Thermal Demonstration of High Density Processor Packages

Submitted By
Peter Houghton Lewis

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

BACHELOR OF SCIENCE IN MECHANICAL ENGINEERING

School of Engineering
Tufts University
Medford, Massachusetts

May 2014

Signature of Author:
Peter Houghton Lewis

Certified By:
Associate Professor Robert D. White
Department of Mechanical Engineering
Tufts University

Committee:
Associate Professor Robert D. White
Department of Mechanical Engineering
Tufts University

Committee:
Dr. Brian Smith
Draper Laboratory

Abstract

As technology continues to miniaturize processor chips, the thermal load of the electronics has become a primary hindrance in the development and implementation of more powerful chips. An internal effort by Draper Laboratory has been made in order to characterize the effect of their high density packaging method on the thermal load. The following thesis provides a thermal characterization of Draper Laboratory's integrated Ultra-High Density (iUHD) material, an extreme miniaturization packaging method. The thesis covers the thermal effects of multiple layers of the iUHD material, the presence of metal traces, the isolation of the testing module's components and the stacking of multiple modules. The thermal contribution of a layer of the iUHD material was determined to be 0.067-0.077 ^{o}C/W depending on the number of layers and the presence of metal traces. The thermal resistance contribution of the components of the second layer of the stack was determined to be 0.24 ^{o}C/W. The isolation of module components was successful and locally raised the temperature in the isolated areas by 4-5 times.

Acknowledgements

I would like to thank Brian Smith for mentoring me through the project and always making the time to meet with me to discuss the experiments. His patience and enthusiasm were infectious, making the project always interesting. I would also like to thank John Burns for the time he spent on various aspects of the project including the laser drilling of the shadow mask for the fabrication of the stack. Doug Gaither was an incredible help, always willing to talk things over or help me find any materials I needed for the research. On the fabrication side of things I would like to thank Dan Frigon and Bill Teynor for taking the time to work with me. Last, I would like to thank Cappy who did much of the background work on the project and whose assistance made the Labview integration work as well as it did.

Table of Contents

Table of Figures

Table of Tables

Chapter 1: Introduction

1.1 Problems to be Addressed/Significance

As technological advances continuously miniaturize processor chips, a prominent problem inhibiting the integration of these chips is the thermal load. The chips become too small for the standard heat sinks to effectively dissipate the heat they produce. Therefore, the chip power is limited by the effectiveness of which heat can dissipate rather than its specifications.

One of Draper Laboratory's Internal Research and Development projects is to quantify the effect their integrated Ultra-High Density material (iUHD), an extreme miniaturization packaging method, has on thermal issues. Many entities in the defense, biotechnology and telecommunications markets have a vested interest in this research as thermal issues impede their developing technologies. Figure 1 shows the difference between the typical high power chip and Draper's iUHD embodiment.

(a) This is a typical high power chip

(b) This is a candidate in Draper's iUHD embodiment

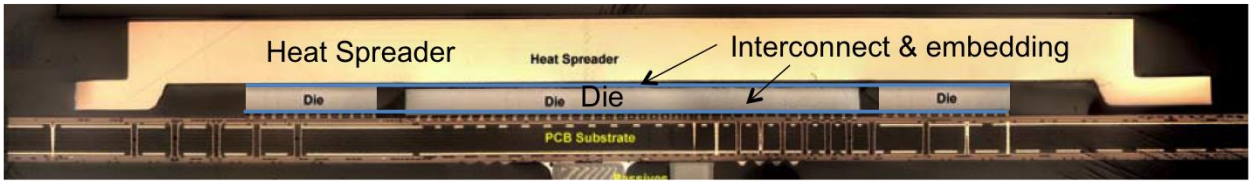

Figure 1: Diagram of difference between typical a high power chip and Draper's embodiment [1]

This thesis provides a thermal characterization of the iUHD material in multi-chip silicon modules. The comparison between modules with bare silicon and varying layers of insulating material and metal traces, which constitute electrical connections between the chips and outside of the module, will help determine the most effective way to promote the cooling of the chip. Potential ramifications of this research include the direct applicability to modules being developed by Draper allowing for their more efficient design through thermal model prediction. Another key inquiry is how the stacking of these modules intensifies these issues. This research is unique in exploring the thermal impact of the novel approach embedding multiple silicon chips (processors, memory) in a silicon carrier wafer and electrically integrating with semiconductor-based dielectric and metal patterning.

1.2 Relevant literature

A wide variety of research has been conducted in the field of heat transfer through high-density modules. The traditional approach employs a variety of methods ranging from air to liquid coolers with higher efficiencies observed in the liquid coolers [2]. The contact between a liquid cooler and the device can also be improved with thermal adhesive. Thermal adhesive improves the conductance between the cooler and module surfaces promoting heat transfer. However, researchers have also investigated a more efficient means of conducting heat away from the chip using various materials and geometries. One recently researched method of managing the thermal burden in high-density electronics is the use of carbon nanotubes, a material that has a thermal conductivity on the order of 3000 W/Km [3]. Other systems have sought to replace expensive high thermal conductivity materials with cheaper options such as thermal spreaders made from heat pipes [4]. The thermal challenge has become a predominate problem resulting from the shrinking of electronics.

The thermal management of three-dimensional stacks of silicon chips has also been extensively investigated. Zanini et. al. developed a controller for the thermal management of such stacks using active control of on-chip switching rates and active interlayer cooling with pressurized fluids. Their research showed energy savings up to 50% of the state-of-the-art liquid cooling techniques [5,6]. Figure 2 depicts the interlayer cooling technology.

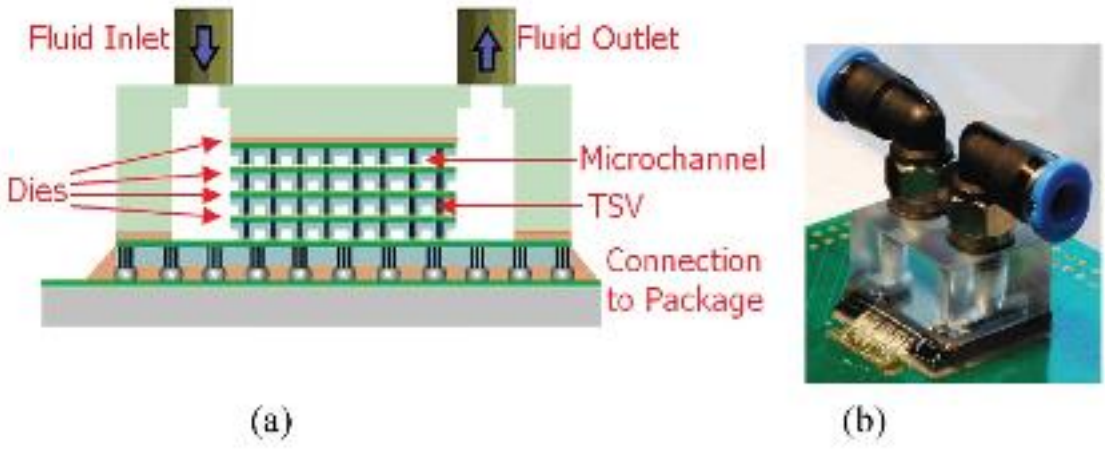

Figure 2: Manufactured prototype and cross section of a test stack with interlayer liquid cooling. (a) Cross Section [5] (b) Prototype [6]

Other efforts to manage the thermal load on multichip devices have revolved around the use of other materials such as SiC instead of Si. SiC has a high range of temperature of operation, approximately 5x higher than that of Si [7].

Packaging technologies are another forefront in which thermal issues are being challenged. Research on package improvement has been taking place over more than two decades. As can be seen in Figure 3, miniaturization of chips has resulted in drastic changes in the Hughes packaging technology from 1970-1990 alone. As is shown, the number of gates per module has increased by a factor of 20 allowing the package to be much smaller than previous packages. This miniaturization of components

increases the thermal load by reducing surface area for the heat to dissipate and be managed through heat sinks.

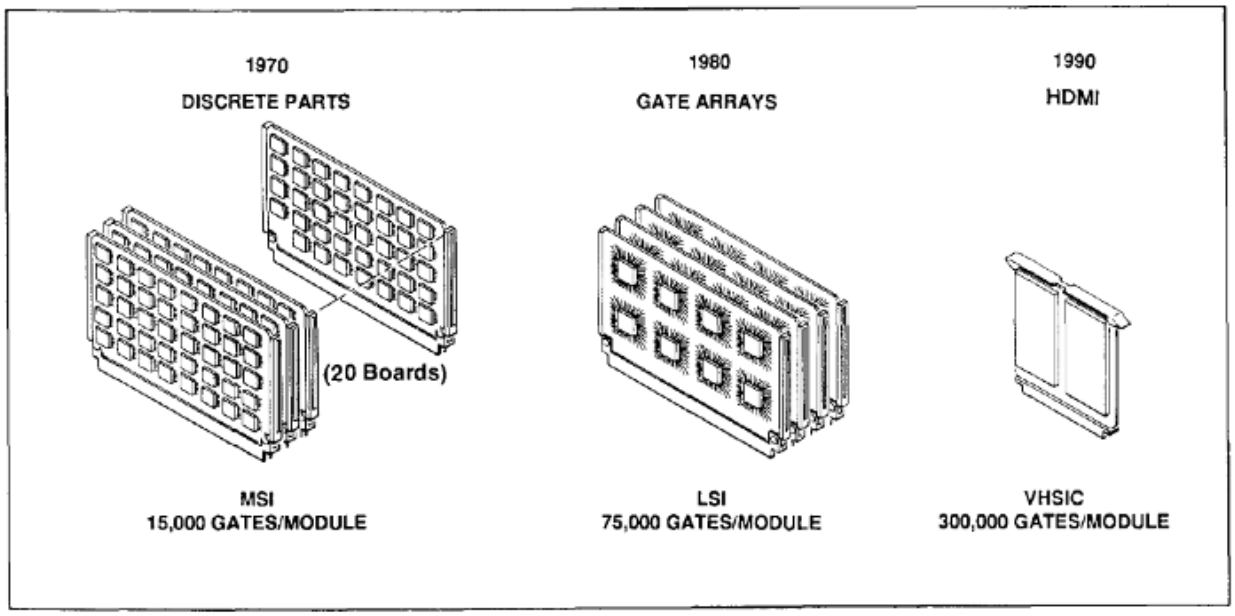

Figure 3: High-Density Multichip Interconnect Modules: Three generations of Hughes packaging technology [8]

Companies such as Lockheed Martin have been working on improving packaging techniques for years. One of their older technologies was the development of Lockheed Martin High Density Interconnect (HDI) Technology, which used a multi-layer routing to eliminate traditional substrates from multi-chip modules [9]. More recent efforts to improve the power levels and functionality of these chips as they continue to shrink include the development of Draper Laboratory's novel Integrated Ultra-High Density (iUHD) multi-chip integration technology. The main motivation of this thesis is determining the thermal effect of the spin-on dielectric material inherent (SOD) to iUHD in conjunction with metal that forms the electrical traces.

The thesis is organized in the following sections. First, a chapter will be devoted to the design of the thermal test system used to investigate the thermal effect of the SOD and the electrical traces. Following this, a chapter will present the experimental results. The thesis will end with a concluding chapter presenting main findings and recommendations for future work. The Appendix includes the supplementary sections covering heater and diode calibration, Labview data interpretation program design, the fabrication of the module stack, the error analysis and more thermal test data.

Chapter 2: Design

This chapter details the design of the thermal test system, which was developed at Draper. The module was designed to emulate the iUHD modules from a thermal perspective. Several key design features include controllable heat dissipation, high spatial density on-chip temperature sensing, a cooling scheme that allows the isolation of the effect of the interconnect layers, and a modular data acquisition (DAQ) that allows easy testing of many different configurations. An "interconnect" layer is defined as a layer of commercial spin-on-dialectric (SOD) insulator plus photopatterned copper traces, roughly 7 μm and 1 μm thick, respectively. One layer is electrically connected to another by drilling vias in the SOD and plating the next layer of interconnect into the holes of the underlying SOD. The following chapter includes sections on the mathematical theory of the thermal resistance network, module fabrication, electrical design, and the design of the module test stand.

2.1 Mathematical Theory of the Thermal Resistance Network

This section discusses the mathematical theory used to determine the thermal resistance of the interconnect layers. First a cross section of the module is shown so the components of the heat path can be observed. This is then generalized into a thermal resistance network, which is then analytically solved.

The high-density chip module is designed to measure the thermal resistance of varying interconnect layers between a heat source and a heat sink as shown in Figure 4. This is done by first calculating the total thermal resistance contribution of the bare silicon module. All the constituents of this module are present in the other chip designs. Therefore, the total thermal resistance of this module

includes the thermal resistance of the thermal interface material (TIM), the cooler and the silicon. Once this value is determined, the contribution of the SOD layers can be determined using the assumption of constant thermal resistance contribution of the TIM, cooler and silicon between the modules. Thus the thermal contribution of the SOD layers in a 4-layer module without metal traces is its total thermal resistance minus the thermal resistance of the bare silicon module case. In theory, the only thermal difference between these two modules is the presence of the SOD layers. Therefore the last unknown is the thermal resistance of the SOD, as can be seen in the generic thermal resistance model shown in Figure 5. Potential error resulting from this assumption is discussed in Appendix A.4. The TIM is Omega Therm 201 thermally conductive paste. The uniform application of the TIM was a challenge due to its viscosity and the lack of uniform pressure across the cooler applied directly to it. A laser profilometer was used to determine the thickness of the TIM. After many individual measurements, it was determined the average thickness of the TIM was $7.08 \cdot 10^{-5}$ m with a 30% variability. Appendix A.4 analyzes the effect of this variability.

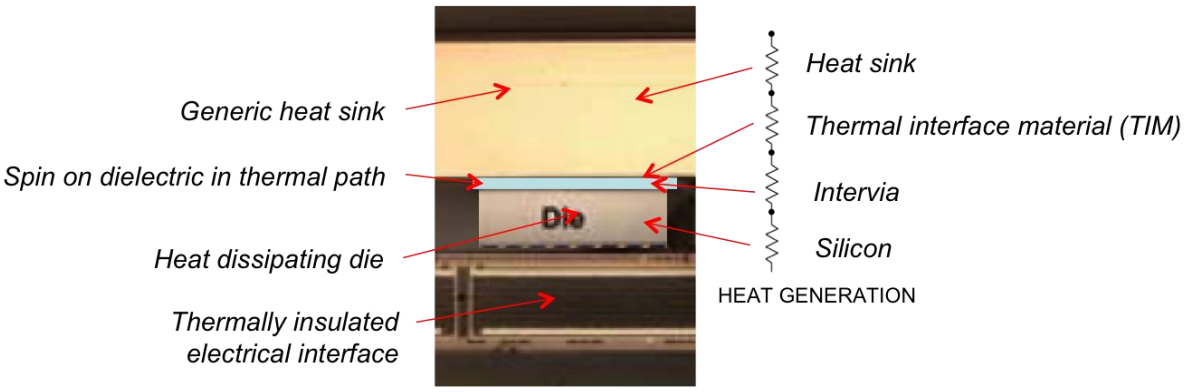

Figure 4: Thermal module cross-section [1]

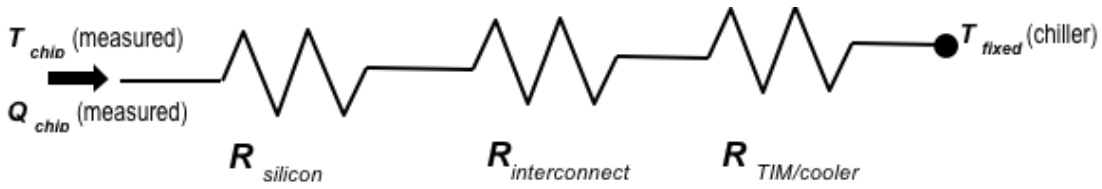

Figure 5: Thermal resistance path for testing module [10]

Another crucial assumption for the mathematical model is the ambient temperature is that of the cooler. Therefore the heat sink temperature is known while the calibrated diodes signify the temperature of the source. Similarly, it is assumed that the heat flow is 1-dimensional. This assumption thus states that heat travels solely from the source to the sink through the material layers.

The following is the mathematical theory for determining the thermal resistance of the contributing layers taking into account the previously stated assumptions. After, the method of lumping together the resistances determined constant between modules and treating them as a bulk resistance will be discussed as this was the method employed during testing.

The governing equation for determining thermal resistance is (1) since the thermal path can be simplified to the path shown in Figure 5.

$$R_{tot} = \sum_i R_i \quad \text{with}$$

$$R_i = \frac{t}{kA} \quad \text{for conduction through a 1D layer} \qquad (1)$$

$$R_i = \frac{1}{hA} \quad \text{for convective resistance at a surface}$$

Where R_{tot} is the total resistance, R_i is the conduction resistance of the i^{th} component, t is the component thickness for a layer component, k is the coefficient of thermal resistance for a layer component, h is the convective resistance at a surface of the cooler, and A is the module surface area.

