Thermal Analysis and Optimization of 2.5-D Integrated Voltage Regulator

Taigon Song[1], Noah Sturcken[2], Krit Athikulwongse[1], Kenneth Shepard[2], and Sung Kyu Lim[1]
[1]Department of ECE, Georgia Institute of Technology, Atlanta, GA 30332, U.S.A.
[2]Department of EE, Columbia University, New York, NY 10027, U.S.A.
{taigon.song}@gatech.edu

***Abstract*—Integrated voltage regulators (IVRs) promise to improve performance-per-watt for microprocessors and systems-on-chip by reducing supply voltage margins, resistive losses in the power distribution network, and by enabling power management with greater temporal resolution. However, the thermal impact of IVRs has not been well studied, and the methodologies for thermal analysis of analog/digital mixed-signal designs, such as voltage regulators on chip, have not yet been developed. In this paper, we present a thermal analysis methodology for 2.5-D IVR. Our results show that (1) the integrated power inductor is the hottest component in the IVR, and (2) the temperature of the IVR rises rapidly when the power inductor and the circuitries of the IVR chip overlap. In order to address these issues, we propose two design optimization techniques: design block relocation and inductor spreading. Related experiments show the effectiveness of these methods.**

***Index Terms*—Integrated Voltage Regulator; Thermal Analysis; Analog; Mixed System; Silicon Interposer;**

I. Introduction

On-chip integrated voltage regulators (IVRs) improve energy efficiency in microprocessors and SoCs by enabling power delivery to ICs at higher voltages, reducing resistive losses in the power distribution network, and by enabling fine grained power management on time scales that are relevant for high-performance digital systems [1], [2], [3]. Recently, an early switched-inductor IVR prototype achieved especially high current density using 2.5-D chip stacking, where an IC was flip-chip attached to a passive 2-D silicon interposer carrying custom fabricated power inductors [4]. In this prototype, on-chip IVRs integrated with other digital and analog components demonstrate the potential for integrated power conversion in a realistic SoC.

However, one of the main issues that was not addressed in [4] is thermal hot spots. The buck converter in the IC and the inductor placed inside the interposer generate significant heat and affect the performance of digital and analog modules placed inside the die mounted on the interposer. Thus, designers must consider the thermal impact of these new heat sources being added to the whole system. Currently, there are several tools that perform thermal analysis for package design. Tools are also available for thermal analysis of digital IC designs. But, none of these tools can handle co-analysis of thermal distribution in the overall 2.5-D chip/package system. In addition, thermal analysis that involves IVRs and mixed signal components has not been studied well [5], [6], [7], [8].

In this paper, we propose a thermal analysis methodology for analog/digital mixed-signal system. By this methodology, we perform thermal analysis for a silicon interposer-based 2.5-D system, where mixed-signal components and an IVR are vertically integrated (see Fig. 1 [4]). The main contributions of this work include the following:

1) Full-chip thermal analysis flow for the mixed-signal 2.5-D design with an IVR: Existing studies handle thermal analysis of digital designs at chip-level, but they do not show how analog designs and off-chip components are considered together. Here, we propose a thermal analysis methodology that handles GDSII-level analog/digital components, an IVR, and an interposer in a holistic manner.
2) Chip-level thermal analysis of IVRs: For the first time, we perform thermal analysis of a 2.5-D IVR and study their thermal impact.
3) Design optimization: We propose design methods to reduce the temperature of our mixed-signal 2.5-D design with an IVR.

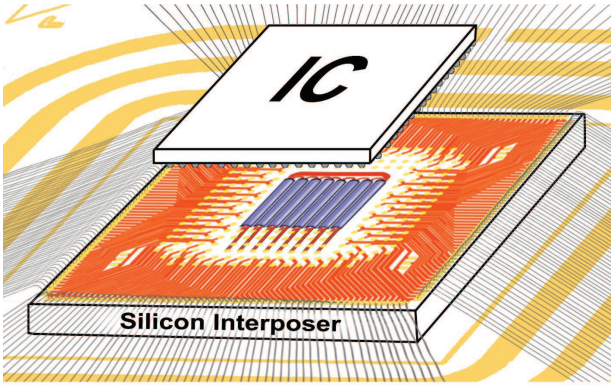

Fig. 1. Diagram of a 2.5-D integrated voltage regulator (IVR) chip stack used in our study. The IC consists of a buck converter and load circuitries. The silicon interposer contains power inductors. The IC is flip-chip mounted on the silicon interposer using ball grid array, and bond wires connect the silicon interposer to the IO.

