

# Calibration Methodology for Interposing Liquid Coolants for Infrared Thermography of Microprocessors

Lewis Hom<sup>1</sup>, Antoine Durieux<sup>1</sup>, Josef Miler<sup>1</sup>, Mehdi Asheghi<sup>1</sup>, Karthik Ramani<sup>2</sup>, Kenneth E. Goodson<sup>1</sup>

<sup>1</sup>Dept. of Mechanical Engineering, Stanford University

440 Escondido Mall, Building 530, Room 224

Stanford, CA, USA, 94305

Phone: (650) 725-2086

Fax: (650) 723-7657

Email: [lrhom@stanford.edu](mailto:lrhom@stanford.edu)

<sup>2</sup>Advanced Micro Devices, Inc.

## ABSTRACT

There have been unprecedented temperature non-uniformities reported in conventional 2D- and emerging 3D-packaged multi-core microprocessors. Therefore, techniques for spatially resolving chip power and temperature profiles in fully operational chips are needed to improve circuit design and enable optimal cooling solutions. This paper presents a high-resolution, infrared (IR) thermography technique for microprocessors operating at fully-operational power levels. A custom, microfluidic heat sink with an IR-transparent working fluid (0.75 LPM) is manufactured to cool an instrumented test chip while permitting optical access for IR thermal imaging. A detailed system calibration is conducted to account for the temperature-dependent optical properties of the chip and heat sink. It is concluded that the IR imaging can be conducted with  $\sim 0.1$  °C error over the temperature range of 45-90 °C if the fluid plenum height is less than 500  $\mu\text{m}$ . For a 2 mm channel, the error can be as high as 43°C due to strong signal attenuation in the fluid.

**Key words:** Infrared Imaging, Heat sink, Liquid Cooling, Microprocessor Thermal Management

## NOMENCLATURE

|    |                                                                            |
|----|----------------------------------------------------------------------------|
| Re | Reynolds Number                                                            |
| Nu | Nusselt Number                                                             |
| T  | Temperature (°C)                                                           |
| h  | Convective heat transfer coefficient<br>( $\text{W}/\text{m}^2/\text{K}$ ) |
| x  | Distance along chip (m)                                                    |

## Subscripts

|      |                           |
|------|---------------------------|
| L    | liquid                    |
| C    | chip                      |
| fd,t | Thermally fully developed |
| max  | Maximum                   |

## INTRODUCTION

The scaling of transistors has resulted in unprecedented levels of power dissipation in microprocessors, which has emerged as a major

design constraint [1-3]. Furthermore, increasingly severe power distribution non-uniformities in 2D and 3D-packaged microprocessors have resulted in large temperature non-uniformities [4-11]. While traditional heat sinks, which constitute metal heat spreaders cooled by forced convection dissipate significant heat, researchers have proposed clock-gating methods to dynamically alter chip processing traffic in order to prevent hotspots from forming in the first place [12,13].

Previous work in [6,14,15] has shown that infrared thermography, combined with modeling, is a useful tool to validate possible dynamic power routing schemes. Measurement of the temperature map resulting from realistic processing traffic can be used to drive chip power models and will lead to more accurate representations of chip temperature for thermal designers. In order to optically probe the chip during operation, an IR-transparent heatsink is required to maintain the chip's temperature below the critical temperature while maintaining high enough granularity [8-10] and contrast in the thermal map to be able to calculate the power map from the temperature map. Previous works [14] and [15] have developed microfluidic heatsinks to support infrared thermography of a microprocessor in operation.

Ardestani et al., in [14], approach the cooling problem using a sapphire heat spreader cooled by a mineral oil flow. The dimensions of the sapphire cooling window and the oil flow conditions are chosen to dissipate the same amount of heat as a conventional heat sink. They focus on determining the thermal response of a single point on a chip due to current pulse modulation. This work shows how a fluidic heat sink can be designed with the same steady state and transient thermal responses as a metal heatsink for a single point. While the approach mimics cooling from a metal heatsink, it does not necessarily give temperature maps that are useful for solving for the chip's power map. The spatial variation in temperature is likely to be convoluted by the variation in convective heat transfer across the chip. This convolution of the temperature field will then complicate the calculation of a spatially-resolved power map.

