Thermal-Power Delivery Network Co-Analysis for Multi-Die Integration

Md Obaidul Hossen
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA, 30332, US
mhossen3@gatech.edu

Yang Zhang
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA, 30332, US
steven.zhang@gatech.edu

Muhannad S. Bakir
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA, 30332, US
mbakir@ece.gatech.edu

***Abstract*—In this paper, we present a thermal-power delivery network (PDN) co-analysis framework to analyze various multi-die integration schemes. In the proposed approach, we capture the interdependencies between temperature distribution of the dice in a package and the supply voltage noise. We use standalone thermal and PDN analyses as references to compare our co-analysis results. Using a multi-die package and a bridge-based 2.5-D package case studies, our analysis shows a 10-12% over-estimation in steady-state temperature and power supply noise.**

***Index Terms*—2.5-D/3-D interconnects and packages, heterogeneous integration, power distribution networks**

I. INTRODUCTION

The pervasive nature of electronics has pushed the need for ever more heterogeneous integration technologies (2.5-D/3-D ICs), which provide high-bandwidth density and low-energy connectivity as well as ultra-small form factors. Fig. 1 presents some of the key heterogeneous integration technologies including interposer/bridge 2.5-D ICs [1], [2], 3-D ICs [3], and fan-out wafer-level based packages including package-on-package (PoP) technology [4] . However, owing

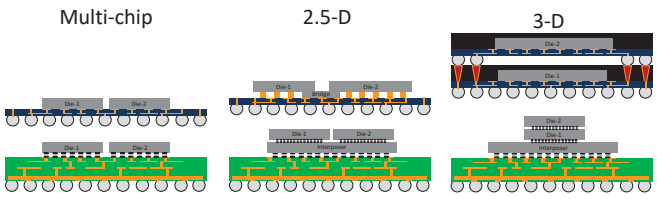

Fig. 1. Evolution of packaging technologies

to these advanced technologies, the total power density is expected to increase beyond 100 W/cm^2 [5]; power delivery becomes a critical challenge, and advanced cooling solutions (for example, microfluidic cooling) are turning into a necessity [6]. Fig. 2 shows the dependencies between power dissipation, temperature, and power delivery network (PDN). The temperature impacts the leakage power and the grid resistivity of the PDN [3]. Conversely, the power supply voltage impacts both leakage and dynamic power [3]. Without considering the interactions between each of the components in Fig. 2 for emerging architectures with increased power density, the results from the standalone or partially integrated models could be overestimated.

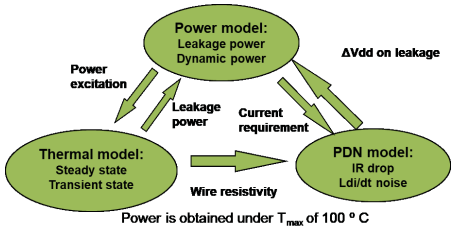

Fig. 2. Thermal-PDN interaction models

In previous efforts [1], [3], we benchmarked our PDN and thermal models to open source IBM benchmarks and finite element based modeling using ANSYS, respectively. Moreover, we presented the PDN results for different 2.5-D integration technologies in [1] and thermal-PDN co-analysis results for 3-D stacked ICs in [3]. In this paper, we present a complete thermal-PDN co-analysis framework for multi-die packages and bridge-based technologies [2]. The rest of the paper is organized as follows: Section II presents the modeling framework. Section III presents the PDN benchmarking results for a multi-die package and a bridge-based 2.5-D package using an approach similar to that shown in [2]. Section IV reports the results from a thermal-PDN co-analysis perspective.