II. Thermal Analysis Flow

In this section, we present our thermal analysis methodology for mixed-signal designs. First, we provide an overview of our methodology, followed by detailed description of individual steps.

A. Overview of the Methodology

Fig. 2 shows the thermal analysis flow developed in this work. Starting from the analog/digital (A/D) mixed-signal netlist, we separate it into an analog part and a digital part. We construct the GDSII layouts of these parts separately and then merge them. We use Cadence Encounter to obtain the layout of the digital part and Cadence Virtuoso for the analog part. In case of power analysis, we create separate test benches for the analog and digital parts. Once we obtain the GDSII layouts and power analysis results, we use our Design Analyzer to create thermal meshes, set up heat equations, and solve those using Ansys FLUENT.

B. Heat Equation

The following heat equation describes the steady-state temperature at a point $\mathbf{p} = (x, y, z)$ inside a 3D structure:

$$\nabla \cdot (k(\mathbf{p})\nabla T(\mathbf{p})) + S_{\mathrm{h}}(\mathbf{p}) = 0 \quad (1)$$

978-1-4673-2538-7/12/$31.00 ©2012 IEEE

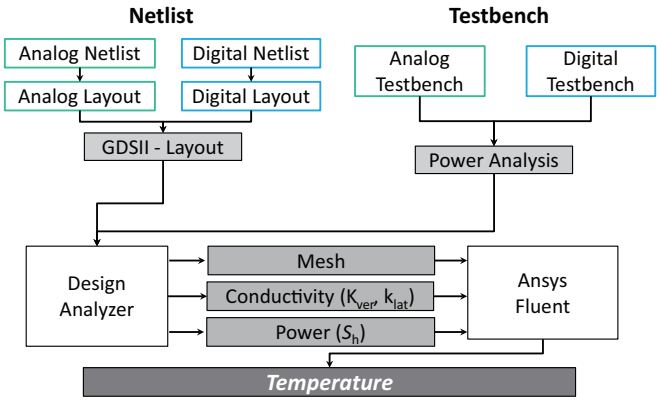

Fig. 2. Proposed thermal analysis flow for GDSII-level analog/digital mixed design.

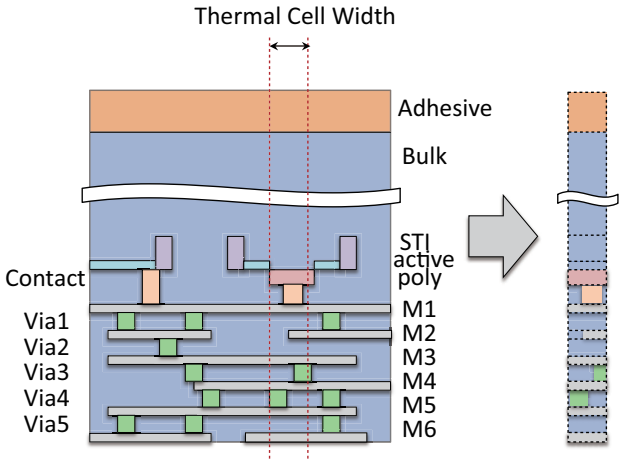

Fig. 3. Example of a thermal cell in 6 metal IC. This IC is modeled using 17 layers of vertical thermal cells (each thermal cell is a dotted square).

where k is $W/m \cdot K$ (thermal conductivity), T is K (temperature), and $S_{\rm h}$ is W/m^3 (volumetric heat source). To solve Eq. 1, we must assign boundary conditions to the six surfaces of the chip stack. Typically, a chip stack is very thin and packaged inside molding materials, which are not good thermal conductors. A majority of the heat flows from the bottom of the chip towards the heat sink. Thus, we apply an adiabatic boundary condition on the bottom and the four sides of the structure. A heat sink exists on the top side, and we apply a convective boundary condition. We implement a thermal model of Eq. 1 by meshing the IC into elements, as shown in Fig. 3. Each element, called a "thermal cell," has a volume of a specific length, width, and height. We generate the thermal model and compute thermal conductivity of the IC using our Design Analyzer.