First, the contribution of the components on the bare silicon chip must be determined. The total thermal resistance of the heat path in a bare silicon module is shown in equation (2). Since the heat flow through bare silicon has been heavily studied, k is known and the thickness was determined from the wafer data sheet. The thermal interface material likewise has well documented thermal properties, so the thermal resistance of the grease is known. In order to determine the thickness of the TIM, a Metralight TLE1 Triangulation Laser Sensor was integrated into the test stand for the measurement of this property. Therefore, the thickness of the TIM is a known value with the uncertainty addressed earlier. Lastly, the test data from the bare silicon chip determines the total thermal resistance of the chip or R_{tot}. Since we already know the area of the module face, every term in equation (2) is known except h, the thermal resistance of the cooler, which is treated as a surface component. It is important to note that the bare silicon module has no interconnect layers so it has no interconnect resistance. Contact resistance is assumed negligible and, furthermore, equal in bare silicon and interconnect cases, so it can be integrated into the cooler resistance estimate.

$$R_{tot} \times A = \left(\frac{t}{k}\right)_{Si} + \left(\frac{t}{k}\right)_{TIM} + \left(\frac{1}{h}\right)_{cooler} \quad (2)$$

So solving for the unknown h results in equation (3).

$$\left(\frac{1}{h}\right)_{cooler} = R_{tot} \times A - \left(\frac{t}{k}\right)_{Si} - \left(\frac{t}{k}\right)_{TIM} \quad (3)$$

Once this value is determined, an assumption is made that the thermal resistance contribution of the cooler will be equivalent across all tests and chips. This assumption means *h* is now a known value and can be used in the determination of the thermal contribution of the interconnect. It is important to recognize that *h* is a complex function that is usually empirically determined. It depends on many factors including the difference in module vs. ambient temperature. However, the assumption that *h* is approximately the same in different cases was made due to the flow rate and temperature of the cooler being constant for all the tests. For the determination of the thermal resistance of the interconnect, R_{INTER}, equation (4) is used. This equation is also based off of Figure 5, however this time there is interconnect layers and therefore an associated resistance.

$$R_{INTER} = R_{tot} \times A - \left(\frac{t}{k}\right)_{Si} - \left(\frac{t}{k}\right)_{TIM} - \left(\frac{1}{h}\right)_{cooler} \quad \text{with}$$

$$R_{INTER} = \left(\frac{t}{k}\right)_{INTER} \tag{4}$$

Since the thickness of the interconnect is known from the processing steps, there is one unknown, k_{INTER} to solve for. This process can be repeated for all the modules.

However a simpler method was used since the thermal resistance value sought was that of the interconnect. Therefore, as discussed earlier, the total resistance of the bare silicon module was subtracted from the other testing cases to see the contribution of the interconnect layers. Figure 6 visually shows the method. In reality, some of the resistances that are assumed to be fixed actually vary. For example, the variation in the thickness of the TIM makes the resistance of the TIM variable between applications. Figure 7 better exemplifies the resistive contributions of the module components. The error analysis in appendix section A.4 discusses the resulting error from this assumption.

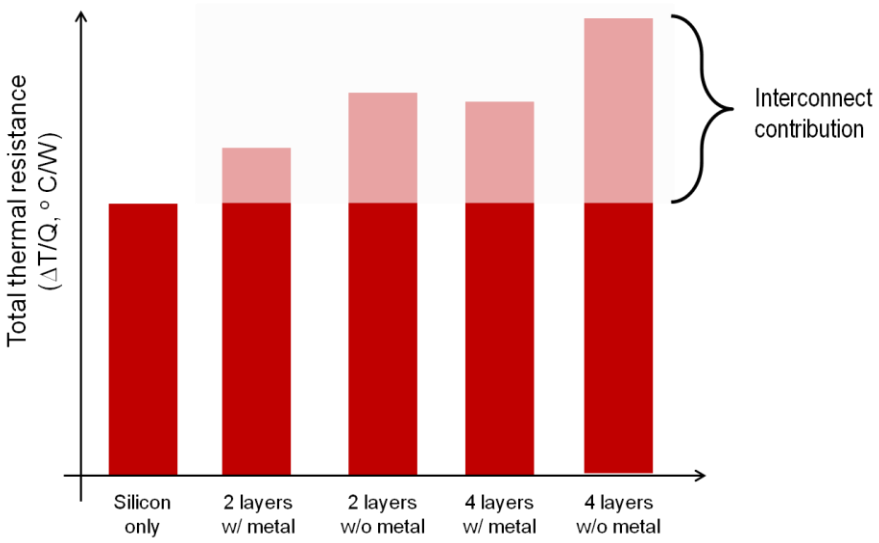

Figure 6: Theoretical contribution of interconnect layers [10]

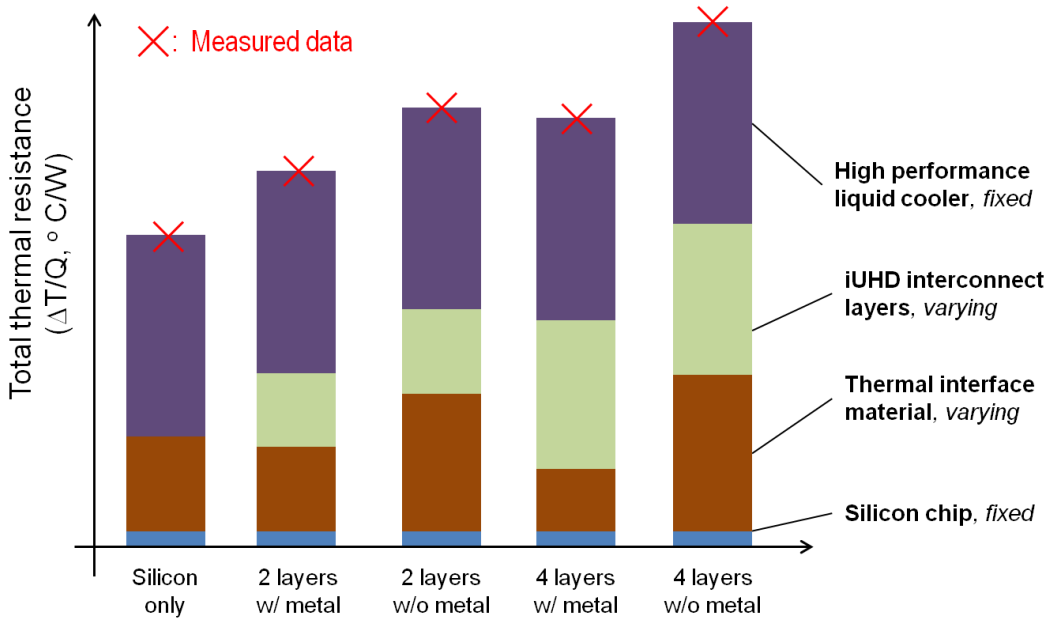

Figure 7: In reality interconnect layers vary along with TIM thickness [10]

2.2 Module Fabrication

This section describes how the previously discussed modules for the thermal testing are created. First the commercial-off-the-shelf (COTS) heater components, fabricated by Thermal Engineering Associates Inc. (TEA), that are built on the wafer are discussed in detail. Then the fabrication process for the development of the module layers over the processed wafer is discussed. Lastly the different types of modules created for testing and how they are identified is discussed.

One of the novel aspects of the system is that an embedded multichip module can be emulated using a single silicon chip due to this high spatial resolution of the heater control and temperature sense. This multichip solution uses commercially available heaters from TEA so that the heat flow through the layers of the module can be analyzed. The utilized TEA heater chips have a part number of TTC-1002. A die map of a single die is shown below in Figure 8. Note that two of these individual die make up the “unit cell” referred to in other sections so that an array of 8x8 dies or 4x8 unit cells constitute a module.

The assembled module consists of several key components: the interposer, the populated silicon wafer, and the interconnect as shown in Figure 9. The interposer is a printed circuit board which functions as the path through which all electrical signals must go through to and from the chip. The TEA heater chip provides the capability to selectively heat parts of the chip for thermal analysis on spreading and isolation. On the heater chip is a thermally sensitive diode, which is probed for voltage to determine temperature change. Lastly there is a silicon layer with or without interconnect. The interconnect consists of SOD and metal layers.

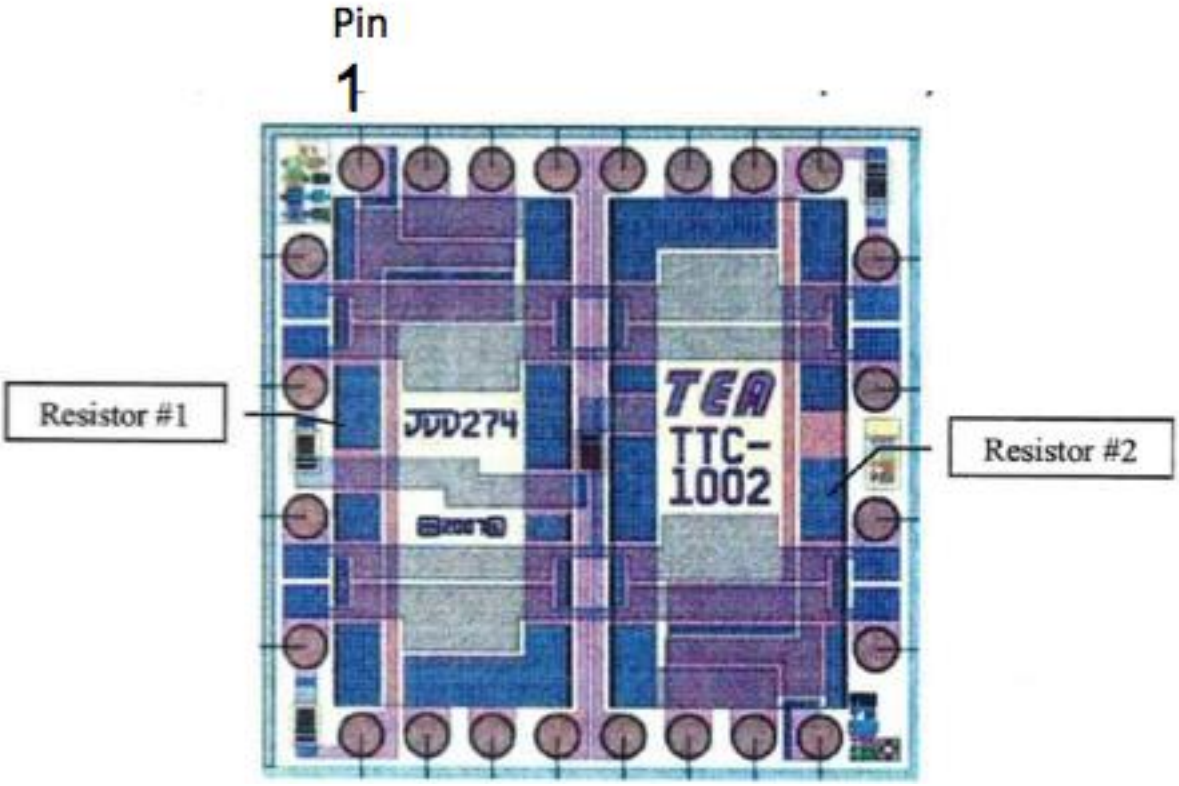

Figure 8: Top View of chip produced by TEA Inc.

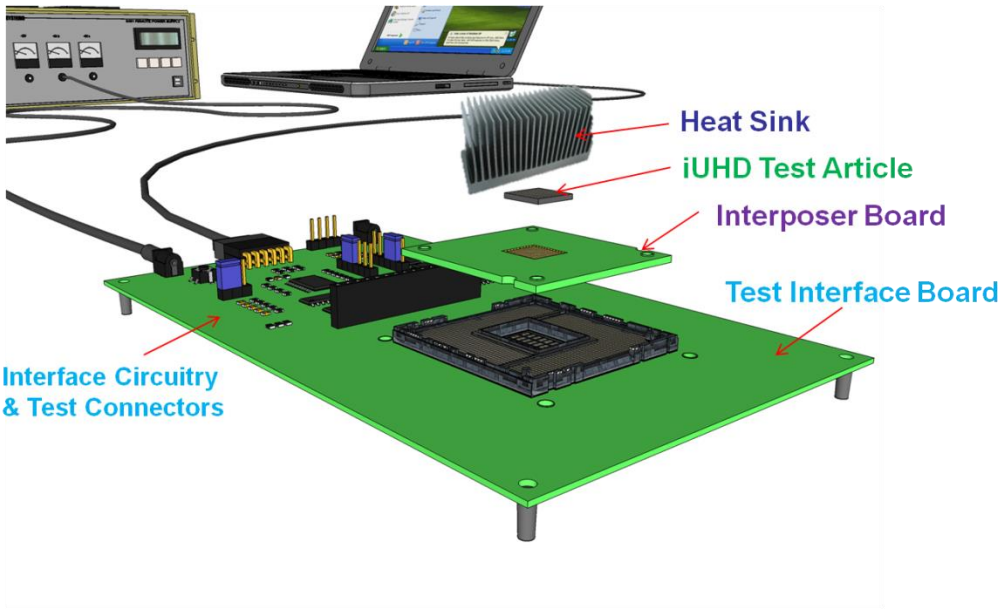

Figure 9: Embodiment of thermal testing stack up [1]

Five different chip types were fabricated as later detailed in Table 1. One module simply had the TEA die with “front side” interconnect to interface the chip to the interpose and no iUHD interconnect in the thermal path. This was used to determine the contribution of the silicon, thermal interface material (TIM), and cooler to the thermal resistivity of the system. This has been previously referred to as the bare silicon case. As shown previously, this calculation was important in order to determine the effect of metal and high-density layers on thermal resistivity. Therefore, the silicon wafer was patterned with different layers to provide the five types of chips that were to be investigated as shown in Figure 10.

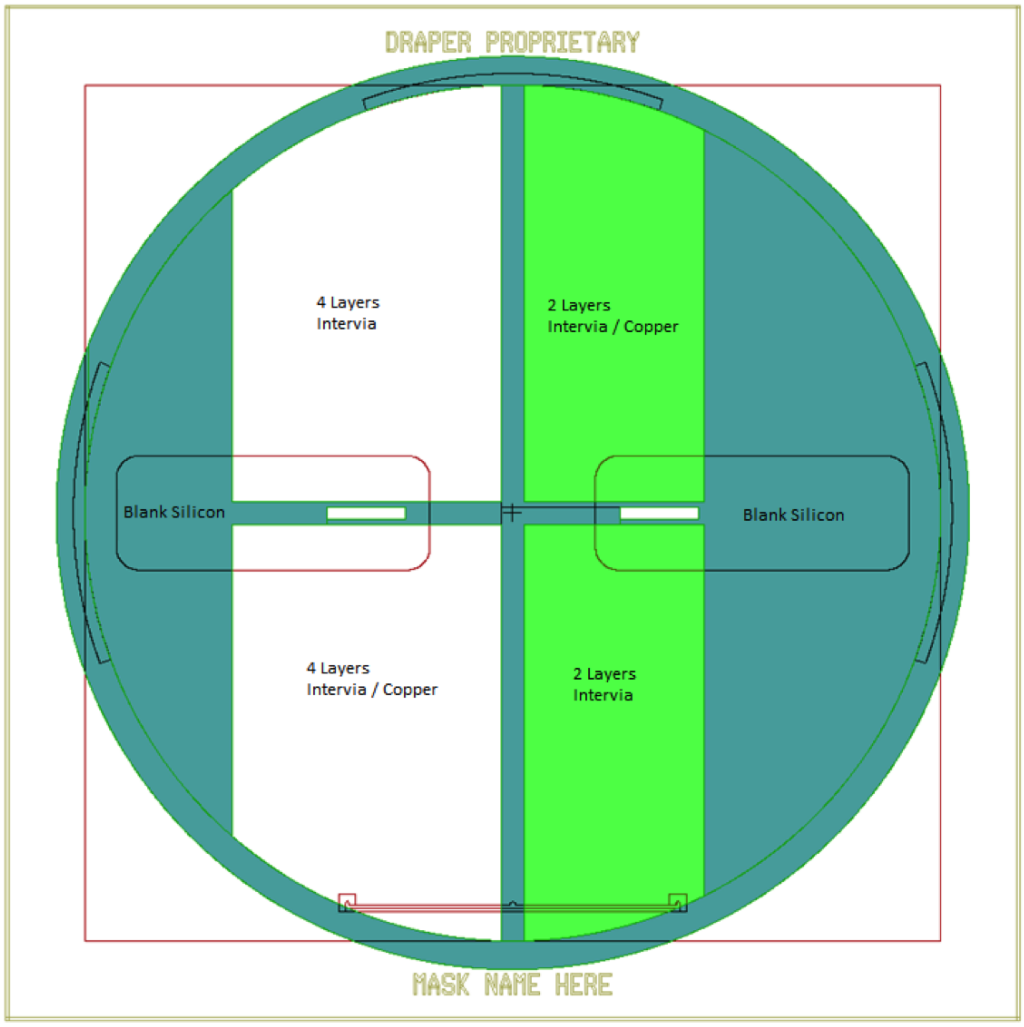

Figure 10: Schematic of the backside of the wafer post-processing created by John Burns IV [11].