II. MODELING FRAMEWORK

In Fig. 3, we present the proposed modeling framework for steady-state analysis. We begin thermal and PDN simulations with a reference power for each die estimated from an architectural tool [7]. Moreover, we use HSPICE to estimate the temperature and supply voltage dependencies of the leakage power. In the subsequent iterations, the power dissipation is updated by the power models that use the updated temperature and supply voltage values. At the end of the simulations, the power dissipation, temperature distribution, and the supply noise of each die become consistent with each other within our interaction models [3]. We consider two different thermal effects as shown in the figure. First, the power estimation of a die from an architectural tool or HSPICE simulations

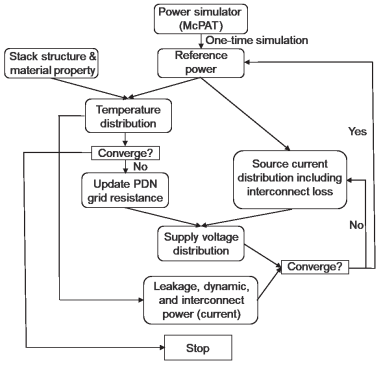

Fig. 3. The flow chart for the thermal-PDN co-analysis

TABLE I
GENERAL PARAMETERS FOR PDN MODEL

On-die global wire **P**itch/**W**idth/**T**hickness (μm)	39.5/17.5/7
On-die decap density (nF / mm^2)	3.35
C4 bump diameter/pitch (μm)	60/130
BGA inner diameter/outer diameter/pitch (μm)	250/300/1000
PCB R/L ($\mu\Omega$/pH)	166/21

is temperature dependent; the outer path in Fig. 3 accounts for this effect. Second, there is self-heating of the PDN where temperature changes the PDN resistivity. Additionally, we included a distributed package model in our co-analysis framework to incorporate irregular packaging structures owing to emerging advanced packaging technologies. In our two die package, Die #1 and Die #2 emulate a 14 nm FPGA die with peak total power of 44.8 W [8] and a 22 nm processor die with peak total power of 74.49 W [1], respectively. We assume uniform power map for both dice with a supply voltage of 0.9 V. Both dice are assumed to be 1 $cm \times 1$ cm and are placed side-by-side with a die spacing of 0.5 mm. For the bridge-based configuration, we assume a 2.5 $mm \times 6$ mm bridge interconnecting the dice. The framework is implemented in MATLAB.

III. PDN BENCHMARKING

Table I presents the PDN specifications for this analysis. For the bridge-based configuration, we assume that the bridge-chip is TSV-less and only contains signaling links for chip-to-chip communication. Hence, the periphery of the dice interconnected by the bridge does not have direct access to the package power/ground planes. It is evident from Fig. 4 that owing to the absence of direct access between the package PDN and the periphery of the chips where the bridge-chip is located, the maximum IR-drop in the CPU die and the FPGA die is 102 mV and 54.5 mV, respectively. Compared to the multi-die package, based on our assumptions, the bridge-based configuration provides 68% and 46% increased IR-drop for the CPU die and the FPGA die, respectively. Some potential solutions to reduce such higher supply noise include putting the critical circuit blocks away from these regions, bridge-chip splitting into multiple smaller bridges [1], adding TSVs in the bridge-chip, and package-level redistribution layers to reroute the PDN in the periphery of the chip.

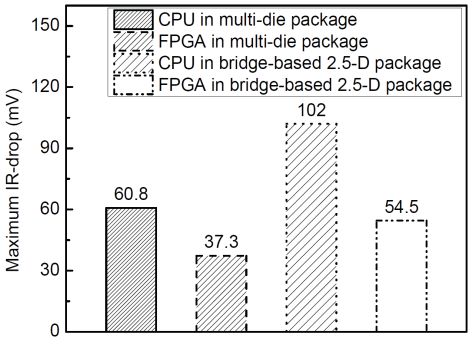

Fig. 4. Steady-state IR-drop results for the comparison of a multi-chip package and a bridge-based 2.5-D package

IV. THERMAL-PDN ANALYSIS

In this section, we analyze the thermal-PDN interactions of different configurations. Table II summarizes the specifications