We calculate the total power dissipated inside a thermal cell P_{cell} from the power dissipation of the logic cells placed in it. Then, the volumetric heat source $S_{\rm h}$ can be computed from

$$S_{\rm h} = \frac{P_{\rm cell}}{W_{\rm cell} \cdot H_{\rm cell} \cdot T_{\rm cell}} \quad (2)$$

where $W_{\rm cell}$, $H_{\rm cell}$, and $T_{\rm cell}$ are the width, the height, and the thickness of the thermal cell, respectively.

Next, we solve Eq. 1 using Ansys FLUENT, a commercial tool for thermal analysis. FLUENT uses the meshed structure generated by our Design Analyzer. In addition, we use user-defined functions in FLUENT to apply $k_{\rm ver}$, $k_{\rm lat}$, and $S_{\rm h}$, because they vary with the position of the thermal cell. Finally, with the boundary conditions described earlier, we run FLUENT to obtain the steady state temperature of all positions inside the 2.5-D chip+interposer stack.

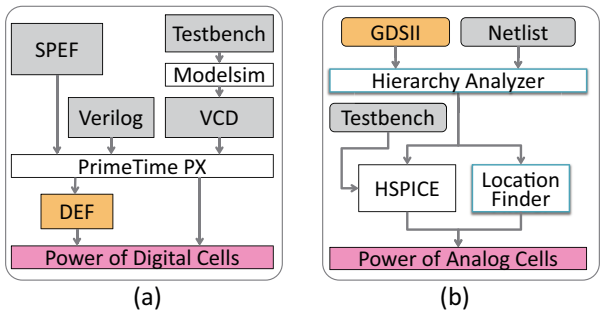

Fig. 4. Power analysis flow of (a) the digital design, (b) the analog design.

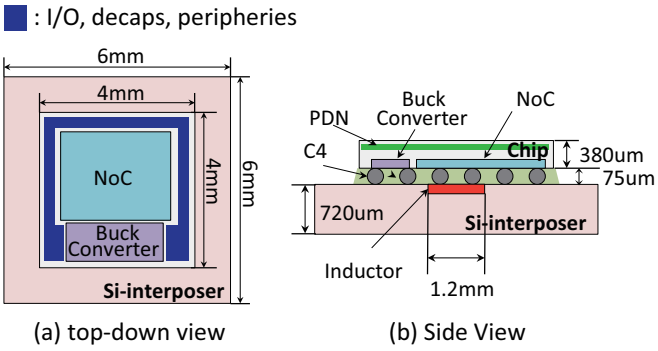

Fig. 5. (a) Top-down view, (b) side view of the IVR.

C. Power Analysis

Fig. 4 describes the proposed power analysis flow performed in our thermal analysis, where each MOSFET is a heat source $S_{\rm h}$. The power analysis is separated into two parts: digital and analog analysis. For the digital part, we first generate the layout in DEF format using Cadence Encounter, and we use the same tool to extract the parasitic resistance and capacitance in SPEF format. Based on the test bench of the digital netlist, we use Mentor Graphics Modelsim to generate the switching activity of each logic cell and obtain its VCD file. Next, we use Synopsys PrimeTime PX to perform static power analysis and obtain the power dissipation of each logic cell. We complete the power analysis by creating a POW file, a compilation of power dissipation and physical location of each cell in the DEF file.

For the power analysis of analog design, we utilize the hierarchy information in the netlist. For each macro in the design, we decide if this macro is reasonably-sized for HSPICE simulation using our Hierarchy Analyzer. If not, we simulate the sub-modules in the macro. Once the decisions are made, we use HSPICE and our Location Finder to compile the power dissipation with location of each macros.

III. Experimental Results

In this section, we perform thermal analysis of the 2.5-D IVR design [4] and study the effectiveness of proposed design optimization techniques. The physical dimensions are described first, followed by the analysis and the optimization techniques.

A. Dimensions and the Power Consumption

Our IVR can operate in many different output voltages and load conditions [4]. In addition, the efficiency of the IVR varies when the

TABLE I
MEASURED POWER CONSUMPTION OF THE MAIN BLOCKS.

	Low Current (1W)	Peak Efficiency (2W)	High Current (5W)
Buck Converter	83.55 mW	111.3 mW	305.5 mW
Inductor	323.7 mW	425 mW	1133 mW
NoC	1000 mW	2000 mW	5000 mW
PDN	45 mW	180 mW	1125 mW

load current changes. Therefore, in this work, we assume the IVR is operating at V_{IN}=1.8V and V_{OUT}=1.0V. Since the maximum output power of the IVR is 5W, we report the temperature when the load current changes from 1A to 5A in V_{OUT}=1.0V (= power range: 1W to 5W).