A high level overview of the process to make the wafers is shown in Figure 11. Note that the intervia referred to in the figure is a commercial spin-on-dialectric (SOD). The silicon wafer is bought from TEA with the heater and diodes already populated on the wafer. The wafer is then processed to obtain the final chip to be tested. First the wafer gets a coat of SOD on the active side of the wafer. Then the SOD is patterned using photolithography. A laser is used to drill around 15,000 vias through the chip. A Ti/Cu/Ti layer is then deposited, patterned and etched followed by another coat of SOD. The SOD is engineered to coat with uniform thickness and not fill in holes. The SOD is then patterned again followed by a hardbake and crosslink process. After these steps are completed, the wafer is flipped over

to process the backside. The backside is coated with SOD then patterned once again. Then the layers are built up on this side creating the five different cases using contact masks. These cases are discussed at the end of the section. Lastly the wafer is flipped back over and the soldermask is laser drilled. The soldermask is a patterned layer of the SOD. The completed wafer is then diced to form and 8x8 array of active die, thus forming the chip size compatible with the thermal testing stand. The wafer has 100 μm streets designed for the dicing operation. Since the wafer was patterned with gold pads that form a ball grid array, adhesion to the interposer board is facilitated after the chip is bumped.

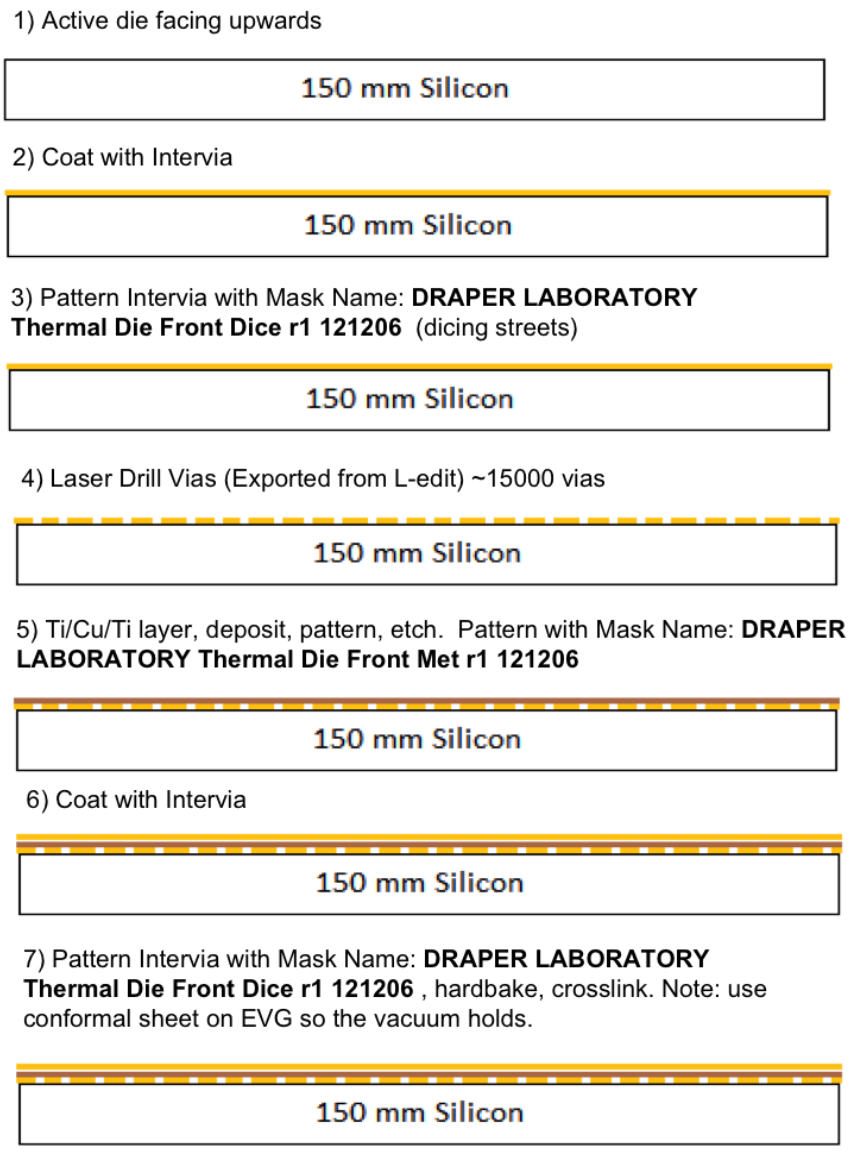

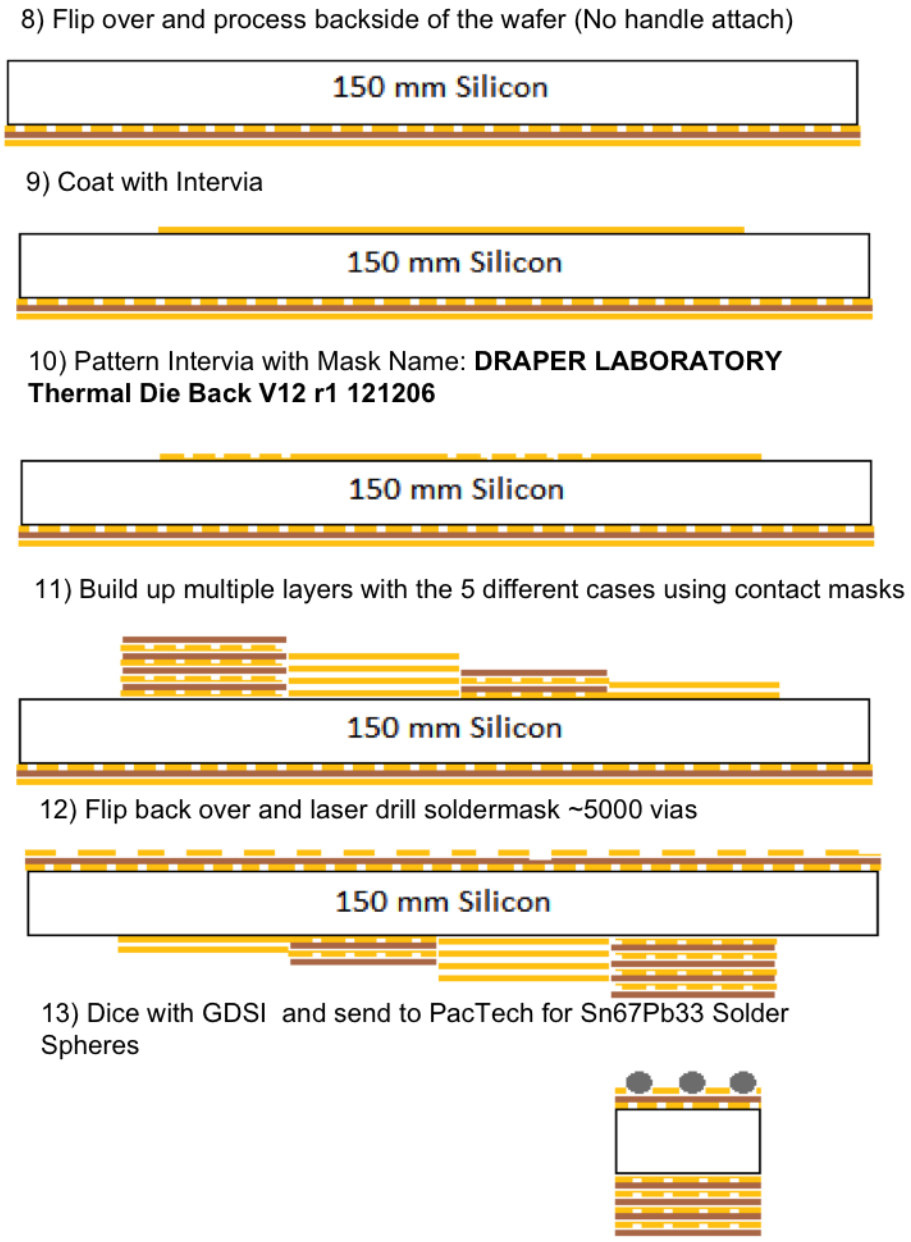

Figure 11: High-level process flow for chip fabrication created by John Burns IV. [11]

After processing, the silicon wafer is diced and bumped with solder balls so that the wafer components can be mounted on the PCB. After alignment between the PCB and the processed wafer, the components are reflowed in an oven resulting in the assembled chip module.

Table 1 describes the composition of the 6 different types of modules used during testing. The differences in the modules are the number of SOD layers, presence of metal traces and whether the module consists of a stack of wafers.

Module	Number of SOD layers	Metal (Y/N)	Number of stacked wafer pieces
1 and 6	0	N	0
2	2	N	0
3	4	N	0
4	4	Y	0
5	2	Y	0
7	0	N	1

Table 1: Composition of the six different wafer types

2.3 Electrical Design

This section discusses the electrical design of the fabricated module. One module contains 64 diodes in an 8x8 array and 32 heaters in a 4x8 array diced from the fabricated wafer in section 2.2. This section includes the operational theory of the device, which involves the use of heaters and diodes to stimulate and read the thermal state respectively. The diodes are linearly sensitive to temperature with a constant current supply, which can be observed in the voltage change.

64 diodes are used for measuring temperature through voltage change. As the chip heats up, the voltage drop across the diodes for a constant current linearly reduces allowing the temperature change to be determined after calibration. Diodes where chosen as the means for measuring temperature due to the ease of monolithic fabrication along with high sensitivity not seen in other temperature sensors such as thermocouples. The implementation of diodes for temperature

measurement is a standard practice and is compatible with standard chip processes. Data is received by DAQ devices and monitored through a Labview program, as detailed in section 2.4. There are 32 controllable heaters groups on each chip. The heater diodes are monolithically fabricated on top of each other in the TEA silicon die using standard semiconductor foundry processes. This allows specific regions of the module to be selectively heated. The diodes and heaters are mapped onto the chip face as shown in Figure 12. This figure shows the adjacent heaters wired together and the access points for the 4-point probe of the thermally sensitive diodes. The white rectangles with a black outline are pads, the circles with red outlines are probing points and the arrows are the diodes. The diodes used are in the center of each square unit. The probing points are for the four point probing method where current is driven from one pad through the diode and to another pad and the voltage drop is read through the other two pads. The diodes at the top left of each unit and at the bottom of each unit are not used because there is no need for that much information especially since it complicates the interface to the DAQ.

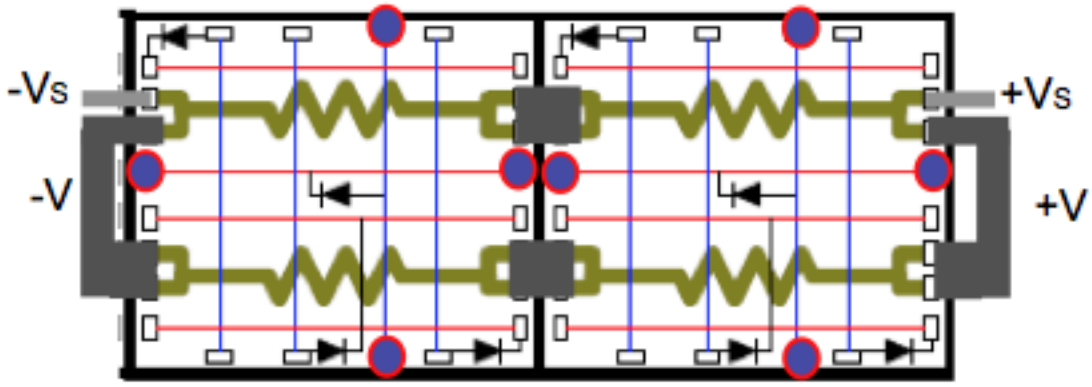

Figure 12: Two diodes with the heaters wired together such that one switch turns on the heater for both chips [1]

Each unit cell of the heater/sensor array contains 2 heaters and 4 probe areas for temperature sensing. Adjacent heaters are connected together resulting in an array of 4x8 heaters and an 8x8 array of thermally sensitive diodes. Through the observation of the voltage drop across the individual diodes

for a constant current, temperature change can be determined. An analysis of the accuracy of this method is discussed in Appendix section A.4.

The voltage drop of the thermally sensitive diode is determined by using the 4-point probe method to supply constant current on two of the wires and measure voltage locally across the diode with the other two. This method increases accuracy of the readings since it effectively eliminates the resistive contribution of sources other than the thermally sensitive diode.

The design and the wiring of the heaters serve a multitude of purposes. For one, the wiring allows the individual control of the heaters, which results in a higher density of information from the chip. Through the individual heater sites, the overall power consumption of the device can be accurately determined and the amount of power that is being generated by a heater local to each pair of diodes can be observed. The accuracy of the power measurements involving the heaters is discussed in Appendix section A.4.The diodes consume negligible power since they draw approximately 0.7 mW. Since each diode is read individually and not simultaneously, 0.7 mW is the total loss of power through the diodes. The power being driven through the chip is over 4 magnitudes higher than this value. There are no more losses from the setup due to the utilization of the 4-point probe method. This eliminates the resistances of the thermal test stand components from the data obtained from the module. However, to improve accuracy, quantifying the resistance of the thermal test stand can be done in future work. The power supplied to the thermal test stand was greater than the power received by the chip heaters for this reason. The potential areas for power loss in the system include the thermal test stand PCB connections, the solder balls, the interposer, and the wires to the DAQ. Anecdotally, the bulk of the resistance appears to be on the thermal test stand board as some electrical lines were fried under high power levels. The investigation into these resistances may indicate whether the cooler is pulling out more heat from the module than is given off by the heaters. A mass flow analysis can be done to see

if the heater is indeed taking away extra heat. Therefore this future work would increase the accuracy and sensitivity of the device.

The TEA chip was designed so that heater resistances optimized current and voltage requirements. If the heater resistance were too low, a high current would be needed to obtain a given power level which would make power delivery instrumentation more difficult and potentially adversely affect the system through electromigration. If the resistance were too high, the voltage required for a given power level would be prohibitive to the power supply. The resistive heating produced by driving a fixed resistor with a DC voltage or current source is described by Ohm's Law and the electric power equation

$$P = \frac{V^2}{R} \tag{5}$$

And

$$P = I^2 R \tag{6}$$

where *P* is power in watts, *V* is the voltage in volts, I is the current in amps and R is the resistance in ohms. Therefore the V^2 and I^2 terms must compensate for all changes in *R* in order for *P* to remain constant.

Figure 13 is an IR image of the un-sinked module that shows the 8x8 grid of diodes at nominal operation near 100°C.

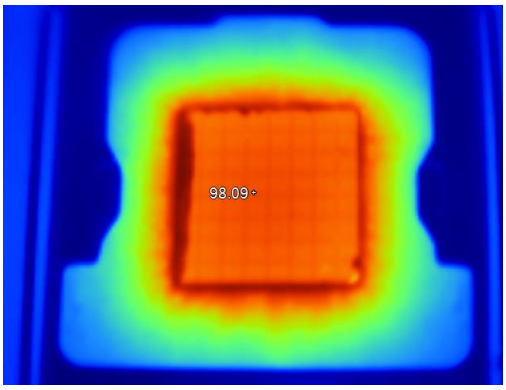

Figure 13: Thermal picture of module with all the heaters on with a surface temperature of 98.09°C

2.4 Design of Module Testing Stand

The testing stand was designed in order to facilitate independent heater control and individual diode measurement. These capabilities result in a high density of information coming from the module. 64 different locations are being actively probed for temperature and 32 heaters can be manually controlled to heat regions of the chip. The setup is shown in Figure 14.

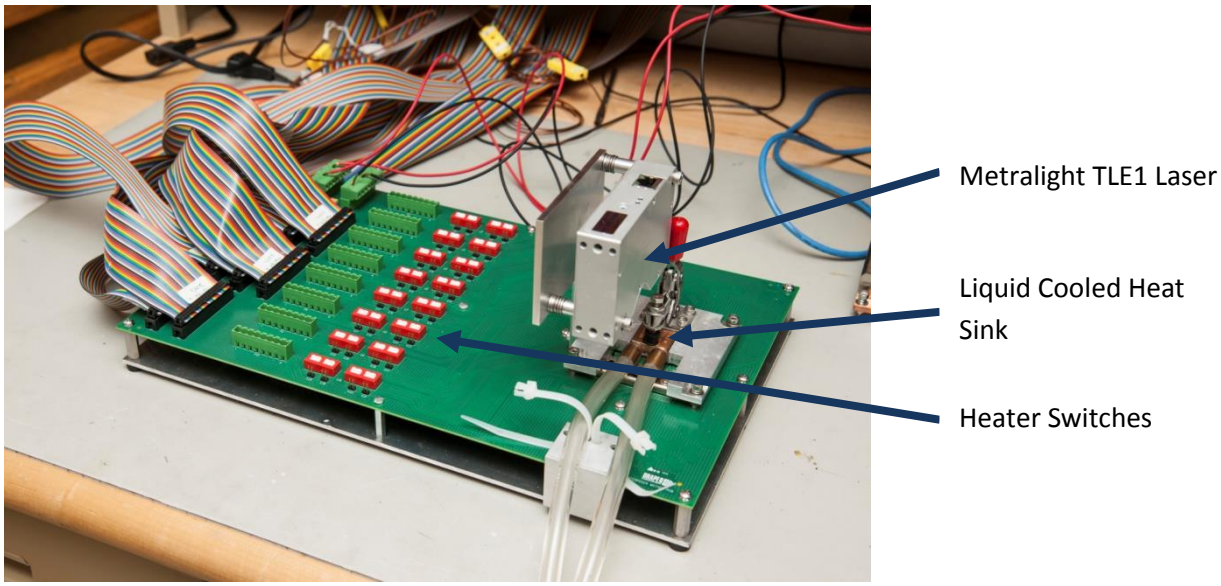

Figure 14: Thermal Test Stand

A Metralight TLE1 Laser was used to provide an estimation of the thickness of the thermal interface material (TIM). The tool has an advertised resolution to 1 µm. The way the instrument was used is shown in Figure 15. The thickness of the TIM was needed in order to reduce the variables in the 1-D thermal equation described in section 2.2. The laser was used to first read the surface height of the module without TIM then again once TIM was applied. Due to its design, the height was read in respect to the surface of the laser mount, giving a consistent standard to measure against. This allows the deduction of the thickness of the TIM. The surface topology of the TIM was not uniform; however an average height across the lasers span was used to estimate the true TIM thickness. Since this technology is in its infancy, the program that operates the laser had glitches making it hard to use. Repeatability of the measurements was often a concern as variations were observed in measurements that should be similar. It was therefore decided that the thickness for the TIM would be averaged over many trials and the resulting average thickness would be used for the thermal calculations.

