-Mascheroni constant. Taking the derivative of the real component of Eq. (2-2) with respect to $\ln(\omega)$ and converting $\mathrm{Z}_{1-sided,low\ \omega}$ into experimentally measured quantities [18], k can be expressed as

$$k = \frac{-1}{2\pi L\dfrac{\partial Re[Z]}{\partial \ln(\omega)}} = \frac{V_{1\omega}I_{1\omega}^2\dfrac{dR}{dT}}{4\pi L\dfrac{\partial V_{3\omega,in-phase}}{\partial \ln(\omega)}}, \tag{2-3}$$

where $I_{1\omega}$ is rms current, R is heater line resistance, and T is temperature. Equation (2-3) is referred to as the "slope method" for determining k in traditional one-sided geometries, and it is independent of the volumetric heat capacity C and heater-substrate contact resistance. In this paper, we generalize the slope method for a bi-directional 3ω geometry to measure k of biological tissue.

2.2.2 Bi-directional 3ω and the Boundary Mismatch Approximation (BMA)

In the bi-directional 3ω technique, the roles of sensor and sample are physically separated, with the sample being supported on a reusable sensor of known k_{sensor}. The sensor consists of a rigid substrate, a microfabricated heater line, and a protective dielectric layer, as illustrated in Figure 2-1(a). When a sample is placed on the sensor, as in Figure 2-1(c) and (d), the heater line temperature contains information about the thermal properties of the sample as well as the sensor.

Thus, measurements in the bi-directional 3ω configuration fundamentally give Z for the combined system of sensor and sample. The challenge is to extract k_{sample}. The most common treatment is the "boundary mismatch approximation" (BMA) [21]–[24]. This treats the sub-domains above ($Z_\uparrow$) and below ($Z_\downarrow$) the heater line as thermal impedances in parallel (Figure 2-1(f)) where $Z_\uparrow = Z_{sample}$ and $Z_\downarrow = Z_{sensor}$), thereby neglecting thermal crosstalk between the sub-domains, except at the heater line itself.

Analytical models of varying complexity have been developed previously to describe bi-directional heat flow from time-periodic heat sources in the contexts of optical thermoreflectance [27]–[29] and resistance thermometry [22]–[24], [30]–[33]. In the common and conceptually

straightforward BMA,[21]–[24] the combined (sensor + sample) thermal transfer function Z_{BMA} is constructed by simply summing $Z_{\uparrow}$ and $Z_{\downarrow}$ in parallel,

$$Z_{BMA}^{-1} = Z_{\uparrow}^{-1} + Z_{\downarrow}^{-1}. \tag{2-4}$$

The BMA thermal circuit for our experiment is shown in Figure 2-1(f). Effects of the dielectric layer and thermal contact resistances are addressed later, with full details given in Appendix 2.7.3. $Z_{\uparrow}$ and $Z_{\downarrow}$ are treated independently and determined analytically from a traditional one-sided 3ω geometry with an adiabatic top-side boundary condition. For example, consider two isotropic semi-infinite solids indexed as 1 and 2 and measured in a bi-directional 3ω geometry. In this case Z_{BMA} is found by summing Z_1 and Z_2 in parallel, where each Z is calculated from Eq. (3-2). In the low-frequency limit, what we will refer to simply as the "slope" throughout this section is shown to be [23]

$$\frac{\partial \mathrm{Re}[Z_{BMA,low\ \omega}]}{\partial \ln(\omega)} = \frac{-1}{2\pi L(k_1 + k_2)}. \tag{2-5}$$

It is clearly appealing to use this simple generalization of the one-sided slope method of Eq. (2-3) to determine k_{sample} in a bi-directional 3ω experiment, as long as the BMA errors are sufficiently small. Once the thermal conductivity of the sensor (k_1) has been calibrated, the thermal conductivity of the supported sample (k_2) can be calculated directly from measured values.

2.2.3 Assessing Errors in the Boundary Mismatch Approximation (BMA)

Figure 2-2 compares numerical solutions of the full thermally coupled problem (left) and the BMA prediction (right) for two semi-infinite materials, obtained using commercial finite element software (COMSOL). Since this model system has no contact resistance or intervening dielectric layer, the temperature, phase, and *y*-directed heat flux must be continuous across the interface, as verified in the full solution (left). The BMA (right) clearly violates the requirements for continuous temperature and phase, and imposes an artificial adiabat at the material interface, all of which call into question the range of validity for the generalized slope method presented in Eq. (2-5), as well as the BMA itself. We also note that the BMA deviation from the exact solution is larger for regions farther from the heater line (compare the two $|x| > 6b$ regions in Figure 2-2(a)). This suggests that errors in the BMA may be worse in a bi-directional extension of a recent "anisotropic 2ω" experiment [34], where a second temperature sensing line displaced in *x* along the interface is utilized, than for the combined heater/thermometer line in isotropic bi-directional 3ω measurements.

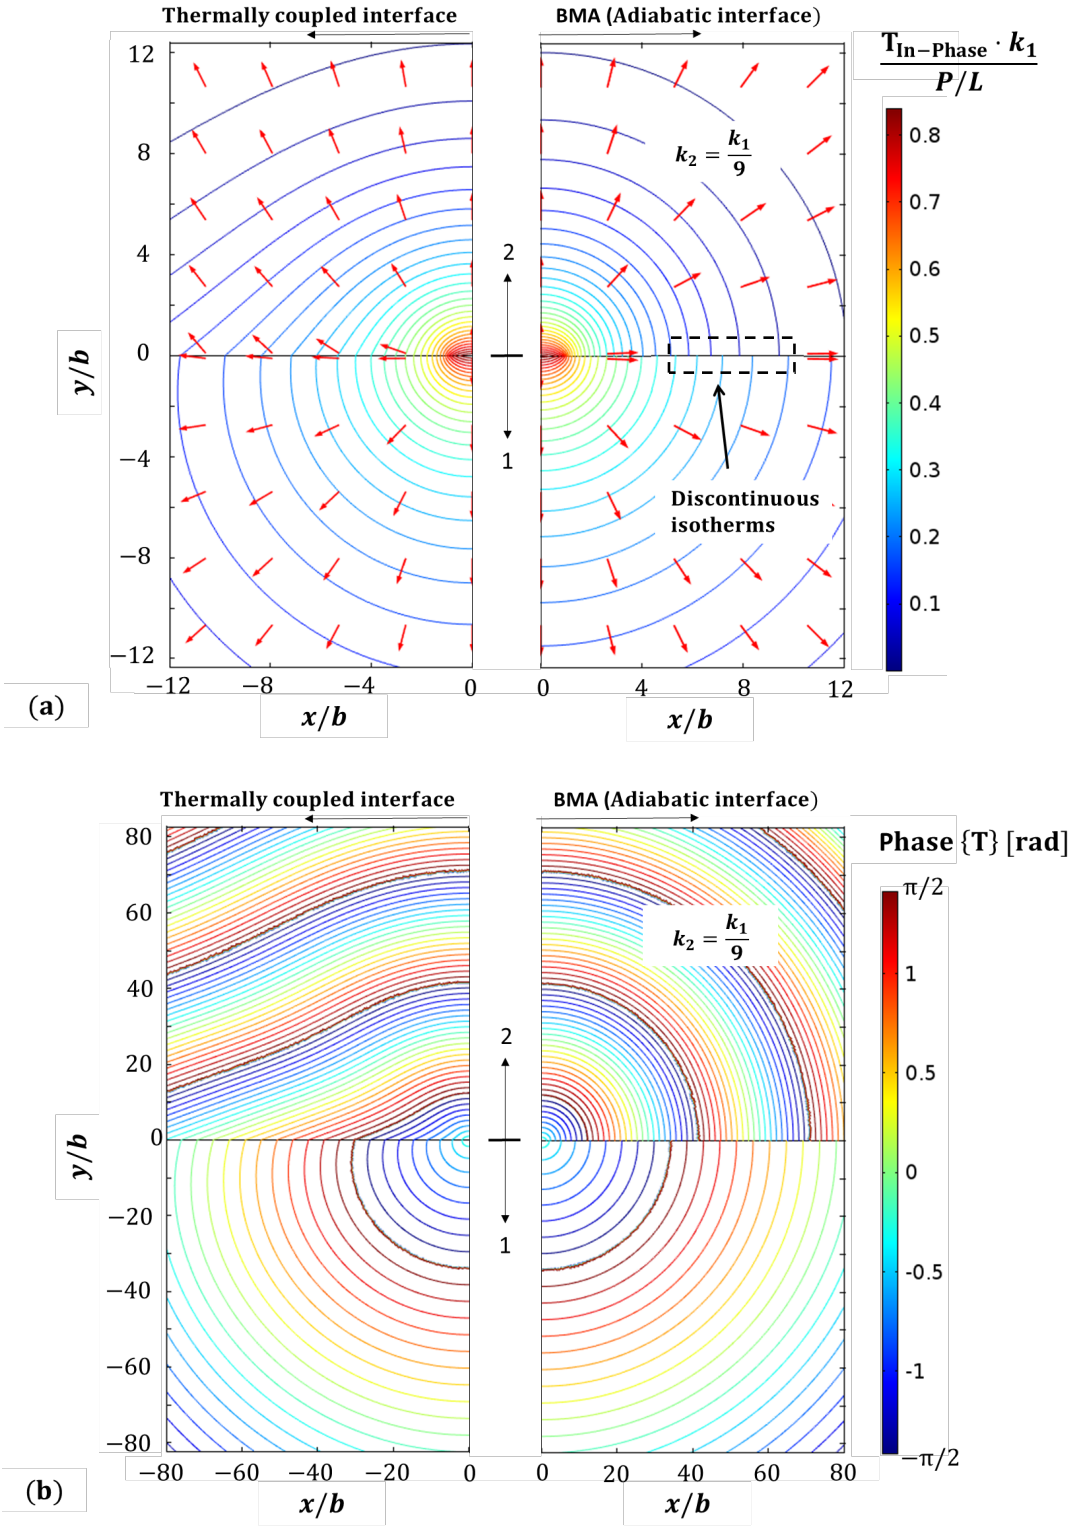

Figure 2-2: Numerical solutions of a representative bi-directional 3ω problem at low frequency, showing contours for (a) in-phase temperature and (b) phase, and comparing exact (left) and BMA (right) treatments. Red arrows in (a) are in-phase heat flux unit vectors (direction only, not magnitude) anchored at their bases. Arrows are drawn for both sides of the interface, i.e. at $y = 0^+$ and 0^-. Compared to the exact solution, the

BMA neglects heat flow through the interface, and therefore fails to match isotherm magnitude and phase. Simulation parameters: $k_2/k_1 = \alpha_2/\alpha_1 = 1/9$, $\lambda_2/b = 9.4$

Motivated by the clear discrepancies between the BMA and the exact solution seen in Figure 2-2, we now present an analytical solution to rigorously quantify these errors. The thermally coupled transfer function for bi-directional 3ω was provided without derivation by Choi et al. [33], and can be derived from the more general Feldman transfer matrix method [35] or by following the method of Shendeleva [36], [37]. For an infinitely thin heater line sandwiched between two semi-infinite isotropic materials (1,2), the exact thermally coupled transfer function Z_{1+2} is

$$Z_{1+2} = \frac{1}{\pi L}\int_0^\infty \frac{1}{\Upsilon_1 + \Upsilon_2}\frac{\sin^2(\xi b)}{(\xi b)^2}d\xi. \tag{2-6}$$

We are concerned with the error in the BMA compared to the exact solution of Eq. (2-6), where we define the error as $(c_{BMA} - c_{1+2})/c_{1+2}$, where c is any scalar, such as Re[Z]. Dimensional analysis shows that this error can be expressed as a function of k_2/k_1, α_2/α_1, and λ_1/b. Figure 2-3 shows the errors for the specific case λ_1/b =10, with extensive additional details given in Appendix 2.7.2.

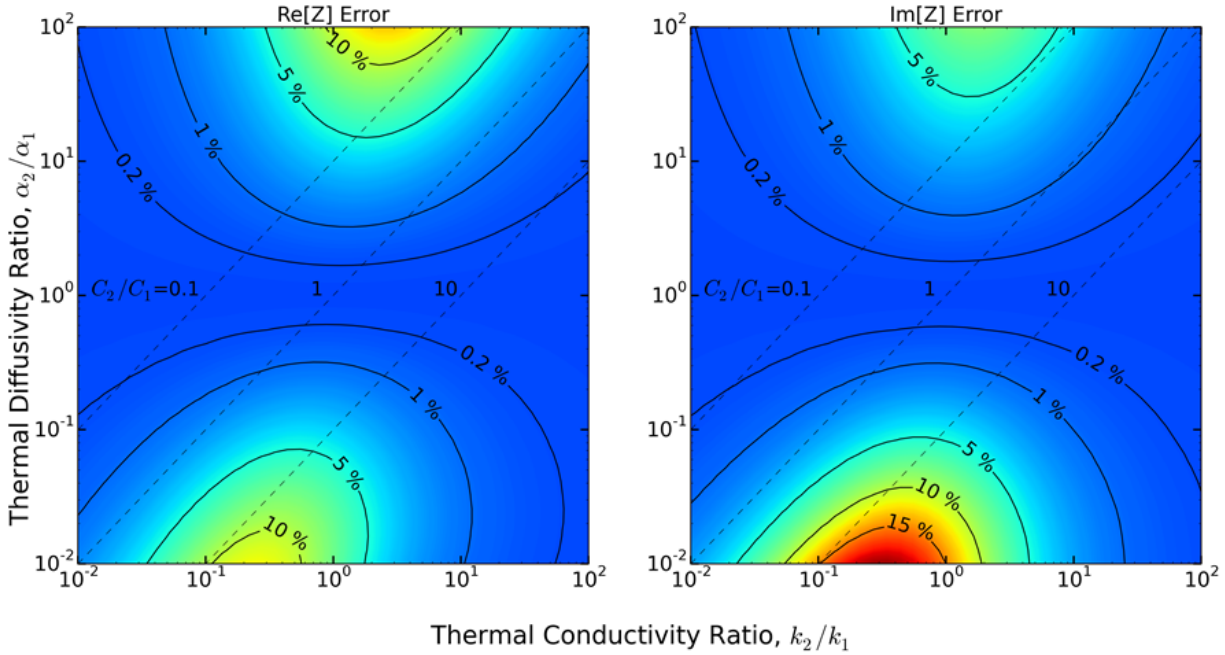

Figure 2-3: Errors (given by shading and contour lines) in the real (left) and imaginary (right) components of the BMA transfer function Z_{BMA} as compared to the exact thermally coupled solution Eq. (2-6), for the particular case λ_1/b =10. Dashed diagonal lines indicate constant volumetric heat capacity ratios.

Figure 2-3 indicates that the BMA becomes exact for matched diffusivities, as well as the limit of one material dominating the heat flow pathway. Both limits are verified analytically for

all frequency ranges in Appendix 2.7.2 and discussed here. First, when $\alpha_1 = \alpha_2$, the penetration depths are identical in both materials. The isotherms automatically align at the interface, so there is no temperature distortion or cross-interface heat flow even in the exact solution. Next, in the scenario that $k_2/k_1 \ll 1$ and $k_2/k_1 \ll C_1/C_2$, i.e., the top-left region of the plot, heat preferentially exits the heater into the higher-k material 1, and heat transfer within the lower-k material 2 becomes negligible, equivalent to an adiabatic surface condition in a single-sided 3ω configuration, in agreement with the BMA model. However, if material 2 has a much larger heat capacity than material 1, i.e., the bottom-left region of the plot, then heat flows across the boundary from the high-k material 1 to the high-C heat sink material 2 and the BMA is in error.

In the upper halves of the plots in Figure 2-3 where $\alpha_2/\alpha_1 > 1$, we have $\lambda_2 > \lambda_1 > b$, and both materials are in a cylindrical heat spreading regime. In the lower halves where $\alpha_2/\alpha_1 \ll 1$ (such that $\lambda_2 < b < \lambda_1$), however, the BMA predicts that material 2 would instead remain in a 1-D planar heating regime. In this scenario, the actual coupled solution has large isotherm shape discrepancies compared to the BMA, which causes $\mathrm{Im}[Z_{BMA}]$ to be in error by as much as 10%-20% for this frequency range and large diffusivity mismatch. Errors in $\mathrm{Re}[Z_{BMA}]$ are less sensitive to this penetration depth dependence and can reach 10% in the range considered, where $\alpha_2/\alpha_1 \ll 1$ and where $\alpha_2/\alpha_1 \gg 1$.

2.2.4 Generalized Slope Method

The generalized slope method introduced through Eq. (2-5) by taking the low frequency limit of the BMA is very convenient for analyzing experimental data. However, the discussions surrounding Figure 2-2 and Figure 2-3 suggest that the BMA can exhibit substantial inaccuracies for certain property combinations, casting doubt onto the accuracy of Eq. (2-5). Nevertheless, because extracting k_{sample} using the generalized slope method of Eq. (2-5) is much simpler than using the amplitudes of Z and numerically integrating Eq. (2-6), we now specifically consider the errors that result when applying the slope method to the exact Z_{1+2}. Previously, Choi et al. [33] considered the low-frequency limit of Z_{1+2} numerically. For the parameters considered they found that $\partial Re[Z_{1+2}]/\partial \ln(\omega)$ is simply $-1/2\pi L(k_1 + k_2)$, consistent with Eq. (2-5). We now consider the same question analytically, and demonstrate this conclusion indeed always holds in the low-frequency limit, for all property combinations.

Following the method of Cahill,[17] when λ_1 and λ_2 are much greater than b, the integral in Eq. (2-6) is dominated by values of $\xi < 1/b$. In this low-frequency limit, the integral is approximated by setting $\sin^2(\xi b)/(\xi b)^2 \to 1$ and imposing an upper limit of integration of $1/b$ to obtain

$$\mathrm{Z}_{1+2,low\,\omega} = \frac{1}{\pi L(\mathrm{k}_1 + \mathrm{k}_2)}\left[-\frac{1}{2}\ln(\omega) + \rho - i\frac{\pi}{4}\right] \tag{2-7}$$

Here ρ is a complicated but purely real function of material properties and b, and is independent of ω. Thus, applying the slope method to Eq. (2-7), which was derived from the exact Z_{1+2} coupled-domain thermal transfer function, proves that the simple generalized slope method $\partial \mathrm{Re}[Z_{BMA,low\,\omega}]/\partial \ln(\omega) = -1/2\pi L(k_1 + k_2)$ is itself exact for all thermal conductivity and diffusivity ratios in the low-frequency limit, even though $\mathrm{Re}[Z_{BMA}]$ often differs from $\mathrm{Re}[\mathrm{Z}_{1+2}]$.

Finally, applying the generalized slope method to our bi-directional tissue experiments, we use

$$k_{sample} = k_{measured} - k_{sensor} \tag{2-8}$$

where k_{sample} is our reported thermal conductivity, $k_{measured}$ is calculated by naively applying the slope method of Eq. (2-3) directly to the bi-directional measurements, and k_{sensor} is known from prior calibration. Analyzing bi-directional 3ω data with Eq. (2-8) is as straightforward as analyzing traditional one-sided 3ω data with the slope method. Because Eq. (2-8) is exact in the low-frequency limit for all property combinations, the only modeling error is in assuming the low-frequency limit itself. We evaluate this error now.

Figure 2-4 shows the percentage error in assuming the low-frequency limit (LFL) slope $-1/2\pi L(k_1 + k_2)$ at a finite frequency (FF), for the representative case $\lambda_1/b = 10$. When $\lambda_2 > \lambda_1 = 10b$, i.e. when $\alpha_2/\alpha_1 > 1$ in Figure 2-4, this error in using the LFL slope is less than 1% for conductivity ratios from 10^2 to 10^{-2}. On the other hand, the slope method errors increase when α_2/α_1 is small because the narrow-heater approximation breaks down in material 2. This effect becomes worse when $k_2/k_1 > 1$ because more of the heat flows through material 2. These finite-frequency errors in the slope method are considered further in Appendix 2.7.2.

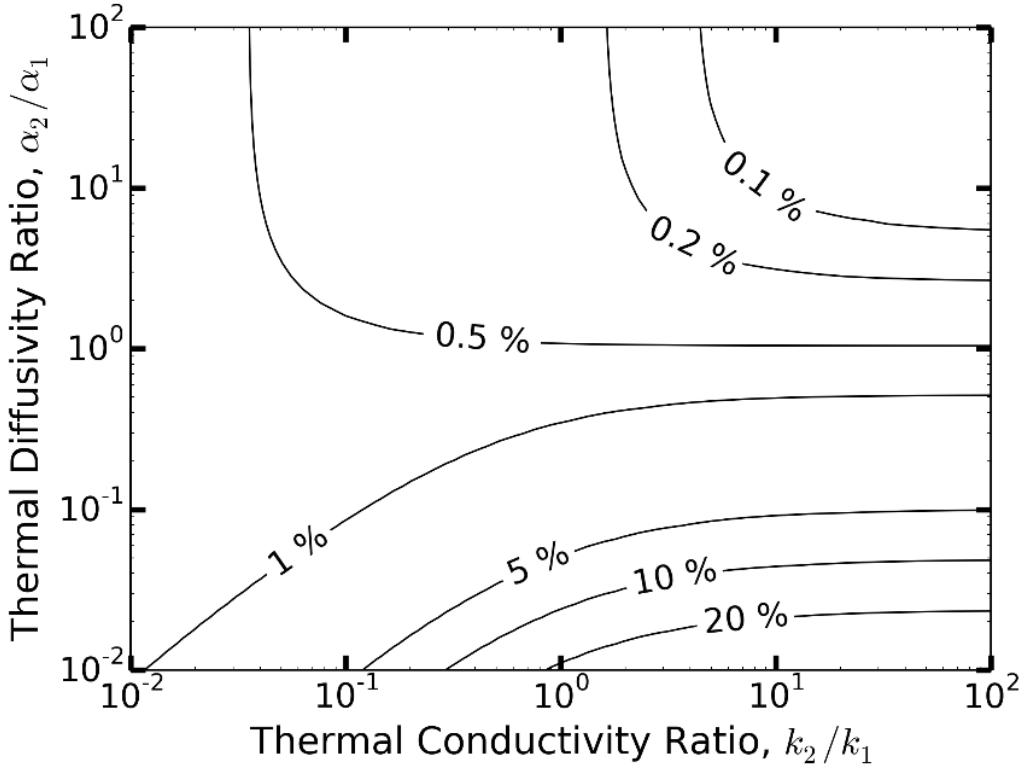

Figure 2-4: Error in assuming the low-frequency limit (LFL) of the slope for the representative finite frequency (FF) case of $\lambda_1/b = 10$. The error in slope is defined as (LFL – FF)/FF, where FF is the exact solution $\partial \mathrm{Re}[Z_{1+2}]/\partial \ln(\omega)$ and the LFL is $\partial \mathrm{Re}[Z_{1+2,low\,\omega}]/\partial \ln(\omega) = \partial \mathrm{Re}[Z_{BMA,low\,\omega}]/\partial \ln(\omega) = -1/2\pi L(k_1 + k_2)$, as can be acquired from Eq. (2-5) or Eq. (2-7).

2.2.5 Sensitivity and Experimental Considerations

The generalized slope method is recommended for bi-directional 3ω experiments, and should employ minimal k_{sensor} to maximize sensitivity to k_{sample}, provided that $\lambda_1 > b$ is still satisfied in the sensor substrate. Although low frequencies are always better for the accuracy of the slope method in semi-infinite samples, for thin samples the frequency cannot be set too low or λ_2 would run into the backside of the sample. As an alternative to the slope method, Eq. (2-7) shows that the out-of-phase voltage at a single low frequency is also convenient for measuring $(k_1 + k_2)$, since $\mathrm{Im}\left[Z_{1+2,low\ \omega}\right] = -1/4L(k_1 + k_2)$ [21]. This out-of-phase method may be preferable for real-time measurements [38], as only a single frequency must be monitored. Plots in Appendix 2.7.2 show the FF deviation of $\mathrm{Im}[Z_{1+2}]$ compared to the LFL of $-1/4L(k_1 + k_2)$. Although ||Z|| or Re[Z] could also be used to extract k_{sample}, this is not recommended since Eq. (2-7) shows that this calculation for k_{sample} now also depends on C, b, and thermal contact resistances, and their corresponding uncertainties.

Many bi-directional 3ω device geometries require a thin dielectric layer between the heater line and the supported sample, both for electrical isolation and to prevent damage to the heater line [30], [31], [38]. Additionally, there is inevitably some finite thermal contact resistance between sample and sensor, although this can be minimized by using a sample that is positioned while hydrated and malleable, and/or by applying clamping pressure [31]. A detailed discussion of dielectric layer and thermal contact resistance effects on the slope method is given in Appendix 2.7.3. The important conclusion is that for the experimental conditions presented below, the total error in using Eq. (2-8) is no more than 3.3% (this worst case error is for non-frozen mouse liver on glass: $k_2/k_1 = 0.5$, $\alpha_2/\alpha_1 = 0.27$, $\lambda_2/b = 8.5$, with a 1 μm thick polystyrene dielectric layer). Without the dielectric layer, the error in using Eq. (2-8) for the same conditions would be 0.8%, so most of the error in the slope method for our experiments can be attributed to the dielectric layer rather than approximating FF as the LFL.

2.2.6 Heat Capacity Measurements

While the focus of this work is on measuring thermal conductivity, we note that this and similar devices can also be used to measure volumetric heat capacity [21], [23], [31], [33], [38], which is important for transient procedures. Unfortunately, there is no tidy analytical expression for C analogous to the simple slope method for k. Therefore, C is typically determined via least-squares fitting of experimental data to a nonlinear analytical [21], [23], [31], [33] or finite element [38] model. As a consequence, fits for C inherit propagated uncertainties from all model input parameters, including the heater line geometry, thermal contact resistance, and the properties and geometry of the sensor substrate and dielectric layer. In general, numerical integration of Eq. (C2-16) from Appendix 2.7.3 can be used for such fitting. The magnitude of Z has very weak sensitivity (typically < 0.05) to C in the low frequency regimes typically used to measure k, and so it is challenging to measure C in this cylindrical heat spreading regime. Alternatively, sensitivity to C increases significantly (typically ~0.25) at higher frequencies, when the conduction is one dimensional along y, provided that λ in any present dielectric layer is still much larger than its thickness. Because Z here becomes sensitive to the thermal effusivity (Eq. (B2-15)), k must also be measured in order to back out C. As an example of these challenges, representative calculations for our system at high frequency show that the uncertainty in k propagates directly into uncertainty in C, and the uncertainty in heater line width is magnified approximately 5-fold in its effect on C.

2.3 Experimental Apparatus and Calibration

The "sensor" (Figure 2-1(a)) is based on the traditional one-sided 3ω method. For each sensor, a gold heater line with typical dimensions 1500 μm in length (L) between voltage probes, 20 μm width ($2b$), and 0.2 μm thickness (plus 10 nm Cr adhesion layer), was deposited on a 1 mm thick glass microscope slide using standard microfabrication techniques. After attaching 76 μm diameter copper wires using silver epoxy, the wiring and heater line were coated in a dielectric polystyrene layer approximately 1 μm thick. This was accomplished by preparing a solution of polystyrene in toluene with concentration 15 mg / mL. Several drops of the solution were applied to the entire surface after wiring, which was then held vertically to allow excess solution to run off, after which the toluene evaporated in 5 – 10 seconds leaving behind a continuous film of polystyrene. The resulting dielectric film thickness, electrical insulation, and repeatability were verified on practice glass slides using a stylus profilometer and digital multi-meter.

Each sensor's temperature dependent electrical resistance was calibrated to extract dR/dT for Eq. (2-3), and to enable measurement of the sample T using the heater line. Typical $R(T)$ calibration parameters were: 6 temperatures from -20 ºC to 30 ºC; 10 current values from 5 mA to 25 mA with R extrapolated to I=0; resistance ~10 Ω at room temperature; ± 1% (std) uncertainty on dR/dT fit; ± 0.06% (std) uncertainty on $R(T$=0°C) fit. The dR/dT fit was one of the primary sources of uncertainty in the final measured thermal conductivity values.

Each sensor was then calibrated for $k_{sensor}(T)$ using the traditional 3ω method. When interpolating between calibrated points later, a linear $k_{sensor}(T)$ fit was used for most sensors, while a 4th order polynomial fit was used for a few low-temperature sensors (-100 ºC < T < -20 °C) to accommodate deviations of $k_{glass}(T)$ from linearity. Typical $k(T)$ calibration parameters were: 12 temperatures spaced from -10 ºC to 30 ºC (regular sensors), or 40 temperatures spaced from -110 ºC to 30 ºC (cryogenic); 20 fitted frequency points spaced logarithmically from 1 Hz to 25 Hz, measurement current 20 mA; ± 6% (std) uncertainty on dk/dT fit; ± 0.2% (std) uncertainty on $k(T$=0°C) fit.

The 1ω heating current was produced by a custom V-to-I circuit, driven by the reference output of the same lock-in amplifier used to measure $V_{3\omega}$. Background subtraction of the $V_{1\omega}$ component of the signal was implemented using a multiplying digital-to-analog converter.[17] A detailed uncertainty analysis for this bi-directional 3ω technique is given in Appendix 2.7.1.

2.4 Measurement Results

2.4.1 Control Experiments: Water and Ice

The system was first validated using ice and liquid water as controls. Some ice samples were made from deionized water that had been boiled for 5 minutes, and others from commercial distilled water (UltraPure, Life Technologies). Thin ice and water samples (< 1 mm) were prepared by injecting liquid water into a sheet-like gap of known thickness over the heater line. Ice samples were then frozen in place. To control the gap thickness and minimize evaporation, rigid transparent plastic sheets were placed over the heater line, peripherally supported by shims of known thickness. The accuracy of the thickness of < 1 mm water and ice samples is believed

to be ± 0.02 mm, and that of ≥ 1 mm to be ± 0.2 mm. The thickness of samples ≥ 1 mm was measured via visual comparison with stack-ups of shims of known thicknesses, without additional constraining top surfaces.

All samples, including liver, were large enough to extend laterally by at least 0.5 mm around the perimeter of the heater line to ensure the longest λ of ~100 μm did not reach the sample boundaries in any direction. Samples not in controlled gaps were held in place by their own weight and the intrinsic adhesion of these moist samples. In order to minimize thermal contact resistances, ice and frozen liver were first deposited on the sensor while in their non-frozen state, so they could form intimate thermal contact and conform to the surface of the sensor, before being frozen in place from the bottom up. Literature values for representative solid-liquid interfaces[39] give thermal contact resistances 2 to 3 orders of magnitude smaller than the effective thermal resistance of the dielectric layer used. These values are also assumed not to change dramatically after freezing in place. Therefore, sample-sensor thermal contact resistance effects were neglected in lieu of the more dominant effect of the dielectric layer, considered in detail in Appendix 2.7.3. Several liver and ice samples were visually inspected from underneath before, during, and after measurement, to check for signs of trapped air pockets, delamination, or thermal-expansion-induced shifting, none of which were observed. During the freezing and measurement process of a few samples, cracking was observed both visually and through large jumps in the raw voltage data. Such samples were excluded from the data set.

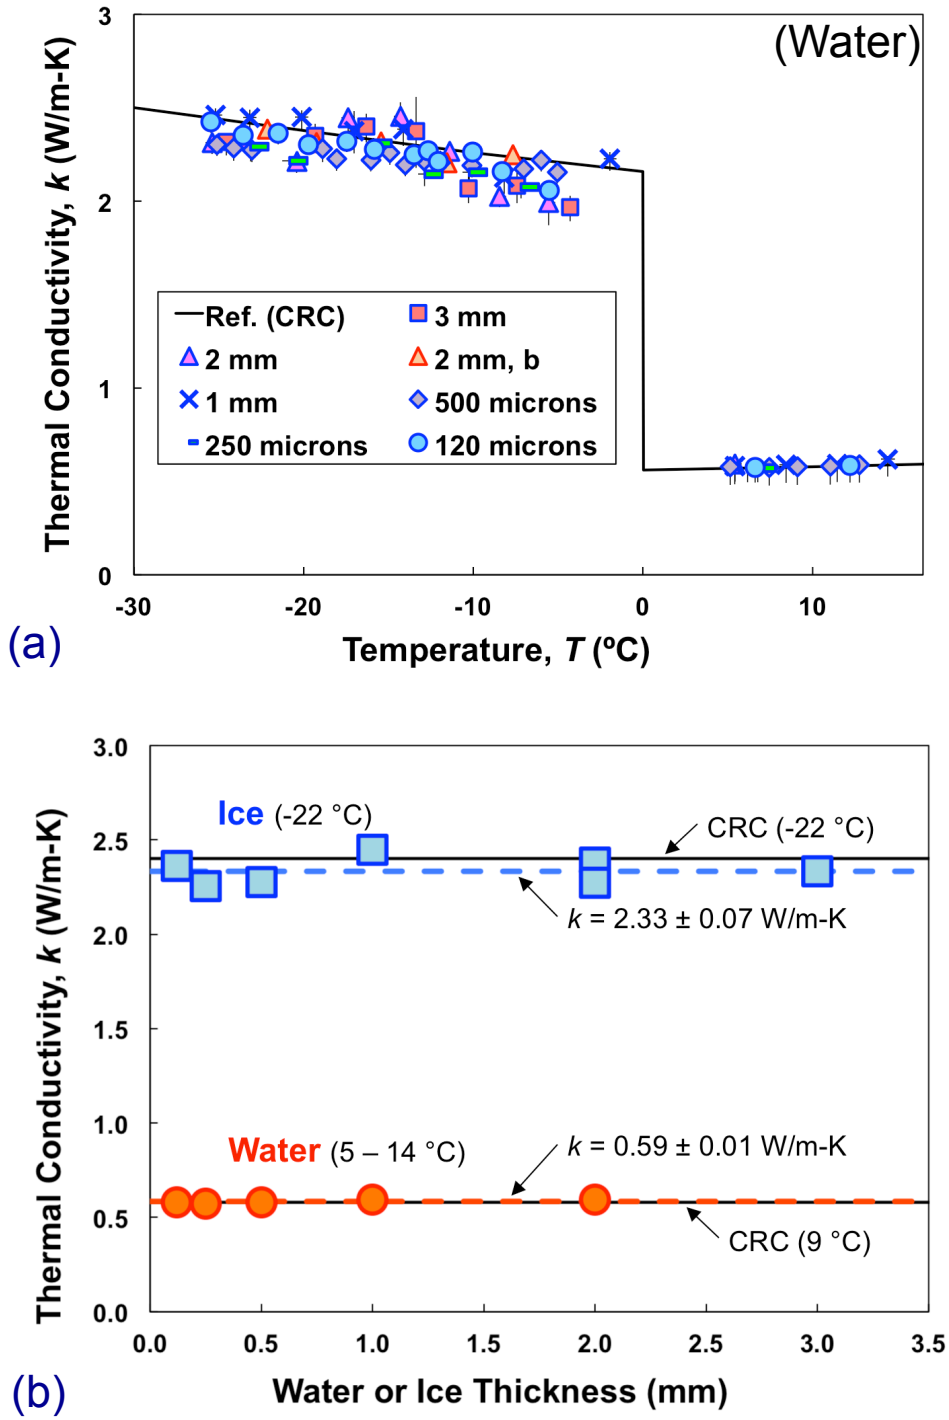

Figure 2-5: Control measurements for ice and water, compared to reference values from.[40] (a) Temperature-dependent measurements. Marker shape corresponds to sample thickness and marker edge color corresponds to the sensor used for the measurement. Error bars (std) represent the combined effects of measurement uncertainty and estimated modeling error, with full details given in Appendix 2.7.1. In most cases, the error bars are smaller than the size of the plotted points. (b) Subset of the data in (a) interpolated to -22 °C for ice and averaged over all T for water. Each point represents a unique sample. No trend with sample thickness is observed. Dashed lines indicate average thermal conductivity with the listed value, and black lines are CRC literature reference data [40].

The results for ice and water are given in Figure 2-5, representing measurements using 2 sensors and 10 samples. The measures of accuracy and precision for these measurements, along with those for liver discussed below, are given in Table 2-1. Here the rms error is measured with respect to the literature data [40], and we define the scatter as the rms deviation of the data from a simple model (linear for suprazero, quadratic for subzero) fit to the data. Full detail regarding the error analysis is given in Appendix 2.7.1.

Figure 2-5(a) shows the measured temperature-dependent k of ice and water. Most of the rms error can be explained by the scatter, and both confirm very good accuracy (< 5% error for ice, < 2% for water). The measured values for ice are slightly below the literature values, which might be explained by impurities, dissolved gases, or grain boundaries caused by freezing from many simultaneous nucleation sites.

Figure 2-5(b) shows a subset of the data from panel (a) re-plotted as a function of sample thickness. To focus on the thickness effect and eliminate the temperature dependence of k_{ice}, we chose to compare all samples at the single temperature of -22 °C. For each ice sample the value of k(-22 °C) was determined by interpolation from the two closest bracketing T points. Because k of water has a very weak temperature dependence, for each water sample we simply plot the average of all $T > 0$ °C points for that sample. Figure 2-5(b) shows that there is no discernible correlation between measured thermal conductivity and sample thickness, for either water or ice. Thus, Figure 2-5 demonstrates the precision, accuracy, and reusability of the bi-directional 3ω sensors, as well as their ability to measure a wide range of sample thicknesses.

Table 2-1: Accuracy and Precision Analysis of Control and Liver Data, at the Level of One std

	Water		Liver	
	Frozen	Non-Frozen	Frozen	Non-Frozen
Experimental uncertainty	2.0%	2.3%	4.9%	5.9%
Experimental scatter	3.6%	1.5%	7.4%	5.7%
Literature scatter	N/A	N/A	14.5%	14.1%
rms Error or deviation	4.6%	1.9%	8.3%	7.4%

2.4.2 Measurements of Mouse Liver

Fresh mouse liver was harvested opportunistically, transported in phosphate buffered saline (PBS) solution, and stored in PBS at 2 °C to 5 °C when not being handled. Samples ≥ 1 mm thick were prepared by manually slicing with a scalpel, and their thicknesses were measured using shim stack-ups (± 0.2 mm). The 0.5 mm sample was prepared using a microtome blade to manually slice across the top of a recessed surface of known depth, with estimated accuracy of ± 0.1 mm. The 100 μm sample was prepared using a Leica CM3050-S cryotome, with accuracy

unspecified but believed to be ± 5 μm or better. The 100 μm sample was necessarily first frozen for cryosectioning and then re-thawed for mounting on the sensor. A layer of plastic wrap was used to cover all non-frozen liver samples and all < 1 mm samples during measurement, to minimize sublimation and evaporative losses. Additional sample preparation and measurement precautions were followed as described for the control experiments above.

A separate study indicated that storage in PBS could result in up to a 10% mass gain over a few days. Most samples were measured within 10 hours of harvesting, with the longest storage time being 33 hours for one of the 1.5 mm liver samples which was stored as bulk (~ 20 mm) and cut to size immediately before measurement, as with all samples. The storage time had no noticeable effect on the k values of these samples. Another study examined the effect of the liver surface being wet vs. dry before mounting, to investigate the adhesion potential as well as possible effects of a thin water layer forming between the sample and sensor. Two sets of measurements differing only in one key protocol step were compared (thickness=2 mm, N=3 of each). The first set was a control following our standard "wet liver" protocol, whereby the samples were placed on the sensor as soon as practical after removal from PBS and slicing. In the second protocol, after slicing, the bottom face of the liver sample was first gently blotted dry before placement on the sensor. In all cases the sensor itself is dry before receiving the sample. The results, shown most clearly in the inset of Figure 2-6(a) (compare green vs. yellow triangles), exhibit no discernible effect on k. This result further confirms the non-frozen liver's ability to form intimate thermal contact with the sensor even in a nominally dry state. Furthermore, the similarity of k values from wet and dry protocols implies that any possible intervening water layer from the wet protocol would have to be substantially thinner than the thermal penetration depths of the measurement (typical λ from 50 - 100 μm), since pure water would bias the apparent k_{sample} upwards.

