Multi-Granularity Thermal Evaluation of 3D MPSoC Architectures

Alain Fourmigue, Giovanni Beltrame, Gabriela Nicolescu, El Mostapha Aboulhamid, Ian O'Connor
Ecole Polytechnique Montreal, University of Montreal, Ecole Centrale Lyon

***Abstract*—Three-dimensional (3D) integrated circuits (IC) are emerging as a viable solution to enhance the performance of Multi-processor System-On-Chip (MPSoC). The use of high-speed hardware and the increased density of 3D architectures present novel challenges concerning thermal dissipation and power management. Most approaches at power and thermal modeling use either static analytical models or slow low-level analog simulations. In this paper, we propose a novel thermal modeling methodology for evaluation of 3D MPSoCs. The integration of this methodology in a virtual platform enables effcient dynamic thermal evaluation of a chip. We present initial results for an architecture based on a 3D Network-On-Chip (NoC) interconnecting 2D processing elements (PE). Our methodology is based on the finite difference method: we perform an initial static characterization, after which high-speed dynamic simulation is possible.**

***Index Terms*—Virtual Platform, 3D IC, MPSoC, Dynamic Evaluation of Performance, Power Estimation, Thermal Analysis**

I. Introduction

Recent advances in three-dimensional (3D) integration technology offer new ways to design high performance 3D architectures of Multi-Processor System-On-Chip (MPSoC). As the first 3D integrated circuits (3D ICs) are being commercialized, the industry is progressively moving towards more complex designs based on through-silicon-vias (TSV) technology [1]. However, tools and methodologies are still needed to enable quick exploration and performance evaluation of more complex 3D architectures.

Current tools are designed for conventional 2D architectures and do not allow to model the peculiarities of 3D ICs. In particular, the research community has proposed several approaches to investigate the power and thermal behavior of 3D architectures, but most of the proposed approaches rely on static analytical models or use analog simulation and do not allow to capture dynamic behaviors under real work conditions.

For conventional 2D designs, virtual platforms have become a widely adopted solution to evaluate performance, and power and thermal dissipation using final or almost final versions of software applications [7]. A critical challenge is to extend existing virtual platforms techniques to enable dynamic evaluation of 3D architectures, integrating high level models of 3D interconnects and considering power consumption and thermal dissipation. Good simulation performance and moderate modeling efforts are also mandatory features [7].

978-3-9810801-7-9/DATE11/©2011 EDAA

We propose a novel thermal modeling methodology based on the finite difference method, which combines high simulation speed (when compared to similar methods) with a user-defined degree of accuracy. We use our methodology with a virtual platform which models a typical architecture for 3D MPSoC, based on a 3D Network-On-Chip (a Stacked Mesh NoC, one of the most promising topologies [4], [5]) interconnecting homogeneous 2D processing elements (PEs, 2D SPARC cores).

The remainder of this paper is organized as follows: Section II discusses related work; Section III formulates the problem and introduces our novel approach to perform thermal analysis; Section IV illustrates the proposed approach with an example chip which models a typical 3D MPSoC, and shows early results. Section V concludes and discusses future work.

II. Related Work

Over the past few years, extensive research has been carried out to predict the performance of 3D ICs. The proposed approaches can be classified into three categories:

- Analog simulation at the transistor level can provide very accurate thermal analysis, but is very computation-intensive [2], [3]. In fact, the performance of these techniques is insufficient for the evaluation of 3D many-core architectures in a feasible amount of time.
- Static analytical models allow for very fast evaluation and are suitable for early exploration of large design spaces [4], [5]. However, these models provide average values and do not allow to capture the dynamic behavior of a system, and are therefore ill-suited to our purpose, as we are interested is the dynamic evaluation of the architecture using targeted applications.
- Virtual platforms offer a convenient solution to capture the dynamic behavior of a SoC [7]. Several approaches have been proposed to integrate power and thermal models of conventional 2D SoC components into virtual platforms [8], [9], [10]. but very few similar tools exist for 3D designs. MEVA-3D [6] is one of these tools, but it can only model and evaluate the micro-architecture of 3D processors. To the best of our knowledge, there is no existing virtual platform with power and thermal models for 3D MPSoC architectures.

