SAND REPORT

SAND2004-0531
Unlimited Release
Printed February 2004

Multiscale Thermal Transport

Edward S. Piekos, Samuel Graham, Jr., and C. Channy Wong

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94-AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone:	(865) 576-8401
Facsimile:	(865) 576-5728
E-Mail:	reports@adonis.osti.gov
Online ordering:	http://www.doe.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone:	(800) 553-6847
Facsimile:	(703) 605-6900
E-Mail:	orders@ntis.fedworld.gov
Online ordering:	http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2004-0531
Unlimited Release
Printed February 2004

Multiscale Thermal Transport

Edward S. Piekos
Microscale Science and Technology

Samuel Graham, Jr.*
Microsystems and Materials Mechanics

C. Channy Wong
Microscale Science and Technology

Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0826

Abstract

A concurrent computational and experimental investigation of thermal transport is performed with the goal of improving understanding of, and predictive capability for, thermal transport in microdevices. The computational component involves Monte Carlo simulation of phonon transport. In these simulations, all acoustic modes are included and their properties are drawn from a realistic dispersion relation. Phonon-phonon and phonon-boundary scattering events are treated independently. A new set of phonon-phonon scattering coefficients are proposed that reflect the elimination of assumptions present in earlier analytical work from the simulation. The experimental component involves steady-state measurement of thermal conductivity on silicon films as thin as 340nm at a range of temperatures. Agreement between the experiment and simulation on single-crystal silicon thin films is excellent. Agreement for polycrystalline films is promising, but significant work remains to be done before predictions can be made confidently. Knowledge gained from these efforts was used to construct improved semiclassical models with the goal of representing microscale effects in existing macroscale codes in a computationally efficient manner.

*now at Georgia Institute of Technology: sam.graham@me.gatech.edu

Acknowledgments

The authors would like to thank the staff at the Stanford center for integrated systems (CIS), Scott Huxtable for help with the 3-ω method, and Ken Goodson and his research group at Stanford for providing data, facilities, and productive discussion. We are also grateful to Michael Gallis and Wahid Hermina for many useful suggestions for the Monte Carlo simulations. Finally, this document has benefited from technical reviews by Michael Gallis, Leslie Phinney, and Steve Kempka.

Contents

1 Introduction **13**
1.1 Theory 13
1.2 Approach 14

2 Thermal Conductivity Measurements **17**
2.1 Available Techniques 17
2.2 Measurements 21
2.3 Results 24

3 Material Model **31**
3.1 Equilibrium Distribution 31
3.2 Dispersion Relation 31
3.3 Phonon Group Speed 34

4 Monte Carlo Simulation **37**
4.1 Baseline Code Description 38
4.2 Typical Simulation Results 57
4.3 Thermal Conductivity Determination 60
4.4 Operational Considerations 60
4.5 Bulk Thermal Conductivity 62
4.6 Thin Film Thermal Conductivity 65
4.7 Grain Boundary Effects 66
4.8 Non-BGK Interphonon Scattering 70

5 Semiclassical Models **79**
5.1 Modeling Strategy 80
5.2 Baseline Model 81
5.3 Boundary Scattering 82
5.4 Grain Boundary Scattering 87

6 Conclusions and Recommendations **91**
6.1 Experiments 91
6.2 Monte Carlo Simulation 92
6.3 Semiclassical Methods 93

Appendix

A Acceptance-Rejection Method **95**

Figures

1 A typical 3-ω surface test device patterned onto a supported thin-film sample 17

2 Estimate showing the ratio of the heat lost by radiation versus conduction through the sample at a heating frequency of 10Hz. 19

3 Diagram showing the steady state electrical resistance thermometry sample with a thin suspended membrane. 20

4 Diagram showing the steady-state thermal conductivity sample used for in-plane thermal conductivity measurements on highly conductive layers. 21

5 TEM image showing the grains in the as-deposited polysilicon. 22

6 Pictures of key components in measurements. 23

7 An SEM photograph showing the thickness of the various layers in the SOI wafer structure used for thermal conductivity testing. 23

8 Thermal conductivity of single crystal Si as a function of temperature. 24

9 Comparison of data from experiments at Sandia and Stanford. 25

10 Data showing a comparison of the thermal conductivity of doped and undoped single crystal silicon. 26

11 Data showing the thermal conductivity of as-deposited and recrystallized undoped polysilicon. 27

12 TEM images of the microstructure of the undoped polysilicon. 28

13 Effect of an 8 hr, 700°C anneal on recrystallized polycrystalline silicon. . . . 29

14 Comparison of computed and measured dispersion relations. 33

15 Comparison of computed and measured specific heats 34

16 Relative probability of TA and LA phonons as a function of wavenumber. . 35

17 Phonon group speed under dispersion model. 36

18 Series test for a linear congruential generator used in previous hypersonic DSMC work. 39

19 Series test for modified ran3() subtractive random number generator from Press *et al.*. 40

20 Energy per unit volume of an equilibrium phonon distribution at a given temperature. 43

21 Histograms of incoming and outgoing particle characteristics at flux boundaries with domain and boundary initialized to the same temperature. 47

22 Consequences of violating detailed balance 50

23 Holland result for the thermal conductivity of a 0.7 *cm* silicon sample 52

24 Comparison of Holland's phonon speeds to those from the current dispersion model. 53

25 Comparison of phase and group speeds under the current dispersion model. 53

26 Modified Holland model result for thermal conductivity of a 0.7 *cm* silicon sample. 54

27 Relaxation times for both the original and the modified Holland model. . . . 55

28 Typical simulation results. 58

29 Illustration of the Poiseuille flow analogy. 59

30 Mean free path for single crystal silicon under the relaxation time model . . 61

31 Domain length effect on thermal conductivity results 63

32 Computed thermal conductivity of bulk silicon 63

33 Computed thermal conductivity of bulk silicon using Holland Model 64

34 Computed thermal conductivity in a 1 μm film. 65

35 Thermal conductivity of 1 μm and 340 *nm* films. 66

36 Simulation results for polysilicon films with varying grain size. 67

37 Comparison of measured grain size distribution with lognormal pdf. 68

38 Comparison of polysilicon thermal conductivities obtained using lognormally-distributed and constant grain sizes. 69

39 Comparison of polysilicon thermal conductivities obtained using physical lognormal and relaxation time-based fixed mean grain scattering. 70

40 The two classes of three-phonon processes. 72

41 Comparison of population histograms before and after a single collision step. 77

42 Predictions of single mode and multimode models compared with experimental data for bulk silicon. 82

43 Comparison of the predicted thermal conductivity of a 1.5 μm single crystal silicon film to measured values. 84

44 Comparison of the predicted thermal conductivity of a 1 μm single crystal silicon film to measured values. 84

45 Comparison of the predicted thermal conductivity of a 0.35 μm single crystal silicon film to measured values. 85

46 Contribution of phonon propagation modes to thermal conductivity of 1.5 μm single crystal silicon films. 86

47 Contribution of phonon propagation modes to thermal conductivity of 1 μm single crystal silicon films. 86

48 Contribution of phonon propagation modes to thermal conductivity of 0.35 μm single crystal silicon films. 87

49 Graph showing the fit of Holland's model to the experimental data for polycrystalline silicon. 89

50 Comparison of the semi-empirical prediction of room temperature thermal conductivity of polysilicon to experimental data. 89

Nomenclature

a	lattice constant, m	5.43 $\AA$ for silicon
b	heater half width, m	
$\mathcal{B}$	Boltzmann constant, J/K	$1.381 \times 10^{-23}\ J/K$
C	specific heat per unit volume, $J/m^3 \cdot K$	
d	grain diameter, m	
$\hbar$	Planck constant / 2π, $J \cdot s$	$1.0546 \times 10^{-34}\ J \cdot s$
k^*	wavenumber, m^{-1}	
k	reduced wavenumber	$k^* a/2\pi$
L	test section gage length, m	
LA	longitudinal acoustic phonon	
L_h	thermal healing length, m	
M	atomic mass, $a.m.u.$	
n	phonon count	
p	phonon polarization	
P	heater power, W	
$\wp$	probability	$0 \leq \wp \leq 1$
R	line resistance	
R_f	random fraction	$0 < R_f < 1$

t	thickness, m	
T	temperature, K	
TA	transverse acoustic phonon	
v_g	phonon group speed, m/s	
x, y	position in domain, m	
Θ	Debye temperature, K	
κ	thermal conductivity, $W/m \cdot K$	
λ	mean free path, m	
σ	Stefan-Boltzmann constant	$5.670 \times 10^{-8}\ W/m^2 \cdot K^4$
τ	relaxation time, s	
ω	angular frequency, rad/s	

1 Introduction

The process of packing increasing functionality into decreasing volume has marched forward with remarkable persistence, particularly in microelectronics and, more recently, microelectromechanical systems (MEMS) and micro-opto-electromechanical systems (MOEMS). In many devices, this trend implies increased heat generation coupled with a smaller area over which to dissipate it. Higher operating temperatures, however, generally cause degraded performance and shortened life. For example, the reliability of transistors is exponentially dependent on the temperature of the junction. Consequently, an increase in the operating temperature of 10-15°C can cut the operating life in half [1]. This is also true of semiconductor lasers, which show a similarly strong temperature dependence in both their lifetime and their output power. As a result, several research efforts have been launched in the area of thermal management, such as DARPA's HERETIC program (FY01: $9M).

Compounding this problem, several experiments have shown that the thermal conductivity of solid materials decreases as their characteristic lengths become very small or characteristic times become very short. For pure single-crystal silicon at room temperature, these effects are observable for (roughly) sub-micron length scales [2] and sub-picosecond time scales, both of which have been reached by a variety of devices already in production. For example, measured thermal conductivities have been shown to depart from predicted values by more than an order of magnitude for superlattices used in vertical-cavity surface-emitting lasers (VCSELs) [3].

From an operational standpoint, this reduction in thermal conductivity makes it more difficult to dissipate self-generated heating. From a design standpoint, it will cause rapidly increasing errors in predictions from macroscale design tools because they are currently unable to represent the governing microscale phenomena. Furthering the modeling difficulty, high aspect ratio features, such as thin films, can exhibit anisotropic thermal resistance. This anisotropy is also a direct consequence of microscale effects and is therefore absent in the traditional Fourier models. These models can therefore yield incorrect results for the heat flow pattern as well as its magnitude.

1.1 Theory

Thermal transport at the microscale is best viewed in terms of the behavior of heat carriers, which are predominantly electrons in conductors and phonons (lattice vibrations) in semiconductors and insulators. In bulk materials at temperatures greater than roughly 20K,

these carriers scatter from each other sufficiently often to cause heat flow to function as a diffusive process. Thermal transport is thus dominated by properties intrinsic to the material. If these properties are measured (generally as a function of temperature) they are considered valid for any geometry.

As the characteristic length (or the temperature) decreases to very small values, carriers become much less likely to scatter from each other than from geometrically-fixed scatterers, such as material boundaries, grain boundaries, defects, or impurities. The location of these scatterers with respect to the heat flow can therefore dominate the observed behavior and thermal transport properties are no longer independent of geometry. For example, in the case of a very thin pure single-crystal material, the heat carriers have a high likelihood of ballistic travel across the film and a low likelihood of ballistic travel along the film. The thermal resistance across the film thus tends to zero and the thermal resistance along the film thus tends to infinity. The thermal conductivity gleaned from a material handbook will therefore be either too high or too low, depending on the heat flow path.

In order to predict these effects, the behavior of individual carriers can no longer be eliminated from the model by assuming a local thermodynamic equilibrium. Their travel and their interactions with each other and with geometric features must be explicitly represented. For domains sufficiently larger than the spatial extent of phonon wavepackets, which generally span a few angstrom-scale unit cells, wave effects such as interference can be neglected [4]. The phonons can then be modeled as a collection of particles (the "phonon gas" assumption). In this case, the Boltzmann transport equation (BTE) provides a good model because it includes the effects of ballistic flight as well as scattering. This equation has seen extensive use in transport modeling, particularly in gases.

Unfortunately, this seven-dimensional integro-differential equation has proven difficult to solve for realistic problems. A number of possibilities have been proposed, ranging from analytic to fully numeric, but all require simplifications to make solution possible. Choosing a technique for a particular problem therefore requires a compromise between accuracy, generality, and computational effort.

1.2 Approach

In this work, a concurrent computational and experimental investigation of microscale thermal transport in semiconductors is undertaken. By using information gleaned from each technique to guide the other, an increased understanding of the governing phenomena can be obtained through a systematic manipulation of subject conditions. In addition, throughout this work, a goal of providing information immediately useful to device design-

ers is maintained. Understanding gleaned from the experiments and microscale simulations is therefore used to develop simplified models that can be incorporated into existing macroscale codes. Because most current, and many proposed, devices fall into a regime where microscale effects are significant but not dominant and where microscale codes are impractical for modeling a complete system, such subgrid models should provide microdevice designers with a valuable predictive capability without a significant increase in burden on users or computers.

2 Thermal Conductivity Measurements

2.1 Available Techniques

The measurement of thermal conductivity in thin-films has relied upon steady-state electrical resistance thermometry [5, 6, 7, 8, 9, 10], harmonic joule heating (3-ω) [11, 12], phase sensitive detection [13], and transient optical methods [14, 15]. Periodic heating methods, such as 3-ω, have become very popular for thin-films due to their ease of implementation and data reduction. In general, this method utilizes periodic surface temperature measurements, in conjunction with an approximate analytical solution, to determine a material's thermal conductivity. The method exploits the interaction between electrical resistance fluctuations, due to the surface temperature changes, and the electrical current which is being used to drive the heating. This interaction produces a small, but measurable, signal at a frequency which is three times the driving frequency (ω), hence the name 3-ω. One great advantage of this approach is that the thermal disturbance and temperature measurement are both implemented using the same surface structure. 3-ω can also be used to determine the conductivity of thin film layers which are situated on the surface of a substrate material.

A typical 3-ω surface structure is shown in Fig. 1. This surface device is constructed of a highly conductive metallic film, which is deposited and then patterned. Obviously, the layer beneath the structure must be a dielectric. It is desirable that the metal have an electrical resistance which is strongly sensitive to temperature. The contact pads provide connections to external electronics. The principal electronic device that makes this method possible is the lock-in-amplifier, which allows one to selectively screen out noise signals which are several orders of magnitude larger than the 3-ω signal.

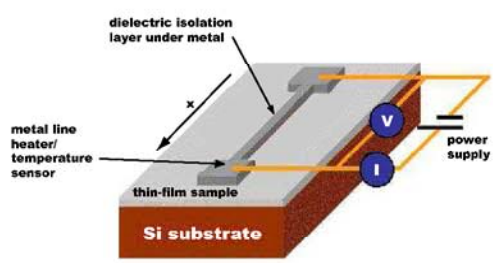

Figure 1. A typical 3-ω surface test device patterned onto a supported thin-film sample

In measuring the thermal conductivity of thin-films, the 3-ω method assumes that film being measured is:

1. of negligible thermal capacity
2. has a much lower thermal conductivity than the substrate beneath (no thermal spreading in the film).

This allows one to assume that the film adds a constant amplitude shift in the 3-ω voltage as function of frequency. In this method, the thermal conductivities of the substrate and film are given by:

$$\kappa_{subs} = \frac{V_{rms}^3}{8R^2L}\frac{\partial R}{\partial T}\frac{1}{(V_{3\omega_{rms}})_{90^\circ}} \tag{1}$$

$$\kappa_{film} = \frac{t_{film}V_{rms}^3}{4bR^2L}\frac{\partial R}{\partial T}\left[\frac{1}{(\Delta V_{3\omega_{rms}})_{film} - (\Delta V_{3\omega_{rms}})_{subs}}\right] \tag{2}$$

where V_{rms} is the applied voltage amplitude, R is the line resistance, L is the test section gage length, $\partial R/\partial T$ is the slope of the resistivity as a function temperature, b is the half width of the heater, t_{film} is the film thickness, and $V_{3\omega_{rms}}$ is the 3-ω voltage amplitude.

Typically, experiments are performed in a thermally controlled vacuum environment. However, the 3-ω method can be performed in standard laboratory conditions with the addition of about 5% error to the measurement results. The method is also resilient in that it is very insensitive to radiation heat loss. This is due to the fact that the heater/temperature sensor in the sample is typically on the order of 20 μm or less in width and 1 mm in length. The overall surface area of the heated region is quite small and emits very little energy. Since the amplitude of the thermal wave decreases with increasing frequency, errors also decrease as higher heating frequencies are used. Figure 2 displays an estimate of the ratio of heat loss by radiation to that by conduction in a typical 3-ω sample operating at 10 Hz. This estimate shows that less than 0.1% of the heat is transferred by radiation for temperatures up to 1000K, assuming a conservative emissivity value of 1 across the temperature regime. Thus, the 3-ω method is very attractive for thermal conductivity measurements in thin-films and requires very little sample preparation.

One drawback in using the 3-ω method to measure thin-film thermal conductivity is that it is difficult to deduce anisotropic properties from this measurement with high accuracy. For this reason, methods such as steady-state and phase shift periodic heating methods

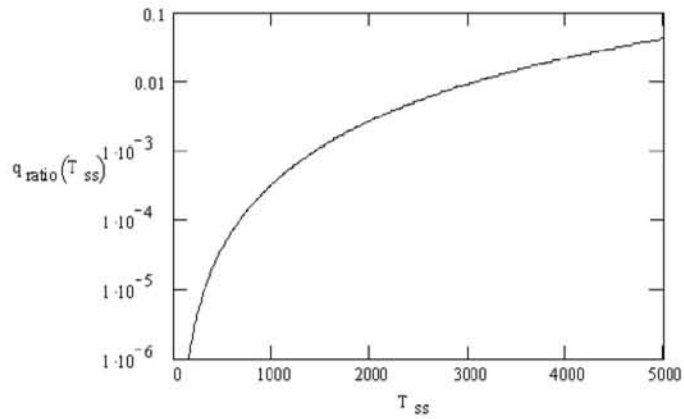

Figure 2. Estimate showing the ratio of the heat lost by radiation versus conduction through the sample at a heating frequency of 10Hz. The estimate show that the radiation contribution to heat loss is less than 0.1% for temperatures up to 1000K.

have been employed to measure in-plane thermal conductivity of thin-films. In general, these methods provide very high accuracy of in-plane thermal conductivity at the price of increased sample preparation.

For the highest accuracy measurements, samples are fabricated using standard semiconductor microfabrication to create a membrane structure as shown in Fig. 3. At least two electrical elements are microfabricated on the membrane to serve as a temperature sensor and heater, independently. Again, electrical isolation layers of dielectric materials are used, typically on the order of 10-30 *nm*. If a steady state heating current is passed through the heater on the membrane, the thermal conductivity is given in a straightforward manner by:

$$\kappa = \frac{P\,L_{heater-sensor}}{2L_{gauge}\,t_{film}\,\Delta T} \tag{3}$$

where κ is the thermal conductivity, P is the power applied to the heater, $L_{heater-sensor}$ is the distance between the heater and temperature sensor, L_{gauge} is the test section length, and t_{film} is the film thickness.

The aspect ratio of the heater design is important. Typically, the gauge length to membrane width should be a factor of 3 of greater in order to reduce membrane boundary effects on the heat flow. The width of the window and spacing of the heaters can be determined by the thermal healing length. The thermal healing length is the distance over which the

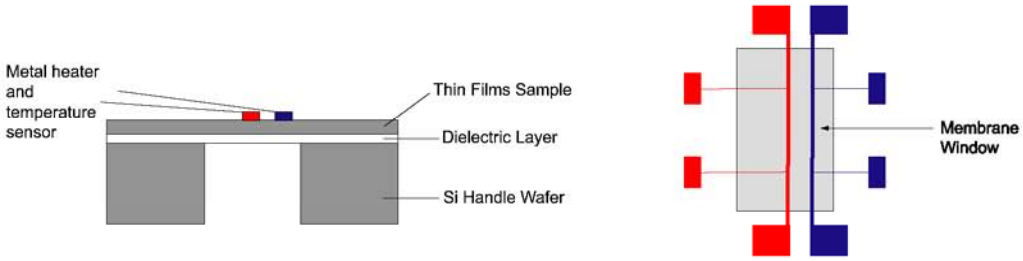

Figure 3. Diagram showing the steady state electrical resistance thermometry sample with a thin suspended membrane. Left diagram displays the side view and right diagram shows a top view.

temperature decays to the ambient temperature of the controlled environment. This can be estimated by using an extended surface analogy and assuming radiation losses along the membrane and conduction through the membrane:

$$L_h = \sqrt{\frac{t_{film}\,\kappa_{film}}{\sigma(T^2 + T_{ambient}^2)(T + T_{ambient})}} \tag{4}$$

Here, κ_{film} is the film thermal conductivity, σ is the Stefan-Boltzmann constant, T is the average film temperature, and $T_{ambient}$ is the surrounding temperature. For materials like polysilicon, this generally results in thermal healing lengths on the order of 1-2 mm. Thus, heater and temperature sensors can be spaced hundreds of microns apart and provide excellent measurements of temperature and thermal conductivity.

In cases where the membrane structure shown in Fig. 3 cannot be fabricated, samples can be measured in the multilayer film configuration shown in Fig. 4. This sample design requires less microfabrication, similar to the 3-ω sample, but is more prone to experimental uncertainty. In this sample, a thin-film of high thermal conductivity is deposited on a thick layer of material with low thermal conductivity. This induces 1-D heat spreading in the thin film layer. The thermal conductivity is still calculated according to Eq. 3. However, the thermal healing length is much smaller in this case. This can be estimated by:

$$L_h = \sqrt{\frac{t_{film}\,\kappa_{film}\,t_{oxide}}{\kappa_{oxide}}} \tag{5}$$

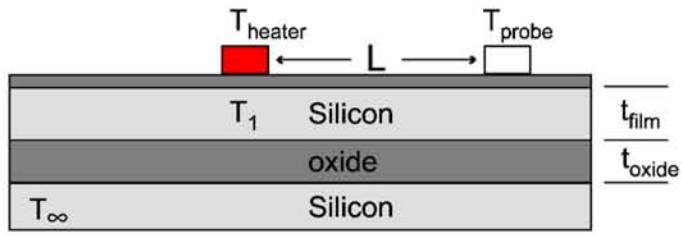

Figure 4. Diagram showing the steady-state thermal conductivity sample used for in-plane thermal conductivity measurements on highly conductive layers.

where t_{oxide} is the oxide layer thickness, and κ_{oxide} is the oxide layer thermal conductivity. This results in thermal healing lengths values on the order of 15 μm or less. Thus, heater-sensor spacing must be on the order of a few microns. With the heater-sensors being on the order of 2 μm in width or greater, measurement of temperature over this distance is no longer a "point" measurement but an average over distance in which the temperature is decaying rapidly. Thus, more uncertainty can be found in these types of measurements.

2.2 Measurements

Thermal conductivity measurements were performed on single crystal and polycrystalline Si thin-films. Silicon was chosen as a model material because information such as its dispersion relationship is already known, it is amenable to microfabrication, and it is of technological importance to the development of surface micromachined MEMS. Thin films of single crystal silicon were obtained in the form of silicon-on-insulator wafers from SOITECH, Inc. Films ranging from 340 nm to 1500 nm were measured. The sample sizes chosen were limited to the availability of wafers from the supplier with the appropriate buried oxide thickness. The buried oxide thickness was kept at 3.0 μm in order to increase the thermal healing length, especially for the thin silicon layers. All samples for the single crystalline silicon layers were amenable to the structure shown in Fig. 4 and were tested under steady state conditions. The healing length ranged from 10-25 μm. For these samples, heaters and sensors were made from evaporated aluminum films. They were etched to be 2 μm wide and between 4-10 μm apart. Aluminum was chosen since it provided excellent sensitivity, with a relatively large $\partial R/\partial T$ (compared to other metals) of approximately 0.4%/K throughout the temperature range of interest.

For polysilicon layers, films were deposited by LPCVD at the Center for Integrated

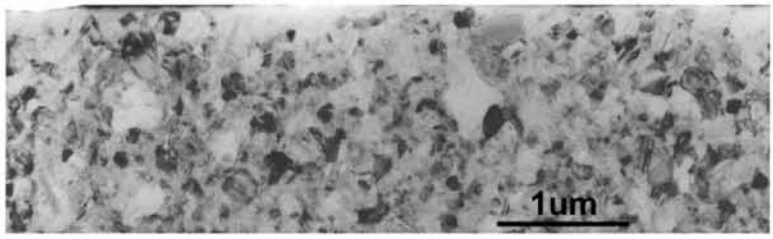

Figure 5. TEM image showing the grains in the as-deposited polysilicon. Average grain size was determined to be approximately 130 *nm*.