Boiadjieva et al., in [15], developed an air-cooled diamond-window heatsink and determined the impact of the chip's power on the heatsink temperature. The highly thermally conductive diamond window provided both optical access and heat spreading. It was able to dissipate 150 W while maintaining the device temperature between 60 °C - 100 °C. They then used a mathematical model to determine the wattage and average chip temperature based on temperature measurements on the heat sink. However, this approach does not recover the spatially-resolved power map, which is necessary to perform hotspot diagnostics and inform clock gating techniques. In particular, a solid heat sink, which is shown to exhibit strong spatial variations in chip heat removal, greatly increases the complexity of the mathematical model required to calculate the power map. In an ideal case, the heat removal rate would be uniform across the chip surface.

Prejean et al., in [18], have developed a calibration methodology to validate to accuracy of the temperature map across the chip. Infrared radiation emitted from the chip backside will be attenuated by the heatsink components, distorting the resulting temperature map. In particular, the cooling liquid has temperature-dependent optical properties which must be accounted for. This methodology accounts for the attenuation to produce an accurate temperature map, but not necessarily an accurate power map.

There has been little work on developing a calibration methodology that produces a temperature map of both sufficient accuracy and contrast to accurately calculate the power map. This methodology must account for attenuation from the heatsink elements in the IR detector's optical path and the spatial variation in heat removal from the chip, which both affect the thermal and power map.

The present study focuses on the design of a microfluidic heatsink and calibration methodology that intends to both minimize temperature map distortion and maximize power map accuracy. First, the work discusses the experiment design parameters that maximize temperature map and power map accuracy. The work will discuss a preliminary experiment where a prototype heatsink with 0.75 LPM flow rate in a 2 mm plenum is used to cool a test chip in operation and develop the calibration methodology. Sources of error in the preliminary experiment and recommendations to further improve the experiment are discussed in the troubleshooting section.

## EXPERIMENT DESIGN REQUIREMENTS

The high power density of the chip and the limitations of the IR microscope impose multiple design requirements on the rig and flow system. The power density of the test chip is expected to be as high as 150 W/cm<sup>2</sup> during periods of high traffic, and in order to preserve chip integrity, the chip must be maintained at a temperature below 130 °C. Any cooling solution must also provide optical access to allow for infrared imaging.

During the IR microscope's calibration process, the chip must be maintained at a uniform and known temperature above 60 °C in order to gather sufficient radiance data to determine the emissivity map of the chip. Furthermore, the optical properties of all interposing media in the optical path between the chip backside (the imaging target of interest) and photodetector must be known and accounted for, since all of these both attenuate the radiation from the backside and add their own emitted radiation to the signal, as shown in Figure 1.




**Figure 1: Cross-sectional schematic of the custom, microfluidic heat sink mounted on a chip. Radiation emitted from the chip and the fluid and is attenuated by intermediate layers; a detailed calibration is required to account for these effects.**

In the proposed cooling solution, the chip forms the floor of a microscale plenum that passes liquid over the chip surface, dissipating heat. The liquid is cooled by a recirculating chiller. We chose Galden HT-170 (see Appendix for properties), a fluorinated liquid used in electronics cooling due to its high boiling point, chemical inertness, thermal stability, and compatibility with metals and plastics. High boiling point and thermal stability are of critical importance in this experiment, as hotspots can lead to local boiling or thermal decomposition, which in turn would drastically vary the heat transfer across the chip. Galden HT-170's boiling point and thermal decomposition temperature (170 °C and 290 °C respectively) both exceed the chip maximum temperature.

