Adaptive Power Blurring Techniques to Calculate IC Temperature Profile under Large Temperature Variations

Amirkoushyar Ziabari, Zhixi Bian, Ali Shakouri

Baskin School of Engineering, University of California Santa Cruz

Abstract

Power Blurring (PB) methods enable calculation of IC temperature profiles from the power dissipation map with calculation speeds hundreds of times faster than finite element methods (FEM). Both static [1,2] and transient distributions can be obtained [3]. Extensions to 3D chips [4] and to the inverse problem, i.e. estimating the power map from the temperature field [5], are available. So far however, the temperature dependence of the material parameters has been neglected. Temperature rises of 40-50°C on the chip will reduce the thermal conductivity of the silicon by 10-20%. This could affect the hot spot temperature by 5-7°C. In this work, we extend the PB approach to account for this effect. We propose two Adaptive Power Blurring (APB) methods based on iterative procedures. In both methods, the PB method provides an initial temperature distribution guess using room temperature Si thermal conductivity. Subsequent iterations take into account the preliminary temperature profile in the chip. The key difference between the two APB methods is the way the thermal masks are selected from a look-up table. The first variant uses one single mask based on the average temperature increase in the silicon, while the second approach employs a different mask for each point to account for the spatial variation of the temperature and according non-uniform thermal conductivity. In either case, the new estimate of temperature profile is acquired from convolution of the thermal masks and the IC Power map. These schemes are then applied iteratively until a final, self-consistent solution is reached. Good convergence is achieved only in 2-3 iterations in both methods. We will demonstrate that these APB methods substantially improve the accuracy under high temperature rise regime, in particular at hot spots, while still being much faster than traditional FEM computations.

Keywords

Adaptive Power Blurring, Temperature-dependent Thermal Conductivity, Power Blurring, Temperature Profile, Iterative Procedure.

1. Introduction

In recent years, newer and faster techniques have emerged as an alternative to traditional grid-based methods used to calculate temperature distribution by solving the heat equation with appropriate boundary conditions. One of these approaches is the Power Blurring method in which the temperature profile is calculated through convolution of the power map of the device under investigation and a thermal mask. This mask conforms to the impulse response (i.e. Green's function) of the system and represents how much of temperature rise occurs on the surface of a die due to a unit point heat source. The significance of the method was not only because of its ability to calculate the temperature profiles by two orders of magnitude (100 times) faster than grid-based methods, but also due to its advantage over the other current Green's function based techniques by considering the real geometry of packaged VLSI chips [1].

The Power Blurring method is further improved to predict IC chip temperature with high spatial resolution, which has been prohibitively expensive with conventional methods [2]. Using PB method, transistor level thermal maps (5×5 μm^2 grid) of a 5×5 mm^2 chip with a computation time of 20 seconds have been obtained.

Chip-level transient thermal simulation is another area that has been tackled by the PB method. With a minor adjustment, the time evolution of the thermal mask resulting from a spatiotemporal impulse is employed for transient simulations. It has been shown that the PB method can estimate the temperature profiles with errors less than 3% with a computation speed hundred times faster than the industry standard finite element tools [3]. All previous works on the PB method have neglected

the fact that the silicon thermal conductivity is temperature dependent and will decrease by around 10-20% with temperature increase of 40-50 degrees. As a result, using the PB method to predict thermal profiles in a chip under high temperature rise will show larger errors (typically 10%, 5-7°C).

In this paper, we improve the PB method, so that it considers the variation in silicon's thermal conductivity, due to temperature change, in a self-consistent manner. We will introduce two methods and call them "Adaptive Power Blurring methods" one and two (APB_I & APB_II). We provide a detailed comparison of the temperature profiles obtained with the APB, the PB and ANSYS Finite Element Modeling software [6]. The remainder of the paper is organized as follows. In section 2 the PB method will be reviewed and the methodology for the two APB methods will be explained. In section 3 the results will be shown and discussed and we will conclude this paper in section 4.

2. PB and APB Methods

Power Blurring is a convolution-based technique. In order to find the temperature profile we need to exploit the IC power map (power dissipation profile) as well as a thermal mask. The latter is basically the impulse response of the system in space domain and is acquired using ANSYS. Then the temperature profile will be obtained from the convolution of the power map with the thermal mask. The idea of using a thermal mask is comparable to Green's function method that has been recently applied to IC thermal analysis [7,8]. However, the Green's function is only available in analytic form for simple geometries, while PB methods consider the real IC chip geometry for estimating the temperature profile [2].