Systems at Stanford University. Polycrystalline silicon films (2 μm thick) were deposited on thermally oxidized silicon wafers at 620°C. This resulted in as-deposited polycrystalline films with an average grain size on the order of 130 *nm* (Fig. 5).

The polysilicon films were annealed in a nitrogen environment at 1100°C for 1 hour. In order to study the effects of multilayer processing found in MEMS devices, a select number of wafers were reinserted into the LPCVD chamber to deposit a second layer of polysilicon for a total thickness of 4 μm. This second layer was also annealed at 1100°C for 1 hour. The in-plane thermal conductivity was measured using electrical resistance thermometry on free-standing membranes of polysilicon created through standard Bosch etch processes. These samples were amenable to testing using the structure in Fig. 3. With a healing length on the order of 2 *mm*, the heaters and temperature sensors could be patterned much further apart than in the single crystal samples. Again, heaters and sensors 2 μm in width were used, but they were placed 300 μm apart. The membrane window was 3.6 *mm* by 1.5 *mm*. The gauge length of the sensor and heater was 700 μm.

For experiments, samples were cleaved from the silicon wafers and packaged in ceramic packages using wirebonding or conductive epoxy (Fig. 6). The samples were placed in a liquid nitrogen cooled cryostat and evacuated to a pressure on the order of 1×10^{-5} torr. The temperature was controlled through a cold pedestal on which a 50W heater was attached. This provided the capability to control the sample from 77-325K. A radiation shield was also employed to reduce radiation heat losses to the environment. The radiation shield was gold coated and controlled to be the same temperature as the sample.

The predominant sources of experimental error in the steady state electrical resistance thermometry method arise from uncertainty in the film thickness and the slope of the resistance versus temperature ($\partial R/\partial T$) of the heater/sensor element. Care was taken to reduce this error by quantifying the film thickness using broadband reflectometers and scanning

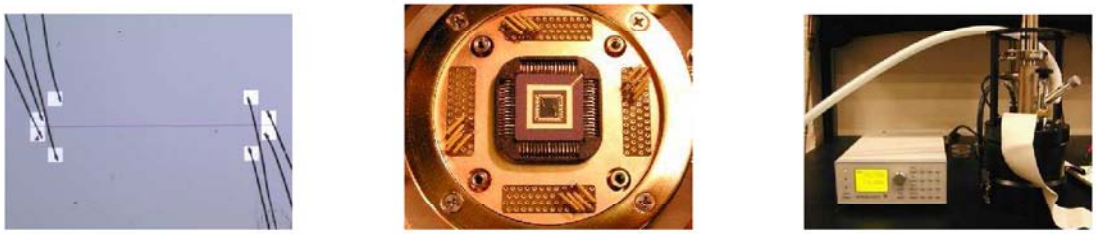

Figure 6. Pictures showing the patterned heater and temperature sensor with wirebond connections (left) packaged and placed in a cryostat (middle) and cooled to the desired test temperature (right).

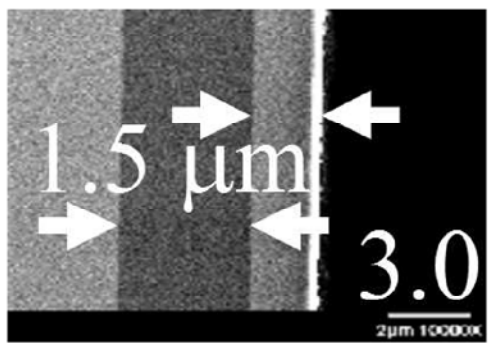

Figure 7. An SEM photograph showing the thickness of the various layers in the SOI wafer structure used for thermal conductivity testing.

electron microscopy (Fig. 7). Values of $\partial R/\partial T$ were carefully calibrated in the cryostat, with a measurement at each test temperature.

While the samples tested in this study primarily used steady-state electrical resistance thermometry, the 3-ω method was also developed and implemented during the project to provide additional capabilities for thermal conductivity measurements. This work was done largely by Brandon Olson, an IGERT and MESA Institute summer intern from the University of Utah. While results from the 3-ω method will not be presented in this report, it should be mentioned that this capability does exist at Sandia as well as some new data reduction algorithms as a result of this LDRD.

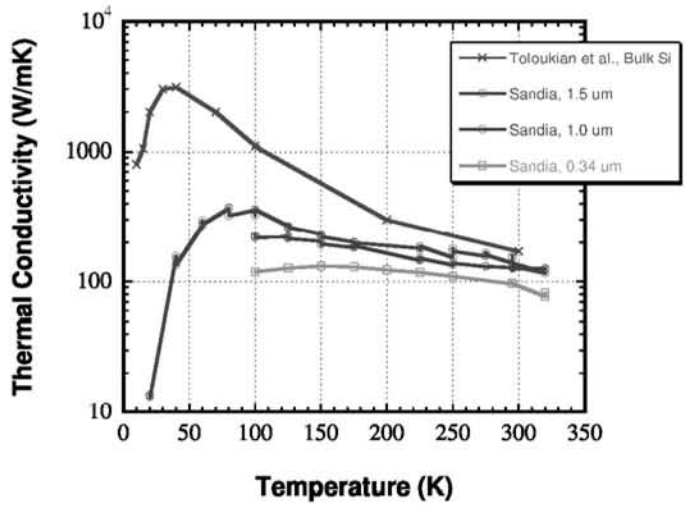

Figure 8. Thermal conductivity of single crystal Si as a function of temperature. Data compare the results of thin-films to bulk silicon.

2.3 Results

The experimental results for single crystal silicon are shown in Fig. 8 for films 1.5, 1.0, and $0.34\,\mu m$ in thickness. Data for the $1.5\,\mu m$ film shows very little reduction in thermal conductivity at room temperature when compared to bulk silicon ($\kappa = 148\,W/m \cdot K$). At lower temperatures however, the data are markedly different. Near 100K, the thermal conductivity of the $1.5\,\mu m$ film is an order of magnitude smaller than the bulk result, with a value of approximately $330\,W/m \cdot K$, and the peak value, which is only slightly larger, occurs near 80K. In analyzing the bulk thermal conductivity data, it can be seen that the peak in thermal conductivity occurs at a much lower temperature (40K) with a peak value of $3000\,W/m \cdot K$. Similar results are also seen for the other two films. For the $340\,nm$ film, there is a reduction of thermal conductivity on the order of 40% at room temperature ($\kappa = 86\,W/m \cdot K$). This shows that length scale effects are still dominant at relatively high temperatures at this size scale.

The reason why the thermal conductivities are so different at lower temperatures and converge at higher temperatures can be attributed to boundary scattering effects. At high temperatures, the dominant phonon scattering mechanism is Umklapp scattering where the scattering mean free path, λ, is dominated by temperature effects ($\lambda \sim \exp(\Theta/T)$) and is much smaller than the thickness of the film (Θ is the Debye temperature, which is 645K

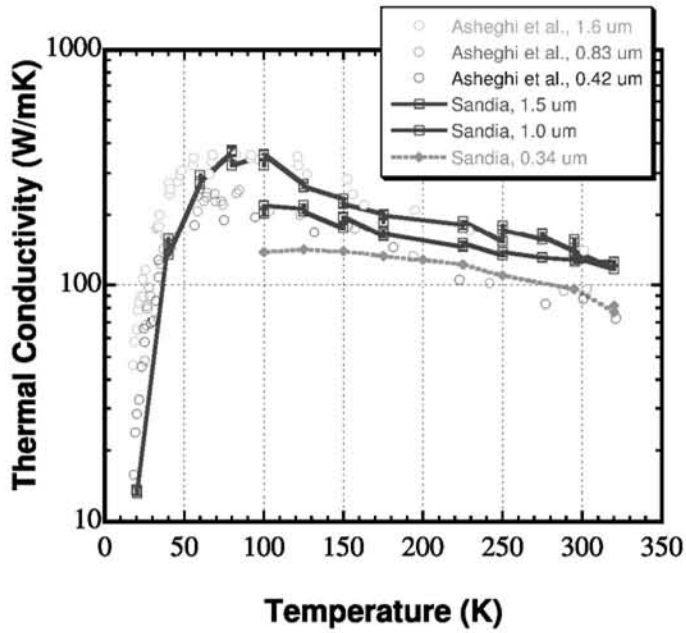

Figure 9. Comparison of data from experiments at Sandia to those of our collaborator's group at Stanford [16]. Good agreement between measurements is seen.

for silicon). While phonons are still scattering from the boundaries, the relative contribution this has to thermal impedance is quite low. As the temperature decreases, boundary scattering becomes more dominant as the mean-free-path for phonon-phonon interactions increases. Since boundary scattering length scales are fixed, this limits the maximum conductivity that can be achieved in the material, thus the observed lower peak.

Figure 9 shows a comparison of the thermal conductivity data measured on single crystal Si to data obtained by Asheghi [16] while a member of Ken Goodson's group at Stanford University. The data show fairly good agreement for this size scale. Data on samples fabricated for this study were taken both at Sandia and at Stanford University to cross check the testing procedure. While this does not ensure the complete accuracy of the data, it provides some measure of benchmarking between two different experimental studies. Thus, we have fairly good confidence in the experimental data. Based on a more formal analysis, uncertainties in the data taken in both studies range between 15-20%.

Experiments on doped single crystal silicon was also performed on films that were $1.5\,\mu m$ in thickness. Films were doped with boron to levels of 1×10^{18}-$10^{20} atoms/cm^3$.

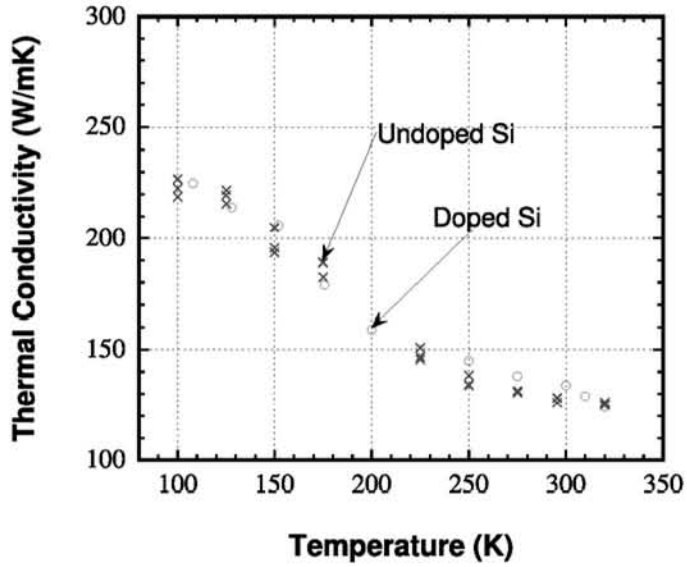

Figure 10. Data showing a comparison of the thermal conductivity of doped and undoped single crystal silicon. Virtually no difference is seen over the temperature range measured.

Results of this experiment are shown in Fig. 10. Overall, no observable difference in thermal conductivity between the doped and undoped films were seen. It is unclear as to why this occurred, but the results may be limited to the temperature regime over which the tests were performed. Differences in the thermal conductivity of doped and undoped silicon have been observed at temperatures below 80K (Ref. [8], Fig. 6) whereas we were limited to testing down to 87K with the experimental setup at the time. This resulted from having a dual stage heater in the cryostat that had to maintain a temperature difference of at least 7K. We found during operation that the sample could be controlled much easier at low temperatures if this were increased to 10K, thus putting us at a minimum test temperature of 87K for liquid nitrogen cooling.

The thermal conductivity of single and double-layer ($2\,\mu m$ per layer) undoped polysilicon films is shown in Fig. 11. The data show that the thermal conductivity ranges from approximately $60\,W/m\cdot K$ at room temperature up to a peak value of $67\,W/m\cdot K$. Data for the double-layer polysilicon film shows a slightly lower conductivity value ranging from $51\,W/m\cdot K$ at room temperature up to $61\,W/m\cdot K$ at 150K.

It is interesting to note that the thermal conduction values reported in this work are quite similar or greater in magnitude to those seen for doped polysilicon [8] films and 4-5

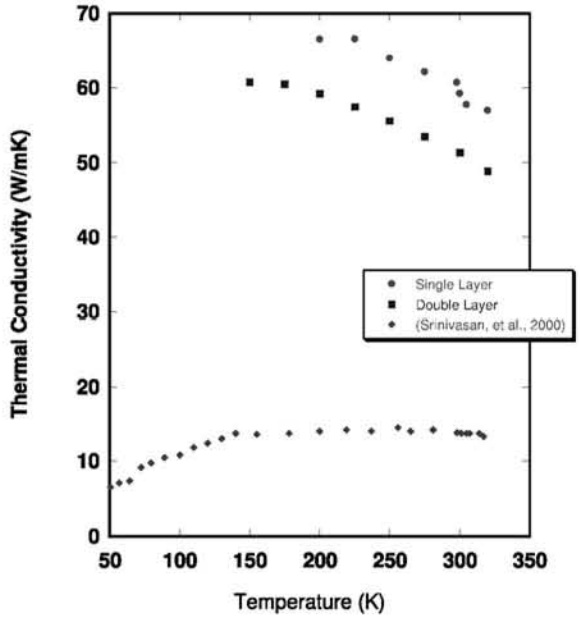

Figure 11. Data showing the thermal conductivity of single and double layer recrystallized undoped polysilicon films compared to previous results on an as-deposited single layer film (Srinivasan *et al.*[17]). Data show that the recrystallized materials posses a larger thermal conductivity as well as a stronger temperature dependence

times greater than previously reported for undoped polysilicon films ($\sim 12\,W/m\cdot K$) [17]. While the same deposition program and facility was used by Uma Srinivasan *et al.* [17], the additional step of the 1100°C anneal induces a large change in the thermal conductivity. The data from Ref. [17] are also included in Fig. 11. It can be seen that the thermal conductivity in the as-deposited film is not only lower, but it has very little temperature dependence above 150K. This is in contrast with the films that were recrystallized in our study. This difference is due to the fact that the Umklapp processes are more dominant in recrystallized materials which possess a smaller defect density. For dielectric materials, the phonon mean free path is determined predominantly by Umklapp scattering, phonon-grain boundary scattering, and phonon-defect scattering. Of these, Umklapp scattering is the only temperature dependent quantity and becomes dominant when the defect density is low. Thus, the behavior of the thermal conductivity is not solely grain size dependent, but also defect density dependent. It is expected that the film boundaries had little effect on the thermal conductivity seen in these films, based on the study of single crystal silicon.

Transmission electron microscopy analysis showed an average grain size on the order

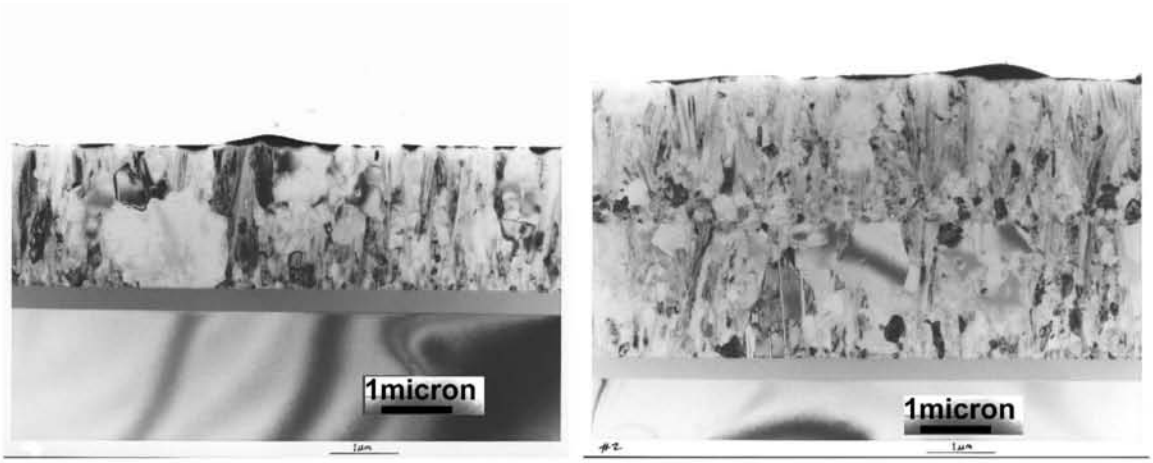

Figure 12. TEM images of the microstructure of the undoped polysilicon. Left picture is a single layer film while the right shows a double layer film

of 150 *nm* in the polysilicon material prior to the high temperature annealing. This was consistent with previous results on undoped polysilicon [17]. However, the 1100°C anneal promoted dramatic recrystallization, with the formation of grains on the order of 580 *nm* in the single layer material (Fig. 12). It is clear to see that some grains are significantly larger than 1 μm. Images with increased magnification also show that these recrystallized regions have few twins or other defects as seen in the regions without significant recrystallization. The presence of these regions provide thermal transport paths with much less resistance. The data also show that the peak in thermal conductivity appears to be shifted to higher temperatures than observed in single crystalline silicon films. This is due to the fact that the phonon mean-free-path is dominated by phonon-grain structure interactions, which are spaced on the order of the grain size (fully recrystallized grains) or smaller. Thus, higher temperatures must be reached than seen in single crystalline Si before phonon-phonon scattering effects become dominant.

TEM micrographs in Fig. 12 show the microstructure of the films with multiple layers of polysilicon. In this case, the grain size in the top and bottom layers were measured. The data show that the grains in the bottom layer are on the order of 580 *nm* while the top layer has a slightly smaller grain size (440 *nm*), with a significant layer of small grains near the interface of the two layers ($\sim$100 *nm* or less in size). The slight difference in grain size may be the cause of the smaller thermal conductivity measured for these two films. It should be

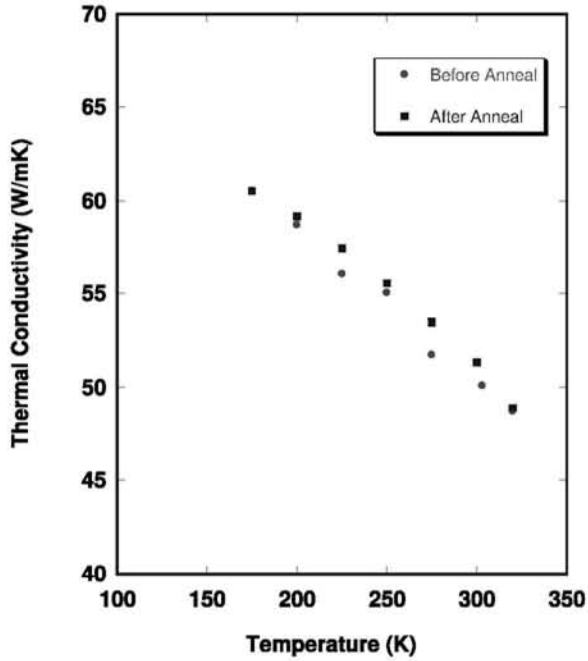

Figure 13. Data showing the thermal conductivity of recrystallized polycrystalline silicon before and after an 8 hour, 700°C anneal. Data show no detectable difference in thermal conductivity for the samples.

noted that the difference in conductivity values was consistent and results were repeatable across several samples taken from each wafer. While it is not clear as to why the top layer has a smaller grain size, it may be due to slight temperature differences between the two depositions. It is well known that polysilicon grain size and morphology are highly temperature dependent and nonlinear between 580-630°C. Thus, a small difference in temperature of a few degrees could be enough to induce this temperature change.

Data in Fig. 13 show the effects of downstream heat treatment on the thermal conductivity of the undoped recrystallized polysilicon material. Both films were annealed at 1100°C for 1 hr after film deposition. This was followed by heat treating one of the wafers at 700°C for 8 hrs. This temperature was chosen in order to be consistent with or slightly greater than many "moderate" temperature CMOS fabrication processes (low temperature oxide, polysilicon deposition, silicon nitride deposition, etc.) and temperatures that are encountered in device service. Both thermal conductivity data and TEM analysis showed no measurable changes in either thermal conductivity or grain size.

3 Material Model

Material models, for the purposes of this work, include phonon characteristics that are independent of the problem at hand except perhaps for the material chosen. These phonon characteristics include the equilibrium distribution, the dispersion relation, and the group speed. The material models described herein are employed, practically verbatim, in all subsequent models discussed in this work so they are grouped into an independent section.

3.1 Equilibrium Distribution

Phonons are members of a particle class known as bosons, which are characterized by overlapping wavefunctions and no exclusion principle. At a given temperature, T, the number of phonons expected to occupy a given state, which is defined by a frequency, ω, and a polarization, p, is given by the Bose-Einstein distribution:

$$< n >= \frac{1}{e^{\frac{\hbar\omega}{\mathcal{B}T}} - 1}. \tag{6}$$

Note that the polarization appears only through the allowed frequencies, given by the dispersion relation described in the next section. For the temperatures treated in this document, the quantization of phonon occupation numbers will be neglected. In that case, Eq. 6 may be viewed as a continuous function representing the number of phonons per unit frequency found in a small region around ω. We will therefore drop the brackets and reference the Bose-Einstein distribution through $n(\omega, T)$.

3.2 Dispersion Relation

A dispersion relation is the key link between the simulation and the material. It relates the wavenumber of a phonon to its frequency, which determines its energy and its phase speed. Also important, the slope of the frequency versus wavenumber relationship at a given wavenumber, k, determines the group speed of a wave packet centered at k, which is the speed at which it will carry information. Because this work concerns energy transport, group speed is a critical parameter.

Dispersion relations for a number of materials have been measured experimentally via neutron scattering, as well as computed using *ab initio* molecular dynamics techniques [18]. For computational convenience, an analytical representation of the dispersion relation

is ideal. The literature contains a number of proposals for such representations, many of which are reviewed in [19]. In this work, a fourth-order polynomial was chosen. The dispersion relation is therefore represented by

$$\omega = A_1 k + A_2 k^2 + A_3 k^3 + A_4 k^4 \tag{7}$$

Fitting Eq. 7 to neutron scattering data for the [100] direction in silicon from Brockhouse [20] and Nilsson and Nelin [21], the coefficients are found to be:

$$\begin{aligned} A_1 &= 16.889222 \\ A_2 &= -2.597880 \\ A_3 &= -4.385738 \\ A_4 &= 2.132736 \end{aligned}$$

for the longitudinal acoustic (LA) branch and

$$\begin{aligned} A_1 &= 11.396625 \\ A_2 &= -5.132308 \\ A_3 &= -6.17036 \\ A_4 &= 2.132736 \end{aligned}$$

for the (degenerate) transverse acoustic (TA) branches to yield ω in radians per second.

The resulting polynomial representation of the dispersion relation is compared to the experimental data in Fig. 14.

For simplicity at this stage, we will use the [100] dispersion relation for all directions, *i.e.* the material is assumed isotropic. Under this assumption, triple integrals over the wavevector, which occur when computing quantities such as total energy, may be replaced by a single integral over the wavenumber by converting to a spherical integral over a sphere of radius k and observing that the integrands of interest are independent of the azimuthal and polar angles. The advantage of this formulation over the more common conversion to frequency integrals is that the density of states is simply $4\pi k^2/a^3$ (under the periodicity assumption) [22] instead of $4\pi k^2/a^3 \frac{dk}{d\omega}$. This removes the numerical difficulties associated with integrating in the region (which is considerable for TA phonons) near the Brillouin zone boundary where $\frac{dk}{d\omega}$ becomes infinite.