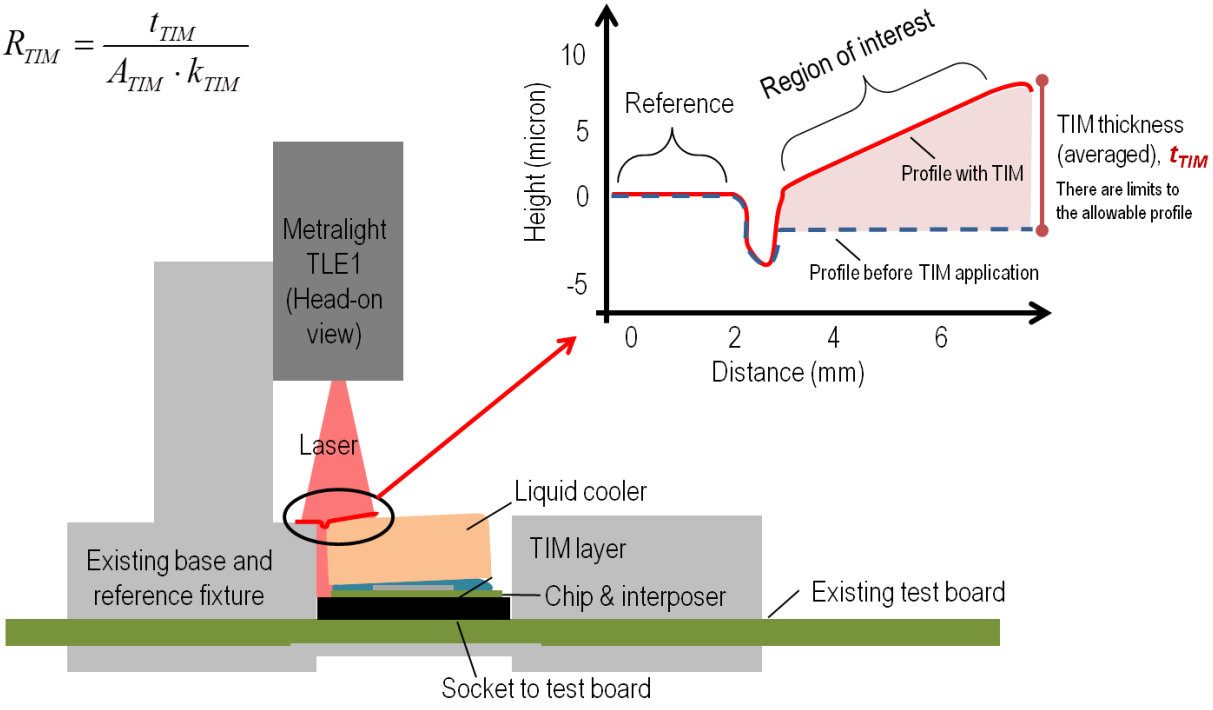

Figure 15: Graphic showing use of Metralight Triangulation Laser [10]

The testing stand also has a bank of red heater switches running down the middle of the board. Each heater corresponds to a region on the chip module. Flipping the switch to the on position causes the associated module region to heat up resulting in heat spread. A mapping of the heater switch to chip surface location is shown in Figure 16. This figure is divided into colored regions to facilitate the ease of reading and representative of the trench locations discussed in section 3.3. This will be utilized later in section A.1 as part of calibration and section 3.3 during the thermal isolation experiment. Note that the heater switches are labeled in Figure 14.

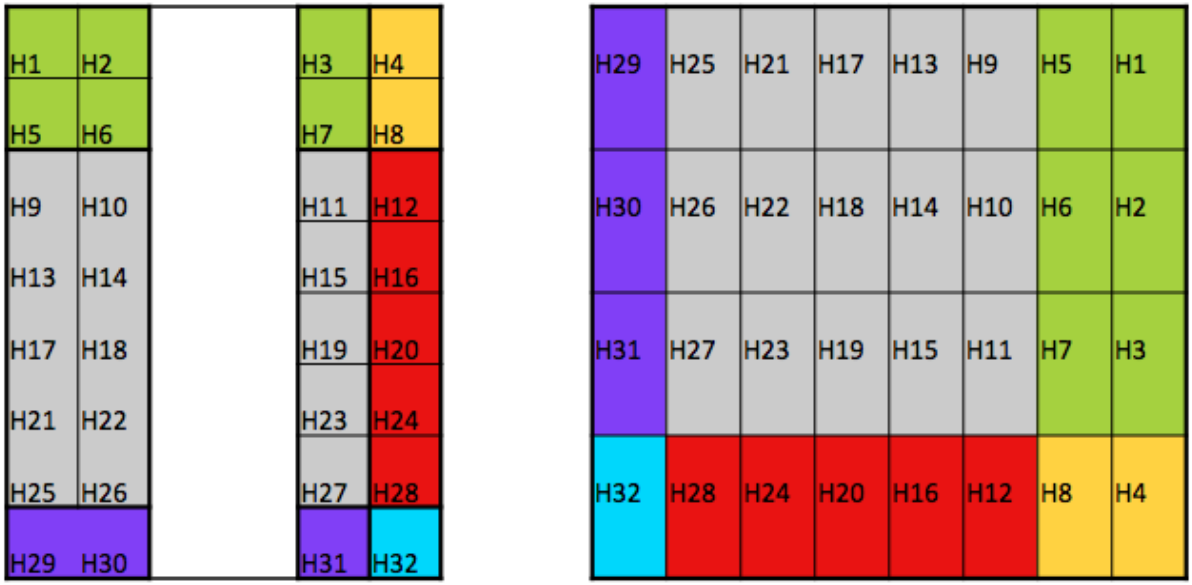

Figure 16: The mapping of heater switches to the module surface. The left diagram represents the heater switches and the right diagram represents the corresponding heater locations looking down at the module face.

Rainbow ribbon cable makes the connections between the test stand PCB and the multimeters. Three three-level multimeters are used to interpret the signals. The multimeters are then hooked up to a Labview program where the data is read and interpreted. Figure 17 shows the overall design of the thermal test stand.

The testing stand was designed in order to facilitate the quick exchange of the five different types of chip modules in the structure. The primary method that the modules were made easily interchangeable was the latch down mechanism shown in Figure 18. The mechanism press fits the back of the chip module onto pins on the thermal test board thus completing the electrical circuit.

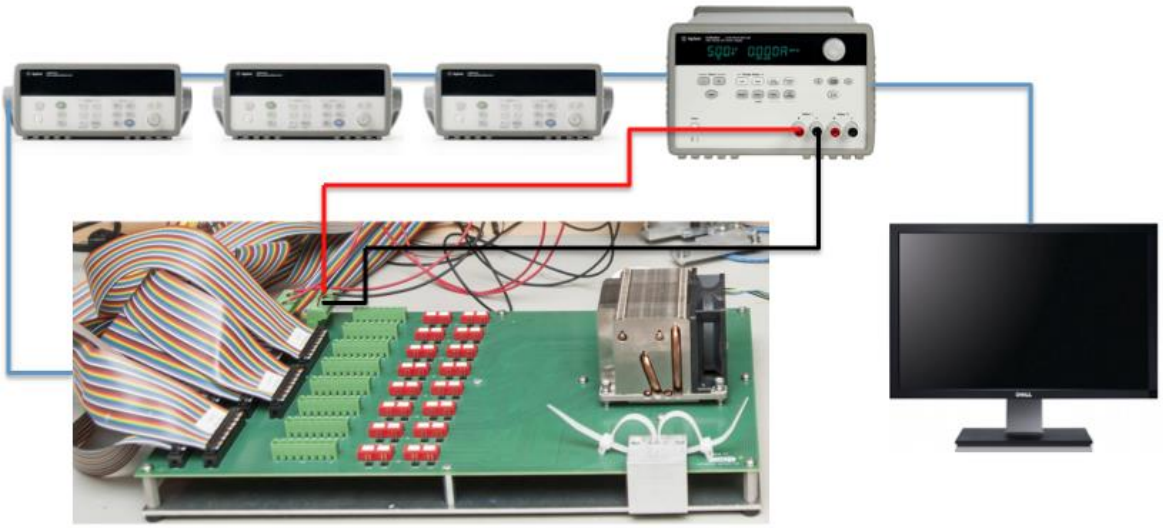

Figure 17: Overview of testing apparatus

Figure 18: The latching mechanism clamps down around the edges of the module creating an electrical path between the module and the PCB

The copper heat sink, shown in Figure 19, is secured on the top of the chip module by a rubber stopper mounted on a metal frame as shown in Figure 20. The metal frame also functions to align the heat sink over the surface of the module. The stopper is mounted on a screw, so the pressure applied to the heat sink can be adjusted. The fixture was also made so that other types of heat sinks could readily be switched onto the Thermal Test Stand board. Two Intel heat sinks, the recommended 2U heat sink with an integrated heat pipe and a fan and the recommended 1U heat sink without a fan, were used in other experiments comparing the effectiveness of heat sinks. This data will not be shown in this report due to repeatability concerns. However, the thermal test stand was designed to incorporate a variety of heat sinks.

Figure 19: The liquid cooled heat sink used on the thermal test stand

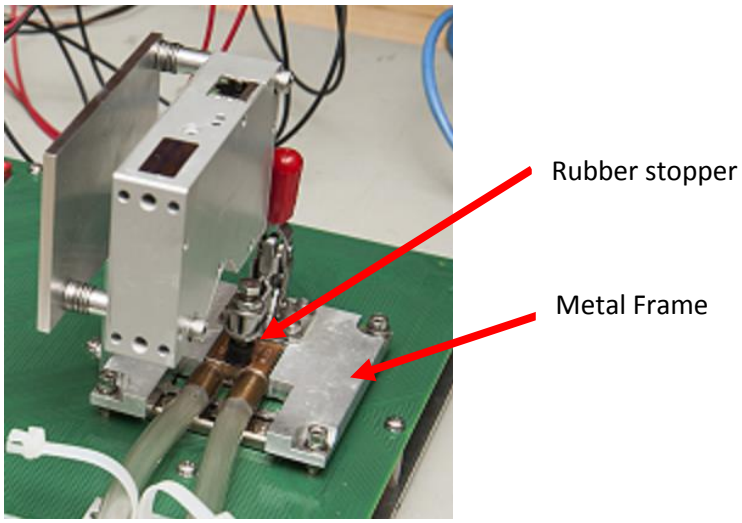

Figure 20: Components of the module attachment to the thermal test stand are shown above. The rubber stopper and the metal frame secure the cooler during testing.

The thermal test stand also has a bank for the power supply to connect to the board. The power supplied was driven by either the voltage set or the current set, depending on the resistance of the system during the test. For example, if one heater was being tested, the current set point would be the actual current but the voltage would only be what the test stand drew. However, if all the heaters are being tested, the voltage set point would be the actual voltage and the current would be driven by the draw of the board.

A Labview program interprets the data obtained by the multimeters. The structure of the Labview program is discussed in Appendix section A.2.

Chapter 3: Experimental Results

3.1 Results of Diode and Heater Calibration

The calibration of the diodes and the heaters on each of the modules was conducted at least twice to ensure consistent results. Each calibration resulted in a slope and intercept for the linear model of the dependence of diode resistance on temperature. Since all the modules have diodes and heaters fabricated by the same process, variation between them should be minimal. However, there are differences that arise and can be seen in the data. Figure 21 and Figure 22 show the diode calibration and the heater calibration respectively. As can be seen, there is a range of calibration curves. Each calibration file for a module contains 32 or 64 slopes and intercepts depending on the calibration type. It should be noted the sensitivity of the heaters to temperature change is significantly smaller than the diodes as they are not designed to vary significantly with temperature change due to their function as a heat source.

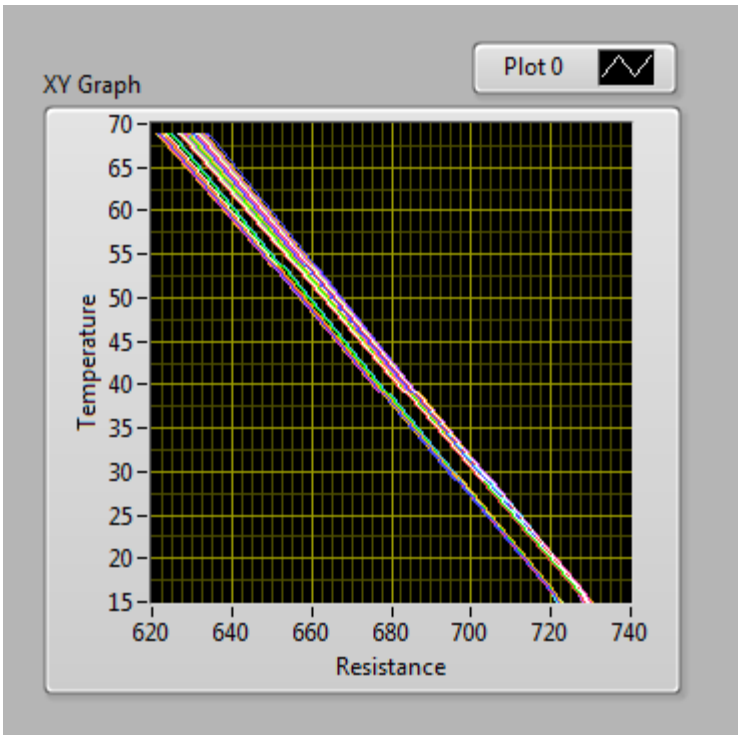

Figure 21: Calibration plot for the 64 diodes of one of the modules. Each diode has a specific calibration line. Each module has a similar looking plot.

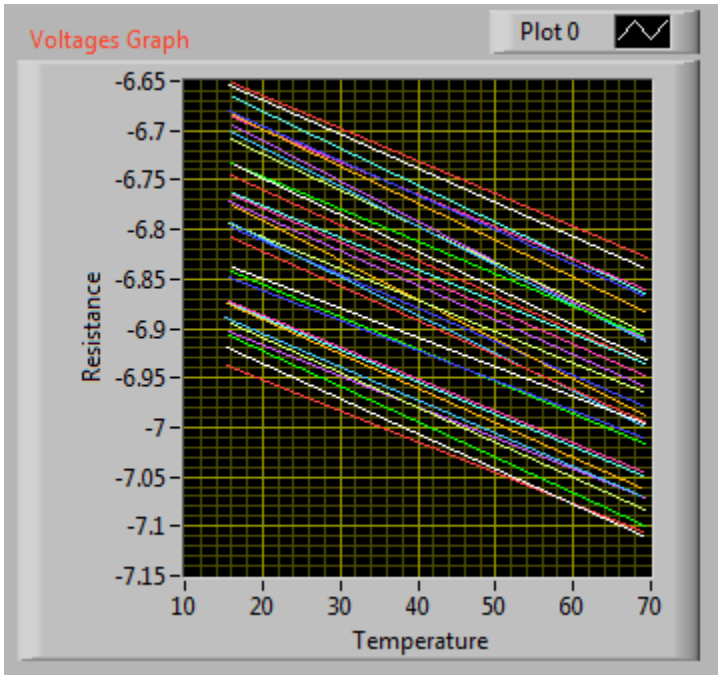

Figure 22: Calibration plot for the 32 heaters of one of the modules. Each heater has a specific calibration line. Each module has a similar looking plot.

In order to determine the variation among the diodes and the heaters, all the calibration data was combined and analyzed. Among all the trials, the average slope of the diode calibration line was determined to be -0.54°C/Ω while the average slope of the heater calibration line was determined to be -0.0038 °C/Ω. The standard deviation of both sets of slopes was determined to be 0.03 °C and 0.0034 °C respectively. Similarly, the variation in the intercept values was analyzed with the mean and the standard deviation away from that mean. This data along with the previously mentioned data can be found in

Calibration Component	Average Slope	Standard Deviation	Average Intercept	Standard Deviation	Average Variance from Best Fit (°C)
Diodes	-0.54 °C/Ω	0.03 °C	465.80 °C	24.34 °C	0.02
Heaters	-0.0038 Ω/°C	0.0034 Ω	-6.87 Ω	0.68 Ω	N/A

Table 2.