TABLE II
THERMAL SIMULATION PARAMETERS

Layer	Conductivity (W/mK)		Thickness (μm)
	In-plane	Through-plane	
TIM	3		30
Heat spreader	400		1000
CPU and FPGA	149		100
Microbump and ILD	1.6		40
Package	30.4	0.38	1000

of the thermal simulations. We assume that the system uses air cooled heat sinks and the case-to-ambient thermal conductance is 0.218 W/K. The secondary heat path is through the PCB. We use an effective heat transfer coefficient of 311 W/m^{2}K as the boundary condition at this interface. The ambient temperature is assumed to be 38^{o}C. The configuration consists of two thermal interface material (TIM) layers: one between the heat sink and the heat spreader, the other between the heat spreader and the dice. We assume a common heat spreader to both dice.

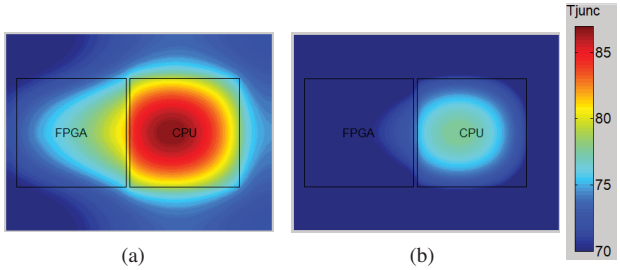

Fig. 5. The temperature distribution for (a) standalone model, and (b) co-analysis model in multi-chip packages

Fig. 5 presents the temperature distribution from thermal-PDN co-analysis for a multi-die package for our two die

system. Fig. 5(a) shows the thermal results from a standalone simulation assuming an ideal supply voltage. The maximum temperature of the CPU and the FPGA dice is 88^{o}C and 81.3^{o}C, respectively. Likewise, Fig. 5(b) presents the temperature distribution accounting all the interactions between the thermal and the PDN simulations. In this scenario, the maximum temperature of the CPU die and the FPGA die is 78.7^{o}C and 73.2^{o}C, respectively. Hence, we see that the standalone thermal simulation overestimates the maximum temperature by 11.3% and 10% for the CPU die and the FPGA die, respectively. Fig. 6 presents the results for a bridge-based

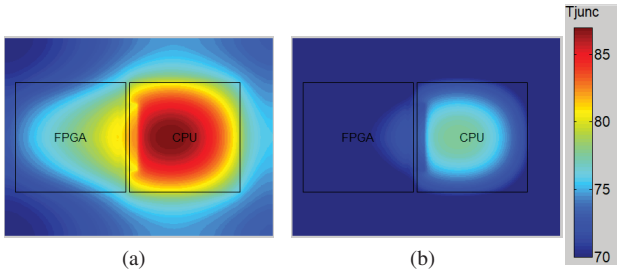

Fig. 6. The temperature distribution for (a) standalone model, and (b) co-analysis model in bridge-based 2.5-D packages

configuration for our two die system. Since there is a silicon bridge interconnecting the dice on the package, there are two thermal coupling pathways from the 'hotter' die to the 'cooler' die (in this case, CPU die to FPGA die). However, since the heat sink is sitting atop the heat spreader, the primary thermal coupling path remains through the heat spreader. Hence, the temperature map is similar to the one observed for our multi-die package.