As a check on the isotropic assumption, the dispersion fit, and the k-space integral formulation, we can calculate the specific heat per unit volume via:

$$C = \frac{\partial U}{\partial T} = \sum_p \frac{4\pi\hbar^2}{\mathcal{B} a^3 T^2} \int_0^1 \omega^2 n^2 e^{\frac{\hbar\omega}{\mathcal{B}T}} k^2 \, dk \tag{8}$$

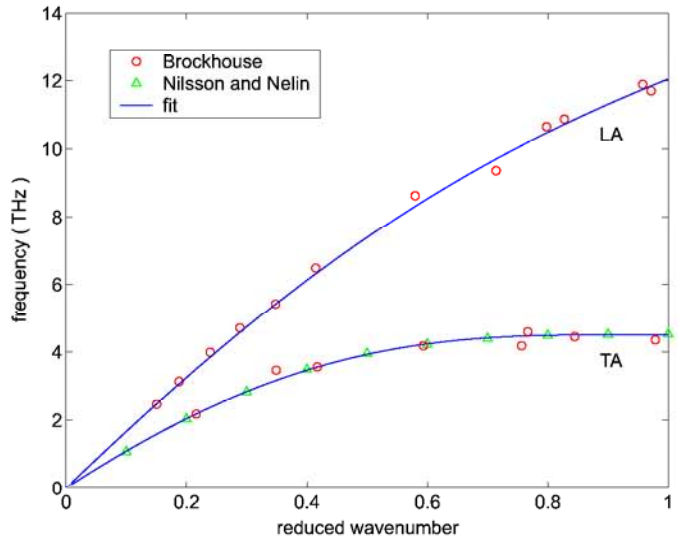

Figure 14. Comparison of the dispersion relation fit in [100] silicon to experimental data [20, 21].

where U is the energy per unit volume, which is obtained by integrating the phonon energy, $\hbar\omega$, over the distribution function.

The result of Eq. 8 as a function of temperature is shown in Fig. 15. While reasonable agreement is obtained at low temperatures, the calculated specific heat departs increasingly from the measured values as T increases. Surmising that the error is due to the omission of optical modes from the calculation, they are included in an approximate manner by assuming that all three modes oscillate at 14THz, which makes evaluation of Eq. 8 trivial. Adding this contribution to the previous result and plotting in Fig. 15, good agreement is obtained over the entire temperature range.

With this representation of the dispersion relation in hand, it is straightforward to evaluate Eq. 6. For example, to compute the total number of phonons per unit volume at a given temperature, we integrate n over k:

$$N = \sum_p \frac{4\pi}{a^3} \int_0^1 n_p k^2 dk \tag{9}$$

which yields 6.3×10^{28} phonons/m^3 at 300K.

To view the relative populations of the modes as a function of wavenumber, we can plot the integrand of Eq. 9, as shown in Fig. 16(a), which shows relative populations of TA

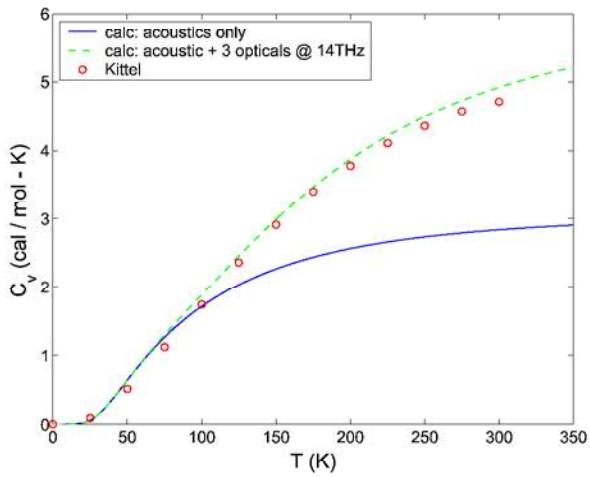

Figure 15. Comparison of measured specific heat to values computed with Eq. 8 under the dispersion model and k-space integrals used in the current work.

and LA phonons at each wavenumber at 300K. It should be noted that the TA populations have been doubled to account for the two degenerate branches. For comparison, Fig. 16(b) shows the same plot at 80K. Comparing these cases, it is clear that more phonons are found at higher frequencies as the temperature increases. At 80K, the high wavenumber LA phonons are nearly absent. The high wavenumber TA phonons are still dominant but a peak is beginning to form at lower wavenumber. This peak becomes dominant as the temperature decreases further.

3.3 Phonon Group Speed

An accurate representation of the carrier speed is critical to any transport simulation. For phonons, information is carried at the group speed of the wavepacket centered at k, which is given by

$$v_g = \frac{d\omega}{dk}. \tag{10}$$

This derivative is trivial to compute under the polynomial approximation to the dispersion relation of Eq. 7, yielding:

$$\omega = A_1 + 2A_2k + 3A_3k^2 + 4A_4k^3 \tag{11}$$

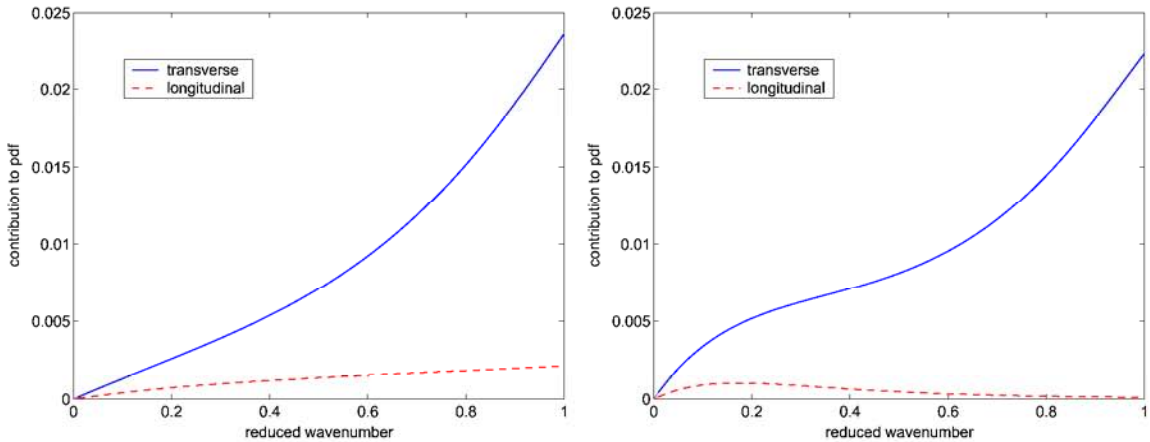

(a) 300K (b) 80K

Figure 16. Relative probability of TA and LA phonons as a function of wavenumber.

The resulting group speeds for the coefficients given in the previous section are shown in Fig. 17. Two weaknesses in the current polynomial dispersion model are apparent in this figure. First, the TA group speed drops slightly below zero at $k = 0.913$, then reemerges and reaches $128\,m/s$ at $k = 1$. In this work, the group speed is manually set to zero above $k = 0.913$. A second issue is that the LA group speed does not go to zero at $k = 1$. Given the good fit to the highest k measured in the experiments, it is likely that the the speed drops very rapidly to zero near $k = 1$, thus the region in wavenumber space with an incorrect speed is relatively small. It may be noted from Fig. 16(a), however, that $k = 1$ is the most probable LA wavenumber at room temperature, though the distribution is rather flat and TA modes are much more numerous. While it is likely that the effects of this shortcoming are minimal, this conclusion cannot be guaranteed, thus alternative dispersion models should investigated in future work.

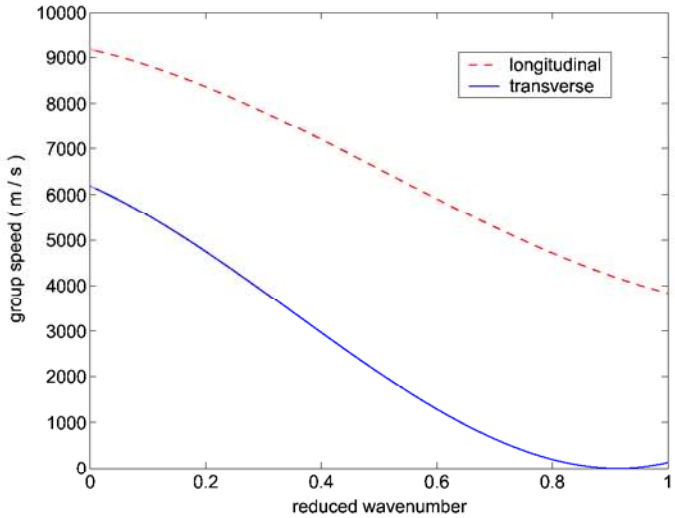

Figure 17. Phonon group speed under dispersion model.

4 Monte Carlo Simulation

The term "Monte Carlo" (MC) appears in a diverse array of generally numerical techniques used in a wide variety of fields. The one feature that binds these techniques is some degree of randomness in their representation, hence the "Monte Carlo" moniker.

The simulation scheme employed in this work most closely resembles the method known as direct simulation Monte Carlo (DSMC). This technique was developed by Bird [23] for simulating gas dynamics in regimes where the continuum assumption breaks down and the Navier-Stokes equations are no longer accurate. Later analysis showed rigorously the relationship between DSMC and the Boltzmann equation [24] and it has been employed in flows ranging from free molecular to fully continuum.

The "direct simulation" portion of the DSMC name comes from the representation of a gas as a collection of loosely interacting particles, as it is in the physical case. These particles move about the domain and interact with each other and with the boundaries in a (generally) physically reasonable way. The "Monte Carlo" portion of the DSMC name is required due to some assumptions required to make the code tractable. First, a smaller number of particles are represented in the simulation than exist in the physical case. This assumption is required due to the vast number of particles in most cases of interest. Functionally, this means that each computational particle represents a volume, rather than a point, in phase space. Second, the move and collide steps are decoupled. This greatly simplifies both steps and can make the code more computationally efficient in terms of pipelining and cache use. A key consequence of these assumptions is that binary collisions defined in terms of two particles arriving at the same point in space and time are no longer possible. Under move/collide decoupling, particles move directly through each other without notice. Furthermore, because the computational particles represent volumes in phase space, several intersection points are possible for a given pair. Collisions are therefore handled in a statistical manner, with stochastic determination of collision pairs as well as the post-collision state.

Functioning under the "phonon gas" assumption (Sec. 1.1), the Boltzmann equation captures the relevant phenomena, thus DSMC provides a reasonable technique for this investigation. Mazumder and Majumdar [25] published the first work in this direction in 2001. Many of the submodels used in the current work build upon the framework developed by these authors.

The code developed for this work will be referred to as MOCAPHTS, for Monte Carlo Phonon Heat Transport Simulator. It was written in C++ and has been run on a number of platforms (Sun, SGI, LINUX) in serial and in parallel. MOCAPHTS includes compile-time

selection between one and two dimensional versions. Details of the collection of models and techniques that make up the code, as well as presentation and discussion of its results are collected in this section.

4.1 Baseline Code Description

4.1.1 Random Number Generation

Because computers generally perform instructions in a consistent manner, it is impossible for them to generate truly random quantities. Algorithms have been written, however, that produce sequences that appear random by satisfying various statistical tests (sometimes referred to as "pseudo-random" sequences). An exhaustive overview of the topic can be found in chapter 3 of Knuth [26].

Due to their stochastic nature, Monte Carlo algorithms can be very sensitive to the quality of the random number generators they employ. Press *et al.* [27] note that the system-supplied rand() is often flawed and assert that "If all scientific papers whose results are in doubt because of bad rand()s were to disappear from library shelves, there would be a gap on each shelf about as big as your fist."

Most system-supplied rand() routines are of a class known as linear congruential generators, which create a sequence according to:

$$R_{f_{n+1}} = (aR_{f_n} + c) \; mod \; m. \tag{12}$$

where a, c, and m are chosen constants.

These generators are popular because they require very few operations, thus they are very fast and can be inlined. In the context of the simulation technique used in this section, however, the linear congruential generators have a key flaw: successive calls can be related in unfortunate ways. For example, if R_{f_n} is exceptionally small, $R_{f_{n+1}}$ *cannot* be large. This type of sequential correlation can be easily visualized by using the generator to choose coordinates in a space discretized into "bins" and placing a marker in each bin when a coordinate is selected within it. In the ideal case, for sufficiently large sequences, all the bins would have an equal number of markers when the algorithm terminates. This test is referred to as the "serial test" by Knuth [26] and code for such a test in 2D is given in Appendix E in Bird [23].

The results of this test for a linear congruential generator with $a = 197, c = 0$, and $m = 9999943$, used in a Ph.D. thesis on simulating hypersonic flows with DSMC, is given

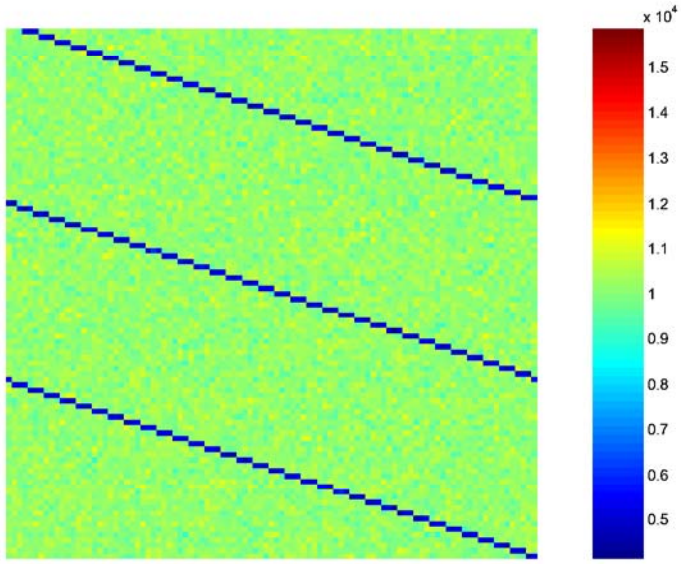

Figure 18. Series test for a linear congruential generator used in previous hypersonic DSMC work.

in Fig. 18. In this case, a 100×100 grid was set up and 10^8 pairs of coordinates were chosen. We therefore expect 10,000 markers in each bin. It can be seen, however, that some contiguous sets of bins have fewer than 5,000 markers in them, implying that certain pairs of numbers tend to be neglected. In general, it can be shown that the coordinates chosen in this manner in t dimensional space will lie on a maximum of about $m^{1/t}$ planes of dimension $(t-1)$ [26]. This can be an issue when the acceptance-rejection method outlined in App. A is used to choose from a distribution. For example, if a small number is improbable, it will be accepted more often than appropriate because the acceptance trial will also choose a small number and will thus pass the subsequent probability test. Because the driving force in MOCAPHTS is much smaller than in a DSMC simulation of hypersonic flow, this generator caused observable anomalous behavior in the current work.

Schemes for shuffling the sequence have been proposed [26] to eliminate the sequential correlation problem. In this work, a subtractive scheme of an entirely different nature was eventually adopted. An implementation of this scheme named ran3() is given in [27]. Taking advantage of C++, the first conditional and the pointer argument to the generator were eliminated by making the seed private. Another conditional was eliminated by taking advantage of unsigned integer rollover to obviate checking for negative numbers. These modifications increased the speed of the scheme by a factor of approximately 2.2. A plot similar to Fig. 18 is given for the new scheme in Fig. 19. In this case, no sequential

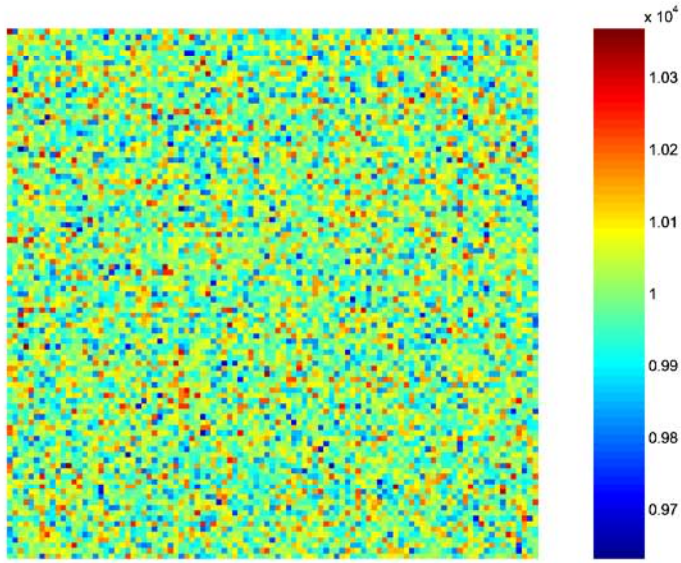

Figure 19. Series test for modified ran3() subtractive random number generator from Press *et al.*.

correlation is observable, even though the color axis has been tightened significantly around the expected value of 10,000.

4.1.2 Weight Factor

In order to make the computation tractable, the number of particles in the simulation is generally smaller than the number present in the physical case. In MOCAPHTS, the user specifies the target particle count for the simulation. A "weight factor" is then computed by calculating physical phonon count and dividing by the target. The physical phonon count is determined by calculating the number of phonons per unit volume at the specified reference temperature (Eq. 9) and multiplying by the domain volume.

4.1.3 Exclusion of Stationary Phonons

It may be observed from Fig. 16(a) that the most probable phonon at room temperature is TA with $k=1$. From Fig. 17, however, it is apparent that these phonons are unable to transport energy because they have zero group speed. During a calculation, any energy

transferred to these phonons remains stationary until it is subsequently transferred to a mobile mode. Because the cases treated in this work are steady-state, this process is not of interest and skipping it by transferring energy only from mobile mode to mobile mode will not change the final result. This is equivalent to setting the relaxation time for these modes to zero. The same argument holds true for the optical modes.

Optical modes and TA phonons with $k > 0.913$ (see Sec. 3.3) are therefore omitted from the calculation. This enables the computational particle count to be reduced by assigning all particles to modes that directly contribute to transport. In effect, for a given number of computational particles, the weight factor is decreased because the phonons in the simulation are used to simulate a subset of the phonons that exist in the physical case. It should be noted that, when using this simplification, the upper limit on wavenumber integrals involving TA phonons, such as Eq. 9, must be changed from 1 to 0.915 in order to properly weight the remaining modes.

4.1.4 Timestep

Due to the cell-based collision scheme, a particle should have an adequate opportunity to collide, and thus make its presence felt, before it leaves the cell. In addition, the timestep must be small compared to the mean collision time to justify decoupling the move and collide phases [23]. To satisfy the first constraint, the timestep is set to some fraction (usually 1/2) of the time it takes the fastest particle (an LA phonon with $k = 0$) to cross a cell at its narrowest point. To satisfy the second constraint, in conjunction with the first, the cells are set to some fraction (roughly 1/4) of the mean free path at the reference temperature.

4.1.5 Initial Conditions

To initiate a calculation from scratch (rather than from the particle field saved from a previous calculation) computational particles must be placed in the domain. For each particle, the position, wavenumber, polarization, and velocity must be chosen. To ensure a reasonably accurate specification of the starting temperature despite the random nature of the particle initialization, particles are added until the target energy, rather than until the target number of particles, is reached. To implement this requirement, the total energy is computed before and after the creation of each particle. If the addition of a given particle brings the total closer (in absolute value) to the target, the creation process is continued. If, however, the particle's addition moves the total energy further from the target, the particle is discarded and the process is terminated.

When creating a computational particle, the wavenumber is chosen first. This selection is performed using the acceptance-rejection method outlined in App. A. To begin the process, a wavenumber is chosen from a uniform distribution. The total number of phonons per unit k at this k (including both polarizations), n, is then calculated at the reference temperature and normalized by the maximum value, n_{max}. A second random number, R_f, is then chosen and compared to this result. The chosen k is accepted if $R_f < n/n_{max}$. If not, the process begins again with another randomly-chosen k.

To choose the polarization, the fraction n_{LA}/n is computed at the target wavenumber. A random fraction is then generated. If $R_f < n_{LA}/n$, then the particle is assigned to the LA branch, otherwise, it is assigned to the TA branch.

Once the wavenumber and phonon type are determined, the frequency and speed of the particle are determined according to equations 7 and 11, respectively. The velocity vector is then randomly oriented according to a spherically uniform distribution to determine the components of the particle speed.

4.1.6 Temperature Determination

In most cases, temperature is a convenient output quantity. For some boundary conditions and scattering routines, it is a required quantity. We therefore require a means of determining the temperature in a given region of the domain, usually a cell, given the microscopic state within it.

It is quite straightforward to measure the total energy in a region by tallying the phonons it contains and multiplying by the weight factor. It is also straightforward to determine the energy per unit volume of an equilibrium state at a given temperature via:

$$E = \sum_p \frac{4\pi}{a^3} \int_0^1 n\hbar\omega\, k^2\, dk. \tag{13}$$

where the temperature dependence appears in n via Eq. 6. A measure of the temperature may therefore be obtained by inverting Eq. 13. Unfortunately, this equation is not easily invertible. While the inversion may be done numerically using standard techniques, this process is rather expensive because it involves numerical integral evaluations at each root finding step. In an early version of the Monte Carlo code developed for this work, it was discovered that 96% of the total computation time was spent on this operation.

A plot of energy as a function of temperature is shown in Fig. 20. At high temperature, the relationship is linear while at very low temperature, it is quartic. One could therefore

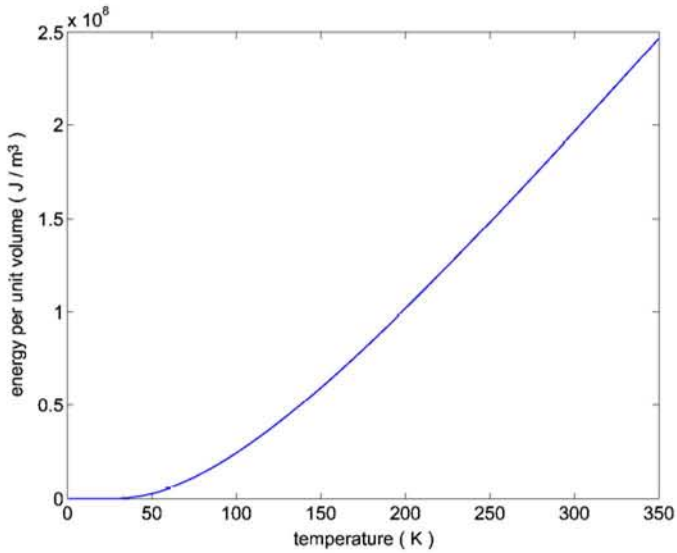

Figure 20. Energy per unit volume of an equilibrium phonon distribution at a given temperature.

conceive of a scheme that switches from quartic to linear at some temperature, with a bridging function in between.

As a simpler option, a lookup table with an adjustable step size linear interpolation scheme was implemented in this work. In this scheme, the step size is adjusted until the "worst case" error in the calculated temperature reaches a target value (usually 0.01%). Because the nonlinearity of the E vs. T curve increases as temperature decreases, it is assumed that the worst case error occurs in the middle of the lowest temperature interval. The step size that satisfies the error constraint may then be found using a Ridders' method root finding scheme [27]. The lookup table is then built by numerical inversion at each point. If, during the simulation, an energy lower than the lowest value in the table is encountered, a new step size is found and the table is rebuilt. If an energy higher than the maximum is encountered, more entries are added to the table using the same step size. Because many simulations involve a relatively small temperature range, this scheme often requires very little memory. In addition, lookups and interpolation are very fast because the step size is constant throughout the table. Speed is important in this operation because the temperature is calculated in every cell during every timestep.

4.1.7 Cell-Finding Scheme

Once a particle is moved, its position on the grid must be determined. In general DSMC codes on unstructured grids, this operation can consume a large percentage of the total CPU time. Several methods have been proposed for treating this problem in the general case [28, 29, 30]. The scheme employed in this work represents a combination of techniques, involving both trajectory tracing and regular Cartesian grids.

The domains in this work are built from a series of "zones", each containing a body-fitted regular grid. The particle movement from zone to zone is handled with a trajectory tracing algorithm, *i.e.* if a particle is noted to have left its zone, the zone adjacent to the intersected face is interrogated or the boundary condition on that face is employed. This operation continues until the particle move is completed and the final zone is identified.

Once the final zone is found, determining the proper cell is then simply a matter of integer division: divide a position component by the cell size in that direction and truncate to the nearest integer, producing the cell index. Due to the finite precision with which the position is represented and the division is performed, this process occasionally yields a cell index that larger than the actual range. Because this only happens at the highest index, and the error is always exactly unity, this condition is trivial to detect and correct.

4.1.8 Boundary Conditions

Two boundary conditions are required to simulate the systems treated in this work: isothermal and adiabatic. The former boundary condition is used to establish a temperature gradient across a domain while the latter is required to represent boundary scattering for in-plane thermal transport simulations in thin films. It is a notable feature of the MC method, compared to the semiclassical methods, that thin film effects are considered separately as boundary conditions instead of lumping them with other scattering mechanisms via Matthiessen's rule [31].