The use of liquid cooling poses another problem for infrared thermography. Since the working liquid will incur a temperature rise as it traverses the chip, the change in optical properties of the liquid due to the temperature rise will distort the thermal map. Thus, we must minimize the variation in liquid temperature across the chip. The silicon and sapphire layers of the stack are much more transparent at the chip operating temperature range are therefore less of a concern than the liquid. A further concern with respect to flow system design is the fact that the accuracy of the power map from the temperature map is greatly increased when the convective heat transfer coefficient ( $h$ ) is constant over the surface of the chip. The flow system requirements to maximize accuracy of the power map are summarized as follows:

- 1) The maximum local temperature rise may not exceed 130 °C in order to ensure device integrity.
- 2) The maximum pressure in the plenum may not exceed 60 psi. Higher pressures may cause failure in the sapphire window and exceed the maximum pressure rating of the pump.
- 3) The liquid temperature rise may not exceed 0.1% of the chip temperature rise in order to minimize emission from the liquid.
- 4) The local Nusselt number ( $Nu_x$ ) must be relatively uniform across the chip to maximize power map accuracy.
- 5) The plenum height may not exceed 500  $\mu\text{m}$  in order to minimize attenuation of the infrared signal by the liquid.

In order to determine a flow condition and plenum geometry that would fulfill all of the above conditions, an analytical model was developed to determine plenum  $Nu_x$ , chip temperature, liquid temperature, and channel pressure drop as a function of flow rate and plenum height. The model uses analytical solutions for  $Nu_x$  for hydrodynamically fully developed, thermally developing flow between two parallel plates with one heated wall (representing the chip) and one adiabatic wall (representing the sapphire window) [16,17]. The model prescribes a uniform heat flux of 150 W/cm<sup>2</sup> for the heated wall, which represents a conventional GPU processor dissipating 135W.

The model results are presented in Figures 2 and 3 and show that the optimal flow condition is for a 500  $\mu\text{m}$  plenum with a flow rate between 3 and 5 LPM: the minimum flow rate required to maintain turbulent flow in the plenum. At lower flow rates, the flow is in the laminar regime, which has a number of disadvantages. Laminar flow decreases the heat




Figure 2: Chip maximum temperature rise as a function of flow rate and plenum height for a chip with uniform heating of 150 W/cm<sup>2</sup>. The operating regime is bounded by dashed lines. Maintaining turbulence drastically decreases chip temperature relative to laminar flow, but has decreasing returns with higher flow rates. The maximum flow rate for a given plenum height is further bounded by the maximum tolerable pressure rise in the system. Increasing the plenum height decreases system pressure at the expense of decreased heat transfer.




Figure 3: Thermal entry length ( $x_{fd,t}$ ) as a function of flow rate and plenum height for a chip with uniform heating of 150 W/cm<sup>2</sup>. The operation is bounded by dashed lines. Thermal entry length in the laminar regime is greater than the channel length and is not shown here. A cap of 4 mm (15% of total chip length) has been imposed to minimize power map distortion. The maximum allowable flow rate for this experiment is bounded by the maximum tolerable  $x_{fd,t}$  and the maximum allowable pressure drop in the channel. Increasing the channel length at a given flow rate decreases the pressure drop at the expense of increased thermal development length.

transfer, causing the chip temperature to exceed the maximum specifications. Furthermore, laminar flow greatly increases the thermal development length, which in turn increases the variation of the  $Nu_x$ . After reaching turbulence, though, higher flow rates will increase the thermal development length and with it, the variation in  $Nu_x$ . Increasing the plenum height decreases the pressure drop in the channel at the cost of decreased heat transfer and transmissivity. Transmissivity tests, detailed in the Troubleshooting section, show a negligible loss of transmissivity for a 500  $\mu\text{m}$  film of liquid. The maximum temperature rise in the channel is also within the acceptable range with a safety factor of  $0.7\Delta T_{max}$ . Therefore, the 500  $\mu\text{m}$  thick plenum is a good, conservative choice that can provide a flow rate that maintains low chip temperature and thermal entry length while allowing for the inevitable pressure head losses in the pump and flow system.

