

HotSpot: A Compact Thermal Modeling Methodology for Early-Stage VLSI Design

Wei Huang, *Student Member, IEEE*, Shougata Ghosh, Siva Velusamy, Karthik Sankaranarayanan, Kevin Skadron, *Senior Member, IEEE*, and Mircea R. Stan, *Senior Member, IEEE*

Abstract—This paper presents *HotSpot*—a modeling methodology for developing compact thermal models based on the popular stacked-layer packaging scheme in modern very large-scale integration systems. In addition to modeling silicon and packaging layers, *HotSpot* includes a high-level on-chip interconnect self-heating power and thermal model such that the thermal impacts on interconnects can also be considered during early design stages. The *HotSpot* compact thermal modeling approach is especially well suited for preregister transfer level (RTL) and presynthesis thermal analysis and is able to provide detailed static and transient temperature information across the die and the package, as it is also computationally efficient.

Index Terms—Compact thermal model, early design stages, interconnect self-heating, temperature, VLSI.

I. INTRODUCTION

AN unfortunate side effect of miniaturization and the continued scaling of CMOS technology is the ever-increasing power densities. The resulting difficulties in managing temperatures, especially local hot spots, have become one of the major challenges for designers at all design levels. High temperatures have several significant impacts on VLSI systems. First, the carrier mobility is degraded at higher temperature, resulting in slower devices. Second, leakage power is escalated due to the exponential increase of subthreshold current with temperature. Third, the interconnect resistivity increases with temperature, leading to worse power-grid IR drops and longer interconnect RC delays, hence causing performance loss and complicating timing and noise analysis. Finally, elevated temperatures can shorten interconnect and device life times and package reliability can be severely affected by local hot spots and higher temperature gradients. For all of these reasons, in order to fully account for the thermal effects, it is important to model temperature for VLSI systems in an accurate but still efficient way. For

Manuscript received November 26, 2004; revised July 19, 2005, and November 13, 2006. This work was supported in part by the National Science Foundation under Grants CCR-0133634 and CCF-0429765, two grants from Intel MRL, and by the University of Virginia through the FEST Excellence Award.

W. Huang and M. R. Stan are with the Charles L. Brown Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA 22904 USA (e-mail: whuang@virginia.edu; mircea@virginia.edu).

S. Ghosh was with the University of Virginia, Charlottesville, VA 22904 USA. He is now with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: shougata@princeton.edu).

S. Velusamy was with the University of Virginia, Charlottesville, VA 22904 USA. He is now with Xilinx Inc., San Jose, CA 95124 USA.

K. Sankaranarayanan and K. Skadron are with the Department of Computer Science, University of Virginia, Charlottesville, VA 22904 USA (e-mail: ks4kk@cs.virginia.edu; skadron@cs.virginia.edu).

Digital Object Identifier 10.1109/TVLSI.2006.876103

example, knowing the across-die temperature distribution at design time permits thermally self-consistent leakage power calculations in an iterative manner, as shown in Fig. 1(a) [1]–[3]. Similarly, an efficient thermal model can also help to close the loop for temperature-aware performance and reliability analysis, as suggested in Fig. 1(b). In particular, it is crucial to take thermal effects into account as early as possible in the design flow, because optimal early and high-level thermally related design decisions can significantly improve design efficiency and reduce design cost.

Obviously, it is impractical to accurately analyze thermal effects and model temperature distribution of a system together with the environment in their full details. Using numerical thermal analysis methods, such as the finite-element method (FEM), is a time-consuming process not suitable for design-time and run-time thermal analysis. In order to gain more insights of the thermal effects during early IC design stages, the tradeoff solution is to build compact thermal models (CTMs) that give reasonably accurate temperature predictions with little computational effort at desired levels of abstraction [5]. Early design stages present unique challenges that we believe require a by-construction compact modeling approach.

Based on the well-known duality between thermal and electrical phenomena,¹ a CTM is a lumped thermal RC network, with heat dissipation modeled as current sources. The resulting thermal RC networks are typically relatively small, can be solved for temperature very efficiently and introduce little computational overhead. Due to this computational efficiency, at pre-register transfer level (RTL) and presynthesis design stages, it is desirable to have compact thermal models for both temperature-aware design and fast simulations of architecture-level dynamic thermal management techniques. Here, temperature-aware design refers to a design methodology that uses temperature as a guideline throughout the design flow. The resulting design can thus be thermally optimized, as it takes into account potential thermal limitations [4].

The major contributions of this work are the following.

- 1) We propose a modeling methodology—*HotSpot*—for generating CTMs that can be used in early VLSI design stages where detailed layout is not available. With this method, reasonably accurate spatial and temporal temperature variations of the silicon die as well as the package can be quickly obtained to help efficient design decisions during

¹In this duality, the heat flow passing through a thermal resistance is analogous to electrical current, and the temperature difference is analogous to voltage. Thermal capacitance, which is based on the material's specific heat and defining the heat absorbing capability, is analogous to electrical capacitance which accumulates electrical charge.

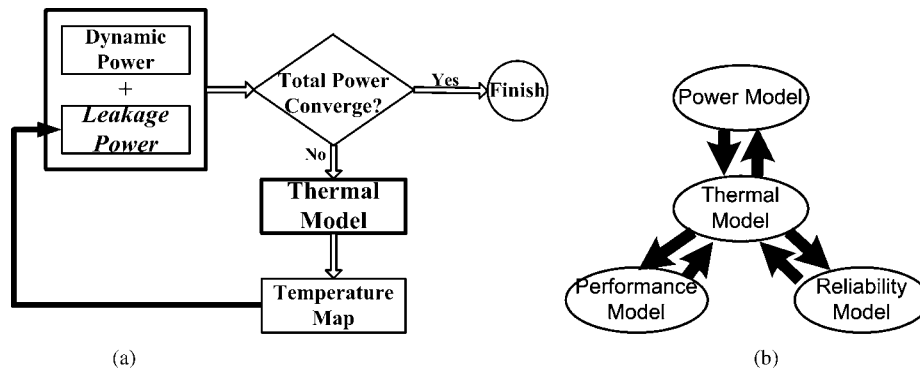


Fig. 1. (a) Thermal model closes the loop for leakage power calculation [3]. (b) The role of temperature in power, performance, and reliability models [4].

early design stages. The modeling method is based on the stacked-layer packaging configuration that is predominant in modern VLSI packaging schemes and is an improvement compared with our previous work [4], [6], [7].

- 2) We analytically investigate the relationship between the number of nodes in the compact thermal model and the accuracy of the model. For thermal analysis during early design stages, it is important to find the right thermal modeling grid density in order to achieve faster computation speed without sacrificing accuracy.
- 3) We also propose a high-level on-chip interconnect self-heating power and temperature model, which can be used for early analysis of thermal impacts on interconnect-related performance, power grid IR drop, and electromigration, again, during early design stages when detailed routing and layout information is not available.

There have been several published efforts in full-chip thermal modeling and compact thermal modeling for microelectronics systems. Wang *et al.* [8] present a detailed and stable die-level transient thermal model based on full-chip layout, solving temperatures for a large number of nodes with an efficient numerical method. The die-level thermal models by Su *et al.* in [9] and Li *et al.* in [10] also provide the detailed temperature distribution across the silicon die and can be solved efficiently, but with no information about the transient behavior. An earlier detailed full-chip thermal model by Cheng *et al.* [11] has an accurate three-dimensional (3-D) model for the silicon and one-dimensional (1-D) model for the package. A significant limitation of the above modeling approaches is the oversimplified thermal package model. For example, the thermal interface material and heat spreader that greatly affect the die temperature distribution are either not included or not properly modeled. The bottom surface of the silicon substrate is treated as isothermal by the above previous works, which significantly deviates from reality and therefore introduces errors. Additionally, these models are not quite suitable for early design stages, since their computation effort is nontrivial while fine-grained thermal analysis is not necessary when detailed layout information is not available. Finally, except for [11], none of these models has shown validation from simulations with detailed numerical models or measurements from real designs.

