

# Parameterized Physical Compact Thermal Modeling

Wei Huang, *Student Member, IEEE*, Mircea R. Stan, *Senior Member, IEEE*, and Kevin Skadron, *Senior Member, IEEE*

**Abstract**—This paper presents a compact thermal modeling (CTM) approach, which is fully parameterized according to design geometries and material physical properties. While most compact modeling approaches facilitate thermal characterization of existing package designs, our method is better suited for preliminary exploration of the design space at both the silicon level and the package level. We show that our modeling method achieves reasonable boundary condition independence (BCI) by comparing a CTM example with a BCI model for a benchmark ball grid array single-chip package under the same standard set of boundary conditions. In essence, the presented CTM method can act as a convenient medium for enhanced interactions and collaborations among designers at the package, circuit and computer architecture levels, leading to efficient early evaluations of different thermally-related design trade-offs at all the above levels of abstraction before the actual detailed design is available. The presented modeling method can be easily extended to model emerging packaging schemes such as stacked chip-scale packaging and three-dimensional integration.

**Index Terms**—Boundary condition independence (BCI), compact thermal model (CTM), package design, parametrization, temperature.

## I. INTRODUCTION

**A**LONG with the continued scaling of complementary metal oxide semiconductors (CMOS) very large scale integration (VLSI) technologies, the ever-increasing power density and the resultant difficulties in managing temperatures have become one of the major challenges for system designers at all design levels. It is well known that the operating temperature affects the performance, power consumption, and reliability of a microelectronic system. Due to the huge computational requirements, it is almost impossible to model temperature and analyze the thermal effects of a system together with the environment in their full details. Using detailed numerical analysis methods, such as finite-element method (FEM), is time-consuming and cost-inefficient, hence, it is not a proper way to model temperature except for special cases. The best trade-off is offered by compact thermal models (CTMs) with reasonably accurate temperature predictions at different levels, e.g., circuit level, die level, package level, etc. [1].

A top-down hierarchy of CTMs would be helpful for designers at different design levels [1]. There are several desired

features that increase the usefulness of such CTMs at a particular design level.

- 1) *Detailed temperature distribution*: A CTM should provide enough thermal information at the desired design level. For example, for package-level CTMs, previous studies [2], [3] have shown that the information of temperature distribution across the package is required. Using only a single junction-to-case thermal resistance will lead to an inferior package design; instead, multiple nodes are needed on the package surfaces. Similarly, a CTM at the silicon die level should consist of enough nodes that give detailed temperature distribution information across the die. In addition, both static and transient temperatures should be modeled.
- 2) *Granularity*: A CTM should model just at the granularity that is needed and hide the details of the lower levels, so that the CTM itself is no more complex than necessary. Modeling at finer granularity introduces unnecessary details and makes the computation slower. For example, package-level CTMs, such as the DELPHI models [4]–[6], hide the lower level details of the package structures, including the die, the thermal attach, the solder balls and so on, mainly because these details are intellectual properties of the vendors, but also because they would just increase the complexity of the model without significantly improving the simulation accuracy. Similarly, a CTM at the die level should also hide the lower level details of the die, such as the actual circuit structures and physical layout.
- 3) *Parametrization*: A fully parameterized CTM allows package designers, circuit designers, and computer architects to explore new design alternatives and evaluate different thermally-related design trade-offs at their corresponding design levels before the actual physical designs are available. More importantly, with the aid of the parameterized CTMs, designers at different design levels can have more productive interactions and collaborations at early design stages of a microelectronic system. This leads to early discovery and considerations of potential thermal hazards of the system. True parametrization requires that the models be constructed based solely on design geometries and material properties.
- 4) *Boundary condition independence (BCI)*: A crucial feature of CTMs is BCI. By achieving BCI, the variation of the environment does not affect the actual model. The package-level CTMs in [4] and [5] achieve BCI by finding a thermal resistance network with minimum overall error when applied to different boundary conditions [5], [7], [8]. In Section IV, we show that our physical CTM is BCI with reasonable accuracy.

Manuscript received October 28, 2004; revised March 18, 2005. This work was supported in part by National Science Foundation Grant CCR-0133634, two grants from Intel MRL, and a University of Virginia FEST Excellence Award. This work was recommended for publication by Associate Editor V. Szekeley upon evaluation of the reviewers' comments.

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).

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

Digital Object Identifier 10.1109/TCAPT.2005.859737

- 5) *Computational speed*: The structure of a CTM needs to be relatively simple. Additionally, efficient algorithms should be developed to simulate the CTM with little computational overhead, so that thermal analyses at all design levels can be carried out efficiently.