### EXPERIMENT SETUP

In the experiment, a Themaltake Bigwater 760 pump-heatsink-fan assembly was used to pump the working liquid and regulate its temperature. Due to the low pump power, a relatively large, 2 mm high plenum with cross-sectional area of 21  $\text{mm}^2$  was used in order to maintain a sufficient flow rate to cool a test chip dissipating 30 W and allow for the development of the calibration methodology. The custom-fabricated fluidic rig is shown in Figures 4 and 5. A more powerful, Thermoflex 3500 recirculating chiller capable of up to 56.8 LPM at 57 psi will operate within the operating region delineated in Figures 2 and 3.




**Figure 4: Side cross-section of the CAD model of microfluidic heat sink mounted on the chip. The working fluid flow path is shown in blue.**

### EMISSION CALIBRATION METHODOLOGY

Accurate infrared thermometry requires an emissivity map of the sample, which allows for the calculation of the temperature map from the measured radiance. The infrared microscope's normal calibration process calculates the emissivity map of the sample and will be referred to as the "emissivity calibration." The emissivity map

establishes a direct relationship between the radiance map, which the microscope's photodetector measures directly, and the temperature map. For this calibration, the sample of interest is heated uniformly to 75  $^{\circ}\text{C}$ . This allows for enough radiation to be emitted to accurately calculate the emissivity. The microscope then measures the radiance map of the sample. Internal software uses a relationship between radiance and temperature to calculate a pixel-by-pixel emissivity map of the sample. This completes the emissivity calibration, and now the microscope can calculate the temperature map based on the measured radiance and the emissivity calibration. When imaging through a stack of interfering media, the resulting emissivity table is the "equivalent emissivity" of the stack.

However, this process is insufficient for the microfluidic heatsink due to the presence of a temperature-invariant liquid layer with strongly temperature-dependent optical properties. During the imaging process, the liquid is maintained at 25  $^{\circ}\text{C}$  in order to increase heat transfer from the chip surface. The equivalent emissivity table of the stack obtained from calibration makes an assumption that the entire stack changes temperature and does not account for the fact that the liquid film, while significantly attenuating the signal, does not change temperature. This necessitates a new calibration procedure that removes the influence of the 25  $^{\circ}\text{C}$  liquid from the temperature measurement: the radiance table calibration.




**Figure 5: Photograph of the microfluidic heatsink mounted on the test chip with pressure sensing ports and sapphire window shown. The liquid is injected and extracted from the rig via the tubes and is thermally regulated by a pump-heatsink-fan assembly.**

## RADIANCE TABLE CALIBRATION METHODOLOGY

In order to calibrate out the effect of liquid film, it is necessary to modify the microscope software's correlation between radiance, emissivity, and temperature. This can be accomplished by modifying the radiance calibration tables to account for the high attenuation of the liquid. The emissivity calibration is insufficient for this because it only calculates the equivalent emissivity of the sample. A conventional, factory radiance table calibration involves measuring the radiance of a temperature-controlled blackbody target and using these measurements to build a relationship between radiance, temperature, and emissivity, which is then stored in the radiance table. The modified radiance table calibration is done by running the factory calibration test and interposing a representative silicon/HT-170/sapphire stack between the blackbody target and the photodetector.