On the other hand, Lasance *et al.* [12], Sabry [5], and Bosch [13] present package-level compact thermal models extracted

from detailed numerical thermal simulations by data fitting. These models are accurate and have the important property of (quasi-)boundary condition independence (BCI). One limitation of these models is that they have only one or a few junction nodes representing die temperature distributions. Also, because the models are constructed by data fitting, they are not physical (i.e., not derived from design geometries and material properties), and, hence, are not parameterizable. Parametrization is important for CTMs to be used in exploring new design alternatives, where empirical fitting is often not possible [14], [15].

There also have been a number of previous works on thermal modeling of on-chip interconnects and vias. For example, Chen *et al.* [16] present an interconnect thermal model that closely considers thermal coupling phenomenon between nearby interconnects. This model is accurate but it is on a per-interconnect basis and is not extended to model multilevel structure at a higher design abstraction. Chiang *et al.* [17] describe an analytical multilevel interconnect thermal model with considerations of via effects. This model copes with the thermal effect of vias by lumping the heat transferred through the vias into an equivalent thermal conductivity for the inter-layer dielectrics (ILDs). However, to make interconnect thermal analysis complete, self-heating power also needs to be modeled. Unfortunately, to the best of our knowledge, we have not seen any previous works providing models on the multilayer interconnect self-heating power at a higher design level.

The remainder of this paper is organized as follows. Section II describes different aspects of the HotSpot compact thermal modeling approach and is divided into three subsections. Section II-A presents the layered thermal modeling approach in detail, Section II-B addresses the issue of grid density versus accuracy of the thermal model, and Section II-C proposes the high-level interconnect self-heating power and thermal model for early design stages. Following that, Section III shows several validation steps that we have performed for the HotSpot CTMs. Then, in Section IV, we show some example applications of the HotSpot CTMs. Finally, Section V concludes the paper and points out future work.

II. MODELING DETAILS

A compact thermal modeling approach must have several features for it to be useful. First, it should provide detailed tem-




Fig. 2. Stacked layers in a typical ceramic ball grid array (CBGA) package [18]

perature distribution at the desired level of abstraction (e.g., a single node representing the die temperature is unacceptable for thermal modeling at the IC level). In addition, both static and transient thermal behavior should be modeled. Second, a CTM should model just at the needed accuracy and hide the details of lower levels, so that the model itself is no more complex than necessary. Third, the model structure should be kept as simple as possible and should introduce little computational overhead. The HotSpot compact thermal modeling methodology proposed in this paper has all of the above desired features.

A. General Methodology

Most modern VLSI systems have a package consisting of several stacked layers made of different materials, as shown in Fig. 2. This is also the package scheme adopted for the HotSpot thermal models used in this paper. Typical layers include heat sink, heat spreader, thermal paste, silicon substrate, on-chip interconnect layers, C4 pads, ceramic packaging substrate, and solder balls. The recently proposed stacked chip-scale packaging (SCP) [19] and 3-D IC designs [20] are also stacked-layer structures and can be easily modeled as extensions of the generic stack structure in Fig. 2.

When deriving a compact thermal model in HotSpot, the different layers, their positions and adjacency are first identified. Each layer is then divided into a number of blocks. For example, in Fig. 3(c), the silicon substrate layer is divided according to architecture-level units or into regular grid cells, depending on what the die-level design requires. Note that only three blocks are shown in Fig. 3(c) for simplicity. Other layers that greatly affect cross-die temperature distribution (e.g., thermal interface material) can be modeled similarly to the silicon substrate. For the analysis of the needed size of regular grid cells, see Section II-B.

For other layers that require less detailed thermal information (such as heat spreader and heat sink), we simply divide that layer as illustrated in Fig. 3(a). The center shaded part in a layer shown by Fig. 3(a) is the area covered by another adjacent layer such as the one shown in Fig. 3(c). This center part can have the same number of nodes as its smaller neighbor layer or can collapse those nodes into fewer nodes, depending on the accuracy and




Fig. 3. (a) Partitioning of large-area layers (top view). (b) One block with its lateral and vertical thermal resistances (side view). (c) A layer, for example, the silicon die, can be divided into an arbitrary number of blocks if detailed thermal information is needed (top view).

computation speed requirements. The remaining peripheral part in Fig. 3(a) is then divided into four trapezoidal blocks, each assigned to one node.

Every block or grid cell in each layer has one vertical thermal resistance connected to the next layer and several lateral resistances to its neighbors in the same layer. Fig. 3(b) shows a side view of one block with both the lateral and the vertical thermal resistances. The vertical thermal resistance is calculated by $R_{vertical} = t/(k \cdot A)$, where t is the thickness of that layer, k is the thermal conductivity of the material of that layer, and A is the cross-sectional area of the block. We see that each layer is not further divided into multiple thinner layers in the vertical direction, i.e., our modeling method is not fully 3-D. This is a reasonable approximation for early design stages since each layer is relatively thin (a millimeter or less), further discretization in the vertical direction would induce more computation while not improving accuracy significantly.

Calculating lateral thermal resistance is not as straightforward as the vertical resistance. This is because heat spreading or constriction in the lateral directions must be accounted for. Basically, the lateral thermal resistance on one side of a block

can be considered as the spreading/constriction thermal resistance of the neighboring part within a layer to that specific block. Lateral thermal resistances are normally much greater than their vertical counterparts due to the fact that the lateral heat-transfer cross-sectional areas are usually much less than vertical ones. We calculate the spreading/constriction resistance based on the formulas given in [21]. The resistance is a spreading one if the lateral area of the source is smaller than the bulk lateral area, and it is a constriction one otherwise.

For each node, there is also a thermal capacitance $C_{th} = \alpha \cdot c_p \cdot \rho \cdot t \cdot A$, connected to ground, where c_p and ρ are the specific heat and density of the material, respectively. The factor $\alpha \approx 0.5$ is a scaling factor accounting for lumped versus distributed thermal RC time constants.²

Finally, the heatsink-to-air convection thermal resistance can be modeled as $R_{convection} = 1/(h \cdot A)$, where A is the surface area and h is the heat transfer coefficient that is boundary condition dependent. For a first-order approximation, this is adequate for thermal analysis during early design stages. Typical values of h for typical heat sinks under different convection conditions usually be found in the heat sink datasheets.

Details of the derivations and formulas of all the above thermal resistances and capacitances can be found in [7] and [21]. An example of how a HotSpot compact thermal model is assembled can be found in [4].

From the above descriptions, it can be seen that our method can model relatively detailed static and transient temperature variations for the silicon die. In particular, different packaging components can be modeled with more detailed temperature distribution information, which is not available in existing works such as [8]–[10]. Additionally, because the thermal models are built as lumped thermal RC networks, the computational overhead for solving the temperatures is small. Therefore, the modeling method is suitable for developing compact thermal models used during early design stages.

Compact thermal models developed from this method are also parameterizable and BCI. The models are parameterizable because they are built only using the physical geometries and material properties. The models are also BCI because the entire internal RC network is built independent of boundary conditions. More discussions on parametrization and BCI of the HotSpot models can be found in [14] and [15].

B. Thermal Modeling Accuracy Analysis

One important aspect of model development is to decide the proper number of nodes at which temperatures are modeled. At one extreme, a single junction temperature for the chip would be enough for board-level designs, but not enough for die-level designs. Taking the die-level temperature modeling as an example, temperature can be modeled at the functional unit level [see Fig. 4(a)] or the die can be divided into regular grid cells to get more detailed temperature distribution as a function of the size of each grid cell [see Fig. 4(b)]. Ideally, we would prefer

²There should be no surprise that the same factor also appears in the analysis of distributed RC electrical interconnect lines. This approximation is legitimate since the lateral thermal resistances are usually much greater and make negligible contribution to the thermal RC time constants compared with the vertical thermal resistances.




Fig. 4. Modeling at the granularity of (a) functional blocks, (b) uniform grid cells, and (c) hybrid-sized grid cells [2].




Fig. 5. (a) A 1-D slab of material with left half dissipating power. (b) Temperature distribution along the length of the slab.

a hybrid grid scheme that combines both the per-function unit model and the uniform-size grid model as in Fig. 4(c) [2]. By doing this, we could still get detailed thermal information for particular blocks under consideration while saving computation effort by introducing fewer nodes inside other blocks. It is clear that the desired accuracy determines the minimum grid cell size needed, i.e., the temperature difference across one grid cell should be less than a certain percentage of the maximum temperature difference across the die. In what follows, we show an analytical method to derive the proper size of a grid cell.