There have been abundant existing contact thermal modeling methods in the literature. For example, the authors of [4], [5], and [9] propose the DELPHI approach and introduce the important concept of BCI. The DELPHI approach extracts and optimizes a thermal resistance network from detailed model simulations under a set of standard boundary conditions. In [1] and [10], the authors independently propose alternative CTM methods considering nonuniform boundary conditions and nonuniform boundary heat flux, resulting in much less optimization efforts compared to the DELPHI approach. There are also other modeling methods using model reduction techniques [11]–[13] or extracting transient thermal  $R$ - $C$  network from real package temperature measurements [14] and detailed model simulations [15]. All these models are accurate to characterize existing designs, but not fully parameterized to perform efficient explorations of new design alternatives.

In this paper, we present our CTM method, which takes a structured assembly approach of constructing a physical CTM by first modeling the silicon die and other packaging components as a collection of simple three-dimensional (3-D) shapes and then assembling them into more complex CTMs according to the overall structure. This modeling approach improves our previous work [16] and makes it fully parameterized and satisfy most of the above desirable features as shown later in the paper.<sup>1</sup>

The paper is organized as follows. Section II describes the general modeling method and discusses the differences between our modeling approach and the existing ones. Section III discusses the importance of CTM parametrization in more detail and gives examples of using our models in this way. In Section IV, we validate that our modeling approach is reasonably BCI by comparing the structural differences and the results of our model with a BCI DELPHI model. Section V discusses some of the advantages and limitations of our modeling method. Finally, Section VI concludes the paper and points out possible future work.

## II. MODELING METHODOLOGY

### A. General Method

Fig. 1 shows a typical modern single-chip package [17]. Heat generated from the active silicon device layer is conducted through the silicon bulk to the thermal interface material, heat spreader and heat sink, then convectively removed to the ambient. In addition to this primary heat transfer path, a secondary heat flow path uses conduction through the interconnect layer, input/output (I/O) pads, ceramic substrate, leads/balls to the printed circuit board (PCB). Our method can model all these layers and both heat flow paths. We also consider lateral heat flow within each layer to achieve greater accuracy of temperature estimation.

<sup>1</sup>Except that, in this paper, we only focus on static CTMs. Dynamic CTM using this approach is still work in progress, especially validation and comparison with other approaches [14] and [15].




Fig. 1. Packaging components in a typical ceramic ball grid array (CBGA) package (adapted from [17]). Primary and secondary heat transfer paths are shown.

In our modeling method, each layer is first divided into a number of blocks. For example, the silicon bulk layer can be divided into an irregular set of blocks, according to architecture-level functional units, or into a regular set of grid cells at lower or finer granularity, depending on what the design requires. Fig. 2 shows an example regular grid primary heat transfer path with  $3 \times 3$  grid cells on the silicon layer. In order to improve accuracy, the thermal interface material and the center part of the heat spreader that is right under the interface material are also divided into the same number of grid cells as the silicon die. The remaining outside part of the heat spreader is divided into four trapezoidal blocks. The heat sink is divided into five blocks: one corresponding to the area right under the heat spreader and four trapezoids for the periphery. Other package layers in Fig. 1 (e.g., C4 pads and ceramic substrate) are not shown but can be modeled similarly. Each block or grid cell maps to a node in the thermal circuit. In addition, each block or grid cell in each layer has one vertical thermal resistance and several lateral resistances, which model vertical heat transfer to the layers above and below, and lateral heat spreading/constriction within the layer itself to the neighboring blocks, respectively. The vertical resistance values are calculated directly based on the shape and material properties as  $R_{\text{vertical}} = t/(k \cdot A)$ , 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. Calculating the lateral thermal resistances is slightly more complicated because heat spreading and constriction must be accounted for. Basically, the lateral thermal resistance on one side of a block or a grid cell can be considered as the spreading/constriction thermal resistance of the neighboring part within a layer to that specific block. Details of lateral thermal resistance derivation and formulas can be found in [18] and [19].

For layers that have surfaces interfacing with the ambient, i.e., the boundaries, we assume that each surface (or part of a surface) has a constant heat transfer coefficient  $h$ . The corresponding thermal resistance is then calculated as  $R_{\text{convection}} = 1/(h \cdot A)$ , where  $A$  is the surface area. Strictly speaking, these convection thermal resistances are not part of the CTM, because they include information about the environment. If the environment changes, i.e., the boundary conditions change, these con-




Fig. 2. Example of CTM with  $3 \times 3$  grid cells for silicon die, thermal interface material and center part of the heat spreader. For clarity, the structure is shown upside-down compared to the primary heat transfer path in Fig. 1. Heat sources are also omitted for clarity [25].

vection resistances also change. On the other hand, for a particular design, the values of all the other internal thermal resistances shown in Fig. 2 will not change if the CTM is BCI. Later we show that this is indeed the case for our models.