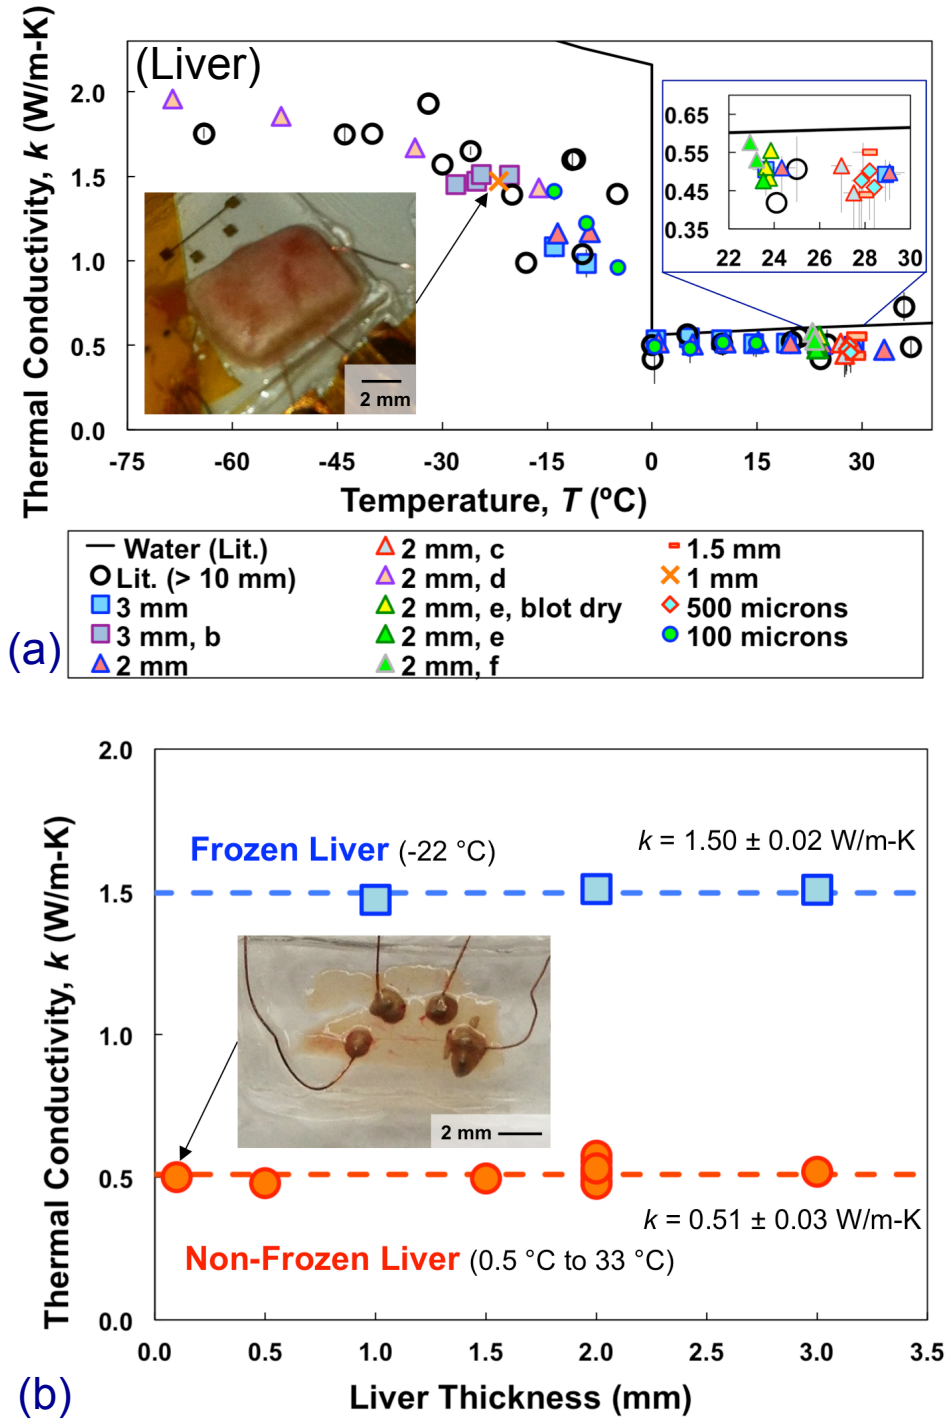

Figure 2-6: Mouse liver thermal conductivity measurements. (a) Temperature-dependent mouse liver measurements (filled, colored points). Marker shape corresponds to sample thickness and marker edge color corresponds to the sensor used for the measurement. Reference liver data (empty circles) are from [12]–[16], and comprise data from pig, cow, human, dog, and rabbit livers. Water reference values are from [40]. Data represent 16 samples measured using 6 sensors in two different laboratories. Error bars (std) represent the combined effects of measurement uncertainty and estimated modeling error, with full details given in Appendix 2.7.1. (b) Subset of the data in (a) interpolated where possible to -22 °C for frozen liver, and averaged over all *T* for non-frozen liver. Each point represents one sample. No trend with sample thickness is observed. Dashed lines indicate average thermal conductivity. Insets: (a) (left) photograph of 1 mm thick

frozen liver sample on a sensor; (right) detail of non-frozen data; (b) photograph of 100 micron thick non-frozen liver on a sensor.

The results for mouse liver are presented in Figure 2-6. These data represent 16 samples measured using 6 sensors in 2 laboratories. The measures of accuracy and precision for these measurements can again be found in Table 2-1. Scatter is defined the same as for water and ice. The error cannot be directly assessed since these are the first measurements of mouse liver, and indeed the most complete data set for liver from any organism. Nevertheless we estimated the overall deviation between our measurements and literature by calculating the rms difference between models fitted to our data and to the literature data, over the temperature range of the data. Further details are given in Appendix 2.7.1.

Figure 2-6(a) shows that our measurements of mouse liver fall within the scatter of available literature data for k_{Liver} of other organisms: pig, cow, human, dog, and rabbit, all > 10 mm thick and measured using traditional macroscopic techniques [12]–[16]. The scatters of our measured liver data are comparable to their estimated uncertainties, and reflect the greater inherent sample-to-sample variability as compared to water. For all k measurements, the dominant sources of uncertainty were found to be the $\partial V_{3\omega,in-phase}/\partial \ln(\omega)$ slope and the sensor's calibrated dR/dT slope, with the latter contributing the most in the case of liver. Table 2-1 shows that the liver measurements had higher uncertainty than the water measurements, but this is simply because the sensors used for the latter were calibrated more accurately, and does not indicate any fundamental issue with biological samples. The scatter of our liver data is still two to three times tighter than the scatter of the literature data. This is attributed to greater inter-organismal variation as compared to intra-organismal variation, and to the superior precision of our technique. The rms deviations between our data and the literature are actually smaller than the internal scatter of the literature dataset itself, while also being slightly larger than the internal scatter of our measurements. This means that comparatively, the best-fit model for the literature data has closer agreement with our data than it does with the literature data itself, again confirming the accuracy and repeatability of the present 3ω method.

Figure 2-6(b) is plotted in the same manner as described for Figure 2-5(b), and confirms that these liver results are essentially independent of thickness over the range studied, from 3 mm down to 100 μm. This is consistent with the design of the experiments, as the frequencies were chosen to probe only the first ~100 μm of each sample. Therefore, this technique is agnostic to the full size of the sample, and could accurately measure the thermal conductivity of very thin tissues such as heart valves (1 - 2 mm), phrenic nerve (0.5 - 1 mm), cornea (0.5 mm), and fascia (0.1 mm).

2.5 Conclusions

These bi-directional 3ω sensors can accurately and precisely measure k of biological tissues and water-based samples at least two orders of magnitude thinner than traditional techniques for biological samples [8]–[16]. Measurements of mouse liver showed scatter of around 6% - 7% and excellent agreement with available literature for other organisms, while control measurements on water and ice prove ultimate accuracy limits better than 2% - 5%. In most previous implementations of the bi-directional 3ω method, the experimental uncertainties included non-negligible contributions from the modeling errors caused by the boundary mismatch approximation (BMA) [21]–[24] and generalized slope methods. Both of these

idealizations are imperfect but common and very convenient, so we have for the first time quantified their errors, and proven analytically that the BMA does become exact in the low frequency limit.

The minimum tissue thickness of 100 μm measured above was limited by sample preparation rather than fundamental thermal issues. In principle much thinner layers could be measured (see also Appendix 2.7.3), and indeed for inorganic solids the traditional single-sided 3ω methods are routinely used to measure films even thinner than 1 μm [19], [20]. The prospects for such extreme thinness in biological systems will be limited by the challenges in preparing thin tissue sheets, the difficulties in obtaining excellent thermal contact when assembling two initially distinct materials (although this may still be achievable [41]), and ultimately by the granularity of individual cells themselves [25].

The 3ω sensors used in this work are reusable, and biological samples can be changed in typically 10 minutes or less. This technique also inherits the traditional 3ω method's robustness against thermal contact resistance and parasitic heat loss issues. Finally, as a microfabricated device this system lends itself to miniaturization, and similar sensors may be amenable to *in-vivo* measurements and devices in the future.

2.6 Suggestions For Future Work

While this dissertation has demonstrated that the bi-directional 3ω method can be easily implemented used to measure soft, wet, and chemically sensitive samples down to 100 μm thicknesses, there is still room for improvement and future work to develop this technique further. Here, we discuss some of these possible avenues.

Many biological tissues are anisotropic in their structure, such as muscle tissues and fascia. Recent work has shown how the traditional 3ω method can be extended to measure an arbitrarily oriented anisotropic thermal conductivity tensor [42]. We therefore recommend that this be applied to the current work, to generalize our sensor so as to be able to measure the full anisotropic k tensor of biological tissues.

While we have demonstrated the efficacy of the scientific principles of the technique, this was all done with bulky "one-off" prototype lab equipment. In order to be of maximal benefit to the biomedical community, we encourage someone to commercialize this technique to create a more modular and self-contained version of the apparatus with an easy user interface, so that medical researchers and hospital staff around the world can use this to quickly and easily measure k of various tissues. Personalized calibrated k measurements could greatly enhance the efficacy and safety of biomedical thermal therapies. Along a similar vein, it could be quite beneficial if the technology could be miniaturized so that *in vivo* measurements were possible.

Our measurements, while as close to *in vivo* as possible, were still not of tissue within a living organism. One key difference here is therefore the lack of blood perfusion and its potential influence on heat transfer and hence the "effective k" that most accurately models the tissue. We therefore recommend miniaturization of our device and *in vivo* measurements where possible. Recent work has made very promising progress toward *in vivo* measurements, accounting for advective thermal effects, real-time measurements, and moving toward miniaturization and commercialization possibilities of this technique as a biomedical device [43].

2.7 Appendices

2.7.1 Appendix A - Uncertainty And Error Analysis

In the usual way, the experimental uncertainties were assumed to be small, normally distributed, and were propagated orthogonally. We also accounted for correlations between dependent parameters from the same fit based on their calculated covariance. The fits for $\partial V_{3\omega,in-phase}/\partial \ln(\omega)$, $dk/d\mathrm{T}$, and $dR/d\mathrm{T}$ were calculated using least-square regression weighted by the inverse of the measured standard deviation of each data point, using the Levenberg-Marquardt method employed via Scipy's optimize.curve_fit module. Uncertainties (std) of fit parameters were calculated from the variance-covariance matrix of the fit. For all k measurements, the dominant sources of uncertainty were found to be the $\partial V_{3\omega,in-phase}/\partial \ln(\omega)$ slope and the $dR/d\mathrm{T}$ calibration, with roughly equal contributions in general.

After calculating the total propagated uncertainty (± std) in k for each measurement, the modeling error was also calculated by numerically evaluating the full coupled, bi-directional geometry solution of Eq. (C2-16) from Appendix 2.7.3, which also accounts for the dielectric layer. All samples were non-frozen and hydrated when originally placed on the sensor and therefore we assumed they were in intimate thermal contact with the dielectric layer and had negligible thermal contact resistance. This assumption is supported by the low values of thermal contact resistance available in the literature [39] for interfaces involving water. Furthermore, the nature of the generalized slope method eliminates error due to purely real thermal contact resistances which have no thermal capacitive effects, when sufficiently within the low frequency limit (shown in Figure 2-8 in Appendix 2.7.3) as was the case for our measurements. The most conservative parameter values are given in Table 2-2 of Appendix 2.7.3, with assumed dielectric properties also given in Appendix 2.7.3, yielding a worst-case modeling error of 3.3%. This error was always positive, indicating that that the generalized slope method applied to a finite frequency experiment will always over-estimate the thermal conductivity for parameters used in our experiments. As such, the error bars in Figure 2-5 and Figure 2-6 are asymmetric, with the - error bar exceeding the + error bar. The modeling error was combined in quadrature with the total propagated uncertainty (± std) to calculate the total error bar.

In order to quantify the precision and accuracy of our data, scatter and deviation metrics were calculated. Because a physics-based model for the temperature-dependent thermal conductivity of liver is not available, we used simple empirical models instead: a linear model was used for suprazero data, and a quadratic model for subzero data. For each experimental dataset, the data's scatter among itself was calculated as the rms relative deviation of the data from its own corresponding fitted model,

$$scatter = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(\frac{k_{meas,i} - k_{model}(T_i)}{k_{model}(T_i)}\right)^2} \tag{A2-9}$$

where i indexes each measured data point, N is the total number of measurements in the dataset, $k_{meas,i}$ is an experimentally measured thermal conductivity, T the temperature, and $k_{model}(T)$ is the temperature-dependent thermal conductivity model (linear for suprazero, quadratic for subzero)

fit to the dataset. To calculate the rms error in our control ice and water measurements, Eq. (A2-9) was applied to our measured data, while k_{model} was fit to the literature data [40]. Because there is no available $k(T)$ literature specifically for mouse liver, a more general agreement metric was devised to compare our data to a collection of $k(T)$ liver data reported for various other organisms [12]–[16]. This was calculated as the rms deviation between a model fit to our data and a model fit to the literature data, integrated over the temperature range spanned by the union of the two data sets,

$$deviation = \sqrt{\frac{1}{T_{max} - T_{min}} \int_{T_{min}}^{T_{max}} \left(\frac{k_{model,exp}(T) - k_{model,lit}(T)}{k_{model,lit}(T)} \right)^2 dT} \qquad \text{(A2-10)}$$

where $k_{model,exp}(T)$ corresponds to our experiments, and $k_{model,lit}(T)$ to the literature. The results of all scatter, error, and deviation analyses for both liver and control samples are summarized in Table 2-1 in the main text.

The scatter of each data set gives a measure of its precision. A comparison of the scatter of our data against its rms error/deviation with respect to the literature gives a measure of the accuracy. The calculated error/deviation gives the combined effect of the scatter (imprecision) and inaccuracy.

2.7.2 Appendix B - Errors in bi-directional data analysis

Figure 2-7 shows further error analysis for two semi-infinite materials in the bi-directional geometry. Each column displays a different heating frequency, and the individual subplot axes are the conductivity and diffusivity ratios. Within the middle column, the topmost, 3rd, and 4th panels are identical to Figure 2-4 and Figure 2-3 in the main text.

The top row in Figure 2-7, like Figure 2-4, shows the percentage error of the low-frequency limit (LFL) compared to the exact finite-frequency (FF) slope of Eq. (2-6). These calculations show that the slope of the numerically integrated Eq. (2-6) is within 1% of the much simpler low-frequency Eq. (2-5) as long as λ_1/b and λ_2/b are both > 10. The second row in Figure 2-7 shows the percentage error in assuming the imaginary amplitude LFL, $-1/4L(k_1 + k_2)$, as compared to the exact FF Im$[Z_{1+2}]$ of Eq. (2-6). The error is defined as (LFL –FF)/FF for both rows. Comparing rows 1 and 2, we see that the slope method has smaller finite-frequency errors than the out-of-phase method in the low and medium frequency cases.

The remaining rows of Figure 2-7 show the percentage errors in the BMA predictions of Re$[Z_{BMA}]$, Im$[Z_{BMA}]$, magnitude $|Z_{BMA}|$, and phase shift angle θ_{BMA} , all compared to the exact Z_{1+2} obtained by numerical integration of Eq. (2-6). From Figure 2-7 we note several trends. First, in the low yet finite frequency scenario of the left-hand column, Im$[Z_{BMA}]$ has a much smaller error than Re$[Z_{BMA}]$, as predicted by Eq. (2-7). Second, for these three frequency cases, Z_{BMA} becomes exact for equal diffusivities and/or large conductivity mismatch. Lastly, the high frequency behavior in the rightmost column indicates that errors in Z_{BMA} are reduced when penetration depths in both materials are of the order b or smaller. In the remainder of this appendix, we analytically study the thermal property ratios where the BMA becomes exact, and also confirm that the BMA is exact for all thermal property ratios in the high-frequency limit.

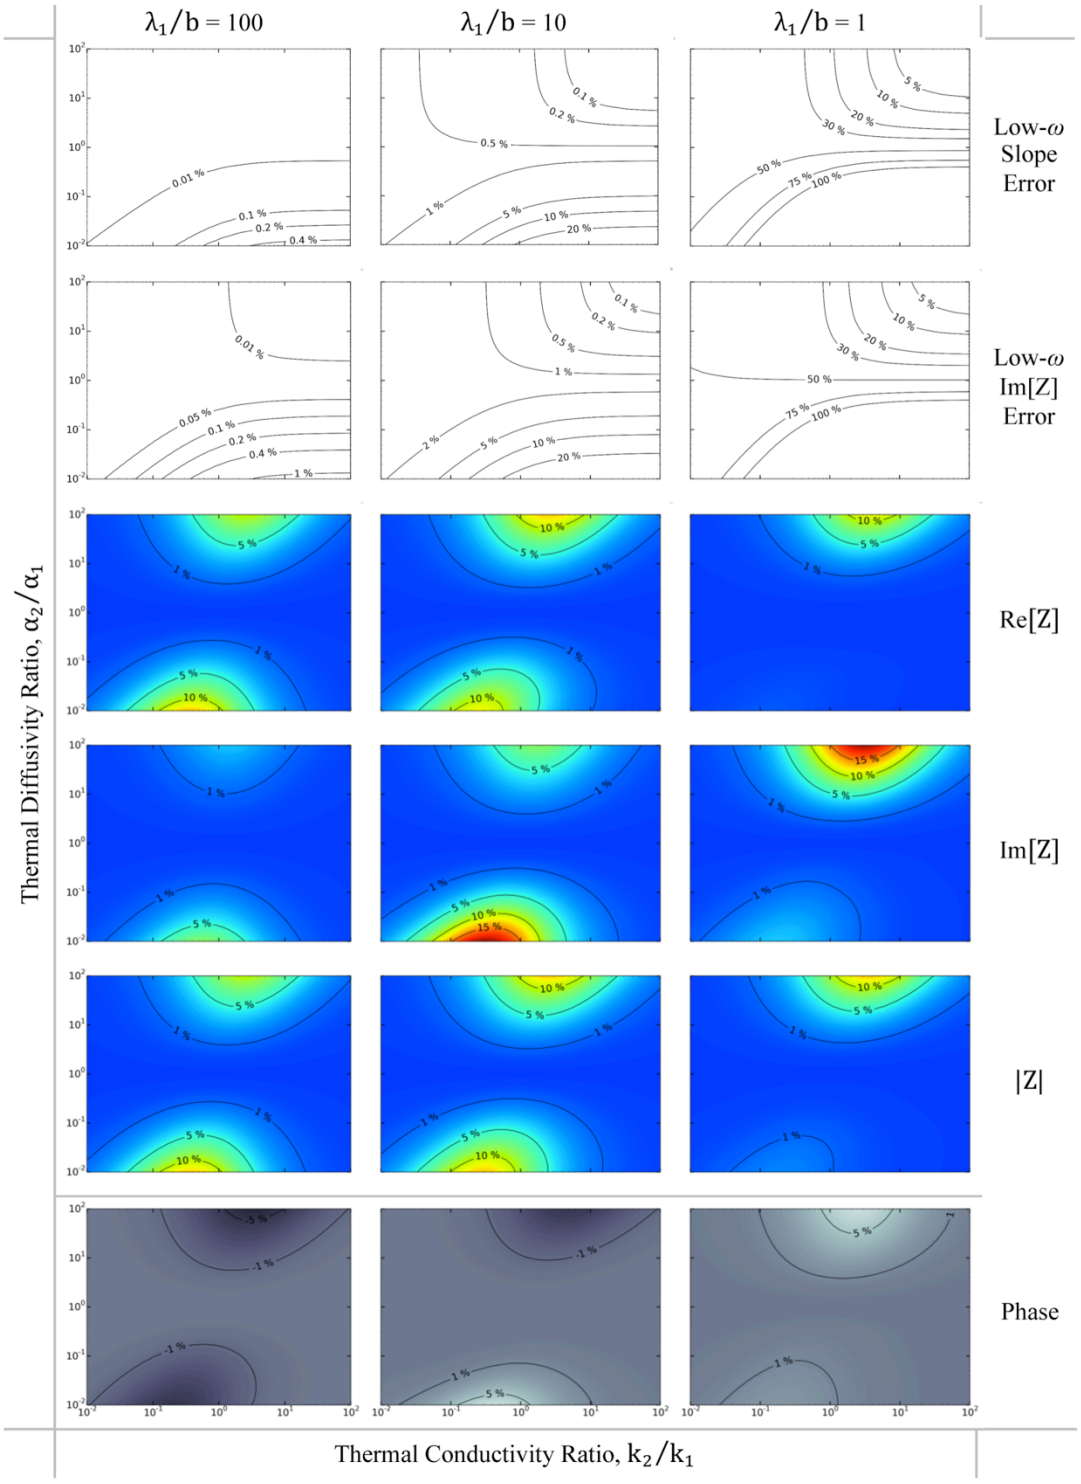

Figure 2-7: Errors in two semi-infinite bi-directional geometry analysis methods. Rows 1-2 give the finite-frequency errors arising from using the respective low-frequency limits for the generalized slope method ($\partial \mathrm{Re}[Z_{1+2}]/\partial \ln \omega$), and the generalized out-of-phase method ($\mathrm{Im}[Z_{1+2}]$), at three frequencies. Rows 3-6 give the errors in Z_{BMA} compared to the full coupled-domain solution, Z_{1+2}.

First, we will analytically show that the BMA is exact for the specific case of $\alpha_1 = \alpha_2$, and for the case of $k_1/k_2 \ll 1$ and $k_1/k_2 \ll C_2/C_1$, for all values of λ_1/b. We rewrite the exact result of Eq. (2-6) as

$$Z_{1+2} = \frac{1}{\pi L k_2}\int_0^\infty \frac{1}{\left(\frac{k_1}{k_2}\right)\sqrt{\xi^2 + i\frac{2\omega}{\alpha_1}} + \sqrt{\xi^2 + i\frac{2\omega}{\alpha_2}}}\frac{\sin^2(\xi b)}{(\xi b)^2}d\xi. \tag{B2-11}$$

With $\alpha_1 = \alpha_2 = \alpha$, this simplifies to

$$\begin{aligned} \mathrm{Z}_{1+2} &= \frac{1}{\pi L(k_1+k_2)}\int_0^\infty \frac{1}{\sqrt{\xi^2 + i\frac{2\omega}{\alpha}}}\frac{\sin^2(\xi b)}{(\xi b)^2}d\xi \\ &= \frac{F(\alpha)}{\pi L(k_1+k_2)} = \left(\frac{\pi L k_1}{F(\alpha_1)} + \frac{\pi L k_2}{F(\alpha_2)}\right)^{-1} = \mathrm{Z}_{BMA}, \end{aligned} \tag{B2-12}$$

where the notation $F(\alpha)$ is shorthand for the value of the integral, and the equality with Z_{BMA} follows directly from Eqs. (3-2) and (2-4). Similarly, to study the behavior for large conductivity mismatches, we define a new variable of integration $\chi = \xi/q_2$, where $q_j = \sqrt{i2\omega/\alpha_j}$, and rewrite Eq. (B2-11) as

$$Z_{1+2} = \frac{1}{\pi L k_2}\int_0^\infty \frac{1}{\left(\frac{k_1}{k_2}\right)\sqrt{\chi^2 + \frac{\alpha_2}{\alpha_1}} + \sqrt{\chi^2 + 1}}\mu d\chi, \tag{B2-13}$$

where $\mu = \sin^2(\chi b)/(\chi b q_2)^2$ is independent of material 1 properties, and we allow for $\alpha_1 \neq \alpha_2$. When $k_1/k_2 \ll 1$ and $k_1/k_2 \ll C_2/C_1$, the terms in the denominator of Eq. (B2-13) containing material 1 properties become small compared to $\sqrt{\chi^2+1}$, and the one-sided solution for material 2 from Eq. (3-2) is recovered. This same limit is also apparent when the BMA is written as

$$\frac{Z_2}{Z_{\mathrm{BMA}}} = \frac{\int_0^\infty \frac{1}{\sqrt{\chi^2+1}}\mu d\chi}{\int_0^\infty \frac{1}{\left(\frac{k_1}{k_2}\right)\sqrt{\chi^2 + \frac{\alpha_2}{\alpha_1}}}\mu d\chi} + 1. \tag{B2-14}$$

When $k_1/k_2 \ll 1$ and $k_1/k_2 \ll C_2/C_1$, the integral in the denominator of Eq. (B2-14) becomes large and Z_{BMA} approaches Z_2. It is intuitive that when $Z_1 \gg Z_2$, the sum of the two impedances in parallel approaches Z_2. The validity of the BMA in scenarios of equal diffusivity and a quasi-adiabatic interface can be observed for all frequency ranges in Figure 2-7. The quasi-adiabatic requirement $k_1/k_2 \ll C_2/C_1$ explains why the error is asymmetric around $k_1/k_2 = 1$ in Figure 2-7.

In the high-frequency limit, 1-D planar transport occurs on both sides of the heater line. The integral in Eqn. (2-6) is dominated by values of ξ such that $\xi \ll Re[q_j]$ for materials 1 and 2, so the transfer function can be approximated as

$$\begin{aligned} \mathrm{Z}_{1+2,high\ \omega} &= \frac{1}{\pi L}\left(\frac{1}{k_1 q_1 + \mathrm{k}_2 q_2}\right)\int_0^\infty \left(\frac{\sin^2(\xi b)}{(\xi b)^2}\right) d\xi \\ &= \frac{1}{2Lb}\left(\frac{\exp\left(-i\frac{\pi}{4}\right)}{\sqrt{2\omega(kC)_1} + \sqrt{2\omega(kC)_2}}\right), \end{aligned} \tag{B2-15}$$

where $\sqrt{(kC)}_j$ is the thermal effusivity of material *j*. As expected, the BMA becomes exact in this high frequency limit for all effusivity ratios because the materials do not thermally communicate outside of the heater line during purely planar heating.[30] If the BMA is used for data analysis in planar heating experiments, care must be taken to ensure that both materials remain in a planar heating regime, because error increases as one material transitions from a planar to a cylindrical heating regime.

2.7.3 Appendix C - Dielectric layer & Interfacial Resistance analysis

Josell et al [44] published a corrected version of the thermal transfer function Z_{1+d+2} that includes a finite-thickness dielectric layer between the heater line and material 2 that was derived by Kuo et al. [44], [45] for an equivalent opto-thermal system. Josell's derivation is equivalent to the Feldman matrix method [35]. Briefly, the Green's function for a periodic line heat source is found by taking the spatial Fourier transform parallel to the material interface for the governing steady periodic Helmholtz heat equations and the interface conditions. Solving for the constants of the temperature field within the three different materials, taking the inverse Fourier transform with respect to space, and using the convolution theorem for a finite heat source yields

$$\mathrm{Z}_{1+d+2} = \frac{1}{\pi L}\int_0^\infty \left(\frac{\cosh(\Upsilon_d R_d) + \sinh(\Upsilon_d R_d)\left(\frac{\Upsilon_2}{\Upsilon_d}\right)}{(\Upsilon_1 + \Upsilon_2)\cosh(\Upsilon_d R_d) + \left(\frac{\Upsilon_1 \Upsilon_2}{\Upsilon_d} + \Upsilon_d\right)\sinh(\Upsilon_d R_d)}\right)\left(\frac{\sin^2(\xi b)}{(\xi b)^2}\right) d\xi, \tag{C2-16}$$

where the subscript *d* indicates the dielectric layer and $R_d = t_d/k_d$ is the area-specific thermal resistance of the dielectric layer with thickness t_d. A nondimensionalized Z_{1+d+2} is a function of 5 nondimensional parameters. We list our selected groups in Table 2-2, along with approximate values for our experiments on non-frozen mouse liver at 20 ^{o}C, representing our most

conservative worst-case measurements. Errors in the slope method are worse for the fresh liver than the higher-diffusivity frozen liver, and we use the high end of the frequency fitting range as the characteristic ω used to calculate values in Table 1, as this is most conservative. As in the main text, subscript 1 refers to the sensor and subscript 2 refers to the sample. Numerically evaluating $\partial \mathrm{Re}[Z_{1+d+2}]/\partial ln(\omega)$, the worst-case error in the $\left(k_{sensor}+k_{sample}\right)$ generalized slope method for our liver experiments is 3.3%. If the dielectric were not present, the error in the slope method would be 0.8% for the same sample and substrate properties and test parameters. Errors in the slope method for the frozen liver properties are less than 0.5% using the same dielectric properties ($k_d = 0.11$ W/m-K[46] and $t_d = 1\ \mu m$).

Table 2-2: Nondimensional Parameters Used In Calculating Z_{1+d+2} With Most Conservative Values

Nondimensional group	Conservative value in mouse liver experiment
k_2/k_1	0.5
α_2/α_1	0.27
$R_d k_1/\lambda_1$	0.11
λ_d/t_d	33
λ_1/b	8.5

Ju and Goodson [47] showed that a thin dielectric layer can be treated as a pure thermal contact resistance when $\lambda_d/t_d \gg 1$ for planar heating in one-sided 3ω. In the bi-directional case of Eq. (C2-16), the dielectric thermal capacitance and the 2D heat spreading inside the dielectric are relatively unimportant when $\lambda_d/t_d >> 1$ and $t_d/b << 1$, respectively. Since these two conditions hold for common thin dielectrics ($t_d/b = 0.1$ in our experiments) and heating frequencies, we Taylor expand the hyperbolic terms of Eq. (C2-16) to first order in $\Upsilon_d R_d$. Furthermore, when $\lambda_1/b \gg 1$, we can follow the method of Cahill[17] to simplify the integral in the low-frequency regime. Neglecting terms of order t_d/b or smaller that would include information about heat spreading and dielectric capacitance, the integral becomes equivalent to assuming a pure thermal contact resistance at the interface:

$$\mathrm{Z}_{1+R_d+2} = \frac{1}{\pi L}\int_0^{\frac{1}{b}} \left(\frac{1+R_d\Upsilon_2}{\Upsilon_1+\Upsilon_2+\Upsilon_1\Upsilon_2 R_d}\right) d\xi. \qquad \text{(C2-17)}$$

For large R_d, the adiabatic one-sided solution for material 1 is recovered, and for small R_d, the perfect contact solution in Eq. (2-6) is recovered. The slopes obtained from Eqs. (C2-16) and (C2-17) using the relevant parameter values from Table 1 agree to within 0.5% for λ_1/b =10.

To gain further insight into the error in using the generalized slope method when a dielectric layer or large thermal contact resistance is present, we study the specific case of equal volumetric heat capacities, $C_1 = C_2$. This is a reasonable first approximation for many fully dense materials. Now, after imposing equal heat capacities, neglecting dielectric capacitance, and neglecting heat spreading within the dielectric, the error in using the low-frequency limit to

approximate $\partial \mathrm{Re}[Z_{1+R_d+2}]/\partial \ln(\omega)$ can be expressed in terms of only three dimensionless groups, selected as λ_1/b, $R_d k_1/\lambda_1$, and k_2/k_1.

For the situation just described, Figure 2-8 shows the error in using the slope method, as a function of the normalized thermal contact resistance for different values of k_2/k_1. The error is again defined as (LFL-FF)/FF. The error in the slope is insensitive to λ_1/b when the low-frequency assumption in the derivation of Eq. (C2-17) is satisfied. As the contact resistance increases, a larger λ_1 is required to access the regime where the slope method is accurate. In the limit that $R_d k_1/\lambda_1$ is large, the slope is only sensitive to material 1. Now looking at the conductivity ratio dependence, we see that the slope method can be used for larger normalized resistances when material 1 and 2 have comparable conductivities. When $k_2/k_1 > 1$, the slope method breaks down at relatively smaller values of the k_1-normalized contact resistance, because the ratio $R_d/\mathrm{Re}[Z_2]$ becomes significant for smaller values of $R_d k_1/\lambda_1$.

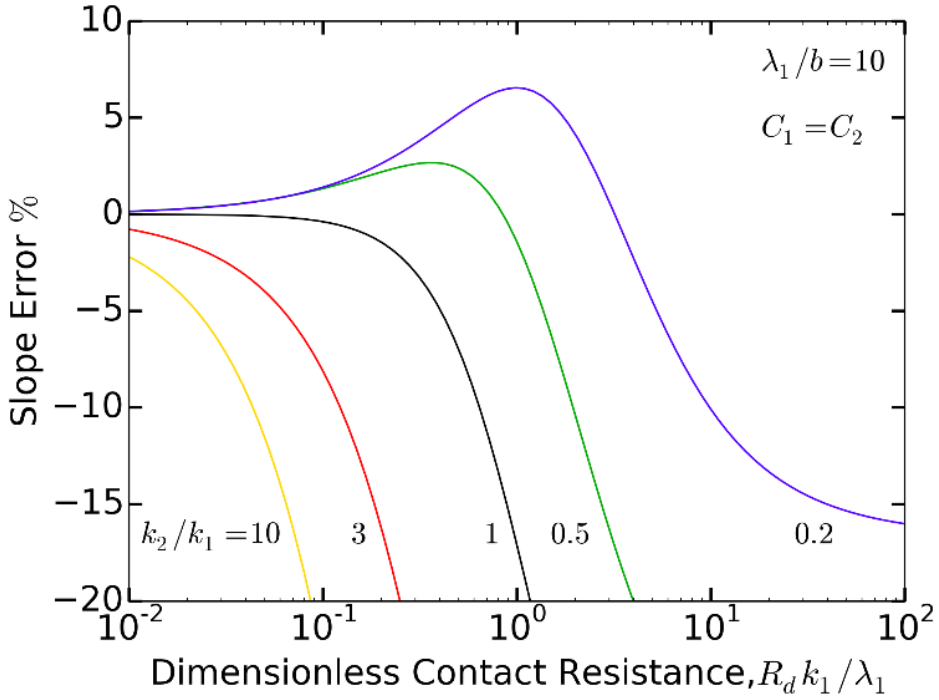

Figure 2-8: Error in the generalized slope method due to a thermal contact resistance layer between the heater and material 2 (the sample). The error is defined as (LFL-FF)/FF. Larger penetration depths are required to reach the range of validity for the generalized slope method when the contact resistance is large. The experimental values for the fresh mouse liver experiment (worst case) correspond to $R_d k_1/\lambda_1 = 0.1$ and $k_2/k_1 = 0.5$ with $C_2/C_1 = 1.6$.

Lastly, we note that these thermal transfer functions for layered bi-directional geometries can be used to extend the single-sided differential 3ω method for measuring thin films [19], [20] to a "bi-directional differential" experiment. The differential measurement would compare the measured bi-directional thermal transfer function when a thin sample is inserted between the dielectric layer and material 2 ($Z_{1+d+f+2}$) to an identical bi-directional geometry without the thin sample (Z_{1+d+2}). In the BMA, the film is treated as an additional thermal resistance that is found by simple parallel/series resistor algebra using the two measured transfer functions. The sensitivity to the cross-plane film conductivity is not improved in the bi-directional differential

method compared with the single-sided differential method because of the parallel heat flow pathways from the heater line, but a bi-directional scheme might still be advantageous for the practical assembly of challenging samples. Optimal sensitivities would be obtained when the substrate (1) has low k, the dielectric layer (d) is as thin and conductive as possible, and the material (2) above the supported film (f) is of high k. The bi-directional differential measurement could potentially be utilized to measure the cross-plane conductivities of solution processed thin films, flakes of bulk materials, or biological samples as thin as 1 μm.

Chapter 3 - Variable Spot Size TDTR

Significant thanks are given to Dr. Christian Monachon for training me on TDTR. Without his extensive experience and mentorship, this project could not have been completed with nearly the speed nor to the level of scientific rigor that was possible.

3.1 Introduction

Traditional engineering was constrained in the diversity of technologies that could be created due to the limited ranges of accessible material properties. With the onset of nano-scale engineering, it is now possible to engineer material properties themselves, opening access to a larger space of possible values of material properties and property combinations, allowing for more complex designs. The family of macroscopic thermal energy properties such as heat capacity and thermal conductivity is governed on the nano-scale by the behaviors of quantized bundles of atomic vibrational energy called phonons. Their velocity, energy, scattering rates, and statistical distributions govern everything from how much energy it takes to keep a building cool in the summer, to why pipes buried deeply enough never freeze in the winter, and even how long the Mars Curiosity rover can power itself.

Detailed knowledge of the behavior of phonons in different materials enables us to calculate their thermal transport properties in all situations, including when the length scales of the material are small enough to invalidate Fourier's law of heat conduction and we need to account for the phononic granularity of the flow of heat. Furthermore, an understanding of how different phonon scattering mechanisms affect the behavior of phonons equips us with the ability to know how the introduction of those different scattering mechanisms will change these thermal transport properties in controllable ways. Taken together, if we can know the baseline phonon properties of a material, how they collectively govern thermal transport properties, and how different introduced scattering mechanisms will change those phonon behaviors, then we can engineer the thermal properties of a material to have extreme and unique behaviors, such as ultra-low thermal conductivity [48] and shielded or inverted thermal fluxes [49].

There are already many pertinent technologies that either already have or could benefit from thermal nano-engineering. Nanostructuring a material can increase phonon scattering to reduce thermal conductivity via the introduction of defects and impurities or by ball milling and hot pressing. This has been used to boost the efficiency with which thermoelectric materials can generate useable electricity from otherwise wasted low-grade heat [50]–[53]. Similar benefits can also be achieved by shrinking materials down to create nanowires [54]–[56] or nanomeshes [57]. Heat dissipation from nanometer scale transistors in modern electronics is one of the primary challenges threatening Moore's law and limiting how small we can make computer chips [58]. The classical equations of heat conduction ignorant of the details of phonon properties fail to predict heat transfer at such small length scales, further hindering our ability to make progress in microelectronics. For the same token, controlled nano-scale heating of magnetic media may hold the key to the next generation of hard disk drive storage with increased information density (up to ~ 375 TB/m^2) via Heat-Assisted Magnetic Recording (HAMR) devices [59], but to reach its full potential, this technology requires writable magnetic media with a precisely engineered nano-scale anisotropic thermal conductivity tensor and effective heat

dissipation from its nanometer scale near field thermal transducer. High-powered lasers are currently limited in their maximum power output by how much energy the lasing medium can dissipate before overheating. Solving this problem requires simultaneously engineering both the optical and the thermal transport properties of a material, and the most promising approaches thus far involve creating anisotropic crystal nano-grains to control directional photon and phonon scattering [60], [61]. LEDs produce much less heat overall than incandescent bulbs, but what heat they do produce is very locally concentrated in nanometer sized quantum wells, and reducing these semiconductor junction temperatures by just 11 ºC would triple the lifetime of LEDs [62] and increase their luminous output by a factor of 5-7 [63], making them far more economically competitive and desirable, hence boosting their adoption rates. If most industrial and residential lighting applications adopted the best LEDs, the total energy consumption of the entire United States would drop by a staggering 5% [64], [65]. Meanwhile, the issue of thermal management in batteries is a direct safety concern influencing how well batteries mitigate thermal runaway shorts that can otherwise cause batteries to explode or burst into flames [66]. Such battery fires have occurred in smart phones, airplanes, and electric vehicles, occasionally prompting widespread product recalls. Improving the heat transport in batteries would also reduce their recharge time [67] and remove the need for much of the often-present bulky external cooling systems [68], [69], leading to safer and smaller personal electronic devices, as well as lighter electric vehicles with increased driving range [70]. All of the controllable internal thermal transport in these battery systems is governed by the interaction of phonons in nanometer to micron scale materials and features [70]. Even promising heated nanoparticle-based cancer treatments depend on nano-scale heat conduction [71].

All of these important problems require a fundamental understanding of nano-scale thermal energy transport coupled with the ability to measure properties of phonon populations in different materials. As all forms of technology advance and move toward smaller minimum feature sizes, such nano-scale thermal engineering will become ubiquitously necessary. A large body of work already exists studying phonon transport properties theoretically and experimentally, with a particular eye toward its application for nano-scale thermal engineering and broader impacts [58], [72]–[78]. This work has resulted in great progress, especially regarding theoretical models and understanding. However, in spite of all this effort, it is only within the last few years that such sub-continuum thermal transport regimes have become more experimentally accessible, especially for bulk materials. There are still only a few number of experimental techniques capable of accessing this regime, and reasonable debate still exists regarding how to properly interpret many of the experimental results [72], [79]–[82].

To contribute to this ongoing effort, in the rest of this thesis we describe the variable spot size time domain thermoreflectance measurement technique and how it can be used as a tool to measure such phonon properties, demonstrating its efficacy by measuring sub-diffusion thermal transport in sapphire, silicon, nano-grained silicon, and silicon germanium samples. We explain the microscopic origins of macroscopic heat conduction in a physically intuitive paradigm and how this understanding can be used to extract information about phonon properties from measurements of bulk samples. The results of our measurements are discussed and related back to this microscopic understanding of thermal energy transport.

3.2 Description Of The Experimental System

3.2.1 Traditional TDTR

Time Domain ThermoReflectance (TDTR) [27], [83] is effective for measuring both high (~1000 W/m-K) and low (~1 W/m-K) thermal conductivity (k) materials [84], thin films and interfaces [73], [85], as well as liquids [28]. Ultimately, the only requirement is that the sample of interest be in intimate thermal contact with a thin, specularly smooth metal transducer layer and be spatially homogeneous on the length scale of the thermal penetration depth as determined by the modulation frequency. The thermal penetration depth for planar heating is given by

$$L_p = \sqrt{D/2\pi f} \tag{3-1}$$

where D is the thermal diffusivity, and f is the frequency of thermal oscillations.

In traditional TDTR experiments, an ultrafast pulsed laser is used in a pump-probe configuration. The raw output of the laser is a constant pulse train of equally spaced high-intensity pulses typically ~200 femtoseconds in duration (FWHM). The repetition rate of a standard Ti:Sapphire pulsed laser is 80 MHz (80.1 MHz in our experiments), resulting in one pulse every 12.5 nanoseconds. The pulse widths are hence approximately 5 orders of magnitude briefer than the time between pulses, and so the pulse train can be very well described mathematically as a series of Dirac delta functions, known as a "Dirac Comb." More details of the mathematical modeling will be given later in section 3.6 TDTR Mathematical Theory. The laser output is divided into two separate beams, a "pump" and a "probe" beam. Traditionally, this beam splitting is done with a regular or polarizing beam splitter cube [86], [87]. The pump beam is used to heat the sample, while the probe beam is used to measure the resulting surface temperature changes of the sample, which depend on the sample's thermal properties.