Let us first start from a simple case. Assume that there is a slab of material with unit width and infinite length. The thickness of the slab is t , and the bottom surface of the slab is isothermal. Half of the slab has a uniform power density of i , while the power density of the other semi-infinite half is 0, as shown in Fig. 5(a). The resulting temperature distribution of the top surface of the slab is approximated in Fig. 5(b). The far end of the left half with power density i has a temperature of T_{max0} , and the far end of the right half with no power dissipated has a temperature of 0, for simplicity. Due to symmetry, the temperature at the boundary between the two halves is $T_{max0}/2$. The temperature at point x can then be derived from the equivalent lumped thermal circuit in Fig. 6(a). The lateral and vertical thermal resistances of an infinitesimal portion of the slab with length of dx are

$$r_1 = \frac{dx}{k \cdot t} \quad \text{and} \quad r_2 = \frac{t}{k \cdot dx} \quad (1)$$

where k is the thermal conductivity of the slab material. Now, the equivalent thermal resistance for the semi-infinite half of the slab should be the same whether or not it includes the first vertical thermal resistance r_2 , i.e.,

$$R_{eq} = r_2 || (r_1 + R_{eq}) \quad \text{and also} \quad R_{eq} = r_1 + (r_2 || R_{eq}). \quad (2)$$




Fig. 6. (a) Thermal resistance network for Fig. 5(a). (b) Thermal circuit at node x for calculating temperature $T(x)$.

Solving the above two equations for R_{eq} leads to

$$\begin{aligned} R_{eq}^2 + r_1 R_{eq} - r_1 r_2 &= 0 \\ R_{eq}^2 - r_1 R_{eq} - r_1 r_2 &= 0. \end{aligned} \quad (3)$$

When $dx \rightarrow 0$, we have $r_1 \rightarrow 0$ and R_{eq} should have a finite value, which means we can neglect the term $r_1 R_{eq}$ in the above two equations of R_{eq} . Therefore, both equations become

$$R_{eq}^2 \approx r_1 r_2 = \frac{1}{k^2}, \quad \text{i.e., } R_{eq} \approx 1/k. \quad (4)$$

Next, to find the temperature $T(x)$, we consider the circuit in Fig. 6(b), in which heat $i(x)$ flows into node x . According to Kirchoff's Current Law, we have

$$\frac{T(x+dx)}{R_{eq}} + \frac{T(x)}{r_2} = i(x) = \frac{T(x)}{R_{eq}}. \quad (5)$$

Substituting R_{eq} with $1/k$, and r_2 with $t/(k \cdot dx)$, and then rearranging both sides of the equation, we obtain

$$\frac{T(x+dx) - T(x)}{dx} = -\frac{T(x)}{t} \Leftrightarrow \frac{dT}{T} = -\frac{1}{t} dx. \quad (6)$$

Taking the integral for both sides from $T_{\max 0}/2$ to $T(x)$ and from 0 to x , respectively, and solving for $T(x)$, we obtain

$$T(x) = \frac{T_{\max 0}}{2} e^{-\frac{x}{t}}, \quad x \geq 0. \quad (7)$$

The above equation shows that the temperature distribution for the right half of the slab is approximately an exponential decay curve with a "spatial" constant of t , which is the thickness from the surface under consideration to the isothermal surface. Furthermore, we can write the temperature distribution of the left half of the slab as a function of position x according to symmetrical nature of the slab structure

$$T(x) = T_{\max 0} - \frac{T_{\max 0}}{2} e^{\frac{x}{t}} \quad (x < 0). \quad (8)$$

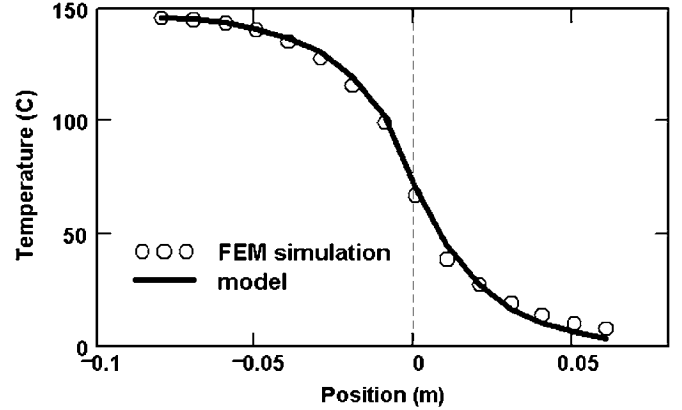


Fig. 7. Comparing FEM simulation result with (7) and (8) for the structure in Fig. 5(a). Power density = 0.5 W/mm^2 , $t \approx 20 \text{ mm}$. Bottom surface of the silicon slab is approximately isothermal.

Fig. 7 confirms the accuracy of the above analysis by comparing with FEM simulations using FloWorks.³

Next, we consider the scenario where heat is dissipated on a finite part of the slab, as shown in Fig. 8(a). The corresponding FEM-simulated temperature distributions within that part of the slab are shown in Fig. 8(b) for different block sizes ($w_2 > w_1$). It is obvious in Fig. 8(b) that, if the size is sufficiently small, the heated part of slab does not actually reach its maximum temperature, as can be seen in Fig. 5(b). This is due to the above mentioned spatial constant, and it means that a block with small size acts as a temperature-spatial low-pass "filter" that prevents the temperature from reaching the maximum possible value. In contrast, with the same power density, a bigger block with its size much larger than the "spatial" constant t can have significantly higher temperature differences. The above analysis explains the "abnormal" observations that although, some tiny structures such as clock buffers in a microprocessor have very high power densities, they do not necessarily cause hot spots, due to this "spatial temperature filtering" effect.

For a particular grid size w , from (8) and Fig. 8(b), the temperature difference within the grid is

$$\Delta T_{\text{grid}} \leq T\left(\frac{-w}{2}\right) - T(0). \quad (9)$$

³FloWorks is an FEM software analyzing computational fluid dynamics and heat transfer. Available at <http://www.nika.biz/index2.htm>




Fig. 8. (a) Part of the slab of material dissipating power. The size of the part is w . (b) FEM simulation results—temperature distribution along the slab with different sizes dissipating power ($w_1 < w_2$). Smaller size “filters” out the temperature difference. (Silicon thickness $t = 20$ mm.)

By setting

$$T\left(\frac{-w}{2}\right) - T(0) = \frac{T_{\max 0}}{2}\left(1 - e^{-\frac{-w}{2t}}\right) = p\% \cdot T_{\max 0} \quad (10)$$

where $p\%$ is the tolerable percentage error, and $T_{\max 0}$ now represents the maximum possible temperature difference across the silicon die.⁴ Solving (10), we get the lower bound for w as

$$w = 2 \cdot t \cdot \ln \frac{1}{1 - 2 \cdot p\%}. \quad (11)$$

Note that t , which is the thickness from the surface under consideration to the isothermal surface, needs to be calculated first. An isothermal surface is an ideal concept that is not found in real packages, but surfaces with negligible temperature differences can be considered as isothermal.⁵ For instance, the thickness t for the silicon surface of the package shown in Fig. 2 can be found by adding up the thickness of silicon substrate, thermal interface material, and heat spreader. An important detail is that, if we use the conductivity of silicon in the above equations, we need to first convert the actual thicknesses of the thermal interface material and the heat spreader to “equivalent” silicon thickness by multiplying their thickness by the ratio of their thermal conductivities to the one for silicon.