Isothermal Boundaries

Isothermal boundaries are based on assuming a stationary equilibrium condition at the specified temperature on the boundary face. Given this assumption, the number of phonons expected to enter the domain through the face, N_f, during each timestep can be calculated as the equilibrium occupation number of each polarization and frequency times its average velocity normal to the boundary. Assuming a spherically-isotropic velocity distribution, this expression becomes:

$$N_f = \frac{\pi \Delta t A_{face}}{Wa^3} \sum_p \int_0^1 k^2 v_p(k) n_p(k) dk \tag{14}$$

Because the result is generally not an integer, but the number of particles is necessarily integral, a means must be chosen for dealing with the fractional part. Two primary options spring to mind: carry the remainder to the next time step, or treat the fractional part as a probability that an additional particle will be introduced in each timestep. Because the first option results in a regular pattern in time of extra particles being introduced, the second option was chosen as slightly more consistent with the stochastic nature of the simulation. It should be mentioned, however, that the first option was implemented in the first version of the code and no striking difference was observed in the results when the options were switched.

It should be noted that Eq. 14 is only strictly valid if the entire domain is in equilibrium at the specified temperature, which would not be an interesting calculation to perform. As temperature gradients in the domain increase, the implicit assumption of a zero mean speed will become an increasing liability. This is equivalent to specifying all macroscopic quantities in a subsonic fluid dynamics calculation even though some of the characteristic directions point upstream (see Ref. [32]), which indicates that some quantities must be calculated from within the domain. In the current work, this error is managed by minimizing temperature gradients to reduce the drift speed and maximizing domain length to reduce the boundary effects on the domain interior until the results become insensitive to these quantities. In the future, more sophisticated treatments may be required, such as the implicit method proposed by Liou and Fang [33] for gas molecules. It should also be noted that the particle velocities must always be chosen according to the rules for three dimensions, even when fewer components are actually used, because Eq. 14 was written for a spherically-symmetric velocity distribution. If the particle velocities are assigned in a plane instead of in a sphere, the outgoing flux will be greater than expected and the domain temperature will decline.

The boundary condition is run in an open loop configuration. In other words, the number and distribution of particles introduced to the domain during any given timestep are independent of the departing particles. The energy of departing phonons is recorded to calculate the net flux, then they are simply deleted from the calculation. Mazumder and Majumdar [25] point out that this boundary condition, considered as a whole, is equivalent to treating the face as a blackbody held at the specified temperature.

A polarization, wavenumber, and direction must be chosen for each particle introduced at the isothermal boundary. Rather than selecting a polarization, Eq. 14 is evaluated for each polarization and the introductions are performed separately. Choosing a wavenumber for an

incoming particle is slightly different than choosing k for a particle in the domain interior. Once again, the phonon group velocity becomes involved. For example, TA phonons with $k = 1$ are most probable in the domain, but they will never be found crossing a boundary because they have zero translation speed. The proper distribution of incoming phonons is therefore formed by multiplying the Bose-Einstein distribution by the phonon group speed. Note that the resulting distribution is a function of polarization as well as frequency.

Once a wavenumber is chosen, the group speed can be calculated from Eq. 11 but a direction must be chosen to determine the velocity. In the plane parallel to the boundary face, all directions are assumed equally probable. A single angle, θ is therefore chosen from a uniform distribution between 0 and 2π to determine the relation between the tangential components. In the plane normal to the boundary face, it was observed that the normal velocity component, normalized by the speed, has a linear distribution by collecting statistics on particles leaving a domain held at equilibrium. This type of distribution is quite easily sampled by taking the square root of a uniform random fraction. The two angles required to set the velocity of an incoming particle are therefore chosen according to:

$$\phi = sin^{-1}\left(\frac{\sqrt{R_f}}{v_g}\right) \tag{15}$$

$$\theta = 2\pi R_f \tag{16}$$

where ϕ and θ measure the angles between the velocity vector and the surface normal and an arbitrary reference in the surface plane, respectively.

In order to test the boundary condition formulation, a domain was initialized at a uniform temperature equal to the boundary temperature. A single timestep was performed, during which the identity and velocity of all incoming and outgoing particles was output in tabular format. This operation was repeated until the table contained more than a million entries. Because the boundary is intended to mimic an interface to a region in equilibrium at the specified temperature, the outgoing particles should match the incoming particles in both number and distribution during this test. Creating unnormalized histograms of the various parameters and comparing the incoming and outgoing distributions therefore provides a sensitive test of all components in the formulation. Such histograms are presented in Fig. 21. In all cases, the agreement is excellent.

Each phonon that enters or leaves the domain through an isothermal boundary during a calculation has its energy tallied. At output, each boundary reports its net flux per unit time per unit area. This quantity is not only a useful result, it presents a means for determining if a calculation has reached steady state. In this work, heat is introduced to, and removed from, the domain only through isothermal boundaries. The calculation may thus be considered steady when the sum of all boundary fluxes is zero. Due to the nature of the

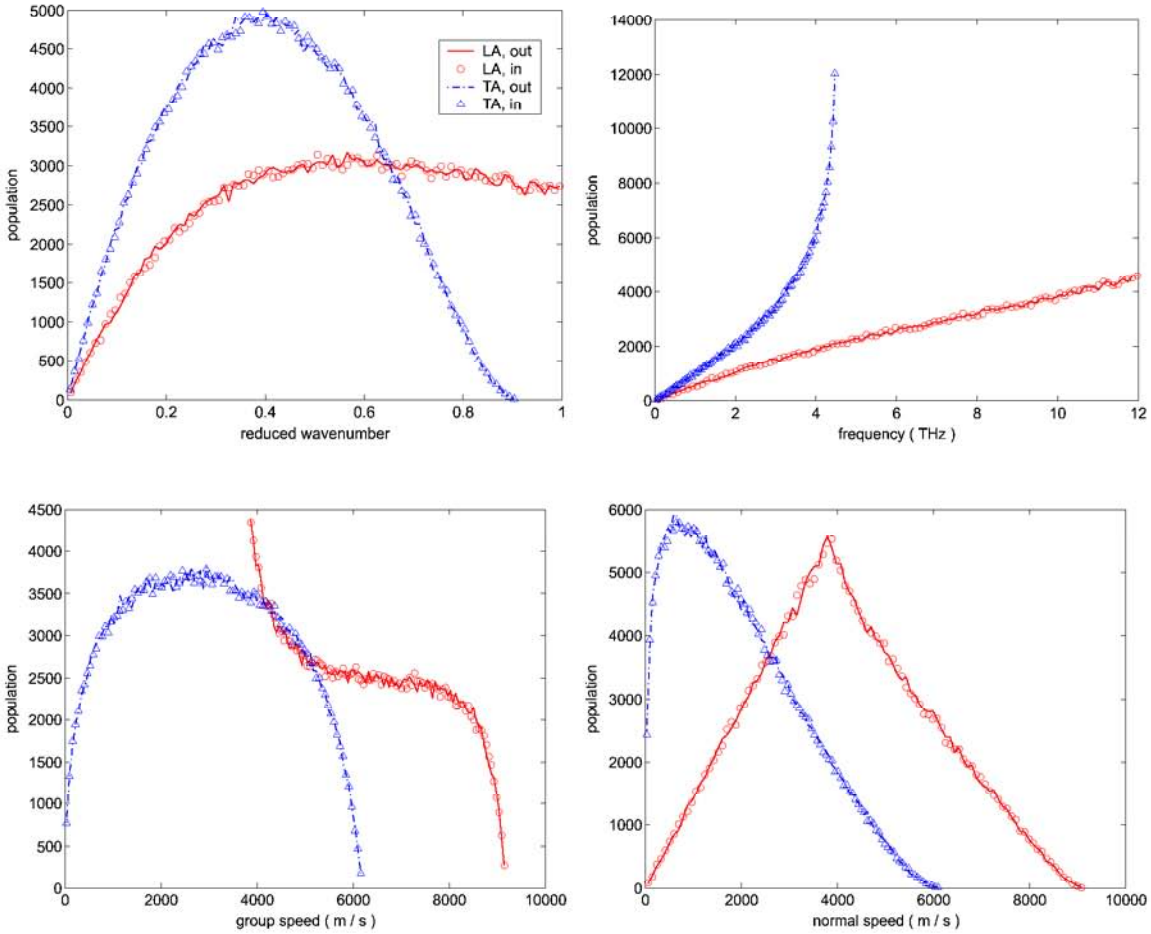

Figure 21. Histograms of incoming and outgoing particle characteristics at flux boundaries with domain and boundary initialized to the same temperature.

calculation, the cutoff below which a quantity is considered zero depends on the statistical convergence of the sample.

Adiabatic Boundaries

Adiabatic boundaries are much more straightforward than their isothermal counterparts. In order to balance the incoming and outgoing energy fluxes exactly, the phonon polarization and frequency are not changed. This leaves only the velocity vector direction open to modification. Two choices are available for choosing the outgoing direction: specular and diffuse reflection. In specular reflection, the velocity component normal to the face is reversed and all remaining components are untouched. This type of reflection represents a symmetry plane and preserves momentum except in the normal direction. In diffuse reflection, the outgoing velocity is chosen using the method described in the previous section without regard to the incoming velocity. This type of reflection represents a fully rough surface in thermal equilibrium with the adjacent region and destroys momentum in all directions. Choosing between these options on a particle-by-particle basis with a specified probability yields a partially-accommodating surface, with the probability of diffuse reflection often referred to as the accommodation coefficient. All these options are available in MOCAPHTS.

4.1.9 Bulk Scattering Routines

While the entries in both this section and the previous section affect phonon count, identity, and direction, often with identical selection mechanisms, the members of this section do so without regard to the particle position relative to the domain geometric features. They will thus be referred to collectively as "bulk scattering" mechanisms because they determine the bulk thermal conductivity, or the thermal conductivity in a sample of infinite extent.

It may be noted that, strictly speaking, impurities and grain boundaries do not quite fall into this category because they are geometrically fixed in the physical case and their distribution may change based on position in the sample. As implemented, however, these scattering mechanisms are statistically based on individual particles, rather than on fixed scattering centers, so they are included here.

Interphonon Scattering

Interphonon scattering is often referred to as anharmonic interaction because it occurs due to third and higher order terms in the lattice potential energy. In simple terms, if the lattice was connected by strictly linear springs, where the potential energy would contain terms

only second order in the displacement, phonons could pass without knowledge of each other's presence. If the springs are not linear, however, the stiffness seen by a passing phonon will change if another phonon is present and has thus already stretched the spring. Interphonon scattering generally becomes more prominent as the temperature increases and there are greater numbers of increasingly-energetic phonons on the lattice. The latter consideration is significant because higher energy phonons do not always conserve crystal momentum in interactions (see nearly any solid state physics text, such as [22], for details). Such interactions are critical to establishing thermal resistance in bulk materials.

Except where noted, the interphonon scattering scheme used in MOCAPHTS is very similar to that described in Mazumder and Majumdar [25]. This scheme, traceable to the analytical BGK scheme of Bhatnagar, Gross, and Krook [34], is a two step process. In the first step, phonons are selected for deletion based on a relaxation time, τ, which relates to the rate at which the relevant segment of the distribution function would return to equilibrium after a small perturbation. This relaxation time is generally a function of the local temperature and the phonon polarization and frequency. In the second step, phonons are introduced until the energy returns to its level before collisions commenced. This process will drive the distribution toward equilibrium, including allowing phonons to switch polarizations.

The collision process is run on a cell-by-cell basis. First, the total energy in the cell is calculated and the corresponding temperature determined. Each phonon is then considered for relaxation. The probability of the phonon scattering during the timestep, Δt, is calculated via:

$$\wp = 1 - exp\left[-\frac{\Delta t}{\tau(T, p, \omega)}\right]. \tag{17}$$

A random fraction is then drawn and the phonon is deleted if the result is smaller than $\wp$.

After each phonon in the cell has been considered for deletion, the addition process commences. In this process, phonons are created from the stationary equilibrium distribution at the cell temperature. This creation process is identical to that used to start the calculation except positions are assigned by copying coordinates from randomly-chosen particles in the cell. This "coordinate copying" was implemented to avoid smearing number density gradients that may exist in the cell, which persist to arbitrarily small cells in the case of linear gradients.

It is important to note that the process described thus far violates detailed balance, *i.e.* the outgoing and incoming fluxes at equilibrium will not match when the relaxation times are frequency dependent (which is true of all proposals in the literature). If the phonons taken out of the cell are selected based on a frequency-dependent relaxation time, the probability of deleting a particular phonon will be the product of its probability of existing

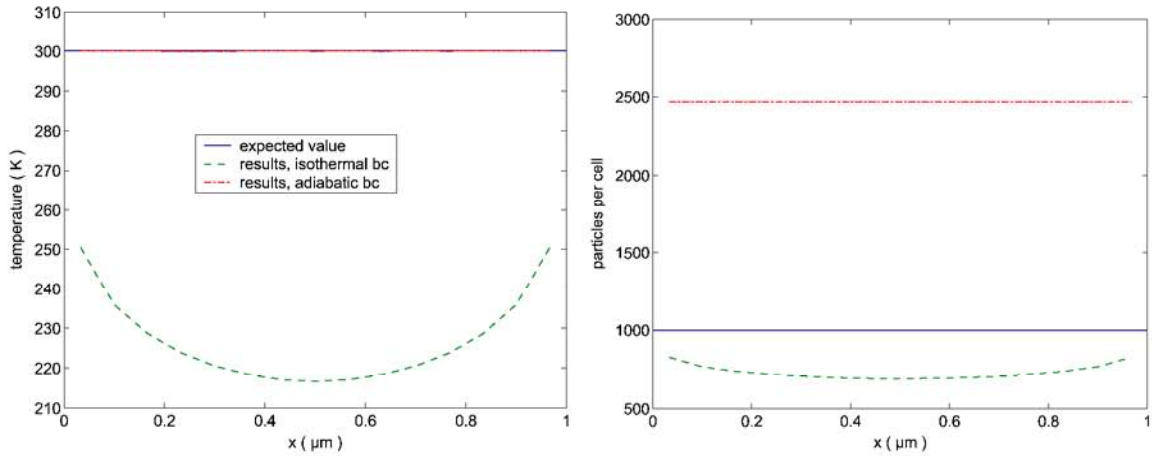

Figure 22. Consequences of violating detailed balance by using frequency-dependent relaxation times on only the outgoing step. The observed behavior depends on whether isothermal or adiabatic boundary conditions are applied.

(from the Bose Einstein distribution) times the frequency-dependent probability given in Eq. 17. The probability of introducing a particular phonon will be simply drawn from the Bose Einstein distribution. For frequency-independent relaxation times, this process works because the probability of all outgoing states are scaled evenly. Adding frequency dependence, however, produces a variable scaling that changes the shape of the distribution. Because the relaxation times are generally smaller as frequency increases, this implies that the algorithm is likely to shift the distribution toward lower frequencies. For closed domains, this causes the phonon count to increase because more phonons are required to reach the pre-collision energy. For domains with open-loop isothermal boundaries, such as those described in Sec. 4.1.8, this causes the phonon count to decrease because the favored phonons tend to have larger group speeds, thus more leave the domain than expected. Examples of the temperature and number density distributions for closed and open 1D domains held at 300K are presented in Fig. 22 compared to the initial condition. In these cases, the relaxation times were set to:

$$\begin{aligned} \tau_L^{-1} &= 5 \times 10^{-24}\, \omega^2 \\ \tau_T^{-1} &= 4 \times 10^{-23}\, \omega^2 \end{aligned}$$

where the coefficients were set such that the maximum probability of each mode is roughly 0.1.

To satisfy detailed balance with frequency-dependent relaxation times, the acceptance-

rejection method was again applied. In this case, a phonon polarization and frequency is chosen in the usual manner, then the relaxation time is calculated using the same expressions as for outgoing phonons. This relaxation time is subsequently used to compute an introduction probability using Eq. 17, although it is scaled by $1/\wp(\tau_{min})$ in this case to set the maximum probability at unity to reduce the rejection rate.

The stopping criterion for the addition process is identical to that used in the domain initialization process described in Sec. 4.1.5. In this case, the difference between the starting and ending energies is stored and added to the starting energy in the next timestep. The process therefore tolerates small changes in the total energy in a given timestep, but the following timestep will work to make up the difference. This scheme is more likely to avoid random walks than simply conserving energy in an average sense over a series of timesteps.

It may be noted that Mazumder and Majumdar [25] use an addition scheme in which phonons are created and destroyed until the energy change falls within a fixed tolerance. The current scheme provides two improvements to this arrangement. First, it is clearly faster because it destroys, at most, only the last created phonon. Second, it avoids the danger of skewing the equilibrium distribution by hunting around for a set of phonons that satisfy the energy constraint, rather than for those that are most probable in the equilibrium distribution.

Several authors [35, 36, 37] have derived expressions for calculating the relaxation times required by this scheme. In this work, the arrangement employed by Holland [38] is adopted. In this arrangement, the relaxation times for LA phonons are calculated according to:

$$\tau_L = B_L \omega^2 T^3. \tag{18}$$

and the relaxation times for TA phonons are calculated using two different expressions:

$$\tau_T = B_T \omega T^4 \tag{19}$$

for low frequencies and

$$\tau_{Tu} = B_{Tu} \omega^2 / sinh\left(\frac{\hbar\omega}{\mathcal{B}T}\right) \tag{20}$$

for high frequencies, where B_L, B_T, and B_{Tu} are empirical constants used to tune the scattering strength. The crossover point from τ_T to τ_{Tu} was set at k=0.5.

Figure 23 shows the result obtained via the method and constants given by Holland, which matches very closely to Fig. 3 in the original paper. In this figure, the contribution to the thermal conductivity, κ, due to LA phonons is shown as κ_L, and the contributions due to low and high frequency TA phonons are shown as κ_T and κ_{Tu}, respectively. From these

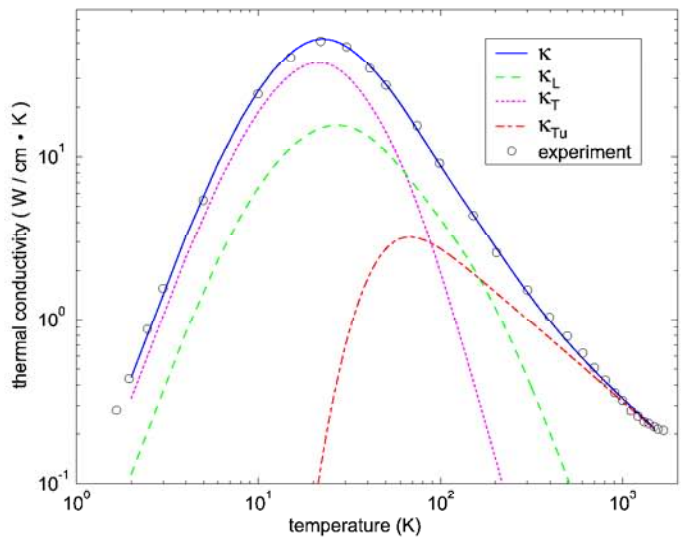

Figure 23. Holland result for thermal conductivity of a 0.7 *cm* silicon sample. The total thermal conductivity is shown as κ and the contributions due to LA, low frequency TA, and high frequency TA phonons are shown as κ_L, κ_T, and κ_{Tu}, respectively

results, Holland concluded that high frequency TA phonons dominate thermal transport at room temperature. [1]

Using Holland's coefficients in the current code, bulk thermal conductivities of silicon were seen to be smaller than the experimental measurements by approximately a factor of two in the temperature range 100-300K. Closer inspection revealed two key assumptions inherent to Holland's model which are not present in the Monte Carlo model: piecewise constant group speeds and group/phase speed interchangeability.

The piecewise constant phonon group speed used by Holland is compared to the output of the dispersion model used in the current work in Fig. 24. This assumption represents a compromise between using realistic dispersion relation and producing an expression for thermal conductivity that can be evaluated analytically.

The group/phase speed interchangeability assumption in Holland's model was pointed out by Sood and Roy [37]. Due to dispersion, the group and phase speeds can differ substantially. A comparison of these speeds under the dispersion model used in this work is

[1] It should be noted that it was necessary to reduce the impurity scattering coefficient given by Holland by an order of magnitude to produce this figure.

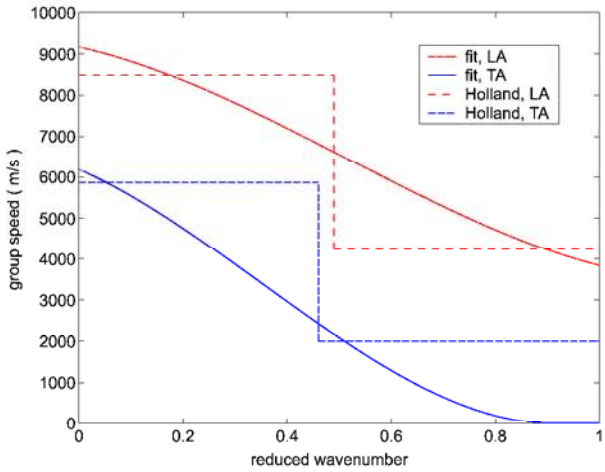

Figure 24. Comparison of Holland's phonon speeds to those from the current dispersion model.

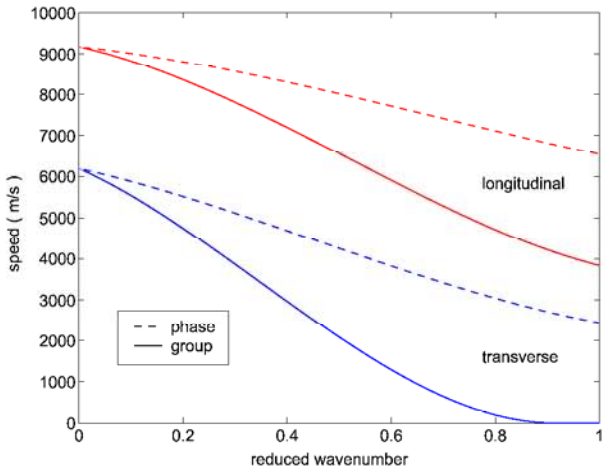

Figure 25. Comparison of phase and group speeds under the current dispersion model.

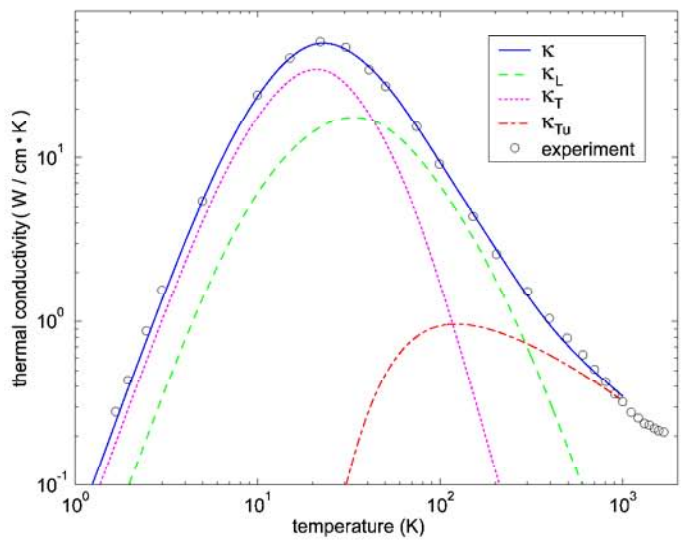

Figure 26. Modified Holland model result for thermal conductivity of a 0.7 *cm* silicon sample. The overall fit is comparable to that of the unmodified model (Fig. 23), but the contributions from individual phonon modes are significantly different.

presented in Fig. 25.