The probe beam passes through a variable length optical path, while the pump beam traverses a constant-length optical path on its way to the sample. By increasing the probe's optical path length by a distance d, its arrival time at the sample is delayed by a time d/c as compared to the pump, where c is the speed of light in air, approximately $3\text{x}10^8$ m/s. The enormity of the speed of light, combined with the brevity of the laser pulse widths, allow for very precise temporal resolution of the measured transient cooling of the sample. For example, increasing the probe's optical path length by 0.3 mm will delay its arrival time by 1 picosecond. This is the core idea of ultrafast optical pump-probe systems. The final major component for traditional TDTR systems is a modulator that chops the pump beam at a high frequency (typically 1 to 10 MHz). This introduces a frequency component at the modulation frequency in the temperature response of the sample, which can be picked out from the probe beam's signal by a lock-in amplifier. Because only the pump beam is chopped, any component of the probe beam's signal at the modulation frequency must have originated from the sample's temperature respond, and hence is the thermal signal of interest. This lock-in technique is similar to that used in the 3-Omega system, and is a great way to increase signal-to-noise by up to a few orders of magnitude [88].

The sample to be measured is typically coated with a thin (~100 nm [86], [89]) metal film to act as an optical transducer layer. The purpose of this metal film is two-fold. First, it serves to absorb the incident pump beam's optical energy over a short distance and convert this energy into heat. This greatly simplifies the thermal modeling as it means the incident heat flux can be well modeled as a boundary condition at the sample's surface, rather than having to account for

the details of a distributed volumetric heat generation penetrating non-negligibly into the sample. Second, the metal's optical reflectivity changes as a function of temperature, allowing us to use the amount of reflected probe beam light as a thermometer for the sample. For small to moderate temperature excursions (typical TDTR experiments have < 1 K of surface temperature oscillations) this relationship between temperature and optical reflectivity can be assumed to be linear.

3.2.2 Variable Spot Size TDTR

Our system is based on the traditional TDTR setup, with a few modifications that will be described in greater detail in section 3.3. Most notably, we vary our laser spot size by changing the objective lens used to focus the pump and probe beams onto the sample. This allows us to create thermometers and heat sources with characteristic lengths on the order of mean free paths of some of the phonons in the sample, resulting in a tool that can study sub-continuum effects of heat conduction at the micro- and nano-scales. A more in-depth explanation of the physics of sub-continuum heat conduction will be given later in Chapter 4 - Effects of Alloying and Nanostructuring on Spectral Phonon Mean Free Path Distributions. Several challenges arise from using spot sizes only a few microns large, especially when trying to precisely overlap two such spot sizes, and prevent either one from changing location, size, or shape throughout a measurement while the probe beam's optical path is elongated by up to a meter. This instrumentation used to achieve such results is now discussed in section 3.3.

3.3 Instrumentation and Experimental Tricks

In this section, I will give an overview of the hardware and instrumentation used in our implementation of TDTR, as well as describe the various tactics employed in our system to reduce noise and generally improve the robustness of our measurements. Most of these tricks of the trade have already been described elsewhere in the literature, but by collecting them all in one place and spending a little bit more time explaining the reasoning behind each, I hope to help make the challenging design of a robust and functional TDTR system more accessible to future researchers.

3.3.1 System Overview

Here I will briefly walk through all components of the variable spot size TDTR system, and then describe key elements in greater detail below. Our specific setup is given in Figure 3-1 below. The laser pulses are originally generated in the Ti:Sapphire laser, whose fundamental wavelength at 800 nm partially passes through as the probe beam. This wavelength is chosen as it is near where aluminum exhibits a local maximum for its temperature-dependent optical reflectivity coefficient, maximizing detected signal strength. Aluminum is used as the optical transducer in all of our measurements, with a typical film thickness of 70 nm. The remaining power from the Ti:Sapphire laser pumps the Optical Parametric Oscillator (OPO) which then feeds a Second Harmonic Generation (SHG) crystal, to output a spatially distinct pulse train beam at 530 nm, used as the pump beam. This wavelength was chosen to maximize the output power of the laser OPO SHG while being sufficiently spectrally distinct from the probe beam. This laser system was purchased from Coherent's Chameleon compact OPO-Vis laser series, and

allows for independent tuning of our pump and probe wavelengths across the full visible spectrum into the near-infrared. This laser system is a shared facility in the Molecular Foundary at the Lawrence Berkeley National Lab, with a large user base.

Independently tuned pump and probe, with large spectral separation, offer some benefits over the more common "two-tint" TDTR system. In two-tint TDTR [87], both pump and probe beams are derived from the same laser, whose output has a spectral width of at least 10 nm or greater. The beam is split using a polarized beam splitter cube, and then different branches are passed through very sharp spectral filters. The top half of the spectrum, from the center wavelength of the original laser output and above, is used for one beam, while the bottom spectral half is used for the other beam. This approach is very effective at reducing noise created by rough samples that diffusely reflect incident light, possibly into different polarizations and angles, limiting the efficacy of polarization or spatial separation to distinguish pump from probe. Our system has these same benefits, but also does not require such ultra sharp spectral filters, owing to our pump and probe beams being spectrally separated by hundreds of nm instead of single digit nm. Much cheaper and readily available filters are sufficient. Thanks to this spectral elbowroom, we can also use a dichroic mirror to combine and separate optical paths of pump and probe. By combining the functions of the polarizing beam splitter cube and the filters into one optical element, we double the overall power transmission of our system, from laser to sample. In the two-tint configuration, the power is divided once by the polarizing beam splitter cube, and then a second time by the sharp spectral filters. Independent tuning of the pump and probe wavelengths over a couple hundred nm of range also allows us to use a variety of transducer film materials. This provides the flexibility to use different metals, as well as plasmonic materials.

The probe beam is directed through a beam expander, before being guided through two passes of a delay stage with the help of three corner-cube reflectors. We can delay the probe beam by a maximum of 1 m in total, corresponding to 3.3 ns of delay with respect to the pump. The probe then passes through an objective that launches it into a single-mode optical fiber. The mouth of the fiber is mounted on a piezoelectric XYZ stage for careful positioning with respect to the objective's focus. The fiber sends the probe through another beam expander and then a 50/50 beam splitter. The transmitted portion is measured by detector B as a pre-sample of the probe power before the sample. The reflected portion is directed through the objective lens (2x, 5x, 10x, or 20x, working distances = 20 mm to 35 mm) focusing it onto the sample. Typical probe laser powers are on the order of 5 mW for spot sizes on the order of 5 μm by the time the beam reaches the sample. The probe beam is then reflected by the sample and sent back to detector A, which has a 715 nm long-pass filter over its entrance to block out all spurious background light signals as well as the pump beam. The pump beam is directed through a long 750 mm focal length lens to converge it through the EOM and avoid internal reflections, and then two mirrors whose position can be precisely translated synchronously to finely adjust the arrival time of the pump. The pump is then also expanded before being sent to the sample via a 700 nm single-edge dichroic beam splitter that transmits the probe but reflects the pump. Typical pump laser powers are on the order of 10 mW for spot sizes on the order of 5 μm by the time the beam reaches the sample. All beam expanders are 4x. Two low-reflectivity (92:8) beam splitters are used, one to inject flood illumination onto the sample for imaging, and the other so that a CCD camera can directly view the sample and laser spots. The sample is mounted on an XYZ stage, and inside an optical cryostat for low-temperature measurements.

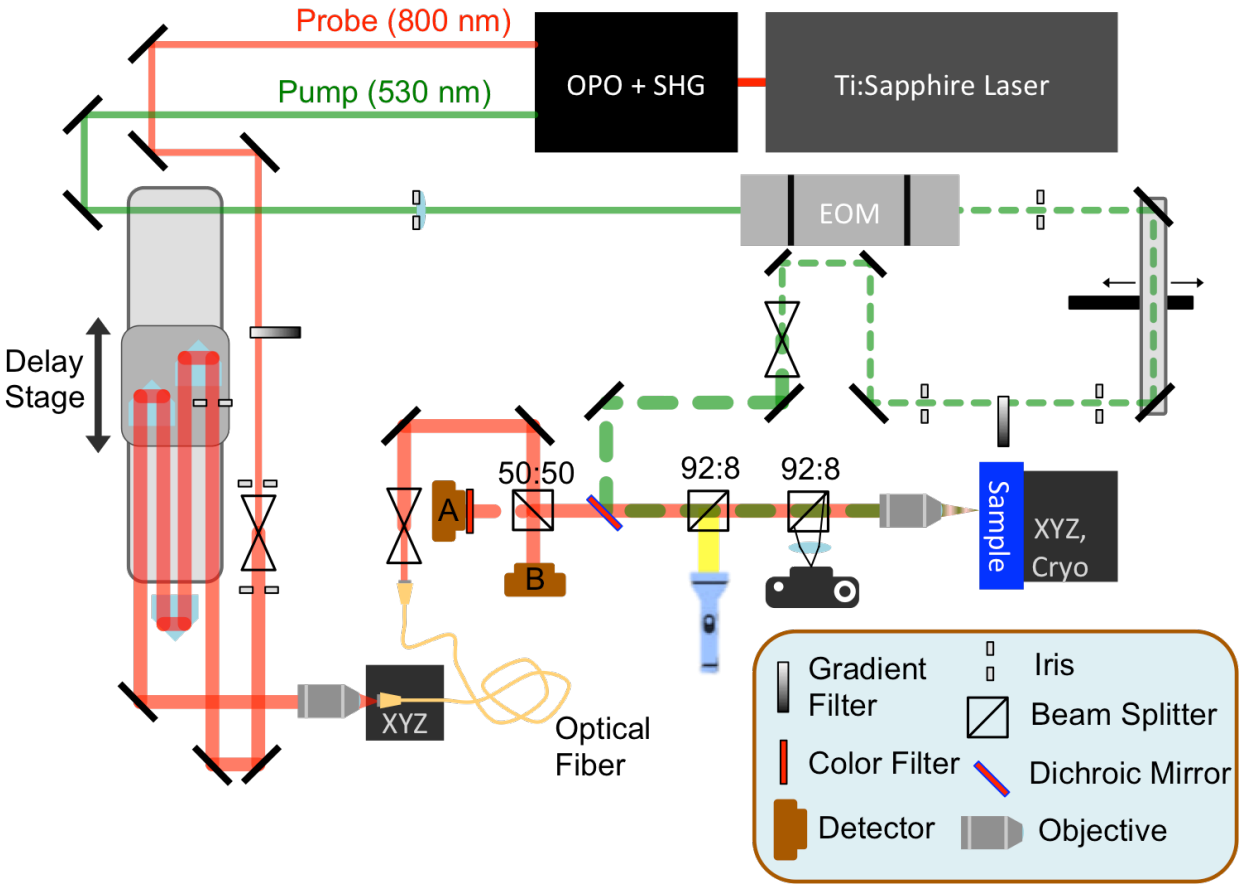

Figure 3-1 Complete TDTR System, in the Molecular Foundry at LBNL

3.3.2 Beam Steering, Spot Size Control, and Beam Expanders

One of the greatest experimental challenges of TDTR is keeping the probe spot location, shape, and size constant on the sample throughout the measurement, as the probe's optical path length is varied by over a meter via motion of the delay stage. We combine several strategies to control probe spot stability during our measurements. A perfectly collimated beam will diverge at a rate approximately inversely proportional to its minimum beam radius [90]. Therefore, to minimize the divergence angle of the probe beam as it travels along the increasing length delay stage, it is first expanded to maximize its radius.

A perfectly aligned probe beam's position on the sample would not drift when its delay is increased. This is impossible in practice, however, and any finite deviation from a perfectly orthogonal incident angle of the probe beam on the sample will result in some slight drift in its position, or "beam steering," when its delay is increased. Additionally, even with the beam expander, there will still be some finite divergence of the probe beam, resulting in a change in its spot size when its delay is increased. A change in the probe spot size or a change in the overlap of pump and probe spots during the measurement will invalidate the thermal model. To mitigate this problem, we follow the approach of Ref. [91] and feed the probe beam into a single-mode optical fiber after the delay stage. Allowing the transmission of only one mode, or spatial intensity distribution, by the fiber ensures that the output spatial intensity profile of the beam is Gaussian regardless of the launch conditions when the probe beam is sent into the fiber. The

output angle of the single-mode fiber is also independent of small variations of the launch angle of the beam entering the fiber. Thus, the probe beam's location on the sample, spot shape, and spot size, remain constant throughout a measurement even as the mechanical delay stage is advanced, thanks to the single-mode fiber.

There are two drawbacks to using the fiber, however. First, a significant amount of power is lost. We routinely achieved 20% to 30% power transmission of our probe beam through optical fibers half a meter in length. Second, even a very small spatial drift on the order of one micron of the position of the mouth of the fiber relative to the launching objective can cause a dramatic loss of coupling efficiency. Similarly, any beam steering coming out of the delay stage, causing the angle at which the probe beam passes through the launch objective, can similarly destroy the fiber coupling. To mitigate this, we attached the mouth of the fiber to a piezoelectric XYZ stage capable of fine (sub-micron) positioning of the fiber in x, y, and z. An automated LabVIEW program would sweep the position of the fiber by +/- 10 microns in x, y, and z, locating the position where the transmitted power was greatest, and then set the fiber mouth to this location. This fiber sweeping re-alignment protocol was performed before the start of every measurement and then again every 200 to 1000 picoseconds of delay during a measurement, to correct for drifts of the fiber launch hardware and of the probe beam exiting the delay stage. This substantially improved the consistency of power coupled through the fiber throughout a measurement. Across the full 3.3 ns of delay, routine power coupling efficiency variations throughout a single measurement were on the order of 20%. Finally, in order to suppress transmission through any cladding modes in the fiber, the alignment of the launching objective was made as parallel to the optical axis of the fiber as possible, and the fiber was deliberately wound into a few loops.

Incorporating an optical fiber also creates a convenient way to add additional static delays to the probe line. By swapping the fiber out for a longer fiber, the total delay time of the probe beam can be statically increased, and then swept out from this new starting point using the delay stage. This can then be repeated to stitch together data for a full pulse-to-pulse cooling curve, as was done in Ref. [92].

A second challenge in our system was to achieve very small spot sizes. We varied our spot sizes by using different magnification objective lenses. We used a 2x, 5x, 10x, and 20x objective lens, corresponding to $1/e^2$ radii of pump and probe beams of approximately 1.8 μm, 3.1 μm, 6.8 μm, and 16.4 μm respectively. The ratios of the spot sizes does not entirely match the ratios of the objective magnifications due to the particular numerical apertures (NAs) of each objective, and due to the effects of other optics in the system that focus the beams prior to their reaching the objective lens. To fully utilize the capabilities of the 2x, 5x, 10x, and 20x objectives used, we expanded both pump and probe beams before sending them into the objective lens. Expanding the size of the incoming beam so as to approach overfilling the backside of the objective lens fully utilizes its numerical aperture, achieving minimal spot sizes.

3.3.3 Modulation, Lock-In, and RLC Filters

It is standard for TDTR systems to use a modulator (we used an electro-optic modulator, or EOM) to chop the pump signal at high (1 to 10 MHz) frequencies (but still lower than the laser's innate pulsed frequency ~ 80 MHz), and then use a lock-in amplifier whose reference is synchronized to this same frequency to pick out the modulation frequency component of the probe beam after it has reflected off the sample. The lock-in amplifier works by internally

multiplying the input detected signal by a sine wave of known amplitude at the reference signal's frequency, f. This mixes with every frequency component, v, of the detected signal, producing a signal at the sum and difference of those two frequencies, $f+v$ and $f-v$, for each frequency component. The total signal is then passed through an adjustable low-pass filter, which permits only those signal components that are very nearly DC, corresponding to $f-v \approx 0$, or $v \approx f$. The cut-off frequency of this low-pass filter is further reduced by increasing the time constant of the lock-in amplifier. In this way, the lock-in amplifier picks out the amplitude of the component of the detected signal at the reference frequency, f.

Ideally, the pump beam modulation would be sinusoidal, but in practice it is usually a square wave. This means that it contains multiple frequencies at all the odd harmonics of the fundamental frequency of the square wave. This would not be a problem if the lock-in amplifier used a pure sine wave mixer when internally multiply the input detected signal. However, the radio frequency (RF) lock-in amplifier used (Stanford Research Systems model SR844) also uses a square wave for its internal multiplication. This means that without intervention, our detected "pure" frequency signal is actually a hodgepodge mix of multiple different frequency components, and is no longer a clean signal as the thermal model assumes. The system is now also highly susceptible to noise at any of the odd harmonic frequencies of the square wave. Our solution, as employed in Ref. [27], was to place an inductor in series with the BNC cable leading the signal from our Si photodiode (detector A in Figure 3-1) to our lock-in amplifier. This couples with the electrical resistances and capacitances of the BNC cable, Si photodiode, and input impedance of the lock-in amplifier to form a resonant RLC filter, whose resonance frequency can be controlled by deliberate selection of the inductor value. The true resonance frequency of each analog filter was characterized after construction using an oscilloscope. The resonant frequencies of our three primarily used filters were measured to be 0.9375 MHz, 2.4615 MHz, and 9.5 MHz (referred to generally as approximately 0.94 MHz, 2.5 MHz, and 9.5 MHz in this dissertation). These resonant filters amplify the signal at the desired fundamental harmonic of the square wave from the EOM, and attenuate the odd harmonics. An alternative approach to eliminating the odd harmonics of the square wave is to carefully adjust the waveform driving the EOM, as employed by Ref. [93].

While the lock-in approach is very effective at rejecting all signal components not at the prescribed modulation frequency, f, it is still susceptible to noise at f. To account for this, a reading of the in-phase and out-of-phase background noise levels at f were taken before each measurement, and subtracted from the measured signal. These background levels were measured as the X and Y outputs of the lock-in for a fully aligned and powered measurement, except with the pump beam blocked. Typical background noise levels were 1% to 10% of the measured signal for the lower 2 modulation frequencies, but could be much higher at 9.5 MHz, which is further discussed later in section 3.7.1 System Modeling and Data Fitting.

We also found that alignment through the EOM is very delicate, and even the slight divergence of the pump beam coming from the laser is enough to cause internal reflections within the EOM and severely distort the beam shape and cause loss of transmission. To mitigate this, we placed a long focal length (750 mm) lens after the iris used to align the pump beam into the EOM, to converge and focus the beam through the EOM. This improved the transmission through the EOM, and considerably improved the spot shape distortion.

3.3.4 Zero-Time Finding

While the mechanical delay stage allows for the continuously advancing delay of the probe beam relative to the pump beam, it is still necessary to align the zero delay position of the probe pulse to closely coincide with the pulse of the pump beam. In our experiments, we typically start our measurement at a probe delay time of -20 ps, or 20 ps before the pump pulse hits the sample. This way, we can see the temperature rise of the pulse and be confident in knowing the zero-time of when the transient temperature decay begins when fitting to the data. However, it can be non-trivial to align the zero-delay point to be near the incident time of the pump pulse. It is not possible to manually "eyeball" the pump beam path during construction, so that it is the same length as the zero-delay probe beam path down to less than 0.3 mm (corresponding to 1 ps of delay). To allow for fine-tuning of the pump path length, two of the mirrors (the two mirrors after the EOM, as shown in Figure 3-1) are mounted on a horizontal beam that is connected to a micrometric screw. This micrometric screw can finely adjust the lateral position of these two mirrors, increasing or decreasing the pump beam optical path length by twice the distance the mirrors move. By iteratively sweeping with the mechanical delay stage and then adjusting this micrometric screw, the zero-time can be found and positioned.

3.4 Spot Size Measurements and Chromatic Aberrations

3.4.1 General Spot Size Measurement Protocol

Spot size measurements are a dominant source of uncertainty in TDTR (see sections 3.7.2 Uncertainty Analysis and 3.8 Sensitivity Analysis). Laser spot sizes are particularly important in our experiments, as we are using small laser spots as a way to measure sub-continuum heat conduction effects. We use a CCD camera in our system as shown in Figure 3-1 for visualizing the sample and measuring spot sizes. We use the camera to record the spot size of pump and probe before every measurement. We placed the camera as close to the sample as possible in order to minimize the number of optical elements between the camera and sample, allowing us to see an undistorted image of our sample with pump and probe spots. We used LED flood illumination in order to see the sample with the camera. Because the flood illumination is injected into the optical path behind the objective lens, it undergoes the same level of magnification as the image viewed on the sample. As such, the total amount of light per viewing area is approximately constant, and the brightness of the flood illumination does not need to be increased with higher magnification objective lenses. This section will describe the protocol used to measure the spot sizes with the camera, and the various steps taken to control and verify these measurements.

In order to use the camera to reliably measure spot sizes, we first had to calibrate the magnification of the camera. This was done using a standard control scanning electron microscope (SEM) grid, which had squares of precisely 10 micron pitch, in 49 x 49 groups of squares, as shown in Figure 3-2. For each of our four objectives, in-focus images were taken of the SEM calibration grid. These pictures were then used to calibrate the camera-pixel-to-micron conversion ratio for each objective. To increase accuracy and attain sub-pixel precision in this conversion coefficient, the largest possible line was taken from the calibration pictures, spanning the greatest integer number of 10 μm x 10 μm squares as possible. The length of this line in

pixels was then divided by 10 μm times the number of squares spanned (both numbers possibly decimals, for the case of a diagonal line) in order to calculate the conversion coefficient.

Figure 3-2 Picture of the standard SEM grid used to calibrate the pixel-to-micron conversion factors for each objective lens. The grid consists of 49 x 49 groupings of 10 μm x 10 μm pitched squares. This picture was taken with our camera using the 20x objective lens.

To measure accurate spot sizes, high quality images must be taken. Each gray scale pixel of the camera can record a value ranging from 0 to 255, with 255 being the brightest. We would like to use as much of this range as possible without ever saturating a pixel. This enables us to use the maximum number of bits in resolving the intensity distributions of our spot pictures without losing information. To achieve this, we used gradient wheel neutral density (ND) filters to independently adjust the overall intensity of pump and probe beams for optical picture taking conditions. Using vertical and horizontal line cuts in the camera's software, we adjusted the filter intensity until the brightest pixel in the laser spot came close to a value of 255 without actually saturating. This also created a consistency in spot size measurements across sample types. When taking a picture, a beam block was used to block whichever laser spot was not currently being photographed, all LED flood illumination turned off, and the camera was outfitted with a long opaque tube to block out most ambient light.

Once photographed, a 40 pixel x 40 pixel array, centered on the laser spot, was cropped out of the original photograph and used to fit an arbitrarily oriented 2-D Gaussian function to the laser spot. This array size included enough surroundings to let the tails of the Gaussian spot decay to zero (i.e. to significantly below the ambient noise level), without including too much of the background to influence fitting to noise. The function used to fit the laser spot had 7 fit parameters: the amplitude of the Gaussian, the center of the Gaussian in x and y, the $1/e^2$ radius

of the Gaussian (w) along its minor and major axis, the rotation angle of the Gaussian, and a constant background offset. These 7 variables were fit to the 40 x 40 array constituting 1600 pixels of data. The fit always converged rapidly, with a typical uncertainty (as measured by the square root of the corresponding elements from the covariance matrix) of less than 1% for the $1/e^2$ radii. This fit method of fitting could properly fit to an elliptical spot.

The TDTR thermal model assumes a perfectly circular Gaussian spot for both pump and probe beams. We consistently found the best fits to TDTR data (as compared to literature values when available, and using the metrics described in section 3.7.2 Uncertainty Analysis) were achieved when taking the $1/e^2$ radius of the minor axis as the isotropic spot radius for the thermal model in cases where our measured spots were irregularly shaped, as opposed to a geometric average of the minor and major axis radii. Typical anisotropies (ratio of major to minor spot radius) for the probe beam spot were ~ 1.1 and for the pump beam spot were ~ 1.3. The pump beam spot anisotropy could be as high as ~ 2.0 for some (~15% of) measurements. We speculate that the minor axis radius yielded the best fit results because it is along this axis that the light (and hence heat source) is most compressed, and so will be the most influential in defining the effective lateral extent of the peak intensity of the heat source, as well as the planar versus spherical geometry heat spreading regime, discussed further in section 3.8 Sensitivity Analysis. Additionally, any slight mismatches in the exact focus of the laser spots on the sample as compared to in the camera would manifest primarily in differences in the major axis, with the minor axis remaining more consistent, as described more completely in section 3.4.2 Chromatic Aberration Corrections.

To separately validate this method of spot size fitting, we also measured select laser spots using a sweeping knife edge technique, which is fully explained in Ref. [86]. In brief, we used a straight edge from a piece of silicon cleaved along a crystal plane, driven by a piezoelectric XYZ stage, to slowly cut across our focused beam in x and y. The intensity of transmitted power as a function of knife-edge position was used to calculate the $1/e^2$ radius of the Gaussian cross section of the beam. Sweeping knife-edge measurements agreed with camera spot size measurements to better than 10%. Therefore, we conservatively set the uncertainty of our laser spot measurements to be 10% (see Table 3-1). This was the dominant source of uncertainty for most thermal conductivity measurements. We verified that neither changing the modulation frequency nor changing the camera brightness (brightest pixel ranging from ~ 80 to ~ 255) affected the measured $1/e^2$ radii.

Finally, changing the sample focus by moving it along the optical axis in either direction decreased the TDTR signal strength, confirming that the laser spots were already properly focused on the sample as a result of the spot picture taking procedure.

3.4.2 Chromatic Aberration Corrections

In our system, there is typically 270 nm of spectral shift between our pump (530 nm) and probe (800 nm) beams. An ideal system would have no chromatic aberrations. However, given the chromatic aberrations and depths of field of our objective lenses, this is sufficient spectral separation to cause one spot to be considerably out of focus when the other is in focus. A defocused laser spot will change the spot size on the sample. If the spot is not a perfectly circular Gaussian, then a defocused spot will also affect the spot shape. If the finite radius beam is not aligned to pass exactly through the center of the focusing objective, exactly parallel to its

optical axis, being out of focus will then affect both the spot shape and the spot position. Such errors can be detrimental to our measurements if not corrected and accounted for.

In order to understand our chromatic aberration correction protocol, it is necessary to understand all focusing elements that affect the pump and probe beams' focus onto the sample and onto the camera. These elements are shown schematically in Figure 3-3. The beam expander lens and sample position can both change the focus of the spots on the sample and in the camera, while the tube lens can additionally change the focus of the spot in the camera. Note that in our actual setup, the pump and probe beams each have their own separate and hence independently adjustable beam expanders. For simplicity of the figure, I have shown both beams passing through the same beam expander. The challenge is to simultaneously have both the pump and the probe beam focused on the surface of the sample and to take in-focus pictures of those spots with the camera for correct spot size measurements, even though all optical elements focus the pump and probe beams differently due to chromatic aberrations. We go through a protocol whereby we effectively lock in a correction using the separate beam expanders for pump and probe that cancels out the chromatic aberration, enabling us to have both beams simultaneously focused on the sample surface. We then tune the focus of the camera by adjusting the tube lens in order to take accurate spot pictures of each beam separately. This protocol is now explained in detail.

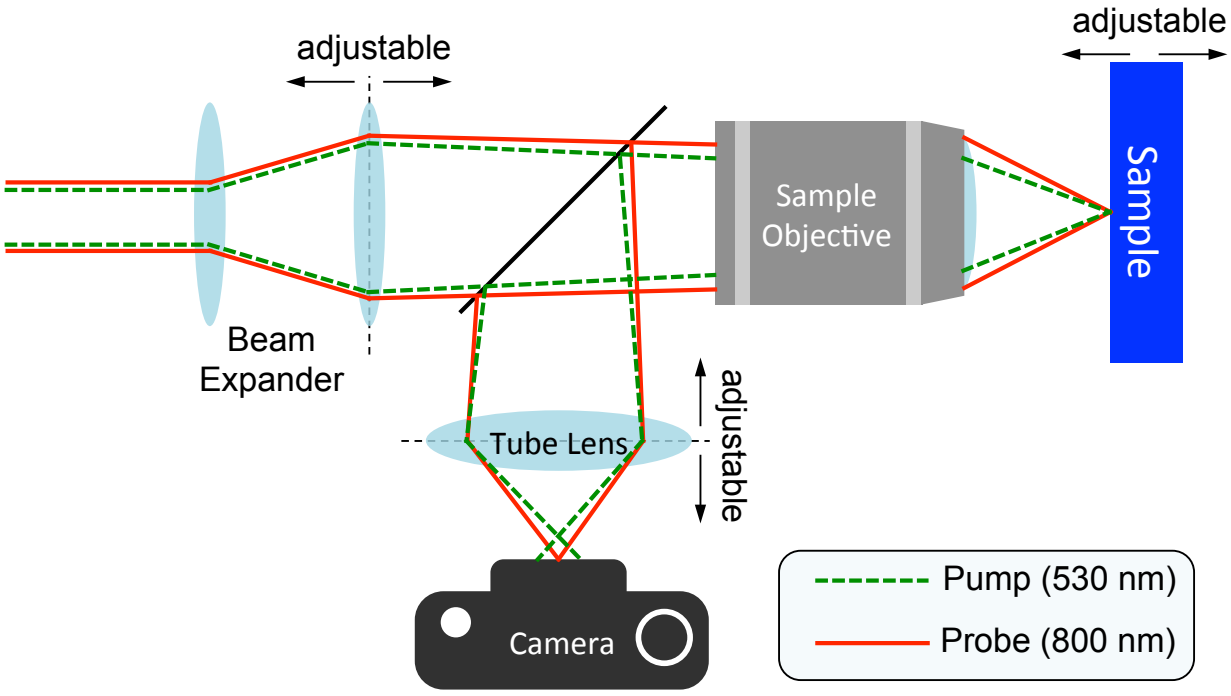

Figure 3-3 Optical elements used in chromatic aberration correction when focusing laser spots onto the sample or camera.

It is important to always be mindful of the fact that we only see what the camera sees, which may be different from what the sample sees. For example, for the configuration depicted in Figure 3-3, both pump and probe spots are in focus on the sample, but only the probe looks to be in focus on the camera, while the pump looks to be out of focus. Furthermore, it is important to remember that samples are typically specularly smooth and so act as mirrors. This means that

it is also possible to accidentally focus the camera on a virtual image of the spot that is slightly above or below the true surface of the sample. In this case, the spot will look to be in focus on the camera, when it is not actually in focus on the sample. These two subtleties make properly focusing the pump and probe a non-trivial process that requires the experimentalist's full attention.

The key to our chromatic aberration correction protocol is the use of wavelength-specific LED flood illumination to image the sample. The diffuse flood lighting injected into the optical path before the camera, as shown in Figure 3-1, is an LED centered on a specific wavelength, which is interchangeable for other LEDs. We use a green LED whose wavelength is exactly matched to the pump beam of 530 nm, and an IR LED at 850 nm, which is close to the wavelength of the probe beam at 800 nm. We ensured that the 50 nm discrepancy between our IR LED and the probe did not cause any measurable error. We confirmed this by performing the focusing protocol described below using the 850 nm LED with the probe beam set to 800 nm, and then afterward shifting the output wavelength of the probe up to 850 nm without changing any optics. No change in the focus of the spot was observed. The chromatic aberration correction protocol is as follows:

1. Using the green LED, adjust the sample position until the sample surface features are in sharp focus on the camera screen, while illuminated by the green LED. We use an Al coated Si sample for this, with very shallow but distinct surface features, so that there is no ambiguity in the focal plane resulting from finite heights of the features used to focus the sample. Step 1 confirms that the sample is in focus for 530 nm light, as viewed by the camera. The exact tube lens focusing ring position is clearly marked with tape to avoid gradual focal plane drift with repeated measurements.

2. We turn off the illumination, block the probe beam, and adjust the beam expander lens of the pump beam until it is in sharp focus on the camera. Because step 1 ensured that any 530 nm light that is in focus on the camera is in focus on the sample, this step ensures that the pump beam is now properly focused on the sample.

3. We now swap out the green LED for the IR LED that is close in wavelength to the probe beam, and illuminate the sample with 850 nm diffuse flood lighting. The sample surface will appear out of focus, because the sample and camera are currently set up to make features on the sample illuminated by 530 nm light be in focus. At this point, the sample position and the pump's beam expander are properly set such that the pump beam is properly focused onto the sample. Therefore, we must not adjust either of those components, or else we will disturb this alignment. Therefore, we adjust the tube lens in front of the camera to bring the sample, now illuminated with IR light, into sharp focus. This does not affect the alignment of the pump, but it means any IR light that appears in focus on the camera is also in focus on the sample.

4. We turn off the illumination, block the pump beam, and adjust the beam expander lens of the probe beam until it is in sharp focus on the camera. Because step 3 ensured that any IR light that is in focus on the camera is in focus on the sample, this step ensures that the probe beam is now properly focused on the sample.

As a result of this protocol, we are now confident that both the pump and probe are simultaneously in sharp focus on the sample, even though it is only possible to view one of them in focus at a time. After the protocol is complete, neither beam expander is touched for the rest of the measurements, until the objective lens is changed and a new chromatic aberration correction must be dialed in. We found that the 2x, 5x, and 10x objectives all had the same level of chromatic aberration, while the 20x required separate calibration.

When it is time to take pictures to measure spot sizes, the tube lens of the camera is adjusted until the beam being photographed is in sharp focus, while the other beam is being blocked. The camera focus must therefore be adjusted between every pump and probe picture. Because the beams are focused on the sample, adjusting the camera to bring the beams into focus is also bringing the sample into focus for that wavelength of light, and so the correct spot size as felt by the surface of the sample is photographed. We would periodically swap either the IR or green LED back in to illuminate the sample, approximately every 10 measurements, after having brought the corresponding wavelength spot into focus with the tube lens. The sample was always still in sharp focus and a discrepancy was never observed, indicating that there was no measurable drift of the alignment configuration over time or with repeated re-focusing of the camera to take pictures for multiple consecutive measurements.

While this approach does mean that the focal length, and hence magnification, for pictures of the pump and probe are slightly different, the difference in focal length is on the order of single microns. This is enough to defocus and distort a spot image if not accounted for when compared to the narrow depths of field that are also on the order of single microns, but compared to the much longer working distances that are on the order of tens of millimeters, the change in magnification is negligible and well below detection threshold. Similarly, there may be concern that with repeated adjusting of the sample position and camera tube lens, the focal plane may gradually "walk" away over time. However, the camera tube lens and beam expander lenses all have a fixed range of travel corresponding to a total possible focal length change of only a few microns. This is insufficient to affect magnification calibrations by a measurable amount. If one finds themselves hitting the ends of the threads on any of the lenses, this is an indication that the focal plane has drifted over time, and one can walk it back to center by alternately over-correcting the focus with the sample position and then bringing the image back into proper focus with the camera tube lens, until all optical elements are back near the middle of their threads and range of travel.

One cautionary warning: before changing to a new location on the sample, or to a new sample, it is important to always first re-focus the camera to either the pump or probe wavelength, and remember for which one it is focused. Then, after moving to a new sample or location, adjust the sample position along the optical axis until the spot is back in focus. This way, you know that you have correctly placed the sample surface at the plane in real space where the pump and probe beams are focused. If the camera focus is accidentally knocked out of focus before changing locations, then there is a risk of adjusting the sample position until the spot appears to be in focus, when in reality the spot is not properly focused on the sample and the camera is focused on a virtual image of the spot reflected off the sample's mirror-like surface. We use tape to clearly mark the position of the tube lens focusing ring for the green light focus to help mitigate such risks of focal plane drifting.

3.4.3 Benefits of Our Spot Size Measurement Technique

While conceptually more complicated, our technique for focusing pump and probe and for measuring spot sizes offers several benefits over the more traditional sweeping knife-edge (SKE) method. These benefits and some drawbacks are now discussed.

Our method does not require removing the sample. The SKE method requires physically removing one's sample and replacing it with the knife-edge, such that the lasers can be focused onto the knife-edge as it is swept through the beam. Because we do not need to remove our sample to measure the laser spot, we can also measure the spot size for each individual measurement. For the SKE, one initial spot size measurement is performed, and then one assumes that there is no drift or variation of the spot size over the course of the subsequent measurements.

Our method is also cheaper, as no additional equipment is necessary. The SKE equipment is not prohibitively expensive, especially as compared to more expensive components of a TDTR system, but off-the-shelf beam profilers can cost a few thousand dollars, as can motorized piezoelectric XYZ stages and drivers for researchers who wish to build their own DIY SKE.

Our method is robust against slight angles in the direction of translation of the driving stages for an SKE. If there is any unintentional angle such that the blade in an SKE is not exactly orthogonal to the optical axis, or in the linear actuator stage driving the SKE, then the blade may have a small component of translation parallel to the optical axis during the sweep. This would then result in an incorrectly measured spot size and shape, as the blade would not remain in the focal plane throughout the measurement. However, such imperfections typically result in negligible spot size measurement error for a carefully aligned SKE implementation.

Our method avoids errors from knife-edge diffraction effects when measuring very small spot sizes. This is of particular concern for systems such as ours that use very small spot sizes that approach the diffraction limited spot size. For such spots, the knife-edge effect can cause the incident light to diffract around the shape knife-edge, potentially missing the detector or causing interference fringes if some of these diffracted rays then reflect off the underside of the finite thickness of the knife-edge. This, or any other effects such as an aggressive convergence angle, which may cause some of the transmitted light to occasionally miss the detector, will result in an incorrect spot size and shape measurement.

Our method also correctly measures the ellipticity of the spot, whereas the SKE approach assumes that the minor and major axes of an elliptical spot are perfectly aligned to the sweeping directions of the knife-edge, which in general will not be true. This can result in considerable error in the effective isotropic spot size that is calculated for the thermal model, as demonstrated in Figure 3-4.

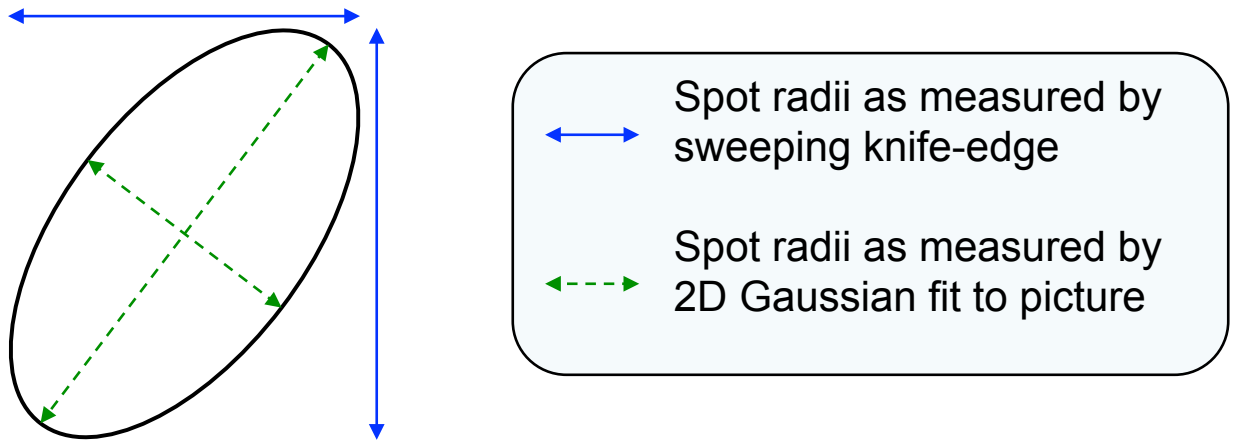

Figure 3-4 Measurement of an arbitrarily oriented elliptical laser spot. Solid blue double arrows indicate the spot radii that would be measured by the sweeping knife-edge approach. Dashed green double arrows indicate the spot radii that are measured using our method of fitting an arbitrarily oriented 2-D Gaussian function to a picture of the laser spot. The effective isotropic $1/e^2$ radius calculated from the sweeping knife-edge approach can substantially over-estimate the laser spot size.

One drawback of our method is that precision of spot size measurement can be limited if the spot fills only a very small fraction of the camera viewing window, and hence view few if any pixels. Fitting a Gaussian to the spot enables sub-pixel resolution in determining the spot size. However, if the $1/e^2$ radius is much less than a single pixel on the camera, then the precision of measurement will be inherently limited. Choosing a tube lens that increases the magnification of the image in the camera can mitigate this problem. In our system, because the camera is positioned behind the objective lens, the number of pixels spanned by the laser spot (~5 – 10 pixels) is approximately independent of the objective used and hence independent of the spot size. Therefore, the accuracy of our spot size measurements as a percentage of the $1/e^2$ radius is constant for all spot sizes, and does not decrease for smaller spots.

3.4.4 Pump-Probe Overlap

For our TDTR experiments, it is necessary to have the pump and probe beams perfectly overlapped. This is greatly helped by the coaxial nature, enabled by the dichroic mirror that injects the pump beam into the same optical path as the probe beam, through the objective lens. However, fine alignment is still necessary. The position of the beams can be controlled with fine adjustment knobs in their alignment mirrors. We again use the camera for precise control of beam position. The camera software allows us to place pixel-precise crosshairs to mark positions on the viewing area and to designate the positions through which intensity profile cuts are to be made (these are the same intensity profile cuts used to ensure that no camera pixels are being saturated prior to taking spot pictures). When the camera is focused on the probe spot, the crosshairs are placed on the central brightest pixel of the probe, and its position is carefully perturbed until the four corner pixels surrounding the central crosshair pixel are illuminated with equal brightness. The digital zoom feature of the camera helps greatly with this process. The variable gradient ND filter is used to suppress the intensity of the spot to the appropriate level such that only these central few pixels are significantly illuminated. Then, the probe beam is blocked, the pump beam is unblocked, and the camera is adjusted until the pump beam is in

sharp focus. The pump beam's position is then adjusted until it meets this same criterion, without moving the crosshairs. Using this method allows for repeatable and precise pump-probe spot overlap to an accuracy better than 1/8 of the $1/e^2$ diameter. This accuracy is limited primarily by our ability to make fine adjustments to the mirrors.

We validated this protocol by perturbing the pump position while continuously monitoring the strength of the TDTR signal, and observing that we were consistently at a global maximum for signal strength after applying the above protocol. This protocol also confers most of the same benefits over using an SKE measurement to overlap pump and probe as were described above in section 3.4.3.