Calibration Component	Average Slope	Standard Deviation	Average Intercept	Standard Deviation	Average Variance from Best Fit (°C)
Diodes	-0.54 °C/Ω	0.03 °C	465.80 °C	24.34 °C	0.02
Heaters	-0.0038 Ω/°C	0.0034 Ω	-6.87 Ω	0.68 Ω	N/A

Table 2: Analysis of all the calibration data

3.2 Results from Modules with Differing Layers

Appendix A contains the tables summarizing the data obtained from the testing modules. Module 6 and Module 1 are both composed of bare silicon. Since these two were from the same wafer

and batch of chips they were given different identification numbers. Module 5, consisting of 2 layers with metal, was not successfully fabricated with this batch. The thermal spreading analysis however utilizes a module 5 fabricated from another batch. The plot below shows the thermal resistance of each of the modules in three trial runs. Table 3 shows the breakdown of the pertinent data obtained from these chips. The average total resistance was taken from all the data points excluding module 6's 1^{st} data points due to the uncharacteristic slope shown. Potential sources of this slope could be the cooler block cooling down the testing fixture rather than just the module. When the module is not giving off much heat, the test fixture may be a significant source of heat. However, as the module starts to give off more heat, the cooling of the thermal test stand becomes minimal in comparison. Another source of the slope could be due to the assumption that the heat transfer coefficient *h* is constant. In reality, *h* changes due to a variety of factors including the change in the delta temperature between the source and the sink. This assumption should be further investigated as more tests are conducted. As can be seen in Figure 23, most of the plotted lines appear to approach a steady state as power increases. Therefore an accurate estimate of the thermal resistance of the modules over the temperature range is provided through the averaging of all the data points in a particular line.

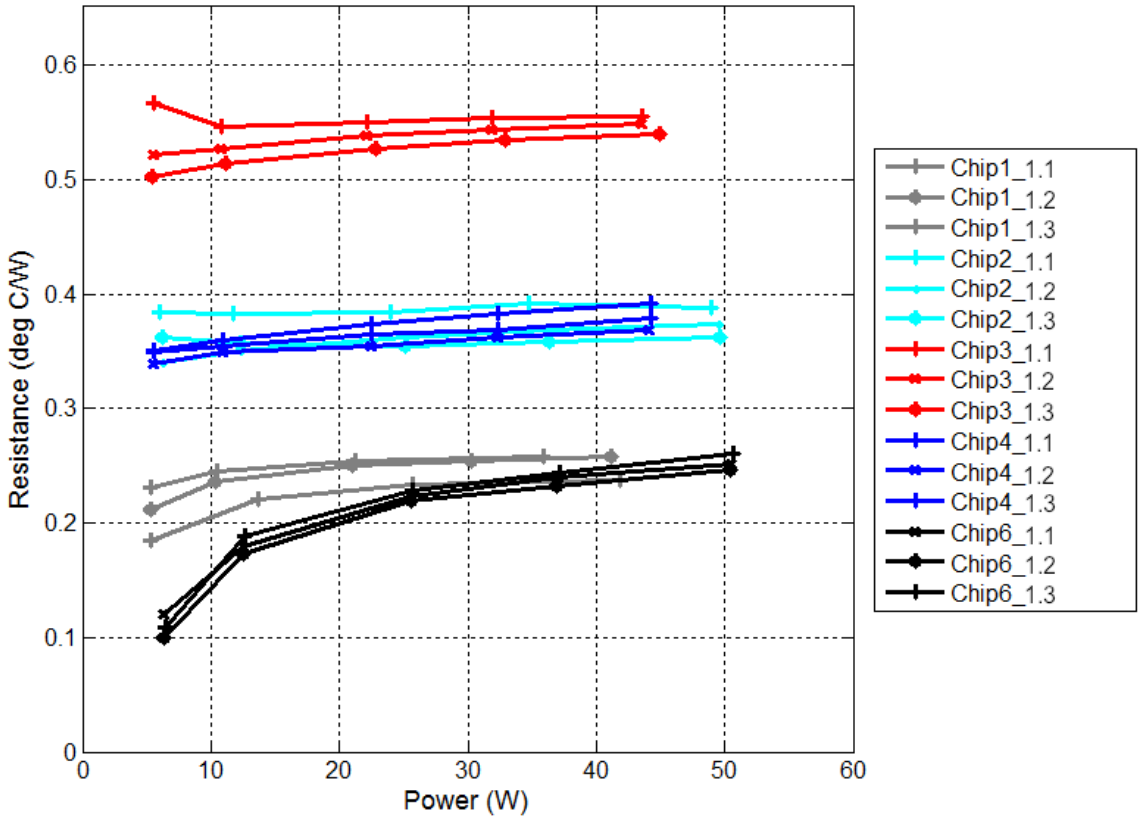

Figure 23: Plot of resistance vs. power for all the tested chips and each of their tests

Case	Average Total Resistance (°C/W)	Standard deviation of measurement (°C/W)	Resistance delta vs. Silicon-only case (°C/W)	Calculated effective conductivity (W/mK)	Resistance increase per layer of interconnect (°C/W)
(3) Four layers, no metal	0.539	0.017	0.308	0.22	0.077
(2) Two layers, no metal	0.368	0.015	0.138	0.24	0.069
(4) Four layers, with metal	0.363	0.015	0.133	0.51	0.033
(1&6) Silicon	0.231	0.025	N/A	N/A	N/A

Table 3: Compilation of thermal test data for each of the four test cases

An important observation of the graphed data is that the thermal resistance of the chips is ranked as follows: Silicon < four layers with metal < two layer no metal < four layers no metal. There are two expected consequences of this. First the silicon only module had the lowest resistance as it had the least material impeding heat flow. As is shown, the addition of layers affects the thermal resistance. Second, the incorporation of metal helps heat flow through the module. From the data shown in Table 3, the metal improves the thermal conductivity by a factor of 2. It appears that the addition of layers of SOD and the contribution of metal can have a 0.133-0.308 °C/W effect on the thermal resistance of the silicon-based module. The determination of this contribution was one of the main goals of this work. It is important to note the consistency between trials of each module. All the tests for a module follow the same trend, helping affirm the validity of the data.

3.3 Results of Heater Isolation

A module was diced with a 70 µm blade in order to isolate heater regions so that a heat-spread study could be conducted. Three dicing cuts were made so that 6 regions of varying sizes were formed. The schematic for the cuts are shown in Figure 24 below. This image is shown as if the metal, silicon and SOD layers covering the TEA module were see-through. The cuts went completely through the module but did not cut through the interposer board so that no electrical connections were disturbed.

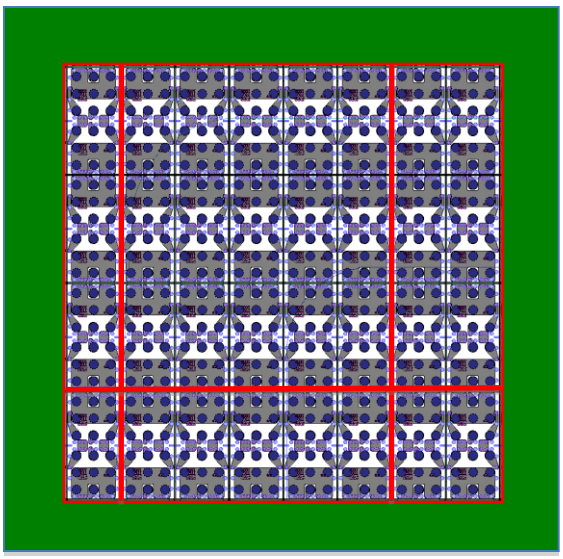

Figure 24: The dicing cut plan for the module in order to isolate regions of heaters. This picture shows the TEA module that is on the backside of the face of the module used during experimentation.

The trench cuts on the module surface are shown in Figure 25. The chip chosen for this study consists of 2 layers of SOD with metal.

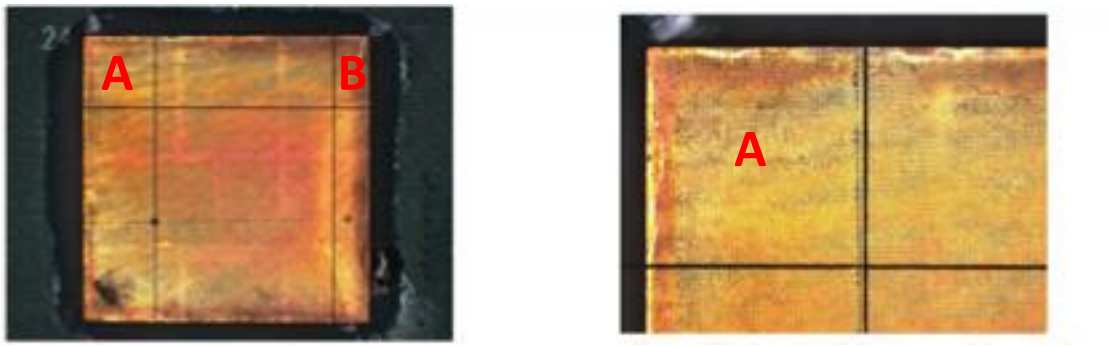

Figure 25: Thermal module with trench cuts for heater isolation

Through the isolation of heater regions, temperature gradients can be observed over the 70 μm gap between the parts of the wafer. This information will help determine the effect heat spread has on the overall thermal resistance of the module.

Care must be taken during the application and removal of the TIM so as not to fill the trenches on the module. This would result in the bridging of the heat path across the surface of the chip. Therefore, after every trial, the module was subjected to an ultrasonic bath in order to remove all remnants of the TIM from the trenches.

The main analysis conducted was on the heat spread of an isolated region with two heater units and a region with one heater unit. These regions are labeled A and B respectively in Figure 25.

The thermal spreading analysis was done by comparing module 5 with trench cuts, as previously stated, and module 2 without trench cuts. The decision was made to use module 2 as a comparison due to only one successful fabrication of module 5 and due to the fact that they both are composed of silicon and the same number of layers of SOD. The testing was done on two sections of the modules. One section isolated one heater while the other isolated two heaters. The results are shown below.

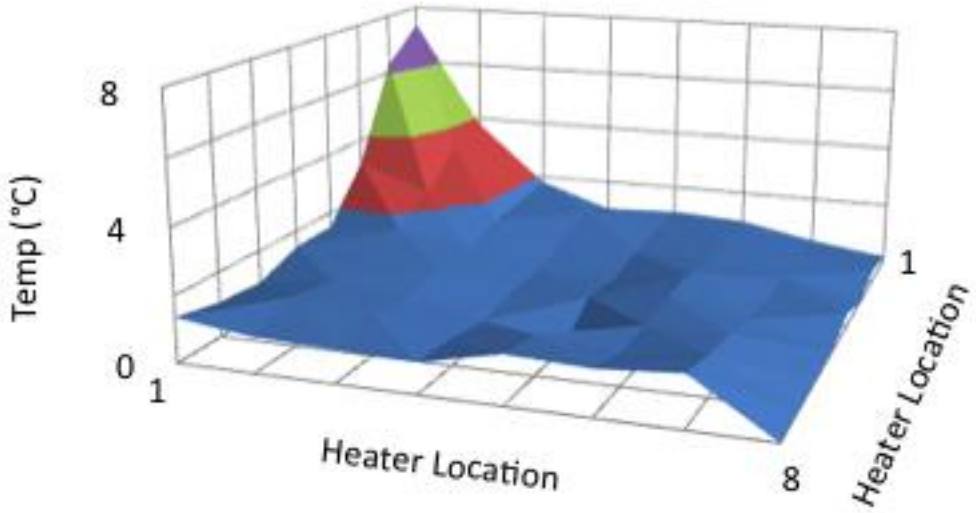

Figure 26: Module 2 with one heater on with a power level of 1.17 W

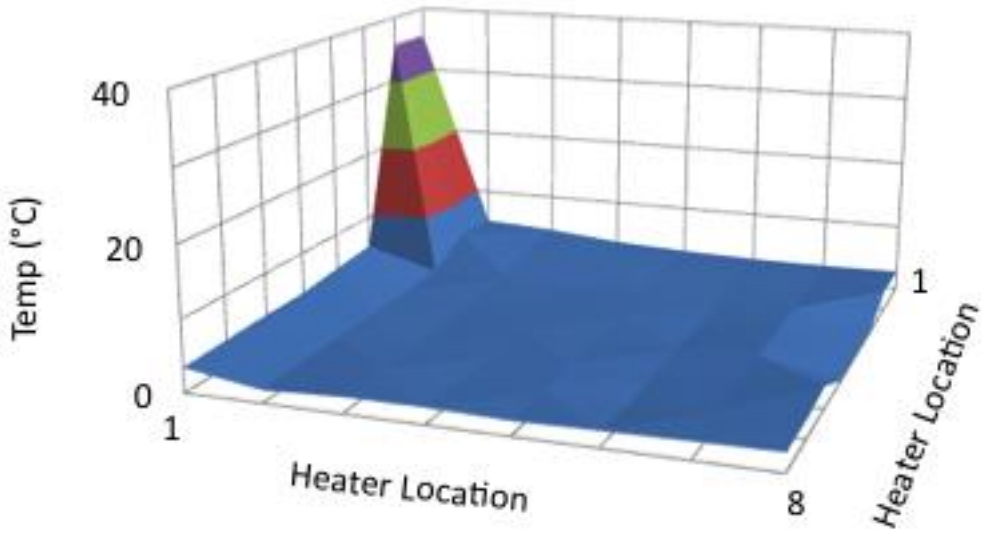

Figure 27: Module 5 with one isolated heater on with a power level of 1.16 W

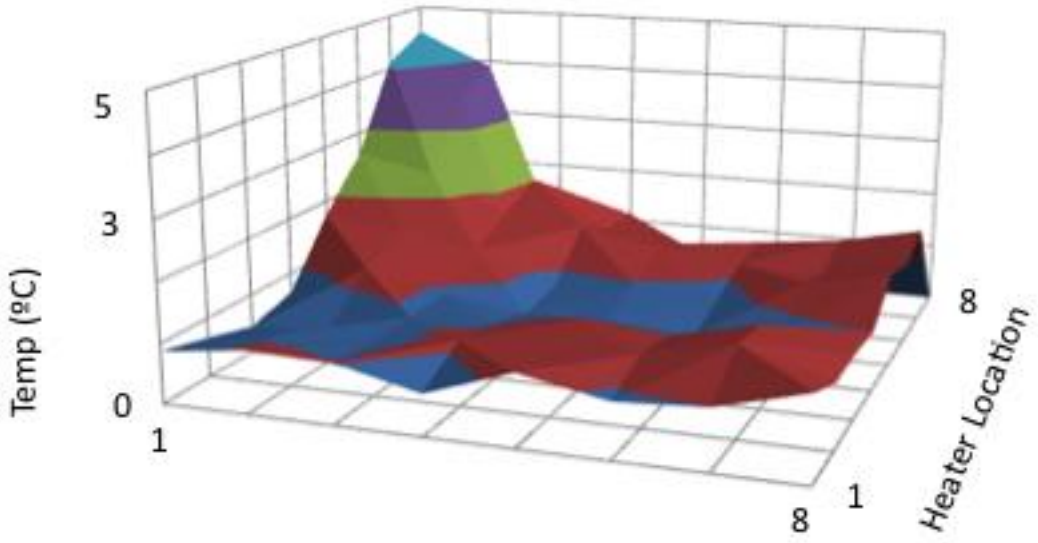

Figure 28: Module 2 with two adjacent heaters on at a power level of 0.97W

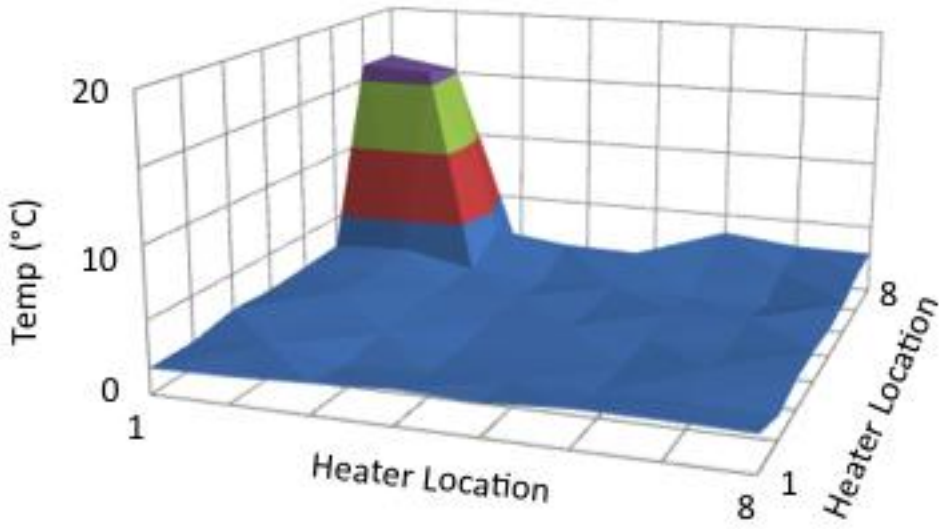

Figure 29: Module 5 with two isolated adjacent heaters on at a power level of 0.81 W

As expected, the trench cuts effectively eliminated the heat spread throughout the chip thus greatly reducing the effectiveness of the heat sink. The isolated 1-heater region had temperatures approximately five times higher as compared with the module without the trench cuts. The isolated 2-heater region had temperatures approximately four times higher as compared with its non-isolated counterpart. It is important to note that the power levels for the 1-heater regions were almost identical with a 0.1 W difference and the power level for the 2 heater regions was 0.17 W different. Since the testing system is current dependent with the required voltage drawn, it is challenging to get the exact same power level between tests.