Even though the Power Blurring method was successfully used to compute temperature profiles of IC's with good accuracy and max 1% error [1,2], it is limited to the case where temperature variation in devices is low. For high temperature changes in ICs, the method will have 7-12% errors. This is due to the fact that for large temperature variations we can no longer neglect the temperature dependency of silicon thermal conductivity. Thermal conductivity of silicon versus temperature is shown in Fig. 1 [9]. In order to overcome this limitation of the PB technique, we propose two algorithms as extensions to the power blurring method so that they can adaptively (self-consistently) include the temperature dependency of silicon thermal conductivity in the calculation of the temperature profiles.

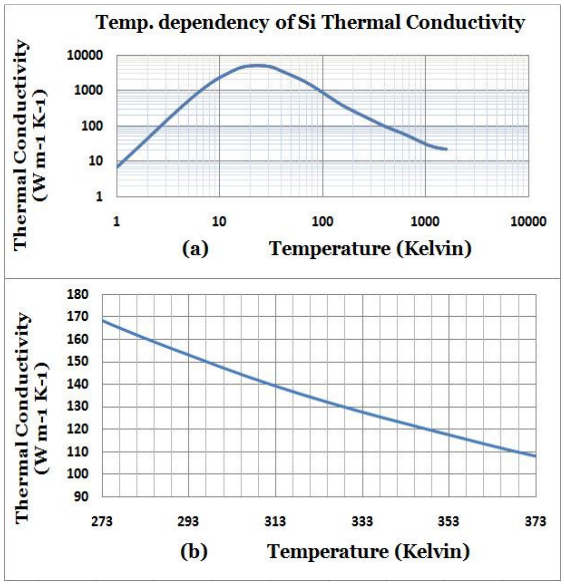

Figure 1. Silicon thermal conductivity vs. temperature for a. 0-1500 K, b. 273-373 K (0-100˚C).

The algorithms for the two Adaptive Power Blurring methods, APB_I and APB_II, are shown in Fig. 2 and Fig. 3, respectively. As it can be seen, in both cases we are using an iterative approach. In order to acquire an initial estimate for the temperature profile of the IC, both methods employ the Power Blurring method using the thermal mask obtained with room temperature thermal conductivity of silicon. The key difference between the two methods manifests itself in the way they analyze the initial temperature profile, and in turn, choose the corresponding thermal mask from a look-up table for the next iterations. In the first APB method (APB_I) the average increase in temperature field in the entire chip is obtained using a weighted averaging of power and preliminary temperature profiles as shown in equation (1).

$$T = \frac{\int\int P(x,y)T(x,y)\,dx\,dy}{\int\int P(x,y)\,dx\,dy} \quad (1)$$

Then, a new thermal mask, based on the average increase in temperature and the according change in

Si thermal conductivity, is selected from the thermal mask's look-up table. The new thermal mask is convolved with the power map to provide a new estimate of the temperature profile. This scheme is then applied iteratively until the temperature profile converges to the final result.

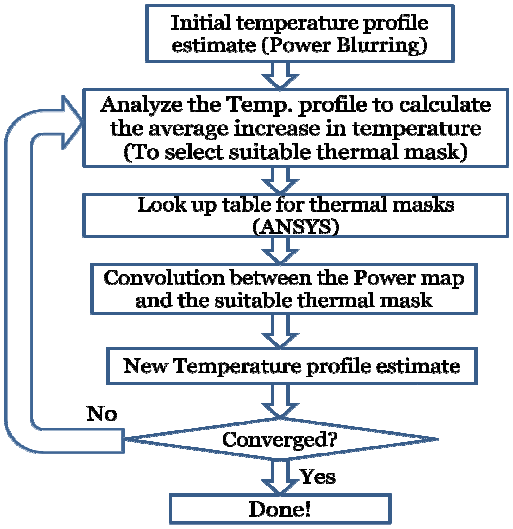

Figure 2. Algorithm for the first Adaptive Power Blurring method (APB_I).