Fig. 9 plots the required grid size for different desired levels of precision according to (11), with equivalent thicknesses $t = 4$ mm and $t = 2$ mm. The horizontal axis is the ratio of ΔT_{grid} to $T_{\max 0}$, i.e., p , in percentage. For example, consider that we have a 20 mm \times 20 mm silicon die, and the maximum possible temperature difference across the die is 30°, then, from Fig. 9, we can find that, if we desire all grid temperature error of less than 3% ($30 \times 3\% \approx 1^\circ$) for $t = 4$ mm, a grid size of approximately 0.5 mm is sufficient. This corresponds to dividing the

⁴ $T_{\max 0}$ is dependent on the power distribution across the die and cannot be known *a priori* without performing thermal analysis. However, one can always start with a reasonable guessed value for $T_{\max 0}$ based on previous design experience and then solve the thermal model and iterate the analysis in Section II-B for a few times to get the needed grid size.

⁵One example is the bottom surface of the heat spreader, since the heat spreader is usually made of materials with high thermal conductivity, such as copper.




Fig. 9. Minimum necessary grid size for different desired levels of precision, with $t = 4$ mm and $t = 2$ mm, respectively. The X-axis is the ratio of ΔT_{grid} to $T_{\max 0}$, i.e., p , in percentage. For example, for a system with $t = 4$ mm, if 3% temperature precision is desired, from the solid line, one finds that a grid cell size of 0.5 mm would be enough.

die into 40×40 grid cells. Any finer grid size is unnecessary in this case.⁶

One assumption that we have made so far is that the power is uniform within each grid cell. This assumption is legitimate if the thermal analysis is performed at early design stages, because detailed layout and power information are not available yet. In later design stages, the structures that are included in one grid cell may turn out to be heterogenous. In this case, we can always first resort to finer grid cells inside which power distribution can be considered as uniform, then perform the above accuracy analysis and decide whether or not that finer grid size is necessary or not. Due to the “spatial temperature filtering effect” mentioned above, often we should find that temperature difference within a finer grid cell is negligible and we need to come back to larger grid cells, unless the power density is extremely high.

⁶It is worth noting that the above granularity analysis is based on simplifications of classical heat transfer equations, which underestimates temperature when applied at size scales less than the phonon–phonon mean free path (about 300 nm for silicon at room temperature) [22]. Thus, for granularity analysis at the transistor level, the phonon Boltzmann transport equation (BTE) should be used instead.




Fig. 10. (a) An example of wire-length distribution at 45nm technology node, with three regions (local, semi-global, and global). (b) Metal-layer assignment by calculating number of metal layers needed for each of the three regions. (c) Metal-layer assignment by filling every two metal layers with signal wires, starting from Metal 1 and Metal 2. The example in (c) is superior to that given in (b) by providing more detailed metal-layer assignment information.

C. Interconnect Self-Heating Power and Thermal Modeling

There are two major heat transfer paths inside an IC package [18]—a primary heat transfer path (e.g., silicon substrate, heat spreader, or heat sink) and a secondary heat transfer path (e.g., silicon substrate, on-chip interconnect layers, C4 pads, ceramic packaging substrate, solder balls, or printed circuit board). On the other hand, the secondary heat transfer path usually removes a nonnegligible amount of total generated heat (up to 30%). Neglecting the secondary heat transfer path can lead to inaccurate temperature predictions. In addition, as part of the secondary heat transfer path, the on-chip interconnect layers are of particular interest, because interconnect temperature information allows designers to perform more accurate electromigration, wire delay, and IR drop analysis. Until now, a high-level interconnect self-heating model has been unavailable for early design stages. Most existing interconnect self-heating power and thermal models are either based on analysis of only a few wires [23] or need full-chip detailed layout information that is not available during early design stages [24].

There are two aspects to be considered in the interconnect model: 1) *the average self-heating power* of interconnects in each metal layer and 2) *the equivalent thermal resistance* for metal wires and their surrounding inter-layer dielectric. Vias also play an important role in heat transfer among different metal layers, and, therefore, need to be included as well.

1) *Interconnect Self-heating Power Model*: The self-heating power of a metal wire can be written as

$$P_{\text{self}} = I^2 \cdot R = I^2 \cdot \rho_m \cdot l / A_m \quad (12)$$

where I is the rms current flowing through the wire, $R = \rho_m \cdot l / A_m$ is the electrical resistance, ρ_m is the metal resistivity (which is temperature-dependent), and l and A_m are the length and cross-sectional area of the individual wire, respectively. Because the model needs to *predict* wire temperatures before physical layout is available, first it has to be able to predict the average wire length and the self-heating current (rms current) for wires in each metal layer. It is also important to notice that, because the routing schemes are significantly different for the *signal interconnects* and the *power distribution network*, the methods of predicting average wire length and self-heating current are also different for signal and power supply wires, and, therefore, we treat them separately.

a) *Average interconnect length in each metal layer for signal interconnects*: We predict the average signal interconnect length in each metal layer by adopting and extending the statistical *a priori* wire-length distribution model presented by Davis *et al.* in [25], which improves the wire-length distribution model by Donath [26]. It is important to note that an interconnect thermal model at high levels of abstraction strongly depends on the *a priori* wire-length distribution model and, hence, is limited by the accuracy and efficiency of the wire-length distribution model.

The model in [25] is based on the well-known Rent's Rule: $T = kN^p$, where k and p are Rent's Rule parameters, N is the number of gates in a circuit, and T is the predicted number of I/O terminals in the circuit. If the interested circuit block is of a heterogeneous nature, i.e., there are different Rent's Rule parameters for different subcircuit blocks, then equivalent Rent's Rule parameters can be found using the heterogeneous Rent's Rule proposed by Zarkesh-Ha *et al.* [27].

Three wire-length regions are considered in [25]—local, semi-global, and global. The model predicts the number of wires of any specific length, which is called the interconnect density function $i(l)$, where l is the wire length in gate pitches. Fig. 10(a) shows an example wire-length distribution based on ITRS data [28] for high-performance designs at the 45-nm technology node, where L_{loc} , L_{semi} , and L_{glob} are maximum local, semi-global, and global wire lengths, respectively.

Using the interconnect density function $i(l)$, one can calculate the average length and number of wiring nets for each region. For example, for the semi-global region, we have

$$\begin{aligned} l_{\text{semi}} &= \chi \text{f.o.} \frac{\int_{L_{\text{loc}}}^{L_{\text{semi}}} i(l) \cdot l \, dl}{\int_{L_{\text{loc}}}^{L_{\text{semi}}} i(l) \, dl} \\ n_{\text{semi}} &= \frac{1}{\text{f.o.}} \int_{L_{\text{loc}}}^{L_{\text{semi}}} i(l) \, dl \end{aligned} \quad (13)$$

where χ is the correction factor that converts the point-to-point interconnect length to wiring net length (using a linear net model $\chi = 4/(\text{f.o.} + 3)$), and f.o. is the average number of fan-outs per wiring net. More details can be found in [25].

However, there is no wire-length distribution information regarding *each metal layer* when using this three-region division method in [25]. For the interconnect CTM, we need the wire-




Fig. 11. Scheme to assign signal interconnects to metal layers. M is the number of signal wires between two power rails and Sp is ratio of the space between every two signal wires to average signal wire length of that metal layer.

length distribution predictions of every metal layer. Because of the predominant usage of Manhattan routing, in general, two metal layers are needed to route one wiring net—one layer for horizontal routing, the other for vertical routing. In this paper, we estimate the pair of metal layers where each wiring net is routed by filling every two metal layers with wiring nets, starting from the shortest wiring nets. We thus assume that the shortest wiring nets of the wire-length distribution in Fig. 10(a) are assigned to Metals 1 and 2. Once the first two metal layers are filled, we proceed to Metals 3 and 4, and so on and so forth, until all the wiring nets are assigned to their corresponding pair of metal layers. Although this is an oversimplification, we expect it to be representative of an actual routing strategy. A useful byproduct of our approach is that we are also able to estimate the total number of metal layers needed for a design. As illustrated in Fig. 11, assuming the length of the shortest and longest point-to-point interconnects that can be assigned to a pair of metal layers are L_{\min} and L_{\max} in gate pitches, we can then find the average length and total number of wiring nets within a pair of metal layers by

$$l_{\text{avg}} = \chi_{\text{f.o.}} \frac{\int_{L_{\min}}^{L_{\max}} i(l) \cdot l \, dl}{\int_{L_{\min}}^{L_{\max}} i(l) \, dl}$$

$$n_{\text{total}} = \frac{1}{\text{f.o.}} \int_{L_{\min}}^{L_{\max}} i(l) \, dl. \quad (14)$$

Furthermore, by assuming the routing structure of Fig. 11, where M is the number of signal wires between two power rails and Sp is ratio of the space between every two signal wires to l_{avg} (both M and Sp are design parameters and are tunable by the designer), we get the following relation:

$$n_{\text{total}} \cdot (Sp + 1) \cdot l_{\text{avg}} \cdot \frac{M + 1}{M} \cdot p = 2 \cdot \text{Area} \quad (15)$$

where p is the wire pitch of a metal layer, and $2 \cdot \text{Area}$ is the available routing area for the pair of metal layers under consideration. Using this relationship, and starting at Metals 1 and 2 with $L_{\min} = 1$, we are able to solve for L_{\max} and L_{\min} for each pair of metal layers. An example metal-layer assignment for the interconnect distribution of Fig. 10(a) is shown in Fig. 10(c).