From the above description, it is worth noting that more package layers can be easily included in the models, due to the structured assembly nature of our method. Therefore, our approach can also accommodate emerging packaging schemes such as stacked chip-scale packaging (SCP) [20] and 3-D integration [21].

### B. Functional Units Versus Regular Grids

As mentioned in Section II-A, the silicon die, the interface material, and the center part of the heat spreader can either be divided naturally according to functional units or be divided into regular grid cells, depending on the needs of the designer. For example, a computer architect may only need to estimate average temperatures for each functional unit, a thermal model at the functional unit granularity being best in that case. However, for a package designer, the temperature gradients across silicon die and other package layers are important metrics to evaluate the reliability of the package. In this case, a grid-like thermal model is more suitable, since it provides more detailed estimations of maximum and minimum temperatures, whereas the functional-unit-level model does not. Our method offers a choice between a possibly irregular functional-unit partitioning and a regular grid of variable granularity.

When modeling the die with regular grid cells, our method is essentially a combination of simplified finite difference method (for the silicon die and the interface material) and FEM (for other package components). On the other hand, if the die is modeled at the architectural functional unit level, our modeling method can be considered as a very simple form of FEM approach. Our modeling method achieves low computational overhead as well as reasonable accuracy by taking advantage of the established architecture-level and package-level information during the simplified finite-element and finite difference processes.

### C. Model Validation

We have validated our models with 3-D FEM simulations [19] and by comparing with measurements on a commercial thermal test chip [22]. In validation with measurements, we neglected

the secondary heat flow path from the die to the PCB, because the test chip is wire bonded and plugged in a plastic socket that has very low thermal conductivity. The CTM is built according to the geometries and material properties of the thermal test chip package whose primary heat transfer path is similar to Fig. 1, except there is no heat sink for the thermal test chip. In the thermal test chip, there is a  $9 \times 9$  array of heaters that can be turned on and off individually. By changing the supply voltage to the test chip, the dissipated heat to the heaters can be tuned to different values. Each heater also has an associated thermal sensor measuring its temperature. In our validation experiments using this thermal test chip, we turned on sets of heaters in the test chip and assigned the same power values at the same locations in our thermal model. Steady-state temperature can be reached after ten or more minutes when no further changes in temperature sensor measurements could be observed. Among all the different validation experiments, we found the ones with the heat sources at one of the die corners have the largest errors. This is due to the larger across-die temperature difference since heat source is at the corner and the complicated heat transfer paths near the die edge that are not perfectly accounted for in our CTM. Fig. 3 shows one example of the validation experiments we have performed with  $3 \times 3$  heaters turned on with power density of  $50 \text{ W/cm}^2$  at the lower-right corner. Other experiments such as the ones with heat sources near the die center have less errors and are not shown here. This figure compares the steady-state thermal plots from temperature sensor measurements on the test chip and from the simulation results of our thermal model. The percentage errors are calculated by  $(T_{\text{model}} - T_{\text{chip}})/(T_{\text{chip}} - T_{\text{ambient}})$  for each of the  $9 \times 9$  cells. As can be seen, our thermal model is reasonably accurate, with the worst-case error values for steady-state temperatures less than 5% in this experiment. The ambient temperature was  $25^\circ\text{C}$ .

### D. Differences to Existing CTMs

From the above it can be seen that our modeling approach is different from other existing CTM methods.

- 1) Existing CTMs are mainly at the package or system level, hence with only a few nodes for the silicon. This is adequate for the package vendors and system-level designers, but not for circuit designers or architects. Our CTMs uses more nodes for the die structures, which is convenient for the silicon-level designers.
- 2) Existing CTMs hide the packaging details due to the requirements of package vendors, while our models need somewhat detailed package-level and silicon-level information.
- 3) The thermal resistances in the existing CTMs are extracted from detailed model simulations or real package temperature measurements, while the thermal resistances of our models are calculated based on dimensions and physical properties of materials.

All these differences lead to different applications for existing CTMs and our CTMs. Existing models are good for thermal analysis and characterization of existing package designs without revealing details of the package, while our models are




Fig. 3. Steady-state validation of our CTM: (a) test chip measurements and (b) results from our model with errors less than 5% [25]. The  $3 \times 3$  lower-right corner dissipates power, with a power density of  $50 \text{ W/cm}^2$ .

better for explorations of new silicon-level and package-level designs. Also, our models are intrinsically parameterized and reasonably BCI, as discussed in Sections III and IV.