Removing these assumptions, as well as converting the integrals to wavenumber space as outlined in Sec. 3.2, a new set of coefficients for Eqs. 18-20 was derived using the least-squares optimization routine in MATLAB®:

$$B_L = 8.6246 \times 10^{-25} \ s/K^3 \tag{21}$$

$$B_T = 1.2449 \times 10^{-12} \ 1/K^4 \tag{22}$$

$$B_{Tu} = 7.1424 \times 10^{-19} \ s \tag{23}$$

The resulting fit is shown in Fig. 26. Again, as in Fig. 23, the agreement is good throughout the temperature range. Notably, however, the contributions of the various modes are now different. For instance, TA phonons are clearly dominant at 300K in Fig. 23

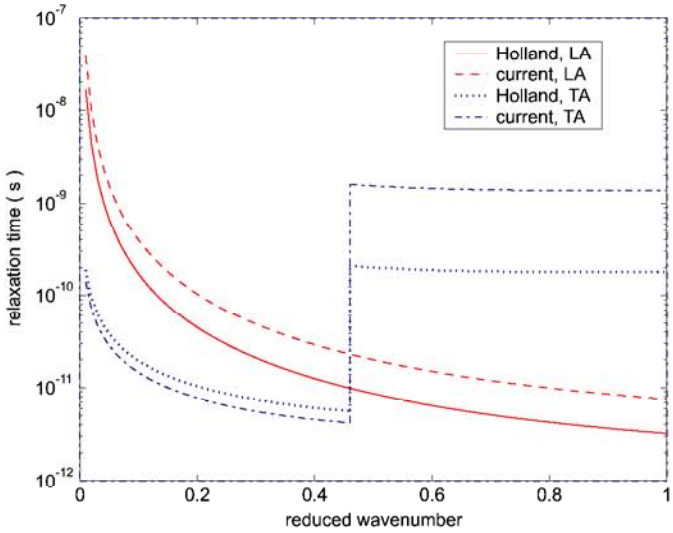

Figure 27. Relaxation times for both the original and the modified Holland model. In both cases, more than an order-of-magnitude jump occurs when switching from the low to the high frequency expression for TA phonons.

whereas TA and LA phonons have comparable strength in Fig. 26. It is not reassuring that one can change assumptions in the model, then choose a set of new coefficients to produce a comparably good fit to the overall thermal conductivity that leads to different conclusions about the underlying physical processes.

Another feature of the model that raises doubt about its physical basis is shown in Fig. 27, where it can be seen that the relaxation time increases by an order of magnitude at the low-to-high frequency transition point for TA phonons. This observation is true for both Holland's original coefficients as well as the new set.

These observations led to an attempt to develop an interphonon collision scheme as part of this work. This attempt will be detailed in Sec. 4.8.

Impurity Scattering

Impurities in the crystal lattice cause phonon scattering through two primary mechanisms: change in the atomic mass, and changes in the local spring constant due to lattice distortion. In general, the spatial extent of the disturbance is only a few atomic volumes. Thus, for the vast majority of phonons, particularly those with large enough group velocities to transport

significant energy, the disturbance may reasonably be considered a point defect. As such, it is common to construct a scattering model based on the arguments made by Rayleigh [39] when computing the scattering of sound waves by a small obstruction.

In this work, as in that of most previous researchers, only the mass difference will be considered. For that case, Berman [40] gives the relaxation time for impurities as:

$$\tau_I = \frac{a^3 k^4}{4\pi} v_g \sum_i c_{p_i} \left[\frac{\Delta M_i}{M} \right]^2 \tag{24}$$

where c_p is the fraction of lattice sites with impurities whose mass differs from the reference mass M by ΔM. The sum is over the individual impurities.

As an example, consider "natural" silicon, which is typically composed of 92.23% ^{28}Si, 4.67% ^{29}Si and 3.10% ^{30}Si. Setting the reference mass as 28 a.m.u., the sum becomes:

$$0.0467 \left(\frac{29-28}{28} \right)^2 + 0.0310 \left(\frac{30-28}{28} \right)^2 = 2.1663 \times 10^{-4}. \tag{25}$$

It may be noted that Ruf *et al.* [41] use the average mass as the reference and compute a sum of 2.01×10^{-4}. Our choice of $M = 28$ a.m.u. is based on an argument that it is more reasonable to base the scattering coefficient upon the mass seen most often by the traveling wavepackets, rather than an imaginary mass never encountered. In the end, the difference between these approaches is small, particularly in view of the other approximations involved.

Preliminary calculations to test this model showed thermal conductivities that were much smaller than expected. Manually reducing the mass difference sum calculated in Eq. 25 by an order of magnitude, as was required to reproduce Holland's results, brought the results into agreement. The reduced coefficient is assumed for the remainder of this report. The physical reason for this reduction should be investigated in future work.

4.1.10 Parallelization

As the scope of the simulations undertaken for this work increased and computation times began to exceed a week, it became clear that a parallel version of the code would be necessary. The parallelization scheme grew directly from the decomposition of the domain into zones (see Sec. 4.1.7). Processor decomposition simply involved assigning zones to different processors, dividing user-defined zones if necessary. Parallel commands for transferring particles between zones and handling communications for global operations were written

using MPI [42]. Each processor has knowledge of all zones, but only the cell structure of the local zones. The cell finding step is thus performed after the particle is transferred to its destination processor. This arrangement combines the scalability advantages of each processor independently determining the final location of all its particles [43] with the memory advantages of making the cell layout and data strictly local.

4.2 Typical Simulation Results

In the current work, we are generally interested in determining the thermal conductivity under various conditions. The typical calculation therefore involves creating a temperature gradient across the domain and examining the results once it reaches steady state. The type of data available from the simulation include the distributions of drift (or mean) velocity, temperature, number density, and collision fraction. The mean flux through each of the isothermal faces is also reported. A demonstration of these quantities from a typical calculation to determine the thermal conductivity of a 1 μm film at 300K is shown in Fig. 28.

From this figure, a similarity can be noted between the flow of phonons in temperature gradient along a film and the flow of gas molecules in pressure gradient along a channel (known as Poiseuille flow). This similarity is made more evident by the comparison shown in Fig. 29.

For the gas molecule case, the pressure is constant in the cross-channel direction and nearly linear in the axial direction. The axial speed assumes a parabolic distribution, which exhibits an increasingly large velocity at the walls as the pressure is decreased. Due to the reduction in density as the pressure falls (the flow is essentially isothermal), the maximum speed also increases with axial distance to maintain a constant mass flow.

For the phonon case, the temperature serves as the pressure analog, which assumes a nearly linear distribution (though the concavity is reversed) that is constant in the cross-channel direction. The temperature gradient, like pressure gradient in the Poiseuille flow, provides the driving force and a parabolic axial flow distribution is established, which again shows a nonzero speed at the walls. The mean phonon speed in this case increases with axial position because the Bose Einstein distribution shifts to lower frequencies at lower temperatures. As shown in Fig. 17, the group speed increases as frequency decreases.

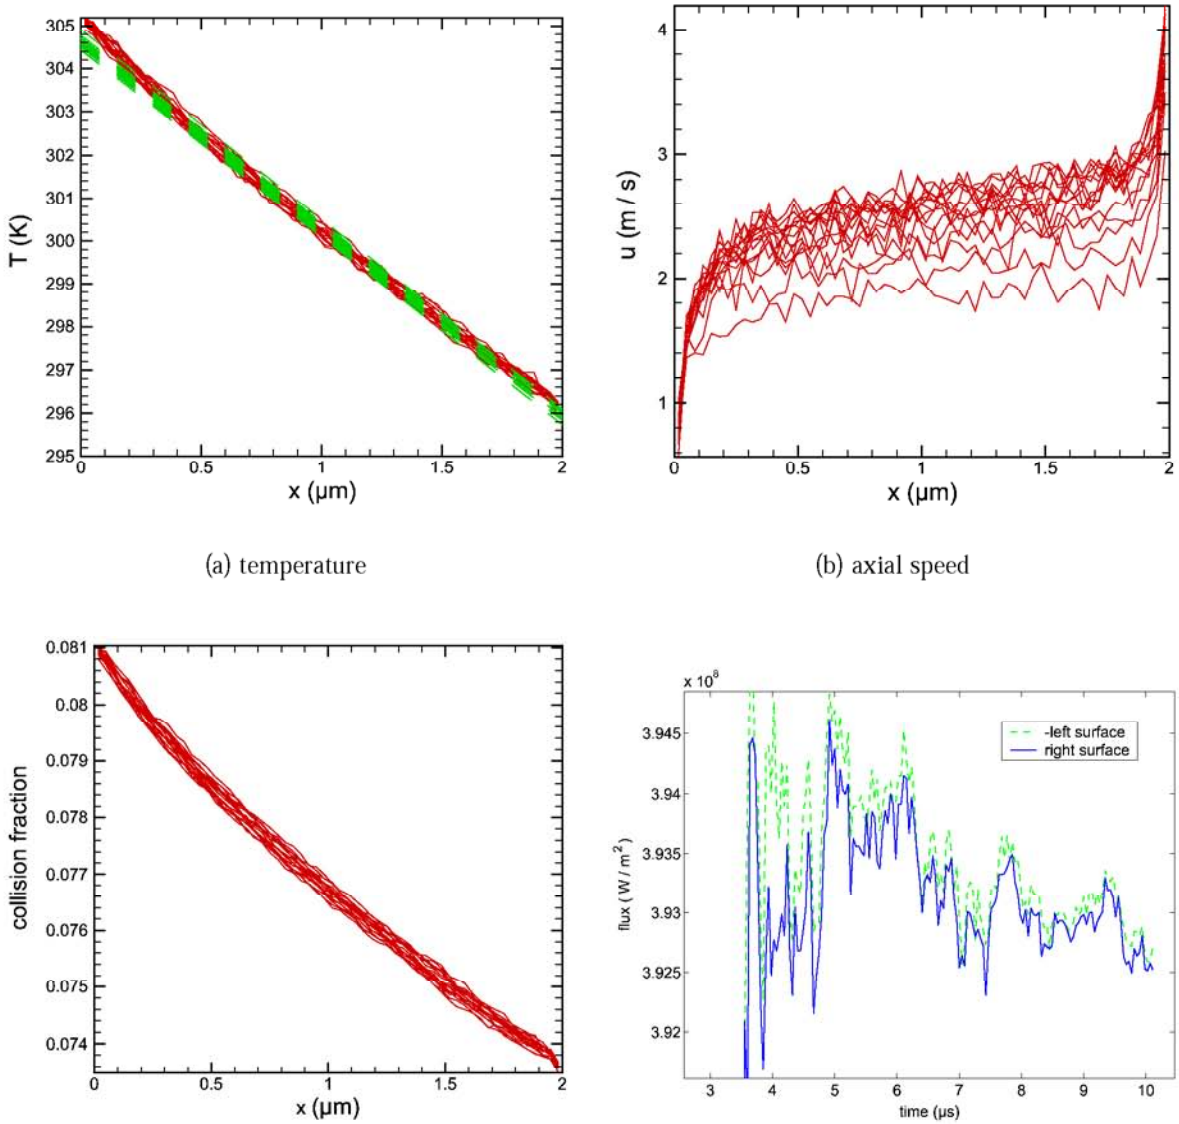

(a) temperature

(b) axial speed

(c) collision fraction

(d) flux at domain ends

Figure 28. Typical simulation results. Case shown is a 1x2 μm film at a reference temperature of 300K. The dashed lines on the temperature plot are the linear fit used for determining thermal conductivity. Each line in plots (a)-(c) represents cross section at a different position on the y axis, which is perpendicular to the temperature gradient.

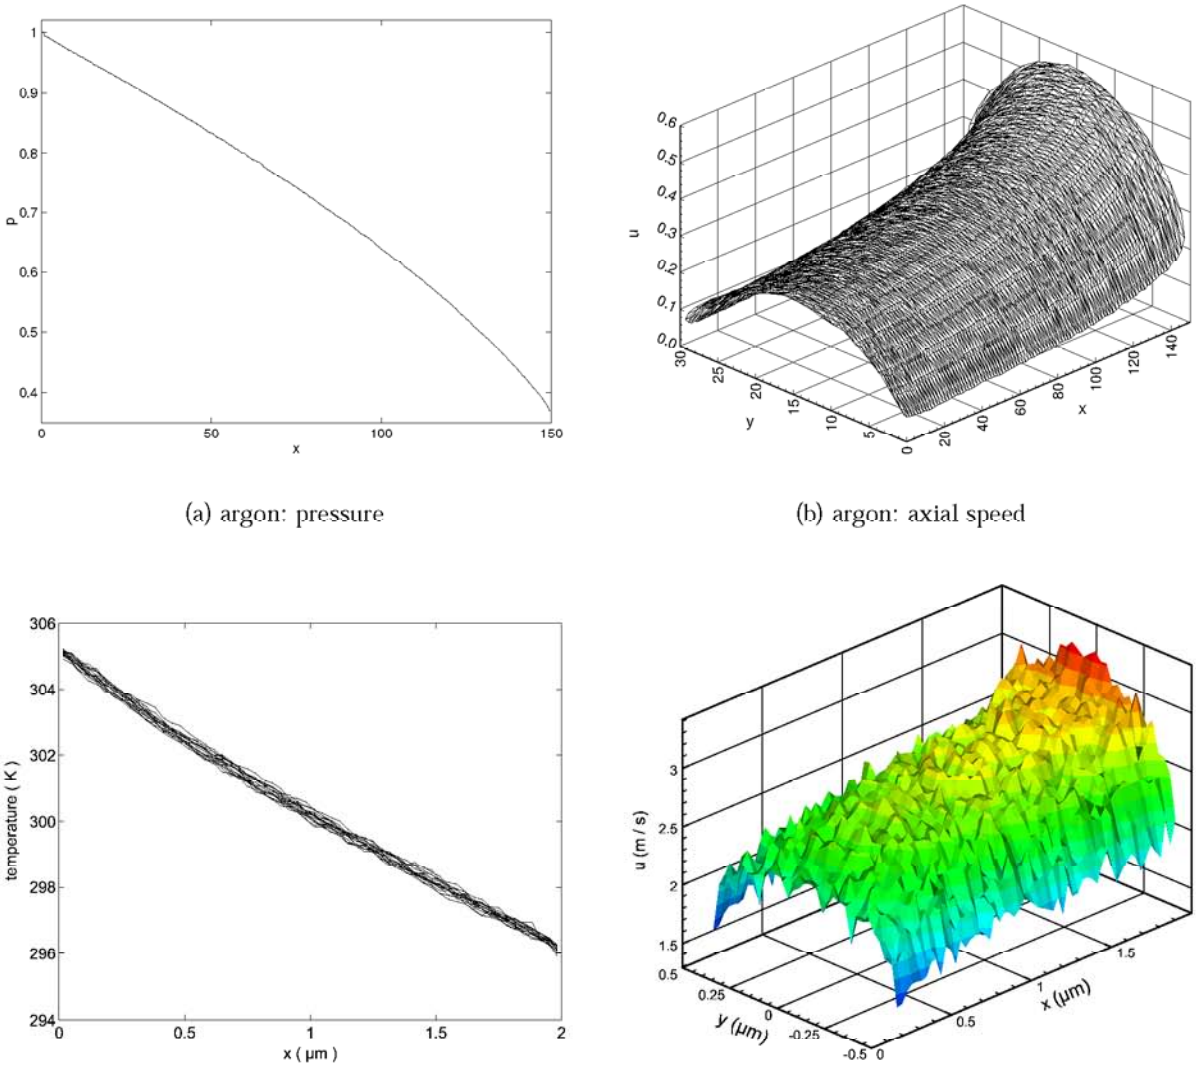

(a) argon: pressure

(b) argon: axial speed

(c) phonon: temperature

(d) phonon: axial speed

Figure 29. Illustration of similarities between argon flow in a pressure gradient (Poiseuille flow) [44] and phonon flow in a temperature gradient. Speed, length and pressure in the argon flow are normalized by the most probable molecular speed, the mean free path, and the pressure at the inlet, respectively. The pressure in the argon flow was averaged across the channel.

4.3 Thermal Conductivity Determination

Because the current work primarily focuses on regimes in which the concept of thermal conductivity is still useful, a means must be defined for defining this quantity from the simulation results. In general, for simulations aimed at determining thermal conductivity, a temperature gradient is established across the domain by imposing isothermal boundary conditions on opposing faces. At the end of the calculation, the net flux through these faces and the temperature distribution are examined.

For reasons outlined in Sec. 4.1.8, the resulting temperature distribution tends to have nonlinear regions near the ends. These areas are thus excluded and a linear curve fit is performed on the remaining data near the domain center, as shown in Fig. 28(a). This curve fit is then evaluated at each of the domain ends. Dividing the difference between these values by the domain length provides an estimate of the temperature gradient. Dividing the calculated flux by the temperature gradient yields the quantity that will be reported in this document as thermal conductivity. By measuring the resulting, rather than using the imposed, temperature gradient, this method will yield the analytically-predicted infinite thermal conductivity in the ballistic regime because the measured temperature difference across the domain will tend to zero.

4.4 Operational Considerations

4.4.1 Number of Particles

The number of particles per cell is a measure of the discretization of the distribution function. This is a somewhat inexact measure, however, due to their stochastic placement in phase space. The "required" number of particles is generally determined via an empirical process similar to grid convergence studies used in CFD: perform a series of calculations on a representative problem and look for an asymptote in the solution versus the number or particles. In DSMC, 30 particles per cell is generally considered a reasonable starting point. While the thermal conductivity was somewhat insensitive to particle counts (above a threshold of roughly 100) the temperature and drift speed measured near the isothermal boundaries tended to show increasingly unsettling behavior as the particle counts were reduced. In particular, for sufficiently small particle counts, the temperature at the higher temperature isothermal boundary exceeded the specified value and the drift speed plunged suddenly to a negative value, though the flux maintained the expected sign. These "end effects" are visible, particularly in the axial drift speed, in Fig. 28. In order to minimize these effects, a target of 500 particles per cell was employed in most of the results presented in

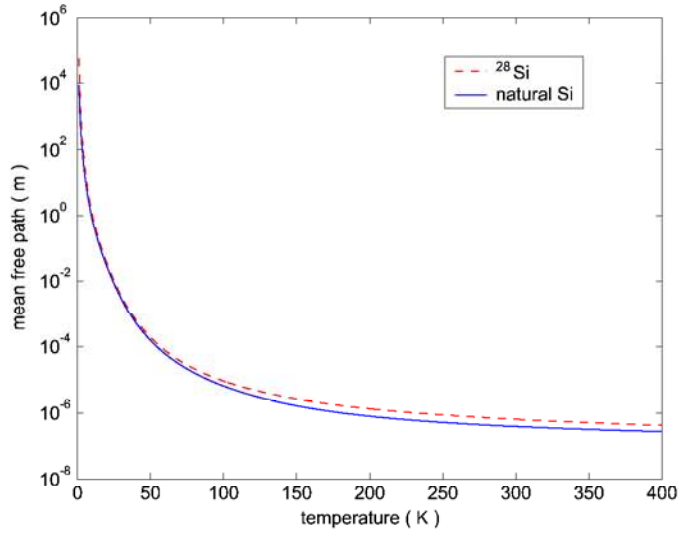

Figure 30. Mean free path for single crystal silicon under the relaxation time model

this report.

4.4.2 Cell Count

Because collisions are performed without regard to a particle's position in the cell, the cell should be small compared to the scale length of any gradients present. In practice, the cell is generally sized such that it is a fraction (1/4 to 1/3) of the local mean free path. When sized in this manner, approximately 10% of the particles in the cell will collide during any given timestep. Under the relaxation time model, it is trivial to calculate the mean free path by averaging the group speed times the relaxation time over the Bose Einstein distribution:

$$\lambda = \sum_p \frac{4\pi}{a^3} \int_0^1 n_p \tau_p v_{g_p} k^2 \, dk \tag{26}$$

The results of this calculation are presented in Fig. 30 for isotope-free ^{28}Si using the collision coefficients given in Sec. 4.1.9. Analogous results for "natural" silicon are also presented.

Because the mean free path becomes significantly longer at lower temperatures, a larger number of cells than required by the above constraint was often used to place a reasonable

number of points on the resulting plots.

4.4.3 Domain Length

Due to the assumptions inherent to the isothermal boundary condition used in the calculations (see Sec. 4.1.8), it was noted that the domain length in the direction of the temperature gradient can affect the calculated thermal conductivity. This effect generally disappeared as the domain length increased. This is similar to the "entrance effects" seen in fluid flow through pipes [45]. As the domain is lengthened, the fraction influenced by these entrance effects decreases and the thermal conductivity reaches an asymptote. This asymptote will be considered the correct value.

A demonstration of this behavior is shown in Fig. 31 for bulk and film thermal conductivities at 100K. The effect tends to be more severe at lower temperatures due to the higher drift speeds and lower interphonon collision rates as the temperature decreases. It also tends to be more severe for the bulk cases because the adiabatic walls present in the thin film cases help reduce the extent of the end effect. This is similar to the case of pipe flows, where the entrance region is usually given as some number of pipe diameters. The domain length required to reach the asymptote therefore tends to be smaller for thinner films.

4.5 Bulk Thermal Conductivity

The thermal conductivity of bulk silicon was calculated following the procedures outlined in the previous sections. To simulate a sample of infinite extent, these simulations were performed with MOCAPHTS compiled for one-dimensional domains. For all cases, the ends were held at the reference temperature plus or minus $5K$. The results of these simulations are presented in Fig. 32 for both isotope-free and "natural" silicon.

For comparison to the isotope-free simulations, experimental measurements of Ruf *et al.* [41] on isotopically-enriched silicon, or silicon that has been processed such that it contains 95.8588% ^{28}Si, are shown. For comparison to the "natural" silicon simulations, data from Holland and Neuringer [46] are shown. In both cases, the agreement between simulation and experiment is excellent, with a maximum error of approximately 14%.

Accurate reproduction of the bulk thermal conductivity by the simulation for isotope-free silicon implies a successful implementation of the interphonon collision model and correct choice of the coefficients. Agreement in the natural silicon case suggests similar

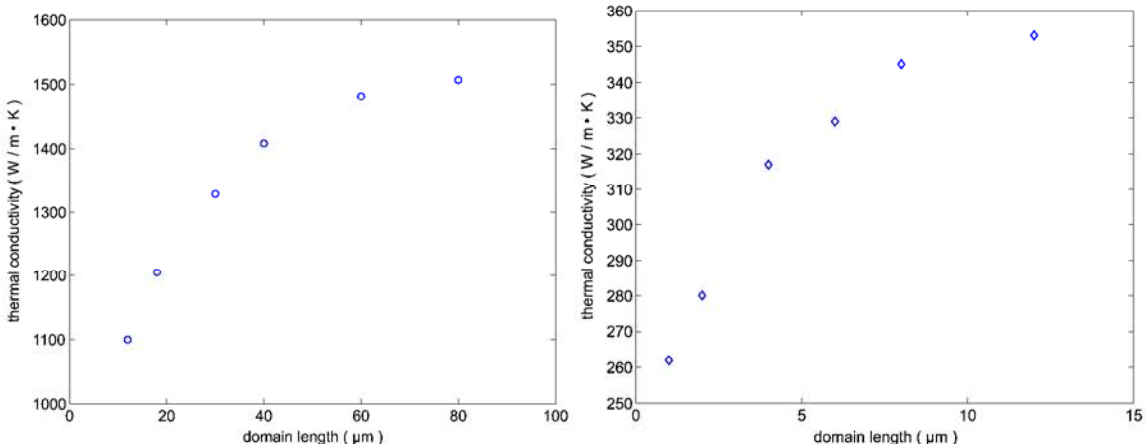

(a) bulk (b) $1\mu m$ film

Figure 31. Domain length effect on thermal conductivity results.

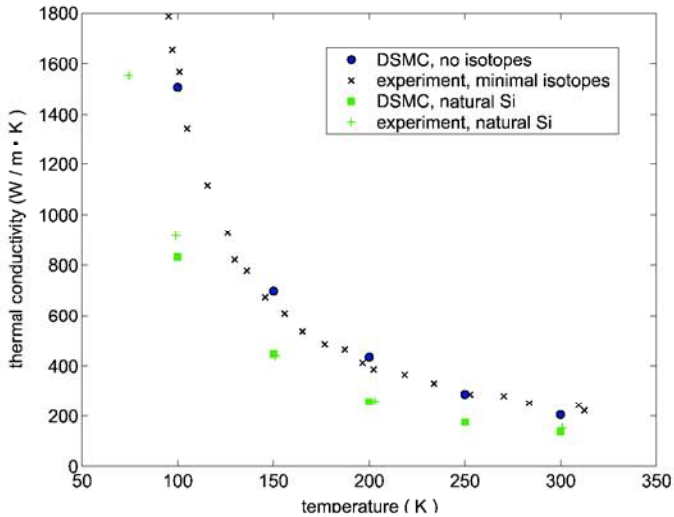

Figure 32. Computed thermal conductivity of bulk silicon with and without naturally-occurring isotopes compared to experimental data. Natural silicon data is from [46] and isotopically-enriched silicon data is from [41].