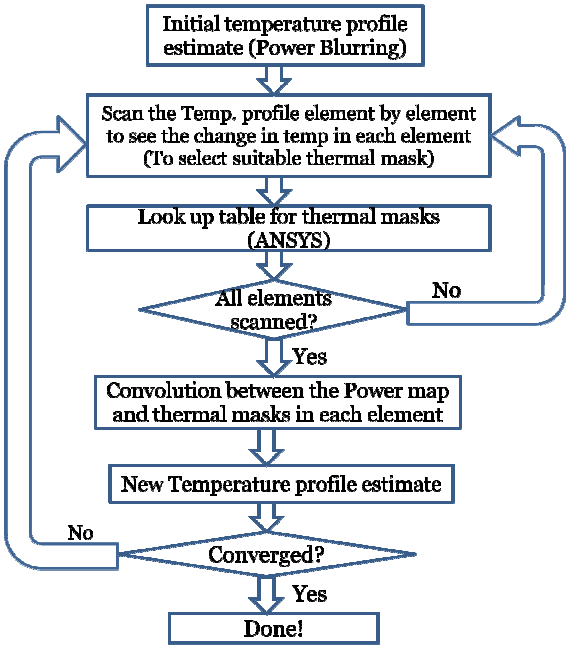

Figure 3. Algorithm for the second Adaptive Power Blurring method (APB_II).

On the other hand, in the second APB method (APB_II) the initial temperature profile is scanned element by element and the local change in the temperature of each element is calculated. Then, based on the increase in temperature in each element and corresponding silicon thermal conductivity, an appropriate thermal mask is chosen. Therefore, for each element on the temperature profile, a new thermal mask is acquired. These thermal masks (impulse responses) are convolved element by element with the power map and yield a new temperature profile estimate. This scheme is again exploited iteratively until the temperature profile converges to the final result. An important point is that for every iterative method, we need an appropriate initial guess, so that the final result could converge correctly. In our iterative methods, we are using the Power Blurring method to give us the initial guess. This by itself already provides a good estimate of the temperature profile, ensuring fast and accurate convergence of the two Adaptive Power Blurring techniques.

3. Results and discussions

3.1. The thermal masks look-up table

The thermal masks look-up table is produced for different thermal conductivities using the FEM software ANSYS. A thermal mask can be obtained by applying a point heat source on a single element at the center of the Si IC. We applied a point heat source with heat flux of 6250 W/cm^2 and set the ambient temperature to 27˚C. Then, by changing the thermal conductivity value based on the data shown in Fig. 1 we obtained different thermal masks for the look-up table.

3.2. Full-chip package model

In order to study PB and APB methods, we used the package model presented in [10]. This flip chip package is shown in Fig. 4. Most of the heat in this type of packages is considered to flow through the bottom surface of the heat sink. The configuration of our flip chip model consisted of a Si IC with a surface area of 1cm × 1cm and a Cu heat sink with a heat spreading layer. Each layer is mounted to the next one by means of a 25µm thick Thermal Interface Material (TIM). The Si IC was orthogonally meshed with element size of 0.025 × 0.025 cm^2 for FEA.

3.3. Case studies

In order to assess the improvement offered by the iterative schemes, we performed the PB and the two APB methods on the power maps shown in Fig. 5. The power distribution shown in Fig. 5(a) ("Edge" power map) was aimed at providing worst case scenario by concentrating all the power on the edges. Fig. 5(b), 5(c), and 5(d) are realistic representations of what a power distribution on a modern-day ASIC

might look like. They aimed at revealing the accuracy of our methods in estimating hot spots.

The results obtained with each of these methods for either of the power maps are compared with the results obtained by ANSYS for the same power map. It should be mentioned that ANSYS also considers the changes in thermal conductivity based on the data shown in Fig. 1.

The temperature profile acquired by ANSYS for the power map shown in Fig. 5(d) is depicted in Fig. 6. As it can be seen in this figure, temperature rises relative to ambient temperature (27˚C) ranges between 20 to 75 degrees. The temperature fields obtained for other power maps have also large variations relative to ambient temperature.