### III. PARAMETRIZATION

Parametrization of CTMs is desirable and has drawn attention from researchers. In [8] and [23], the authors point out that achieving a sensible parametrization of CTMs is next to impossible for the chosen simple structure of some of the existing models, such as the DELPHI ones. This is because the DELPHI model structure consists of only a few thermal resistances which makes it impossible to parameterize the actual very complex package structure, together with the variations of thermal conductivities and the heat spreading/constriction effects within the die and the package. On the other hand, our modeling approach can be better parameterized due to its physically-based nature. The cost for parametrization is that our models are usually more complex than, and not as accurate as, the DELPHI-like and other existing models.

The importance of achieving full parametrization of CTMs is obvious. Fully parameterized CTMs allow designers at all design levels to freely explore all the possible thermally-related design spaces. For example, the heat spreader and heat sink are two important package components for high-performance VLSI designs. While a large heat spreader and heat sink made from high thermal conductivity materials can reduce the temperature of silicon die, they also significantly increase the total price of the system. Therefore, exploring design trade-offs between hot spot temperatures and package cost is crucial. With parameterized CTMs, this exploration can be done easily and efficiently by simply sweeping the size of the heat spreader and heat sink with different material properties to achieve the desired package design point. On the other hand, building package prototypes or detailed thermal models greatly slows the design process and increases the design cost.

In addition, fully parameterized CTMs can also provide a more productive communication channel among designers at different design levels. Therefore, potential thermal hazards in the design can be discovered and dealt with in early design

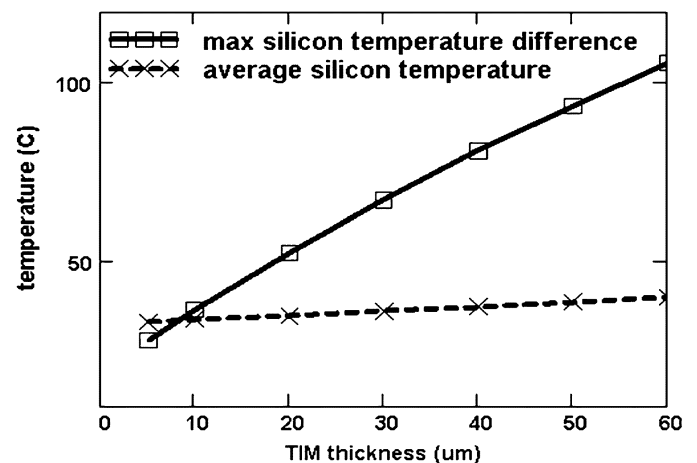


Fig. 4. Impact of thermal interface material (TIM) thickness to silicon die temperature difference. Average silicon die temperature is also plotted as reference.

stages. For example, a typical design scenario would be: circuit designers come up with estimations of power consumptions of each circuit block, computer architects come up with a floorplan, while package designers come up with a proposed package. Using a parameterized CTM, these designers can together easily evaluate the combined system design from a thermal point of view. If it appears that there are unacceptable hot spots on silicon, or the temperature difference across the package is too high and degrades the reliability of the design, different design decisions can be made in this early design stage—circuit designers may need to develop novel circuits with lower power consumption, computer architects may apply different dynamic thermal management techniques and rearrange the floorplan for better across-silicon temperature distribution, or the package designer can improve the proposed package design by using a more advanced heat spreader or heat sink, etc.

In fact, in [11], the authors have proposed CTMs which include the parametrization of certain package parameters such as the PCB size, PCB thermal properties, heat sink dimensions and thermal vias. However, full parametrization was not achieved in