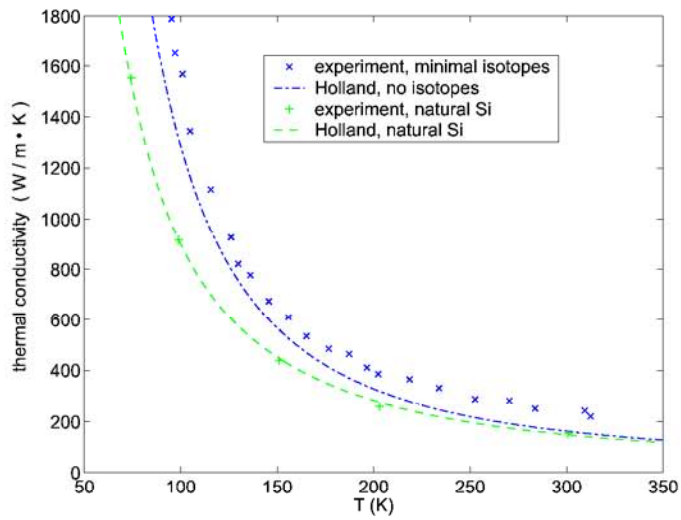

Figure 33. Computed thermal conductivity of bulk silicon with and without naturally-occurring isotopes using the Holland model.

conclusions for the impurity scattering model.

Because the phonon-phonon and impurity relaxation time expressions are identical to those used by Holland, except for the coefficients in the former case, it is instructive to prepare a figure similar to Fig. 32 using Holland's analytical model to evaluate the gains made via the added complexity and computer time required for the Monte Carlo model. In the resulting figure (Fig. 33), the analytical model can be seen to agree very well with the natural silicon data, which is expected because the coefficients are fit to this data. Setting the impurity scattering coefficient to zero in the analytical model, however, does not produce very good agreement with the isotopically-enriched silicon data. It may therefore be argued that the separate treatment of these collision mechanisms in the Monte Carlo simulation produces a more physically representative model. The inability of the Holland model to match the isotope-free data with the impurity scattering coefficient set to zero implies a "smearing" of effects, *i.e.* the impurity scattering model is contributing to the curve fit in a manner not physically associated with the presence of impurities.

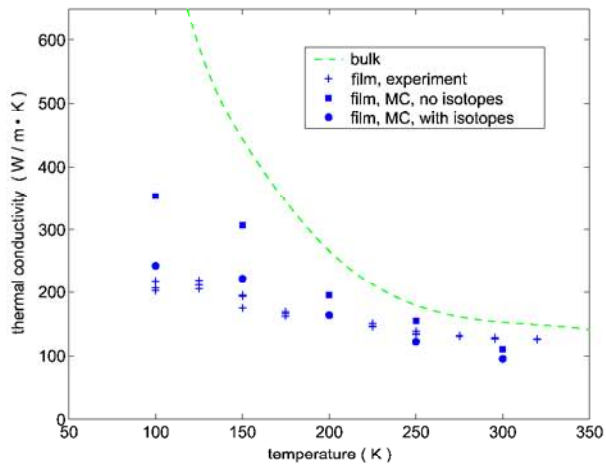

Figure 34. Computed thermal conductivity in a 1 μm single crystal silicon film, with and without isotopes, compared to bulk values and experimental results.

4.6 Thin Film Thermal Conductivity

Switching to two-dimensional simulations and introducing an adiabatic boundary on one of the additional faces and a specular boundary on the opposite face allows the simulation of heat conduction in a film of single crystal material twice as thick as the domain (as noted in Sec. 4.1.8, specular boundaries represent symmetry planes).

Simulation results for a 1 μm silicon film with and without impurity scattering are shown in Fig. 34 compared to experimental measurements taken in this work. The data for bulk silicon are also shown for reference. It is immediately evident that the simulations show the size effect seen in the experiments. It is also clear that scattering due to naturally-occurring isotopes plays a significant role in determining the thermal conductivity, particularly at low temperature, with κ overestimated by 69% at 100K when isotope scattering is neglected. With the isotope model running, the simulation reproduces the experimental data reasonably well. In this case, the simulation tends to overpredict thermal conductivity at low temperature and underpredict at high temperature. In either case, the maximum error is less than 25%.

To evaluate the predictive capability of the Monte Carlo model against the Holland model, the measured thermal conductivity of 1 μm and 340 nm films are compared to model

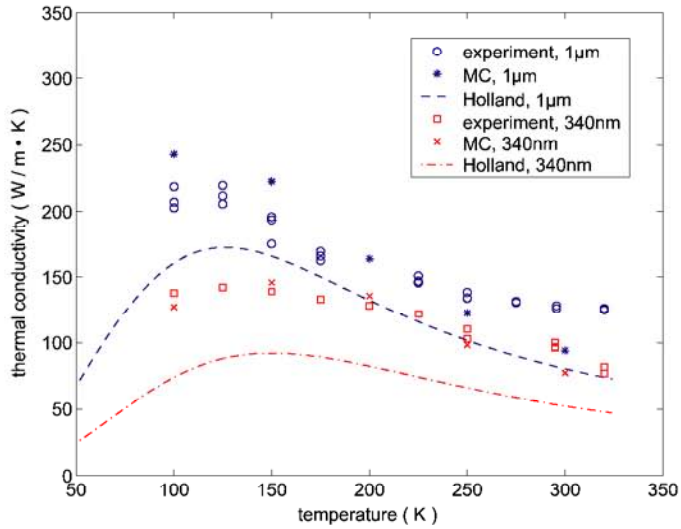

Figure 35. Thermal conductivity of 1 μm and 340 nm single crystal silicon films, computed with Monte Carlo and Holland models, compared to experimental results.

results in Fig. 35. The Monte Carlo model shows good agreement for both films without adjusting any parameters while the Holland model performs poorly in both cases. The Holland model was calibrated for a 0.716 cm film, thus Fig. 35 provides additional evidence that the relaxation times appearing in the model are performing curve fit functions outside of their supposed physical meaning. In the Monte Carlo method, however, the boundary scattering model is entirely separate from the interphonon and impurity scattering models, being invoked only when a computational particle physically intersects a boundary.

4.7 Grain Boundary Effects

Surface micromachined MEMS, such as those produced using Sandia's SUMMiT process [47], generally use polysilicon as a structural material. It is therefore important to develop a model for phonon scattering by grain boundaries because this can be a dominant factor in determining the thermal transport characteristics of the material.

The simplest grain boundary scattering model is based upon relaxation times calculated by dividing the mean grain size by the mean phonon speed. This model can be improved by calculating a relaxation time for each particle based on its group speed and generat-

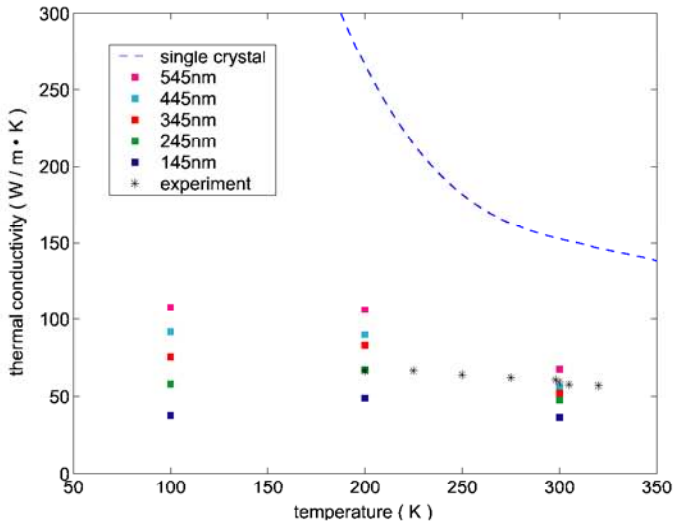

Figure 36. Simulation results for polysilicon films with varying grain size. Experimental data is for the $2\,\mu m$ film in Fig. 12(a), which has an estimated grain size of $580\,nm$.

ing grains stochastically according to a distribution. Building upon these improvements, columnar grain structures can be represented by making the relaxation time dependent on the phonon direction, speed, and position.

In the following calculations, only the phonon direction was changed if it was selected for relaxation due to grain boundary scattering. This procedure guarantees energy conservation in the current formulation but it cannot represent partial transmission/reflection at the boundary. Partial transmission, in addition to being expected intuitively, has been observed in molecular dynamics simulations of a boundary between two materials whose mass differs by a factor of two [48], so it should be considered in future work.

Using the relaxation time model with an average grain size, but a group speed based on each particle, the thermal conductivity for a range of grain sizes between 145-545 nm was calculated at 100-300K. No difference was observable between the 1D and 2D calculations for a $2\,\mu m$ film with $545\,nm$ grains at 100K, so the remaining calculations were performed in 1D. The results are presented in Fig. 36, along with single-crystal data and experimental measurements from the $2\,\mu m$ film in Fig. 12(a) (average grain size: $\sim 580\,nm$). In this case, the simulations were able to bracket the experimental measurements, but the simulations show a stronger temperature dependence for a given grain size than was observed experimentally.

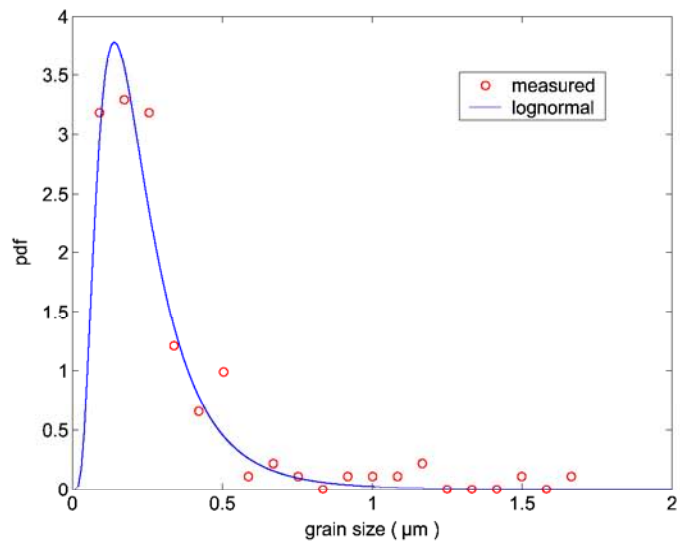

Figure 37. Comparison of measured grain size distribution with lognormal pdf of Eq. 27 with $\sigma = 0.62\,\mu m^{-1}$ and $\mu = -1.58\,\mu m$.

It can be noted from Fig. 36 that the temperature dependence in the simulations becomes weaker as the grain size decreases. Noting also that a wide range of grain sizes are visible in Fig. 12(a), it was surmised that the discrepancy between simulation and experiment may be narrowed by increasing the number of parameters in the model by using a grain size distribution, rather than a simple average.

Palmer *et al.* [49] investigated possible distribution functions for grain size in germanium and concluded that a lognormal distribution provides a good overall description. The lognormal distribution of grain diameters, d:

$$\wp = \frac{1}{\sigma d\sqrt{2\pi}} e^{-(ln(d)-\mu)^2/2\sigma^2} \tag{27}$$

with $\sigma = 620\,nm^{-1}$ and $\mu = 1580\,nm$ appears to fit a grain size distribution obtained from Fig. 12(a) by measuring the distance between grain boundary intersections on a series of horizontal lines. A comparison of the lognormal fit to the measured distribution is shown in Fig. 37. The constants required for this fit imply a mean of $251\,nm$ and a variance of $172\,nm$.

Simulations were then performed in which a grain size was chosen according to Eq. 27 for each cell during each timestep for calculating the grain boundary scattering relaxation time. The result is compared to the previous, constant grain size, calculations in Fig. 38.

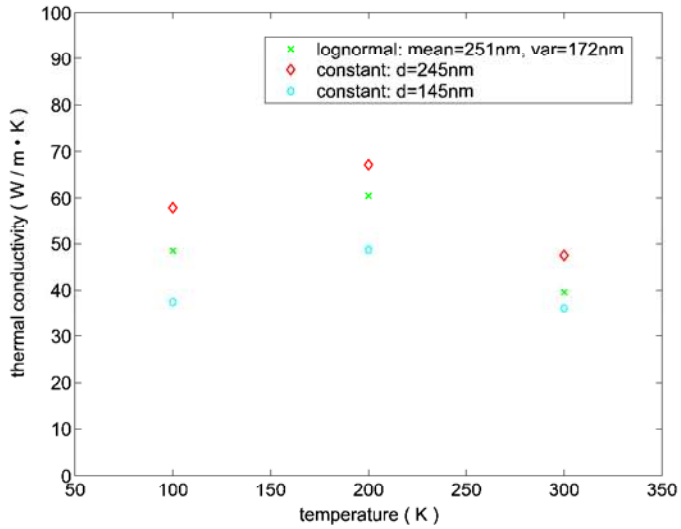

Figure 38. Comparison of polysilicon thermal conductivities obtained using lognormally-distributed and constant grain sizes.

Not surprisingly, simulations performed using a lognormal grain size distribution yield a lower thermal conductivity than the constant grain size case using a comparable mean. Unfortunately, the temperature variation in the lognormal case is similar to that of the constant size case, so this does not appear to be the dominant cause of the disagreement visible in Fig. 36.

Departing altogether from relaxation time-based models, which effectively assign a new grain every timestep, a third scheme for modeling polysilicon was attempted in this work. This scheme involves assigning a grain size to each particle. When the particle intersects its virtual grain boundary, a scattering operation is performed. If the particle's direction with respect to the surface normal of the intersected face still has the same sign after scattering, it is considered to have passed out of its current grain and a new grain is generated from the distribution. This scheme allows grains of various shapes to be implemented, though cubic grains were assumed for the current work.

The results of these calculations are presented in Fig. 39. In this case, the temperature departs significantly from the specified reference and the thermal conductivity is reported at the resulting mean. The increase in mean temperature on the domain indicates an interaction between the grain scattering model and the isothermal boundary model. In particular, it implies a significant disruption of the distribution function. Unfortunately, time did not permit a detailed investigation of this issue.

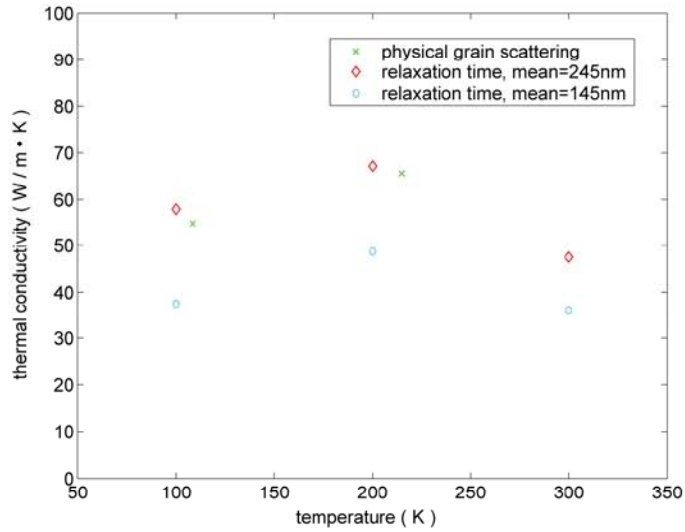

Figure 39. Comparison of polysilicon thermal conductivities obtained using physical lognormal and relaxation time-based fixed mean grain scattering.

The results of this model are reminiscent of those obtained with the previous models proposed in this section and, consequently, the disagreement between the simulation and the experiments remains an open question. Evidence presented in Sec. 5.4 implies that the issue is likely related to defects at the grain boundary. A systematic study of this issue is recommended for future work.

4.8 Non-BGK Interphonon Scattering

Due to the deficiencies of the single-mode relaxation time (SMRT) based collision scheme discussed in Sec. 4.1.9, as well as its inherent limitation to small departures from equilibrium, a more general scheme is desirable to improve the predictive ability of the simulation over a wider range of conditions. Toward this end, this section details efforts to construct a collision scheme that is, ideally, dependent only on the individual computational particles present in a cell. This scheme was modeled after that used in DSMC for gases.

In DSMC the basic collision algorithm is: choose two particles in the cell, compute a collision probability based on their properties (such as species identity and relative velocity), and perform a collision if a random number less than the collision probability is

chosen. If a collision is to be performed, the post collision state is computed according to a process that satisfies the relevant conservation laws (in this case mass, energy, and momentum). Multiple proposals appear in the literature for both the selection rules and the post-collision state calculation [23].

Boiled down to its essence, there are three quantities that must be reproduced properly by the overall DSMC collision scheme to yield a physically reasonable simulation:

1. collision rate
2. probability of a particular collision
3. post-collision deflection angle

In a typical simulation, the collision rate is set by a combination of the number of pairs considered each timestep and the probability of particular pairs colliding. The probability of a particular collision is set by the collision cross section and the relative speed. When constructing a model for gases, the macroscopic quantities of diffusion and viscosity are used as an indicator of the appropriateness of a given microscopic model. For example, the variable hard sphere (VHS) model [50] was developed in order to bring the simulations into agreement with experiment on the viscosity temperature dependence and the variable soft sphere (VSS) model [51] was subsequently developed to correct the VHS model's deficiencies in reproducing diffusion processes.

Following the nomenclature of Srivastava [36], the most basic type of interphonon interaction, known as three-phonon processes, may divided into two classes: class I processes, where two phonons coalesce into one phonon, and class II processes, where one phonon splits into two phonons. These classes are illustrated schematically in Fig. 40.

Class I collisions immediately appear compatible with the basic DSMC algorithm because they begin with two phonons, thus the idea of selecting a pair from the cell and deciding whether or not they will interact based on their individual and pairwise characteristics should be possible. Class II collisions, however, bear no such resemblance to basic DSMC collisions.

Previous work [35, 36, 52] generally gives the collision probability in terms of occupation numbers of all three phonon modes (which, strictly, makes it a state transition probability) and creation/annihilation operators that depend on the polarization and frequency of the participating phonons. In the current framework, occupation numbers are time consuming and error-prone to compute due to the relatively coarse representation of the distribution. In addition, the creation/annihilation operators, and even their functional form, are largely unknown.

(a) class I (b) class II

Figure 40. The two classes of three-phonon processes.

In this section, new collision schemes are constructed to evaluate the hypothesis that the occupation numbers and creation/annihilation operators in the transition probabilities can be replaced by a DSMC-like process. In this process, population contributions are replaced by the probability of selecting a particular phonon state and the creation/annihilation operators are replaced by the probability of a selected collision type and partners successfully negotiating a set of constraints built around the conservation laws and the dispersion relation.

4.8.1 Constraints

The conservations laws generally considered valid for phonon interactions spring directly from the state transition probabilities, appearing as "resonance factors" that cause the only significant contributions to be centered on processes that satisfy the laws:

$$\omega + \omega' + \omega'' = 0 \tag{28}$$

$$k + k' + k'' = \mathbf{G} \tag{29}$$

where the unprimed, prime, and double prime exponents refer to the three phonons participating in the process and the quantities are negative for outgoing phonons.

The first expression represents a statement of the conservation of energy. The second expression represents a statement of conservation of momentum, where the term is used rather loosely in this context because phonons have no mass (note that the three phonon processes all violate number conservation). The rather fuzzy nature of phonon momentum can be seen in the term **G** in Eq. 29. This term may be either the reciprocal lattice vector or

zero. In cases where **G** is zero, it is said that the process conserves "crystal momentum" and it is referred to as a normal process, or simply an N-process. In cases where **G** is nonzero, the crystal momentum is lost (generally considered gained by the crystal as a whole) and the process is referred to as an Umklapp process or a U-process.

In addition to the frequency and wavevector constraints expressed by Eqs. 28 and 29, all participating phonons must also be located on the dispersion relation. This additional constraint can be shown through a number of graphical constructs [35, 52, 53] to limit the possible polarization combinations to:

$$TA + TA \leftrightarrow LA \tag{30}$$

$$TA + LA \leftrightarrow LA \tag{31}$$

These graphical constructs, particularly under the isotropic dispersion relation used in this work, make it clear that very few phonon combinations will satisfy all constraints. In fact, once a particular phonon is chosen for collision, often only one polarization/wavevector combination for the second phonon will yield a realizable third phonon. Randomly choosing pairs of phonons in a given cell, particularly given the reduced number of computational particles over the physical case, is therefore very unlikely to yield a combination that will strictly satisfy the constraints. Few to no collisions would therefore be performed even if all pairs were considered (which would require computational work scaling as the number of particles squared).

The strategy attempted in this section for assembling collisions under these constraints is to always conserve energy exactly and allow, but discourage, departures from Eq. 29 by including a term in the collision probability that has a relatively narrow maximum around situations where it is satisfied. This strategy is consistent with the transition probabilities proposed by many authors, which do not go to zero immediately outside of the constraints imposed by Eqs. 28-29. In this manner, selecting particles randomly from the cell mimics the occupation numbers in previous transition probabilities and the collision probability function, which is of the form:

$$\wp = exp(-\eta k_{err}^2) \tag{32}$$

where η is some coefficient that determines the width of the peak, mimics the resonance factors. Unless otherwise noted, $\eta = 100$ in this section.

4.8.2 Method

To evaluate the collision schemes, 10 million computational particles are placed on a domain from the equilibrium distribution at a temperature of 300K. A single collision step is

then performed and the rates for each process and the resulting distribution are examined. This procedure is constructed to check a necessary condition for a reasonable collision scheme: that the distribution should not change when starting from equilibrium, which implies that the forward and backward rates of each creation/destruction process should be equal. Because this condition should hold regardless of the collision rate, the number of iterations through the collision routine was simply set to the number of particles in the cell for these tests.

Several permutations of the basic collision scheme were evaluated. In all cases, the stages can be broken into:

1. classification
2. computation
3. decision

If this process indicates that a particular collision should proceed, a further stage is required: determination of the post-collision state.

Classification Stage

One key issue is how to deal with the differing number of pre- and post-collision phonons in the two collision classes. Three schemes were considered. In each scheme, the goal is to choose the collision class, as well as the polarization of the outgoing phonons if a class II collision is chosen.

The first classification scheme chooses two particles from the cell. If they are both TA, the only process available is class I and both phonons are used on the pre-collision side of the equation. Conversely, if they are both LA, only class II processes are possible and one phonon is used on each side of the equation (in a manner to be determined at stage 3). If one of each polarization is chosen, either class is possible and each is given a 50% probability. If the chosen collision type is subsequently ruled impossible, the scheme aborts.

This scheme, somewhat independent of decisions made in the remaining stages, resulted in a vast number of $TA + TA \rightarrow LA$ collisions due to the preponderance of TA phonons at 300K (71% of the chosen pairs were $TA - TA$). The net result was a large increase in the LA population and a twice-as-large decrease in the TA population. Modifying this scheme by creating an acceptance probability on $TA + TA \rightarrow LA$ of (N_{LA}/N) and on $LA \rightarrow TA + TA$ of (N_{TA}/N) helped close this gap.

The second classification scheme chooses three phonons from the cell for each sample. The resulting polarization combination determines which of Eqs. 30-31 are possible (or neither if all three phonons had the same polarization). The class is then chosen randomly, with each having a 50% probability.

This scheme was constructed in an attempt to force the forward and reverse processes of each class to occur with equal probabilities, thus guaranteeing detailed balance. This turned out to not be the case once the subsequent stages of the process were considered. In particular, due to dispersion, there are a larger number of frequency-wavevector combinations possible when forming an LA phonon than a TA. The class I processes therefore tend to have a higher success rate than class IIs. This problem can be corrected somewhat by allowing the class I/class II decision to be made after considering downstream conditions and by using weighting functions. This scheme was ultimately abandoned, however, when it was realized that it would completely fail in a case in which the domain was initialized to contain only phonons of one polarization. In this situation, this scheme would yield no collisions, thus it would never reach the proper equilibrium.

The third classification scheme combined elements of the first and second schemes: two phonons are chosen from the cell, but the class is selected by subsequent tests. If neither class is ruled out, both are considered possible and probabilities for each, as well as neither, are computed and considered in the decision stage.

It may be noted that one of the phonons on the product side of the process in class II collision may reasonably be selected from the equilibrium distribution instead of from the cell. The relative merits of this option cannot be evaluated by the current test case because the phonons in the cell are also selected from the equilibrium distribution.