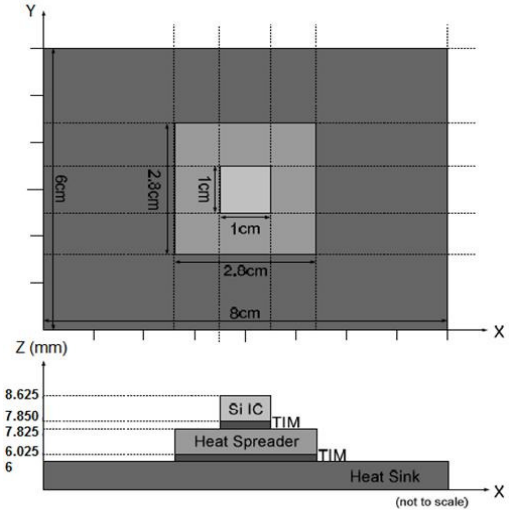

Figure 4. Schematic of packaged IC chip and its dimensions

The average, the maximum error and that in the hottest spot is calculated by comparison between the three PB, APB_I and APB_II methods relative to simulations done by ANSYS. These relative errors are obtained for all four power maps and are tabulated in Table 1. As it can be seen in the table both the APB_I and APB_II are decreasing the maximum, average error and that in hot spots. The APB_II in general provides more accurate results in comparison with the APB_I, because it scans the entire temperature profile and chooses the appropriate thermal mask at each location. Therefore, temperatures at hot spots are calculated using the APB_II with order of 2-20 times better (less error) than the APB_I method and 8-100 times better than the PB method. As an example, for power map shown in Fig. 5(c), the error in computing hot spot temperature relative to ANSYS using APB_II is 0.25%, while this error is 9 times (2.21%) and 25 times (6.4%) larger for the APB_I and PB methods, respectively. On the other hand, the speed of calculation and convergence is the major advantage of the APB_I over the APB_II method. This is due to the fact that the second iterative method (APB_II) needs to scan the entire temperature profile and perform the convolution for each element. By looking at the results obtained for the APB_I method, it can be seen that while the method has error rate close to what is obtained from the APB_II method, it provides calculation speed 5-8 times better than the APB_II. For instance, the execution time for the three PB, APB_I, APB_II methods and the ANSYS simulation performed on the power map shown in Fig. 5(c) were 0.04s, 0.15s, 0.67s, and 19.8s, respectively. In each case for both iterative methods it took just 3-4 iterations to converge. This fast convergence is due to the fact that we used the PB method to obtain the initial guess. In order to stop iterations we selected the point in which the mean square difference between the last two results produced in each of the APB method is less than 10^{-9}.

Fig. 7(a), through 7(d) feature the cross section of the temperature profiles calculated with the PB method, the two APB methods and simulated by ANSYS. As it can be seen both iterative procedures provide results in good agreement with what has been calculated by ANSYS.

4. Conclusion

In this paper, Adaptive Power Blurring methods are presented as an extension to the Power Blurring method to account for the temperature dependency of the silicon thermal conductivity. It has been shown that for large temperature variations in ICs, APB methods can estimate the temperature with much less error than the PB method. In addition, even though these are iterative methods, they converge fast and accurately (in just 3 or 4 iterations). This can be attributed to the fact that the PB method, giving a good estimate of temperature distribution, is used to get the initial guess for our iterative procedure. Finally, we conducted a comparison between two APB methods. The APB_II method, which selects an appropriate thermal mask point by point based on the local temperature, provides a better assessment of temperature particularly at hot spots. Though it runs slower due to the multiple numbers of convolutions, it is still considerably faster than conventional FEM calculations.