3.4 Results of Module Stack

The thermal analysis of the stack was done by comparing the total thermal resistance of the stack to that of a single bare silicon chip. As discussed in Appendix A.3, the stack was fabricated with two bare silicon chip pieces stacked and adhered together by solder balls. The top chip was not electrically connected and served the purpose of simulating the heat path of a stacked module. The attachment

process mirrors the methods used in actual application very closely so this setup can be an accurate thermal analog to a functional chip stack. Both the single wafer silicon module and the stacked silicon module were tested three times to establish repeatability. The results for the thermal resistance vs. power applied are shown in Figure 30. Potential reasons for the observed slope are discussed in section 4.2. A plot showing the temperature rise of the module vs. power applied is shown in Figure 31.

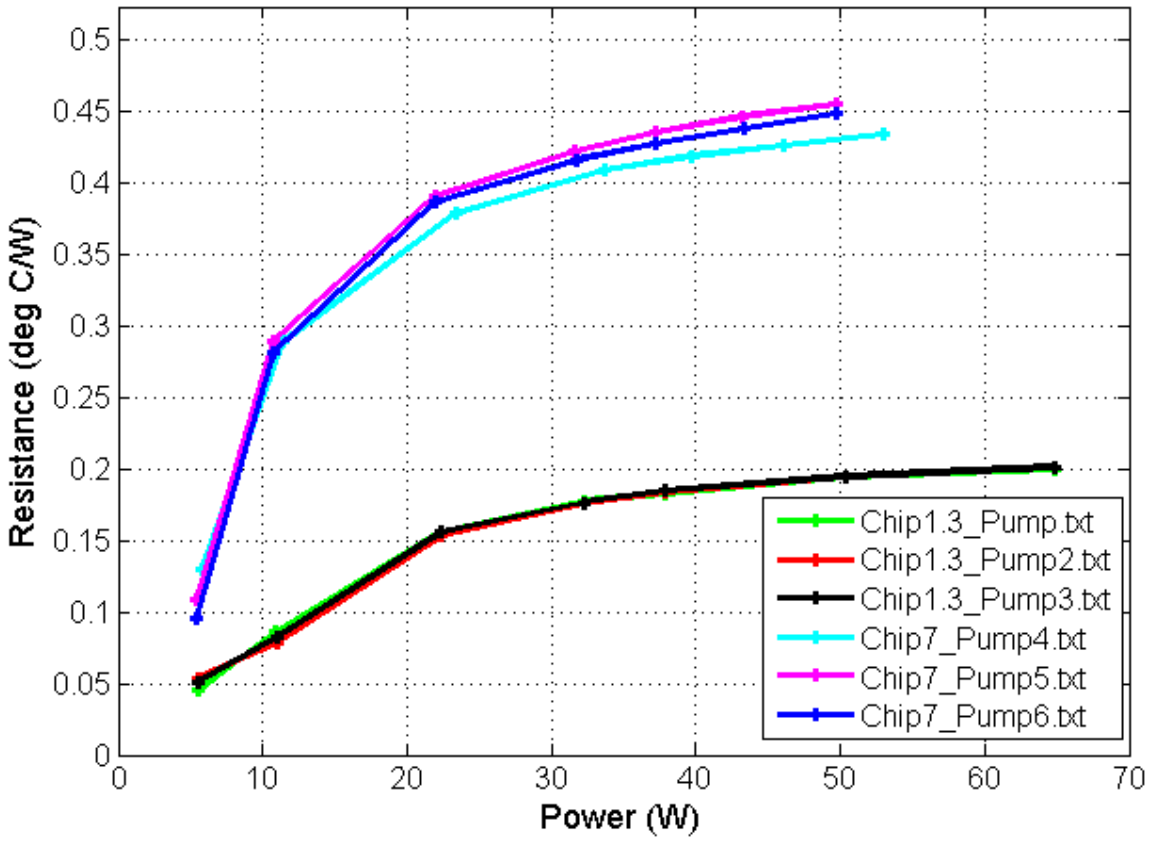

Figure 30: The thermal resistance vs. power of the stacked module (7) and the single wafer silicon module (1).

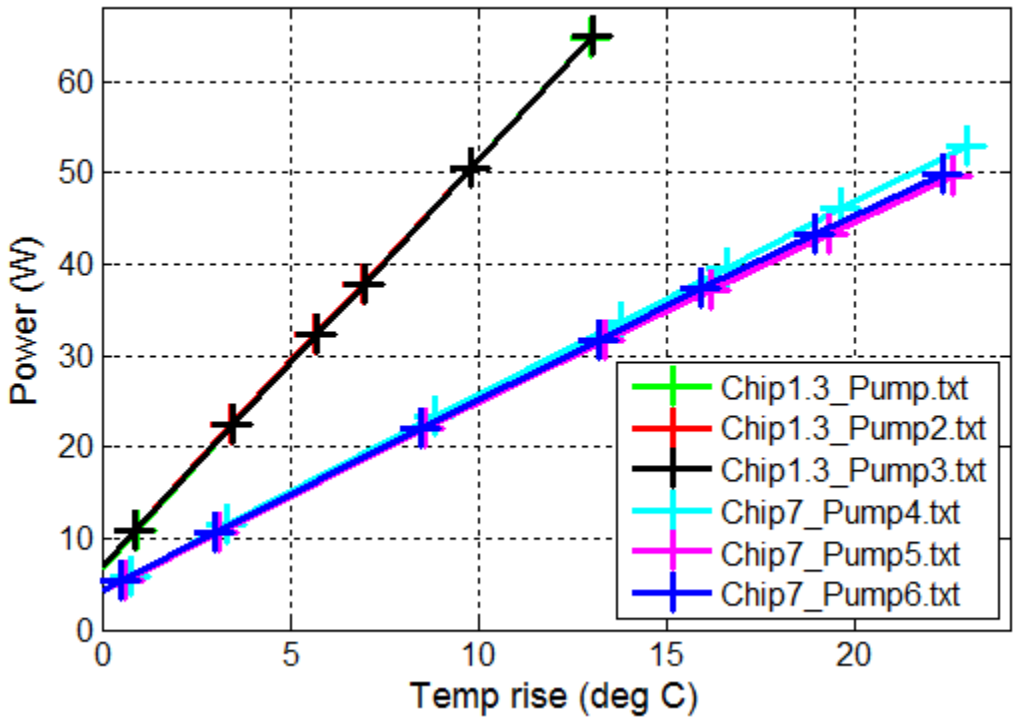

Figure 31: Power vs. temperature rise of the single wafer silicon module (1) and the stacked silicon module (7)

Table 4 shows the breakdown of the pertinent data for the two cases. As can be observed, the addition of the second wafer to the thermal path affects the thermal resistance by a magnitude greater than 2. This is likely because the addition of the second wafer layer not only adds thermal resistance by doubling the amount of silicon in the thermal path, but also adds thermal resistance by forcing the heat path through the solder balls which connect the two chips. The SnPb solder balls are high conductivity balls, but have a thickness of roughly 375µm after reflow and cover approximately 27% of the surface area of the chip face. It is this thickness and low surface area that likely contributes to the thermal resistance of the stack.

Case	Average Total Resistance (°C/W)	Standard deviation of measurement (°C/W)	Resistance delta vs. Silicon-only case (°C/W)
2 silicon wafer stack	0.422	0.023	0.24
Silicon	0.182	0.016	N/A

Table 4: The thermal data obtained from the 2 silicon wafer stack and its single wafer silicon comparison.

Chapter 4: Conclusions

4.1 Contributions to Field

One of the most significant contributions this research provides is isolating the effect of SOD layers in controlled environments. This provides data that can be used for more complicated thermal models which would be more challenging and more expensive to test. This data also directly addresses the concerns of the increased thermal load associated with miniaturization of chips as it quantifies the thermal effect of a high density material. This information can be used to assist in the design of chips to reduce this thermal load.

The results show the impact of adding SOD layers, metal traces, and stacking chips to the thermal resistance. All of these operations are currently being used without quantification of their impact. Also the effect of reducing heat spreading was examined through heater isolation. This simulated the challenges in sinking heat from a source as the surface area is decreased. Although there are sources of error in the measurements, there is confidence in the measurements as discussed in Appendix section A.4.

4.2 Future work

First the assumption of the thermal coefficient being constant throughout the test must be investigated. As can be seen in the results, the thermal resistance appears to initially have a steep slope over low powers. This may be because the delta temperature between the source and the sink is low at these points changing the thermal coefficient. In a similar vein, analysis should be conducted to see if the cooler is taking away extra heat from the module. There is a big disparity between the power supplied by the power source and the power supplied to the heaters, meaning a heat source could be

present in another part of the module. Investigating this by analyzing the resistance of the thermal test stand components and the mass flow of the cooler would increase the accuracy and sensitivity of the results.

An important next step in the determination of the thermal resistance of the high-density processor packaging method is determining variability on a larger scale. This involves testing modules from different wafers to insure the fabrication process doesn't have any inherent variability. Preliminary tests with the second batch of modules indicate very low variability. However, to insure this, more than two wafers need to be fabricated. Due to time and cost constraints, this hasn't been previously investigated.

More modules will also allow the variability of the thermal spreading analysis to be quantified. Due to lack of materials and time, more diced wafers could not be produced. Also a direct comparison between the same types of module (trench cut vs. not) would increase the accuracy of the results on the effect of the trenches on heat spreading. This will involve dicing more trenches into the different types of modules to see if have a varying effect. Another way the thermal spreading analysis can be improved is by implementing a method to get identical power levels into the modules so that their thermal response can be compared more accurately.

As more modules are being produced, the tests for the stacking of the wafers can be furthered. This involves doing more tests on different combinations of stacked wafers. One stack planned is a 2-wafer stack of silicon wafers with SOD layers. This will then be able to be directly compared to the bare silicon 2-wafer stack and the thermal effect to the SOD on a larger scale can be analyzed. Another area to be investigated is increasing the stack height. There has been interest at Draper to make a stack as tall as three wafer pieces. The method for fabricating this stack would be identical to the fabrication of the 2-wafer stack, a method that has proved successful.

Lastly, one important implication of this research is how it affects current Draper projects. As such, the module can be redesigned to directly mimic Draper modules in order to determine their thermal characteristics and establish methods for thermal management. This is a primary motivation for the research, and this future work would increase the fidelity and applicability of the results.

References

[1] Smith, Brian. “Charmin, iUHD Thermal Test Capability for TP.” Design Review. 4 October 2012. Draper Laboratory internal document

[2] Kandlikar, Satish G. and Hayner II, Cliford N. (2009) “Liquid Cooled Cold Plates for Industrial High-Power Electronic Devices—Thermal Design and Manufacturing Considerations,” Heat Transfer Engineering, 30: 12, 918-930.

[3] Lingbo Zhu; Hess, D.W.; Wong, C.P., "Assembling Carbon Nanotube Films as Thermal Interface Materials," *Electronic Components and Technology Conference, 2007. ECTC '07. Proceedings. 57th* , vol., no., pp.2006,2010, May 29 2007-June 1 2007

[4] Avenas, Y.; Ivanova, M.; Popova, N.; Schaeffer, C.; Schanen, J.-L.; Bricard, A., "Thermal analysis of thermal spreaders used in power electronics cooling," *Industry Applications Conference, 2002. 37th IAS Annual Meeting. Conference Record of the* , vol.1, no., pp.216,221 vol.1, 13-18 Oct. 2002

[5] Brunschwiler, T.; Paredes, S.; Drechsler, U.; Michel, B.; Wunderle, B.; Reichl, H., "Angle-of-attack investigation of pin-fin arrays in nonuniform heat-removal cavities for interlayer cooled chip stacks," *Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2011 27th Annual IEEE* , vol., no., pp.116,124, 20-24 March 2011

[6] Brunschwiler, T.; Paredes, S.; Drechsler, U.; Michel, B.; Cesar, W.; Toral, G.; Temiz, Y.; Leblebici, Y., "Validation of the porous-medium approach to model interlayer-cooled 3D-chip stacks," *3D System Integration, 2009. 3DIC 2009. IEEE International Conference on* , vol., no., pp.1,10, 28-30 Sept. 2009

[7] Hornberger, J.M.; Cilio, E.; Schupbach, R.M.; Lostetter, A.B.; Mantooth, H.A., "A High-Temperature Multichip Power Module (MCPM) Inverter utilizing Silicon Carbide (SiC) and Silicon on Insulator

(SOI) Electronics," *Power Electronics Specialists Conference, 2006. PESC '06. 37th IEEE* , vol., no., pp.1,7, 18-22 June 2006

[8] Licari, J.J., "Fabrication and packaging of high-density multichip interconnect (HDMI) substrates," *Aerospace and Electronics Conference, 1989. NAECON 1989., Proceedings of the IEEE 1989 National* , vol., no., pp.1682,1688 vol.4, 22-26 May 1989

[9] Dufour, B.; McNulty, M.; Miller, S., "Microwave multi-chip module utilizing aluminum silicon carbide with in-situ cast components and high density interconnect technology,"*Multichip Modules, 1997., International Conference on* , vol., no., pp.309,314, 2-4 Apr 1997

[10] Smith Brian. “Thermal Stand Test Philosophy.” Internal document. 19 June 2013.

[11] Burns, John IV. “FT-Live Dicing Request.” Internal document.

Appendix

A.1 Individual Diode and Heater Calibration

In order to obtain accurate information from the module, both the diodes and the heaters were calibrated. This was done through the implementation of a Labview® program with a process for diode calibration and another for heater calibration. The diode calibration is run first as the heater calibration is dependent upon its resulting data. The diodes were individually calibrated in order to minimize error resulting from the variation of resistance and thermal sensitivity of the diodes. Therefore the diode calibration resulted in 64 equations for the conversion of resistance to temperature. As stated previously, this calibration was based on the assumption that the ambient temperature is that of the liquid cooled heat sink. Two thermocouples, one at the inlet and one at the outlet of the cooling block, were averaged in order to determine this temperature. The thermocouples used were Cole-Palmer 2.5" Type K thermocouple probes. Since the probe tips were designed for penetration, one probe went into the tubing directly before entering and one right directly after exiting the liquid cooled heat sink. The probes were placed to be as close to the liquid cooled heat sink as possible. The temperature of the heat sink is controlled by a Cole-Palmer Polystat bath and circulator. Since the resistances of the diodes are designed to vary linearly with temperature, the calibration equation is of the following form:

$$T = m \times r + b \tag{7}$$

Where T is the temperature of the diode, m is the slope constant determined from calibration, r is the resistance read at the diode, and b is the vertical offset at the Y-axis determined from calibration.

Calibration data was also analyzed for the r^2 value of the linear fit, which proved the efficacy of the linear model.

The resistance of the diodes was taken three times at five temperatures: 16°C, 30°C, 40°C, 55°C and 70°C. This resulted in plot of 15 points for every diode. Each set of data was then used to determine the diode's best-fit line.

The calibration of the heaters was done to account for the thermal variation of the heating element. Since these resistors were not designed to be thermally sensitive, the test was done at only two temperatures: 16°C and 70°C. Therefore the heater resistance change observed in respect to temperature changed will be much less significant than with the diodes. However, this calibration is important in order to determine the amount of power going to each heater so that the actual power being dissipated by each region and in the entire module can be analyzed.

The heater calibration process is done through the application of a specific current and drawn voltage to each individual heater at each of the temperature levels. Since the resistance of the heater resistors cannot be directly determined, the following analysis was done.

$$R_H = \frac{V_H}{I_H} \tag{8}$$

Where R_H is the resistance, V_H is the voltage drawn and I_H is the current applied to the heating element. This allows the use of a linear fit to model the equation thus resulting in the following.

$$R_H(T_H) = m_H \times T_H + b_H \tag{9}$$

Where T_H is the temperature of the heating element, m_H is the slope of the fit and b_H is the y-intercept of the best fit line.

Since the switch to turn on the heaters is physical, the process is a little tedious as it involves the flipping of 64 switches. A waiting period allows the module to reach steady state before data is taken. Therefore each heater will have two data points, one at each temperature. The linear fit of the heater resistance change with temperature shows a very small slope, however this data is still used in the Labview program designed for power tests of the modules since it helps to reduce error. The temperature of the heater resistors were determined by averaging the temperature determined from the two diodes each heater directly impacted.

Both the diode and heater calibration were run multiple times in order to ensure repeatability of the results.

A.2 Labview Data-Interpretation Program Design

Labview controlled the thermal testing stand in order to automate as much of the process as possible. Depending on whether the module was being calibrated or tested, different Labview programs were developed and used.

The diode calibration program was designed to determine diode temperature as outlined in section A.1. The user interface for the diode resistance of the calibration VI is shown in Figure 32. As can be seen, the resistance of the diodes are read out in raw format and shown on the screen under "Readings." This array gives the meter, resistance and the location in the meter. The "thermocouple average" data field shows what the estimated surface temperature of the chip is for calculation. Directly below, the "data points taken" box counts the total number of data points recorded so it can be determined where in the process the program is in data collection. As the data points are taken, the "Diode slope, int, error" fields are populated with the slope, intercept and error of the best fit line though the data. This data is then graphically represented in the XY plot so that the any outlier behavior among the diodes can be identified. The "diode resistance and voltage values" array outputs all the raw data being read by the multimeters in a row with each subsequent reading populating the next row. The slope and intercept data along with the raw data are then saved to separate .txt files.