3.5 Sample Preparation and Characterization

Before depositing the transducer layer, samples were polished using Chemical Mechanical Polishing (CMP) to achieve a specular, flat surface. For our samples, we used 70 nm thick Al as our transducer layers. The transducer layer was deposited using an electron beam evaporator for uniform and pin-hole free layers. Prior to deposition, samples were carefully cleaned by means of 30 minutes of sonication in an acetone bath, followed immediately (without letting the sample dry) by a methanol wash and then an IPA wash. The samples were blown dry with compressed dry nitrogen gas. Half an hour before deposition, a hydrofluoric acid (HF) dip was performed to remove native oxide layers on the samples. The HF dip was performed using a 1:10 solution of HF:water, dipped for 30 seconds at room temperature, immediately followed by a DI water wash and then air dried. The evaporation chamber (the "CHA" evaporator at the UC Berkeley Marvell Nanofabrication Lab) was pumped down over 1.5 hours to a base pressure of approximately $8x10^{-7}$ mTorr. The aluminum melt was thoroughly degassed before opening the shutter to begin deposition. This was accomplished by monitoring the chamber pressure spike after the metal began to melt, and waiting until this pressure had stabilized back down to base pressure before beginning deposition. The deposition was initiated with a high deposition rate, typically around 10 Å/sec, to promote metal condensation nucleation sites and encourage uniform deposition and smooth films. The deposition rate was then kept between 5 to 10 Å/sec until the desired film thickness was reached, as measured by a crystal monitor in the evaporation chamber.

Film thicknesses were characterized using AFM scans in tapping mode across scratches made in the metal film using tweezers. The scratch was inspected using an optical microscope to ensure that it penetrated through the entire metal layer, and had not damaged the sample surface below. The AFM scans further confirmed this by showing the uniform flat sample at the bottom of the scratch. Typical surface roughnesses on top of the Al film were in the range of 1 to 5 nm rms for nominally 70 nm thick films. Our samples did not have a strong or clear enough acoustic echo signal to be able to use picosecond acoustics [89], [94] to measure film thicknesses. The aluminum layers were assumed to have developed a 4 nm oxide layer from being stored in atmosphere [95], [96].

Bulk thermal conductivity cannot be assumed for such a thin layer of Al, whose thermal conductivity will be reduced from boundary scattering of the electrons. Instead, the electrical sheet resistance of the Al films was measured using a van der Pauw method. The sheet resistance was then converted into a thermal conductivity using the Wiedemann-Franz law [97], and taken to be isotropic. Only the electrically conducting portion of the film's thickness (i.e. AFM measured thickness minus ~ 4 nm oxide layer thickness) was used in this calculation. Al

film thermal conductivities ranged from ~100 to ~160 W/m-K at room temperature, consistent with the literature [98].

Bulk values of gravimetric heat capacity were assumed for our nano-grained silicon samples, as even the smallest grains were not small enough to affect the silicon density of states. The density was multiplied by 1-p, where p is the porosity of the sample. These are the same nano-grained silicon samples as were measured in [99], and their full characterization is explained there. For the SiGe samples, the virtual crystal approximation [100] was used to estimate the volumetric heat capacity as a simple average between that of silicon and germanium, weighted by their stoichiometric ratio.

3.6 TDTR Mathematical Theory

A complete derivation of the mathematical theory for the thermal model of a standard TDTR geometry is given in Refs. [27], [83]. Here, I will briefly review the key concepts and the physical picture behind the derivations. My primary goal in this section is to give the reader a physical intuition for how to interpret and understand each equation, and the assumptions and models on which they are built. Interested readers can then be directed to Refs. [27], [83] for the full formalism and mathematical details of each step. As the mathematical formalism of TDTR is quite involved, my hope is to provide a qualitative guide that new researchers can couple with the published literature to produce a more manageable first introduction to the theory.

I will first give a high-level overview of the derivation approach: one begins with the heat conduction Green's function for a time-periodic point heat source in a semi-infinite medium. One then uses superposition of this Green's function separately in time, t, in the through-plane direction, z, and in the in-plane direction, (cylindrical) r, in order to create a full three-dimensional steady-periodic solution for the TDTR geometry. One primarily works in the Fourier domain for the time dimension, Hankel transforms are used for the in-plane spatial dimension, and the Feldman algorithm [35] is used to decompose the solution into counter-propagating plane waves along the (real-space) z-direction in order to handle the multi-layered material geometry. Once a full frequency-domain response for a steady-periodic, Gaussian-distributed, pumped-probed laser system has been established, one incorporates the details of the probe and modulated pump laser pulse trains via a convolution of their respective frequency domain expressions, resulting in a Fourier series expansion to provide complete temporal detail.

I will now provide some further details on key steps. The Green's function for a time-periodic point heat source in a semi-infinite medium, as shown in Figure 3-5, is given by

$$g(r,\omega) = \frac{\exp(-qr)}{2\pi kr} \tag{3-2}$$

where r is the (spherical) radial coordinate, ω the (angular) heating frequency, k the thermal conductivity, and $q = \sqrt{i\omega/D}$ is the complex thermal wave vector with D being the thermal diffusivity. The function g can be thought of as a thermal transfer function in the Fourier domain, giving the amplitude of the periodic temperature rise at any point a distance r away from the periodic point heat source located at the origin. Later, r is separated into an in-plane cylindrical r and a through-plane Cartesian z, as shown in the coordinate system of Figure 3-5.

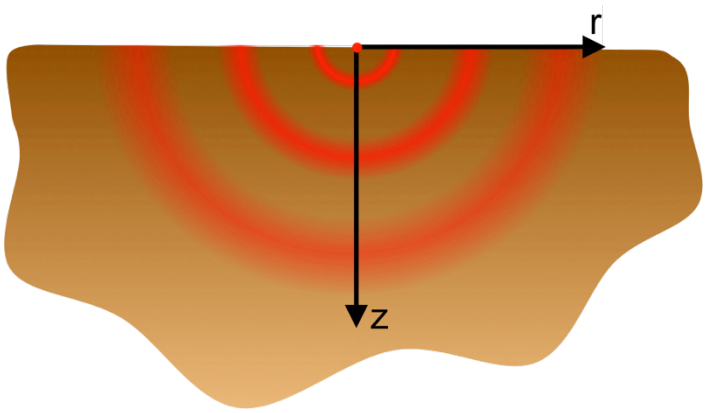

Figure 3-5 Time-periodic point heat source in a semi-infinite medium

This Green's function is non-trivial to derive, and can be found on page 263 of Ref. [101]. Note that Ref. [101] has divided out the heat liberation rate of $\rho c e^{i\omega t}$, which must be added back in to the listed equations to maintain proper units. The solution in Ref. [101] is for an infinite medium. To adapt this to a semi-infinite medium, the heat input is halved, as only half of the heat diffuses down into the bottom semi-infinite domain under consideration. Symmetry means that no other part of the solution must be altered, as we divide the infinite domain into semi-infinite domains by cutting along an adiabatic plane through the origin across which no heat is transferred.

In standard TDTR, the pump and probe beams are coaxially aligned with circular Gaussian intensity distributions. Because Eq. (3-2) gives the Green's function temperature response for a unit heat input, we can convolve it with a normalized Gaussian function in order to get the temperature response from laser heating. Because convolution in real space is equivalent to multiplication in reciprocal space, we take the Hankel transform of Eq. (3-2),

$$G(K,\omega) = \frac{1}{k\sqrt{4\pi^2 K^2 + q^2}} \tag{3-3}$$

where G represents the Hankel transform of the function g, and K (not to be confused with k for thermal conductivity) is the Hankel transformed conjugate variable to the in-plane real-space cylindrical coordinate r. Eq. (3-3) is then multiplied by the Hankel transform of a normalized Gaussian in order to paint out the response to a Gaussian heating profile from the Green's function of Eq. (3-2) and complete the convolution.

Because we measure temperature by the total integrated intensity of a reflected probe beam also with circular Gaussian intensity, we want to weight this convolved surface temperature response with another Gaussian in real-space to accurately predict the signal strength experimentally measured. After both convolving with a normalized Gaussian intensity distribution with the pump beam's $1/e^2$ radius and then also weighting with another normalized Gaussian with the probe beam's $1/e^2$ radius, we have [27]

$$\Delta T(\omega) = 2\pi A(\omega) \int_{K=0}^{\infty} G(K,\omega) e^{-\pi^2 K^2 (w_0^2 + w_1^2)/2} K\, dK \tag{3-4}$$

where A is the amplitude of pump power absorbed at frequency ω, G is from Eq. (3-3), and w_0 and w_1 are the pump and probe $1/e^2$ radii, respectively. Eq. (3-4) therefore gives the frequency domain amplitude of the surface temperature oscillations as measured by the probe laser in response to pump laser heating at the single frequency ω, and accounts for the full three dimensional details of the distributed laser heat source and its conduction into the semi-infinite substrate.

There are a couple reasons why we choose to take this Hankel transform instead of simply doing a regular convolution over r in real space. The first is that our system exhibits cylindrical symmetry, making it a natural fit for Hankel space, as Hankel transforms project functions using zero-order Bessel functions of the first kind, which can be thought of as our regular friendly sinusoidal functions in cylindrical coordinates. This therefore tends to produce well-behaved expressions when transforming functions that exhibit cylindrical symmetry (such as our Gaussian laser spots), which can then be handled analytically. The second reason is that by taking the Hankel transform of the Green's function, which represents a point heat source in real-space, we transform the fundamental heat input from a point source into a plane source. The point source that is infinitely "localized" at the origin in real space is infinitely "spread out" in reciprocal-space, as shown in Figure 3-6. This is analogous to how the Fourier transform of an origin-centered Dirac delta function in real space yields a constant value that is uniform over all of Fourier space.

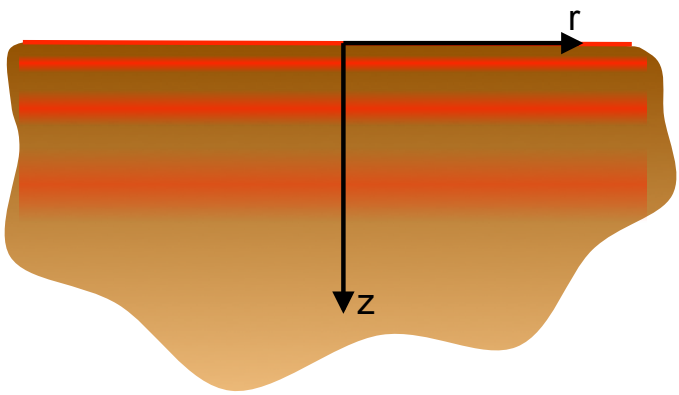

Figure 3-6 Time-periodic planar heat source in a semi-infinite medium as a result of taking the Hankel transform of a point source.

This transformation therefore has considerable utility, as we are now able to apply Feldman's algorithm [35], which only applies to one-dimensional heat conduction. Even though our solution is still three-dimensional in space, the two in-plane spatial dimensions are wrapped up in the Hankel transform, and so we can elegantly decompose the solution into one-dimensional plane waves propagating in z as required by the algorithm. Feldman's algorithm allows us to extend our solution of the heat diffusion equation in a single semi-infinite medium to an arbitrary layered stack of dissimilar materials, each of which can have its own thermal diffusivity, thickness, and thermal boundary conductance to above and below layers. The solution easily incorporates simple anisotropy between the through-plane and in-plane thermal conductivities. We can also place planar heat sources anywhere we like throughout the layered stack, although in our case we place the laser heat source as a boundary condition at the top of

the metal transducer layer, at the top of the stack. This also gives the flexibility to model layers both above and below the transducer layer such as is useful when measuring liquid samples [28]. In practice, applying the algorithm amounts to multiplying a series of second rank transfer matrices, corresponding to forward and backward propagating one-dimensional waves. The first type of matrix accounts for wave propagation through a material, and consistent progression of the phase and amplitude of the waves according to that layer's material properties. The second type of matrix accounts for the matching of waves at layer interfaces, which can include heat sources and boundary resistances. The mathematical details of this process for the particular application to TDTR are given in Ref. [27], and the derivation for the general case is in Ref. [35].

The temperature response of Eq. (3-4) is for heating at only a single frequency. In order to incorporate details of the pulsed nature of the pump and probe beams, their respective pulse trains are modeled as Dirac Combs as justified in section 3.2.1, with the pump pulse train multiplied by an oscillating envelope function at the modulation frequency. These temporal distributions are multiplied together (in practice, convolved together in the frequency domain), to yield a final solution that has all spatial and temporal details of the TDTR system properly modeled. The output of the lock-in amplifier is thus given by [27], [83],

$$Z(t) = \frac{\beta Q_{pump} Q_{probe}}{\tau^2} \frac{dR}{dT} \sum_{m=-\infty}^{\infty} \Delta T(\omega_0 + m2\pi/\tau) e^{im2\pi t/\tau} \tag{3-5}$$

where β is a constant representing the gain of the electronics, Q_{probe} is the probe laser intensity incident on the sample, Q_{pump} is the pump laser power absorbed by the sample, τ is the time between laser pulses (~12.5 ns), dR/dT is the Al film thermoreflectance (~10^{-4} K^{-1} for 800 nm light), ΔT is the frequency-domain amplitude of the temperature response for a given heating frequency described above, ω_0 is the pump modulation frequency, and t is the time by which the arrival of the probe is delayed relative to the arrival of the pump (note that t is not regular laboratory time).

Eq. (3-5) is complex valued and so carries information about both the amplitude and phase of the system's response to the pump laser excitation. The infinite sum in Eq. (3-5) originates from the Dirac Combs of the pump and probe pulse trains. Eq. (3-5) can be understood as taking ΔT, the fundamental temperature response in the frequency domain, and projecting it onto the delay-time (t) domain by means of a Fourier series expansion, with the time between laser pulses (τ) as the fundamental period of the expansion. It is important to remember that the variable t in Eq. (3-5) is not regular time, but is specifically the delay time of the probe pulse with respect to the pump pulse (which is typically chosen to arrive at a laboratory time of zero). In the limit where the sample completely returns back to thermal equilibrium before the arrival of the next pulse, the delay time (t) becomes equal to the laboratory time. This would require either a sample with very high thermal diffusivity and/or a laser with very long times (τ) between pulses. In practice, samples rarely fully return to equilibrium between pulses, and so there is an accumulation effect when the end of the thermal decay of a preceding pulse bleeds into the next pulse [83].

Although complicating to the modeling, this pulse accumulation effect is useful as it means the sample feels the effects of the modulation frequency envelope. This can be directly visualized in Figure 3-7, which is taken from Ref. [83].

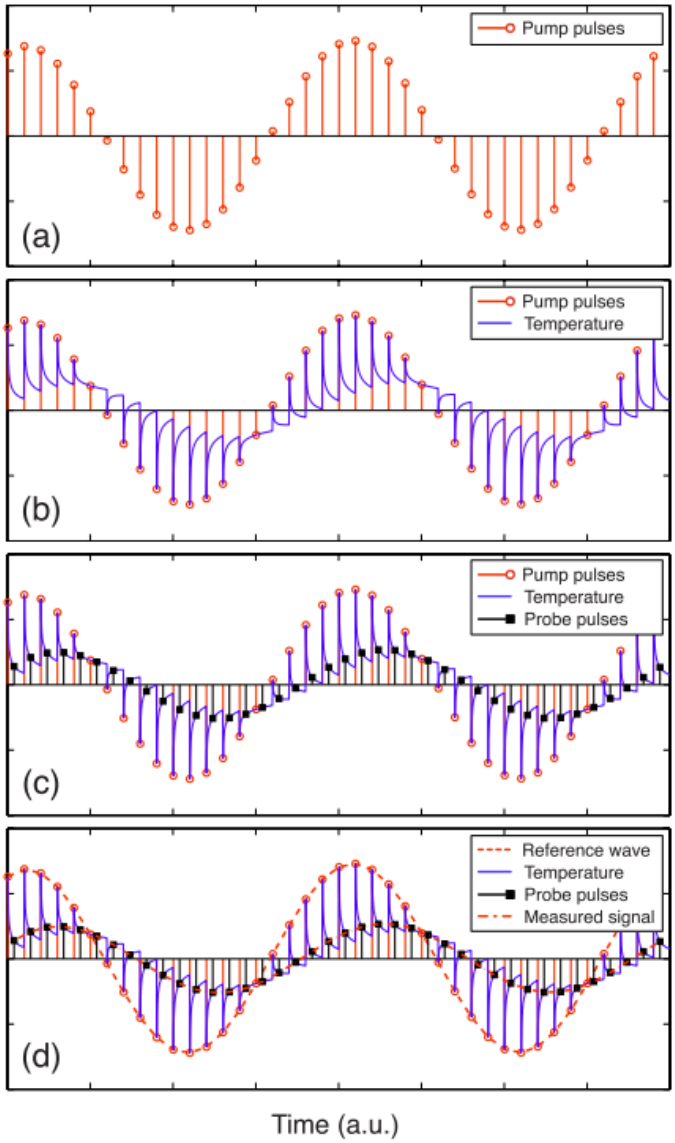

Figure 3-7 (a) The pump beam input to the sample modulated by the fundamental component of the EOM. (b) The surface temperature of the sample in response to the pump input. (c) The probe pulses arrive at the sample delayed by a time, *t*, and are reflected back to a detector with an intensity proportional to the surface temperature. (d) The fundamental harmonic components of the reference wave and measured probe wave. The lock-in amplifier records the amplitude and phase difference between these two waves at every delay time. This figure and most of its caption taken from Ref. [83]**.**

Figure 3-7(d) makes it easy to see that the probe response function measured by the lock-in amplifier and given by Eq. (3-5) has a different amplitude and phase lag relative to the pump laser modulation frequency. It is acceptable to only consider the fundamental (sinusoidal) component of the modulation (as the modulation is a square wave) as well as ignore DC offsets, because the lock-in amplifier in combination with the RLC filters discussed in section 3.3.3 pick out only this fundamental frequency component of the signal and reject all others.

3.7 Experimental Modeling and Data Analysis

3.7.1 System Modeling and Data Fitting

In principle it is possible to fit Eq. (3-5) directly to the output of the lock-in amplifier in Volts. In practice, however, this opens up susceptibility to uncertainty in all of the coefficients of Eq. (3-5). Additionally, the pump and probe power can drift throughout a measurement due to instabilities in the laser or varying coupling strengths of the probe beam throughout the range of motion of the delay stage. It is therefore usually recommended that one instead fits to the negative ratio of the in-phase (X) and out-of-phase (Y) components of the lock-in amplifier [85]. Mathematically, the in-phase signal is the real component of Eq. (3-5), while the out-of-phase signal is the imaginary component, as was a similar case for the 3-Omega system. This offers the considerable advantage of canceling out all of the coefficients in Eq. (3-5), which therefore no longer need to be known and whose uncertainties do not affect the final measured results. This is much like fitting to the phase as a function of delay time. Therefore, we fit our data by equating experimental and theoretical quantities, according to

$$\frac{-V_{in}}{V_{out}} = \frac{-\mathrm{Re}\{Z(t)\}}{\mathrm{Im}\{Z(t)\}} \tag{3-6}$$

where V_{in} and V_{out} are the in-phase (X) and out-of-phase (Y) outputs from the lock-in amplifier, respectively, and $Z(t)$ is given by Eq. (3-5). Taking the negative is not strictly necessary for data fitting, but it makes the resulting plotted data positive-valued and decaying toward zero, as is consistent with our intuition for a decaying thermal transient.

We fit our data using a multi-variable non-linear least squares protocol, as implemented by the Python "curve_fit" function from the "optimize" library in SciPy version 0.17.1. This function uses the Levenberg-Marquardt algorithm when applied to optimizing unconstrained problems as in our case. The thermal conductivity of the sample and the thermal boundary conductance (TBC) between the Al transducer and sample are both simultaneously fit, with all other parameters being fixed inputs to the thermal model. Data was generally fit in the range of 0.5 ns to 3 ns of delay, although the exact interval was selected for based on a sensitivity analysis used to find the optimal region that maximized the signal-to-noise ratio for the desired measured quantities, and subject to the assessed goodness of fit discussed below. Extensive details of the sensitivity analysis are given in section 3.8 Sensitivity Analysis.

To assess the goodness of fit, we employed the Durbin-Watson statistic (DW) [102] in conjunction with the mean relative magnitude of the residuals, ε. The Durbin-Watson statistic is essentially the normalized autocorrelation of the fit residuals with a lag of 1 time step between consecutive observations. Conceptually, it quantifies in a consistent manner how much left over physics in the data leaked through to the residuals because the fitted model failed to capture it. For a perfectly fitted data set, the fit residuals should be completely randomly distributed with no predictive power. A value for the Durbin-Watson statistic less than 1 implies a positive autocorrelation of the residuals, meaning noticeably non-random trends in the residuals, suggesting that our model has failed to properly fit the data and is not correctly accounting for all of the physics that systematically affected the data. A value between 1 and 2 is ideal, and means the residuals are fairly randomly distributed without significantly discernible patterns. A value

over 2 typically indicates a negative autocorrelation of the residuals, and can be an indicator of over-fitting to noise in the case of polynomial or other general high-order fits. This is not a concern for our system, as we use a physics-based model with a fixed number of few parameters, and so cannot over-fit to noise in this way. Therefore, for our system, a Durbin-Watson statistic greater than 2 is still considered acceptable and an indicator of a good fit. We also utilized the relative mean of the residuals, ε, designating the average error of the model in relation to the data, as a second metric. The average value was taken of the absolute value of each residual, normalized by the magnitude of its corresponding experimental data point. Specifically,

$$\varepsilon = \frac{1}{N}\sum_{i=1}^{N}\left|\frac{y_i - \hat{y}_i}{y_i}\right| \tag{3-7}$$

where N is the number of measured data points, y_i is a specific measurement value, and $\hat{y}_i$ is the model-predicted value for that measurement. We took $\varepsilon \leq 1.5\%$ as signifying a good fit. An example fit is shown in Figure 3-8.

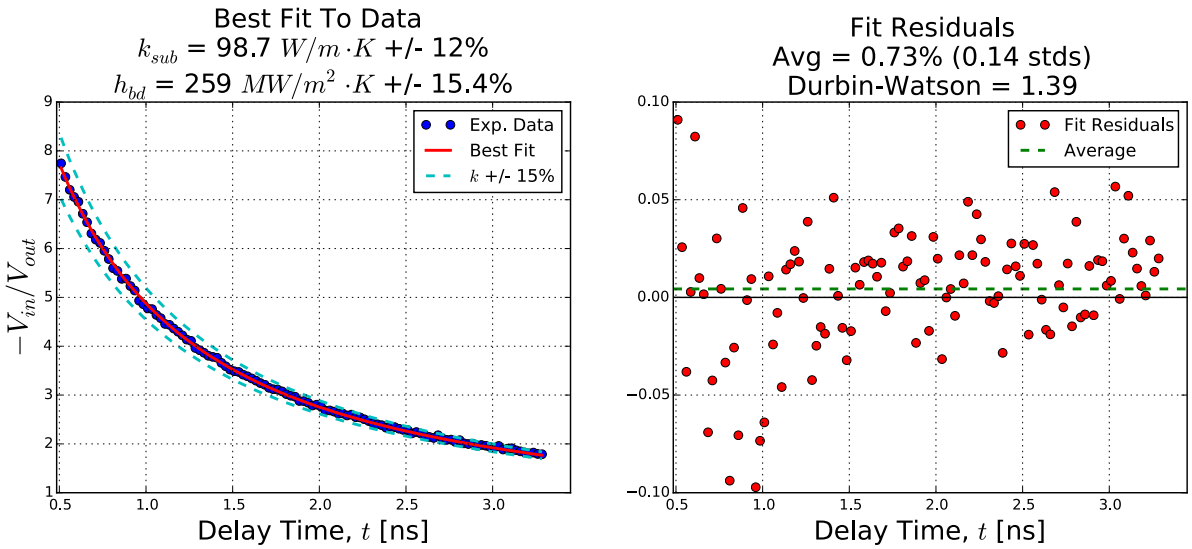

Figure 3-8 Representative good fit showing raw data, best fit line, theoretical curves for a k ± 15% as compared to the best fit value, residuals of the fit, and their calculated values of DW and ε. This fit is for single crystal silicon at T = 297 K and w = 2.9 µm.

DW and ε are complimentary metrics that compensate for the shortcomings of the other. The Durbin-Watson statistic does not take into consideration the variance of the residuals nor their magnitude as compared to the fitted signal magnitude. As a consequence, DW can be unduly strict for very low-noise and clean data, and overly forgiving of data that has a lot of random noise. On the other hand, ε specifically quantifies the relative magnitude of noise and the mismatch between fitted model and data, but is ignorant of any systematic trends in this mismatch. In brief, DW is poor at accounting for random error, while ε poor at accounting for systematic error. Therefore, measurements that failed both requirements of DW $\geq$ 1 and $\varepsilon \leq$ 1.5% were considered poor fits, and excluded from the results. Additionally, any measurements

that passed one metric, but were still especially poor according to the other metric (i.e. had either DW < 0.3 or ε > 3%) were also excluded from the results.

The system was modeled as a stack of two physical layers, with two additionally simulated effective layers to account for the interfacial TBC as well as the Al_2O_3 oxide layer and optical absorption depth of the laser. The bottom layer, being the sample, was modeled as semi-infinite in thickness. This is a valid assumption, as even the thinnest samples (~ 500 μm) were still approximately 2 orders of magnitude thicker than our largest thermal penetration depths (~ 5 μm), ensuring no part of the thermal signal was influenced by the backside boundary conditions of the sample. Heat capacity values were taken from literature and interpolated at the sample's temperature (which included a calculated background DC temperature rise due to the pump and probe lasers), and the thermal conductivity was one of two fitting parameters. The top layer modeled the Al film using temperature-dependent literature values for the heat capacity, and with thermal conductivities calculated from measured electrical sheet resistances. More complete details on these input parameters are given in section 3.5 Sample Preparation and Characterization.

The thermal boundary conductance (TBC) between the Al film and sample was modeled using a layer with vanishing thickness (1 nm) and vanishing heat capacity (1 J/m^3-K). The latter ensured the layer contributed no "thermal inertial" effects (appropriate as a true interface has no "thermal mass"). The thermal conductivity of this TBC layer was the second fitting parameter, as is standard [27]. Holding this layer's thickness constant and fitting k_{TBC} is therefore effectively fitting the TBC (W/m^2-K) as it is the ratio of k_{TBC} and the layer thickness.

The thermal model used to fit the data treats the impinging laser heat flux as a boundary condition at the top of the Al film. In reality, the laser energy is volumetrically absorbed over the top few nm of the film, spread out over the optical absorption depth of Al for the wavelength of pump light used. To harmonize these two scenarios, the top 10 nm of the Al film was collapsed into a 1 nm thick effective layer, with ten times the true heat capacity and thermal conductivity, following the recommendation of Ref. [27]. This effectively compresses the heat absorption into a boundary condition, while preserving the true heat capacity per unit area and lateral thermal diffusivity of the Al layer.

The thickness of Al layer was also reduced by 4 nm, representing the 4 nm thick Al_2O_3 native oxide layer that grows on Al. The heat capacity from this 4 nm of oxide was added to the 1 nm effective absorption layer described above in order to preserve the correct total heat capacity per unit area (J/m^2-K) of the combined Al film + oxide layer. This areal heat capacity is an important parameter for the system as discussed in section 3.8 Sensitivity Analysis, and failure to account for the increased heat capacity due to the oxide layer can skew the measured sample thermal conductivity value by as much as 10%.

The out-of-phase (aka. imaginary) component of the temperature response is constant across the zero delay time (t = 0) point when the laser pulse hits. In general, however, the experimentally detected phase will be incorrect due to phase contributions from the electronics that are variable and difficult to quantify. The best solution is therefore to correct the phase in post-processing by forcing this condition to be true [27]. We corrected the phase of our data so that the average value of the out-of-phase component over the range -10 ps to -5 ps is equal to the average value over the range +5 ps to +10 ps (i.e. just before and just after the zero delay time point).

Aluminum (our transducer material of choice) has an optical thermoreflectance around 10^{-4} K^{-1} for 800 nm light (our probe beam wavelength) [103]. Our measurements often produced

surface temperature oscillations with amplitude ~ 0.1 K at the modulation frequency. Therefore, we routinely measure a signal on the order of 10 ppm (10^{-5}), corresponding to 1 μV – 1 mV, depending on the laser powers and objectives used. Due to the globally reduced sensitivities, reduced relative sensitivity of k_{sample} with respect to C_{Al} and L_{Al} (discussed in section 3.8 Sensitivity Analysis), and lower absolute magnitude of sample temperature oscillations at high modulation frequencies, combined with the poorer quality of the high frequency RLC resonant circuit, the signal for the 9.5 MHz measurements was consistently more than an order of magnitude weaker than for the two lower frequencies while experiencing the same magnitude of noise. While curves could be fit to these measurements, they were deemed unreliable due to the poor signal-to-noise ratio for a majority of measurements. Therefore for consistency, all 9.5 MHz measurement data was excluded from the results. This is an area for improved instrumentation for future measurements.

3.7.2 Uncertainty Analysis

Analyzing TDTR data requires a multi-parameter non-linear fit to a complicated thermal model with many input parameters. Therefore, there are many possible sources of uncertainty in the measurement and it is important to take great care in considering all of them. The uncertainties used in the error propagation for all of our input parameters are summarized in Table 3-1. Conventional error propagation uses partial derivatives of the measured quantity with respect to the parameter under consideration to quantify the uncertainty contribution of the latter to the measured quantity. However, this approach assumes all sources of uncertainty are independent and it otherwise quickly becomes complicated and tedious to account for covariances among all fitting variables and system parameters. This approach also assumes normally distributed noise, and often fails for large noise levels or other situations where this assumption breaks down. Therefore, we chose to propagate uncertainties by performing multiple fits to the same data set with perturbed inputs. For each fit, we perturbed a single input parameter by its level of uncertainty while holding all other inputs constant, and observed the effect this had on the values of the fitted variables. Although computationally more expensive, this approach captures the effects of covariance, non-normally distributed noise sources, multi-parameter fitting, and noise on the fitting itself, and is a more realistic estimate of the true uncertainty in our measured quantities. We also calculated the uncertainty (standard deviation) of the fit itself from the square root of the diagonals of the corresponding covariance matrix. All calculated uncertainties were combined in quadrature to calculate the total uncertainties of each fitted variable.

Table 3-1 Input Parameter Uncertainties in TDTR Measurements

Input Parameter	Symbol	Uncertainty	Quantification of Uncertainty
Laser Spot Sizes	w	10%	Comparison to sweeping knife edge measurements
Sample Temperature	T	2%	Room temperature consistency
Signal Phase	ϕ	0.2 degrees	Phase correction consistency and sensitivity to noise
Al film thermal conductivity	k_{Al}	3% to 8%	Propagated Van der Pauw sheet resistance measurement uncertainties
Al film volumetric heat capacity	C_{Al}	5%	Best guess
Al film thickness	L_{Al}	2% to 6%	AFM measured surface roughness
Sample heat capacity	C_{sample}	5%	Best guess

The relative contribution of uncertainty from the different input parameters varied for measurements in different regimes of spot sizes, modulation frequencies, and sample thermal conductivities. Generally speaking, however, the dominant source of uncertainty was the laser spot size, followed by the heat capacity and thickness of the Al transducer layer. This is to be expected, given the results of the sensitivity analysis presented in Figure 3-9 and its related discussions in section 3.8.

In addition to careful uncertainty propagation, all measurements were repeated up to 9 times. Repeated measurements were completely independent. Although they were performed under consistent experimental conditions (i.e. same modulation frequency, spot size, and sample), they were performed at different physical locations on the sample, each with independently measured laser spot sizes, and sometimes on different days. The standard deviation of all repeated measurements was calculated as an empirical measure of the measurement uncertainty. Between this empirically measured standard deviation and the theoretically propagated uncertainty described above, the larger of the two on a case-by-case basis is reported as the measurement uncertainty and is used for all error bars for reported measurements in Chapter 4 - Effects of Alloying and Nanostructuring on Spectral Phonon Mean Free Path Distributions (where each point is the average of all repeated measurements). This maximum error was typically around 10% to 20%, and got as high as ~40% for a select few of the low frequency, low thermal conductivity, small spot size measurements that are extremely susceptible to uncertainties in laser spot sizes. This approach offers a realistic estimate of the true measurement uncertainty, erring on the side of being overly conservative. By using the larger of the two errors between the empirically and theoretically derived quantities, we avoid under-estimating the uncertainty due to accidentally under-predicting the uncertainty of input parameters, taking too few repeated measurements, or statistical anomalies causing non-representatively tight bunching of repeated measurements. We also increase our chances of capturing both random and systematic sources of error. In the majority of cases, the empirically measured and theoretically predicted uncertainties were within a factor of 2 of each other,

indicating that we have not overlooked any dominant sources of uncertainty, nor left out any crucial physics in our thermal model.

3.8 Sensitivity Analysis

An important step to designing a TDTR experiment and optimally analyzing its data is generating a sensitivity plot. Here, I define sensitivity the same as in Ref. [104],

$$S_x^R = \left| \frac{\partial \ln(R)}{\partial \ln(x)} \right| = \left| \frac{x_0}{R_0} \frac{\partial R}{\partial x} \right| \tag{3-8}$$

where S_x^R is the sensitivity of measured value R to parameter x, and 0 subscripts denote property values for the particular measurement. This definition yields the relative (always dimensionless) sensitivity, meaning that $S_x^R = 0.3$ implies that if the value of parameter x in an experiment were doubled, the resulting measured value of R would increase by 30%. We take R to be the negative ratio of the real to the imaginary component of the voltage signal, as this is the measured quantity in our experiment. We calculate sensitivities by taking numerical partial derivatives of Eq. (3-6). By taking the absolute value, we lose information about the sign of the sensitivity, meaning we do not really know if doubling x made R increase or decrease by 30%. However, for the purposes of designing an experiment, usually only the magnitude of the sensitivity is relevant. Taking the absolute value makes comparing relative magnitudes of sensitivities to various parameters in a plot much easier, as is evident in Figure 3-9, and so we choose this convention.

An ideal experiment has maximum sensitivity to the parameters under investigation, and minimum sensitivity to all other parameters. If the sensitivity of an experiment to a parameter under investigation is too low, then the influence of that parameter on the measured signal may be comparable to or weaker than background noise levels. In this case, it will not be possible to resolve the influence of the parameter and accurately measure it. Estimating these values before preparing samples and attempting a futile experiment can save tremendous time. Such estimates, for the purposes of quantitative comparison to noise levels, should be done using semi-absolute sensitivity rather than relative sensitivity. This is calculated by omitting the natural logarithm function surrounding R in Eq. (3-8). The resulting sensitivity will yield an absolute ("dimensionful") change in R for a relative change in parameter x. For example, let us take R to be the in-phase voltage response, and x to be the thermal boundary conductance (TBC) of the sample. A calculated semi-absolute sensitivity of 5 x 10^{-6} V means that if the value of the TBC were doubled, we would expect our in-phase voltage signal to change by 5 μV. If the measured in-phase voltage signal (i.e. after lock-in detection) contains random noise with magnitude of order 2.5 μV, then at best we might hope to determine the TBC to within +/- 50%, and realistically worse than that.

Alternatively, if the sensitivity to the other input parameters is too high compared to the sensitivity of the parameter under investigation, then uncertainties in the other input parameters can be significantly amplified in their effects on the measurement uncertainty of the parameter under investigation. This can be seen explicitly by considering

$$\frac{\sigma_x}{\sigma_p} = \frac{\partial\left(x/x_0\right)}{\partial\left(p/p_0\right)} = \frac{\partial \ln\left(x\right)}{\partial \ln\left(R\right)} \frac{\partial \ln\left(R\right)}{\partial \ln\left(p\right)} = \frac{S_p^R}{S_x^R} \tag{3-9}$$

where x is the parameter under investigation, p is any other input parameter of the model, σ denotes the relative uncertainty of a value, and all other quantities are as defined for Eq. (3-8). The chain rule has been used to show that the amount of measurement uncertainty (σ_x) in x, given an amount of uncertainty (σ_p) in input parameter p, is simply given by the ratio of their relative sensitivities. Contributions from multiple input parameters can be added up in quadrature as with standard uncertainty propagation. For example, if an experiment's sensitivity to parameter p is equal to twice its sensitivity to parameter x, then a 10% uncertainty in the input value of p used to fit a model to experimental data in order to calculate x will result in a 20% uncertainty in the measured value of x. This is a second measure by which a sensitivity analysis can be used to help predict the efficacy of an experiment *a priori*.

Figure 3-9 provides a comprehensive set of systematic TDTR sensitivity calculations representative of the majority of realistic TDTR measurements, including our own. Figure 3-9 can therefore be used as a reference tool in designing TDTR experiments. I will devote considerable space in this section to discussing different trends present in Figure 3-9 in order to shed light on their physical origin. A solid physical intuition for the complicated interplay of these various parameters is crucial for planning and designing effective TDTR experiments, and it is worthwhile for prospective researchers to spend some time familiarizing themselves with these physical relationships.