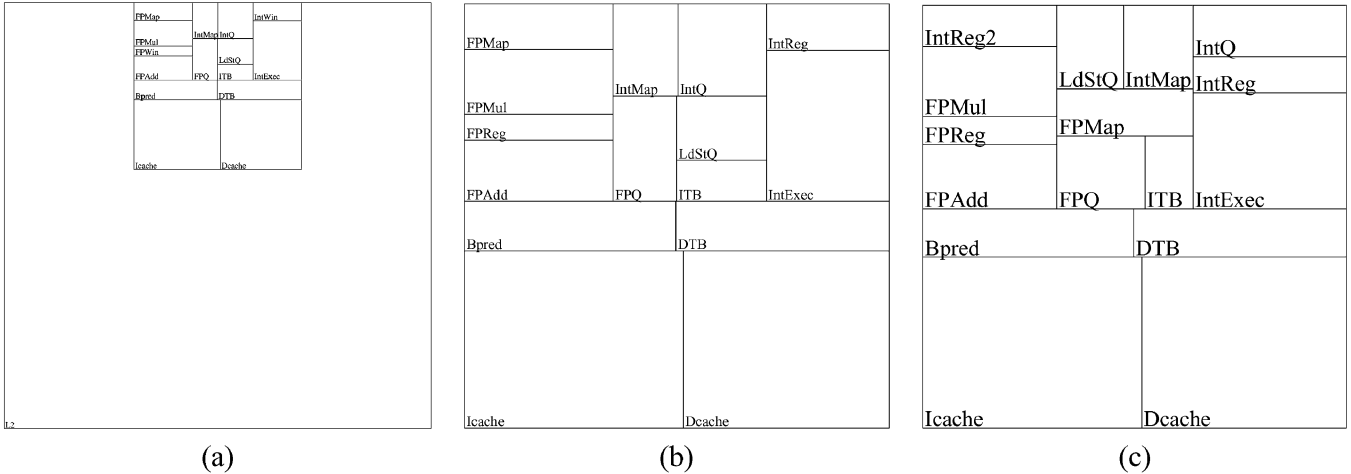


Fig. 5. (a) Approximated floorplan of Alpha 21364 microprocessor. (b) Close-up look of the 21364 core, with only one register file on the top-right corner. (c) 21364 core with an extra copy of register file at the top-left corner.

[11] because the die itself is not parameterized. On the other hand, our modeling approach can be considered as fully parameterized, including the silicon die itself. The parametrization of the die is a crucial requirement to explore different preliminary die-level designs before the prototypes are available.

For illustration, we first present an example analysis to show the strength of using parameterized CTMs to efficiently investigate the impact of thermal interface material on across-silicon temperature differences. Fig. 4 shows the relationship between the thickness of the thermal interface material (TIM), which glues the silicon die to the heat spreader. We plot the temperature readings from a CTM similar to Fig. 2 with  $40 \times 40$  grid cells on silicon. This analysis is based on an Alpha 21364-like microprocessor floorplan. Average silicon die temperature is also plotted in Fig. 4 for reference. The total heat generated from the silicon surface is 40.2 W, the die size is  $15.9 \text{ mm} \times 15.9 \text{ mm} \times 0.5 \text{ mm}$ , and the thermal conductivity of the thermal interface material is  $1.33 \text{ W/(m-K)}$ .

As can be observed from Fig. 4, although the TIM thickness doesn't have obvious impact on the average die temperature, thicker TIM results in poor heat spreading which leads to large temperature differences across the die. Such large temperature differences may be disastrous to circuit performance and die/package reliability. Using a better heat sink will only lower the average silicon temperature but will not help to reduce the temperature difference. From this analysis, which has been easily performed by our parameterized model, we can reach the conclusion that using as thin as possible thermal interface material is one of the key issues for package designers to consider. In some recent work [5], [24], the importance of measuring and modeling thermal interface's impacts on the entire package has been also discussed, although not at the same silicon die level as the above example shows.

Another example of utilizing our parameterized CTM is the investigation of a dynamic thermal management technique (DTM) called migrating computation (MC), at the microarchitecture level [19]. It is obvious that two silicon-level functional units that run hot by themselves will tend to run even hotter when adjacent. On the other hand, separating them will in-

roduce additional communication latency that is incurred regardless of operating temperature. This suggests the use of spare units located in cold areas of the chip, to which computation can migrate only when the primary units overheat. In [19], the floorplan of the Alpha 21364 core is carefully changed to include an extra copy of integer register file (see Fig. 5), which is usually the hottest spot on the silicon die for this design. We also model the fact that accessing the secondary integer register file entails extra power and extra access time due to the longer distance. With this new floorplan, we can shift the workload back and forth between the two register files when the one in use overheats, with a little performance overhead (11.2% slower). The changes in silicon floorplan can be easily adapted into corresponding parameterized CTMs, thus the temperatures of functional units can be analyzed efficiently to investigate the usefulness of the migrating computation DTM technique. By doing this, packaging complexity and cost can be significantly reduced and we can still get almost the same operating temperature and performance as if we use a much more expensive and complicated thermal package.

#### IV. BOUNDARY CONDITION INDEPENDENCE

Another important aspect of CTM is BCI [4], [5], [7]. Achieving BCI is essential to CTMs. If the model changes whenever the boundary conditions change, the model would be almost useless. Traditionally, researchers in the package CTM community usually adopt the DELPHI approach to achieve BCI, that is, finding a thermal resistance network with minimum overall error when applied to different boundary conditions. The resistance values are extracted from detailed thermal simulations with the same package structure. Such simulations can be performed in numerical analysis tools.

When using our modeling approach, because there is no data extraction procedure and all the resistance values are calculated from the physical dimensions and properties of the materials, the model itself should be reasonably BCI if the major heat transfer paths are properly modeled. At the present state, in order to validate that our modeling method can indeed achieve reasonable BCI, we compare our CTMs with the DELPHI models. Since




Fig. 6. Thermal resistances network for (a) the BCI DELPHI model and (b) our thermal model of a DELPHI BGA benchmark chip, extracted from [6].

