

A Stochastic Coolability Analysis For A Microprocessor Considering Chip Power Fluctuation

Sai Ankireddi
Sun Microsystems Inc.
4120 Network Circle, Santa Clara, CA 95054
sai.ankireddi@sun.com

Abstract

As process technology continues to evolve, the shrinking gate sizes in today's microprocessors are accompanied by a notable increase in the die power. While increased power is a challenge to the cooling design in its own right, another significant aspect is that the distribution of power on the die (silicon) can vary dynamically depending on the software application that is running. As multi-core, multi-threaded chips become more prominent, these problems are compounded even further. There is another aspect of package cooling analysis which is traditionally ignored, and that is the part-to-part variability in terms of materials and stack up geometry. These can significantly influence the cooling design, and it is important to be aware of their impact. In this study we focus on a stochastic analysis of a hypothetical multi core microprocessor. The power in each block is assumed to a random variable, as are the thickness and conductivities of the die, TIM1, lid (heatspreader), TIM2 and the sink thermal performance. A detailed Monte Carlo analysis of the package coolability is conducted, considering one scenario where power fluctuation is enabled and the other where power is treated as being static. The findings highlight the significance of considering the realistic power fluctuations in the cooling solution design.

Keywords

stochastic, coolability, chip power fluctuation

1. Introduction

In the high performance computing arena, the use of multi-core processors which support hardware threading is becoming more common [1]. These advanced microprocessors are used in high end servers that perform intensive data manipulation, graphical visualization and floating point computations.

As transistor technology evolves and gate lithography shrinks, chip power has continued to rise. One key observation is that the distribution of the power on the die can vary quite a bit depending on the software application that is running, the core activity migration policy in place, and the power management paradigm that was implemented in the design. The transient power distribution on the die can make for significant cooling challenges [2].

There is another aspect of package cooling analysis which is traditionally ignored, and that is the part-to-part variability in terms of materials and stack up geometry. These can significantly influence the cooling design, and it is important to be aware of their impact. With this in mind, in this study we focus on a stochastic analysis of a hypothetical multi core microprocessor. The power in each block is

assumed to a random variable, as are the thickness and conductivities of the die, TIM1, lid (heatspreader), TIM2 and the sink thermal performance. A detailed Monte Carlo analysis [3] of the package coolability is conducted, considering one scenario where power fluctuation is enabled and the other where power is treated as being static. The findings highlight the compelling need to consider the realistic power fluctuations in the cooling solution design.

2. Modeling Setup

The hypothetical CPU package under consideration has a 20 mm square die with four cores. Each core has a power of 30 W, and it is assumed that the IO power is fixed at 5 W- this results in a total die power of 125 W.

The silicon die was assumed to have a uniform distribution for both its thickness and conductivity, with the range of variation being 20u and 20 W/mK respectively. The TIM1 material was assumed to have a bondline thickness that was uniformly distributed with a 18u range of variation. The lid was assumed to have a uniform distribution for both its thickness and conductivity, with the

range of variation being 100u and 35W/mK respectively. The TIM2 material was assumed to have a bondline thickness that was uniformly distributed with a 30u range of variation, consistent with what was seen in empirical tests. Finally, the sink thermal resistance was modeled to be a normally distributed random variable, with a 6-sigma range of variability at ± 0.01 C/W.

In one set of computational models, the power was held fixed but all other random variables above were modeled as described. To study the effects of power distribution in the die, these models were then re-run assuming that the block powers were independent and uniformly distributed, each with a range of variability of 50%.

All computational models were simulated using standard Monte Carlo methods and the behavior of the sample means examined over 4000 iterations. The gathering of the simulation statistics was terminated when 4000 sample iterations were complete.

3. Convergence of Simulation Statistics

Figures 1-3 show the convergence of the sample statistics as the sample count increased. It is noted that by the time 4000 tuples were gathered and assimilated into the statistics summary the results had converged. In light of this, all simulations were terminated at 4000 tuples without loss of generality.




Fig. 1. Convergence of Rjc (package junction-case resistance) for both fixed & variable power cases.

4. Findings

Figure 4 shows a comparison of the model Rjc (package junction-case thermal resistance) for both fixed power and variable power cases in a histogram format. What is shown on the x-axis is the range of variability, in percent, and what is on the y-axis is the number of tuples that correspond to a given value on the x-axis.




Fig. 2. Convergence of Rcs (package case-sink resistance) for both fixed & variable power cases.




Fig. 3. Convergence of Rsa (package case-sink resistance) for both fixed & variable power cases.




Fig. 4. Rjc trends

As Fig. 4 shows, considering power variability the package thermal resistance R_{jc} goes up by 12% on average, and at the extremum is as high as 40% above that in the fixed-power case.

Figures 5 & 6 depict the trends in R_{cs} , the case-sink thermal resistance and R_{sa} , the sink-ambient thermal resistance.




Fig. 5. R_{cs} trends




Fig. 6. R_{sa} trends

It is noted that neither R_{cs} nor R_{sa} are significantly different, from a statistical viewpoint considering both the fixed and varying chip power scenarios.

At this time it is instructive to examine the histogram of power fluctuations that cause the extreme variation in the package R_{jc} noted earlier. This is depicted in Fig. 7. From this histogram it is apparent that the cumulative change in the die power (from a range standpoint) is only about $\pm 15\%$.




Fig. 7. Total die power fluctuations, including the effects of block level power variability

Figure 8 shows the distribution of the total path thermal resistance for the cooling solution (ie R_{ja} , or junction-ambient thermal resistance) for both fixed power and varying power cases. It is noted that on average, the junction-ambient thermal resistance is increased by 4%, and at the higher end, by about 11%.




Fig. 8. R_{ja} trends

From a cooling standpoint, it is now instructive to look at the histogram of the coolable power for the solution considered (Fig. 9). It is noteworthy that the 3-sigma low coolable power drops by almost 11%, and the average by almost 4%. These are significant quantities to keep in mind when finalizing a cooling solution.




Fig. 9. Coolable power trend

5. Analysis & Interpretation

For the microprocessor cooling design shown here, if we were to only consider the effects of material and geometry variability, the 3-sigma-low coolable power is roughly 7% lower than the nominal value, as Fig. 9 shows. However, once we include the effects of power fluctuation in the die, this number climbs to as high as 14%. For many chip cooling solutions, this level of drop in coolability may be unacceptable, but the thermal designer may not be aware of the fact that the coolable power can be compromised to such an extent by die power distribution variability.

6. Concluding Remarks

The current study highlights the need for thermal designers to acknowledge the potentially significant losses to coolable power if die power variability is considered in the cooling design. In the fictitious example shown here, the 3-sigma coolable power dropped by almost 14%. In margin-sensitive applications, such a significant drop in coolable power may be unacceptably high. This body of work showcases the need for making such considerations a standard part of cooling design practice to avoid potential field recalls.

7. References

- [1] W. Huang, M. R. Stan, K. Sankaranarayanan, Robert J. Ribando, and K. Skadron. "Many-Core Design from a Thermal Perspective", Proceedings of the 45th ACM/IEEE Conference on Design Automation (DAC), June 2008.
- [2] Etessam-Yazdani, K., Hamann, H.F., & Asheghi, M., "Impact of Power Granularity On Chip Thermal Modeling", Proceedings of ITherm 2006, The Tenth InterSociety Conference on Thermal and Thermomechanical Phenomena in Electronics Systems.

- [3] Jun S. Liu, Monte Carlo Strategies in Scientific Computing. Springer Series in Statistics, 2001.