Another way to assign signal wiring nets to different layers is to calculate the number of metal layers needed for each of the three regions, namely, local, semi-global, and global, as in [25]. The resulting metal-layer assignment is shown in Fig. 10(b). As can be seen, the results in Fig. 10(c) and (b) are similar, but Fig. 10(c) provides detailed metal-layer assignment estimations for every two metal layers without considering the three regions, while the information provided in Fig. 10(b) is coarser. Therefore, we prefer the approach used in Fig. 10(c). On the other hand, if the total number of metal layers is fixed, the parameters Sp and M can be adjusted accordingly to fit all of the signal interconnects into the metal layers.

b) Average interconnect length in each metal layer for power and ground: So far, we have considered the average signal interconnect length in each metal layer. We also need to find the average wire length for the power and ground networks, which are usually grid-like. This is relatively simple: we only need to find the length of the power grid section in each metal layer. The assumption here is that the power grid for each metal layer is uniformly distributed, which is a reasonable assumption for early high-level design stages.

With this, we are done with estimating wire length. Next, we need to use this information to estimate interconnect self-heating power.

c) Average interconnect rms self-heating current in each metal layer for signal interconnects: For each switching event, half of the energy drawn from the power supply is dissipated in the form of heat on the charging/discharging transistor and on the output signal interconnect. The average current flow through the interconnect during a switching event can be solved from the following equation:

$$I_{\text{RMS}}^2 (R_{\text{tr}} + R_{\text{wire}}) t_d = \frac{1}{2} \alpha C_L V_{\text{dd}}^2 \quad (16)$$

where I_{RMS} is the self-heating current per wire in each metal layer. R_{tr} is the on-resistance of the transistor, R_{wire} is the wire resistance, α is the switching activity factor, C_L is the load capacitance, and t_d is the delay of the switching event. For long interconnects, repeaters are inserted in order to achieve optimum delay, and these need to be also taken into account. The critical wire-length between repeaters (L_{crit}), the delay for one section of buffered interconnect (τ_{crit}), the optimal number of repeaters ($N r_{\text{crit}}$), and the optimal size of repeaters (s_{crit}) for interconnects in each region can be found using the repeater insertion model proposed in [29]. The calculations of R_{tr} , R_{wire} , C_L , and t_d are different for wires with or without inserted repeaters—the wire length is either the total wiring net length or the length of a wire section between repeaters; the driving and load gates are either gates with average transistor size or repeaters with size of s_{crit} . Finally, the delay of the switching event t_d can be approximated as τ_{crit} for interconnects with

repeaters or as `clock_cycle_time/logic_depth` for interconnects without repeaters.

d) Average interconnect rms self-heating current in each metal layer for power and ground: To calculate average rms currents for power supply grid sections, we can use one of two methods.

The first method is to build a grid-like resistive network model for VDD and GND, somewhat resembling the grid-like die-level thermal model as in Fig. 4(b). Each resistor connecting two nodes in the same metal layer is now the electrical resistance of one power supply grid section. Resistors connecting power grid nodes of different metal layers represent the vias. The topology of the network is obtained by knowing the pitch between power rails in each metal layer, average length, and number of power grid sections between power grids. Next, by applying voltage sources to the top-layer C4 pad sites and adding current loads at Metal-1 endpoints, the resistive network is solved to find the average self-heating current of the power grid in each metal layer.

The other method to calculate average rms self-heating current of power grid section in a metal layer is quite straightforward—we can simply divide the total current delivered to a metal layer by the number of power grid sections. This method is suitable for high-level design stages but is not as accurate as the first method.

e) Total interconnect self-heating power in each metal layer: With all the above information of average interconnect length and rms self-heating current in each layer (for both signal interconnects and power grid sections), we calculate the average self-heating power per interconnect in each metal layer as

$$P_{\text{wire}} = I_{\text{RMS}}^2 \cdot R_{\text{wire}} = I_{\text{RMS}}^2 \cdot \rho_m \frac{l_{\text{wire}}}{A_{\text{wire}}} \quad (17)$$

where A_{wire} and l_{wire} are the cross-sectional area and the average length of signal interconnects or power grid sections in each metal layer, respectively.

Finally, we calculate the self-heating power for each metal layer. For example, we calculate the self-heating power of metal layer i as

$$P_{\text{self}_i} = P_{\text{wire_sig}_i} \cdot n_{\text{sig}_i} + P_{\text{wire_pwr}_i} \cdot n_{\text{pwr}_i} \quad (18)$$

where $P_{\text{wire_sig}_i}$ and $P_{\text{self_pwr}_i}$ are the self-heating power of each individual signal interconnect and power supply wire for metal layer i , respectively. n_{sig_i} and n_{pwr_i} are the number of signal interconnects and power supply sections in Metal i .

So far, we are done with the first aspect of interconnect thermal modeling—self-heating power calculation of metal layers. Next, we need to calculate the equivalent thermal resistance of wires and the surrounding dielectric, together with the thermal resistance of vias.

2) Equivalent Thermal Resistance of Wires/Dielectric and Vias: In order to derive a model, we consider the case in Fig. 12, where two wires (Wire1 and Wire2) are adjacent to each other. On top of and beneath them are orthogonal wires in neighboring metal layers. All wires are surrounded by ILDs. We want to find the equivalent thermal resistance (R_0) from Wire1 to $d/2$




Fig. 12. Interconnect structures for calculating equivalent thermal resistance of wires with surrounding dielectric.

above Wire1, where d is the thickness of the ILD between two metal layers. The other half of d belongs to the metal layer above Wire1 and is considered when calculating equivalent thermal resistance for wires in that layer. Since we have assumed that all of the wires in the same metal layer are the same, Wire1 and Wire2 are two identical wires dissipating the same power at the same time. Consequently, Wire1 and Wire2 also have the same temperature. We approximate the isothermal surface by the outer dashed area in Fig. 12. This isothermal surface is used for the calculation of R_0 and is $d/2$ away from the wires. Also, it does not overlap with similar isothermal surfaces for the perpendicular wires in neighboring layers. The effective heat-conducting angle which is used for the calculation of R_0 can be approximated by $\theta = 2 \cdot \arctan(D/(d+H))$, as shown in the figure.

There is also a lateral thermal resistance between Wire1 and Wire2— R_{lat} . However, because Wire1 and Wire2 are identical and have the same temperature, there is no heat transfer in the lateral direction and R_{lat} can be removed.

For the calculation of R_0 , we first calculate the thermal resistance of the dark slice of ILD shown in Fig. 12, which can be written in the form of the integral

$$dR_0 = \int_0^{d/2} \frac{1}{k_{\text{ins}} (r+x) d\phi \cdot l} dx = \frac{1}{k_{\text{ins}} \cdot l \cdot d\phi} \ln\left(\frac{d+2r}{2r}\right) \quad (19)$$

where x is the integral variable, k_{ins} is the thermal conductivity of ILD, ϕ is the angle of the slice, $r = \sqrt{WH/\pi}$ is the equivalent radius of the wire, and l is the length of the wire.