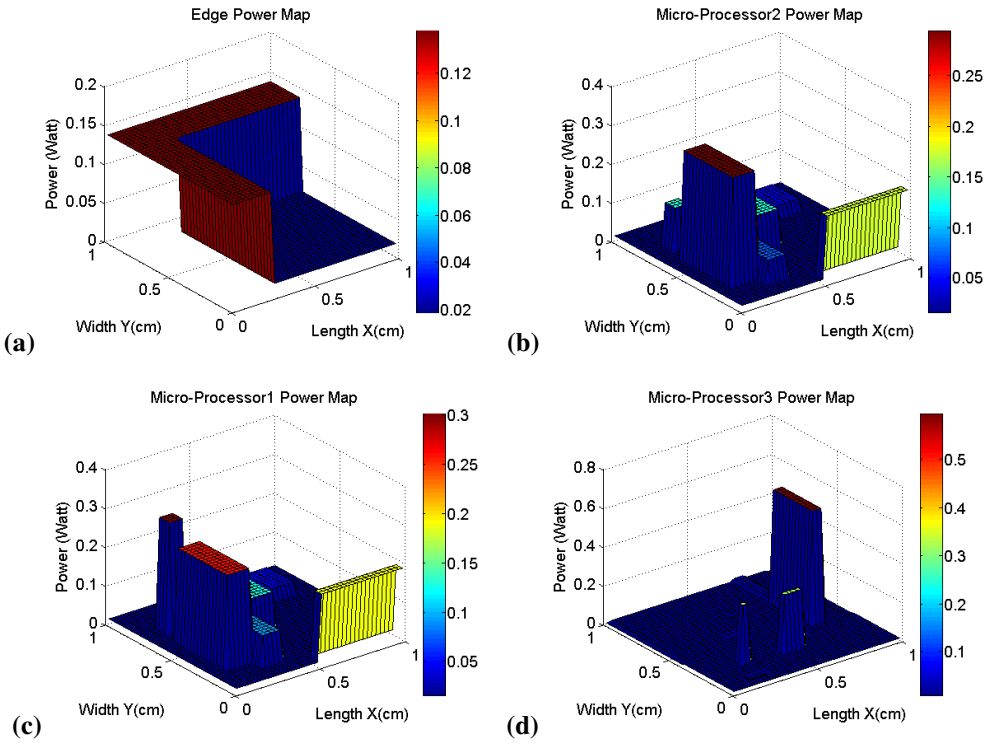

Figure 5. 2D view of the four power dissipation maps selected for our study. (a) "Edge" Power map (b) "µprocessor 1" power map (c)"µprocessor 2" power map, and (d) "µprocessor 3" power map.

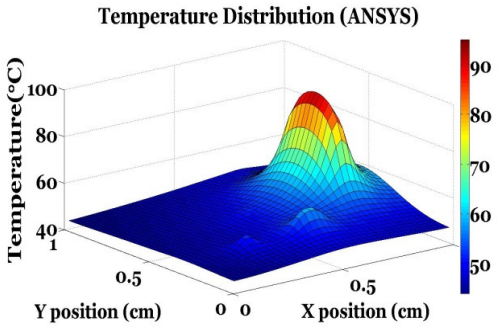

Figure 6. Temperature profile obtained from ANSYS for the power map shown in Fig. 5(d).

5. Acknowledgements

I would like to thank Dr. Xi Wang and Dr. Bjorn Vermeersch for their help and support on this work. This work was supported in part by the Interconnect Focus Center, one of the research centers funded under the Focus Center Research Program, a DARPA and Semiconductor Research Corporation program.

		PB	APB_I	APB_II
"Edge" Power Map	Max. Error	4.54%	3.47%	3.6%
	Average Error	2.36%	0.97%	0.92%
	Hot Spot Error	4.53%	0.77%	0.04%
"µprocessor 1" Power Map	Max. Error	7.32%	3.05%	2.16%
	Average Error	1.54%	0.78%	0.62%
	Hot Spot Error	7.04%	2.46%	0.96%
"µprocessor 2" Power Map	Max. Error	6.69%	2.69%	2.09%
	Average Error	1.55%	0.74%	0.59%
	Hot Spot Error	6.4%	2.21%	0.25%
"µprocessor 3" Power Map	Max. Error	7.9%	1.84%	1.78%
	Average Error	1.14%	0.56%	0.61%
	Hot Spot Error	7.9%	1.83%	1.02%

Table 1. Error rates with respect to ANSYS for the different methods and power maps under study.