Fig. 5 describes the top-down view and the side view of the IVR. In the IVR system, the size of silicon interposer is 6mm×6mm, and the thickness is 720μm. The silicon interposer consists of a set of 8 cross-coupled inductors. The inductor footprint is 1.2mm×1.2mm, and the metal of the inductor is 5μm thick and 40μm wide. The size of the chip mounted at the center of the silicon interposer is 4mm×4mm, and the thickness is 380μm. NoC (network on chip) module is placed in the middle of the chip, and a buck converter is placed next to the NoC. Notice that the inductor is beneath the NoC due to the issues with I/O routing. The I/Os, decoupling capacitors, and peripheries are placed on the boundaries of the NoC and the buck converter. The silicon interposer and the chip are connected by 75μm C4 bumps, and epoxy underfill materials fill the empty space between the chip and the silicon interposer. On the top of the chip, we assume a 2mm-thick copper heat sink. The most power-consuming blocks in the 2.5-D system are the buck converter, the NoC, the PDN (power distribution network) of the chip, and the power inductor. Therefore, we focus our analysis on these four important blocks. Table I reports the measured power consumed by these main blocks.

B. Thermal Analysis of Essential Design Blocks

We start our thermal analysis by calculating the temperature of each main block assuming they operate separately. Fig. 6 shows the thermal map of each block. For the NoC, the maximum temperature rises up to 70.82°C when consuming high power (5W). Each NoC tile shows a similar temperature map because NoC connects 64 symmetric digital blocks. For the inductor, one set consists of eight coupled power inductors, and the maximum temperature rises up to 77.3°C. Because the inductor consumes high power in a relatively small footprint, the inductor is the hottest block of this 2.5-D system. For the buck converter, the maximum temperature is 54.49°C. The buck converter consists of one controller, and eight power drivers that are connected to the eight coupled inductors. The hot spots in the buck converter are from the eight power drivers. Between the power drivers, decoupling capacitors and other circuitries exist. Therefore, a temperature valley is created between the hot spots of the power drivers. These temperature valleys reduce the temperature of the buck converter.

Notice that there is a factor which can contribute to the global temperature rise in the IVR. PDN of the IVR chip is a path where the input power (V_{IN}, I_{IN}) must flow before reaching the buck converter. Since the buck converter and the NoC are integrated together, the PDN exists above both the NoC and the buck converter. However, the temperature rise by the PDN is not severe, because the size of the PDN is same as the total chip (4mm×4mm). Fig. 7 shows a temperature map when all components are assembled together. The maximum temperature of the entire 2.5D system is 114.96°C. A

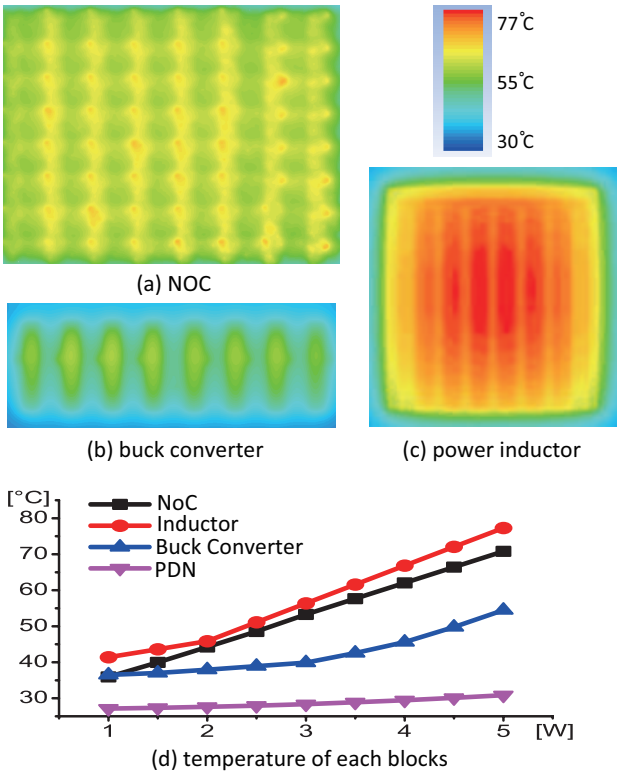

Fig. 6. (a,b,c): Thermal maps, (d) temperature of each blocks. (a) NoC, (b) buck converter, and (c) power inductor when consuming 5W.