The radiance table must be calibrated twice: once for the emissivity calibration, when the liquid is maintained at 75 °C ("hot calibration"), and once for the temperature measurement, when the liquid is maintained at 25 °C (the "cold calibration"). Once the radiance calibration is complete, the attenuation effect of the liquid is theoretically removed for each

case as long as the liquid does not change temperature during the measurement. Afterwards, the emissivity calibration can be completed as normal using the hot radiance tables. The resulting emissivity map is the emissivity map of only the chip backside, since the effect of the silicon, sapphire, and liquid are fully accounted for in the radiance tables. With the emissivity map in hand, the liquid can be regulated at 25 °C and temperature measurements can commence under the cold radiance table. The process of swapping radiance calibration tables for emissivity mapping and temperature mapping is illustrated in the flowchart in Figure 6.

## PRELIMINARY RESULTS

In an experiment to validate the calibration methodology, a test chip was uniformly heated at three different power levels and thermal images were compared to onboard temperature diode measurements. The flow parameters of the system are summarized in Table 1. A representative thermal image of the test chip is shown in Figure 7 and the thermal measurements from the infrared microscope and temperature diodes are compared in Figure 8.

The IR and diode temperatures differ by a maximum of 42.6 °C and the slope of the IR-




Figure 6: Flow chart depicting the emissivity calibration procedure. Radiance is measured at 75 °C (red region) and 25 °C (green region) to account for the temperature-dependent optical properties of the liquid.

measured temperature is 0.58 °C/W versus the diode-measured temperature slope of 1.92 °C/W. Furthermore, the heaters, which are shown in Figure 7 as metallization patterns, incorrectly register as cool spots in the infrared image. Significant troubleshooting, described in detail in the next section, shows that this discrepancy is due to heavy attenuation of the chip's radiation by the liquid layer, which results in a low-contrast emissivity map.

**Table 1: Flow parameters for the experiment used to develop the calibration methodology.**

|                                              |       |
|----------------------------------------------|-------|
| Channel Area [mm <sup>2</sup> ], Height [mm] | 21, 2 |
| Flow Velocity [m/s]                          | 0.28  |
| Volumetric Flow rate [LPM]                   | 0.75  |
| Re <sub>L</sub>                              | 585.2 |




**Figure 7: Representative thermal map of the test chip dissipating 25.6 W. The test chip is instrumented with uniformly spaced heaters, which are shown in the image as metallization patterns. The top, cool region of the chip corresponds to an area of the chip without instrumented heaters. The metallization patterns are significantly cooler than the surrounding passivation layers, which is an artifact of an emissivity calibration done while the chip's radiation was heavily attenuated by a thick liquid layer.**




**Figure 8: IR temperature reading vs. those measured using onboard temperature diode measurement. The large error in the IR measurement is caused by signal attenuation in the working fluid.**

## TROUBLESHOOTING

Extensive steps were taken to troubleshoot both the temperature offset and temperature /power slope measurement errors. A constant temperature offset alone would indicate an inconsistency in the temperature set points during the emissivity calibration process. However, the fact that the infrared-measured temperature and on-chip diode temperature follow different slopes suggests an error in the emissivity map. To investigate this, emissivity maps of the chip were measured under both open air and conditions, with the appropriate radiance table calibration used in acquiring the emissivity. These emissivity maps are shown in Figure 9.




**Figure 9: (Top) Emissivity map of a chip measured in open air. (Bottom) Emissivity map of a chip measured through silicon, Golden HT-170, and sapphire using the 75 °C hot liquid calibration table. Reducing the Golden HT-170 thickness from 2 mm to 500 um dramatically improved results.**

The emissivity maps should theoretically be the same, since the radiance table calibration is intended to remove the influence of the hot liquid. However, the emissivity map resulting with the interposing hot liquid film has decreased contrast between backside features and has a nonphysical average emissivity of 1.15.




**Figure 10:** Schematic of the test rig to determine the temperature measurement error as a function of liquid temperature and liquid thickness. A blackbody target of known temperature is imaged through layers of interfering media representative of GPU backside measurement. By measuring the radiance and temperature of the blackbody using the IR camera, the attenuation due to the interposing media can be determined.