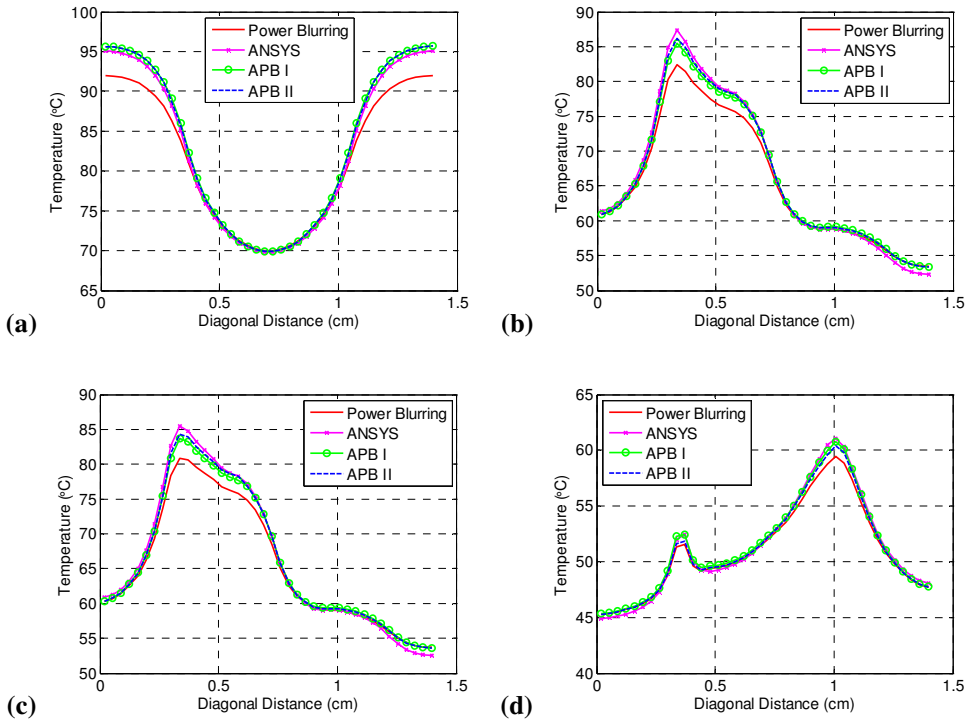

Figure 7. Diagonal cross section of the temperature profiles for different power distributions and algorithms. (a)"edge" power map, (b)"µprocessor 1" power map, (c)"µprocessor 2" power map, (d)"µprocessor 3" power map.

References

[1] V.M. Heriz, J.H. Park, T. Kemper, S.M. Kang, A. Shakouri, "Method of Images for the Fast Calculation of Temperature Distributions in Packaged VLSI Chips." International Workshop on Thermal Investigation of ICs and Systems (Therminic), pp.18-25, 2007

[2] J.H. Park, X. Wang, A. Shakouri, S.M. Kang, "Fast Computation of Temperature Profile of VLSI ICs with High Spatial Resolution" Semiconductor Thermal Measurement, Modeling, and Management Symposium (Semi-Therm 24), pp. 50-55, 2008.

[3] J.H. Park, A. Shakouri, and S.M. Kang, "Fast Evaluation Method for Transient Hot Spots in VLSI ICs in Packages" 9th International Symposium on Quality Electronic Design (ISQED 08), San Jose, pp. 600-603, 2008.

[4] J.H. Park, A. Shakouri, S.M. Kang, "Fast Thermal Analysis of Vertically Integrated Circuits (3D ICs) Using Power Blurring Method," InterPACK'09, San Francisco, July 19-23, 2009.

[5] X. Wang, S. Farsiu, P. Milanfar, A. Shakouri, "Power Trace: An Efficient Method for Extracting the Power Dissipation Profile in an IC Chip From Its Temperature Map" IEEE Transaction on Components and Packaging Technologies, VOL. 32, NO. 2, pp. 309-316, JUNE 2009

[6] ANSYS R11.0, Swanson ANSYS Inc., 2007.

[7] Y.K. Cheng and S.M. Kang, "An Efficient Method for Hot-spot Identification in ULSI Circuits," Proceedings of the IEEE/ACM international Conference on Computer-aided design (ICCAD), San Jose, CA, pp. 124–127, 1999.

[8] Y. Zhan, and S. Sapatnekar, "A High Efficiency Full-Chip Thermal Simulation Algorithm," Proceedings of the IEEE/ACM International conference on Computer-aided design (ICCAD), San Jose, CA, pp. 635-638, 2005.

[9] http://www.efunda.com/materials/elements/TC_Table.cfm?Element_ID=Si

[10] J.H. Park, V.M. Heriz, A. Shakouri, and S.M. Kang, "Ultra Fast Calculation of Temperature Profiles of VLSI ICs in Thermal Packages Considering Parameter Variations", IMAPS 40th International Symposium on Microelectronics, San Jose, Nov. 11-15, 2007.