If we define thermal conductance G_0 as the reciprocal of thermal resistance R_0 , we have

$$dG_0 = \frac{k_{\text{ins}} \cdot l \cdot d\phi}{\ln\left(\frac{d+2r}{2r}\right)} \Rightarrow G_0 = \int_0^\theta dG_0 = \frac{\theta \cdot k_{\text{ins}} \cdot l}{\ln\left(\frac{d+2r}{2r}\right)} \quad (20)$$

and thus the total equivalent thermal resistance is

$$R_0 = \frac{1}{G_0} = \ln\left(\frac{d+2r}{2r}\right) / (\theta \cdot k_{\text{ins}} \cdot l). \quad (21)$$

Inter-layer heat transfer also happens through vias. A simplistic approximation of the number of vias for signal interconnect is to assume that each wiring net has two vias, one connected to




Fig. 13. Estimating the number of vias for signal interconnects. A wiring net with fan-out 3 is shown in this figure. The number of vias is $(2 \cdot \text{f.o.} + 2)$.




Fig. 14. Estimating the number of vias for power supply wires. An array of vias are put in the intersection of power wires at two metal layers. W_{wire} and W_{via} are the widths of the power wire and the via, respectively.

the upper metal layer, and another one connected to the lower metal layer. A more accurate approximation is to assume that each wiring net has $(2 \cdot \text{f.o.} + 2)$ vias, where f.o. is the average fan out number of each gate. As illustrated in Fig. 13, $(\text{f.o.} + 2)$ vias are at the ends of the wiring net and connecting the wiring net to lower metal layers and eventually to the device layer at the silicon surface. The other vias are used to aid the routing of the wiring net between the pair of metal layers in which the wiring net resides. For the power supply grid, in order to increase the reliability and because the wires are typically wider than minimum size, designers usually use multiple vias at the intersection of two power rails between different metal layers. As illustrated in Fig. 14, the number of vias at an intersection of power rails can be estimated by $(1/4)((W_{\text{wire}}/W_{\text{via}}) - 1)^2$, where W_{wire} and W_{via} are the widths of the power wire and the via, respectively. The thermal resistance of each via is approximately calculated as $R_{\text{via}} = t_v / (k_v A_v)$, where k_v is thermal conductivity of via-filling material, and t_v and A_v are the thickness and cross-sectional area of the via.

All thermal resistors of wires and vias between two metal layers can be considered parallel to each other. Thus, combining all of the thermal resistors between two metal layers, we obtain the total equivalent thermal resistance between two metal layers.

Now, we are almost done with the interconnect thermal modeling. One last step is to stack the thermal resistances for each layer to construct the whole thermal circuit for all interconnect layers. Thermal capacitances can also be calculated for each metal layer and the ILD based on dimensions and material properties using an equation similar to the one in Section II-A.

3) *Interconnect Power and Thermal Model at Different Granularities:* Although the above interconnect thermal modeling approach was presented at the entire die level, in principle, it is also applicable at other granularities. For example, Rent's Rule can also be applied at the functional unit level to estimate intra- and inter-functional-unit wire-length distribution for metal layers above each functional unit. The total self-heating power and the equivalent metal-layer thermal resistance can

then be calculated for each functional unit using similar methods as described before. As part of our future work, the power and temperature estimations for each metal layer at the functional unit level or other abstraction levels will be investigated.

4) *Accuracy Concerns About the Interconnect Power and Thermal Model:* From the above descriptions of the proposed interconnect power and thermal model for early design stages, one might raise some concerns about accuracy. Here we address them one by one.

- 1) Usefulness of the interconnect model—Because interconnect layers usually have much higher absolute temperatures and greater temperature differences than silicon, reliability issues such as thermo-mechanical stress between metal layers, thermo-electromigration of long wires, are increasingly more important. If reasonably accurate early-stage wire-temperature estimations are available, they will be helpful for the designers to discover and deal with such thermally related reliability hazards early in the design flow, hence greatly expediting the design convergence process. For example, the architect needs a way to reason about the thermal and reliability properties among different architectural choices at the pre-RTL architecture determination stage. These kinds of choices do not necessarily need high degrees of precision, it is enough only to know what combination of choices might be problematic.
- 2) Accuracy concern of Rent's Rule—For a mature circuit design style of a specific functional unit along a micro-processor family, Rent's Rule parameters derived from ancestor designs can be used to predict future designs' wire-length distributions with good accuracy, as indicated by Rent's Rule validation data presented in previous works on both traditional and improved Rent's Rules, such as [25], [27], [30], [31], [32]. Rent's Rule is indeed inaccurate for any *individual* wires, but is quite accurate about aggregate average wire behavior for mature circuit design styles. This is also true about other applications of Rent's Rule.
- 3) Concern about current loading accuracy—We think that reasonably accurate average/rms current estimations for typical signal wires are achievable as presented earlier in this section. This is because power estimations at this level (dynamic power with switching factors and static power) are available from tools such as Watch [33]. Average current loading in the power/ground network can also be roughly estimated by solving a coarse VDD/GND mesh (which is similar to the regular-grid-cell thermal resistive network in HotSpot) without loss of much accuracy. It is also obvious that this kind of approach is consistent with the needs of pre-RTL architectural modeling.

D. Computation Speed of Hotspot Compact Thermal Models

The computation speed of HotSpot thermal models to obtain steady-state and transient solutions for several different simulated time intervals at different granularities are in the order of milliseconds to minutes, depending upon the number of blocks/grids, number of material layers, and the simulated transient

TABLE I
COMPUTATION SPEED OF A HOTSPOT MODEL, RUNNING ON A DUAL-PROCESSOR (AMD MP 1.5 GHZ) SYSTEM. (CONVERGING METHOD FOR TRANSIENT SOLUTIONS IS DIFFERENT FROM THAT FOR STEADY-STATE SOLUTIONS)

simulated time interval	CPU time
0.1 ms	20 ms
20 ms	100 ms
2 sec	7.9 sec
20 sec	78.5 sec
steady-state	14.9 sec

time interval. Table I shows the CPU time used to simulate a HotSpot model with 40×40 grid cells.

The small overhead is due to the relatively small and manageable number of nodes in the lumped thermal *RC* circuit, together with the use of first-order difference equations to iteratively solve the *RC* network. The computational efficiency of HotSpot models means there is little computation overhead for existing design methodologies to incorporate the compact thermal models for temperature-aware design or dynamic thermal management simulations.

III. VALIDATION

The HotSpot modeling approach was first validated through detailed FEMs in FloWorks [34]. It was also validated by comparing with real temperature measurements from a commercial thermal testing chip [4]. Validation of the interconnect thermal model also can be found in [4]. Here, we present another step that we have taken recently to further validate HotSpot models.

We designed an field-programmable gate array (FPGA)-based system that monitors the temperature at various locations on a Xilinx Virtex-2 Pro FPGA [35]. The system is composed of a controller interfacing to an array of temperature sensors that are implemented on the FPGA fabric. We use ring oscillators as temperature sensors by exploiting the fact that the frequency of oscillation is approximately proportional to temperature [36]. Calibrations are done for six different sensors placed near the center of each unit on the die. Power consumption for different units is extracted through various methods. Using the floorplan shown in Fig. 15, we compare the sensor readings with values obtained from the corresponding HotSpot model. The results are in Table II. We see that, on average, the temperatures predicted by the HotSpot thermal model and those obtained from the sensors differs by less than 0.2°C .

The low temperature rise (4.1°C maximum) and small temperature difference across the FPGA chip (0.7°C maximum) are due to the fact that typical operating powers for the PPC and MB blocks on the FPGA are not significant enough to heat up the chip (because of this, the FPGA chip is not equipped with a heatsink). In order to achieve greater across-die temperature differences, we have intentionally left two “zero-power” blank blocks (blank1 and blank2). Regardless of the relatively cool die temperature, the errors between the HotSpot model and the thermal sensor measurements are within 10% of the measured temperatures, for example, for “MB” in Table II, the percentage error is $(4.1 - 3.96)/4.1 = 3.4\%$. This confirms the validity of the HotSpot model, although the FPGA application itself does




Fig. 15. Floorplan with six functional blocks implemented in an FPGA for HotSpot primary heat transfer path model validation.