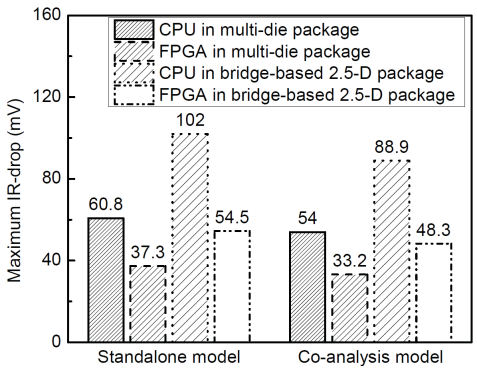

Fig. 7. Steady-state IR-drop comparison for different configurations

Finally, in Fig. 7, we summarize the steady-state IR-drop results from these two configurations. The first half of the figure is the same as shown in Fig. 4 and is included for clarity. As stated previously, the leakage power is dependent on both the temperature of the die and the supply voltage. In each iteration of the analysis, we use a fitting function to determine the effective leakage power. We assume a worst case temperature as our initial condition (100^{o}C). However, since the temperature is lower than the initial condition, the estimated leakage power decreases. Likewise, our dynamic power estimation is based on a perfect supply voltage. When we incorporate the supply voltage fluctuations, the overall estimated power decreases. Moreover, the resistivity of the metal layers in the PDN is temperature dependent. Hence, in Fig. 7, we see that for both the multi-die package and the bridge-based package, there is a significant overestimation in the standalone model. For the multi-die package case, compared to the standalone modeling, both the CPU die and the FPGA die overestimate the maximum IR-drop by almost 11%. For the bridge-based case, the maximum IR-drop follows a similar trend where both dice overestimate the maximum IR-drop by approximately 12%. However, compared to the multi-die package configuration, the increase in IR-drop for the bridge-based configuration is 64% and 45% for the CPU die and the FPGA die, respectively. This increase in IR-drop is similar to what we observed in the standalone models.

V. CONCLUSIONS

In this paper, we present a thermal-PDN co-analysis framework that incorporates impact of the thermal distribution of the dice on the supply voltage and vice versa. From steady-state co-analysis, we observe approximately 11% overestimation in the maximum temperature and 11-12% overestimation of the supply voltage for each die compared to the standalone models. While the standalone models can be adequate for pre-design exploration and for conventional packages, the co-analysis model provides added accuracy for 2.5-D/3-D architectures with increased power density and higher temperature gradients within and between dice.

VI. ACKNOWLEDGEMENT

This work is supported in part by Intel Corporation.

REFERENCES

[1] Y. Zhang, M. O. Hossen, and M. S. Bakir, "Power delivery network benchmarking for interposer and bridge-chip-based 2.5-d integration," *IEEE Electron Device Letters*, vol. 39, no. 1, pp. 99–102, Jan 2018.

[2] A. Podpod, J. Slabbekoorn, A. Phommahaxay, F. Duval, A. Salahouedlhadj, M. Gonzalez, K. Rebibis, R. Miller, G. Beyer, and E. Beyne, "A novel fan-out concept for ultra-high chip-to-chip interconnect density with 20-μm pitch," in *2018 IEEE 68th Electronic Components and Technology Conference*, Oct 2018.

[3] Y. Zhang and M. S. Bakir, "Integrated thermal and power delivery network co-simulation framework for single-die and multi-die assemblies," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 7, no. 3, pp. 434–443, March 2017.

[4] C. F. Tseng, C. S. Liu, C. H. Wu, and D. Yu, "Info (wafer level integrated fan-out) technology," in *2016 IEEE 66th Electronic Components and Technology Conference (ECTC)*, May 2016, pp. 1–6.

[5] ITRS, "International Technology Roadmap for Secmiconductors, 2013." [Online]. Available: http://www.itrs.net/

[6] T. E. Sarvey, Y. Zhang, C. Cheung, R. Gutala, A. Rahman, and M. S. Bakir, "Monolithic integration of a micropin-fin heat sink in a 28 nm fpga," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, Accepted, 2016.

[7] S. Li, J. H. Ahn, R. Strong, J. Brockman, D. Tullsen, and N. Jouppi, "Mcpat: An integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Proc. Annual Int. Symp. Microarchitecture*, Dec 2009, pp. 469–480.

[8] Altera, "Leveraging HyperFlex Architecture in Stratix 10 Devices to Achieve Maximum Power Reduction." [Online]. Available: https://www.altera.com/products/fpga/stratix-series/stratix-10/overview.html