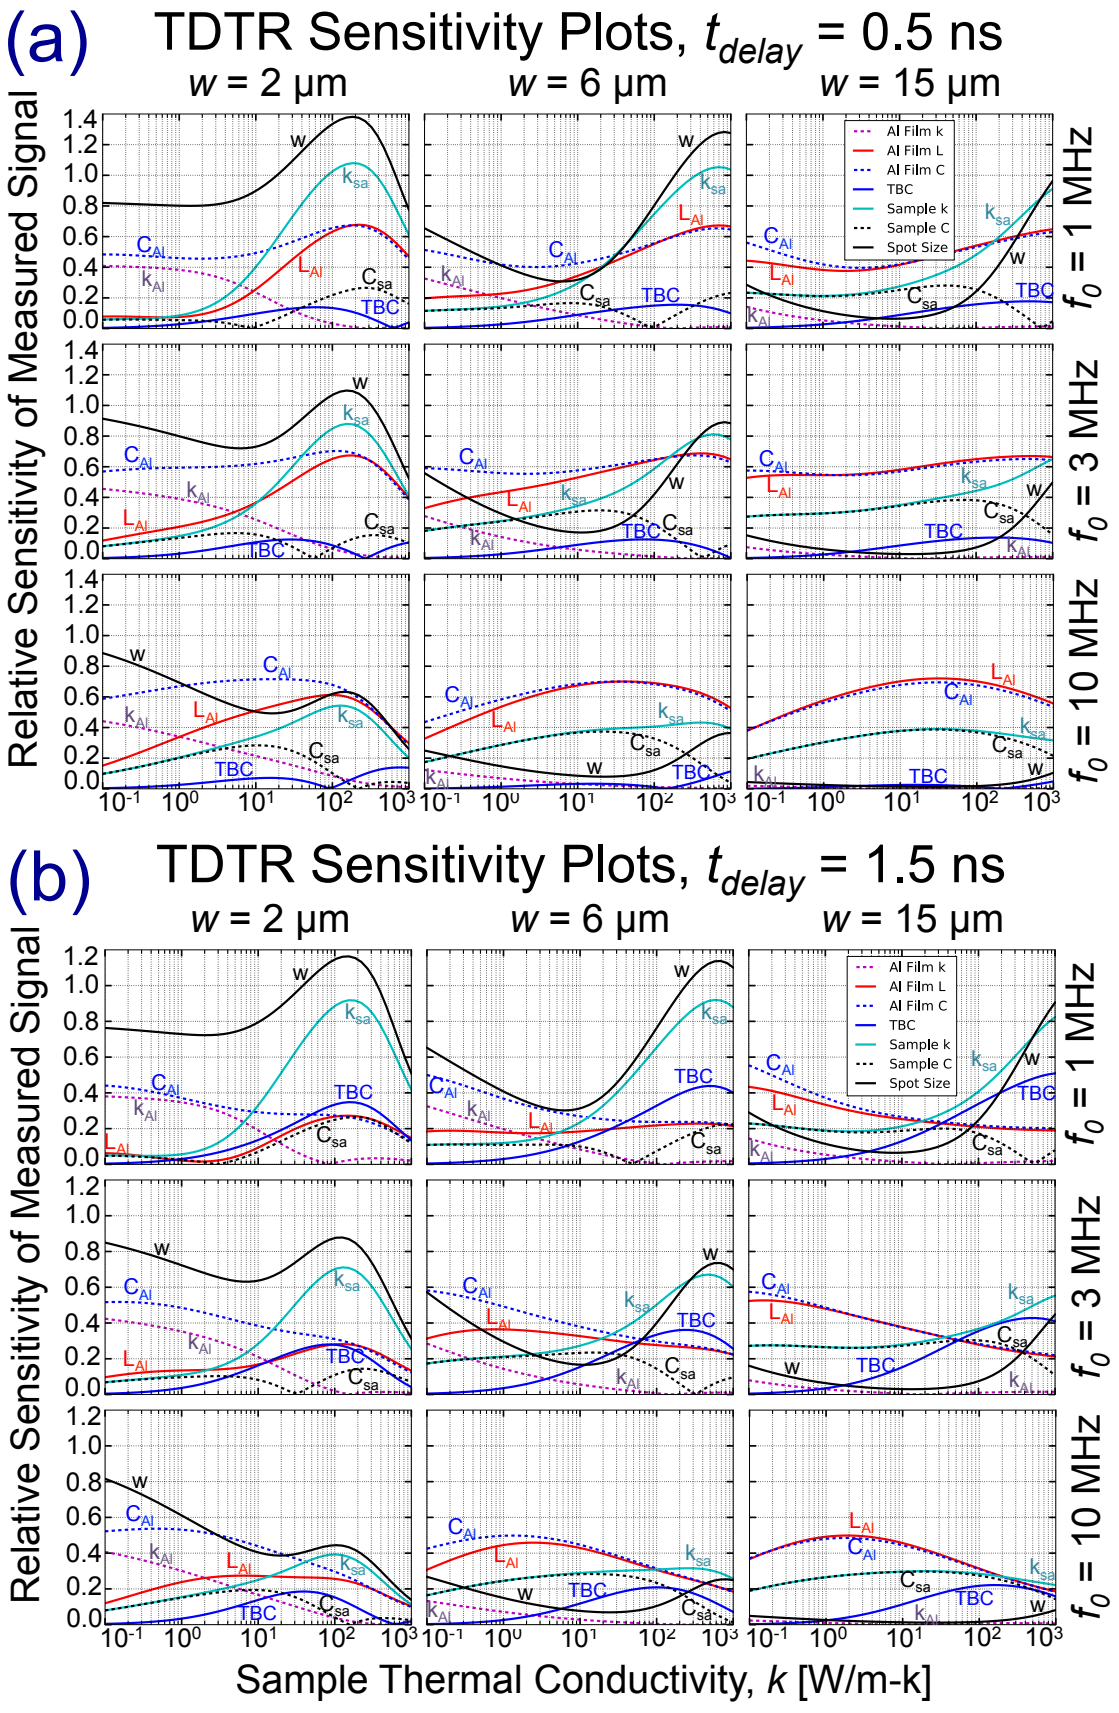
(a)
TDTR Sensitivity Plots, t_{delay} = 0.5 ns
w = 2 μm
w = 6 μm
w = 15 μm
Relative Sensitivity of Measured Signal
f_0 = 1 MHz
f_0 = 3 MHz
f_0 = 10 MHz
Al Film k
Al Film L
Al Film C
TBC
Sample k
Sample C
Spot Size
(b)
TDTR Sensitivity Plots, t_{delay} = 1.5 ns
w = 2 μm
w = 6 μm
w = 15 μm
Relative Sensitivity of Measured Signal
f_0 = 1 MHz
f_0 = 3 MHz
f_0 = 10 MHz
Sample Thermal Conductivity, k [W/m-k]

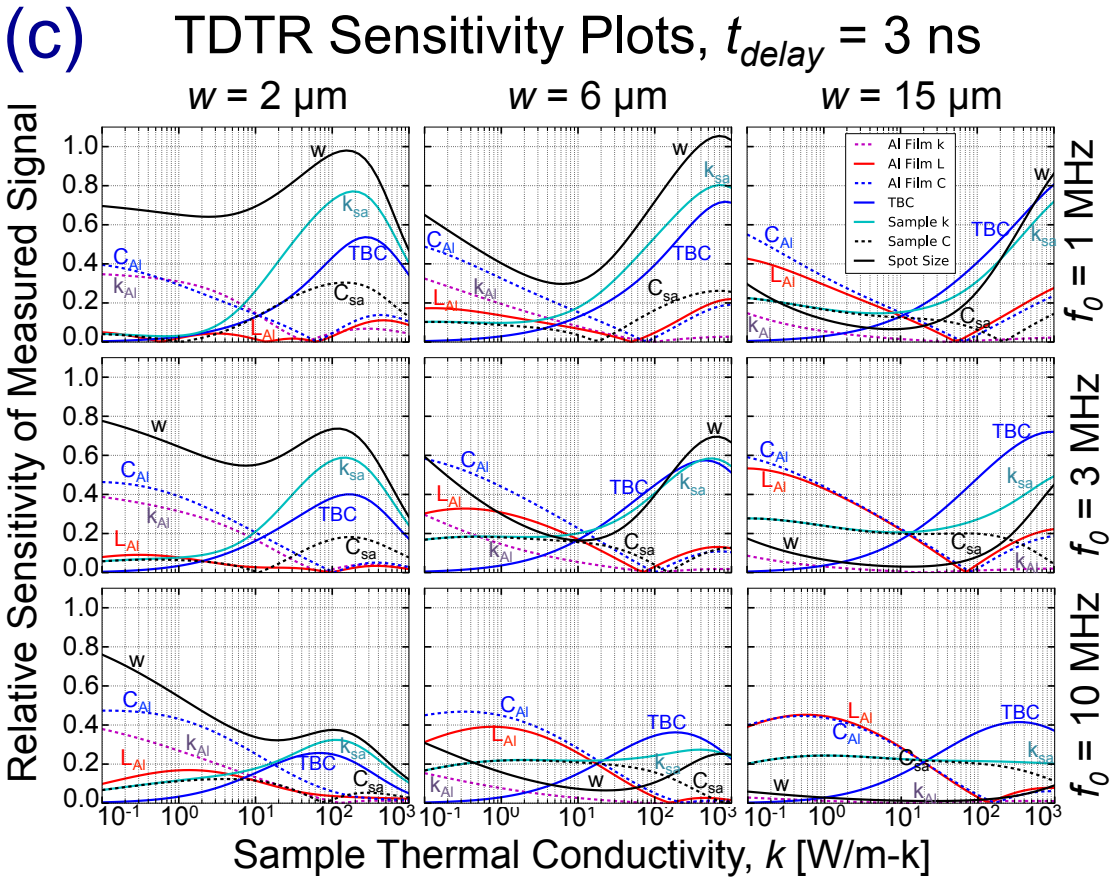

Figure 3-9 TDTR sensitivity plots at probe delay times of (a) 0.5 ns, (b) 1.5 ns, and (c) 3 ns. Sensitivities are calculated using Eq. (3-8) with *R* always taken as the ratio from Eq. (3-6) for a 70 nm thick Al transducer layer (k_{Al} = 130 W/m-K) with a 4 nm Al_2O_3 oxide layer; the sample is a semi-infinite material at 300 K with a volumetric heat capacity equal to that of Silicon (C = 1.6 MJ/m^3-K) and variable thermal conductivity; the thermal boundary conductance between the Al layer and sample is set to 100 MW/m^2-K. Sensitivities are calculated for a system with pump beam modulation frequencies, f_0, of 1 MHz, 3 MHz, and 10 MHz, for pump and probe beam 1/e^2 radii, *w*, of 2 μm, 6 μm, and 15 μm, and for a sample thermal conductivity, *k*, ranging from 0.1 W/m-K to 1000 W/m-K. Lines are plotted showing *R*'s sensitivity to *L*, *k*, and *C* of the Al layer, *k* and *C* of the sample, *w*, and the TBC. Note the slightly different y-axis scales among (a), (b), and (c).

One of the most immediately apparent features of Figure 3-9 is that the sensitivity to the thermal conductivity of the sample is often one of the highest sensitivities, particularly for high values of *k*. This is consistent with the historical success of TDTR as a general thermal conductivity metrology technique, particularly for high-*k* materials such as silicon. Another prominent feature is that for most of our measurements, the laser spot size has a strong if not dominant sensitivity. This is consistent with the laser spot size being among the largest sources of uncertainty in our measurements, particularly for our smallest spot sizes, and justifies our meticulous and labor-intensive spot size measurement protocol.

Several of the trends in sensitivities can be understood by considering heat spreading regimes. Specifically, considering whether heat conduction is primarily spherical or planar in nature, as determined by the relative magnitude of the laser spot size (w) as compared to the thermal penetration depth (L_p) in a material at the pump laser's modulation frequency. These two regimes are shown schematically in Figure 3-10. When the laser spot size is much larger than the thermal penetration, edge effects around the perimeter of the spot are negligible and the

heat conducts into the substrate primarily downward in a 1-D planar fashion. In the reverse limit, edge effects dominate and, from the thermal waves' point of view, the spot looks more like a point source and so the heat conducts into the substrate in a spherical fashion. In general, moving from the top-left panel toward the bottom-right panel in Figure 3-9 corresponds to moving from spherical to planar heat spreading.

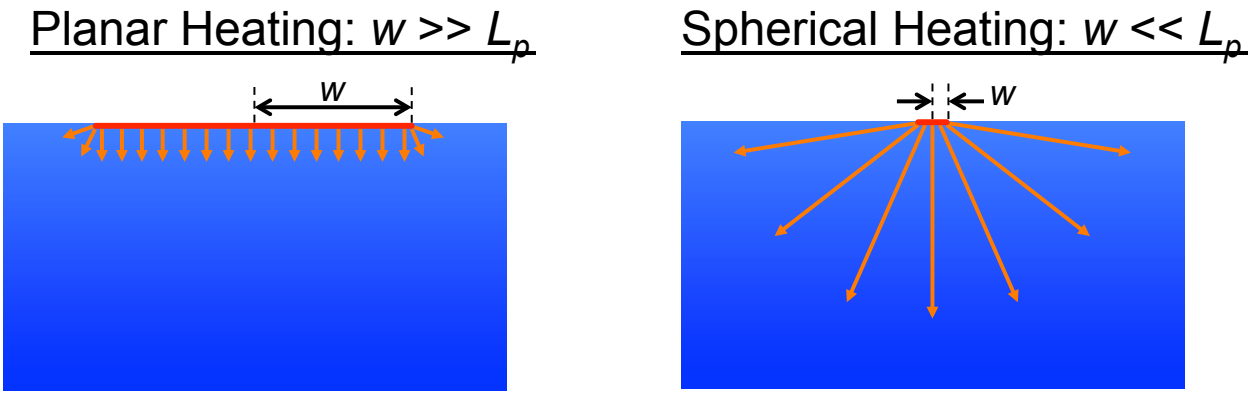

Figure 3-10 The two heating regimes for a TDTR experiment. Sample cross section view. Orange arrows denote the thermal penetration depth (L_p) into the substrate at the pump laser's modulation frequency, and w denotes the radius of the heating laser spots.

These two limits can be understood more quantitatively by considering the high and low frequency limits of the frequency-domain surface temperature response given by Eq. (3-4). The high frequency limit is given by [27]

$$\Delta T_{HF}(\omega) = \frac{A(\omega)}{\pi w^2 \sqrt{i\omega kC}} = \frac{A(\omega)}{\pi w^2 \sqrt{\omega kC}} e^{-i\frac{\pi}{4}} \tag{3-10}$$

where A is the amplitude of pump power absorbed at frequency ω, w is the $1/e^2$ laser spot radius (with pump and probe spots set to the same size), k is the substrate thermal conductivity, and C is the substrate volumetric heat capacity (J/m^3-K). The low frequency limit is given by [27]

$$\Delta T_{LF} = \frac{A_{DC}}{2\sqrt{\pi} kw} \tag{3-11}$$

where A_{DC} is the DC power absorbed from the pump. In the high frequency limit, the thermal penetration depth is much smaller than the spot size and so we expect the heat conduction to follow a planar geometry. Eq. (3-10) has been re-written to explicitly show the $\pi/4$ phase shift between surface heating and temperature response that is characteristic of planar periodic heating of semi-infinite media. Indeed, Eq. (3-10) is equivalent to the solution for planar periodic heating by a uniform heat flux equal to $A(\omega)/\left(\pi w^2\right)$. Similarly, in the low frequency limit, the thermal penetration depth is much larger than the spot size and so we expect the heat conduction to follow a spherical geometry, and indeed Eq. (3-11) takes the form of steady spherical heat conduction. Note that we would not be able to obtain a finite converged temperature response

for steady heating without spherical heat spreading, which is also consistent with the phase shift vanishing for spherical heating.

We can now examine the sensitivity trends through the lens of heating regimes. For all panels in Figure 3-9, the sensitivity to the sample thermal conductivity, k_{sa}, and volumetric heat capacity, C_{sa}, collapse into the same curve when increasing modulation frequency, increasing spot size, or decreasing sample thermal conductivity, while holding all other parameters constant. This happens because increasing w or f_0, or decreasing k_{sa} pushes the system more toward a planar heating regime. In this regime, the temperature response is only sensitive to the thermal effusivity, $\sqrt{k_{sa}C_{sa}}$, as shown in Eq. (3-10). Perturbing either k_{sa} or C_{sa} has the same effect on the signal, and so they have the same sensitivity. This means that uncertainty in C_{sa} propagates directly into uncertainty in k_{sa}, and also means that it becomes impossible to experimentally distinguish k_{sa} from C_{sa} and fit to both simultaneously. The sensitivity to k_{sa} itself exhibits a global maximum as function of varying k_{sa}. In either limit of very high or very low k_{sa}, the interface between the Al layer and the sample becomes a thermal boundary condition (adiabatic B.C. for very low k_{sa}, fixed temperature for very high k_{sa}). Therefore, the details of k_{sa} itself no longer significantly affect the heat transport, and so sensitivity to k_{sa} vanishes at either extreme. The thermal boundary conductance (TBC) and Al layer thermal conductivity dictate what qualifies "very low" or "very high" k_{sa} values, as it is with respect to heat transport in and through the Al layer that heat transport in the sample must become negligible.

The sensitivity to laser spot size is a more complicated non-monotonic function of sample thermal conductivity due to the competing effects of pulse accumulation and heat spreading regime. Lower thermal conductivities reduce L_p and push the system more toward planar heating. At the same time, lower thermal conductivities also increase the thermal diffusion time of the system, and hence cause greater pulse accumulation. This means that after absorbing a heating pulse, the sample does not fully return back to equilibrium before the next pulse hits. In this sense, a component of the surface temperature response results from the cumulative effects of consecutive heating pulses. This pulse accumulation effect is what lets the sample feel the effects of the laser modulation frequency (which is a longer time scale than the inter-pulse time), as shown in Figure 3-7. In general, being in a more planar heating regime reduces sensitivity to spot size because edge effects (and hence the finite size of the laser spot) are diminished, as shown in Figure 3-10. From the thermal wave's point of view, the spot size grows to look like an infinite plane heat source regardless of its true size. Conversely, sensitivity to laser spot size is increased for spherical heat spreading, because edge effects dominate and the thermal waves can see and feel the finite size of the laser spot.

For low thermal conductivities (the left third to half of most panels in Figure 3-9), the system is already in a predominantly planar heating regime and so further decreasing k_{sa} primarily influences pulse accumulation and has little additional effect on heating regime. By increasing pulse accumulation effects, the influence of the modulation frequency is enhanced and hence the influence of thermal penetration depths and their associated edge effects is enhanced. Although predominantly in a planar heating regime, because further decreasing k_{sa} increases sensitivity to edge effects via a boosted pulse accumulation effect faster than it decreases the edge effects via a stronger push toward planar heating, the net effect is increased sensitivity to spot size with reduced k_{sa} in the low thermal conductivity regime. For mid-range thermal conductivities, a local minimum in spot size sensitivity occurs where the system begins to transition to a spherical heating regime. In this regime, the boost to spot size sensitivity via the shift toward spherical heating outweighs the decreased sensitivity from thermal penetration depth

related effects via reduced pulse accumulation. The trend shifts to a net increase in spot size sensitivity with increasing k_{sa}. Finally, for very high thermal conductivities (the right quarter of the 2 μm spot size panels) the pulse accumulation effect begins to disappear entirely, and so the sensitivity to spot size decreases again. Even though the system would be in a completely spherical heating regime for thermal penetration depths at the modulation frequency for these ultra-high k_{sa} systems, the system is numb to the effects of the modulation frequency due to lack of pulse accumulation. Sensitivity to spot size is also increased for smaller spot sizes, as this also pushes the system toward spherical heat spreading without reducing pulse accumulation effects. Similarly, increasing modulation frequency pushes the system toward the planar heating regime without affecting pulse accumulation, and so increasing modulation frequency reduces sensitivity to spot size.

Sensitivity to TBC is greatest for later delay times, when the heat is in the process of transferring from the Al layer, through the boundary, and into the substrate. At very early delay times, the heat is still spreading within the Al layer, and has not yet begun to significantly conduct through the boundary, and hence sensitivity to TBC is low. Conversely, sensitivities to the heat capacity and thickness of the Al layer, C_{Al} and L_{Al}, are greatest for early delay times when the dominant heat transport is still taking place through the Al film, before most of the heat has left the Al layer and is diffusing through the sample. The reduction in sensitivity to C_{Al} and L_{Al} at longer delay times is most pronounced for high k_{sa} samples, which suck the heat down and out of the Al film more rapidly. The global maximum in TBC sensitivity as a function of k_{sa} is a matter of thermal impedance matching. For very low or very high k_{sa}, the thermal transport is primarily limited by conduction through the substrate or the Al layer, respectively. In this case, changing the TBC does not a strong effect on thermal transport, as TBC is not the bottle-neck, hence the sensitivity is low. The magnitude of the TBC itself as well as the thermal conductivities of the sample and Al transducer layer will all influence this impedance matching. Similarly, the heating regime will influence what shape factor governs the heat transport through the substrate (planar heating is more resistive than spherical heating), and hence will also affect the thermal impedance matching.

In general, very high frequencies globally reduce sensitivities to physical parameters. This happens because as the system moves more toward the high frequency limit of Eq. (3-10), the phase shift of the temperature response becomes a constant, $\pi/4$, regardless of physical parameters. Because our measured signal is the ratio of the real and imaginary components of the temperature response, which is essentially the phase of the temperature response, we lose sensitivity to all parameters at very high frequencies. In practice, however, Nyquist sampling limitations from the 80 MHz laser repetition rate would cause modulation frequencies > 80 MHz to "wrap around" and behave like < 80 MHz modulation frequencies, analogous to how phonons with spatial frequencies greater than that of the atomic lattice "wrap around" and are folded back into the First Brillouin Zone.

The sensitivities to C_{Al} and L_{Al} collapse into the same curve in the planar heating regime, because the heat conduction becomes one dimensional through the Al layer. This effectively "collapses" the affect of the thermophysical properties of the layer, and is the reason why the heat capacity per unit area of the film is an important quantity in TDTR experiments. This is more pronounced for larger spot sizes and higher modulation frequencies, where the heat conduction in the film is more planar. Unlike for k_{sa}, reducing the sample thermal conductivity makes heat conduction in the Al layer effectively more spherical (and so separates the C_{Al} and

L_{Al} curves). This happens because the heat has more time to spread out radially within the Al layer before it is sucked down into the sample.

Finally, sensitivity to k_{Al} is greatest when the heat transport in the Al film is spherical. As explained above, this happens for small spot sizes, low modulation frequencies, and high sample thermal conductivities. As explained previously, when heat conduction through the Al film is primarily planar and 1-D, the layer is effectively "collapsed" and so there is little sensitivity to k_{Al}, which is only important when the shape of heat conduction within the layer has interesting features. C_{Al} and L_{Al} still have higher sensitivities than k_{Al} in the planar Al heating regime, because unlike k_{Al}, the film areal heat capacity still has relevance in this regime by serving as a the thermal mass from which the absorbed heat is released over time into the sample below.

Chapter 4 - Effects of Alloying and Nanostructuring on Spectral Phonon Mean Free Path Distributions

4.1 Macro-to-Nano Heat Conduction

4.1.1 Connecting the Micro and Macro Worlds

A comprehensive discussion of nano-scale energy transport and conversion is given in Ref. [105]. In section 4.1, I will present a high-level overview to provide the reader with the key concepts necessary to understand the results of the rest of this thesis. In certain places, I will augment existing models and comment on the prevailing physical interpretations of select classes of experiments.

The macroscopic transport of heat through solids via conduction is governed by Fourier's law [106],

$$q'' = -k\nabla T \tag{4-1}$$

where *q"* is the heat flux [W/m^2], k is the thermal conductivity [W/m-K], and ∇T is the spatial temperature gradient [K/m].

Eq. (4-1) dictates that the flow of thermal energy is directly proportional to the temperature gradient, and that the proportionality constant connecting these two is the thermal conductivity, k, which is a material property. The negative sign ensures that heat flows from hot to cold. In general, ∇T and *q"* are vectors and k is a tensor, which can also be a function of temperature and position within the material. Eq. (4-1) can be plugged into the 1st Law of Thermodynamics (energy conservation) [106],

$$\rho c \frac{\partial T}{\partial t} + \nabla \cdot q'' = 0 \tag{4-2}$$

assuming constant properties, in order to yield the heat conduction equation (with zero heat generation) [106],

$$\frac{\partial T}{\partial t} = D\nabla^2 T \tag{4-3}$$

where the ρc product gives the volumetric heat capacity, C [J/m^3-K], and $D = k/\rho c$ is the thermal diffusivity [m^2/s].

In this form, Eq. (4-3) can be recognized as Fick's second law, also known as a diffusion equation, with the thermal diffusivity (D) as the diffusion coefficient with the telltale m^2/sec units. In this light, Eq. (4-1) can also be understood as a form of Fick's law of diffusion. This tells us that fundamentally, macroscopic heat conduction through solids is a form of diffusion.

The units of D have specific physical significance, and as a consequence the time t it takes for heat to conduct (i.e. diffuse) through a given length L scales according to $t \propto L^2$. This offers clues regarding how the macroscopic and microscopic pictures of heat conduction are tied together. During a random walk, a particle on average displaces a net distance of L from its original starting position after traveling an absolute distance (i.e. non-straight path including all the random jitter) of L^2. Modeling the particle as traveling at a constant average speed, the time it takes the particle to displace a net distance L from its starting position therefore scales with L^2. This parallel scaling physics is no coincidence, and implies that macroscopic diffusion processes are the result of microscopic Brownian motion and the random walks of ensembles of individual particles. Indeed, heat conduction in nonmetallic solids is the result of the random motion of individual energy carriers called phonons.

Formally, phonons are quasi-particles, meaning they are not physical particles but are eigenstates of the total dynamical Hamiltonian of a collection of bound atoms that accounts for the kinetic and potential energy of the massive ionic cores of the atoms. Phonons are collective excitations of the atomic lattice, or quantized bundles of vibrational atomic energy. One can approximately visualize a phonon propagating through an atomic lattice as analogous to a pulse traveling down the length of a stretched out slinky (if the vibrational energy were quantized into discrete chunks). Although phonons are fundamentally quantum mechanical objects with wave nature, for length scales at least ~10x larger than their wavelength (wavelengths ~1 – 3 nm for Si at 300 K [76]), phonons can be treated as discrete particles carrying chunks of thermal energy, which obey Bose-Einstein statistics [76], [107], [108]. Doing so lets us apply kinetic theory, which is a classical formalism that treats phonons like hard spheres in order to calculate macroscopic average transport properties for large collections of phonons. A key result of kinetic theory is the magnitude of net thermal energy current carried by phonons in response to a spatial temperature gradient, which in 1-D is given by [105], [109], [110],

$$q''_x = -\left(\frac{1}{3}Cvl\right)\frac{dT}{dx} \tag{4-4}$$

where C (= ρc) is the volumetric heat capacity of the phonon population at constant volume [J/m^3-K], v is the group velocity of the phonons [m/s] (i.e. the speed of sound in the material in the case of low frequency phonons), and l is the mean free path of the phonons [m]. The mean free path (MFP) is the average distance that a phonon travels before scattering and exchanging energy. There are many different mechanisms that scatter phonons, such as collisions with other phonons, impurities or defects in the lattice, boundaries around the perimeter of the sample and between grains, or interfaces between dissimilar materials. Such phonon scattering is the primary resistive process that reduces the thermal energy flux and creates thermal resistance in a material by redirecting the flow of heat-carrying phonons.

Eq. (4-4) connects the macroscopic description of heat transfer (q'', ∇T) with the microscopic description of phonons (C, v, l). The quasi-random jiggling of atoms often used to describe the microscopic physical picture of heat is really the superposition of all the different phonon vibrational waves zipping around in all directions. The random motion of phonons, which themselves are the heat, tends to smooth out any non-uniformities in the distribution of phonons throughout a material (i.e. temperature gradients) and hence lead to the conduction of heat. This smoothing is governed by the diffusion processes dictated by Eq. (4-3). By

comparing Eq. (4-4) to Eq. (4-1) we can identify the microscopic origins of the thermal conductivity, k,

$$k = \frac{1}{3} C v l \tag{4-5}$$

This is the well-known kinetic theory result for the thermal conductivity [105]. In an electrically conducting material, C, v, and l can refer both to the free electrons and the phonons. The total thermal conductivity in this case will be the combination of contributions from electrons and phonons, $k_{total} = k_{electron} + k_{phonon}$, with $k_{electron}$ typically dominating for metals around 300 K. In general, Eq. (4-5) governs how any energy carrier (electrons, phonons, gas molecules, photons, magnons, etc.) contributes to their host material's total thermal conductivity. In the rest of this thesis, we restrict our discussion to nonmetallic crystalline solids for which phonons are the only significant carriers of heat.

By explicitly showing the dependence of k on various microscopic properties, Eq. (4-5) reveals how to engineer the thermal conductivity of a material. Even for structures on the scale of the phonon MFP, C and v are mostly fixed once a material selection has been made, as they are more fundamental properties of the choice of atoms, how they are bonded together, and their harmonic vibrational frequencies. They usually retain bulk values unless confined in structures small enough to approach the phonon wavelengths. The mean free path, l, is an easier target for manipulation. Introducing additional forms of scattering, such as by inserting grain boundaries (e.g. from ball milling and hot pressing) or lattice defects (e.g. from doping, alloying, or ion bombardment), will shorten the phonon mean free paths and consequently reduce the bulk thermal conductivity [51], [52], [76], [77], [99], [100], [111]. Similarly, purifying a material so as to remove grain boundaries, isotopes, and other impurities, will increase mean free paths and increase the innate bulk thermal conductivity. Later in this chapter, we will experimentally investigate the effects of two such approaches to reducing thermal conductivity by examining grain boundary scattering in nano-grained silicon and alloy scattering in silicon germanium.

Alternatively, the physical dimensions of a sample can be reduced in order to reduce l by means of increased scattering events at the boundaries of the sample (e.g. in nanowires [54]–[56], [112], [113] or thin films [114]). These forms of reducing k are often referred to as "classical size effects" [105], because the change in energy transport resulting from a reduction in material external dimensions can be explained using a particle picture of phonons. When the characteristic length scale, L_{char}, of a micro or nanostructure approaches the MFP of the bulk material, l_{bulk}, it will increase phonon scattering, reducing l and hence resulting in a suppressed measured thermal conductivity as compared to bulk. When this occurs, it means that many phonons are scattering primarily at the structure's boundaries rather than within its interior. The restricted dimension can be orthogonal or parallel to the direction of thermal transport, or both.

An important distinction should be emphasized between classical size effects and decreasing k_{bulk} by introducing additional intrinsic bulk scattering mechanisms, as we first discussed. Both will result in a lower measured k, but only in the latter case is the intrinsic k as a material property innately reduced. In the case of many classical size effects, the material's innate k as a material property has not been reduced, and if brought back up to macroscopic dimensions, samples would recover k_{bulk}. The apparent reduction in k is a direct consequence of the size of the sample and not changes to its material properties. This distinction may seem like pure semantics at first, but consider the following two situations. First consider the case of heat

transfer perpendicular to a thin dielectric film [115]. L_{char} for this system is the film thickness. It initially seems intuitive and expected that reducing the film thickness until $L_{char} \leq l$ would increase phonon scattering at the film boundaries, reducing the phonon MFPs and hence reducing k. And, indeed, experiments do measure a reduced k [96], [116]. But now consider the case of heat transport from a small heated isothermal sphere embedded in an infinite medium [117]. L_{char} for this system is the sphere's radius. Reducing L_{char} as before, until $L_{char} \leq l$ (where l is the phonon MFP in the host material, not in the sphere) will again result in measuring a reduced k of the host material. However, this cannot be explained in the same way as with the thin film. The presence of the sphere does not reduce l in the material. There is no special way in which L_{char} suddenly and significantly increases phonon scattering when it shrinks to approach l. If anything, a larger sphere would scatter more phonons due to its larger surface area. All phonons in the host material still propagate with the same average l_{bulk} as before, including those phonons in the immediate vicinity of and emitted by the sphere. So, why then is k reduced, if not due to truncated phonon MFPs? And, has the true k actually been reduced (as happens when we increase materials' bulk scattering rates), or are we just measuring a reduced apparent k? If the latter, does a reduced apparent k have any physical consequence on the actual heat transfer from tiny spheres, or is it just a measurement artifact that we need to correct? We will return to these questions once we have secured a couple more important concepts under our belt. The answers to these questions hold immediate importance regarding all phonon MFP spectroscopy techniques to date (including the variable spot size TDTR measurements performed for this thesis). These techniques fundamentally work to measure MFP properties by using tiny heaters on the order of l, akin to a small, embedded, heated sphere.

4.1.2 Diffusive Versus Ballistic Transport, and *k* Suppression

When external dimensions approach the phonon mean free path (MFP) and phonons scatter predominantly at the boundaries of the material, we say we have transitioned from a "diffusive" thermal transport regime to a "ballistic" thermal transport regime. This progressive transition is depicted in Figure 4-1, and will be discussed in some detail below.

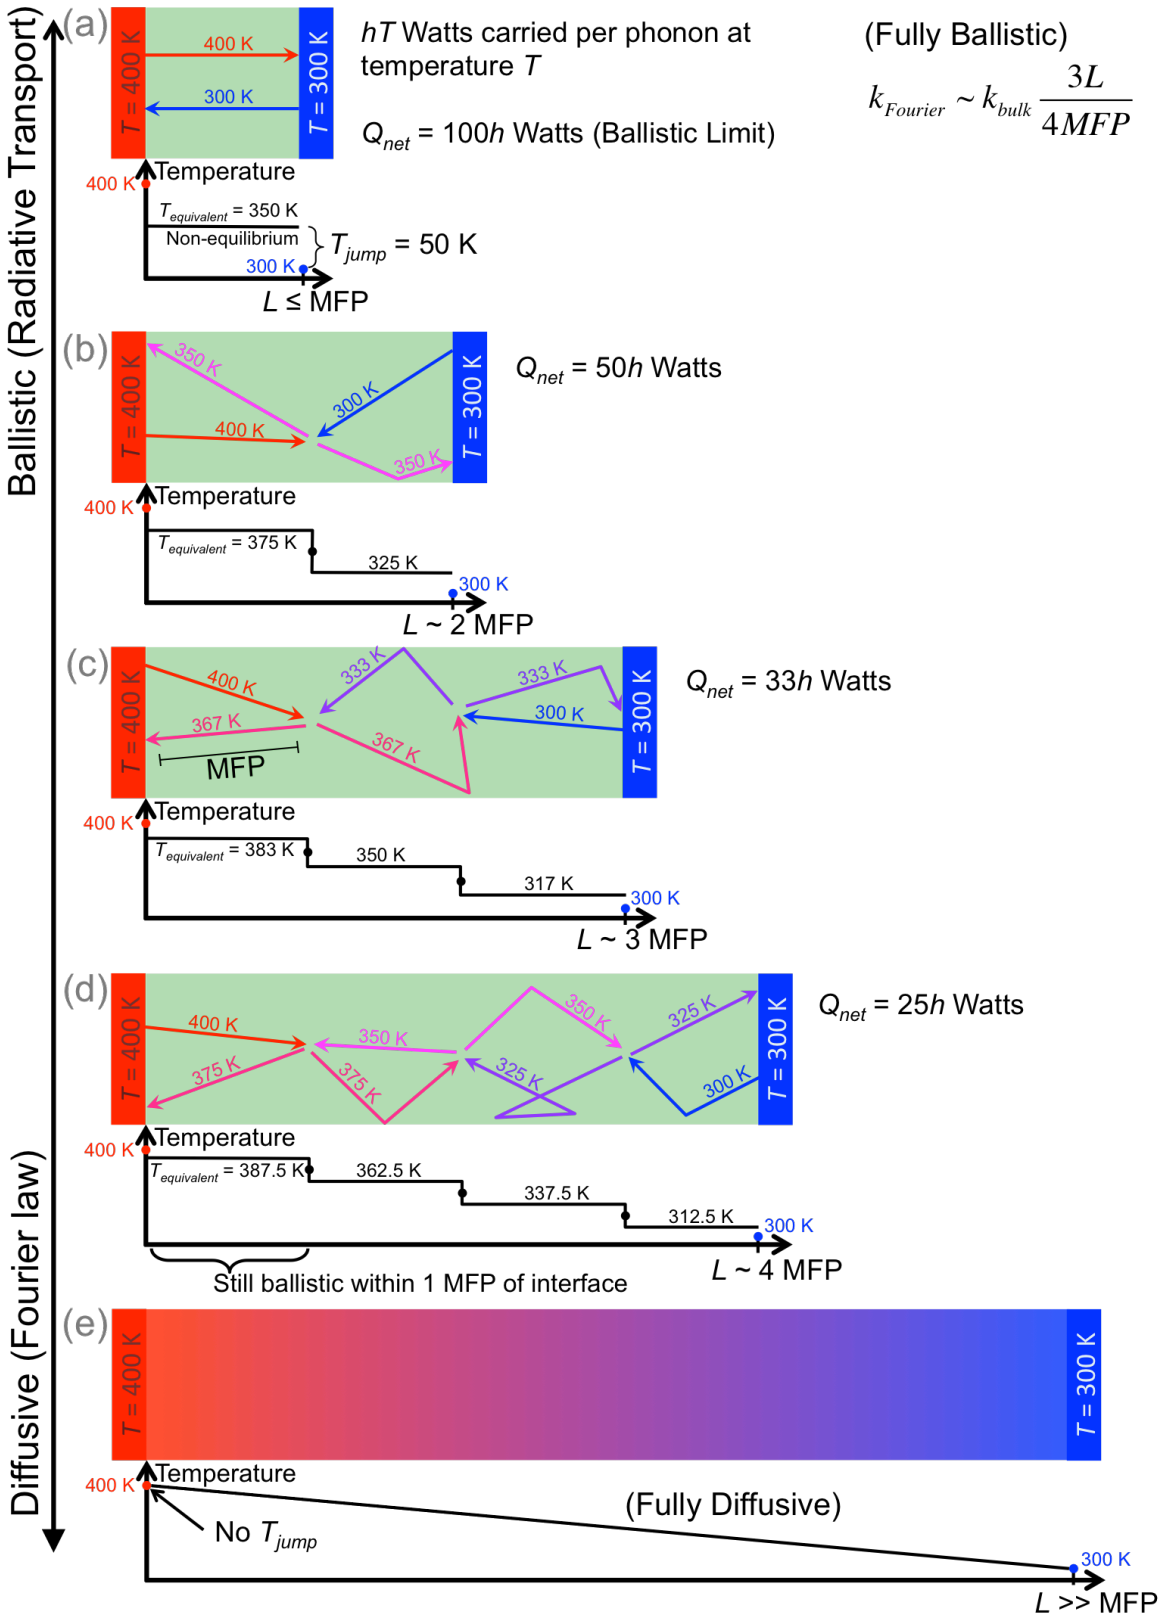

Figure 4-1 Ballistic to Diffusive 1-D Phonon Heat Transfer, Concept Schematic. When transitioning from ballistic to diffusive transport, temperature jumps at interfaces disappear, a continuous temperature gradient

is established, and the net heat flux scales with the length of the bar. Fully ballistic heat flux is independent of *L*. A more complete explanation is provided in the text.

Figure 4-1 depicts a 1-D bar with fixed-temperature hot and cold thermal reservoirs at either end undergoing heat conduction as its length is varied from $L \leq l$ to $L >> l$. Arrows indicate the trajectories of phonons within the material, with colors and labels giving the temperature of that phonon (individual phonons do not have a well defined temperature, but think of each arrow as a representative phonon from an ensemble of similar phonons). Perfect inelastic scattering among phonons is assumed for clarity, as are perfectly absorbing and emitting interfaces at the thermal reservoirs. All scattering events conserve energy. Below each heated bar schematic is its temperature profile, showing fixed hot and cold reservoir temperatures and equivalent temperatures within the bar (discussed below). To the right of each schematic is the *net* heat flux, Q_{net}, being carried through the bar by the phonons, in units of h, the "irradiance per phonon" [W/K per phonon]. hT is the irradiated power [Watts per phonon] carried by a representative phonon at temperature T. Therefore, the net heat flux leaving the hot surface equals the heat out minus the heat in, $Q_{net} = Q_{out} - Q_{in}$. For example, the hot reservoir in (d) has a representative phonon at 400 K (and hence carrying $400h$ Watts of power) leaving, and a representative phonon at 375 K (carrying $375h$ Watts) entering, so the net thermal flux leaving the reservoir is $400h - 375h = 25h$ Watts. This can also be calculated by examining the cold reservoir interface. The thermal conductivity is also given to the right of each schematic, as defined by Eq. (4-1). The effective k for ballistic transport in panel (a) will be derived later. The schematics show the complete progression from fully ballistic phonon transport when $L \leq l$ in (a), to fully diffusive phonon transport when $L >> l$ in (e). In practice, of course, l is a *mean* free path, and the actual free paths of the phonons will follow a Poisson distribution and will be spread out evenly throughout the bar. In Figure 4-1, phonon-phonon collisions are depicted as happening at regular isolated intervals for clarity, rather than smeared out throughout the bar as actually happens in real materials.

Diffusive transport, as named because energy is transported via the diffusion of energy carriers as previously discussed, is what governs macroscopic every-day heat conduction. Fourier's law and the heat conduction equation, Eq. (4-1) and Eq. (4-3), are based on the assumption of diffusive transport. Diffusive transport is a valid model when phonons scatter within a material's bulk frequently enough to follow random walks. Sufficiently frequent inelastic scattering among phonons must also occur such that a local thermal equilibrium can be established at different points in space, and such that the concept of a temperature gradient is well defined. The bottom panel, Figure 4-1(e), depicts heat transfer via fully diffusive phonon transport.

When external dimensions are comparable to or smaller than bulk mean free paths, the phonons primarily travel ballistically from one boundary to the other, without scattering much in transit. In this regime, the phonons do not follow random walks but rather straight-line trajectories, and so a diffusion model is no longer valid and Fourier's law breaks down, as will be further discussed later. This regime is sometimes also referred to as "sub-continuum" heat transfer because the fact that heat is carried by individual phonons with individual trajectories—the details of which are usually washed out in the diffusion process—becomes relevant. The top panel, Figure 4-1(a), depicts heat transfer via fully ballistic phonon transport.

Ballistic phonons will travel directly from their emitting thermal reservoir, without scattering, until being absorbed by the receiving thermal reservoir on the opposite side. Because

of these ballistic trajectories, the total Q_{net} thermal flux carried from hot to cold will be completely independent of the length of the bar, provided the transport remains ballistic. This constant Q_{net}, termed the "ballistic limit," is the maximum physically possible rate of net heat transfer between the two reservoirs (as mediated by phonons), and is limited only by the rate at which phonons can be emitted from the reservoirs. Ballistic transport is sometimes also called "radiative" transport in analogy to heat transfer by thermal radiation, because photons typically travel balistically and do not interact with each other. In ballistic transport, non-interacting hot and cold phonons exist in the same region of space, representing two distinct thermal distributions of phonons simultaneously. Because these phonons do not scatter with one another, local thermal equilibrium cannot be established at interior points within the bar, and the bar does not have a classically defined temperature anywhere within its interior. We can, however, define an equivalent average temperature, $T_{equivalent}$, based on the temperature that all the phonons instantaneously occupying a local region of space would reach if they were hypothetically allowed to interact and adiabatically come to a single thermal equilibrium. $T_{equivalent}$ is really a measure of the local energy density, and in the case of a temperature-independent heat capacity (assumed here for simplicity) it is the simple average of the rightward and leftward propagating phonon temperatures [105]. Because the phonons never interact for fully ballistic transport, $T_{equivalent}$ is uniformly 350 K everywhere within the bar in Figure 4-1(a).

This non-equilibrium among ballistic energy carriers also gives rise to a temperature "jump" at each interface, as can be seen in Figure 4-1(a). The thermal reservoirs are at 400 K and 300 K, while $T_{equivalent}$ = 350 K everywhere within the bar including adjacent to the thermal reservoir interface. This results in a discontinuous 50 K jump at the interface, and is sometimes also referred to as temperature "slip." This jump in T at the interface is a routine phenomenon in the study of radiation heat transfer [118], and is a direct consequence of the non-equilibrium of non-interacting ballistic energy carriers. All phonons emitted from the hot reservoir are at 400 K, while all cold phonons ballistically arriving at the hot reservoir from the cold reservoir are at 300 K. Therefore, the hot reservoir is at 400 K, but $T_{equivalent}$ immediately within the material, adjacent to the reservoir, instantly drops to 350 K because of the averaging with the incoming cold phonons.

As we progress through the other panels, Figure 4-1(b)-(e), the length of the bar (L) is increased to be greater than the phonon MFP, l, while holding the total temperature difference, ΔT, constant. This results in increased phonon scattering events. Each phonon-phonon collision is depicted as perfectly inelastic for clarity, and it exactly averages the energy of the two incoming phonons between the two outgoing phonons. In practice, phonon-phonon collisions can span the full spectrum from elastic to inelastic, need not conserve phonon number (such as in Umklapp processes), and need only conserve crystal momentum and not physical momentum [119]. Figure 4-1(b)-(e) show why the 1-D Fourier law conduction thermal flux decreases as L increases. Mathematically, this direct relationship between q'' and L (for fixed k and ΔT) is captured in Eq. (4-1) (remembering that $\nabla T = \Delta T / L$ for our 1-D bar). But now we can understand its physical origins. As L is increased, the number of phonon-phonon scattering events increases proportionally, which reduces the flow of heat. On average, these scattering events establish a local thermal equilibrium and average out the temperature of the phonons arriving from either end of the rod. By examining this energy averaging effect in detail in Figure 4-1(b)-(d), we see that as L increases from $2l$ to $3l$ to $4l$, Q_{net} decreases to ½, ⅓, and ¼ of the ballistic limit thermal flux proportionally. Here, we did not need to invoke Eq. (4-1) to calculate this decrease in Q_{net} proportional to increasing L, but correctly deduced it based on purely

physical arguments and careful book-keeping of the energy averaging from phonon-phonon scattering events. Also note that each time L increases by one MFP, the number of scattering events per representative phonon increases by 1 proportionally.

With increased energy averaging from phonon-phonon scattering, the temperature profile also approaches a smooth gradient, and the temperature jump at the interfaces progressively decreases. In the extreme limit when $L >> l$ and there are a tremendous number of phonon-phonon scattering events occurring between hot and cold reservoirs (Figure 4-1(e)), the temperature profile is a completely smooth gradient with no jumps, as predicted by Eq. (4-1). Again, note that these features all arise naturally from simple considerations of phonon-mediated energy transfer.