Computation Stage

Once possible partners and collision classes are selected, the computation stage begins. In this stage, information from the known phonons is used to construct the unknown phonon. In classification schemes 1 and 3, two phonons are chosen from the cell at stage 1. For class I collisions, the wavevectors of these phonons are combined to create the target wavevector for the outgoing phonon. For class II collisions, one of the phonons chosen from the domain will determine the wavevector and polarization of one of the outgoing phonons. The unknown phonon will be calculated by subtracting this phonon's wavevector from the other phonon chosen at the classification stage. Classification scheme 2 functions in a similar manner except a third phonon is chosen in the first step and used to determine the polarization of the second outgoing phonon in class II collisions.

To compute the wavevector of the unknown phonon, the two known phonon wavevec-

tors are broken into components along the Cartesian directions. The unknown wavevector components are then determined according to the type of collision class using Eq. 29.

At this stage, a procedure for U processes must be chosen. In this work, the components were examined individually. Any component that exceeded unity was added to a two unit vector pointed in the opposite direction. This operation reverses the component and yields the proper wavenumber in terms of reflection across the Brillouin zone boundary (most easily seen in Peierls [52], Fig. 3). For example, a phonon with $k=1.4$, directed along $+X$ should become a phonon with $k=0.6$ directed along $-X$. This process does not guarantee that the overall magnitude will be less than unity, however, so this must also be checked after the component-wise checks are completed. In this case, a U process is performed with a wavevector in the opposite direction as the overall wavevector. In Peierls' construction, it can be seen that $TA+TA \leftrightarrow LA$ transitions are only available through U processes, so a flag is set if a U process is performed to allow enforcement of this constraint.

The frequency of the unknown phonon is now calculated according to Eq. 28. The target wavenumber (energy-conserving k) is then determined from the dispersion relation and k_{err}^2 is calculated as the square of the difference between the result and that found by the vector operations described in the previous paragraph, normalized by the energy-conserving k.

Decision Stage

With the potential collision (or collisions, in the case of classification scheme 3) defined, a decision must be made as to whether or not it is performed. For classification schemes 1 and 2, this requires simply evaluating a probability, $\wp$ from Eq. 32 and choosing a random fraction, R_f. If $R_f < \wp$, the collision is performed. For classification scheme 3, Eq. 32 is evaluated for each potential collision. A random fraction is then chosen and a potential collision, or none at all, is chosen based on its value. For example, if the probability of collision A is 0.2 and the probability of collision B is 0.3, collision A is performed if $0 < R_f < 0.2$ and collision B is performed if $0.2 < R_f < 0.5$, and neither is performed if $R_f > 0.5$ [43].

Post-Collision State Calculation

If a collision is accepted, any phonons on the left side of the corresponding process shown in Fig. 40 are deleted from the simulation. If the process is class II, a new phonon is created with the same polarization and wavenumber as the test phonon chosen from the cell. To avoid cloning the existing phonon, which can adversely affect the statistics by creating a highly unlikely configuration of two particles occupying exactly the same point in phase space, the wavevectors of both created phonons are displaced a random amount on either side of their vector average.

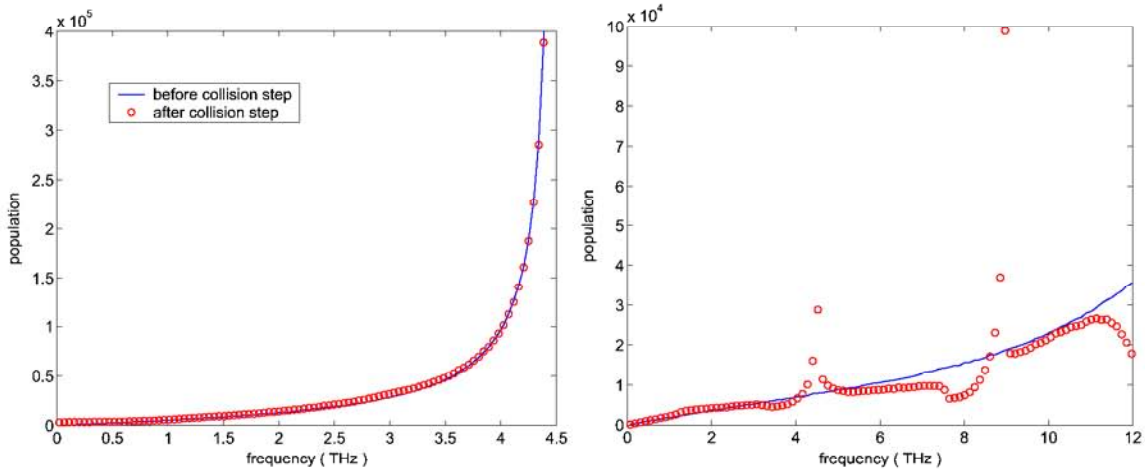

(a) transverse

(b) longitudinal

Figure 41. Comparison of population histograms before and after a single collision step.

4.8.3 Results

Many permutations of the new collision scheme were evaluated. At first, the goal was to arrange the decision tree, which includes the selection process, tests for constraints, and calculation of a collision probability, such that no net change in phonon population or distribution was observed in the equilibrium test case. Due to the complexity of the decision tree, particularly in the presence of dispersion, this proved to be an extremely difficult task. The goal was eventually relaxed by introducing coefficients, similar to forward and backward reaction coefficients, to force the processes to balance. This left maintenance of the proper distribution the only remaining hurdle.

One of the more successful outcomes is shown in Fig. 41. It can be seen from this figure that the distributions before and after the collision process are similar, particularly in the transverse case. The primary differences are peaks in the LA distribution near 4.5 and 9 THz. The latter peak corresponds to a frequency equal to twice the maximum frequency of the TA mode. Considering the dispersion relation and the Bose Einstein distribution together, it can be seen that 4.5 THz is, by far, the most probable frequency for TA phonons. A great number of candidates therefore exist to feed the $TA + TA \rightarrow LA$ process in equilibrium, but not the reverse process. The smaller 4.5 THz LA peak may be akin to a harmonic

of the 9 THz LA peak: as the process progresses, many 9 THz LA phonons are available. If one is chosen for a collision, it is very likely a $LA \rightarrow LA + TA$ process will choose a 4.5 THz TA and, thus, a 4.5 THz LA product.

Through the use of differing forward-backward rate coefficients, this problem can be partially corrected. The uniform application of these coefficients to all reactions of a given type, however, causes problems elsewhere on the distribution. While time did not permit investigation of some possible remedies, this observation may indicate that the frequency and polarization-dependent operators we set out to avoid may be necessary after all.

5 Semiclassical Models

This section describes simplified models, most of which grow from classical theory, that were modified to improve their predictive capability based on data and understanding gained through the experiments and phonon-scale simulations performed under this LDRD. These models were constructed with the goal of providing microscale information to macroscale codes in a computationally convenient manner. This goal is critical to successful modeling of microdevices because, when considered on a system scale, most are far too large for microscale techniques to be practical using current computers. A successful tool will therefore be capable of identifying dominant mechanisms affecting thermal transport, which will vary with temperature and characteristic length scale. If mechanisms associated with microscale effects, such as boundary scattering, become significant, the subgrid model will provide the necessary corrections to the macroscale model.

For most current devices, which have a limited number of areas where microscale effects are significant, but not dominant, this goal can be accomplished by developing the "effective thermal conductivity" concept. Under this concept, microscale effects are communicated to a macroscale code by modifying the local thermal conductivity, which may now be anisotropic, heterogeneous, and a nonlinear function of temperature. The microscale model can thus be supplied in the form of a subroutine for calculating κ and changes to the overall code are minimized. This scheme works well within the general scope of the ASCI program that promotes the use of modeling and simulation for performance analysis and qualification of microdevices as well as integrated systems with scales reaching meters.

This approach has the advantages of:

- ease of implementation
- better understanding
- robustness

when compared to performing a coupled continuum-noncontinuum simulation. However the "effective thermal conductivity" approach has a few drawbacks:

- assumption of near-equilibrium
- required validation of heat flux - temperature gradient relationship
- explicit temperature dependence.

5.1 Modeling Strategy

In a thermal non-equilibrium state, where a temperature gradient exists, the phonon distribution will deviate from the Bose-Einstein distribution given by Eq. 6. The degree of deviation from the B-E distribution depends on the heat flux and boundary conditions. To obtain the new phonon distribution, one needs to solve the Boltzmann equation. However this is a difficult task because of limitations in computer power and the complexity of modeling the phonon scattering process, which highly depends on the phonon propagation mode and local temperature.

A classical approach to this problem is to linearize the Boltzmann equation using perturbation theory and to develop an analytical closed-form solution for the phonon distribution function. A key assumption in this approach is that the nonequilibrium phonon distribution is only a small deviation from the equilibrium condition. All the phonon scattering processes are assumed to be ways to restore the phonon distribution back to the thermal equilibrium condition. The restoration rate is proportional to the departure of the distribution function from its equilibrium condition. The time scale to characterize the rate is the relaxation time, τ.

Under these assumptions, a thermal conductivity can be defined as [54, 38]:

$$\kappa = \frac{1}{3(2\pi)^3} \sum_p \int v_g^2 \tau \, C(p,\omega) \, d\mathbf{k} \tag{33}$$

where $C(p,\omega)$ is the contribution of the phonon mode with polarization p and frequency ω to the specific heat and the integral is taken over the wavevector, $\mathbf{k}$ and summed over all polarizations.

If a constant average phonon speed and relaxation time are assumed, the integral may be easily evaluated to yield:

$$\kappa = \frac{1}{3} C v_g^2 \tau \tag{34}$$

which may also be derived via kinetic theory [31, 22]

Calculating the thermal conductivity under this model therefore becomes a question of how the group speed and the relaxation time are treated. In the latter case, another assumption is required: that each scattering process (and its rate) is independent from each other. That implies the existence of quantum chaos. The chaotic nature permits us to use Matthiessen's rule [31] to approximate the total relaxation time in a manner similar to calculating a series resistance in an electrical circuit. That is, a single relaxation time can be defined for each class of scattering process. The reciprocal of the total relaxation time is then the sum of the reciprocals of the relaxation time for each phonon scattering process.

Next it is important to study each class of scattering process and determine the relaxation time and its variation with temperature, phonon propagation mode, and polarization. A moderate amount of work toward understanding and characterizing the dynamics of phonon transport and scattering in materials and quantitatively predicting the thermal conductivity of materials exists. However, most of this work was driven by an interest in material behavior in the low-temperature or cryogenic conditions. Very little work addresses the microscale heat transfer problem. The difference is that at the lowest temperatures, the boundary scattering is dominant. Hence the thermal conductivity depends on the size and shape of the crystal or crystallites. On the other hand, for structures in microsystems above room temperature, both boundary scattering and three-phonon scattering will dominate. This makes it difficult to predict the thermal conductivity behavior of materials.

5.2 Baseline Model

Among existing models, the Callaway [55] and Holland [38] models were chosen to be evaluated in detail. These models were selected because they have been used by many researchers to predict thermal conductivity of bulk silicon and other semiconductors. The Callaway model uses a single mode for phonon transport, while the Holland model utilizes three modes (one longitudinal and two transverse modes). These are the baseline models that may be further developed to incorporate any essential microscale effects to predict thermal response in microsystems. Another interest is to determine the tradeoff between complexity and accuracy of employing a single dominant mode versus multiple modes of phonon propagation. Obviously, including multiple modes of phonon propagation will produce a more accurate value of thermal conductivity, however it may require significant additional computation time.

Figure 42 shows a comparison of the thermal conductivity of silicon predicted by both the Callaway (single mode) model and the Holland (multiple mode) model with measurements by Holland and Neuringer [46] (1.7 to 300 K) and Slack and Glassbrenner [56] (300 to 1683 K). The Callaway model fails at high temperature because it focuses on capturing the normal and Umklapp scattering processes at low temperature. At high temperature, the dependence of normal and Umklapp processes is quite different. In addition, we need to be cautious with the Callaway model because it predicts the temperature dependence is T^{-n}, where $n < 1$, while experimental data implies the temperature dependence is $T^{-1.2}$ from 200K to well above 658K [38]. For these reasons, the Holland model was selected as the baseline for further work.

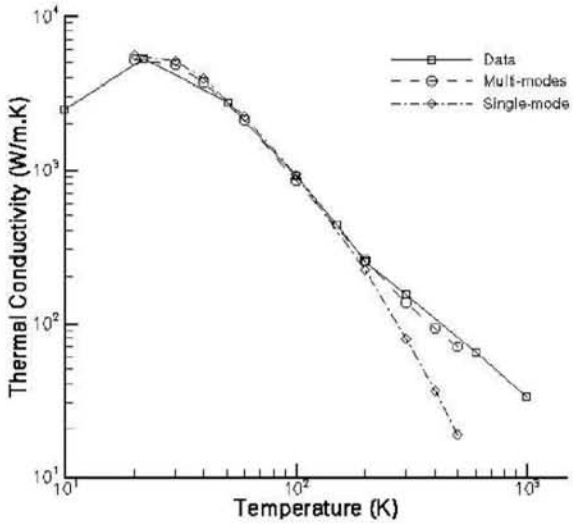

Figure 42. Predictions of single mode and multimode models compared with experimental data for bulk silicon.

5.3 Boundary Scattering

As shown in Fig. 35, the Holland model performs poorly in predicting the thermal conductivity of thin silicon films. Because it performs well for bulk silicon, this shortcoming can reasonably be attributed to its boundary scattering model.

In the Holland model, the relaxation time for boundary scattering is determined by dividing the thickness of a structure by the phonon propagation speed and multiplying by a weighting factor, F:

$$\tau_{boundary} = F\frac{t}{v_g} \tag{35}$$

This weighting factor is intended to characterize the effect of surface topology. In Holland's work, a constant value of $F = 0.8$ is used. This constant value is derived from measurements of the thermal conductivity of bulk silicon at low temperature. At low temperature, even for bulk samples, the boundary scattering dominates because the phonon mean free path increases with decreasing temperature.

Our approach to improve the boundary scattering model is to consider the surface topol-

ogy at the boundary and how its scattering effect will vary with the wave vector of an incident phonon. This improved model follows the work of Ziman [31] and McConnell *et al.* [8]. Since six polarization modes of phonon propagation are considered in our model, a constant value cannot reflect the very different interactions between individual phonon wave vectors and the surface topology. In order to capture this variation in interaction, we must treat phonon propagation as waves that impinge upon, and reflect from, the wall. Modeling in this way will allow us to compare the phonon wave vector with the roughness of the thin-film surface. By comparing the reflection of two parallel waves as a result of scattering from the boundary, one can derive the probability of specular reflection as follows:

$$\wp = exp\left[-4\pi\left(\frac{\eta\omega}{v_g}\right)^2\right] \tag{36}$$

where η is the surface roughness.

By tracing the propagation path of phonons scattering from the boundary, for several collisions, one can express the weighting factor in a different form as follows:

$$F = \frac{1+\wp}{1-\wp} \tag{37}$$

thus $F = \infty$ for zero roughness and $F \rightarrow 1$ as the roughness increases. Note that now F is a function of frequency as well as polarization. If the roughness is small compared to the incident wavelength, the wave is likely to scatter specularly, increasing F and, consequently, the relaxation time due to boundary scattering. The boundary effect on thermal conductivity in the current model is thus reduced for that phonon state.

Figures 43 to 45 show predictions of the thermal conductivity of silicon, for three different film thicknesses, by several variations on the semiclassical model and their comparison with measurements made by Asheghi at Stanford [16] and measurements made as part of the current work, which were described in Sec. 2.

The models used to produce Fig. 43, in which the film thickness is 1.5 μm, differ in their treatment of boundary scattering. As noted previously, the original Holland model underpredicts the thermal conductivity significantly. The first modified model uses a single weighted average phonon propagation speed, as in the original model, but a nonconstant weighting factor calculated from Eq. 37 using a relatively rough surface (rms roughness ~20nm). Again, the predicted results do not compare well with experimental data. If we repeat this calculation with the rms surface roughness changed to 0.1nm, which is more reasonable for single crystal silicon, the result compares much better with experimental data. At room temperature, however, it still underpredicts the thermal conductivity somewhat. Finally, when we use the group velocity for each mode from the dispersion model

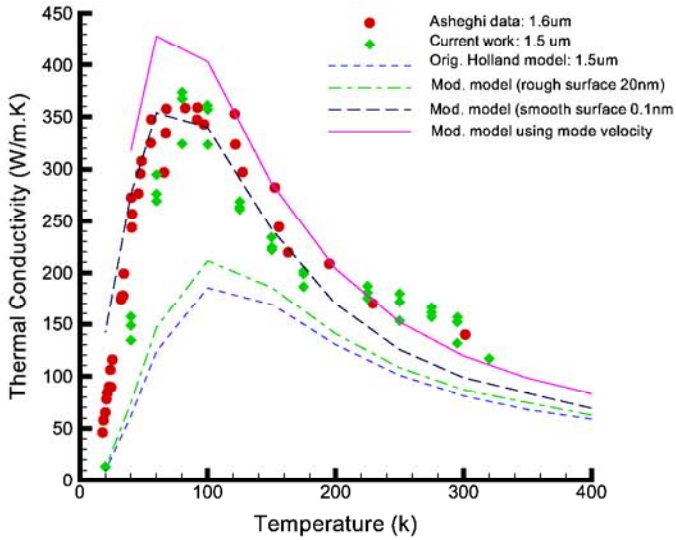

Figure 43. Comparison of the predicted thermal conductivity of a 1.5 μm single crystal silicon film to measured values.

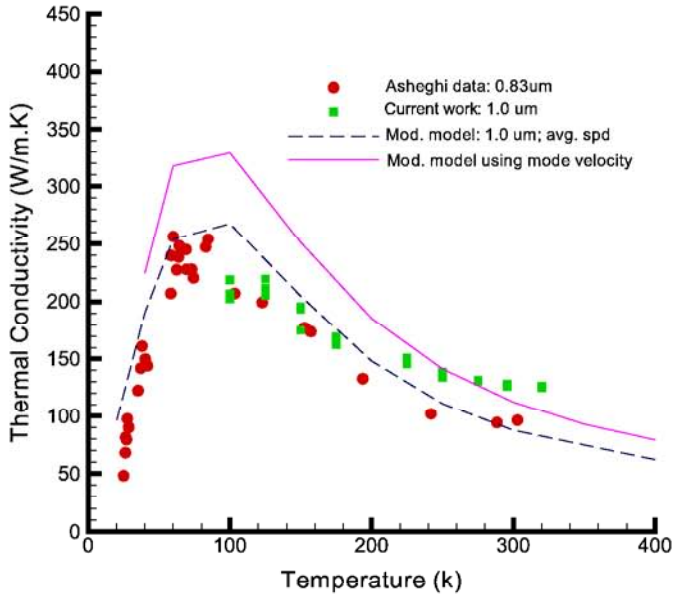

Figure 44. Comparison of the predicted thermal conductivity of a 1 μm single crystal silicon film to measured values.

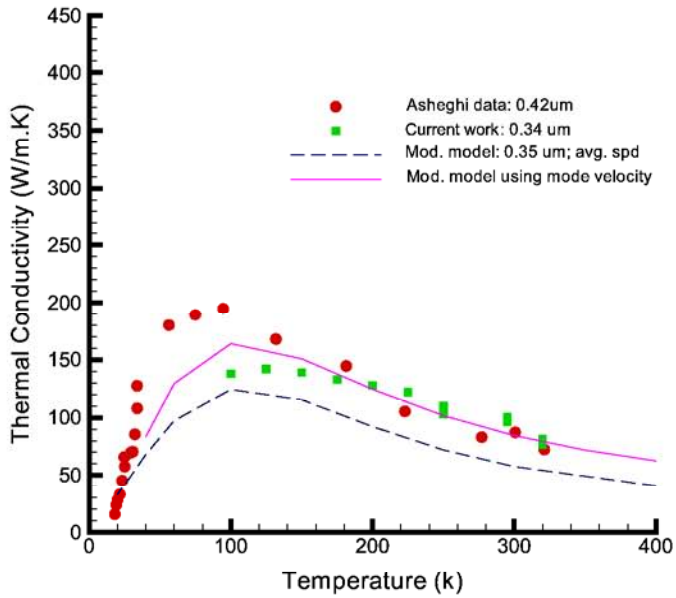

Figure 45. Comparison of the predicted thermal conductivity of a 0.35 μm single crystal silicon film to measured values.

described in Sec. 3.3, the results compare reasonably well with the data near room temperature.

Figure 44 presents a similar comparison for a 1 μm film. As in the 1.5 μm film, the modified model with a smooth surface and an averaged phonon speed compares well with experimental data at low temperature (<200K), while the modified model with mode-specific speeds performs better at room temperature or above. For the 0.35 μm film (Fig. 45), the mode speed model matches the data better than the average speed model down to the lowest temperature (100K).

Figures 46-48 show the contribution of different phonon propagation modes to the total thermal conductivity of the thin-film single crystal silicon with different thicknesses. Predictions show that at room temperature, the upper transverse mode is the dominant mode for phonon propagation.

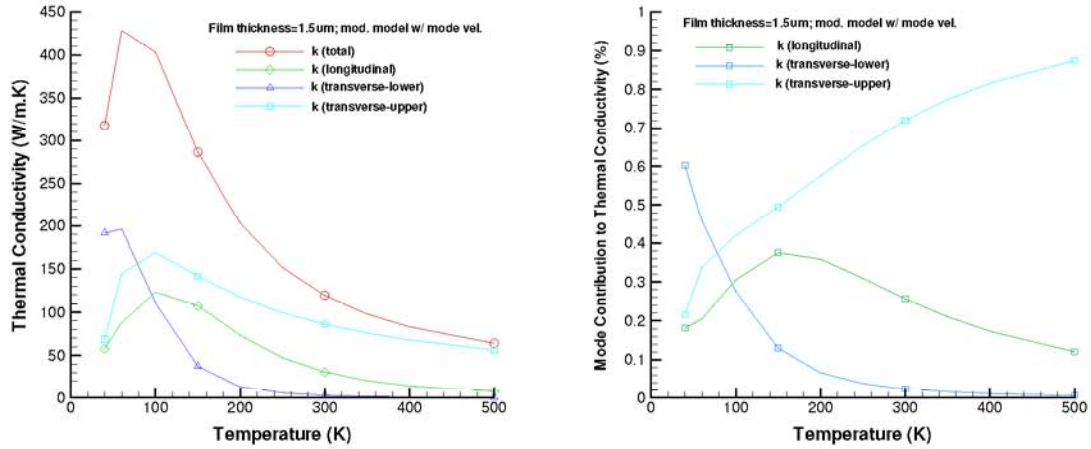

Figure 46. Contribution of phonon propagation modes to thermal conductivity of 1.5 μm single crystal silicon films.

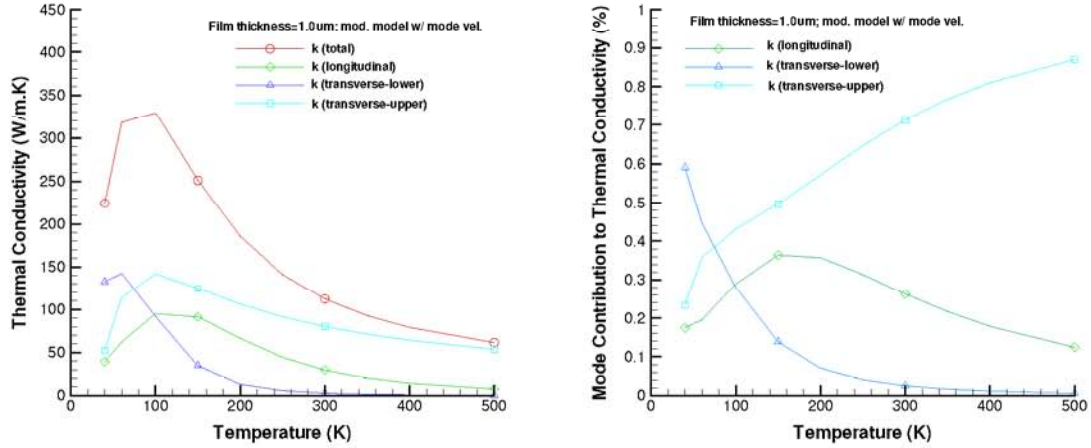

Figure 47. Contribution of phonon propagation modes to thermal conductivity of 1 μm single crystal silicon films.