The decreased contrast suggests that the signal attenuation and radiation emitted by the rig, as opposed to the chip backside, is so high that it eliminates emissivity contrasts of the circuitry. This means that the image depicts the interfering media itself, rather than the chip backside. To find the source of this attenuation, numerous temperature readback tests were done to estimate the attenuation effect of each part of the rig being imaged: the silicon die, the Galden HT-170, and the sapphire window.

The silicon die and sapphire window were tested individually. They were tested by imaging a temperature-controlled blackbody through a window of the tested material, performing an emissivity calibration, and then attempting to read back the temperature. Both the polished silicon and sapphire, being highly transparent in the MWIR range (3.4 to 5.1  $\mu\text{m}$  for this particular photodetector), yielded less than a 1% error in readback temperature. After establishing that sapphire had a negligible effect on the temperature measurement, the blackbody radiance was measured through a temperature-controlled, 2 mm thick Galden-HT170 film. The schematic of this experiment is shown in Figure 10.

The A/D radiance counts measured through the liquid as a function of the liquid temperature and the blackbody target temperature is shown in Figure 11.

According to the results of this experiment, the measured radiance is extremely dependent on the liquid temperature. At a given liquid temperature, the dynamic range available is extremely low: on average, the photodetector only reads 50 digital




**Figure 11:** Radiation incident on photoreceptor (measured in microscope analog/digital “counts” units) from a blackbody calibration target for varying liquid and blackbody temperatures for a 2 mm thick film of liquid. The measured radiation is heavily affected by the liquid temperature: the radiance has a sensitivity of 27.3 counts per  $^{\circ}\text{C}$  change in liquid temperature, as opposed to 1.1 counts per  $^{\circ}\text{C}$  change in blackbody target temperature.

counts of difference between a 45  $^{\circ}\text{C}$  and 90  $^{\circ}\text{C}$  blackbody. Typical infrared measurements require roughly 1500 digital counts of range for accurate emissivity calibration and thermal maps. The radiance calibration then is ineffective because the liquid proved to be essentially opaque in the photosensitive range of the IR microscope.

The attenuation due to the liquid was confirmed by a radiance test taken through two different liquid thicknesses. In this test, a temperature-controlled blackbody target was measured with the IR microscope through a 0.5 mm thick film of liquid, a 2 mm thick film, and in the open air. As shown in Figure 12, the 2 mm layer of liquid drastically decreased the dynamic range of the camera while the 500  $\mu\text{m}$  film displayed a large dynamic range. These results suggest the use of a 500  $\mu\text{m}$  film of liquid because the increased dynamic range will allow for the acquisition of more accurate emissivity maps, which in turn will provide more accurate temperature maps. To confirm this hypothesis, another simple test was designed to read back the temperature of a temperature-controlled blackbody through both 2 mm and 500  $\mu\text{m}$  thick films of liquid. This test was successful: the 500  $\mu\text{m}$  thick film of liquid could read back the temperature of the blackbody target to within 0.1  $^{\circ}\text{C}$  accuracy at temperatures above 60  $^{\circ}\text{C}$ . The result of this test is shown in Figure 13.

The high accuracy of temperature measurements through a 500  $\mu\text{m}$  thick film makes it an attractive candidate for achieving higher temperature resolution




**Figure 12: Radiance measurement of a temperature-controlled blackbody through varying quantities of Golden HT-170. The 0.5 mm film is almost transparent, whereas the 2 mm film is effectively opaque.**




**Figure 13: IR-measured temperature of a temperature-controlled blackbody target. A 2mm liquid film causes large temperature errors; a 500 μm thick liquid shows less than a 0.1 °C error at temperatures above 60 °C.**

future experiments. However, the trade-off in using a thinner channel is that there will be a higher variability of  $Nu_x$ . The analytical model results in Figure 2 show that the thermal entry length is between 3 and 4 mm in the ideal operating regime – roughly 15% of the chip. The power map is expected to be slightly skewed in this area, and future work can be done to quantify this distortion.