Because Fourier's law breaks down in the non-diffusive regimes, a new mathematical approach is required to model heat transfer for temperature gradients and structure sizes comparable to l. Since photons are well-studied energy carriers that typically travel ballistically, one approach is to borrow much of the mathematical formalism from thermal radiation heat transfer and apply it to ballistic phonons [120], [121]. For the general case, the Boltzmann Transport Equation (BTE) is used for a rigorous mathematical model of phonon heat transport that is valid in all regimes from fully ballistic to fully diffusive [105], [115], [121]. Unfortunately, the full BTE is a high-dimensional integro-differential equation of time, real space, and phase space, and closed fully analytical solutions are scarce. Here, we will examine simpler toy models with readily available analytical solutions that still provide all of the physical intuition we will need.

The fully ballistic limit can be modeled using the exact formalism of radiation heat transfer. Therefore, we begin examining ballistic phonon transport by solving the thermal radiation problem analogous to our conducting bar. Specifically, two diffuse, gray, parallel plates held at fixed temperatures. In general, the radiative heat flux for this system is given by [106]

$$q_{b,p} = \frac{E_{b_1} - E_{b_2}}{\frac{1-\varepsilon_1}{A_1\varepsilon_1} + \frac{1}{A_1 F_{12}} + \frac{1-\varepsilon_2}{A_2\varepsilon_2}} \tag{4-6}$$

where $q_{b,p}$ is the total radiative heat flux in [W] for ballistic phonons carrying heat between parallel plates, E_b is the phonon radiative emissive power for each plate [W/m^2], A the surface area of each plate [m^2], ε the phonon emissivity for each plate, and F_{12} is the view factor from plate 1 to plate 2.

We can make a number of simplifying approximations to Eq. (4-6). To model the 1-D linear heat transfer of Figure 4-1, we choose infinite parallel plates. This means $A_1 = A_2 = A$, which can be divided through to give $q/A = q''$, and also means $F_{12} = 1$ (infinite plates). We will also neglect thermal contact resistances at the interfaces, and so model the plates as black, dictating that $\varepsilon_1 = \varepsilon_2 = 1$. Finally, by approximating the phonon volumetric heat capacity, C, as being independent of temperature, we can approximate E_b as [117], [122]

$$E_b \approx \frac{Cv}{4} T \tag{4-7}$$

where C is the phonon volumetric heat capacity at constant volume, v is the phonon group velocity, and T is the temperature of the plate emitting the phonons. Combining all of these approximations and substituting in Eq. (4-7), Eq. (4-6) becomes

$$q''_{b,p} = \frac{1}{4} C v \Delta T \tag{4-8}$$

where ΔT is the temperature difference between the two plates. Note that L never entered the problem, because the ballistic heat flux is a constant and is independent of length.

Eq. (4-8) gives the heat transfer physics for heat conduction by fully ballistic phonons. Physically, this ballistic limit is an upper bound to the possible heat flux and represents the net rate at which phonons can be emitted and radiate away from the surface. To study the opposite limiting case, we solve the diffusive heat conduction equations. Applying Eq. (4-4) to the case of a 1-D bar of length L, we acquire the diffusive heat flux,

$$q''_{f,p} = \frac{1}{3} C v \Delta T \left(\frac{l}{L_{char}} \right) \tag{4-9}$$

where $q''_{f,p}$ is now the total Fourier's law heat flux in [W/m^2] for diffusive phonons carrying heat between parallel plates, ΔT is redefined as $\Delta T = T_{hot} - T_{cold}$ to eliminate the negative sign, l is the phonon MFP as per usual, and L_{char} is the length of the 1-D bar in this case. The dimensionless quantity l/L_{char} is the Knudsen number (Kn). For large values of Kn > 1, sub-continuum effects become important and the validity of Fourier's law is questionable. We set Eq. (4-9) equal to Eq. (4-8) to examine the cross-over point from diffusive to ballistic transport and find that it occurs when $L_{char} = \frac{4}{3} l$. Next, we plot Eq. (4-9) and Eq. (4-8) together in non-dimensional form in Figure 4-2 as functions of L_{char}, representing the fully diffusive and fully ballistic phonon transport regimes, respectively.

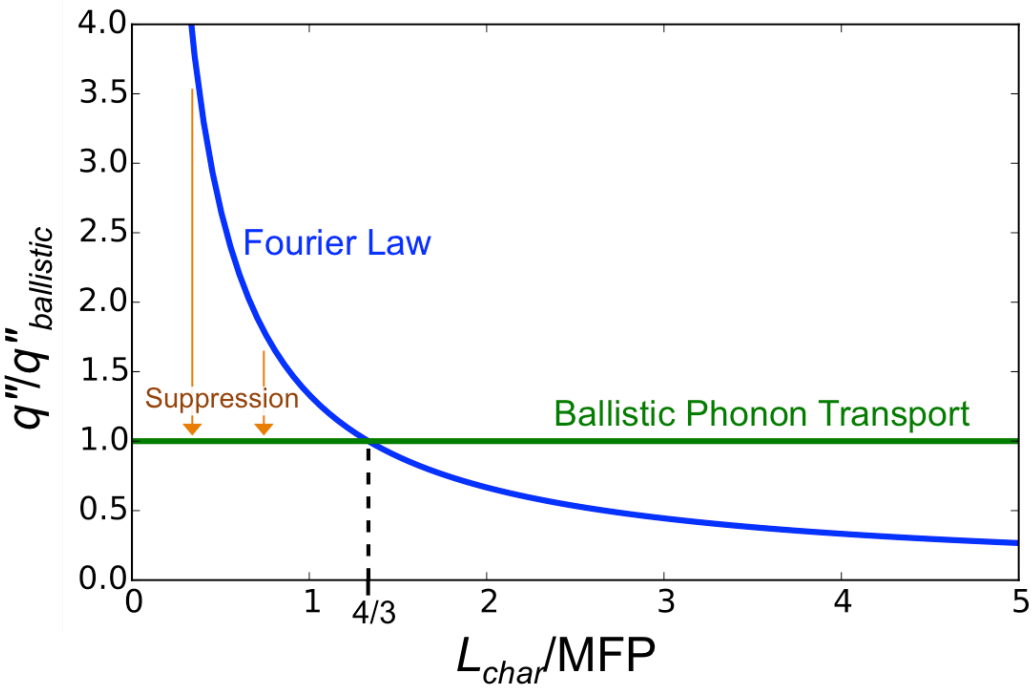

Figure 4-2 Fully diffusive versus fully ballistic phonon transport heat fluxes through a 1-D bar of length L_{char} or an embedded sphere of radius L_{char}, with fixed ΔT and fixed phonon MFP.

For $L_{char} > \frac{4}{3}l$, the heat transport is predominantly diffusive, and Fourier's law gives the limiting heat flux. For $L_{char} < \frac{4}{3}l$, the heat transport is predominantly ballistic, and the radiation model gives the limiting heat flux. The observation that the ballistic heat flux is a constant while the diffusive heat flux continues to vary as $1/L_{char}$ for all L_{char} is quite pertinent. Therein lies the key to answering the questions posed earlier with regards to reducing k by shrinking the bar length. If a system is in the ballistic regime (approximately Kn $\geq$ 3/4) and an experiment is performed to determine k, then a Fourier's law model is being used to fit to the data (since the concept of thermal conductivity only exists for Fourier's law, and is not properly defined for ballistic transport due to the lack of temperature gradients and local thermal equilibrium). However, Fourier's law breaks down in the non-diffusive regimes, and it over-predicts the true heat flux for a given ΔT as is evident in Figure 4-2. The true heat flux in the ballistic regime—what the real experiment will actually measure—is given by Eq. (4-8), and can be significantly less than what Fourier's law predicts if L_{char} is much smaller than l.

Analysis of the experimental data entails varying the value of k in the thermal model until the best possible agreement between the model and the data is achieved. Because the experimentally measured heat flux is less than the q'' that Fourier's law would predict using k_{bulk}, and because artificially decreasing k in the model decreases the predicted heat flux for fixed geometry and fixed ΔT (see Eq. (4-1)), the end result will be an apparent suppressed thermal conductivity that is less than k_{bulk}. This apparent suppressed thermal conductivity, $k_{apparent}$, will provide the best agreement between the ill-suited model and the data, and so the researcher would claim to have observed a suppression of the sample's thermal conductivity as a result of shrinking its external dimension in the direction of heat transport down to be on the order of the phonon MFP. Alternatively, in an experiment where the heat flux is controlled and held constant, Fourier's law would under-predict the resulting temperature difference, ΔT. The

measured true ΔT for fixed q'' would be higher than what Fourier's law would predict using k_{bulk}, and so again the experiment would measure a suppressed $k_{apparent} < k_{bulk}$ in order to make model predictions match experiments.

Because of the linear relationship between k and q'' for diffusive transport, the amount of suppression of $k_{apparent}$ will equal the amount of suppression necessary to get $q''_{Fourier}$ to equal $q''_{ballistic}$, as shown in Figure 4-2. We can calculate the scaling of this suppression by comparing Eq. (4-9) to Eq. (4-8), and remembering the definition of k_{bulk} given by Eq. (4-5),

$$\frac{k_{apparent}}{k_{bulk}} = \frac{3}{4}\left(\frac{L_{char}}{l}\right) = \frac{3}{4Kn} \tag{4-10}$$

This is the same result as derived from the full BTE solution for an equivalent problem in the same regime [117], and is given for the ballistic case in Figure 4-1(a). Physically, Eq. (4-10) results from forcibly casting the ballistic heat flux in the form of Fourier's law, to see what $k_{apparent}$ a non-diffusive heat transfer experiment would interpret if analyzing data using Fourier's law.

Eq. (4-10) is only valid around $L_{char} \leq \frac{4}{3}l$, and it shows that the suppression of $k_{apparent}$ will scale directly with L_{char} (for fixed l). An expression that is valid for all values of L_{char} can be derived using the full BTE [117], and will converge to Eq. (4-10) in the ballistic limit, but it is far less trivial to derive and such closed analytical solutions do not exist for most geometries. One can similarly work out the limiting ballistic and diffusive expressions for the net heat flux [W/m^2] leaving the surface of a black sphere at T_1 with radius L_{char} embedded in an infinite medium with a far-field temperature of T_2. Fortuitously, the expressions for both the diffusive and ballistic heat fluxes work out to be identical to those expressions for the 1-D bar of length L_{char} that we just derived, as given by Eq. (4-9) to Eq. (4-8). Therefore, the full preceding discussion also applies to the case of the heated embedded nano-sphere. It is also important to note that for the Fourier's law solution for an embedded heated sphere, its temperature profile falls off as $1/r$, where r is the radial distance away from the sphere. Therefore, 50% of the total temperature drop from T at the surface of the sphere to T infinitely far away in the medium occurs within one radius of the spherical surface, and 75% of the total ΔT occurs within 3 radii of the spherical surface. This rapid temperature decay is a unique property of the geometry of spherical heat spreading and is the reason why the thermal resistance between two concentric spheres remains finite even when the radius of the outer sphere goes to infinity [106], in contrast to analogous planar and cylindrical systems that diverge. This means that the Fourier's law model of a heated embedded sphere is primarily sensitive to the nature of the heat transfer occurring within just a couple radii of the surface of the sphere.

We now have all the tools necessary to properly explain why embedding a heated sphere of radius $L_{char} \leq \frac{4}{3}l$ in an infinite medium will result in a measured $k_{apparent} < k_{bulk}$. While it is tempting to apply similar logic from other nanostructuring or classical size effects and say that because the sphere's radius is on the same order as or smaller than the host medium's phonon MFPs, the MFPs are being clipped and the reduction in l is causing a reduction in k. However, this is incorrect. Many people fall into this trap and false line of reasoning. As previously explained, the sphere has no effect on l, which remains at its bulk value in the host medium. The thermal conductivity of the host medium is still k_{bulk}, and has not been affected by the presence of the sphere. What actually happens is the diffusive Fourier's law model used to examine the

heat transfer process is predominantly sensitive to the heat transfer occurring within a couple radii of the surface of the sphere. When the radius of the sphere is comparable to the host medium's MFP, then the heat transfer within a couple radii of the sphere is mostly ballistic, which results in a lower heat flux for fixed ΔT (or a higher ΔT for fixed heat flux) than what is predicted by Fourier's law. As a result, a diffusive model forces its assumed value of k to decrease in an effort to match the experimentally observed values with predictions. Therefore, in reality the thermal conductivity of the host medium has not actually been reduced, but the physical heat flux is lower for fixed ΔT than what Fourier's law predicts using k_{bulk}, because the dominant heat transfer is non-diffusive (or vice-versa for a fixed heat flux). Note that for strongly ballistic transport it may be entirely impossible to satisfactorily fit a Fourier's law model to the data, for any assumed value of $k_{apparent}$, due to qualitative differences in the fundamental nature of the heat transfer process (especially for complicated geometries or boundary conditions). For example, the shape of the cooling curve in a transient grating experiment stops being exponential in the strongly ballistic regime, making it impossible to model using Fourier's law regardless of the choice of $k_{apparent}$ [123].

Because of the frequent confusion these subtle concepts provide, it is worth reiterating the distinction between actually reducing k_{bulk} by introducing additional phonon scattering mechanisms or by reducing the external dimensions of diffusely scattering boundaries orthogonal to the direction of heat transfer, versus measuring a reduced apparent thermal conductivity by using a small heater comparable to or smaller than a material's phonon MFP. Both result in measuring a suppressed k, and both can be due to experimental characteristic length scales approaching l. However, only the former actually reduces k_{bulk}. The latter only measures an apparent $k_{apparent} < k_{bulk}$ as a result of trying to force an ill-suited diffusive thermal model to fit data from experiments where the relevant heat transfer is predominantly non-diffusive. To stress this point, note that if hypothetically Fourier's law happened to under-predict q'' in non-diffusive regimes instead of over-predicting it, then such experiments would yield an enhanced $k_{apparent} > k_{bulk}$, instead of a suppressed $k_{apparent}$ as is actually observed. But nanostructuring a material would still always result in a suppressed true thermal conductivity. If one were to be pedantic when examining the thermal transport through a thin film whose thickness is comparable to l, its cross-plane thermal conductivity is not reduced, it is just not properly defined due to the non-diffusive nature of the transport. We can accurately describe the thermal transport by using an effective $k_{apparent} < k_{bulk}$, but strictly speaking the concept of thermal conductivity is only rigorously valid for fully diffusive systems, and the observed reduction in $k_{apparent}$ is really the result of forcing a diffusive model to describe quasi-ballistic transport. Similarly, even the concept of temperature is not rigorously valid for the non-equilibrium ballistic transport across a thin film, and we really can only talk about an effective $T_{equivalent}$ that describes local energy density. In contrast, if we nanostructure a bulk material to contain grains with $L_{char} \leq l$ that scatter phonons at their boundaries, then we really have reduced k_{bulk}. The bulk-scale sample is still perfectly well governed by diffusive transport, and so the concept of a thermal conductivity is rigorously valid, and its value will be reduced. This is also true of the axial thermal transport down rough-walled nanowires, whose boundaries will diffusely scatter phonons as a result of a restricted diameter. This will genuinely decrease the axial phonon MFP, and hence reduce k for long nanowires whose axial thermal transport is still predominantly diffusive.

Just because $k_{apparent}$ is not a true measure of a material's actual thermal conductivity, does not mean that such experiments are without utility. Because the degree of suppression of

$k_{apparent}$ is a function of the relative magnitude of L_{char} compared to l, such experiments are excellent approaches to experimentally measuring phonon MFPs of bulk materials, *without* needing to nanostructure those materials or otherwise alter the very phonon mean free path spectrum that you are trying to measure. This is a powerful concept that we will return to later, and is the underlying principle behind variable spot size TDTR as a phonon MFP spectroscopy tool.

4.1.3 Understanding the Breakdown of Fourier's Law

At this point, it is worth briefly discussing why Fourier's law over-predicts heat flux in non-diffusive transport regimes. The discussions of section 4.1.2 explain why Fourier's law can sometimes be in error in the first place, but many find it confusing that Fourier's law should predict a greater heat flux than ballistic transport. After all, for any given transport problem, the maximum energy flux always occurs when the energy carriers are traveling ballistically. When compared to a ballistic model, a diffusive model assumes the existence of imaginary scattering events that are not actually there, and scattering events add resistance to transport. So, why does Fourier's law, which assumes the presence of additional non-existent resistive scattering events, predict a higher heat flux than the ballistic limit—the maximum possible heat flux?

One common stated explanation is that Fourier's law, by means of additional scattering mechanisms, assumes the existence of local thermal equilibrium and hence a temperature gradient that does not actually exist. In reality, when the hot and cold reservoirs are brought too close together as in Figure 4-1(a), the concept of local thermal equilibrium and hence temperature gradients disappears. However, if Fourier's law still assumes the existence of a continuous temperature gradient, then it would become increasingly steep as the distance separating the fixed ΔT became increasingly small. Eq. (4-1) shows that the diffusive heat flux scales directly with the temperature gradient, so as the temperature gradient diverges to infinity with a vanishing separation distance between temperature boundary conditions, so too does the predicted diffusive heat flux diverge to infinity. This is a completely fair and mathematically consistent explanation, but it is a macro-scale explanation and it does not provide a physical understanding on the deepest level for why, from a micro-scale heat conduction perspective, the axioms of diffusive transport would predict an enhanced heat flux. Fundamentally, why would assuming local thermal equilibrium and temperature gradients from imagined scattering events require a hot surface to emit more phonons than it actually can, or cause those emitted phonons to travel with less resistance than if they traveled ballistically?

A compelling micro-scale understanding is possible by revisiting Figure 4-1 in its entirety. For each panel, the total net heat flux (Q_{net}) leaving the hot reservoir is the result of all energy emitted as hot phonons minus all energy absorbed as colder phonons. In panel (d), Q_{net} = 25h, because the reservoir emits phonons at 400 K while absorbing phonons at 375 K, for a net emission of 400h - 375h = 25h Watts. Following the trends up the panels as the transport becomes more ballistic, Q_{net} increases even though the hot thermal reservoir continues to emit phonons at the same temperature (400 K) and at the same rate. The increase in Q_{net} is therefore a result of the decreasing temperature of the cold phonons absorbed by the hot reservoir. Less and less of the hot reservoir's emitted energy is canceled out by absorbed energy as scattering events are removed. The absorbed phonons get progressively colder because there are fewer inelastic scattering events between the thermal reservoirs. Each inelastic phonon scattering event averages out the temperature, bringing it to a local thermal equilibrium, and hence reducing the

local temperature difference between hot and cold phonons. It is ultimately this temperature difference between hot and cold phonons at the interface that dictates the magnitude of Q_{net}. One could therefore describe the increased diffusive heat flux with decreasing L_{char} as resulting from the assumed fewer and fewer scattering events that occur as L_{char} is reduced. It then follows that this principle will break down exactly as the transport becomes ballistic. At this point, there are no more scattering events left to get rid of in order to further increase the local temperature difference between hot and cold phonons at the interface. However, Fourier's law is agnostic to the granularity of phonons and the finite number of scattering events that exist. As a result, even as L_{char} continues to shrink past l and a ballistic transport regime is entered, Fourier's law continues to assume that a shrinking L_{char} effectively means a reduced number of scattering events. Because there are no more actual scattering events to get rid of, in the ballistic regime Fourier's law effectively assumes the presence of *negative* scattering events that artificially increase the temperature difference between hot and cold phonons at the interface beyond what is physically possible. This translates to effectively requiring the absorbed cold phonons to continue to drop in temperature, below that of the cold reservoir and even 0 K. These fictionally cold phonons, born out of fictional negative scattering events, begin to suck heat out of the hot reservoir when absorbed, leading to a diverging Fourier's law heat flux.

4.1.4 Non-Gray Mean Free Paths

Eq. (4-5) makes the simplifying assumption that C, v, and l are the same for all phonons in the material. This is referred to as the "gray medium approximation" [115] because it ignores the fact that there are phonons of different vibrational frequencies (analogous to photons of different colors) that can each have different values of frequency-dependent C, v, and l. We are primarily concerned with the breadth of values of l (which itself is a function of ω). All discussions up until now have assumed a gray phonon model regarding mean free paths. This is a reasonable approximation for ideal gases [124] and free electron gases that span only a narrow band of energies centered around the Fermi energy. However, it is a poor approximation for phonons in real crystalline materials, where the spectrum of MFPs can span more than two orders of magnitude [76], [125]–[127] (see Figure 4-3). If this approximation is relaxed, then Eq. (4-5) can be applied to each type of phonon individually, each of which contributes to the material's total k, and may have its own unique value of C, v, or l. Traditionally, different types of phonons are indexed by their vibrational frequency, ω, and polarization. All other properties such as C, v, l, and the phonon's energy, can be calculated as a function of the phonon frequency. In practice, however, it is far more useful from an engineering perspective for relating k to nanostructure effects to index different kinds of phonons by their mean free path, l, instead [125]. In the case of materials with isotropic phonon dispersion relations and mean free paths, l uniquely identifies a given phonon for a particular polarization branch. Typically, l is a smooth and monotonically decreasing function of ω. The material's total thermal conductivity can be expressed as an integral over the contributions from phonons of all mean free paths. Integrating up the contribution of Eq. (4-5) for each kind of phonon, we get [76], [125],

$$k_{bulk} = -\int_0^{\infty} \sum_s \frac{1}{3} C v l_{bulk} \left(\frac{dl_{bulk}}{d\omega} \right)^{-1} dl_{bulk} \tag{4-11}$$

where "*bulk*" subscripts denote properties as they would exist in a bulk, macroscopic, sample of the material, s indexes different phonon polarizations, and a formal change of variable from phonon frequency (ω) to phonon mean free path (l) has been applied. C, v, and l_{bulk} are all functions of ω, so that l_{bulk} now represents a whole spectrum of mean free paths that exist in the bulk material, rather than just a single value as in the discussions of previous sections. The negative sign ensures that k remains a positive quantity when swapping the bounds of the integral during the change of integration variable so that the integral begins by adding up the smallest MFPs first (small values of ω correspond to large values of l, and vice-versa).

Eq. (4-11) shows explicitly that a material's total bulk thermal conductivity is the result of adding up the contributions from all phonons in the material, which form a continuous spectrum. The integrand of Eq. (4-11) gives the spectral thermal conductivity as a function of l (i.e. the thermal conductivity per l), which is important for understanding later results and is qualitatively depicted in Figure 4-7 in section 4.5, Summary of Variable Spot Size TDTR Findings. This spectral k as a function of l is referred to as the "phonon MFP spectrum." It is also possible to terminate the integration in Eq. (4-11) before $l = \infty$ to calculate the k due to a fraction of the total MFP spectrum. If one instead stops the integration at some specific $l = l_\alpha$, then the result is the thermal conductivity due to the contributions from all phonons with mean free paths $\leq l_\alpha$. When normalized to k_{bulk}, this is termed the thermal conductivity accumulation function [125], α, and it ranges from 0 to 1,

$$\alpha(l_\alpha) = \frac{-1}{k_{bulk}} \int_0^{l_\alpha} \sum_s \frac{1}{3} C v l_{bulk} \left(\frac{dl_{bulk}}{d\omega} \right)^{-1} dl_{bulk} \tag{4-12}$$

Calculated k accumulation functions from Refs. [75], [76] for Si and $Si_{90}Ge_{10}$ are plotted in Figure 4-3. The 60% MFP, $l_{0.6}$, is defined such that $\alpha(l_{0.6}) = 0.6$, and so on. This means that phonons with mean free paths $\leq l_{0.6}$ collectively account for 60% of the total heat carried in the bulk material. The paradigm introduced by Eq. (4-12) can be very instructive when calculating heat transfer in micro and nano structures. If, for example, additional scattering mechanisms are introduced that effectively eliminate all phonons with $l > l_{0.6}$, so that only phonons with $l \leq l_{0.6}$ actively participated in transporting heat, then 40% of the thermal conductivity will have been removed and the material will now have a k only 60% of its original value. Alternatively, we say that we have suppressed the thermal conductivity of the material by 40%.

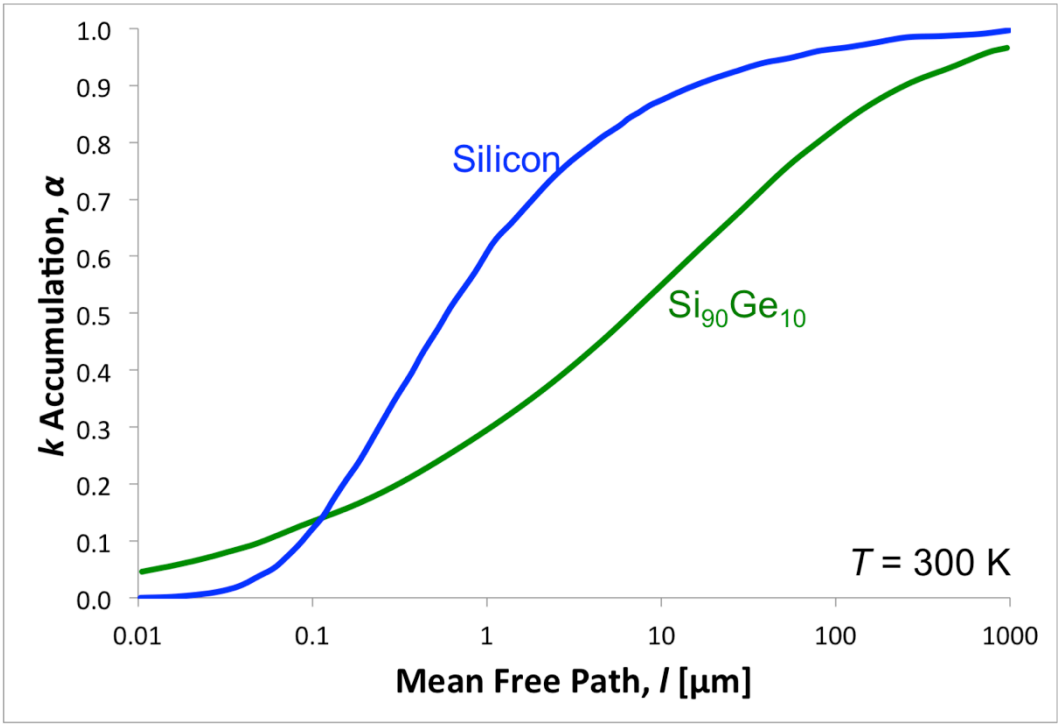

Figure 4-3 Thermal conductivity accumulation functions (Eq. (4-12)) of Si and $Si_{90}Ge_{10}$. Plots are theoretical calculations from Refs. [75], [76]**.**

In practice, however, nanostructures that target $l > l_{0.6}$ MFPs will not simply suppress k by 40%, because the real effect of such structures is not to literally eliminate phonons with $l > l_{0.6}$ from the material, but to reduce the MFPs of those phonons to be $\leq l_{0.6}$. Therefore in reality, all types of phonons (i.e. phonons of all vibrational frequencies) will still participate in heat transfer (even if negligibly). However, most of the phonons that used to have $l > l_{0.6}$ will now have $l \leq l_{0.6}$ as a result of the introduced scattering mechanisms and so will contribute less to the total k via Eq. (4-5). There are now a greater number of types of phonons with $l \leq l_{0.6}$ carrying heat. The overall k is always reduced, but whether or not this equates to a greater or lesser absolute quantity of heat now carried by phonons with $l \leq l_{0.6}$ depends on how much the MFPs of the phonons that previously had $l \leq l_{0.6}$ were reduced, as compared to how much heat is now carried by the converted phonons that used to have $l \geq l_{0.6}$ but now have a reduced $l \leq l_{0.6}$. A much larger *fraction* of the heat is now also carried by phonons with $l \leq l_{0.6}$, and so $l_{0.6}$ of the original material may be equal to $l_{0.95}$ of this new modified material. This is the case for grain boundary scattering introduced by nano-grains within a material. If the average grain size in a material is 550 nm, then phonons whose mean free paths were already < 50 nm hardly notice the difference. However, phonons whose MFPs were previously on the order of a few microns now experience significantly more scattering and have a considerably reduced MFP as a result. Therefore, this kind of nanostructuring primarily attacks long MFP phonons.

Alternatively, it is also possible to introduce scattering mechanisms that attack the short end of the MFP spectrum, and further reduce l for phonons with, for example, $l < l_{0.6}$. This would have the effect of decreasing the amount of heat carried by the short MFP phonons, causing the heat carried by the longer MFP phonons to now constitute a larger fraction of the total thermal conductivity. If this is the case, then $l_{0.6}$ of the original material may be equal to $l_{0.4}$ of the new modified material, because now a larger fraction of the heat is carried by longer MFPs

than before. This is the case for point defect scattering. If a point defect is on the order of or much smaller than phonon wavelengths, it obeys Rayleigh scattering physics. Rayleigh phonon scattering strength scales as the fourth power of phonon frequency, and so aggressively attacks the high frequency phonons. High frequency phonons correspond to short MFP phonons, and so point defect scattering targets short MFPs. One example of point defect scattering is the alloy scattering from isolated alloyed atoms in alloys such as SiGe. These examples demonstrate that the types of scattering mechanisms introduced will not only affect how strongly k is reduced, but also how the transport of heat is distributed among the different phonons. Two materials could have the same k_{bulk}, while having very different distributions of heat among their phonons. These two materials would then respond very differently to nanostructuring, as we will investigate.

4.1.5 Phonon Mean Free Path Spectroscopy

Once a material has been modified by the addition of bulk phonon scattering mechanisms or by constricting external dimensions orthogonal to the direction of heat transport, it will have reduced MFPs and a new distribution of heat among its phonons. This new k_{bulk} will be less than the original k_{bulk}, and so to avoid confusion we call it k_{nano}. We can replace l_{bulk} with l_{nano} in Eq. (4-11) and multiply and divide by l_{bulk} in order to rearrange the equation back into the same form as Eq. (4-11), preserving the same integration variable with one added multiplicative term [125],

$$k_{nano} = \int_0^\infty \left[-\sum_s \frac{1}{3} C v l_{bulk} \left(\frac{dl_{bulk}}{d\omega} \right)^{-1} \right] \frac{l_{nano}}{l_{bulk}} dl_{bulk} \tag{4-13}$$

Note again that unlike $k_{apparent}$, k_{nano} represents a real reduction of the true bulk thermal conductivity of the nanostructured material. Recall that both l_{nano} and l_{bulk} are functions of ω, and represent full spectra of phonons. Here, l_{bulk} refers to the frequency-dependent MFPs as they existed in the original bulk material, and l_{nano} refers to the new spectrum of MFPs in the nanostructured material. The ratio l_{nano}/l_{bulk} is often called the suppression function, S, and it is a function of the ratio of the characteristic length scale of the system, L_{char}, and the bulk mean free paths, l_{bulk} [72], [125], [128], [129]. Eq. (4-13) lets us easily calculate the thermal conductivity of a nanostructured system given knowledge of the bulk phonon mean free paths and characteristic lengths of the system. Eq. (4-13) is a very powerful tool, as it means that we can measure the bulk phonon mean free path spectrum once, and from it predict how the same material will transport heat under most kinds of arbitrary nano-structuring using the much more analytically solvable Fourier law equations (as opposed to BTE), provided that we know its suppression function. There is therefore great utility in being able to measure a bulk MFP spectrum.

But measuring a phonon MFP spectrum is tricky. We could try nanostructuring a material in a known way and correlating the characteristic lengths used with the resulting reduced k_{nano}, but as Eq. (4-13) shows, the new nano-material will have a different MFP spectrum than the original material whose spectrum we wish to measure. Therefore, we would not have really measured the target MFP spectrum, but rather just messed up our original sample. We need a way to measure the MFP spectrum without changing it. Herein lies the beauty of $k_{apparent}$. Unlike k_{nano}, $k_{apparent}$ is an apparent measured thermal conductivity whose measurement does not actually change any of the MFPs in the host material, and whose suppressed magnitude

depends in a known way on the relative sizes of L_{char} and l. Therefore, a set of measurements of $k_{apparent}$ as a function of L_{char} provides the quantitative information about the MFP spectrum of the bulk material that we seek. This makes $k_{apparent}$ the perfect tool for phonon mean free path spectroscopy. As was discussed at length in sections 4.1.2 and 4.1.3, a heated embedded sphere with radius on the order of l yields a measurement of a reduced $k_{apparent}$. Embedding, heating, and measuring a bunch of nano-spheres of varying radii is not an easy task. Fortunately, the concept of an embedded heated nano-sphere can be extended to the generalized case of any nano-scale point-like heat source, and can exist at the boundary of a sample. A laser beam spot directed onto the surface of a sample closely approximates a heated sphere for penetration depths much larger than the radius of the beam, and so satisfies all of these requirements, provided that its radius can be made very small. For very short penetration depths, the laser approximates a thin film, which works equally well for acquiring measurements of $k_{apparent}$. This is the principle behind using variable spot size TDTR as a phonon MFP spectroscopy technique.

We use high-magnification microscope objectives to concentrate our laser spots down to very small radii, on the order of phonon MFPs. This results in non-diffusive phonon heat transport in the immediate vicinity of the heated laser spot; the region to which the measurement is most sensitive. The thermal model used to fit our data, derived in section 3.6 TDTR Mathematical Theory, is based on Fourier's law and the diffusive heat conduction equation. Therefore, as a consequence of using Fourier's law to interpret data from non-diffusive thermal measurements, we measure a reduced $k_{apparent}$ without actually affecting k_{bulk} of our measured samples. Unfortunately, recovery of the exact MFP spectrum from this series of $k_{apparent}(L_{char})$ measurements would require exactly knowing the suppression function, S, for our system. This would require an exact solution of the 3-D, time-dependent, spectral BTE for the quite involved TDTR system, which has very complicated spatial and temporal distributions of heat. The solution would also need to rigorously include material interfaces and be valid for finite domains. While there are some recent BTE solutions in the literature that come close to meeting these conditions [122], [123], [130], [131], there is no existing solution that satisfies the precisely stated requirements. Deriving such a solution or performing the full numerical analysis would be highly nontrivial, and is beyond the scope of this work. Even if a solution were known, its fully accurate implementation would necessitate knowledge of the spectral phonon transmission coefficients between the Al transducer layer and the sample, as well as the phonon specularity parameter, neither of which are known or measurable with current techniques. Therefore, instead of trying to recover a full MFP accumulation function from our samples, we note from our simpler models that $k_{apparent}$ shows suppression below k_{bulk} when L_{char} is approximately equal to $\frac{4}{3}l$ and shows continued suppression as L_{char} is further decreased below l, scaling with L_{char} (Eq. (4-10)). From these facts, we can draw general conclusions about MFP spectra based on our measurements, as well as identify consistent trends. The next sections will examine the results of our MFP spectroscopy experiments.

4.2 Control Measurements: Silicon and Sapphire

We first validated our system on control samples of single crystal silicon and sapphire to confirm that we were able to recover established literature values of thermal conductivity. After validation with relatively large (w = 6.1 μm $1/e^2$ radius) spot sizes, we reduced spot sizes to w = 2.9 μm and then 1.6 μm, in order to see suppression of the bulk thermal conductivity due to non-

diffusive heat transfer in the vicinity of our laser spots when $w \sim l$. The results of these measurements are summarized in Figure 4-4.

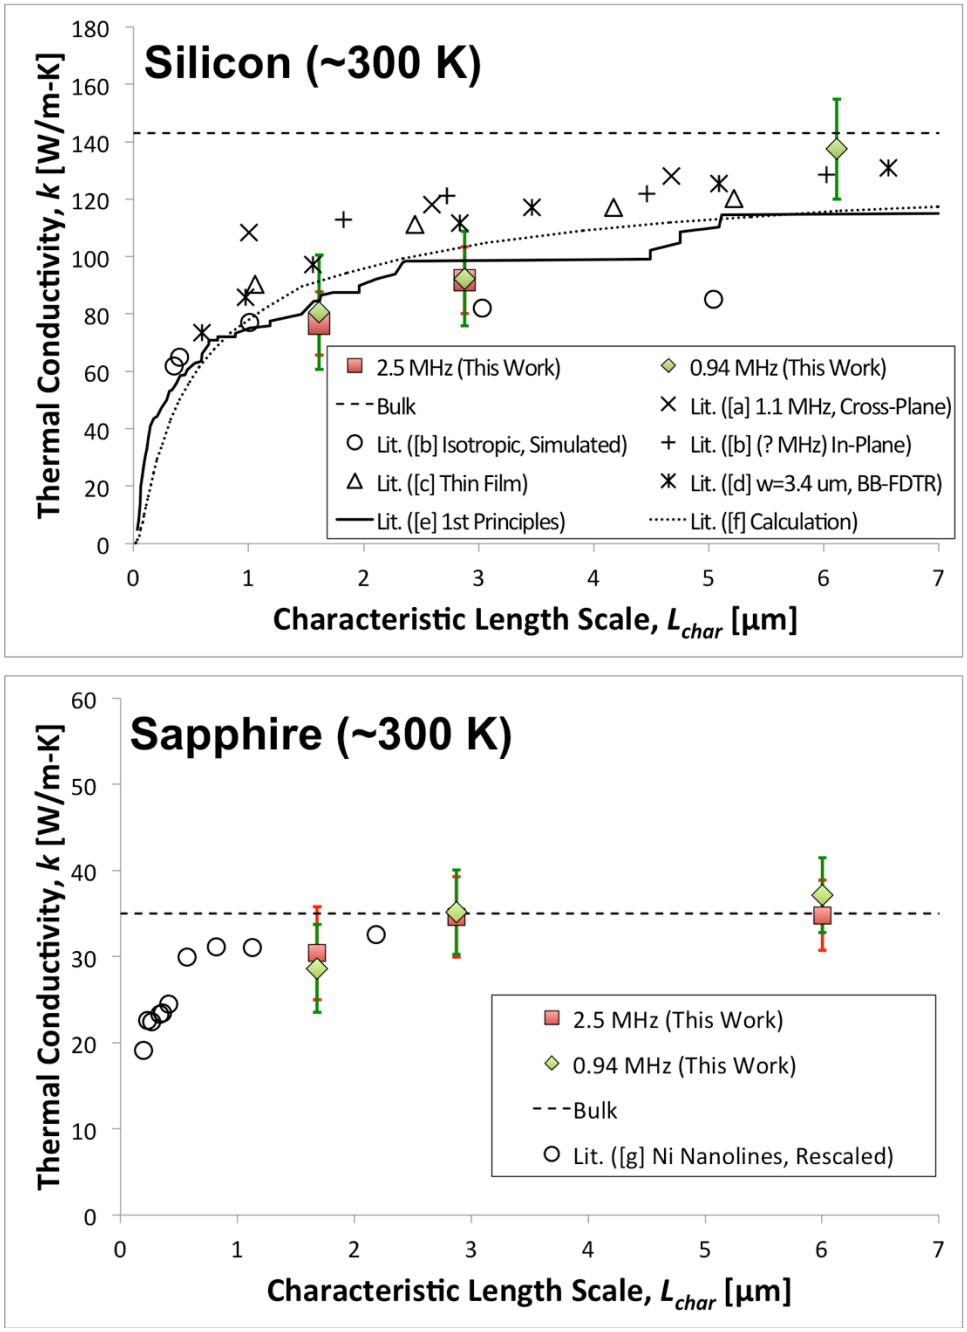

Figure 4-4 Control variable spot size TDTR measurements on single crystal silicon and sapphire at 300 K. Bulk thermal conductivity values shown by dashed lines. Apparent thermal conductivity is measured as a function of a characteristic micro length scale of heating, L_{char}. For our measurements, performed at pump modulation frequencies of 2.5 MHz and 0.94 MHz, L_{char} is the $1/e^2$ pump and probe radius, w. Points are the average of 1 to 4 repeated measurements. Error bars are conservative estimates of uncertainty, fully explained in section 3.7.2. The L_{char}, type of measurement, and reference for plotted literature data are: [a] – pump and probe $1/e^2$ spot radius, beam-offset anisotropic TDTR [79], [b] – $1/e^2$ pump (anisotropic TDTR) or pump and probe spot radii (isotropic Monte Carlo Boltzmann Transport Equation simulation) [132], [c] – thickness of Si film, large-spot TDTR [96], [d] – thermal penetration depth, broadband FDTR [80], [e] –

phonon MFP*4/3 (See Eq. (4-10)), first principles calculations [125], [127], [f] – phonon MFP*4/3 (See Eq. (4-10)), theoretical calculation [75], [76], [g] – heater line width, ultrafast soft coherent X-rays [133].

We successfully recover bulk room temperature thermal conductivity for our largest spot sizes for both silicon and sapphire, with agreements to literature of 4.2% and 6.0% respectively (low modulation frequency data). Agreement is within estimated measurement uncertainty, as shown by the dashed lines in Figure 4-4 passing within the bounds set by the error bars.

Sapphire only shows thermal conductivity suppression for the smallest spot size, for which we measure an 18% suppression of apparent thermal conductivity as compared to the bulk value. This degree of suppression is consistent with measurements of thermal conductance suppression using ultrafast soft coherent X-ray pulses [133], whose data is also plotted for comparison. The data from Ref. [133] gives the measured thermal conductance from nano-scale Ni heater lines into a sapphire substrate as a function of the heater line width. To facilitate comparison to our measurements, the thermal conductance data was normalized to its bulk value, and then rescaled to sapphire's bulk thermal conductivity. While not directly equivalent to our experiments, the data from Ref. [133] captures thermal transport suppression effects due to small heater sizes on the order of long MFPs creating non-diffusive thermal transport within their immediate vicinity, and so is a fair comparison for the kind of physics we are targeting. No other measurements of sub-diffusion thermal transport in sapphire exist in the literature to our knowledge.