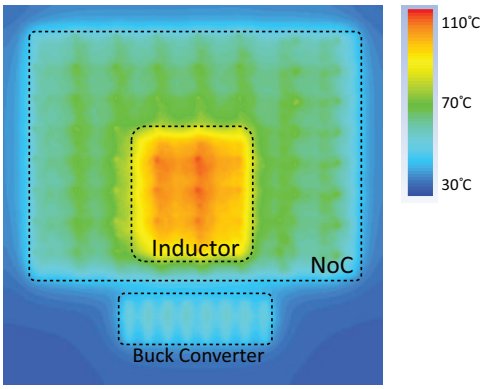

Fig. 7. Thermal map of the full chip when operating at 5W (maximum temperature: 114.96°C).

significant thermal coupling occurs between the inductor and the NoC. Therefore, this thermal coupling is our prime target for design optimization.

C. Design Optimization of 2.5-D IVR

Design block relocation is a key technique to reduce the temperature in the IVR system. We assume that the IVR chip design is fixed, but the power inductor design in the silicon interposer can be modified. The inductor cannot be placed on the periphery of the silicon interposer and should be placed beneath the chip because of the routing issues. Considering all these factors, instead of placing the inductor in the middle of the silicon interposer, we place it on the bottom right corner to avoid overlap with the NoC and the buck converter as much as possible. As shown in Fig. 8, by changing the

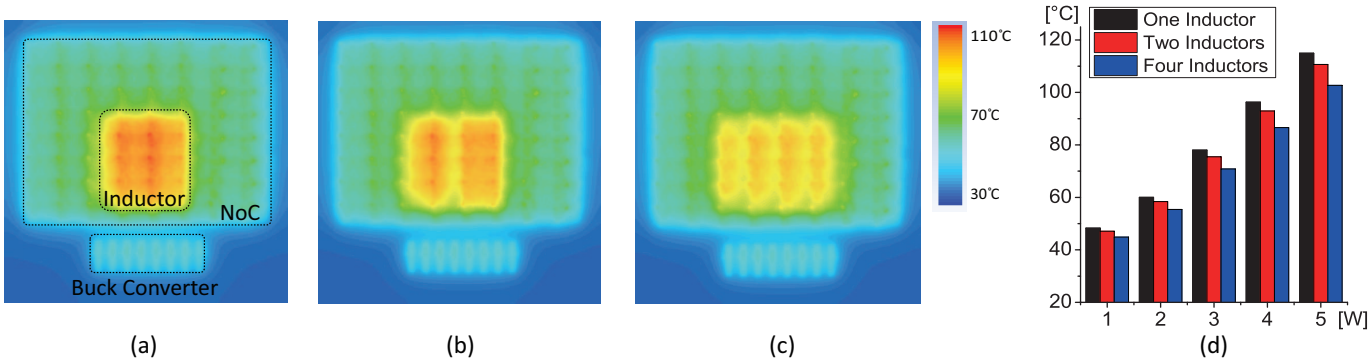

Fig. 9. (a,b,c): Temperature maps of the proposed inductor spreading technique, (d) inductor spreading results. (a) One inductor set, (b) two inductor sets, and (c) four inductor sets.

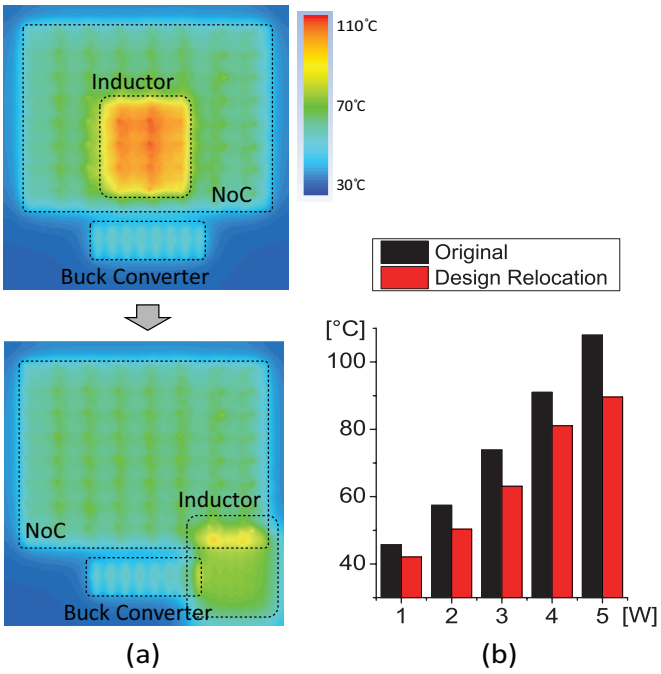

Fig. 8. Proposed design block relocation technique. (a) Inductor relocation to minimize the overlap, (b) design block relocation results.