The program also uses serial communication with a cooler to set the surface temperature of the chip. Since the program is automated, steady state is ensured through the implementation of a three-minute wait time.

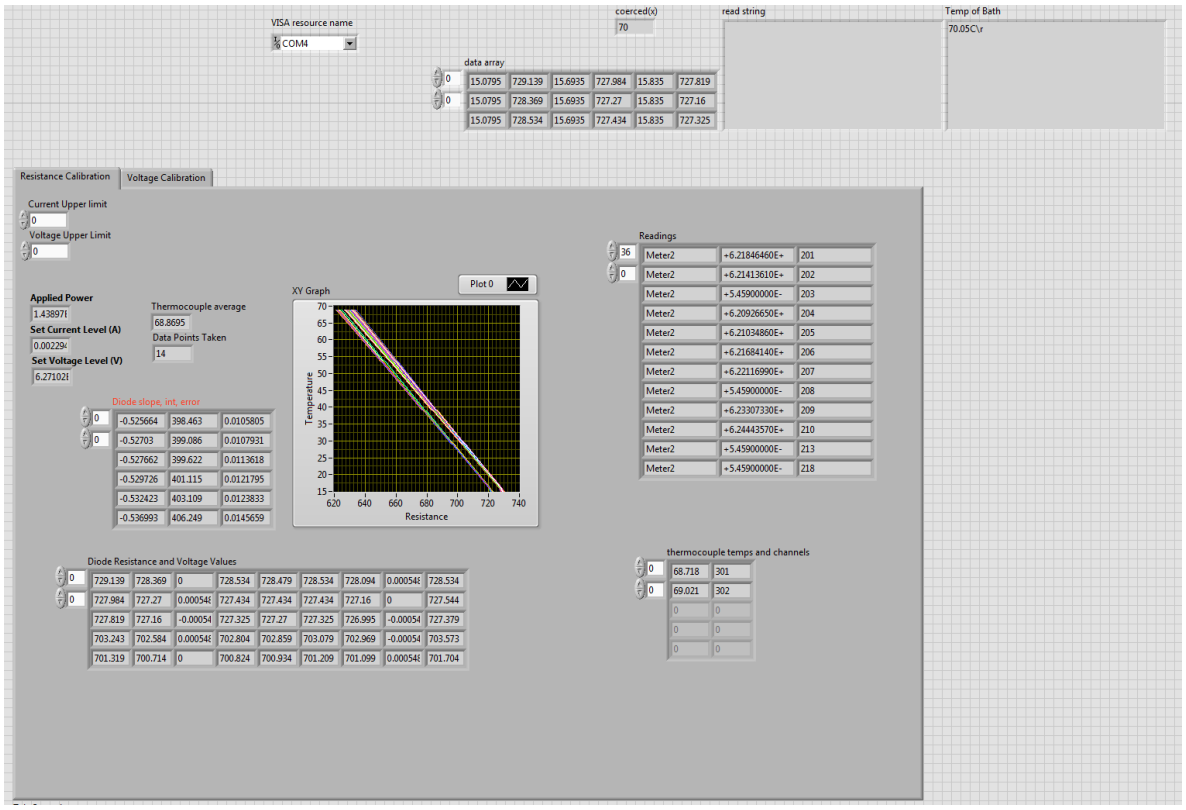

Figure 32: Diode Calibration Labview® Program Operation Panel

Since the calibration of the heater resistors cannot be taken directly as described in section A.1, the voltage and current must be analyzed. Figure 33 shows the user interface for the VI. The voltage array is displayed in 1d format along with a 2d format which represents the face of the module**.** When a heater is turned on and a data point is taken, the" I, V, T array" is populated. The voltage is taken at chip level with the 4-point probe technique mentioned earlier. This provides the actual voltage the heater is experiencing. The current was driven by the power supply which was calibrated by Draper Laboratory. The thermocouple average is also displayed so the heat sink temperature can be observed. Once all the data points are taken, the "heater cal slope" array is populated with the slope and the y intercept of each of the best-fit lines for the heaters. Once all this data is obtained, a graphical representation of all the heater slopes is plotted as can be seen on the right side of Figure 33. This enables rapid comparison

of resistance versus temperature of the heaters. The slope and intercept data along with the raw data is saved in separate .txt files.

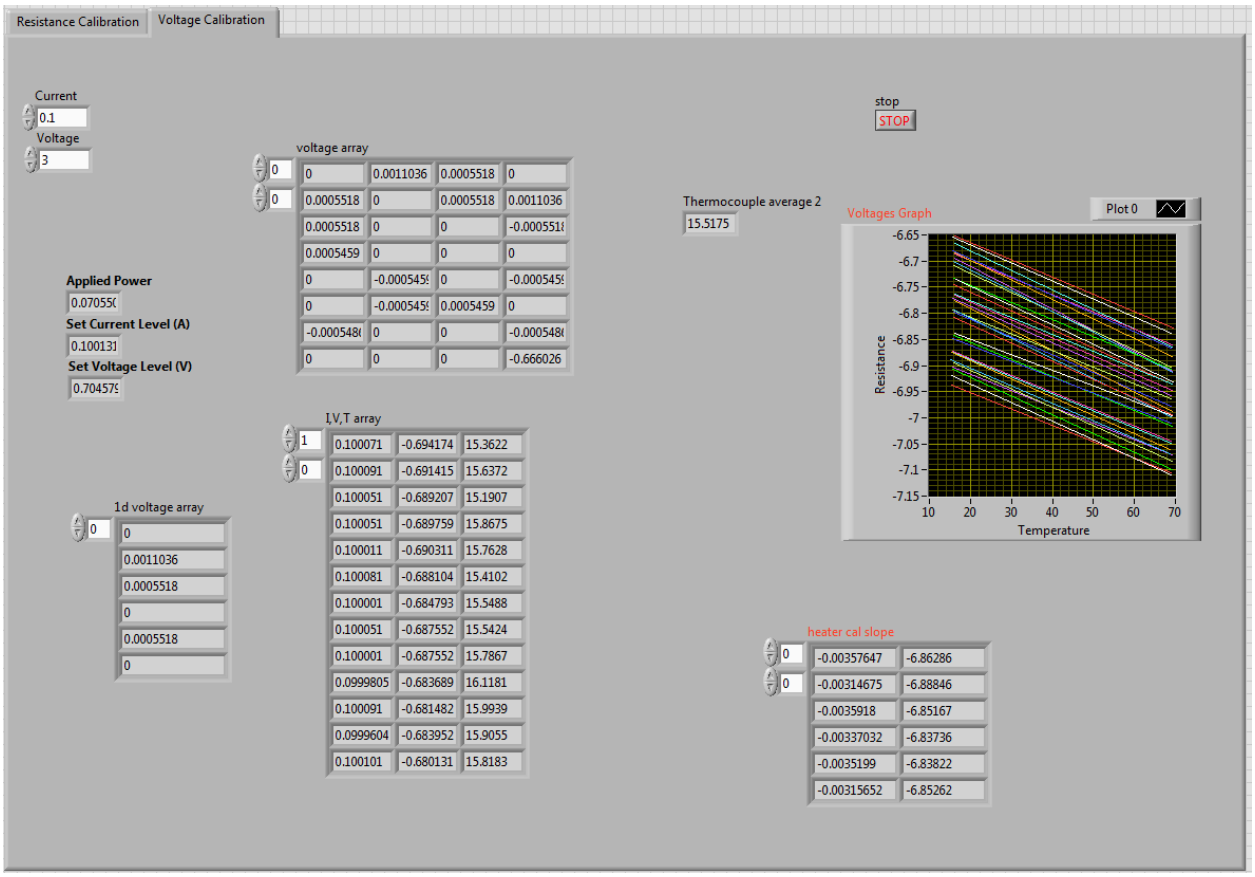

Figure 33: Heater Calibration Labview® Program Operational Panel

The thermal resistance analysis program operates using the calibration data obtained from the heaters and the diodes. Instead of varying the temperature of the cooler, the power supplied to the module is varied resulting in the heating of the diodes. Each of the five different module types responds differently to the application of similar power levels. Figure 34 depicts the user interface for the VI. The resistance intensity plot and the temperature intensity plot are shown to facilitate quick reading of the surface temperatures**.** These intensity plots are graphical representations of the numerical arrays, which are also shown to the operator**.** Both these plots have a changeable range for graphing data points so

dead diodes, diodes that no longer respond linearly or predictably to temperature change, do not affect the representation of the data. Since there is loss in the thermal test stand board, the heater calibration data and the observed heater voltages are used to determine the actual power delivered to the module. The actual power dissipation in respect to heater location is shown in a numerical array along with its corresponding intensity graph. The average diode temp is also calculated so that outliers can be programmatically determined during the automatic operation of the program. Thermocouple readings are shown so that any inconsistencies in operation can be identified**.** Dead diodes are removed from the data before thermal resistance is calculated. The program has the option to alter the cooling method of the module so that it is compatible with fin heat sinks with and without fans. The arrays for “1D T*”, “1D Temperature heater orientation zeroed values” and the analysis of standard deviation are all used to eliminate diodes that are malfunctioning. The piece of data the program is designed to evaluate is the thermal resistance, which is shown at the top of the figure. This value is constantly being updated throughout the operation of the program, reflecting the real time thermal resistance of the module. In the calculation of the thermal resistance, the thickness of the TIM previously discussed is used. A safety shutoff is built in to the program to ensure the module doesn’t overheat and damage the thermal testing stand.

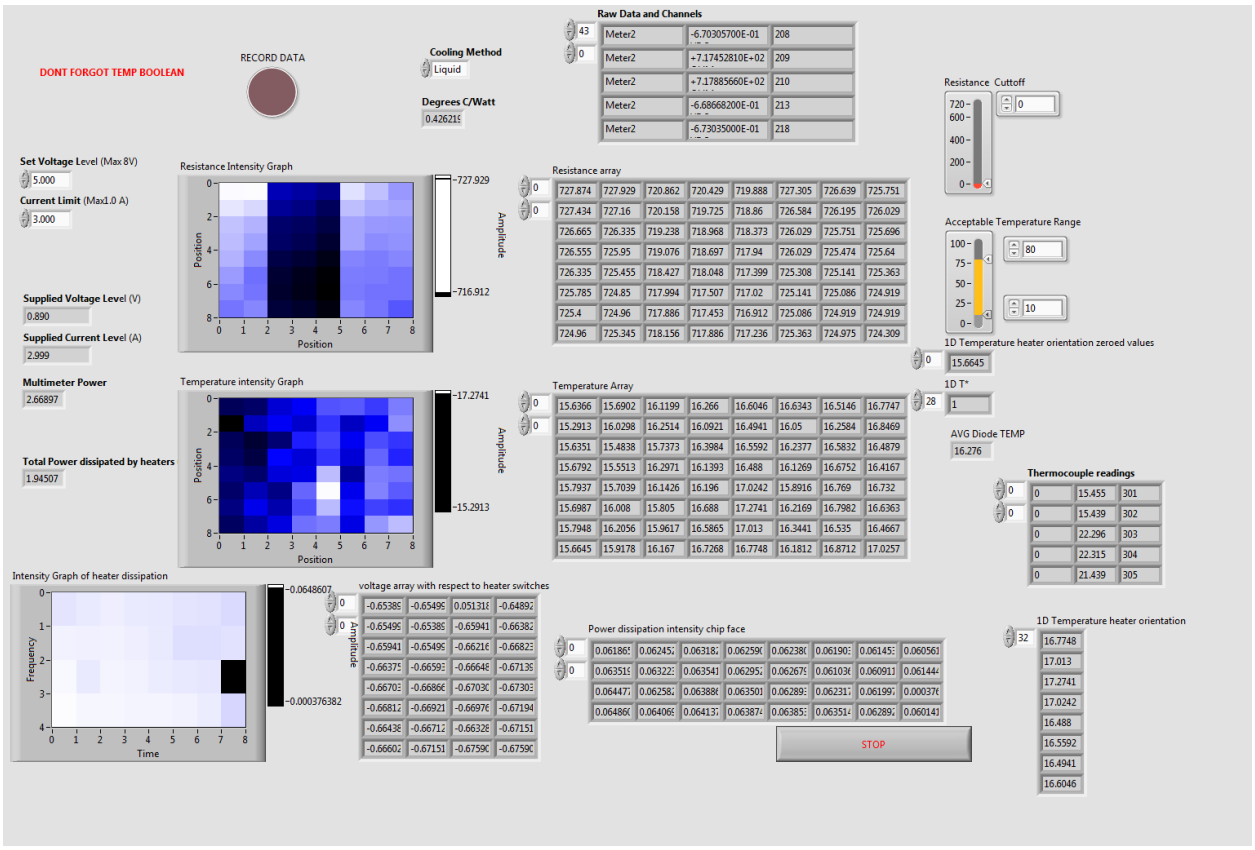

Figure 34: Operation Panel of Labview Data Interpretation Program

A.3 Fabrication of Stack

Due to the usage of high-density chip modules in a stacked fashion, a fabrication method to simulate such a stack was designed in order to determine the effect on the thermal properties of the module. This stack consists of two wafer pieces bonded together through the same means the bottom wafer is bonded to the interposer board. This way the thermal path between the two wafers mimics the path between the bottom wafer and the interposer. The top wafer is a purely mechanical component and therefore neither diode nor heater data is obtained from it. As shown in Figure 35, the presence of the wafer stack only slightly alters the thermal path described in the mathematical theory section.

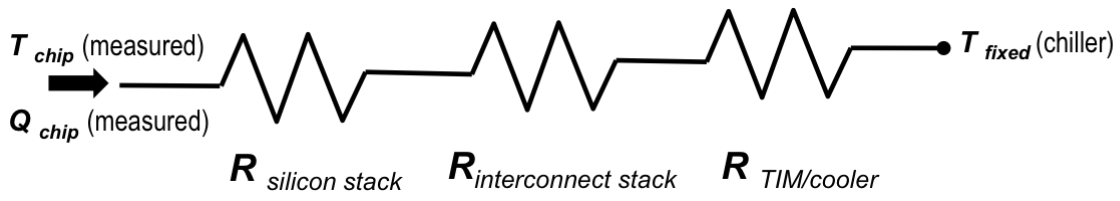

Figure 35: Thermal path for the stack of modules

A bare silicon wafer, populated with the TEA chips that had undergone similar processes to the other thermal wafer, was diced into 8x8 arrays of diodes. The diced wafer pieces were then bumped with 16mil SnPb solder balls with a reflow temperature of 183 °C. After this, some of the wafer pieces were set aside so that the adhesion layer for bonding between the wafer pieces could be created.

As discussed previously, the bonding layer on the surface of the bottom wafer pieces was designed to simulate the heat path of the bonding layer between these wafer pieces and the interposer board. Therefore, an array of bonding pads for the solder balls to adhere to was designed. First, a shadow mask made out of 170°C release REVALPHA® Tape had an 8x8 array of square holes laser cut into it. Then the shadow mask was carefully stuck to the surface of the module to create zones for the pads to be selectively deposited. The sides of the wafer were then covered in Kapton® Tape to protect it

from getting metalized and shorting connections. Subsequently, the module pieces were sputtered with a 300 Å TI/50,000 Å Au mixture. 5 µm of copper was used in order to ensure that the solder balls did not consume the pad thus ruining the adhesion layer. Post metallization, the shadow mask was removed through the application of heat to the REVALPHA® leaving an 8x8 array of pads. In order to determine the true amount of material sputtered, the pad heights were then measured using a profilometer. Two pads were measured for each wafer piece, and the average value is shown in Table 5. The profile of the surface of one pad and the space next to it is shown in

Figure 36. This bumped wafer piece with pads was then aligned with the interposer board, and another bumped wafer piece without pads was aligned on top of the sputtered pads. The solder balls were then reflowed forming a stacked module. To provide for mechanical stability, the solder ball regions were under-filled. The overview of the process is shown in Figure 37.