Silicon shows much stronger suppression, implying that at room temperature the dominant heat carrying phonons in silicon have much longer mean free paths than in sapphire. We measure a maximum thermal conductivity suppression in Si of 44% as compared to its bulk value for a laser spot $1/e^2$ radius of w = 1.6 μm. This is consistent with theoretical calculations from first principles [125], [127] and traditional models [75], [76] (plotted for comparison) of Si thermal conductivity at 300 K. These models are plotted as a function of $L_{char} = \frac{4}{3}l$ for consistent comparison to L_{char} of our and other experiments (see Eq. (4-10) and Ref. [115]). For a cutoff MFP of $\frac{4}{3}l \sim 1.6\ \mu m$, these models predict apparent Si k suppressions of 41% and 36%, respectively. Our measurements fall on the lower end of the scatter of literature data, meaning that we measure around a 10% to 20% stronger k suppression than measured by others with different techniques. Note that the literature Si data plotted in Figure 4-4 constitutes a variety of kinds of geometries and measurement techniques, whose characteristic thermal length scales (L_{char}) can therefore only be approximately compared among each other. As a consequence, clustering of data from dissimilar experiments is partially coincidental. In total, the experiments represented by the plotted Si literature data include broadband frequency domain thermoreflectance (BB-FDTR) (L_{char} is the thermal penetration depth) [80], Si thin films (L_{char} is the film thickness) [96], Monte Carlo simulations of TDTR using the Boltzmann Transport Equation (L_{char} is the simulated $1/e^2$ pump and probe radius) [132], first principles and theoretical calculations of Si thermal conductivity (L_{char} is 4/3 the phonon MFP) [75], [76], [125], [127], and various flavors of TDTR (L_{char} is various combinations of pump and/or probe $1/e^2$ radii) [79], [132].

For the other TDTR experiments, which are the most directly comparable, some used an anisotropic thermal conductivity model to fit their data and reported the cross-plane value, while we used an isotropic model. Recent literature predicts that the breakdown of Fourier Law at micro and nano length scales may happen anisotropically, depending on the dimension that is being constricted [79], [132]. However, we have inconsistent sensitivities to anisotropy across

all of our samples. For some samples, the sensitivity is low enough that fitting an anisotropic model would introduce more error than might already exist from using the isotropic model in the first place. Therefore for consistency, and following the recommendation of Ref. [132] for our geometry, we have used an isotropic model to fit all of our data. It should also be noted that some of the plotted TDTR literature data primarily reduced the size of their pump beam while keeping the probe beam spot size mostly fixed. TDTR *k* suppression measurements from different groups have previously had difficulty finding close agreement when one group shrunk both pump and probe while the other group shrunk only one spot size [79], [126].

4.3 Measurements of Nano-Grained Silicon

We now present our measurements of nano-grained silicon, which are some of the same samples that were measured in bulk by 3-Omega in Ref. [99]. To encapsulate the full effect of grain boundaries, we chose to measure the samples with the largest and smallest average grain sizes that were still fully dense (i.e. both samples had only 1% porosity) from that sample set. This corresponded to the 550 nm and 76 nm average grain size samples. To our knowledge, these are the first direct sub-continuum thermal transport measurements of a nanostructured bulk crystal, and there is no prior literature against which we can compare our results. The results of these measurements are summarized in Figure 4-5.

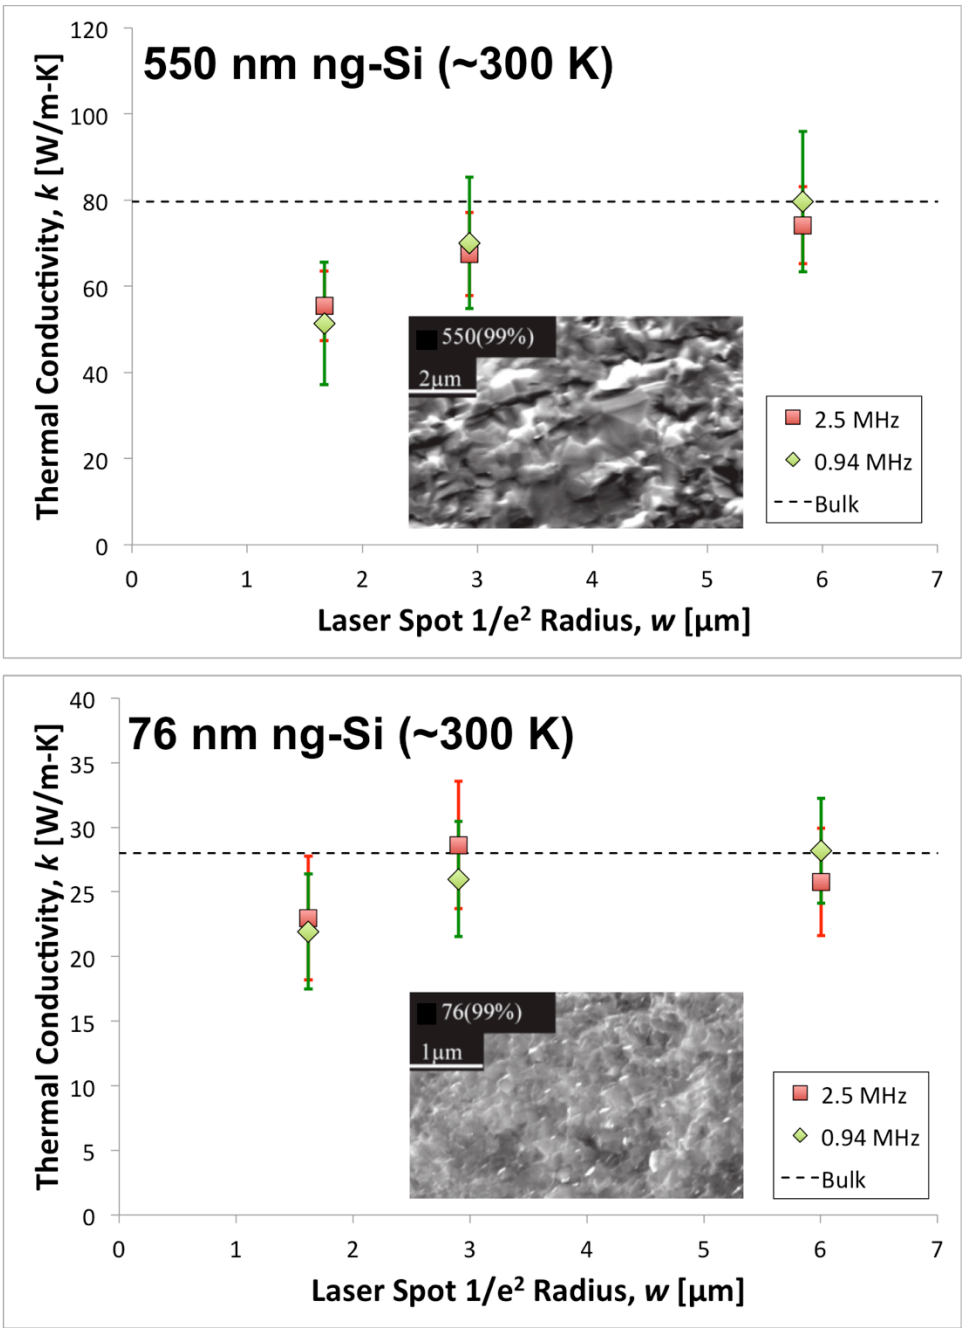

Figure 4-5 Variable spot size TDTR measurements of nano-grained silicon samples at 300 K. Bulk thermal conductivity values, taken from 3-Omega measurements performed in Ref. [99], are shown by dashed lines. Samples have average grain sizes of 550 nm and 76 nm, and both are 1% porous. Insets show SEM micrographs of fracture surfaces. Measurements were performed using pump modulation frequencies of 2.5 MHz and 0.94 MHz. Points are the average of 1 to 9 repeated measurements. Error bars are conservative estimates of uncertainty, fully explained in section 3.7.2.

For both samples, we successfully recovered bulk thermal conductivities with our largest (w = 5.9 µm 1/e^2 radius) spot sizes with agreements of 0.4% and 0.7% for the 550 nm and 76 nm samples respectively (low modulation frequency data). This implies that, like with sapphire and single crystal silicon, the vast majority of heat in nano-grained silicon samples at room

temperature is carried by phonons with $\frac{4}{3}l < \sim 6$ μm. Bulk values were taken from the 300 K 3-Omega measurements of these same physical samples performed by Ref. [99]. References [79], [132] predict that increasing modulation frequency should cause k suppression in the cross-plane direction, while shrinking spot sizes will cause suppression in the radial direction. Experiments will be most sensitive to cross-plane suppression for large spot sizes, when there is not also simultaneous radial suppression. For this reason, we always use our lowest frequency data to evaluate bulk thermal conductivity with our largest spot sizes, and to calculate the spot size-induced suppression of k. While it is true that for our largest spot sizes, the 2.5 MHz data shows a lower k than the 0.94 MHz data, it is still within the measurement uncertainty and the trend is not consistent with smaller spot sizes. Therefore, we can neither confirm nor reject Ref. [79]'s frequency-dependent k suppression hypothesis from this data. Ref. [79] predicts that frequency-dependent suppression effects should be weak in silicon.

The 550 nm sample is most similar to single-crystal silicon. In comparison, it shows a slightly lower maximum thermal conductivity suppression of 36% for w = of 1.7 μm. This implies that around 8% less of the 550 nm sample's heat is carried by phonons with $\frac{4}{3}l > \sim 1.7$ μm as compared to single crystal silicon. This is consistent with the nanostructuring reducing both k and phonon MFPs. However, it is surprising that our smallest laser spot size, with a $1/e^2$ radius still three times larger than the average grain size in the sample, is still able to suppress the thermal conductivity by as much as 36%. This implies that in spite of the nanostructuring, a considerable fraction of the 550 nm sample's phonons still have MFPs longer than ~550 nm. This is consistent with the results of Ref. [99], which predicts a frequency-dependent transmissivity of grain boundaries, or with a gray transmissivity model with transmissivity considerably greater then zero. Ref. [99] predicts that the scattering strength of grain boundaries scales as the inverse of the phonon frequency, meaning that high frequency phonons scatter more strongly while low frequency phonons will scatter more weakly and thus still have MFPs considerably longer than the grain sizes.

It should be noted that because the 1.7 μm spot is only a few times larger than the grain sizes, the precise placement of the spot in relation to individual grains may be of importance. For example, a measurement where the laser spot is in the center of a large grain may measure $k_{apparent}$ closer to that of pure Si as compared to a measurement where the spot happens to fall directly on top of or close to a grain boundary. The laser spot will average over far fewer grains than with the 76 nm sample, which can be treated as homogeneous from the laser's point of view. To mitigate this concern, we took the greatest number of repeated measurements (9) of all samples for the 550 nm 20x objective (i.e. w = 1.7 μm) 2.5 MHz and 0.94 MHz data points.

The 76 nm sample continued all trends set by the 550 nm sample. The 76 nm sample showed a weaker suppression of only 22% for our 1.6 μm $1/e^2$ radius spot. This is consistent with our expectations from both preceding trends. Further reducing the average grain size further reduces phonon MFPs, resulting both in reduced bulk k_{nano} and weakened suppression of $k_{apparent}$. Because the phonon MFPs are all reduced as compared to the 550 nm sample, shrinking w to 1.7 μm does not lead to heat transport near the spot that is as ballistic as it was for the 550 nm sample because the now much smaller MFP phonons are still diffusive at these length scales, and so the k suppression is weaker. Additionally, as explained toward the end of section 4.1.4, nanostructuring will specifically attack the long MFP phonons, shifting the spectral distribution of heat toward shorter MFP phonons. As a result, one would expect to see less suppression of $k_{apparent}$ for fixed w when nanostructuring a material and shifting its heat more toward MFPs $\frac{4}{3}l < w$. The nanostructuring thus has the effect of restoring diffusive type thermal transport for the

heating by $w > \frac{4}{3}l$, consistent with our observation of a weakened thermal conductivity suppression relative to bulk.

At the same time, the fact that there is any suppression at all when our spot is now an order of magnitude larger than the average grain size is further support to the phonon-frequency-dependent grain boundary scattering proposal of Ref. [99], or a gray transmissivity model with high transmissivity.

4.4 Measurements of Silicon Germanium

We now present our measurements of silicon germanium alloys. In analogy with the nano-grained silicon samples, we measured a sample with both high and low germanium concentration in order to examine the effect of mass defect phonon scattering from alloying. Because of the saturation of the alloying effects for the middle ~ 90% of the range of possible Ge concentrations [100], [134], it is important to choose one sample within this range and one dilute sample outside this range with only light alloying in order to properly study the effect of alloying. We chose samples with germanium concentrations of 1 atomic % ($Si_{99}Ge_1$), and 9.9 atomic % ($Si_{90.1}Ge_{9.9}$). Acquiring bulk, single crystal SiGe samples with controlled alloy concentration is very difficult. The $Si_{90.1}Ge_{9.9}$ sample was grown and characterized by N. Abrosimov of the Leibniz Institute for Crystal Growth, while the $Si_{99}Ge_1$ sample is the same as was used in Ref. [79]. The results of our measurements are summarized in Figure 4-6.

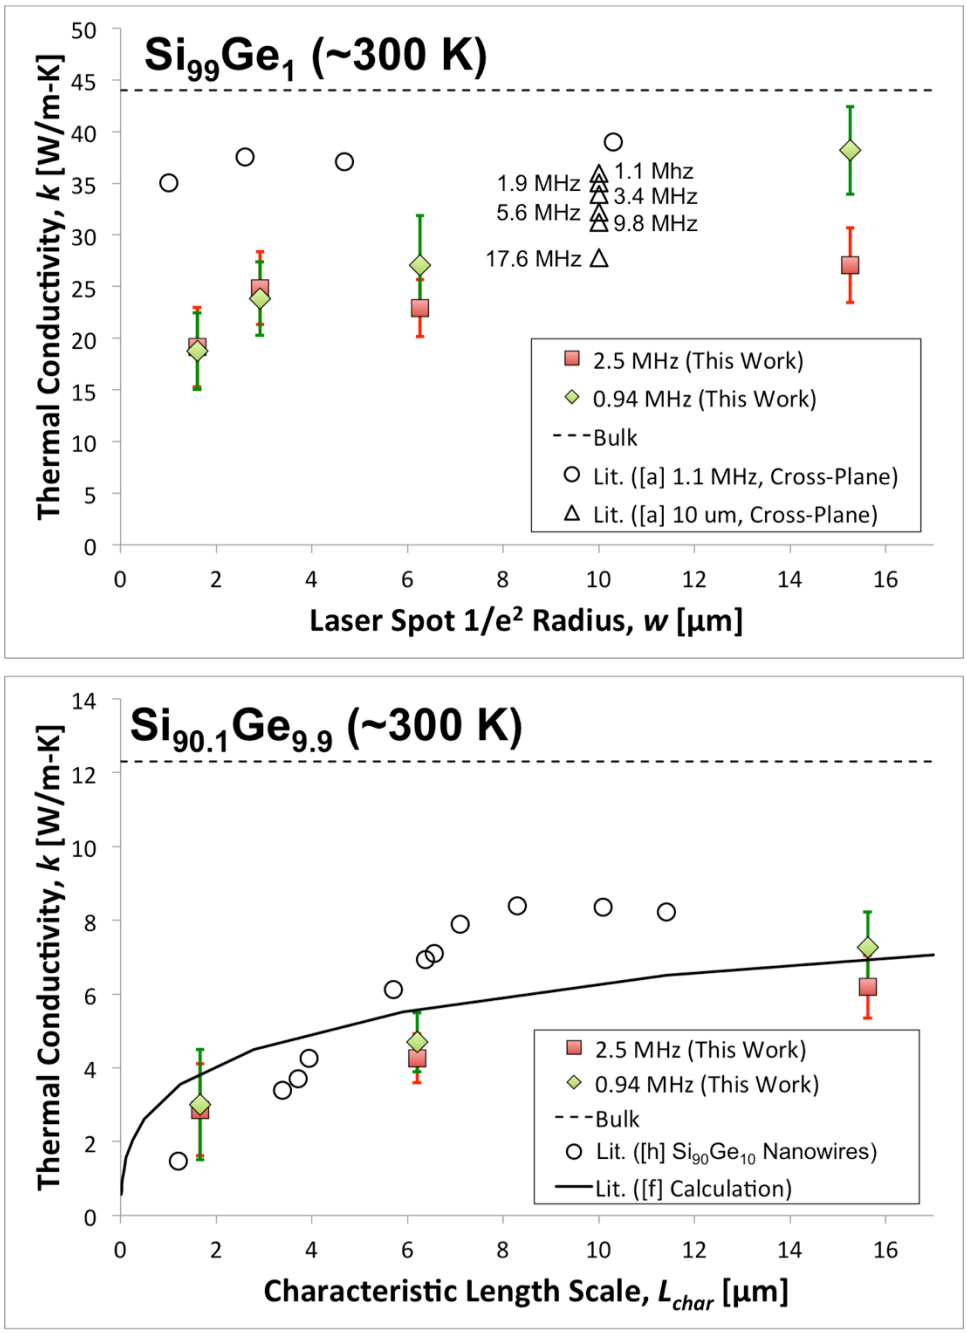

Figure 4-6 Variable spot size TDTR measurements of silicon germanium samples at 300 K. Bulk thermal conductivity values are shown by dashed lines. The bulk *k* value for $Si_{99}Ge_1$ is taken from Ref. [79], while we measured the bulk *k* value for $Si_{90.1}Ge_{9.9}$ via the 3-Omega method, using an 80 µm x 1500 µm heater line (thermal penetration depth ranging from 75 µm to 350 µm). Samples have 1.0 and 9.9 atomic percent Ge concentrations. Measurements were performed using pump modulation frequencies of 2.5 MHz and 0.94 MHz. Points are the average of 2 to 6 repeated measurements. Error bars are conservative estimates of uncertainty, fully explained in section 3.7.2. The L_{char}, type of measurement, and reference for plotted literature data are: [a] – pump and probe $1/e^2$ spot radius, beam-offset anisotropic TDTR [79], [f] – phonon MFP*4/3 (See Eq. (4-10)), theoretical calculation [75], [76], [h] – $Si_{90}Ge_{10}$ Nanowire length, suspended heated platforms [135].

For comparison of our $Si_{99}Ge_1$ measurements we have included TDTR data from Ref. [79] (also for $Si_{99}Ge_1$) for both fixed w = 10 μm / variable modulation frequency (1.1 MHz to 17.6 MHz), and fixed modulation frequency (f = 1.1 MHz) / variable w. To our knowledge, no other direct sub-continuum thermal measurements of similar SiGe alloys exist in the literature. We measure considerably stronger $k_{apparent}$ suppression for small w than Ref. [79], although we appear to converge to similar values for large w. However, their measurements appear to fail to converge all the way to bulk values (not shown is their measurement for w = 26 μm, which still measures the same k value as our w = 15.3 μm point, and the same as their w = 10.3 μm point). Our largest spot size is 15.3 μm, so we cannot confirm whether we would also fail to recover bulk k values. This may be due to persistent modulation frequency-induced suppression, independent from spot size suppression. Beyond the facts that Ref. [79] used an anisotropic model to fit their data and report a cross-plane $k_{apparent}$ while we used an isotropic model, and they used a sweeping knife-edge approach to measure their spot sizes for data fitting while we used a 2-D Gaussian fit, we cannot identify any other significant differences between our respective measurement protocols that would explain such different results. The disagreement between their two 1.1 MHz points (one at w = 10 μm, one at w = 10.3 μm), taken from different data sets of the same paper corresponding to different parametric sweeps, gives an indication of the size of their error bars.

As previously mentioned, Ref. [79] predicted a modulation frequency dependent suppression of $k_{apparent}$, with stronger suppressions expected for higher frequencies. This frequency-dependent suppression of $k_{apparent}$ is demonstrated in their plotted frequency sweep data at w = 10 μm. Ref. [79] predicted this effect would be weak for Si, but strong for SiGe, especially for large spot sizes where it is not competing with radial spot size-induced $k_{apparent}$ suppression. We do observe a stronger apparent frequency dependence of $k_{apparent}$ for SiGe samples than we did for previous samples, even beyond the uncertainty of our measurements for our largest spot size in $Si_{99}Ge_1$. Our tentative frequency-dependent suppression is stronger than what is observed by Ref. [79], consistent with our spot size-dependent suppression also being stronger. This may again be due to differences in anisotropic vs. isotropic measurements, as well as how spot size is calculated both for reporting results and for fitting raw data. However, we have insufficient data across all spot sizes given our measurement uncertainty to confidently comment on the pump laser modulation frequency-dependence trend. Our higher frequency data also shows a less consistent trend with spot size. A more comprehensive data set including a larger frequency sweep would be necessary to more rigorously investigate this claim.

Both SiGe alloys fail to recover bulk thermal conductivity values for large spot sizes, and so we cannot place an upper bound on the maximum phonon MFP relevant to heat transfer at room temperature in either sample based on our variable spot size TDTR data. The bulk k value for $Si_{99}Ge_1$ of 44 W/m-K is taken from Ref. [79]. For comparison, literature values for $Si_{99}Ge_1$ k_{bulk} are 33 W/m-K [100] and 35 W/m-K [134], indicating that our sample taken from Ref. [79] is either a higher k than usual, or more likely a slightly lower germanium concentration than what was reported. The thermal conductivity of SiGe is a very steep function of alloying concentration for dilute alloys, and a value of k_{bulk} = 44 W/m-K would correspond more closely to a 0.7% germanium concentration instead of 1% [134]. Note that using the lower literature values for k of $Si_{99}Ge_1$ would not change any of the major findings of our results discussed below. We measured the bulk k of the $Si_{90.1}Ge_{9.9}$ sample via a traditional 3-Omega method, using an 80 μm wide x 1500 μm long x 0.2 μm thick Au heater line (thermal penetration depth ranging from 75 μm to 350 μm) on top of a ~ 1 μm thick spin-coated electrically insulating

photoresist layer. Our 3-Omega measurements yielded a bulk k of 12.3 ± 0.2 W/m-K for the $Si_{90.1}Ge_{9.9}$ sample, agreeing with literature values to within 7.5% [100] and 0.8% [134]. This implies that the vast majority of heat in $Si_{90.1}Ge_{9.9}$ is carried by phonons with $\frac{4}{3}l$ less than at most ~80 μm at room temperature. Even with w ~ 15.5 μm $1/e^2$ radius spot sizes, more than twice the size of the largest spot sizes used on other samples, we still did not recover bulk k values for either SiGe sample. This implies that at room temperature, a considerable fraction of the heat in SiGe alloys is carried by phonons with $\frac{4}{3}l$ > 15.5 μm. These are very large mean free paths for 300 K, and is a somewhat striking result. This "long tail" in SiGe is still consistent with the theory in literature, which predicts that in order to account for 90% of the heat carried in SiGe, one needs to consider phonon MFPs as high as 100 μm [76], [111] (see Figure 4-3).

Our measurements of the $Si_{90.1}Ge_{9.9}$ sample are in reasonable agreement with the theoretical calculations of $Si_{90}Ge_{10}$'s k accumulation function from [75], [76] (plotted in Figure 4-6 for comparison, and in semilog scale in Figure 4-3). We have also plotted literature data for $Si_{90}Ge_{10}$ nanowire measurements [135] in Figure 4-6. Note that L_{char} for the nanowire measurements is their length, and not their diameter as is more conventional in nanowire studies. Ref. [135] somewhat surprisingly did not see a difference in measured k values for nanowires of varying diameters, provided they had the same length. Therefore, it is a reasonable assumption that the nanowire length was the dominant length scale responsible for diffusive vs. ballistic axial heat transport regimes. Therefore, this is an appropriate comparison. This could be explained if the nanowires had very smooth surfaces (as is indicated by the STEM images in Ref. [135]), as specular phonon scattering by nanowire walls does not reduce axial phonon momentum and therefore has very minimal impact on axial heat conduction [76]. Such an effect is also more likely in alloy systems such as SiGe, where long wavelength phonons (which see nanowire walls as smoother) are more dominant as a result of the alloy Rayleigh scattering. In this case, the nanowire length would be the dominant length scale responsible for increased thermally resistive phonon scattering, even though the length was much greater than the diameter. The data from Ref. [135] does plateau to a value around k = 8 W/m-K (less than the bulk value of 12.3 W/m-K) for nanowire lengths ≥ 8 μm. We speculate that at there was a regime switch for lengths ≥ 8 μm and the nanowire diameter became the dominant length scale suppressing axial heat transport, which then became diffusive leading to a constant value of k_{nano} with increasing L_{char}. Beyond nanowire lengths of ~5μm, Ref. [135] did not measure nanowires with varying diameters anymore, and so this speculation can be neither confirmed nor disproved (one would expect the k leveling-off to occur at shorter lengths for smaller diameter nanowires according to our "regime shift" hypothesis). Consequently, we feel it is only appropriate to compare our results to data from Ref. [135] for nanowires with lengths ≤ 8 μm, for which we believe the nanowire length was still the dominant length scale causing non-diffusive phonon transport, and therefore still the appropriate choice of L_{char} for plotting. We have still plotted all data for completeness, including for nanowire length > 8 μm. Ref. [135] measured nanowires with both 10% and 60% Germanium concentrations, and as with the varying diameters, did not see this influence measured k values. This implies that for nanowire lengths ≤ 8 μm, neither the alloying concentration nor the diameter (perpendicular to the heat transport direction) was the primary resistive phonon scattering mechanism for axial transport. It would therefore be interesting to perform variable spot size TDTR measurements on higher Ge concentration SiGe alloys, to see if they collapse down to the same k values for small spot sizes. This would be consistent with the "saturation effect" on the influence of mass point defect phonon scattering

with increased alloying (i.e. that that k_{bulk} quickly levels off and is roughly constant for the middle ~80% of alloying concentrations [100], [134]).

Both SiGe alloys show very strong $k_{apparent}$ suppressions for w = 1.6 μm. We measure a 57% and 76% thermal conductivity suppression with w = 1.6 μm for the $Si_{99}Ge_1$ and $Si_{90.1}Ge_{9.9}$ samples, respectively. This implies that the majority of the heat in both samples at room temperature is carried by phonons with $\frac{4}{3}l$ > ~1.6 μm, consistent with theoretical predictions [75], [76] (see Figure 4-3). This could have significant relevance to microelectronics and microstructured thermoelectric materials that use SiGe alloys [136].

At first, these two observations seem contradictory. On the one hand, increasing the alloying concentration of SiGe reduces thermal conductivity, due to the increased phonon scattering and hence shorter phonon MFPs. On the other hand, increasing the alloying concentration of SiGe also increases the level of k suppression for fixed w, seemingly implying that the phonon MFPs have become larger in relation to the (fixed) laser spot size. How can phonon MFPs simultaneously be getting longer and shorter with increasing alloying concentration? The paradox is resolved by remembering that it is possible to change the *relative* distribution of heat among different kinds of phonons when introducing new scattering mechanisms, as was discussed at the end of section 4.1.4. Imagine introducing a new scattering mechanism that selectively scatters short phonon MFPs, making them even shorter, while leaving the long MFPs relatively unscathed. This concept is portrayed qualitatively in Figure 4-7 in section 4.5 Summary of Variable Spot Size TDTR Findings. The net result would be to reduce the amount of heat carried by the short MFP phonons, globally shifting the "center of mass" of the spectral k distribution to the right toward longer MFPs so that a larger *fraction* of the total heat were carried by long MFP phonons as compared to before, even if they now carried a slightly reduced *absolute* quantity of heat. Because a larger fraction of heat is carried by long MFP phonons than before, it means that a larger fraction of the heat coming from a heat source with L_{char} smaller than those long MFPs will be transported non-diffusively, and so a stronger suppression of $k_{apparent}$ will be measured. Simultaneously, the new bulk thermal conductivity of the material, k_{bulk}, will be reduced due to the further MFP reduction of short MFP phonons. This combination of a reduced k_{bulk} but an increased suppression of $k_{apparent}$ for fixed w is the precise combination of effects we observe in our SiGe data, with the trends growing stronger for increased alloying concentration.

Accomplishing this would require a very selective scattering mechanism that primarily attacked short MFP phonons. Indeed as previously discussed, alloy scattering is a result of point mass defects much smaller than the phonon wavelengths, hence it obeys Rayleigh scattering physics, and the scattering rate scales with the fourth power of phonon frequency [75], [137]–[139]. Phonon MFP, l, is generally a monotonically decreasing function of phonon frequency, ω, meaning that high phonon frequencies correspond to short MFPs and vice-versa. Therefore, alloying very selectively scatters the high frequency (short MFP) phonons, while minimally affecting the low frequency (long MFP) phonons. This is the exact required behavior necessary to explain our results. Additionally, the fact that even very light alloying (by only a few percent) significantly reduces k_{bulk} (by over an order of magnitude [134]), means that alloy scattering is the dominant form of high frequency phonon scattering for SiGe at room temperature. Thus the fourth power physics of alloy scattering governs the phonon frequency-dependence of high frequency phonon MFPs in our samples (Umklapp scattering governs the low frequency phonon scattering). Our results are fully explained by the increased alloy scattering in SiGe.

4.5 Summary of Variable Spot Size TDTR Findings

In this section, I will concisely reiterate the primary conclusions from all measured samples in one unified discussion, and make some final observations. The bulk values of k for all samples, their reduced apparent values of k as measured using the smallest pump and probe laser spot $1/e^2$ radii of w ~1.6 μm, and the corresponding apparent relative suppressions in k, are summarized in Table 4-1. Arrows indicate important and differing trends in the data. Full data are plotted in Figure 4-4, Figure 4-5, and Figure 4-6.

Table 4-1 Summary of Variable Spot Size TDTR Measurements of All Samples. Arrows designate differing trends in k_{bulk} and $k_{apparent}$ vs. k suppression.

	Sapphire	76 nm ng-Si	550 nm ng-Si	Single Crystal Si	$Si_{99}Ge_1$	$Si_{90.1}Ge_{9.9}$
k_{bulk} [W/m-K]	35	28	80	143	44	12.3
		← Nanostructuring (~ D_{grain}/MFP)			Alloying (~ ω^{-4}) →	
$k_{apparent}$ ($w \approx 1.6$ μm) [W/m-K]	29	22	51	81	19	3.0
k suppression ($w \approx 1.6$ μm)	18%	22%	36%	44%	57%	76%
		←				→

The fact that suppression does not track monotonically with bulk k indicates the rich physics governing our results. We have examined two families of silicon-based samples, joined in the center by intrinsic single crystal silicon. Both families introduce additional phonon scattering mechanisms that reduce bulk k from that of silicon, but the way in which the different scattering mechanisms reshape the distribution of heat among phonons leads to opposite trends in observed k suppression, as depicted in Figure 4-7.

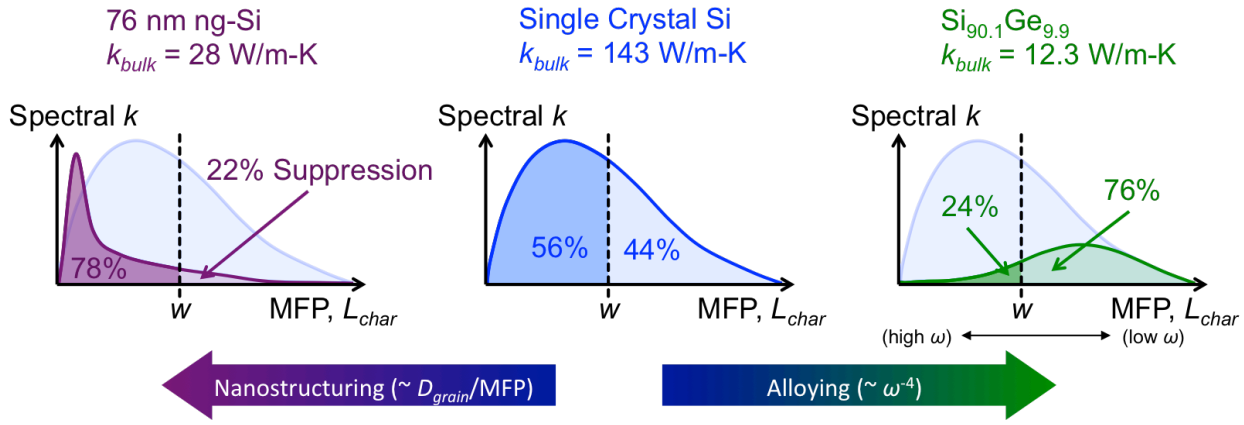

Figure 4-7 Qualitative concept schematic, not to scale, depicted values chosen to mirror Table 4-1. Total k is calculated by integrating the area under the spectral k curves over all contributing MFPs. Difference between the effect of alloying versus nanostructuring on the distribution of heat among different phonon MFPs shown schematically, consistent with the findings of Table 4-1. Alloy scattering strength scales with ω^{-4} and so attacks high frequency (short MFP) phonons, while nanostructuring scattering strength depends

on how large the MFP is compared to the nanograin diameter with a D_{grain}/MFP scaling, and so attacks long MFP phonons.

For the nano-grained silicon (ng-Si) family, grain boundaries scatter phonons with $l > D_{grain}$ significantly more than they scatter phonons with $l < D_{grain}$. This selectively reduces the amount of heat carried by long MFP phonons, shifting the heat toward shorter MFPs, as shown on the left of Figure 4-7. As a result, a fixed size small laser spot, w, will measure less suppression of $k_{apparent}$, because a larger fraction of the heat leaving w is transported diffusively due to the additional grain boundary scattering, and so Fourier's law is still valid. This is illustrated qualitatively in Figure 4-7 (left), where phonons with $l > w$ do not contribute to the measured $k_{apparent}$, because they travel ballistically and so are interpreted as carrying negligible heat flux by the fitted Fourier law model that over-estimates q'' for ballistic phonons. The increased scattering also reduces the bulk thermal conductivity (i.e. reduced area under the curve as compared to single crystal Si). The phonon transmissivity through grain boundaries is quite high for some or all phonon frequencies, as evidenced by the fact that significant suppression can still be measured even for w more than 10x larger than D_{grain}. Note that some of the low MFP phonons in nano-grained silicon can actually carry more total heat than the corresponding low MFP phonons in single crystal Si. This happens because the collection of all low MFP phonons in ng-Si includes all the phonons that used to have long MFPs in single crystal Si, but were moved to the "low MFP bin" by the grain boundary scattering. The reverse is not true for SiGe, because alloy scattering of the short MFP phonons shortens rather than lengthens their MFPs, and so cannot move them up into the "long MFP bin."

For the SiGe alloy family, point mass defects introduced by Ge alloying reduce the bulk thermal conductivity. The increased density of point scatterers in the $Si_{90.1}Ge_{9.9}$ sample leads to a stronger reduction of bulk k as compared to the relatively dilute $Si_{99}Ge_1$ sample. The Ge atoms act as point scatterers that are much smaller than the phonon wavelengths and hence lead to Rayleigh scattering. The predicted frequency dependence of the Rayleigh scattering strength (scaling as ω^4) selectively scatters the high frequency (short MFP) phonons more than the low frequency (long MFP) phonons [75], [137]–[139]. Alloying thus redistributes the carried heat toward long MFP phonons, leading to an increase in the observed k suppression for fixed w, as shown in Figure 4-7 (right).

Our measurements compare favorably with some literature values where available for silicon, sapphire, and $Si_{90.1}Ge_{9.9}$, and we observe a 10% to 20% stronger k suppression effect than what some others measured with different techniques for silicon and $Si_{99}Ge_1$. Our measurements are in fairly good agreement with model theory predictions and first principles calculations.

A final interesting observation, best seen in Table 4-1, is that while k_{bulk} of the 76 nm nano-grained Si sample (28 W/m-K) is much less than k_{bulk} of the $Si_{99}Ge_1$ sample (44 W/m-K), the suppressed $k_{apparent}$ of the 76 nm sample (22 W/m-K) is greater than the suppressed $k_{apparent}$ of the $Si_{99}Ge_1$ sample (19 W/m-K), using the same laser spot size of w = 1.6 μm and even though they are both silicon-based materials whose phonons have nearly identical properties other than MFP. Because $k_{apparent}$ is a tool to effectively measure the spectral k of a material without changing its MFP spectrum in the process, this means that given the choice between the two materials, 76 nm ng-Si would be a superior choice over $Si_{99}Ge_1$ for the thermal management of microstructures, despite having a much lower bulk thermal conductivity than $Si_{99}Ge_1$. This can be understood by again considering Figure 4-7. $Si_{99}Ge_1$ might have a larger bulk value of k

when including all MFPs (total area under the curve), but because most of its heat is carried by long MFP phonons with lower limiting ballistic fluxes than short MFP phonons, $Si_{99}Ge_1$ loses heat carrying capacity more rapidly than ng-Si when microstructured. This is a concrete example of the importance of understanding how the heat is distributed among the phonons in a material, as values of k_{bulk} alone do not tell the whole story and can lead to very incorrect design decisions for microstructured systems.

There is again a subtlety in properly understanding this result. To make the explanation clearer, imagine a material, "SG," whose spectral k distribution among different phonons looks like that of the $Si_{99}Ge_1$ sample (i.e. skewed toward long MFPs), and a material, "ngSi," whose spectral k distribution looks like that of the 76 nm nano-grained Si sample (i.e. skewed toward short MFPs). Except in our example, SG and ngSi have identical values of k_{bulk} while SG still has a much lower value of $k_{apparent}$ than ngSi for the same w. The reason SG has a lower suppressed $k_{apparent}$ than ngSi is not because a larger fraction of its phonons is being excluded by w. These long MFP phonons still participate in heat transfer, only now for small w they participate ballistically instead of diffusively. No matter how large a phonon's bulk MFP may be, it will still participate in heat transfer for any size of w, but it will participate ballistically for comparatively small w. Therefore, it is really the magnitude of a phonon's ballistic limit heat flux that determines how well the material will still transport heat when in a non-diffusive regime. However, ballistic heat fluxes represent a maximum possible flux, so one would therefore expect this to cause an *increase* in the heat carried by SG at small sizes relative to ngSi, because a larger fraction of its phonons will be carrying heat ballistically. But this is incorrect, and here is the subtlety. Because most of the heat in the bulk SG sample is carried by long MFP phonons, and because the total contribution to k from each phonon (indexed by its frequency, ω) is given by $\frac{1}{3}Cvl$ (Eq. (4-5) – each parameter itself a function of ω), it means that those phonons with large values of l must compensate by having smaller values of their Cv product in order to still only contribute the same amount to k as the same phonons (i.e. same ω phonons) in the ngSi sample that have smaller l. This Cv product is what determines the magnitudes of the phonons' ballistic heat fluxes (Eq. (4-8)). Therefore, when comparing the relative magnitudes of suppressed $k_{apparent}$ between SG and ngSi, it is largely the fact that the long MFP phonons in SG must have smaller ballistic heat fluxes than the corresponding phonons in ngSi that causes $k_{apparent}$ of SG to be less than $k_{apparent}$ of ngSi. It is not just simply because w "excludes" a larger chunk of spectral k at long MFPs in SG (although, this way of thinking fortuitously aligns with the fact that low ω phonons happen to simultaneously have longer MFPs and smaller Cv products in general, and vice-versa for high ω phonons). While most of the phonons in SG have longer MFPs than in ngSi, the phonons with the highest Cv products have longer MFPs in ngSi than in SG, because these are the phonons with very high ω that the alloy scattering attacks most aggressively. Said another way, SG gets most of its diffusive bulk k from its long MFPs, while ngSi gets most of its diffusive bulk k from its large Cv products. Therefore, when MFP leaves the picture in the case of ballistic transport, ngSi phonons have the advantage, even if some of them are still travelling diffusively while most of the SG phonons are traveling ballistically.

In conclusion, our results agree with the model that grain boundaries primarily scatter long MFP phonons and have high transmissivities for some or all phonon frequencies, while alloying primarily scatters short MFP phonons giving SiGe "long tails" in their k accumulation functions. At room temperature, most of the heat in Si and its nanostructured cousins is carried by phonons with MFPs smaller than a few microns, while an appreciable fraction of the heat in SiGe is still carried by phonons with MFPs up to a few tens of microns. More generally, we

show that different scattering mechanisms can have pronounced differences in how they redistribute the carried heat among different kinds of phonons, which in turn can have real and tangible consequences regarding the efficacy of different materials' thermal management for nano and microstructures, sometimes contrary to expectations based on bulk values of k. In particular, SiGe was found to be a poor choice for the thermal management of nano and microstructures.

4.6 Suggestions For Future Work

Phonon mean free path spectroscopy is a rich landscape that is still largely uncharted. There are many opportunities for both experimental and theoretical developments that could immediately resolve controversies and ambiguities. Here, we will discuss a few of the future research directions that we believe warrant time and investigation.

First and foremost, the kinds of sub-diffusion measurements presented in this dissertation should be repeated at varying temperatures, ideally going down to low cryogenic (< 10 Kelvin) temperatures. Reducing T changes k of silicon in very well studied and predictable ways [105] by reducing C of the phonons and reducing Umklapp phonon scattering rates. Therefore, comparing the expected shifts in phonon spectra and their corresponding shifts of k with different samples' actual observed k shifts can be used to further confirm or refute the way in which alloying and nanostructuring redistributes heat among phonons. For example, initial reductions in T from 300 K should increase phonon MFPs due to reduced Umklapp scattering, without yet decreasing C because samples will still be in the saturated Dulong and Petit limit for C [105], [140]. Therefore, we would expect to see a relatively small change in the $k_{apparent}$ suppression of ng-Si samples when reducing T because the nanograins keep phonon MFPs from getting too long even after the removal of Umklapp scattering, and so heat will not be significantly redistributed to longer MFPs to increase suppression. On the other hand, because most of the heat in SiGe is carried by long MFP phonons and the alloying only attacks short MFP phonons, we would expect reduced T and hence reduced Umklapp scattering to further increase phonon MFPs, shifting the heat toward even longer MFPs, and hence further increasing the $k_{apparent}$ suppression of SiGe at reduced T. At very increased T, we would expect Umklapp scattering to become dominant in both classes of samples, greatly reducing phonon MFP, reducing k_{bulk}, and reducing the $k_{apparent}$ suppression.