TABLE II
COMPARISONS OF TEMPERATURE READINGS FROM THE FPGA AND THE HOTSPOT THERMAL MODEL. TEMPERATURES ARE WITH RESPECT TO AMBIENT TEMPERATURE. ERRORS ARE WITHIN 0.2°C

Unit	Power(mW)	Sensor Temperature	HotSpot Temperature
blank1	0.1	3.4	3.37
left_ppc	75	3.5	3.69
bott_ppc	75	3.4	3.67
ppc	45	3.5	3.66
mb	313	4.1	3.96
blank2	0.1	3.4	3.38

not show much interesting “hot” temperatures and temperature gradients.

IV. HOTSPOT APPLICATIONS

The HotSpot thermal models can be utilized to achieve accurate preliminary design estimations and precise run-time thermal management techniques. As an example, die-level temperature estimations from HotSpot can be used as a guideline for temperature-aware design during the entire design flow [4]. Other example applications are: HotSpot compact thermal models have been used to close the loop of leakage power calculation [3], explore different architecture-level run-time dynamic thermal management (DTM) techniques [34], [37], aid the analysis of state-of-the-art computer architectures [38], and perform temperature-aware electromigration (EM) analysis for more accurate interconnect lifetime predictions [39].

Apart from the above published applications of HotSpot, here we show the importance of modeling packaging components (e.g., thermal interface material (TIM), heat spreader and heat sink) in greater detail by comparing across-die temperature difference for different TIM thicknesses. TIM is a thin layer of material that glues the silicon die to the heat spreader, made of material with a much lower thermal conductivity than silicon and metals. Therefore, this thin layer of TIM plays an important role in preventing effective heat spreading and resulting in higher temperature gradient at the bottom of the die. Table III shows the across-die temperature difference from a HotSpot thermal model built for a POWER4-like microprocessor with

TABLE III
IMPACT OF TIM THICKNESS ON ACROSS-DIE TEMPERATURE DIFFERENCE

TIM thickness (μm)	across-die T difference ($^{\circ}\text{C}$)
0	12.3
10	18.5
20	24.1
30	31.8

a thermal package similar to that in Fig. 2. Power dissipation for each on-chip functional unit is estimated from IBM's cycle accurate Turandot performance simulator [40] and PowerTimer power modeling tool [41]. Here, we can see that TIM thickness significantly affects the temperature difference across the die. Therefore, it is inaccurate to simply model the bottom surface of the silicon as isothermal, as many of the previous researchers have done.

V. CONCLUSION AND FUTURE WORK

In this paper, HotSpot, a generic by-construction compact thermal modeling methodology for VLSI systems, has been presented. HotSpot models have been shown to be efficient for early design stages and run-time thermal analysis. We also analyzed the needed grid density for a desired precision of temperature estimations. As part of the modeling process, we have also presented a high-level interconnect self-heating power and thermal model to estimate the temperatures of interconnect layers for early design stages.

As topics of future work, the HotSpot modeling approach can be further extended to model 3-D integrated circuits, multichip modules, and active and liquid-cooling structures. Regarding practical applications, we plan to incorporate HotSpot models into real VLSI designs for thermally self-consistent leakage power calculations, power grid IR drop analysis, and interconnect/device/package lifetime analysis to further demonstrate the benefits of using the HotSpot modeling approach.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their helpful suggestions.

REFERENCES

- [1] K. Banerjee, S. C. Lin, A. Keshavarzi, and V. De, "A self-consistent junction temperature estimation methodology for nanometer scale ICs with implications for performance and thermal management," in *Proc. Int. Electron Device Meeting (IEDM)*, 2003, pp. 36.7.1–36.7.4.
- [2] L. He, W. Liao, and M. R. Stan, "System level leakage reduction considering the interdependence of temperature and leakage," in *Proc. 41st Design Automation Conf.*, Jun. 2004, pp. 12–17.
- [3] W. Huang, E. Humenay, K. Skadron, and M. Stan, "The need for a full-chip and package thermal model for thermally optimized IC designs," in *Proc. Int. Symp. Low Power Electron. Design*, Aug. 2005, pp. 245–250.
- [4] W. Huang, M. R. Stan, K. Skadron, K. Sankaranarayanan, S. Ghosh, and S. Velusamy, "Compact thermal modeling for temperature-aware design," in *Proc. 41st Design Automation Conf.*, Jun. 2004, pp. 878–883.
- [5] M.-N. Sabry, "Compact thermal models for electronic systems," *IEEE Trans. Compon. Packaging Technol.*, vol. 26, no. 1, pp. 179–185, Mar. 2003.
- [6] M. R. Stan, K. Skadron, M. Barcella, W. Huang, K. Sankaranarayanan, and S. Velusamy, "Hotspot: A dynamic compact thermal model at the processor-architecture level," *Microelectron. J.*, vol. 34, pp. 1153–1165, 2003.
- [7] K. Skadron, K. Sankaranarayanan, S. Velusamy, D. Tarjan, M. R. Stan, and W. Huang, "Temperature-aware microarchitecture: Modeling and implementation," *ACM Trans. Architecture Code Optimization*, vol. 1, no. 1, pp. 94–125, Mar. 2004.
- [8] T.-Y. Wang and C. C.-P. Chen, "3-D thermal-ADI: A linear-time chip level transient thermal simulator," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 21, no. 12, pp. 1434–1445, Dec. 2002.
- [9] H. Su, F. Liu, A. Devgan, E. Acar, and S. Nassif, "Full chip estimation considering power, supply and temperature variations," in *Proc. Int. Symp. Low Power Electron. Design*, Aug. 2003, pp. 78–83.
- [10] P. Li, L. Pileggi, M. Asheghi, and R. Chandra, "Efficient full-chip thermal modeling and analysis," in *Proc. Int. Conf. Comput.-Aided Design*, 2004, pp. 319–326.
- [11] Y. Cheng, P. Raha, C. Teng, E. Rosenbaum, and S. Kang, "ILLIADS-T: An electrothermal timing simulator for temperature-sensitive reliability diagnosis of CMOS VLSI chips," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 17, no. 8, pp. 668–681, Aug. 1998.
- [12] C. J. M. Lasance, "Two benchmarks to facilitate the study of compact thermal modeling phenomena," *IEEE Trans. Compon. Packag. Technol.*, vol. 24, no. 4, pp. 559–565, Dec. 2001.
- [13] E. G. T. Bosch, "Thermal compact models: An alternative approach," *IEEE Trans. Compon. Packaging Technol.*, vol. 26, no. 1, pp. 173–178, Mar. 2003.
- [14] W. Huang, M. R. Stan, and K. Skadron, "Physically-based compact thermal modeling—achieving parametrization and boundary condition independence," in *Proc. 10th Int. Workshop THERMAL Investigations of ICs Syst.*, Oct. 2004, pp. 287–292.
- [15] —, "Parameterized physical compact thermal modeling," *IEEE Trans. Compon. Packaging Technol.*, vol. 28, no. 4, pp. 615–622, Dec. 2005.
- [16] D. Chen, E. Li, E. Rosenbaum, and S. Kang, "Interconnect thermal modeling for accurate simulation of circuit timing and reliability," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 19, no. 2, pp. 197–205, Feb. 2000.
- [17] T. Y. Chiang, K. Banerjee, and K. Saraswat, "Analytical thermal model for multilevel vlsi interconnects incorporating via effect," *IEEE Electron Device Lett.*, vol. 23, no. 1, pp. 31–33, Jan. 2002.
- [18] J. Parry, H. Rosten, and G. B. Kromann, "The development of component-level thermal compact models of a C4/CBGA interconnect technology: The motorola PowerPC 603 and PowerPC 604 RISC microprocessors," *IEEE Trans. Compon. Packag. Manuf. Technol. A*, vol. 21, no. 1, pp. 104–112, Mar. 1998.
- [19] *Intel's stacked chip scale packaging*, [Online]. Available: <http://www.intel.com/research/silicon/mobilepackaging.htm>
- [20] K. Banerjee, S. J. Souri, P. Kapur, and K. C. Saraswat, "3-D ICs: A novel chip design for improving deep-submicrometer interconnect performance and systems-on-chip integration," *Proc. IEEE*, vol. 89, no. 5, pp. 602–633, May 2001.
- [21] S. Lee, S. Song, V. Au, and K. Moran, "Constricting/spreading resistance model for electronics packaging," in *Proc. American-Japan Thermal Eng. Conf.*, Mar. 1995, pp. 199–206.
- [22] E. Pop, K. Banerjee, P. Sverdrup, R. Dutton, and K. Goodson, "Localized heating effects and scaling of sub-0.18 micron CMOS devices," in *Proc. Int. Electron Devices Meeting*, 2001, pp. 677–680.
- [23] K. Banerjee and A. Mehrotra, "Global (interconnect) warming," *IEEE Circuits Devices Mag.*, vol. 17, no. 5, pp. 16–32, Sept. 2001.
- [24] T.-Y. Wang, J.-L. Tsai, and C. C.-P. Chen, "Thermal and power integrity based power/ground networks optimization," in *Proc. Design Autom. Testing of Europe (DATE)*, Feb. 2004, pp. 830–835.
- [25] J. A. Davis, V. K. De, and J. D. Meindl, "A stochastic wire-length distribution for gigascale integration (GSI)—part I: Derivation and validation," *IEEE Trans. Electron Devices*, vol. 45, no. 3, pp. 580–589, Mar. 1998.
- [26] W. E. Donath, "Wire-length distribution for placements of computer logic," *IBM J. Res. Develop.*, vol. 2, no. 3, pp. 152–155, May 1981.
- [27] P. Zarkesh-Ha, J. Davis, and J. Meindl, "Prediction of net-length distribution for global interconnects in a heterogeneous system-on-a-chip," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 8, no. 6, pp. 649–659, Jun. 2000.
- [28] *The International Technology Roadmap for Semiconductors (ITRS)*, 2003.
- [29] R. H. J. M. Otten and R. K. Brayton, "Planning for performance," in *Proc. Design Autom. Conf. (DAC)*, Jun. 1998, pp. 122–127.
- [30] P. Christie and D. Stroobandt, "The interpretation and application of rent's rule," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 8, no. 6, pp. 639–648, Jun. 2000.