TABLE I  
COMPARISON OF OUR COMPACT THERMAL MODEL AND DELPHI MODEL FOR THE DELPHI BGA BENCHMARK CHIP UNDER THE SAME SET OF BOUNDARY CONDITIONS. TEMPERATURES ARE IN CELSIUS AND WITH RESPECT TO AMBIENT TEMPERATURE

#	b.c.	our model (HotSpot)	DELPHI	error
1	DGP-1	16.79	16.68	0.66%
2	DGP-2	19.94	20.00	-0.30%
3	DGP-3	66.42	62.78	5.80%
4	DGP-4	2960.00	3070.00	-3.58%
5	infinite	10.20	10.56	-3.41%

the DELPHI models have been extensively validated to be BCI for a large set of boundary conditions, we regard the DELPHI models as proper benchmarks for the BCI validations of our modeling approach. Comparisons with real designs will be part of our future work. One of our BCI validations is done by comparing with a DELPHI BGA benchmark chip in [6]. The dimensions of the BGA benchmark chip and package and the set of boundary conditions are both taken from the specifications in [5] and [6]. The model structures of both the DELPHI model and our model are shown in Fig. 6. In this comparison, the notion of quarter symmetry can be applied because there is only one node needed for the die in both models. Therefore, only a quarter of the package is sketched for our model and the DELPHI models in Fig. 6. The first four standard boundary conditions in Table I correspond to different typical settings of heat transfer coefficients for the top, bottom, side, and the lead of the package. The last boundary condition accounts for the ideal case where the heat transfer coefficients of all the surfaces are infinite [5], [6].

The temperature readings from both models are also listed in Table I. The heat generated at the die surface is 2.5 W, which is used as the input to both models. As can be seen from Table I, our model achieves reasonable BCI. For the listed five standard boundary conditions, it yields almost the same temperature readings as the DELPHI model. The worst case percentage error is 5.8%. One possible reason of the error is that the top surface division ratio is fixed according to the area of the smaller neighbor layer, in this case, it is the die area. This division ratio might not be exactly the optimal ratio, but it is sufficiently near the optimal ratio. In a previous work [10], the author argues that

the surface division ratio should be determined by the heat flux distribution on a particular surface. He also shows that the heat flux distribution function  $f(-)$  of the top surface develops a peak just above the die area. Therefore, using the die area to divide the top surface as in our model is reasonable.

## V. ADVANTAGES AND LIMITATIONS

### A. Advantages

So far, the modeling method and major characteristics of our CTM approach have been presented. Despite a few limitations, our modeling approach has several significant advantages; the advantages are mainly due to the fact that it is parameterized and BCI.

- 1) Parametrization is useful because a variety of design explorations can be carried out by only changing the dimensions and material parameters without reconstructing the whole CTM through detailed simulations. For example, using our models, one can easily find the optimum die thickness by simply sweeping the die thickness parameter and keeping all the other parameters constant. Another example would be investigating the effect of different types of heat spreaders or heat sinks. One can easily add/change the layers of heat spreader or heat sink by following our modeling method in Section II.
- 2) It is also important to notice that our modeling method can be used to study hypothetical systems for which physical implementations and thermal measurements cannot yet be obtained. One example is the investigation of emerging 3-D integrated circuits. Prototyping 3-D integrated circuits (ICs) would introduce prohibitive cost due to the drastic change of existing fabrication process. But with our models, designers can easily model heat transfer and temperature rise in 3-D IC structures, thus evaluate the feasibility of this new design paradigm from a thermal point of view.
- 3) Because our modeling approach has been validated to be reasonably BCI, designers can focus more on their design efforts without worrying about the thermal model's validity under different boundary conditions.
- 4) There have been several successful applications of our modeling approach in different design areas. For example,

it has been used to build CTMs in research areas such as dynamic thermal management (DTM) techniques for microprocessors [19] and die-level thermal-aware computer-aided designs [25]. Similarly, we expect the presented modeling method will stimulate more research collaborations among package designers, VLSI circuit designers, and computer architects.

### B. Limitations

Our modeling approach also has a few limitations.