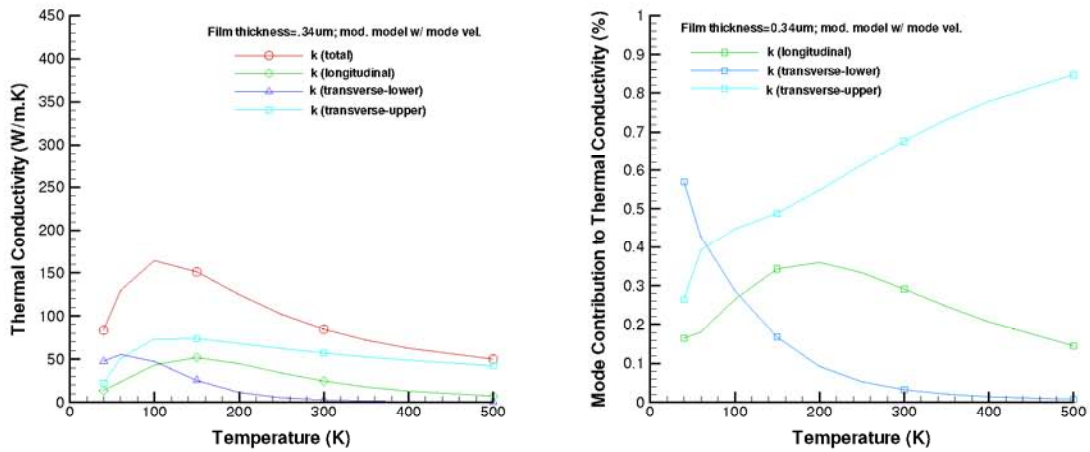

Figure 48. Contribution of phonon propagation modes to thermal conductivity of $0.35\,\mu m$ single crystal silicon films.

5.4 Grain Boundary Scattering

To model polysilicon, the Holland-type phonon transport model was modified to include phonon-grain boundary scattering effects in a manner similar to that used in the previous section for phonon-boundary scattering. This modification uses the method suggested in McConnell *et al.* [8] with group speeds from the dispersion model discussed in Sec. 3.3.

Results from this model are shown in Fig. 49 compared to experimental data for single and double $2\,\mu m$ polysilicon layers from Fig. 11. The model shows an excellent fit to the temperature dependence of the thermal conductivity. Analysis of the effects of phonon scattering mechanisms show that grain boundary and defect scattering dominate the response. It should be noted that without defect scattering included, the shape of the thermal conductivity versus temperature response could not be captured.

While the Holland model can be used to predict the effect of grain boundaries and other phonon scattering effects on thermal transport, a simplified semi-empirical formula was developed based on an electrical resistance analogy. In this case, the thermal resistance of the material is broken down into an intrinsic resistance of single crystal silicon plus the

resistance due to grain boundaries:

$$\frac{1}{\kappa_{film}} = \frac{1}{\kappa_{bulk}} + \rho_{grain-boundary}\left(\frac{S}{V}\right), \tag{38}$$

where S/V is the surface to volume ratio of the grain and $\rho_{grain-boundary}$ is the specific grain boundary thermal resistivity. Film boundary effects are assumed negligible based on previous measurements and MC simulations, thus the results are independent of film thickness. The specific grain boundary thermal resistivity for the undoped samples was calibrated at room temperature and used to estimate the effect of grain size on the thermal conductivity at the calibrated temperature. In this case, for room temperature, the specific grain boundary resistivity was determined to be $2.07 \times 10^{-9} Km^2/W$ with a reference resistivity of $6.667 \times 10^{-3} Km^2/W$. Using this value of resistivity, the semi-empirical estimates of room temperature thermal conductivity value are plotted in Fig. 50 versus experimental data taken in our study and others from our collaborator's group [8, 17]. The fit gives an excellent estimate of the thermal conductivity as a function of grain size. The largest error is seen for the as-deposited film, which has the largest defect density. An additional term taking into account defect scattering resistivity should therefore be applied to Eq. 38 in order to make it more physically meaningful.

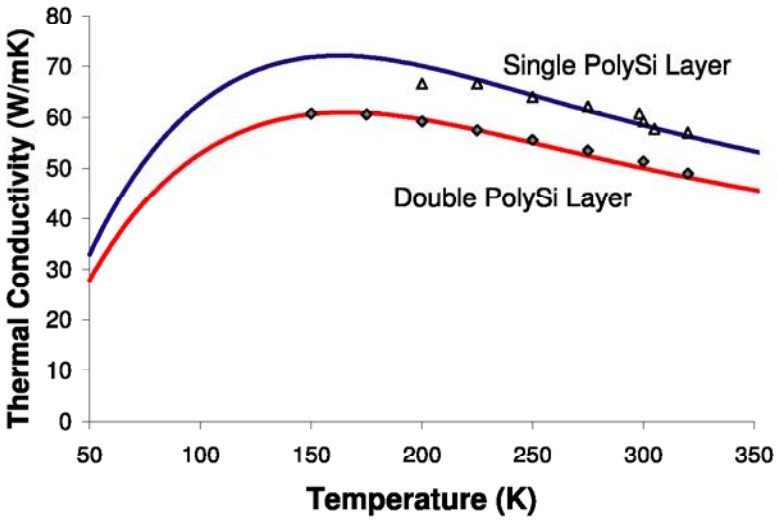

Figure 49. Graph showing the fit of Holland's model to experimental data for the polycrystalline silicon samples shown in Fig. 12. The model shows an excellent correlation with the data.

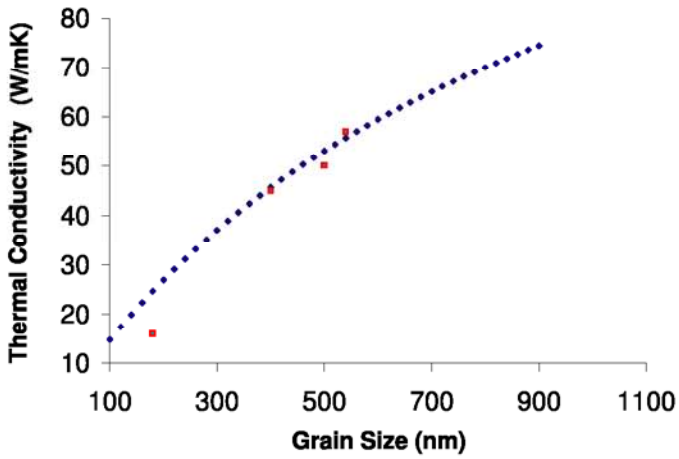

Figure 50. Comparison of the semi-empirical prediction of room temperature thermal conductivity of polysilicon to experimental data from the current work, Ref. [8], and Ref. [17].

6 Conclusions and Recommendations

6.1 Experiments

The experimental portion of this LDRD added a new dimension to the measurement capabilities at Sandia. This work was the first to explore the thermal transport properties of semiconductor thin-films down to thicknesses of 340 *nm*. The results showed a very strong temperature and size scale dependence which helped to elucidate the pertinent factors controlling thermal conductivity in silicon materials.

Key findings in this work were as follows:

- The length scale effects on the thermal conductivity of single crystalline silicon do exist at room temperature for films which are on the order of 300 *nm* and less. A reduction in thermal conductivity was seen on the order of 40% at 300K and increased to be an order of magnitude at temperatures below 100K. For thicker films ($>1\,\mu m$), deviation from bulk behavior was insignificant at room temperature but was manifested as the temperature decreased.
- The effect of doping on the thermal conductivity of single crystalline silicon was not observed over the temperature range in which the experiments were performed.
- The thermal conductivity of undoped polycrystalline Si showed a strong dependence on grain size. Grain boundary scattering along with defect scattering were determined to be the primary factors controlling the reduction of thermal conductivity and its temperature dependence. Future work should be extended to doped polycrystalline silicon fabricated by the SUMMiT V process. Experiments must also be extended to higher temperatures in order to provide data over the temperature range of many thermal MEMS applications. Such experiments were planned but not executed in the LDRD due to logistical obstacles in obtaining samples.

Overall, these experiments proved to be difficult due to the sample fabrication requirements. In future programs, dedicated fabrication assistance or outsourcing specimen fabrication may be prudent. From the experiences of the authors, Sandia's current facilities are not amenable to the fabrication of highly specialized and tailored films and samples.

Finally, experimental tools were developed during this work which enabled the measurement of the thermal conductivity of thin-films. Both steady-state and transient techniques were implemented at Sandia through the course of the LDRD and will remain a

platform for future studies. These tools can cover the temperature range of 4-320K, and a new custom system was designed with Janis Research to extend the temperature range up to 800K. This will provide additional resources to fill the needs of thermal property measurements for Sandia's MEMS programs in the future.

6.2 Monte Carlo Simulation

The Monte Carlo simulation tool developed for this work has proven valuable for investigating phonon transport. In particular, its ability to employ a realistic dispersion relation and to enable scattering mechanisms to be modeled independently provides a more physically reasonable representation than was possible in past work. The simulations were thus able to successfully capture the thermal conductivity, over the temperature range of 100-300K, of isotopically-enriched and natural bulk silicon, as well as films as thin as 340 *nm*. Most notably, these simulations were performed in a predictive mode, *i.e.* without adjusting any empirical parameters when changing conditions or geometries.

Due to its limited time frame, this project can best be described as providing a basis for future work. A short list of examples would include:

- While impurity scattering model performed well for naturally-occurring isotopes, an investigation of doped materials over a wide parameter space should be performed. It is possible that some dopants may stress the lattice sufficiently to require additional terms in the current model, which accounts only for scattering due to mass difference. It should be kept in mind, however, that the accuracy of the model as it stands will decline when the doping reaches a level such that charge carriers are contributing significantly to thermal transport.
- Significant additional work is required to place the grain boundary scattering model on a firm footing. The process of fitting the semiclassical model to the experimental data indicated that defects near the boundaries may play a significant role. A model for these defects should therefore be developed.
- While this project ended before a firm conclusion could be reached, it appears that a frequency and polarization-dependent creation/annihilation operator may be required to achieve the proper equilibrium distribution in a non-BGK collision process.

It may be possible to glean information on many of the unknowns cited above through molecular dynamics simulation. Regardless of the final outcome, the process of performing such simulations would produce valuable insight into aspects of the collision process that are unclear in the theoretical analyses to date.

6.3 Semiclassical Methods

While experiments and Monte Carlo simulation have proven valuable as investigative tools for gaining insight into phonon transport, they have limited value by themselves to a device designer. Particularly in the case of microfabricated devices, the time and monetary cost of design iterations, as well as the difficulty in obtaining measurements with sufficient resolution, generally preclude exclusive reliance on testing to guide the development phase. Similarly, 3D Monte Carlo simulations of full devices, particularly with the moving parts found in many MEMS, are currently impractical in most cases. The semiclassical methods investigated in this work are therefore critical to providing timely, cost-effective input to the design process.

Continued work on these models is necessary to include more detailed physics and to establish quantitative bounds on their regimes of applicability. Such work is critical to providing predictive models that can confidently be applied to devices that operate outside of conventional limits, which is a common situation at Sandia in general and in microdevices in particular.

A Acceptance-Rejection Method

The "acceptance-rejection method" is a means for choosing a quantity randomly when its distribution function is not invertible. This method, detailed in Appendix C of Bird [23], involves selecting a value for the quantity of interest using a random number generator. The distribution function is then evaluated at this value. A second random number is then generated and compared to this result divided by the maximum value of the distribution function. If the random number is smaller than this quotient, the selected value is accepted and returned. Otherwise, the selected value is rejected and the entire process begins again.

Two pieces of information are required for a successful implementation of this technique: the distribution maximum and the bounds on the quantity of interest.

If the distribution maximum is set artificially high, the resulting quantities will be properly distributed, but several otherwise reasonable values will be discarded, increasing the cost of the technique. If the distribution maximum is set too low, its peak will be "flattened." In practice, for distributions with unknown maxima, the maximum is compared to the function evaluation each time it is called and increased when necessary. This technique practically eliminates the danger of reshaping the distribution, but it introduces a conditional to every function evaluation, slowing execution.

For selecting from the Bose-Einstein distribution, the maximum occurs at $k = 1$ for temperatures above approximately 100K. Below 20K, this maximum shifts drastically to approximately $k = 0.05$.

Choosing the bounds on the quantity of interest is straightforward except in cases, such as choosing gas molecule speeds from a Maxwellian distribution, where one or more of the bounds is infinite. In this case, a compromise must be reached between the number of rejected values and the fraction of the quantity's domain that is excluded. Because the codes in this work function with the reduced wavenumber as the primary descriptive variable, the bounds are simply 0 and 1.

References

[1] R. Viswanath, V. Wakharkar, A. Watwe, and V. Lebonheur. "Thermal performance challenges from silicon to systems," *Intel Technology Journal*, 3rd Quarter 2000.

[2] M. I. Flik, B. I. Choi, and K. E. Goodson. "Heat transfer regimes in microstructures," *Journal of Heat Transfer*, 114:666–674, Aug 1992.

[3] G. Chen and M. Neagu. "Thermal conductivity and heat transfer in superlattices," *Applied Physics Letters*, 71(19):2761–2763, Nov 1997.

[4] P. K. Schelling and S. R. Phillpot. "Multiscale simulation of phonon transport in superlattices," *Journal of Applied Physics*, 93(9):5377–5387, May 2003.

[5] E. Jansen and E. Obermeier. "Thermal conductivity measurements on thin films based on micromechanical devices," *Journal of Micromechanics and Microengineering*, 6(1):118–121, Mar 1996.

[6] O. Paul and H. Baltes. "Thermal conductivity of CMOS materials for the optimization of microsensors," *Journal of Micromechanics and Microengineering*, 3(3):110–112, Sep 1993.

[7] A. D. McConnell and K. E. Goodson. "Thermal conduction in silicon micro and nanostructures," *Journal of Microelectromechanical Systems*. in press.

[8] A. D. McConnell, S. Uma, and K. E. Goodson. "Thermal conductivity of doped polysilicon layers," *Journal of Microelectromechanical Systems*, 10(3):360–369, Sep 2001.

[9] S. Graham, B. Olson, C. Wong, E. S. Piekos, A. D. McConnell, and K. E. Goodson. "The effects of processing conditions on the thermal conductivity of polycrystalline silicon films," In *Proceedings of ASME-IMECE*, Washington, D.C., 2003.

[10] M. von Arx, O. Paul, and H. Baltes. "Process-dependent thin-film thermal conductivities for thermal CMOS MEMS," *Journal of Microelectromechanical Systems*, 9(1):136–145, Mar 2000.

[11] D. Cahill, M. Katiyar, and J. Abelson. "Thermal conductivity of *a*-Si:H thin films," *Physical Review B*, 50(9):6077–6082, Sep 1994.

[12] Y. S. Ju, K. Kurabayashi, and K. E. Goodson. "Thermal characterization of anisotropic thin dielectric films using harmonic joule heating," *Thin Solid Films*, 339(1-2):160–164, Feb 1999.

[13] X. Zhang and C. P. Grigoropoulos. "Thermal conductivity and diffusivity of free-standing silicon-nitride thin-films," *Review of Scientific Instruments*, 66(2):1115–1120, Feb 1995.

[14] D. Cahill, K. E. Goodson, and A. Majumdar. "Thermometry and thermal transport in micro/nanoscale solid-state devices and structures," *Journal of Heat Transfer*, 124(2):223–241, Apr 2002.

[15] M. N. Touzelbaev, P. Zhou, R. Venkatasubramanian, and K.E. Goodson. "Thermal characterization of Bi_2Te_3/Sb_2Te_3 superlattices," *Journal of Applied Physics*, 90(2):763–767, Jul 2001.

[16] M. Asheghi. *Thermal Transport Properties of Silicon Films*. PhD thesis, Stanford University, 1999.

[17] S. Uma, A. D. McConnell, M. Asheghi, K. Kurabayashi, and K. E. Goodson. "Temperature-dependent thermal conductivity of undoped polycrystalline silicon layers," *International Journal of Thermophysics*, 22(2):605–616, Mar 2001.

[18] P. Giannozzi, S. de Gironcoli, P Pavone, and S. Baroni. "*Ab initio* calculation of phonon dispersions in semiconductors," *Physical Review B*, 43(9):7231–7242, Mar 1991.

[19] J. D. Chung, A. J. H. McGaughey, and M. Kaviany. "Role of phonon dispersion in lattice thermal conductivity modeling," *ASME Journal of Heat Transfer*, 2003. submitted.

[20] B. N. Brockhouse. "Lattice vibrations in silicon and germanium," *Physical Review Letters*, 2(6):256–258, Mar 1959.

[21] G. Nilsson and G. Nelin. "Study of the homology between silicon and germanium by thermal-neutron spectrometry," *Physical Review B*, 6(10):3777–3786, Nov 1972.

[22] C. Kittel. *Introduction to Solid State Physics*. Wiley, New York, 7th edition, 1996.

[23] G. A. Bird. *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*. Oxford Engineering Science. Oxford University Press, New York, NY, 1994.

[24] W. Wagner. "A convergence proof of Bird's direct simulation Monte Carlo method for the Boltzmann equation," *Journal of Statistical Physics*, 66(3-4):1011–1044, Feb 1992.

[25] S. Mazumder and A. Majumdar. "Monte Carlo study of phonon transport in solid thin films including dispersion and polarization," *Journal of Heat Transfer*, 123(4):749–759, Aug 2001.

[26] D. E. Knuth. *Seminumerical Algorithms*, volume 2 of *The Art of Computer Programming*. Addison-Wesley, Reading, MA, 3rd edition, 1998.

[27] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, Cambridge, 2nd edition, 1992.

[28] S. Dietrich. "Efficient computation of particle movement in 3-D DSMC calculations on structured body-fitted grids," In A. E. Beylich, editor, *Rarefied Gas Dynamics: Proceedings of the 17th International Symposium on Rarefied Gas Dynamics*, Aachen, Germany, July 1990. VCH.

[29] R. N. Nance, R. G. Wilmoth, and H. A. Hassan. "A comparison of grid-definition schemes for DSMC," Paper AIAA-96-0604, American Institute of Aeronautics and Astronautics, Jan 1996.

[30] T. J. Otahal, M. A. Gallis, and T. J. Bartel. "An investigation of two-dimensional CAD generated models with body decoupled Cartesian grids for DSMC," Paper AIAA-2000-2361, American Institute of Aeronautics and Astronautics, Jun 2000.

[31] J. M. Ziman. *Electrons and Phonons; the theory of transport phenomena in solids*. The international series of monographs on physics. Clarendon Press, Oxford, 1960.

[32] E. S. Piekos and K. S. Breuer. "Numerical modeling of micromechanical devices using the direct simulation Monte Carlo method," *Journal of Fluids Engineering*, 118(3):464–469, Sep 1996.

[33] W. W. Liou and Y. C. Fang. "Implicit boundary conditions for direct simulation Monte Carlo method in MEMS flow predictions," *Computer Modeling in Engineering & Sciences*, 1(4):119–128, 2000.

[34] P. L. Bhatnagar, E. P. Gross, and M. Krook. "A model for collision processes in gases. I. small amplitude processes in charged and neutral one-component systems," *Physical Review*, 94(3):511–525, May 1954.

[35] P. G. Klemens. "Thermal conductivity and lattice vibrational modes," In F. Seitz and D. Turnbull, editors, *Solid State Physics*, volume 7, pages 1–98. Academic Press, NY, 1958.

[36] G. P. Srivastava. *The Physics of Phonons*. Adam Hilger, Ltd., Bristol, 1990.

[37] K. C. Sood and M. K. Roy. "Longitudinal phonons and high-temperature heat conduction in germanium," *Journal of Physics: Condensed Matter*, 5(3):301–312, Jan 1992.

[38] M. G. Holland. "Analysis of lattice thermal conductivity," *Physical Review*, 132(6):2461–2471, Dec 1963.

[39] Lord Rayleigh. *Theory of Sound*. MacMillan & Company, 1896.

[40] R. Berman. *Thermal Conduction in Solids*. Oxford Studies in Physics. Clarendon Press, Oxford, 1976.

[41] T. Ruf, R. W. Henn, M. Asen-Palmer, E. Gmelin, M. Cardona, H. J. Pohl, G. G. Devyatych, and P. G. Sennikov. "Thermal conductivity of isotopically enriched silicon," *Solid State Communications*, 115(5):243–247, 2000.

[42] M. Snir, S. Otto, S. Huss-Lederman, D. Walker, and J. Dongarra. *MPI-The Complete Reference*, volume 1. The MIT Press, Cambridge, 2nd edition, 2000.

[43] M. A. Gallis. personal communication.

[44] E. S. Piekos. "DSMC modeling of micromechanical devices," Master's thesis, Massachusetts Institute of Technology, Cambridge, Sep 1995.

[45] F. M. White. *Viscous Fluid Flow*. McGraw-Hill, New York, NY, second edition, 1991.

[46] M. G. Holland and L. J. Neuringer In *Proceedings of the International Conference on the Physics of Semiconductors, Exeter 1962*, page 475, London, 1962. The Institute of Physics and the Physical Society, London.

[47] J. J. Sniegowski and M. S. Rodgers. "Manufacturing micro-systems-on-a-chip with a 5-level surface micromachining technology,". Presented at the 2nd International Conference on Engineering Design and Automation, Aug 9-12, 1998.

[48] P. K. Schelling, S. R. Phillpot, and P. Keblinski. "Phonon wave-packet dynamics at semiconductor interfaces by molecular-dynamics simulation," *Applied Physics Letters*, 80(4):2484–2486, Apr 2002.

[49] J. E. Palmer, C. V. Thompson, and H. I. Smith. "Grain growth and grain size distributions in thin germanium films," *Journal of Applied Physics*, 62(6):2492–2497, Sep 1987.

[50] G.A. Bird. "Monte Carlo simulation in an engineering context," In *Proceedings, 12th International Symposium on Rarefied Gas Dynamics, Charlottesville, VA, 1980*, volume 74 of *Progress in Aeronautics and Astronautics*. AIAA, Washington, DC, 1981.

[51] K. Koura and H. Matsumoto. "Variable soft sphere molecular model for inverse-power-law or Lennard Jones potential," *Physics of Fluids A*, 3:2459–2465, 1991.

[52] R. E. Peierls. *Quantum Theory of Solids*. The International Series of Monographs on Physics. Oxford at the Clarendon Press, London, 1965.

[53] C. Herring. "Role of low-energy phonons in thermal conduction," *Physical Review*, 95(4):954–965, Aug 1954.

[54] A. Majumdar. "Microscale heat conduction in dielectric thin films," *Journal of Heat Transfer*, 115:7–16, Feb 1993.

[55] J. Callaway. "Model for lattice thermal conductivity at low temperatures," *Physical Review*, 113(4):1046–1051, Feb 1959.

[56] C. J. Glassbrenner and G. A. Slack. "Thermal conductivity of silicon and germanium from 3K to the melting point," *Physical Review*, 134(4A):A1058–A1069, May 1964.

Distribution

External

Stanford University
Building 520
Stanford, CA 94305-3032
Attn: Kenneth E. Goodson

Georgia Institute of Technology
771 Ferst Drive
Atlanta, GA 30332-0405
Attn: Samuel Graham (5)

Center for Microtechnology
c/o Lawrence Livermore National Laboratory
7000 East Ave.
P.O. Box 808
Livermore, CA 94550-9234
Attn: Ray Mariella, L-222

Internal

1	MS 0323	D. L. Chavez, LDRD Office, 1011
1	MS 0825	W. L. Hermina, 9110
1	MS 0826	S. N. Kempka, 9113
5	MS 0826	E. S. Piekos, 9113
5	MS 0826	C. C. Wong, 9113
1	MS 0827	M. Gallis, 9113
1	MS 0834	M. R. Prairie, 9112
1	MS 0841	T. C. Bickel, 9100
1	MS 1310	L. M. Phinney, 9112
1	MS 1411	E. B. Webb III, 1834
1	MS 1421	J. P. Sullivan, 1112
1	MS 9161	C. J. Kimmer, 8763
1	MS 9403	J. M. Hruby, 8700
1	MS 9409	J. R. Garcia, 8754
1	MS 9018	Central Technical Files, 8945-1
2	MS 0899	Technical Library, 9616