location of the inductor, we obtain a significant temperature reduction. The temperature maps based on different inductor placement are shown in Fig. 8 (a), and the temperature results of this technique with varied power consumptions are shown in Fig. 8 (b). By minimizing the overlap between these design blocks, we achieve a maximum of 18.41°C reduction in temperature. If we can fully avoid the overlap between the inductor and other design blocks, the temperature will decrease further.

In the IVR design, one set of eight coupled power inductors is placed in the silicon interposer. Assuming that the location of the inductors is fixed but the design of inductors can be modified, spreading the inductors to a larger footprint is an effective method to reduce heat coupling and thus temperature. Therefore, we propose an inductor spreading technique, where we spread the eight inductors into two sets of four coupled inductor or four sets of two coupled inductor. We fix the distance between the inductor sets to 200μm, and we compute the maximum temperature. Fig. 9(a)-(c) show the thermal map of inductor spreading, and Fig. 9 (d) shows the temperature reduction with varied power consumptions. By spreading the inductor from one set to four sets, we observe a clear trend of temperature decrease. In particular, we achieve a maximum of 12.27°C temperature reduction by spreading the inductors.

IV. Conclusions

In this work we proposed a thermal analysis methodology for silicon interposer based mixed-signal 2.5D integrated voltage regular designs. We observed that the power inductor is the hottest component in the system, and the temperature of the inductor alone rises up to 77.3°C. When the IVR generates $5W$, the maximum temperature of the system rises up to 114.96°C because of the heat coupling between the power inductor inside the silicon interposer and the NoC module in the chip. We proposed two design techniques for thermal optimization: design block relocation and inductor spreading. These techniques led to 18.41°C and 12.27°C temperature reduction, respectively.

References

[1] S. A. et al, "A multi-stage interleaved synchronous buck converter with integrated output filter in a 0.18/spl mu/ sige process," in *Solid-State Circuits Conference, 2006. ISSCC 2006. Digest of Technical Papers. IEEE International*, feb. 2006, pp. 1398 –1407.

[2] P. H. et al, "A 233-mhz 80%-87% efficient four-phase dc-dc converter utilizing air-core inductors on package," *Solid-State Circuits, IEEE Journal of*, vol. 40, no. 4, pp. 838 – 845, april 2005.

[3] W. K. et al, "System level analysis of fast, per-core dvfs using on-chip switching regulators," in *High Performance Computer Architecture, 2008. HPCA 2008. IEEE 14th International Symposium on*, Feb. 2008, pp. 123 –134.

[4] N. S. et al, "A 2.5d integrated voltage regulator using coupled-magnetic-core inductors on silicon interposer delivering 10.8a/mm2," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2012 IEEE International*, feb. 2012, pp. 400 –402.

[5] S. P. et al, "A new methodology for ic-package thermal co-analysis in 3d ic environment," in *Electrical Design of Advanced Packaging Systems Symposium (EDAPS), 2010 IEEE*, dec. 2010, pp. 1 –4.

[6] M.-C. H. et al, "Thermo-mechanical simulative study for 3d vertical stacked ic packages with spacer structures," in *Semiconductor Thermal Measurement and Management Symposium, 2010. SEMI-THERM 2010. 26th Annual IEEE*, feb. 2010, pp. 47 –54.

[7] C. X. et al, "Fast 3-d thermal analysis of complex interconnect structures using electrical modeling and simulation methodologies," in *Proc. IEEE Int. Conf. on Computer-Aided Design*, nov. 2009, pp. 658 –665.

[8] N. A. et al, "Thermalscope: Multi-scale thermal analysis for nanometer-scale integrated circuits," in *Proc. IEEE Int. Conf. on Computer-Aided Design*, nov. 2008, pp. 603 –610.