- 1) It is not as “compact” as other existing CTMs, but the number of nodes are still within a manageable amount, and the computational overhead is also negligible compared to detailed numerical models.
- 2) When it comes to analyze or release a fixed CTM for an existing design or a final product, our models are not as friendly as most existing models to the users. This is due to the complexity of our model and the revealing of package design details.
- 3) At the same level of complexity, our model is not as accurate as other existing models. This is because the existing models are extracted from detailed model simulations or real package temperature measurements, which are still the most accurate ways to model thermal effects, while our model is essentially a simplified version of the detailed model, and therefore cannot achieve the same level of accuracy as the detailed simulations, or the derived DELPHI-like models. In addition, some lumped thermal resistances (e.g., the ones in the peripheral parts of heat spreader and heat sink, see Section II-A) do not fully account for all the possible heat transfer paths, thus, not really representing the exact thermal resistance according to the analysis in [4] and [26].
- 4) Our modeling approach is not as BCI as the DELPHI and other existing modeling approaches. This is because the surface area division method used by our model is not exactly the optimal one [23], although it is proved to be a reasonable one as shown in Section IV and [10]. It is also partly due to the fact that heat spreading in the package cannot be as well considered in our modeling approach as in the detailed models.
- 5) Our modeling approach needs to be improved to account for more different types of packages. For example, if one layer of the package contains more than one material or the package itself is a multichip module (MCM), our current thermal models are not flexible enough to cover these cases. This is just a limitation of our current implementation since our general modeling method can be easily extended to account for more complicated structures.

## VI. CONCLUSION

In this paper, we have presented a physical CTM approach, which is parameterized and BCI. Our modeling approach is more suitable for exploring new designs, while most existing modeling approaches are more suitable for accurately analyzing and characterizing existing designs. Achieving parametrization

for CTMs is a significant contribution of our method to the research area of package level design and analysis. Important future work consists of extending our modeling to include the secondary heat transfer path, modeling transient temperature, validating, and analyzing our modeling method’s boundary and initial condition independence (BICI) with real package designs. We also need to accommodate more types of packages in our modeling method, such as MCMs. Modeling active cooling effects and different types of interfacing surfaces are also interesting topics.

## REFERENCES

- [1] M.-N. Sabry, “Compact thermal models for electronic systems,” *IEEE Trans. Comp. Packag. Technol.*, vol. 26, no. 1, pp. 179–185, Mar. 2003.
- [2] H. Rosten and C. Lasance, “Delphi: The development of libraries of physical models of electronic components for an integrated design environment,” in *Proc. Conf. Int. Electronic Packaging Soc.*, 1994.
- [3] A. Bar-Cohen, T. Elperin, and R. Eliasi, “ $\theta_{jc}$  characterization of chip packages—Justification, limitations and future,” *IEEE Trans. Comp., Hybrids, Manufact. Technol.*, vol. 12, no. 4, pp. 724–731, Dec. 1989.
- [4] H. I. Rosten, C. J. M. Lasance, and J. D. Parry, “The world of thermal characterization according to DELPHI—part I: Background to DELPHI,” *IEEE Trans. Comp. Packag. Technol.*, vol. 20, no. 4, pp. 384–391, Dec. 1997.
- [5] C. J. M. Lasance, H. I. Rosten, and J. D. Parry, “The world of thermal characterization according to DELPHI—part II: Experimental and numerical methods,” *IEEE Trans. Comp. Packag. Technol.*, vol. 20, no. 4, pp. 392–398, Dec. 1997.
- [6] C. J. M. Lasance, “Two benchmarks to facilitate the study of compact thermal modeling phenomena,” *IEEE Trans. Comp. Packag. Technol.*, vol. 24, no. 4, pp. 559–565, Dec. 2001.
- [7] —, “The influence of various common assumptions on the boundary-condition-independence of compact thermal models,” *IEEE Trans. Comp., Packag., Manufact. Technol. A*, vol. 27, no. 3, pp. 523–529, Sep. 2004.
- [8] H. Vinke and C. J. M. Lasance, “Compact models for accurate thermal characterization of electronic parts,” *IEEE Trans. Comp., Packag., Manufact. Technol. A*, vol. 20, no. 4, pp. 411–419, Dec. 1997.
- [9] C. J. M. Lasance, H. Vinke, H. Rosten, and K. L. Weiner, “A novel approach for the thermal characterization of electronic parts,” in *Proc. 11th IEEE SEMI-THERM Symp.*, 1995, p. 1.
- [10] E. G. T. Bosch, “Thermal compact models: An alternative approach,” *IEEE Trans. Comp. Packag. Technol.*, vol. 26, no. 1, pp. 173–178, Mar. 2003.
- [11] M. Furmanczyk, A. Napieralski, K. Szaniawski, W. Tylman, and A. Lara, “Reduced electro-thermal models for integrated circuits,” in *Proc. 1998 Int. Conf. Modeling Simulation Microsystems, Semiconductors, Sensors Actuators*, Apr. 1998, pp. 139–144.
- [12] B. Shapiro, “Creating compact models of complex electronic systems: An overview and suggested use of existing model reduction and experimental system identification tools,” *IEEE Trans. Comp. Packag. Technol.*, vol. 26, no. 1, pp. 165–172, Mar. 2003.
- [13] L. Codecasa, D. D’Amore, and P. Maffezzoni, “An Arnoli based thermal network reduction method for electro-thermal analysis,” *IEEE Trans. Comp. Packag. Technol.*, vol. 26, no. 1, pp. 186–192, Mar. 2003.
- [14] M. Rencz, G. Farkas, V. Sz’ekely, A. Poppe, and B. Courtois, “Boundary condition independent dynamic compact models of packages and heat sinks from thermal transient measurements,” in *Proc. Electronics Packaging Technology Conf.*, 2003, pp. 479–484.
- [15] H. Pappe, D. Schweitzer, J. H. Janssen, A. Morelli, and C. M. Villa, “Thermal transient modeling and experimental validation in the european project PROFIT,” *IEEE Trans. Comp. Packag. Technol.*, vol. 27, no. 3, pp. 530–538, Sep. 2004.
- [16] 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.
- [17] 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. Comp., Packag., Manufact. Technol. A*, vol. 21, no. 1, pp. 104–112, Mar. 1998.