- [31] J. Dambre, P. Verplaetse, D. Stroobandt, and J. V. Campenhout, "A comparison of various terminal-gate relationships for interconnect prediction in vlsi circuits," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 11, no. 1, pp. 24–34, Jan. 2003.
- [32] Y. Cao, C. Hu, X. Huang, A. B. Kahng, I. L. Markov, M. Oliver, D. Stroobandt, and D. Sylvester, "Improved a priori interconnect predictions and technology extrapolation in the GTX system," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 11, no. 1, pp. 3–14, Jan. 2003.
- [33] D. Brooks, V. Tiwari, and M. Martonosi, "Wattch: A framework for architectural-level power analysis and optimizations," in *Proc. Int. Symp. Comput. Architecture*, Jun. 2000, pp. 83–94.
- [34] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan, "Temperature-aware microarchitecture," in *Proc. Int. Symp. Comput. Architecture*, Jun. 2003, pp. 2–13.
- [35] *Xilinx virtex-2 pro user guide*. [Online]. Available: <http://direct.xilinx.com/bvdocs/publications/ds083.pdf>.
- [36] S. Lopez-Buedo and E. Boemo, "Making visible the thermal behavior of embedded microprocessors on FPGAs. a progress report," in *Proc. FPGA*, Feb. 2004, pp. 79–86.
- [37] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan, "Temperature-aware computer systems: Opportunities and challenges," *IEEE Micro*, vol. 23, no. 6, pp. 52–61, 2003.
- [38] Y. Li, D. Brooks, Z. Hu, and K. Skadron, "Performance, energy, and thermal considerations for SMT and CMP architectures," in *Proc. High Performance Comput. Architecture*, Feb. 2005, pp. 71–82.
- [39] Z. Lu, W. Huang, J. C. Lach, M. R. Stan, and K. Skadron, "Interconnect lifetime prediction under dynamic stress for reliability-aware design," in *Proc. Int. Conf. Comput.-Aided Design*, Nov. 2004, pp. 327–334.
- [40] M. Moudgill, J. D. Wellman, and J. H. Moreno, "Environment for PowerPC microarchitecture exploration," *IEEE Micro*, vol. 19, no. 3, pp. 15–25, 1999.
- [41] D. Brooks, P. Bose, V. Srinivasan, M. Gschwind, P. G. Emma, and M. G. Rosenfield, "New methodology for early-stage microarchitecture-level power-performance analysis of microprocessors," *IBM J. Res. Devel.*, vol. 47, no. 5/6, pp. 653–670, 2003.



Siva Velusamy received the B.S. degree from Anna University, Chennai, India, and the M.S. degree from the University of Virginia, Charlottesville, both in computer science.

He is a Software Engineer with Xilinx, Inc., San Jose, CA. His research interests include embedded systems, computer architecture, and operating systems.



Karthik Sankaranarayanan received the B.E. degree in computer science and engineering from Anna University, Chennai, India, in 2000, and the M.S. degree from the University of Virginia, Charlottesville, in 2003. He is currently working toward the Ph.D. degree at the University of Virginia.

He is a member of the LAVA Laboratory, University of Virginia, and his area of interests include computer architecture in general and thermal and power-aware microarchitectures in particular.



Kevin Skadron (SM'05) received the B.S.E.E. and the B.S. degree in economics from Rice University, Houston, TX, and the M.A. and Ph.D. degrees from Princeton University, Princeton, NJ.

He joined the Department of Computer Science, University of Virginia, Charlottesville, in 1999 and is now an Associate Professor. His research interests focus on the implications of technology trends and physical constraints for computer architecture.

Dr. Skadron is a member of the Association for Computing Machinery. He was recipient of three Best Paper Awards, National Science Foundation ITR and CAREER Awards, co-chaired MICRO 2004 and PACT 2002, has served on numerous program committees, and is founding Associate Editor-in-Chief of *IEEE Computer Architecture Letters*.



Wei Huang (S'01) received the B.S. degree in electrical engineering from the University of Science and Technology of China, Hefei, China. He is now working toward Ph.D. degree in electrical and computer engineering at the University of Virginia, Charlottesville.

His research interests include low-power circuit design, modeling and analysis of temperature and other parameter variations at the circuit and microarchitecture levels, and the design of variation-aware VLSI circuits and systems. He was an intern with

the IBM T. J. Watson Research Center, in the VLSI Design Department during the summer of 2004.

Mr. Huang was the recipient of the Best Student Paper Award at ISCA 2003.



Shougata Ghosh received the B.S. degree in electrical and computer engineering from the University of Virginia, Charlottesville. He is currently working toward the Ph.D. degree in computer engineering at Princeton University, Princeton, NJ.

His research interests are in low-power computing and thermal aspects of microprocessors. He is also interested in in-network cache coherence architectures in CMPs.



Mircea R. Stan (SM'98) received the Diploma in electronics and communications from "Politehnica" University, Bucharest, Romania, and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Massachusetts at Amherst.

Since 1996, he has been with the Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, where he is now an Associate Professor. He is teaching and doing research in the areas of high-performance and low-power VLSI, temperature-aware circuits and architecture,

embedded systems, and nanoelectronics. He has more than eight years of industrial experience, was a visiting scholar at the University of California, Berkeley, in 2004–2005, a Visiting Faculty Member with IBM in 2000, and with Intel in 2002 and 1999.

Dr. Stan is a member of the Association for Computing Machinery, Usenix, Eta Kappa Nu, Phi Kappa Phi, and Sixma Xi. He was a Distinguished Lecturer for the IEEE Circuits and Systems Society for 2004–2005. He was the recipient of the National Science Foundation CAREER Award in 1997 and was a coauthor on Best Paper Awards at ISCA 2003 and SHAMAN 2002. He was a program chair for ISLPED 2005, the general chair for GLSVLSI 2003, has been on technical committees for several conferences, has been a Guest Editor for the IEEE COMPUTER Special Issue on Power-Aware Computing in December 2003 and an Associate Editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS since 2004 and for the IEEE TRANSACTIONS ON VERY LARGE-SCALE INTEGRATION (VLSI) SYSTEMS in 2001–2003.