Discussions with Quantum Focus Instruments, the manufacturers of the infrared microscope, yielded the suggestion that Galden HT-170 is not simply IR transparent in the photodetector's range. This motivates a change in the working liquid. Two additional electronics cooling liquids were used, including perfluorooctane and Fluorinert FC-70. Perfluorooctane was the IR-transparent liquid of choice in previous work in [6], and Fluorinert FC-70 exhibits a >90% spectral transmissivity in the desired MWIR range. However, each of these proved to also

appear opaque to our photodetector. These troubleshooting results imply that the problem of the low-contrast emissivity map is caused by a combination of the photodetector and the working liquid. The addition of the working liquid (including those known to be transparent at IR spectrum strongly decreases the dynamic range of the camera. The liquid transmissivity is even higher than silicon, which is only 55% transmissive in the 1.5 to 6.0 μm range and yet can be interposed between the photodetector and blackbody target with minimum temperature measurement error. These tests overall show that despite high transmissivity in the MWIR regime, a 2 mm thick liquid film can prove to be opaque to the photodetector. This motivates further testing to determine the liquid thickness-dependence of transmissivity in order to estimate the maximum liquid film thickness for future experiments.

## CONCLUSION

A methodology is developed for back-side IR thermal imaging of fully operational microprocessors, which accounts for the temperature-dependent attenuation of the liquid by modifying the infrared microscope's radiance tables. The conclusion is that the current methodology does not produce accurate temperature maps when the dynamic range of the infrared microscope is drastically decreased by the radiation attenuation and emission from the cooling liquid.

Experiments confirm that the dynamic range of the infrared microscope can be greatly enhanced by decreasing the thickness of the liquid layer. It is possible to achieve greater than 0.1 °C accurate IR temperature measurements through the 0.5 mm thick liquid film. Analytical simulation has confirmed that 0.5 mm is sufficiently thick of a plenum to maintain the chip temperatures below the failure temperature. However, this flow system will lead to higher variability in the Nusselt number throughout the channel, which may lead to error in the power map calculation. Future experiments will involve determining the thickness-dependence of transmissivity of the liquid in order to determine the optimum liquid film thickness with respect to power map accuracy.

## ACKNOWLEDGEMENTS

This research was supported in part by a grant from AMD. In addition, Lewis Hom is supported by an NSF Graduate Research Fellowship.

## APPENDIX: GALDEN HT-170 PROPERTIES

All properties, except boiling point, are specified at 25 °C.

|                              |                         |
|------------------------------|-------------------------|
| Boiling Point [°C]           | 170                     |
| Kinematic Viscosity [cSt]    | 1.8                     |
| Density [kg/m <sup>3</sup> ] | 1.77                    |
| Thermal Conductivity [W/m/K] | 0.0065                  |
| Volume Resistivity [Ω-cm]    | 1.50 x 10 <sup>15</sup> |
| Specific Heat [cal/g/K]      | 0.232                   |

## REFERENCES

[1] M. Horowitz, E. Alon, D. Patil, S. Naffziger, R. Kumar, and K. Bernstein, "Scaling, power, and the future of CMOS," in Proc. *IEDM*, Dec. 2005 [Online]. Available: [http://mos.stanford.edu/papers/mh\\_iedm\\_05.pdf](http://mos.stanford.edu/papers/mh_iedm_05.pdf)

[2] Xiuyi Zhou, Jun Yang, Yi Xu, Youtao Zhang, and Jianhua Zhao, "Thermal-Aware Task Scheduling for 3D Multicore Processors," *IEEE Transactions on Parallel and Distributed Systems*, Vol. 21, No. 1, pp. 60-71, 2010.

[3] Pop, E. and Goodson, K. E., "Thermal phenomena in nanoscale transistors," Proc. The Ninth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM), Las Vegas, NV, USA, Vol. 1, pp. 1-7, 2004.