- [18] S. Lee, S. Song, V. Au, and K. Moran, "Constricting/spreading resistance model for electronics packaging," in *Proc. AJTEC*, Mar. 1995, pp. 199–206.
- [19] K. Skadron, K. Sankaranarayanan, S. Velusamy, D. Tarjan, M. R. Stan, and W. Huang, "Temperature-aware microarchitecture: Modeling and implementation," *ACM Trans. Architect. Code Optimiz.*, vol. 1, no. 1, pp. 94–125, Mar. 2004.
- [20] *Stacked chip scale packaging*, [Online] Available: <http://www.intel.com/research/silicon/mobilepackaging.htm>, 2005.
- [21] 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.
- [22] V. Székely, C. Márta, M. Rencz, G. Végh, Z. Benedek, and S. Török, "A thermal benchmark chip: Design and applications," *IEEE Trans. Comp., Packag., Manufact. Technol. A*, vol. 21, no. 3, pp. 399–405, Sep. 1998.
- [23] C. J. M. Lasance, "Recent progress in compact thermal models," in *Proc. 19th IEEE SEMI-THERM Symp.*, Mar. 2003, pp. 290–299.
- [24] —, "The urgent need for widely-accepted test methods for thermal interface materials," in *Proc. 19th IEEE SEMI-THERM Symp.*, 2003, pp. 123–128.
- [25] W. Huang, M. R. Stan, K. Skadron, K. Sankaranarayanan, S. Ghosh, and S. Velusamy, "Compact thermal modeling for temperature-aware design," in *Proc. 41st DAC*, Jun. 2004, pp. 878–883.
- [26] C. J. M. Lasance, "Thermal characterization of electronic parts with compact models: Interpretation, application, and the need for a paradigm shift," in *Proc. 13th IEEE SEMI-THERM Symp.*, 1997, pp. 40–48.



**Wei Huang** (S'01) received the B.S. degree in electrical engineering from the University of Science and Technology of China, Beijing. He is currently pursuing the Ph.D. degree in electrical and computer engineering at the University of Virginia, Charlottesville.

He was an Intern with the VLSI Design Department, IBM T. J. Watson Research Center, Yorktown Heights, NY, in 2004. His research interests include modeling and analysis of temperature and other parameter variations in circuits and computer architecture, and the design of variation-aware VLSI circuits and systems.


Mr. Huang received the Best Student Paper Award at ISCA 2003.



**Mircea R. Stan** (SM'98) received the M.S. degree in electronics and communications from Politehnica University Bucharest, Bucharest, Romania, and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Massachusetts, 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 has more than eight years of industrial experience, was a Visiting Scholar at the University of California, Berkeley from 2004 to 2005, a visiting faculty at IBM in 2000, and at Intel in 1999 and 2002. 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.

Dr. Stan received the NSF CAREER Award in 1997 and Best Paper Awards at ISCA'03 and SHAMAN'02. He has been a Guest Editor for the IEEE TRANSACTIONS ON COMPUTERS 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 VLSI SYSTEMS from 2001 to 2003. He is a member of ACM, Usenix, Eta Kappa Nu, Phi Kappa Phi, and Sigma Xi. He was a Program Chair for ISLPED'05, the General Chair for GLSVLSI'03, has been on technical committees for several conferences, and is a Distinguished Lecturer for the IEEE Circuits and Systems Society for 2004–2005.



**Kevin Skadron** (SM'05) received the B.S. degree in electrical engineering 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 received three Best-Paper Awards, an NSF ITR Award, and a CAREER Award. He is Founding Associate Editor-in-Chief of IEEE COMPUTER ARCHITECTURE LETTERS and a member of ACM. He co-Chaired MICRO'04 and PACT'02 and has served on numerous program committees.