On the theoretical side, it would be very beneficial to the field if a full closed form solution to the 3-D spectral BTE could be attained for the full TDTR geometry and temporal distributions, capable of rigorously incorporating interfaces and finite domains. This would enable the calculation of suppression functions in order to directly convert measurements of $k_{apparent}$ into k MFP accumulation functions. Full implementation of such a solution would also require knowledge of both the phonon transmission and specularity coefficients for the Al/sample interface. There has already been considerable effort toward this goal [122], [123], [130], [131], [141], and it is likely that a fully implementable solution with known coefficients may be realizable soon.

Ref. [135] and the sharp "U" shape of the Si_xGe_{1-x} k vs. x plot [100], [134] indicate a possible saturation effect of k of SiGe alloys as a function of alloying concentrations. It would therefore be interesting to carry out phonon MFP spectroscopy measurements on SiGe of varying alloying concentrations varying from very dilute (almost pure Si or Ge) to very heavily alloyed ($\geq$ 35%), to see if this resulted in a loss of influence of alloying on suppression for high alloying

concentrations. Additionally, the rate at which $k_{apparent}$ suppression (relative to a lowering k_{bulk}) changed for different heater sizes, especially combined with temperature-dependent measurement, would yield spectral information regarding how the alloying affects different phonons and progressively shifts the phonon spectrum toward longer MFPs. This will be challenging, however, as growing single crystal bulk SiGe samples with controlled and varying alloy concentrations is non-trivial.

We also believe that it would be beneficial to the field to have an independent measurement of the pump laser modulation frequency-dependence of $k_{apparent}$ suppression. At only two separate frequencies, we had insufficient data to draw any definitive conclusions regarding possible trends. Ref. [79] hypothesizes and presents data in support of a frequency-dependent $k_{apparent}$ suppression that should be stronger for SiGe alloys (i.e. phonon spectra shifted toward long MFPs) while weaker for Si. It would also be beneficial in general to repeat our measurements using a different phonon MFP spectroscopy technique to verify that the results are repeatable. We measured around a 10% - 20% stronger $k_{apparent}$ suppression in Si at 300 K as compared to other groups using different techniques, and so it would be helpful to be able to compare measurements for other samples spanning a variety of phonon spectral distributions.

Because the heat diffusion equation and the BTE under relaxation time approximation are linear differential equations, they can be described using transfer functions and Green's functions. This creates an exchange symmetry between heaters and thermometers [79], implying that swapping the heater and thermometer in a system should not change the thermal transfer function. However, as discussed in this dissertation, $k_{apparent}$ suppression requires a non-equilibrium and at least quasi-ballistic transport of the energy carriers, which in turn means significant temperature gradients on the order of dominant phonon MFPs. In the case of optical thermoreflectance thermometers such as used in TDTR, there is no heat flow into the thermometer for detection. Rather, the local temperature of the metal transducer layer can be sampled via its temperature-dependent reflectivity. In contrast, there *is* heat flux out of the heater, as the pump beam deposits energy in the transducer layer (the probe beam also deposits energy, but this is not at the modulation frequency of the lock-in and so is not picked up as a part of the final signal or represented in the temperature measurement, which is really a measure of the magnitude of temperature oscillations at the modulation frequency). Therefore, it is reasonable to expect that $k_{apparent}$ suppression effects may only be observed with shrinking heater sizes and not shrinking thermometer sizes. If the pump spot is on the order of l, then some of the radiated phonons travel ballistically away from the heated region, resulting in suppression of $k_{apparent}$. However, in the case of a very large pump spot beam, but a probe spot on the order of l, there is no reason why there should exist any more or less ballistic transport or non-equilibrium in the vicinity of the probe spot, as there is no heat transfer to or from the probe spot (at the modulation frequency) to set up ballistic transport and a non-equilibrium among energy carriers. This prediction, that suppression effects in TDTR (or FDTR) might only be observable for constrained pump spot radii and not probe radii, is in contradiction to the expectations of a linear BTE and some literature [79]. This makes it a very intriguing experiment to conduct. In our set up, our pump and probe laser spots are the same size and we lack the ability to easily vary the size of one significantly as compared to the other. Such capabilities, however, are not difficult to implement, making this somewhat of a low-hanging fruit. If proven true, this would also mean that the correct way to compare spot size-based MFP spectroscopy data among different experiments and research groups, as well as theoretical k accumulation functions, would be

based on the pump spot size, not the probe size or some combination of pump and probe spot sizes.

There has been increasing attention regarding understanding thermal interfaces using detailed phonon frameworks [73], [141]. In light of the expectation for the existence of a strong non-equilibrium among energy carriers in the immediate vicinity of an interface experiencing directional heat flow, it would be interesting to investigate whether there exists a correlation between the spectral transmissivity of an interface, the breadth of the k accumulation function of the material into which heat is flowing, and the magnitude of the thermal boundary conductance (TBC). If the heat receiving material has very long MFPs, particularly for those phonons that transmit through the interface most easily, then it means that ballistic transport will persist much farther past the interface than for a material with very short MFPs. This thermalization length may offer an additional "ballistic resistance," and effectively decrease the apparent TBC measured by an experiment based on Fourier's law. On the other hand, the same may not be true for a material in which the phonons that carry most of the heat across the interface have very short MFPs.

Dissertation Summary and Conclusions

In this thesis, I have presented new theoretical formulations, measurement techniques, and experimental data in support of measuring thermal conductivity in challenging regimes. Measurements of soft, chemically sensitive, and hydrated samples; single-micron scale measurements of samples with nanometer scale features; and the very definition and physical origins of thermal conductivity itself are all explored.

A novel reusable bi-directional 3ω sensor that is capable of accurately measuring the thermal conductivity of hydrated biological tissues down to 100 μm thicknesses is presented. This device is an electrothermal measurement technique based on simple microfabrication processes. The measurement technique is rigorously validated on control samples of water, ice, and agar gel, and then demonstrated on both frozen and non-frozen mouse liver tissues, ranging in thickness from 100 μm to 3 mm. Such a technique is capable of measuring not only biological tissues, but any potentially delicate sample ill-suited to direct microfabrication or other more invasive or destructive measurement techniques. A single sensor can be used repeatedly and rapidly to measure multiple different samples in succession.

A simple data analysis scheme for the bi-directional 3ω sensor, the generalized slope method, is presented in conjunction with the hardware design in order to make the experiment easily accessible and very computationally inexpensive. Comprehensive numerical as well as analytical studies of the full system are used to quantify and bound the error of the generalized slope method as well as all experimental parameters. Extensive error quantification maps and guides are presented to aid future researchers in optimally designing sensors and experiments. Exact analytical solutions for the full geometry are presented for situations where the generalized slope method is insufficient. This technique holds the promise for fast and easy thermal conductivity measurement of biomedically relevant tissues and similarly challenging samples.

A laser-based variable spot size TDTR technique is presented to measure phonon spectral properties of inorganic solids. Extensive documentation of the system hardware is given, along with explanations for the inclusion and selection of each component and how it improves the measurement. A robust method for consistently focusing and measuring the laser spot sizes, even in the presence of elliptical distortions and considerable chromatic aberrations, is presented and compared to traditional spot size measurement techniques. This is of particular utility to small spot size measurements where uncertainty in the laser spots can be the overwhelming dominant source of error in thermal conductivity measurements. Extensive sensitivity plots for all experimental parameters are presented and discussed, and consolidated into "quick reference" plots that can aid future researchers in optimal experimental design.

A careful explanation for the microscopic physical origins of thermal conductivity suppression at small length scales is discussed. A subtle but often confused distinction between the reduced thermal conductivity of nanostructured materials compared to an observed apparent reduction in thermal conductivity as measured by tiny heaters with characteristic lengths on the order of the phonon mean free paths is also discussed, and related back to physically intuitive principles.

Sub-continuum thermal conductivity measurements are performed on sapphire, Si, nano-grained Si, and SiGe samples. Results are compared favorably with literature measurements and theoretical predictions where available. Our results show that the thermal conductivity

accumulation function of SiGe has a "long tail," with much of the heat still being carried by phonons with mean free paths up to tens of microns even at room temperature. In contrast, at room temperature most of the heat in sapphire, Si, and nano-grained Si is found to be carried by phonons with mean free paths of a couple microns or smaller. Measurements show that nanostructuring and alloying scatter phonons so as to redistribute the heat in opposite ways. Nanostructuring primarily scatters long mean free path phonons, shifting the heat toward shorter mean free paths, while alloying does the opposite and pushes the heat toward the long mean free path end of the phonon spectrum. While both scattering mechanisms reduce thermal conductivity of bulk materials, the differences in how they distribute heat throughout the spectrum of phonons means SiGe is a worse choice than nano-grained silicon for thermal management of nano and microstructures, even when it sometimes has a higher bulk k.

The results and measurement capabilities of the bi-directional 3ω sensor can directly benefit thermal biomedical procedures such as thermal- and cryo-ablation to treat cancer and atrial fibrillation, and cryogenic preservation for organ transplants and skin grafts. The variable spot size TDTR results and phonon MFP spectrum measurement capabilities can directly benefit the heat dissipation and thermal engineering of microelectronics, batteries, LED lighting, HAMR disk drive technology, nanoparticle-based thermal cancer treatments, and advanced thermoelectric materials.

References

[1] S. D. Lubner, J. Choi, G. Wehmeyer, B. Waag, V. Mishra, H. Natesan, J. C. Bischof, and C. Dames, "Reusable bi-directional 3ω sensor to measure thermal conductivity of 100-μm thick biological tissues," *Rev. Sci. Instrum.*, vol. 86, no. 1, p. 014905, Jan. 2015.

[2] S. Rossi, M. Di Stasi, E. Buscarini, P. Quaretti, F. Garbagnati, L. Squassante, C. Paties, D. Silverman, and L. Buscarini, "Percutaneous RF interstitial thermal ablation in the treatment of hepatic cancer.," *AJR. Am. J. Roentgenol.*, vol. 167, no. 3, pp. 759–768, 1996.

[3] B. Rubinsky, "CRYOSURGERY," *Annu. Rev. Biomed. Eng.*, vol. 2, pp. 157–187, Nov. 2000.

[4] C. J. Diederich, "Thermal ablation and high-temperature thermal therapy: Overview of technology and clinical implementation," *Int. J. Hyperth.*, vol. 21, no. 8, pp. 745–753, May 2011.

[5] X. He and J. C. Bischof, "Quantification of Temperature and Injury Response in Thermal Therapy and Cryosurgery," *Crit. Rev. Biomed. Eng.*, vol. 31, no. 5–6, pp. 355–422, 2003.

[6] B. Han and J. C. Bischof, "Engineering Challenges in Tissue Preservation," *Cell Preserv. Technol.*, vol. 2, no. 2, pp. 91–112, Jul. 2004.

[7] E. van Sonnenberg, T. Livraghi, P. R. Mueller, W. McMullen, L. Solbiati, and S. G. Silverman, *Tumor Ablation: Principles and Practice*. Springer Science & Business Media, 2008.

[8] J. Choi and J. C. Bischof, "Review of biomaterial thermal property measurements in the cryogenic regime and their use for prediction of equilibrium and non-equilibrium freezing applications in cryobiology.," *Cryobiology*, vol. 60, no. 1, pp. 52–70, Feb. 2010.

[9] H. F. Bowman, "Theory, Measurement And Application Of Thermal Properties Of Biomaterials," *Annu. Rev. Biophys. Bioeng.*, vol. 4, no. 1, pp. 43–80, 1975.

[10] J. Chato, "Measurement of thermal properties of biological materials," *Heat Transf. Med. Biol.*, vol. 1, pp. 167–192, 1985.

[11] J. W. Valvano and B. Definitions, "Tissue Thermal Properties and Perfusion," *Opt. Response Laser-Irradiated Tissue*, vol. 12, pp. 455–485, 2011.

[12] J. Choi and J. C. Bischof, "A quantitative analysis of the thermal properties of porcine liver with glycerol at subzero and cryogenic temperatures.," *Cryobiology*, vol. 57, no. 2, pp. 79–83, Oct. 2008.

[13] T. A. Balasubramaniam and H. F. Bowman, "Thermal Conductivity and Thermal Diffusivity of Biomaterials: A Simultaneous Measurement Technique," *J. Biomech. Eng.*, vol. 99, no. 3, p. 148, Aug. 1977.

[14] T. E. Cooper and G. J. Trezek, "A Probe Technique for Determining the Thermal Conductivity of Tissue," *J. Heat Transfer*, vol. 94, no. 2, p. 133, May 1972.

[15] J. W. Valvano, "Low temperature tissue thermal properties," in *Low Temperature Biotechnology Emerging Applications and Engineering Contributions*, American Society of Mechanical Engineers, J.J. McGrath, K.R. Diller, 1988, pp. 331–345.

[16] H. Zhang, S. Cheng, L. He, A. Zhang, Y. Zheng, and D. Gao, "Determination of Thermal Conductivity of Biomaterials in the Temperature Range 233–313K Using a Tiny Detector Made of a Self-Heated Thermistor," *Cell Preserv. Technol.*, vol. 1, no. 2, pp. 141–147, Jun. 2002.

[17] D. G. Cahill, "Thermal conductivity measurement from 30 to 750 K: the 3ω method,"

Rev. Sci. Instrum., vol. 61, no. 2, pp. 802–808, 1990.
[18] C. Dames and G. Chen, "1ω, 2ω, and 3ω Methods for Measurements of Thermal Properties," *Rev. Sci. Instrum.*, vol. 76, no. 12, p. 124902, 2005.
[19] C. Dames, "Measuring the thermal conductivity of thin films: 3 omega and related electrothermal methods.," in *Annual Review of Heat Transfer*, 16th ed., vol. 16, no. 16, Begel House Inc., 2013, pp. 7–49.
[20] D. G. Cahill, M. Katiyar, and J. R. Abelson, "Thermal conductivity of a-Si:H thin films," *Phys. Rev. B*, vol. 50, no. 9, pp. 6077–6082, 1994.
[21] B. K. Park, J. Park, and D. Kim, "Note: Three-omega method to measure thermal properties of subnanoliter liquid samples.," *Rev. Sci. Instrum.*, vol. 81, no. 6, p. 066104, Jun. 2010.
[22] F. Chen, J. Shulman, Y. Xue, C. W. Chu, and G. S. Nolas, "Thermal conductivity measurement under hydrostatic pressure using the 3ω method," *Rev. Sci. Instrum.*, vol. 75, no. 11, p. 4578, 2004.
[23] I. Moon, Y. Jeong, and S. Kwun, "The 3ω technique for measuring dynamic specific heat and thermal conductivity of a liquid or solid," *Rev. Sci. Instrum.*, vol. 67, no. January, pp. 29–35, 1996.
[24] D.-W. Oh, A. Jain, J. K. Eaton, K. E. Goodson, and J. S. Lee, "Thermal conductivity measurement and sedimentation detection of aluminum oxide nanofluids by using the 3ω method," *Int. J. Heat Fluid Flow*, vol. 29, no. 5, pp. 1456–1461, Oct. 2008.
[25] B. K. Park, N. Yi, J. Park, and D. Kim, "Thermal conductivity of single biological cells and relation with cell viability," *Appl. Phys. Lett.*, vol. 102, no. 20, p. 203702, 2013.
[26] J.-Y. Duquesne, D. Fournier, and C. Frétigny, "Analytical solutions of the heat diffusion equation for 3ω method geometry," *J. Appl. Phys.*, vol. 108, no. 8, p. 086104, 2010.
[27] D. G. Cahill, "Analysis of heat flow in layered structures for time-domain thermoreflectance," *Rev. Sci. Instrum.*, vol. 75, no. 12, pp. 5119–5122, Dec. 2004.
[28] A. J. Schmidt, M. Chiesa, X. Chen, and G. Chen, "An optical pump-probe technique for measuring the thermal conductivity of liquids.," *Rev. Sci. Instrum.*, vol. 79, no. 6, p. 064902, Jun. 2008.
[29] P. E. Hopkins, B. Kaehr, L. M. Phinney, T. P. Koehler, A. M. Grillet, D. Dunphy, F. Garcia, and C. J. Brinker, "Measuring the Thermal Conductivity of Porous, Transparent SiO[sub 2] Films With Time Domain Thermoreflectance," *J. Heat Transfer*, vol. 133, no. 6, p. 061601, 2011.
[30] X. J. Hu, A. a. Padilla, J. Xu, T. S. Fisher, and K. E. Goodson, "3-Omega Measurements of Vertically Oriented Carbon Nanotubes on Silicon," *J. Heat Transfer*, vol. 128, no. 11, p. 1109, 2006.
[31] M. L. Bauer and P. M. Norris, "General bidirectional thermal characterization via the 3ω technique.," *Rev. Sci. Instrum.*, vol. 85, no. 6, p. 064903, Jun. 2014.
[32] J. H. Kim, A. Feldman, and D. Novotny, "Application of the three omega thermal conductivity measurement method to a film on a substrate of finite thickness," *J. Appl. Phys.*, vol. 86, no. 7, p. 3959, 1999.
[33] S. R. Choi, J. Kim, and D. Kim, "3Omega Method To Measure Thermal Properties of Electrically Conducting Small-Volume Liquid.," *Rev. Sci. Instrum.*, vol. 78, no. 8, p. 084902, Aug. 2007.
[34] A. T. Ramu and J. E. Bowers, "A '2-omega' technique for measuring anisotropy of thermal conductivity.," *Rev. Sci. Instrum.*, vol. 83, no. 12, p. 124903, Dec. 2012.

[35] A. Feldman, "Algorithm for solutions of the thermal diffusion equation in a stratified medium with a modulated heating source," *High Temp. Press.*, vol. 31, no. 3, pp. 293–298, 1999.
[36] M. L. Shendeleva, "Temperature oscillations produced by an impulsive interfacial line heat source," *Appl. Phys. Lett.*, vol. 82, no. 12, p. 1983, 2003.
[37] M. L. Shendeleva, J. a. Molloy, and N. N. Ljepojevic, "Periodic line heat source at the interface," *Rev. Sci. Instrum.*, vol. 74, no. 1, p. 427, 2003.
[38] S. R. Choi and D. Kim, "Real-time thermal characterization of 12 nl fluid samples in a microchannel.," *Rev. Sci. Instrum.*, vol. 79, no. 6, p. 064901, Jun. 2008.
[39] Z. Ge, D. G. Cahill, and P. Braun, "Thermal Conductance of Hydrophilic and Hydrophobic Interfaces," *Phys. Rev. Lett.*, vol. 96, no. 18, pp. 1–4, May 2006.
[40] *CRC Handbook of Chemistry and Physics, 90th Edition*. CRC Press, 2009.
[41] J. Yang, Y. Yang, S. W. Waltermire, X. Wu, H. Zhang, T. Gutu, Y. Jiang, Y. Chen, A. a Zinn, R. Prasher, T. T. Xu, and D. Li, "Enhanced and switchable nanoscale thermal conduction due to van der Waals interfaces.," *Nat. Nanotechnol.*, vol. 7, no. 2, pp. 91–5, Feb. 2012.
[42] V. Mishra, C. L. Hardin, J. E. Garay, and C. Dames, "A 3 omega method to measure an arbitrary anisotropic thermal conductivity tensor," *Rev. Sci. Instrum.*, vol. 86, no. 5, p. 054902, 2015.
[43] H. Natesan, W. Hodges, J. Choi, S. Lubner, C. Dames, and J. Bischof, "A Micro-Thermal Sensor for Focal Therapy Applications," *Sci. Rep.*, vol. 6, no. October 2015, p. 21395, 2016.
[44] D. Josell, E. J. Gonzalez, and G. S. White, "Correcting errors in the theory for mirage-effect measurements," *J. Mater. Res.*, vol. 13, no. 5, pp. 1117–1119, 1998.
[45] P. K. Kuo, E. D. Sendler, L. D. Favro, and R. L. Thomas, "Mirage-effect measurement of thermal diffusivity. Part II: theory," *Can. J. Phys.*, vol. 64, no. 9, pp. 1168–1171, Sep. 1986.
[46] H. F. Mark, D. F. Othmer, C. G. Overberger, and G. T. Seaboarg, *Encyclopedia of Chemical Technology*. 1989.
[47] Y. S. Ju and K. E. Goodson, "Process-dependent thermal transport properties of silicon-dioxide films deposited using low-pressure chemical vapor deposition," *J. Appl. Phys.*, vol. 85, no. 10, p. 7130, 1999.
[48] Z. Chen and C. Dames, "An anisotropic model for the minimum thermal conductivity," *Appl. Phys. Lett.*, vol. 107, no. 2015, p. 193104, 2015.
[49] S. Narayana and Y. Sato, "Heat flux manipulation with engineered thermal materials," *Phys. Rev. Lett.*, vol. 108, no. 21, pp. 1–5, 2012.
[50] K. Biswas, J. He, I. D. Blum, C.-I. Wu, T. P. Hogan, D. N. Seidman, V. P. Dravid, and M. G. Kanatzidis, "High-performance bulk thermoelectrics with all-scale hierarchical architectures.," *Nature*, vol. 489, no. 7416, pp. 414–8, 2012.
[51] 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," *Adv. Funct. Mater.*, vol. 19, no. 15, pp. 2445–2452, 2009.
[52] G. Joshi, H. Lee, Y. Lan, X. Wang, G. Zhu, D. Wang, R. W. Gould, D. C. Cuff, M. Y. Tang, M. S. Dresselhaus, G. Chen, and Z. Ren, "Enhanced thermoelectric figure-of-merit in nanostructured p-type silicon germanium bulk alloys," *Nano Lett.*, vol. 8, no. 12, pp. 4670–4674, 2008.

[53] A. J. Minnich, M. S. Dresselhaus, Z. F. Ren, and G. Chen, "Bulk nanostructured thermoelectric materials: current research and future prospects," *Energy Env. Sci*, vol. 2, no. 5, pp. 466–479, 2009.
[54] D. Li, Y. Wu, P. Kim, L. Shi, P. Yang, and A. Majumdar, "Thermal conductivity of individual silicon nanowires," *Appl. Phys. Lett.*, vol. 83, no. 14, pp. 2934–2936, 2003.
[55] A. I. Boukai, Y. Bunimovich, J. Tahir-Kheli, J.-K. Yu, W. A. Goddard, and J. R. Heath, "Silicon nanowires as efficient thermoelectric materials.," *Nature*, vol. 451, no. 7175, pp. 168–71, 2008.
[56] A. I. Hochbaum, R. Chen, R. D. Delgado, W. Liang, E. C. Garnett, M. Najarian, A. Majumdar, and P. Yang, "Enhanced thermoelectric performance of rough silicon nanowires," *Nature*, vol. 451, no. 7175, pp. 163–167, 2008.
[57] B. Piccione and D. S. Gianola, "Tunable thermoelectric transport in nanomeshes via elastic strain engineering," *Appl. Phys. Lett.*, vol. 106, p. 113101, 2015.
[58] E. Pop, "Energy dissipation and transport in nanoscale devices," *Nano Res.*, vol. 3, no. 3, pp. 147–169, 2010.
[59] W. a. Challener, C. Peng, a. V. Itagi, D. Karns, W. Peng, Y. Peng, X. Yang, X. Zhu, N. J. Gokemeijer, Y.-T. Hsia, G. Ju, R. E. Rottmayer, M. a. Seigler, and E. C. Gage, "Heat-assisted magnetic recording by a near-field transducer with efficient optical energy transfer," *Nat. Photonics*, vol. 3, no. 5, pp. 303–303, 2009.
[60] A. T. Wieg, Y. Kodera, Z. Wang, T. Imai, C. Dames, and J. E. Garay, "Visible photoluminescence in polycrystalline terbium doped aluminum nitride (Tb:AlN) ceramics with high thermal conductivity," *Appl. Phys. Lett.*, vol. 101, no. 11, pp. 1–6, 2012.
[61] A. T. Wieg, Y. Kodera, Z. Wang, C. Dames, and J. E. Garay, "Thermomechanical properties of rare-earth-doped AlN for laser gain media: The role of grain boundaries and grain size," *Acta Mater.*, vol. 86, pp. 148–156, 2015.
[62] N. Narendran and Y. Gu, "Life of LED-based white light sources," *IEEE/OSA J. Disp. Technol.*, vol. 1, no. 1, pp. 167–170, 2005.
[63] A. A. Efremov, N. I. Bochkareva, R. I. Gorbunov, D. A. Lavrinovich, Y. T. Rebane, D. V Tarkhin, and Y. G. Shreter, "Effect of the joule heating on the quantum efficiency and choice of thermal conditions for high-power blue InGaN/GaN LEDs," *Semiconductors*, vol. 40, no. 5, pp. 605–610, 2006.
[64] *Adoption of Light-Emitting Diodes in Common Lighting Applications*. U.S. Department of Energy, 2015.
[65] *Estimated U.S. Energy Use in 2014*. Lawrence Livermore National Lab, 2015.
[66] T. M. Bandhauer, S. Garimella, and T. F. Fuller, "A Critical Review of Thermal Issues in Lithium-Ion Batteries," *J. Electrochem. Soc.*, vol. 158, no. 3, p. R1, 2011.
[67] S. J. Drake, M. Martin, D. A. Wetz, J. K. Ostanek, S. P. Miller, J. M. Heinzel, and A. Jain, "Heat generation rate measurement in a Li-ion cell at large C-rates through temperature and heat flux measurements," *J. Power Sources*, vol. 285, pp. 266–273, 2015.
[68] D. Chen, J. Jiang, G. H. Kim, C. Yang, and A. Pesaran, "Comparison of different cooling methods for lithium ion battery cells," *Appl. Therm. Eng.*, vol. 94, pp. 846–854, 2016.
[69] K. Shah, S. J. Drake, D. A. Wetz, J. K. Ostanek, S. P. Miller, J. M. Heinzel, and A. Jain, "Modeling of steady-state convective cooling of cylindrical Li-ion cells," *J. Power Sources*, vol. 258, pp. 374–381, 2014.
[70] V. Vishwakarma, C. Waghela, Z. Wei, R. Prasher, S. C. Nagpure, J. Li, F. Liu, C. Daniel, and A. Jain, "Heat transfer enhancement in a lithium-ion cell through improved material-

level thermal transport," *J. Power Sources*, vol. 300, pp. 123–131, 2015.
[71] D. P. O'Neal, L. R. Hirsch, N. J. Halas, J. D. Payne, and J. L. West, "Photo-thermal tumor ablation in mice using near infrared-absorbing nanoparticles," *Cancer Lett.*, vol. 209, no. 2, pp. 171–176, 2004.
[72] A. J. Minnich, "Advances in the measurement and computation of thermal phonon transport properties," *J. Phys. Condens. Matter*, vol. 27, p. 053202, 2015.
[73] C. Monachon, L. Weber, and C. Dames, "Thermal Boundary Conductance□:A Materials Science Perspective," *Annu. Rev. Mater. Res.*, vol. accepted, no. March, 2016.
[74] M. Maldovan, "Sound and heat revolutions in phononics.," *Nature*, vol. 503, no. 7475, pp. 209–17, 2013.
[75] E. S. Toberer, L. L. Baranowski, and C. Dames, "Advances in Thermal Conductivity," *Annu. Rev. Mater. Res.*, vol. 42, no. 1, pp. 179–209, 2012.
[76] C. Dames and G. Chen, "Thermal Conductivity of Nanostructured Thermoelectric Materials," in *Thermoelectrics Handbook Macro to Nano*, vol. 80, no. 10, D. M. Rowe, Ed. CRC Press, 2005, p. 1014.
[77] D. G. Cahill, W. K. Ford, K. E. Goodson, G. D. Mahan, A. Majumdar, H. J. Maris, R. Merlin, and S. R. Phillpot, "Nanoscale thermal transport," *J. Appl. Phys.*, vol. 93, no. 2, pp. 793–818, 2003.
[78] L. Shi, C. Dames, J. R. Lukes, P. Reddy, J. Duda, D. G. Cahill, J. Lee, A. Marconnet, K. E. Goodson, J.-H. Bahk, A. Shakouri, R. S. Prasher, J. Felts, W. P. King, B. Han, and J. C. Bischof, "Evaluating Broader Impacts of Nanoscale Thermal Transport Research," *Nanoscale Microscale Thermophys. Eng.*, vol. 19, no. 2, pp. 127–165, 2015.
[79] R. B. Wilson and D. G. Cahill, "Anisotropic failure of Fourier theory in time-domain thermoreflectance experiments," *Nat. Commun.*, vol. 5, p. 5075, Oct. 2014.
[80] K. T. Regner, D. P. Sellan, Z. Su, C. H. Amon, A. J. H. McGaughey, and J. A. Malen, "Broadband phonon mean free path contributions to thermal conductivity measured using frequency domain thermoreflectance.," *Nat. Commun.*, vol. 4, p. 1640, Jan. 2013.
[81] Y. Koh and D. G. Cahill, "Frequency dependence of the thermal conductivity of semiconductor alloys," *Phys. Rev. B*, vol. 76, no. 7, p. 075207, Aug. 2007.
[82] K. T. Regner, A. J. H. McGaughey, and J. A. Malen, "Analytical interpretation of nondiffusive phonon transport in thermoreflectance thermal conductivity measurements," *Phys. Rev. B*, vol. 90, p. 064302, 2014.
[83] A. J. Schmidt, X. Chen, and G. Chen, "Pulse accumulation, radial heat conduction, and anisotropic thermal conductivity in pump-probe transient thermoreflectance," *Rev. Sci. Instrum.*, vol. 79, no. 11, pp. 1–9, 2008.
[84] J. Yang, E. Ziade, and A. J. Schmidt, "Uncertainty analysis of thermoreflectance measurements," *Rev. Sci. Instrum.*, vol. 014901, 2016.
[85] R. Costescu, M. Wall, and D. Cahill, "Thermal conductance of epitaxial interfaces," *Phys. Rev. B*, vol. 67, no. 5, p. 54302, 2003.
[86] R. J. Stevens, A. N. Smith, and P. M. Norris, "Signal analysis and characterization of experimental setup for the transient thermoreflectance technique," *Rev. Sci. Instrum.*, vol. 77, no. 8, 2006.
[87] K. Kang, Y. K. Koh, C. Chiritescu, X. Zheng, and D. G. Cahill, "Two-tint pump-probe measurements using a femtosecond laser oscillator and sharp-edged optical filters," *Rev. Sci. Instrum.*, vol. 79, no. 11, 2008.
[88] "About Lock-In Amplifiers," *Stanford Research Systems*. [Online]. Available:

http://www.thinksrs.com/downloads/PDFs/ApplicationNotes/AboutLIAs.pdf.
[89] D. G. Cahill and F. Watanabe, "Thermal conductivity of isotopically pure and Ge-doped Si epitaxial layers from 300 to 550 K," *Phys. Rev. B - Condens. Matter Mater. Phys.*, vol. 70, no. 23, pp. 1–3, 2004.
[90] E. Hecht, *Optics*. New York: Addison Wesley Publishing Company, 1997.
[91] W. S. Capinski and H. J. Maris, "Improved apparatus for picosecond pump-and-probe optical measurements," *Rev. Sci. Instrum.*, vol. 67, no. 8, p. 2720, 1996.
[92] K. C. Collins, A. a. Maznev, J. Cuffe, K. a. Nelson, and G. Chen, "Examining thermal transport through a frequency-domain representation of time-domain thermoreflectance data," *Rev. Sci. Instrum.*, vol. 85, p. 124903, 2014.
[93] A. J. Schmidt, R. Cheaito, and M. Chiesa, "A frequency-domain thermoreflectance method for the characterization of thermal properties," *Rev. Sci. Instrum.*, vol. 80, no. 9, pp. 0–6, 2009.
[94] G. Tas and H. J. Maris, "Electron diffusion in metals studied by picosecond ultrasonics," *Phys. Rev. B*, vol. 49, no. 21, pp. 15046–15054, 1994.
[95] T. J. Campbell, G. Aral, S. Ogata, R. K. Kalia, A. Nakano, and P. Vashishta, "Oxidation of aluminum nanoclusters," *Phys. Rev. B - Condens. Matter Mater. Phys.*, vol. 71, no. 20, pp. 1–14, 2005.
[96] P. Jiang, L. Lindsay, and Y. K. Koh, "The role of low-energy phonons with mean-free-paths >0.8 um in heat conduction in silicon," *J. Appl. Phys.*, vol. 119, p. 245705, 2016.
[97] N. Ashcroft and D. Mermin, *Solid State Physics*, 1st ed. New York: Harcourt College, 1976.
[98] A. J. Schmidt, R. Cheaito, and M. Chiesa, "Characterization of thin metal films via frequency-domain thermoreflectance," *J. Appl. Phys.*, vol. 107, no. 2, 2010.
[99] Z. Wang, J. E. Alaniz, W. Jang, J. E. Garay, and C. Dames, "Thermal conductivity of nanocrystalline silicon: importance of grain size and frequency-dependent mean free paths," *Nano Lett.*, vol. 11, no. 6, pp. 1–8, Jun. 2011.
[100] B. Abeles, "Lattice thermal conductivity of disordered semiconductor alloys at high temperatures," *Phys. Rev.*, vol. 131, no. 5, pp. 1906–1911, 1963.
[101] H. S. Carslaw and J. C. Jaeger, *Conduction of Heat in Solids*, 2nd ed. New York: Oxford University Press, 1959.
[102] J. Durbin and G. S. Watson, "Testing for serial correlation in least squares regression I," *Biometrika*, vol. 37, no. 3–4, pp. 409–428, 1950.
[103] R. B. Wilson, B. a Apgar, L. W. Martin, and D. G. Cahill, "Thermoreflectance of metal transducers for optical pump-probe studies of thermal properties.," *Opt. Express*, vol. 20, no. 27, pp. 28829–38, 2012.
[104] B. C. Gundrum, D. G. Cahill, and R. S. Averback, "Thermal conductance of metal-metal interfaces," *Phys. Rev. B - Condens. Matter Mater. Phys.*, vol. 72, no. 24, pp. 1–5, 2005.
[105] G. Chen, *Nanoscale Energy Transport and Conversion*. New York: Oxford University Press, 2005.
[106] T. L. Bergman, A. S. Lavine, F. P. Incropera, and D. P. Dewitt, *Fundamentals of Heat and Mass Transfer*, 7th ed. John Wiley & Sons, 2011.
[107] R. Prasher, T. Tong, and A. Majumdar, "Approximate analytical models for phonon specific heat and ballistic thermal conductance of nanowires," *Nano Lett.*, vol. 8, no. 1, pp. 99–103, 2008.
[108] N. Mingo and L. Yang, "Phonon transport in nanowires coated with an amorphous

material: An atomistic Green's function approach," *Phys. Rev. B*, vol. 68, no. 24, 2003.
[109] M. I. Flik, B. I. Choi, and K. E. Goodson, "Heat Transfer Regimes in Microstructures," *J. Heat Transfer*, vol. 114, no. 3, p. 666, 1992.
[110] W. M. Rohsenow and H. Y. Choi, *Heat, mass, and momentum transfer*. Englewood Cliffs, NJ: Prentice-Hall, 1961.
[111] C. Bera, N. Mingo, and S. Volz, "Marked effects of alloying on the thermal conductivity of nanoporous materials," *Phys. Rev. Lett.*, vol. 104, no. 11, pp. 1–4, 2010.
[112] K. Hippalgaonkar, B. Huang, R. Chen, K. Sawyer, P. Ercius, and A. Majumdar, "Fabrication of microdevices with integrated nanowires for investigating low-dimensional phonon transport," *Nano Lett.*, vol. 10, no. 11, pp. 4341–4348, 2010.
[113] W. Li, N. Mingo, L. Lindsay, D. A. Broido, D. A. Stewart, and N. A. Katcho, "Thermal conductivity of diamond nanowires from first principles," *Phys. Rev. B - Condens. Matter Mater. Phys.*, vol. 85, no. 19, pp. 1–5, 2012.
[114] Y. S. Ju and K. E. Goodson, "Phonon scattering in silicon films with thickness of order 100 nm," *Appl. Phys. Lett.*, vol. 74, no. 20, p. 3005, 1999.
[115] A. Majumdar, "Microscale Heat Conduction in Dielectric Thin Films," *J. Heat Transfer*, vol. 115, no. February 1993, pp. 7–16, 1993.
[116] Q. Fu, J. Yang, Y. Chen, D. Li, and D. Xu, "Experimental evidence of very long intrinsic phonon mean free path along the c-axis of graphite," *Appl. Phys. Lett.*, vol. 106, p. 031905, 2015.
[117] G. Chen, "Nonlocal and Nonequilibrium Heat Conduction in the Vicinity of Nanoparticles," *J. Heat Transfer*, vol. 118, no. August 1996, 1996.
[118] R. Siegel and J. R. Howell, *Thermal Radiation Heat Transfer*, 3rd ed. Taylor & Francis, 1993.
[119] C. Kittel, *Quantum Theory Of Solids*, vol. 1. New York: John Wiley & Sons, 1987.
[120] T. Klitsner, J. E. Vancleve, H. E. Fischer, and R. O. Pohl, "Phonon radiative heat transfer and surface scattering," *Phys. Rev. B*, vol. 38, no. 11, pp. 7576–7594, 1988.
[121] A. A. Joshi and A. Majumdar, "Transient ballistic and diffusive phonon heat transport in thin films," *J. Appl. Phys.*, vol. 74, no. 1, p. 31, 1993.
[122] C. Hua and A. J. Minnich, "Analytical Green's function of the multidimensional frequency-dependent phonon Boltzmann equation," *Phys. Rev. B*, vol. 90, p. 214306, 2014.
[123] C. Hua and A. J. Minnich, "Transport regimes in quasiballistic heat conduction," *Phys. Rev. B*, vol. 89, no. 9, p. 094302, Mar. 2014.
[124] J. D. Chung, a. J. H. McGaughey, and M. Kaviany, "Role of Phonon Dispersion in Lattice Thermal Conductivity Modeling," *J. Heat Transfer*, vol. 126, no. 3, p. 376, 2004.
[125] F. Yang and C. Dames, "Mean free path spectra as a tool to understand thermal conductivity in bulk and nanostructures," *Phys. Rev. B*, vol. 87, no. 3, p. 035437, Jan. 2013.
[126] 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," *Phys. Rev. Lett.*, vol. 107, no. 9, p. 095901, Aug. 2011.
[127] K. Esfarjani, G. Chen, and H. T. Stokes, "Heat transport in silicon from first-principles calculations," *Phys. Rev. B - Condens. Matter Mater. Phys.*, vol. 84, no. 8, pp. 1–11, 2011.
[128] F. Yang and C. Dames, "Heating-frequency-dependent thermal conductivity: An analytical solution from diffusive to ballistic regime and its relevance to phonon scattering

measurements," *Phys. Rev. B - Condens. Matter Mater. Phys.*, vol. 91, no. 16, pp. 1–14, 2015.
[129] A. J. Minnich, "Determining Phonon Mean Free Paths from Observations of Quasiballistic Thermal Transport," *Phys. Rev. Lett.*, vol. 109, no. 20, p. 205901, Nov. 2012.
[130] A. J. Minnich, "Phonon heat conduction in layered anisotropic crystals," *Phys. Rev. B*, vol. 91, no. February, p. 085206, 2015.
[131] C. Hua and A. J. Minnich, "Semi-analytical solution to the frequency-dependent Boltzmann transport equation for cross-plane heat conduction in thin films," *J. Appl. Phys.*, vol. 117, no. 17, p. 175306, 2015.
[132] D. Ding, X. Chen, and A. J. Minnich, "Radial quasiballistic transport in time-domain thermoreflectance studied using Monte Carlo simulations," *Appl. Phys. Lett.*, vol. 104, no. 14, p. 143104, Apr. 2014.
[133] M. E. Siemens, Q. Li, R. Yang, K. A. Nelson, E. H. Anderson, M. Murnane, and H. C. Kapteyn, "Quasi-ballistic thermal transport from nanoscale interfaces observed using ultrafast coherent soft X-ray beams.," *Nat. Mater.*, vol. 9, no. 1, pp. 26–30, Jan. 2010.
[134] J. Garg, N. Bonini, B. Kozinsky, and N. Marzari, "Role of disorder and anharmonicity in the thermal conductivity of silicon-germanium alloys: A first-principles study," *Phys. Rev. Lett.*, vol. 106, no. 4, pp. 1–4, 2011.
[135] T.-K. Hsiao, H.-K. Chang, S.-C. Liou, M.-W. Chu, S.-C. Lee, and C.-W. Chang, "Observation of room-temperature ballistic thermal conduction persisting over 8.3 μm in SiGe nanowires.," *Nat. Nanotechnol.*, vol. 8, no. 7, pp. 534–8, 2013.
[136] Y. Ezzahri, G. Zeng, K. Fukutani, Z. Bian, and A. Shakouri, "A comparison of thin film microrefrigerators based on Si/SiGe superlattice and bulk SiGe," *Microelectronics J.*, vol. 39, no. 7, pp. 981–991, 2008.
[137] P. G. Klemens, "The Scattering of Low-Frequency Lattice Waves by Static Imperfections," *Proc. Phys. Soc. Sect. A*, vol. 68, no. 12, pp. 1113–1128, 1955.
[138] P. G. Klemens, "Thermal Conductivity and Lattice Vibrational Modes," *Solid State Phys. - Adv. Res. Appl.*, vol. 7, no. C, pp. 1–98, 1958.
[139] G. A. Slack, "The Thermal Conductivity of Nonmetallic Crystals," *Solid State Phys. - Adv. Res. Appl.*, vol. 34, no. C, pp. 1–71, 1979.
[140] C. Kittel, *Introduction to Solid State Physics*, 8th ed. John Wiley & Sons, 2004.
[141] C. Hua, X. Chen, N. K. Ravichandran, and A. J. Minnich, "Fresnel transmission coefficients for thermal phonons at solid interfaces," *arXiv*, pp. 1–24, 2015.