[4] Hamann, H. F., Weger, A., Lacey, J., Cohen, E., and Atherton, C., "Power distribution measurements of the PowerPCTM970MP microprocessor," Proc. *IEEE International Solid-state Circuits Conference (ISSCC)*, 2006.

[5] Hamann, H. F., "Spatially resolved imaging of microprocessor power," Proc. *IEEE International Solid-State Circuits Conference (ISSCC)*, p. 16., 2005.

[6] Hamann, H. F., Lacey, J., Weger, A., and Wakil, J., "Spatially-resolved imaging of microprocessor power (SIMP): hotspots in microprocessors," Proc. *The Tenth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronics Systems (ITHERM)*, pp. 121-125, 2006.

[7] Yazdani, K., Hamann, H. and Asheghi, M., "Impact of Power Granularity on Chip Thermal Modeling," *IEEE Transactions on Components and Packaging Technologies*, *IEEE Transactions on Components and Packaging Technology* Volume: 31 Issue: 1, pp. 211-15, 2008.

[8] Yazdani, K., and Asheghi, M., "Spatial Frequency Domain Heat Transfer Analysis of Hot Spot

Spreading in Convectively Cooled Microprocessors," Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM), May 28 - 31, 2008, Orlando, Florida.

[9] Yazdani, K., Hamann, H. and Asheghi, M., "Spatial Frequency Domain Analysis of Heat Transfer in Microelectronic Chips with Applications to Temperature Aware Computing," *Interpack '07 Conference*, July 8-12, 2007, Vancouver, British Columbia, CANADA.

[10] Yazdani, K., Hamann, H. and Asheghi, M., "Fast and Accurate Simulation of Heat Transfer in Microarchitectures Using Frequency Domain Techniques," *Interpack '07 Conference*, July 8-12, 2007, Vancouver, British Columbia, CANADA.

[11] Kai, X. and Ketchen, M., "Thermal modeling of a small extreme power density macro on a high power density microprocessor chip in the presence of realistic packaging and interconnect structures," Proc. *54th Electronic Components and Technology Conference*, Las Vegas, NV, USA 2004, Vol.1, pp. 918-923.

[12] Skadron, K., Abdelzaher, T., and Stan, M. R., "Control-theoretic techniques and thermal-RC modeling for accurate and localized dynamic thermal management," Proc. *Eighth International Symposium on High-Performance Computer Architecture*, Cambridge, MA, USA, 2002, pp. 17-28.

[13] Skadron, K., Stan, M. R., Huang, W., Sivakumar, V., Karthik, S., and Tarjan, D., "Temperature-aware microarchitecture," Proc. *30th International Symposium on Computer Architecture*, San Diego, CA, USA, 2003, pp. 2-13.

[14] E. K. Ardestani, F. Mesa-Martinez, J. Renau. "Cooling Solutions for Processor Infrared Thermography," Proc. *of the Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM)*, Santa Clara, CA, USA, 2010, pp. 187-190.

[15] Boiadjieva, N., Kelleher, C., Dajee, G., Nadig, S., "Novel Heat Spreader for the Optical Probing of PC Board Mounted Flip Chips," Proc. *Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM)*, San Jose, CA, USA, 2003, pp. 284-289.

[16] Shah, R.K., London, A.L., *Laminar Flow Forced Convection in Ducts*. New York: Academic Press, 1978, pp. 309.

[17] Hatton, A.P., Quarmby, A., "The Effect of Axially Varying and Unsymmetrical Boundary Conditions on Heat Transfer with Turbulent Flow between Parallel Plates". *International Journal of Heat and Mass Transfer*. Vol. 6, pp. 903-914, 1963.

[18] Prejean, S., Santana, M., Potok, R. "Infrared Sensor and Method of Calibrating the Same." U.S. Patent 7 795 589 B2, Sept. 14, 2010.