Wafer Piece #	Average Pad Height (µm)
1	3.93
2	4.36
3	4.37
4	4.57

Table 5: Average pad height on wafer surface after sputtering operation

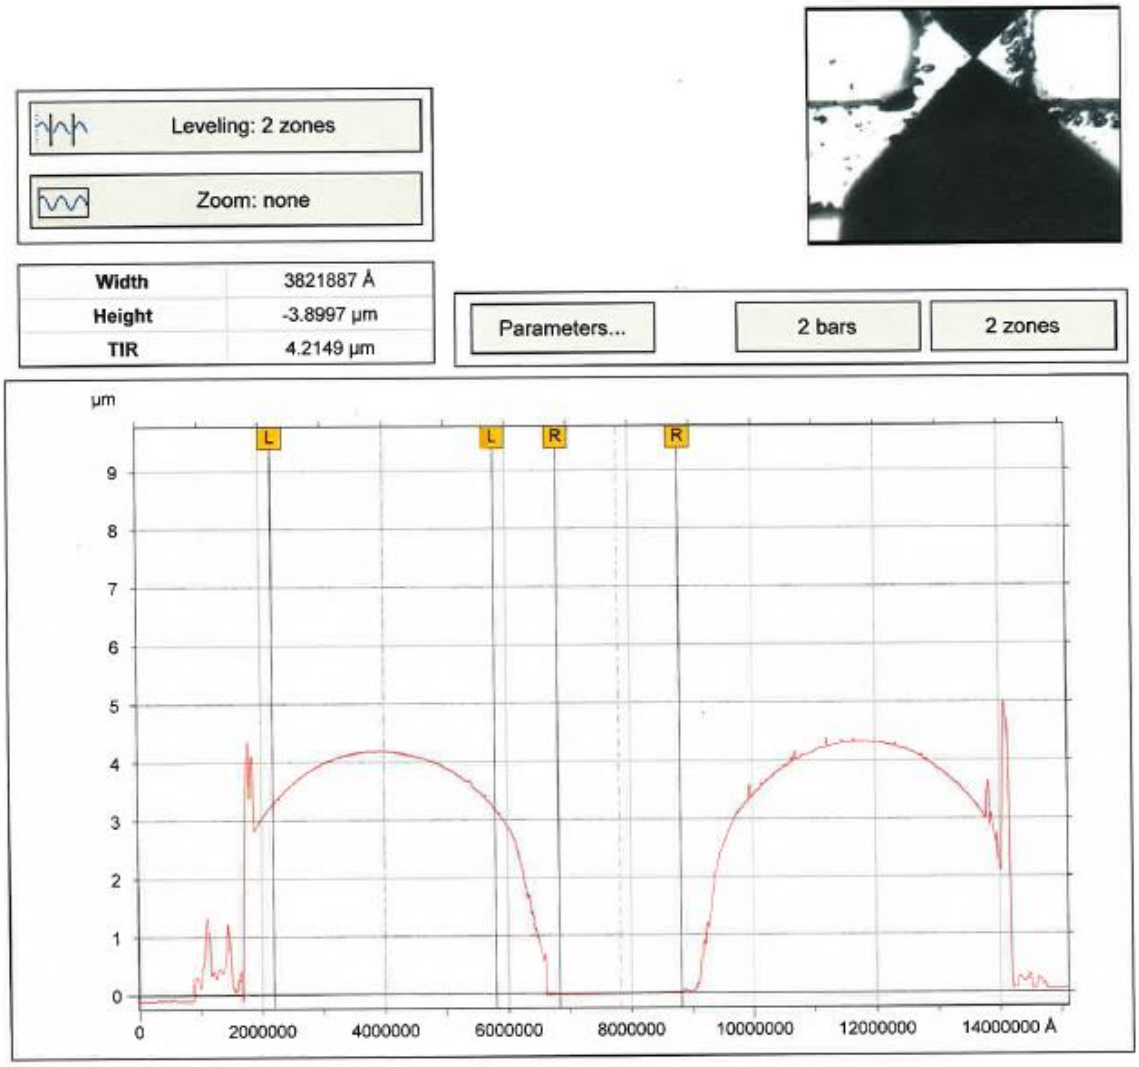

Figure 36: A Profilometer Data sheet

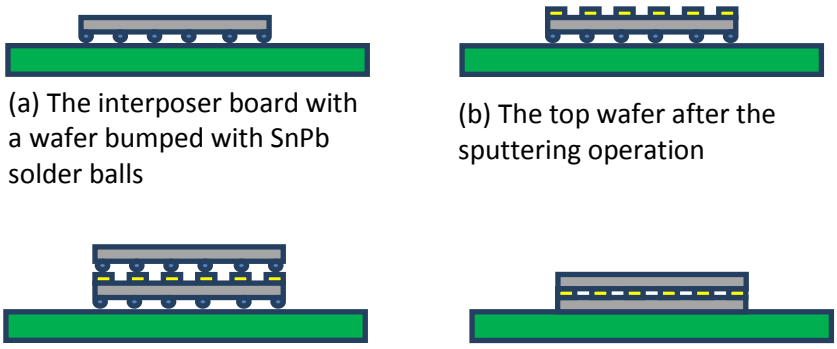

(a) The interposer board with a wafer bumped with SnPb solder balls

(b) The top wafer after the sputtering operation

(c) The upper bumped wafer aligned on the lower metalized wafer. Both are populated with SnPb solder balls

(d) The stack after reflow

Figure 37: Process overview of fabrication of stack

The final results from the fabrication are shown in Figure 38 and Figure 39.

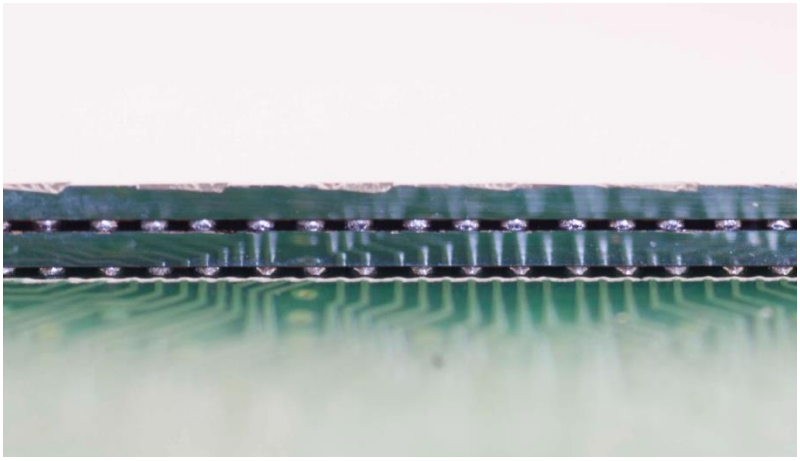

Figure 38: Sideview of stacked wafer after fabrication

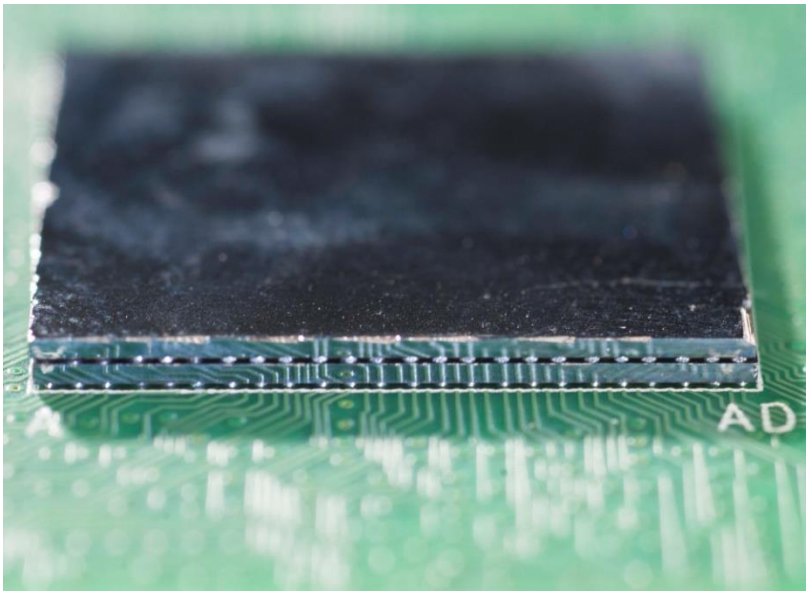

Figure 39: Stacked wafer after fabrication

A.4 Error Analysis

As mentioned throughout the thesis, sources of error are present throughout the experiments. Figure 35 shows the thermal resistance network that is being analyzed. Each component of the thermal resistance network (the silicon chip, the TIM, the interconnect, the cooler) have a respective contribution to the total thermal resistance. Some of these components have error or variability in their values. This section quantifies and addresses those sources of error.

The error resulting from the 30% variability in TIM thickness can be quantitatively addressed. The resistance of the TIM is approximately 14% of the total resistance of module 3 while for modules 2 and 4 it is approximately 20%. Therefore a 30% variation in thickness would only have a 4-6% error on the resistance measurement. Since these tests are conducted dozens of times, it is assumed that the TIM thickness averages to the mean and thus this analysis provides a conservative guess for the resulting error.

The error associated with the diodes and heaters also is small. The diodes and heaters themselves are very stable at a given temperature. From observation, there seemed to be very little, if any, drift in the values over tens of minutes. Since each of the diodes and heaters were individually calibrated, different best fit lines were able to be used for each of the 64 diodes and 32 heaters reducing the error. The average of all the mean square errors for the calibration of the diodes, depending on the module, ranges from 0.003-0.034 °C. Since each of the diodes is shown to have a sensitivity of approximately 0.5 °C/Ω, the error resulting from calibration is not significant. It is important to note that the resistive value is simply a representative value based on the current applied and voltage drop observed through the diode. Also since the calibration is run several times with essentially identical results, the confidence in their accuracy is increased. Therefore there is very little error in the diode measurements. Since the heaters were calibrated at only two temperatures, there is no data on the

mean square error for the calibration line. However, the heaters were designed to have consistent electrical properties over large heat ranges. This can be observed from the slope of the heater calibration data which ranges from 0.00306-0.00374 Ω /°C. Once again it is important to note that the resistance is a derived value from the voltage drop over the heater and the current applied to it. From the observation of these slope values it can be deduced that the resistance of the heaters only changes by tenths of an ohm over the entire temperature operating range of the module (15-70°C).

The thermocouples used during the calibration were calibrated by INNOCAL® before use. Through the calibration results, adjustments were made to the temperature readings to reduce the error.

The test setup was designed so that the accuracy of the power source was not a concern. Since the voltage drop is taken at the individual heater level and the temperature based resistance is calibrated, the amount of power lost in the lines is accounted for in the thermal calculations. By observation however, the power supply at outputs of around 70W lost around 20W of power before reaching the module so not measuring at the heater level and calibrating would have introduced a significant error. Consistency of the current supplied to the test stand can be investigated in future work.

A harder to define source of error results from the assumption that the *h* of the cooler is constant throughout testing. As was previously discussed, *h* is a complex function which is usually empirically derived. One of the fundamental properties of *h* its change in result of the module temperature compared to the ambient. Therefore it changes as the devices temperature changes. However, since the flow rate from the chiller is monitored by a flow meter and the temperature of the cooler is known within hundreds of microns of the module surface, *h* was estimated to be approximately the same. An investigation into this should be done in future work.

A.5 Thermal Test Data

Module 1: Bare Silicon (Batch#.Trial#)	Total Power dissipated by heaters (W)	Average Diode Temp (°C)	Standard Deviation	Total Thermal Resistance (°C/W)	Number of Working Diodes (0-64)	Average Thermocouple Temp (°C)
1.1	0	21.67	0.24	N/A	58	21.9
	5.33	23.06	0.480355	0.23049	56	21.83
	10.43	24.358046	0.793213	0.245232	56	21.798
	21.260343	27.202724	1.4828	0.253958	56	21.8035
	35.95292	31.196409	2.447855	0.25789	56	21.9245
1.2	0	21.583263	0.257738	N/A	58	21.88
	5.324358	22.962003	0.459447	0.211857	55	21.834
	10.343037	24.282511	0.750471	0.236634	55	21.835
	21.009371	27.075824	1.414508	0.25007	55	21.822
	30.245593	29.576875	2.029829	0.25362	55	21.906
	41.207194	32.581918	2.81986	0.257951	55	21.9525
1.3	0	21.624322	0.264474	N/A	58	21.9
	5.362535	22.819312	0.401048	0.185325	56	21.8255
	13.644931	24.86657	0.809675	0.22104	56	21.8505
	25.730066	27.877212	1.459035	0.233412	55	21.8715
	41.85463	31.853821	2.333513	0.237998	55	21.8925

Table 6: Summary of module 1 thermal testing data

Module 2: 2 Layers no metal (Batch#.Trial#)	Total Power dissipated by heaters (W)	Average Diode Temp (°C)	Standard Deviation	Total Thermal Resistance (°C/W)	Number of Working Diodes (0-64)	Average Thermocouple Temp (°C)

1.1	0	21.96181	0.252771	N/A	62	21.97
	5.978509	24.201941	0.679663	0.384032	62	21.834
	11.729189	26.376548	1.197838	0.382554	62	21.8895
	23.998595	31.229069	2.334714	0.384838	62	21.9935
	34.697791	35.62761	3.392466	0.392103	62	22.0225
	48.915512	40.971157	4.736351	0.387559	62	22.0135
1.2	0	21.903459	0.251449	N/A	62	21.95
	6.244144	24.000213	0.644596	0.340273	62	21.8755
	12.253011	26.135102	1.121445	0.351065	62	21.8335
	25.090398	30.929869	2.222004	0.363022	64	21.8215
	36.219386	35.175087	3.207579	0.368341	64	21.834
	49.485948	40.368827	4.432407	0.373901	64	21.866
1.3	0	21.822056	0.274462	N/A	59	21.74
	6.271804	23.95648	0.633249	0.362412	59	21.6835
	12.306799	26.112736	1.055574	0.356854	64	21.721
	25.181569	30.701619	2.004233	0.354629	64	21.7715
	36.348602	34.795994	2.892828	0.358597	64	21.929
	49.606489	39.739371	3.980245	0.3624	64	21.762

Table 7: Summary of module 2 thermal testing data

Module 3: 4 Layers no metal (Batch#.Trial#)	Total Power dissipated by heaters (W)	Average Diode Temp (°C)	Standard Deviation	Total Thermal Resistance (°C/W)	Number of Working Diodes (0-64)	Average Thermocouple Temp (°C)
1.1	0	21.727698	0.36472	N/A	63	21.8
	5.506351	24.856657	0.457006	0.567555	63	21.7315
	10.806284	27.70396	0.593583	0.546993	63	21.793
	22.124476	33.916753	1.117997	0.550104	63	21.746

	31.927338	39.479792	1.635669	0.554064	63	21.79
	43.590955	46.064534	2.289991	0.556022	63	21.827
1.2	0	21.516784	0.388359	N/A	64	21.81
	5.500738	24.649038	0.439664	0.522664	64	21.774
	10.798139	27.508452	0.538437	0.527725	64	21.81
	22.1143	33.684623	0.990171	0.539249	64	21.7595
	31.939387	39.17964	1.496023	0.544332	63	21.794
	43.594235	45.80237	2.145195	0.549634	63	21.8415
1.3	0	21.588138	0.423458	N/A	61	21.8
	5.387594	24.593331	0.4261	0.502326	61	21.887
	11.137748	27.620921	0.610224	0.514819	60	21.887
	22.801763	33.926911	1.164058	0.527302	60	21.9035
	32.911493	39.506276	1.758159	0.535065	59	21.8965
	44.975325	46.234663	2.491465	0.539433	59	21.9735

Table 8: Summary of module 3 thermal testing data

Module 4: 4 Layers with metal (Batch#.Trial#)	Total Power dissipated by heaters (W)	Average Diode Temp (°C)	Standard Deviation	Total Thermal Resistance (°C/W)	Number of Working Diodes (0-64)	Average Thermocouple Temp (°C)
1.1	0	21.782747	0.223014	N/A	64	21.854
	5.59519	23.744983	0.680354	0.34994	64	21.787
	10.973145	25.677994	1.191275	0.354547	64	21.7875
	22.449745	29.920452	2.342944	0.364969	64	21.727
	32.399106	33.778116	3.406579	0.368609	64	21.8355
	44.201548	38.515796	4.742907	0.378783	64	21.773
1.2	0	21.892836	0.258262	N/A	64	21.92
	5.595259	23.71328	0.693753	0.339355	64	21.8145

	10.974822	25.620074	1.216603	0.349124	64	21.7885
	22.446802	29.788591	2.370463	0.354242	64	21.837
	32.391501	33.598554	3.455237	0.362257	64	21.8645
	44.184661	38.193159	4.790348	0.369261	64	21.8775
1.3	0	21.721147	0.275834	N/A	59	21.84
	5.597458	23.79996	0.766543	0.350777	64	21.8365
	10.977824	25.84846	1.353529	0.359403	64	21.903
	22.451129	30.317796	2.678061	0.374271	64	21.915
	32.394843	34.313583	3.891815	0.382563	64	21.9205
	44.201178	39.302114	5.418638	0.392402	64	21.9575

Table 9: Summary of module 4 thermal testing data

Module 6: bare silicon (Batch#.Trial#)	Total Power dissipated by heaters (W)	Average Diode Temp (°C)	Standard Deviation	Total Thermal Resistance (°C/W)	Number of Working Diodes (0-64)	Average Thermocouple Temp °C)
1.1	0	20.900965	0.271106	N/A	59	21.85
	6.396304	22.527694	0.598714	0.120725	59	21.7555
	12.53349	24.100846	1.056485	0.179467	59	21.8515
	25.633339	27.688607	2.097262	0.223639	59	21.956
	36.985482	30.862657	3.03138	0.240058	59	21.984
	50.515957	34.738323	4.196604	0.25158	59	22.0295
1.2	0	20.806105	0.27936	N/A	59	21.81
	6.390728	22.404297	0.605883	0.099331	59	21.7695
	12.514735	23.971669	1.0497	0.17289	59	21.808
	25.581808	27.433863	2.057549	0.219604	59	21.816
	36.906347	30.444771	3.126144	0.232461	59	21.8655
	50.39546	34.335404	4.324602	0.246231	59	21.9265

1.3	0	20.766765	0.276671	N/A	59	21.78
	6.434592	22.484457	0.670764	0.108936	59	21.7835
	12.601478	24.17688	1.214841	0.189214	59	21.7925
	25.760453	27.769645	2.370151	0.228806	59	21.8755
	37.140784	30.975356	3.446371	0.243973	59	21.914
	50.7214	35.16139	4.812609	0.260884	59	21.929

Table 10: Summary of module 6 thermal testing data