A well-known framework for thermal modeling is HotSpot [10], which allows to develop thermal models of processors (as lumped thermal RC networks) that can be incorporated in cycle-accurate simulators. This framework is

not specifically designed to address the problem of modeling 3D ICs and has a limited ability to model vertical integration. In particular, layers are connected to each other with a single vertical resistance Unlike HotSpot, our approach allows to accurately model multiple stacked layers using a refined mesh. In addition, HotSpot has to solve large matrix equations while our thermal approach is specifically designed to provide high simulation performance.

Though we use the finite difference method to perform a detailed thermal analysis of the 3D MPSoC elements, our approach differs from traditional numerical analysis: we perform a fine-grained analysis only once to identify the parts of the 3D MPSoC which can be represented with less detail without sacrifying accuracy, and then we generate a non-uniform grid model and we perform a full-chip dynamic thermal analysis with acceptable runtime.

III. Proposed Thermal Analysis Method

Macroscopically, local thermodynamic equilibrium is maintained through the chip [11] and heat transfers are governed by the heat equation:

$$\int_V \rho c_p \frac{\partial T}{\partial t} dv = \int_S \lambda_{th} \overrightarrow{\bigtriangledown T}.\vec{n}_{out} ds + \int_V P_{vol} dv \qquad (1)$$

where ρ, c_p,λ_{th} are the density, the specific heat, and the thermal conductivity of the material, respectively. Also, V is the volume of the considered section of the chip, S is the surface of said volume, $\vec{n}_{out}$ is the normal to the surface S, $T(x, y, z, t)$ is the temperature field, and P_{vol} is the considered section's power density per volume. Solving Equation 1 gives the temperature field in space and time[1]. We define the *steady-state* behavior as a state of equilibrium when temperature doesn't change in time. On the other hand the system's *transient response* corresponds to the process before the equilibrium temperature is established.

As opposed to most literature, we propose to perform a transient analysis rather than steady-state analysis. While the steady-state approximation greatly simplifies the heat equation, its assumptions may not be valid for modern MPSoC. In fact, the time constant of conductive heat transfers in the materials used for chip production is in the order of milliseconds [11], whereas the clock period in today's MPSoC is in the order of nanoseconds. Consequently, several thousands of cycles will elapse before the steady-state can prevail in the chip, while the power density distribution can change significantly, as new instructions are issued at each cycle. Especially when considering task migration, the steady-state approximation does not hold, as the processor can be suddenly put to the idle state.

A. Proposed Approach

Our approach consists of using a numerical method (the finite difference method) on a non-uniform grid to solve the non-stationary heat equation for the whole target 3D MPSoC. The finite difference method is based on the discretization in space and time of the differential operators: the chip is discretized into small increments of space Δl and the time is discretized into small increments of time Δt. This leads to Equation (2) which explicitly expresses the temperature at a node m at time $t + \Delta t$ in terms of the temperature at neighboring nodes at time t.

$$T_m^{t+\Delta t} = \frac{\lambda_{th}\Delta t}{\rho c_p (\Delta l)^2} \sum T_{neighbor}^t + (1 - \frac{6\lambda_{th}\Delta t}{\rho c_p (\Delta l)^2}) T_m^t + \frac{P_m^t \Delta t}{\rho c_p} \qquad (2)$$

This method can precisely model the heterogeneous structure of 3D circuits, also taking into account their large variety of boundary conditions. However, the computational requirements of such method are too high for modeling a complete and realistic system: even for small circuits. The simulation times might become unfeasible. To accelerate the thermal analysis we propose a smart method to determine a non-uniform mesh. This allows to greatly reduce the number of nodes to be considered, while maintaining sufficient accuracy. Our experiments show that the temperature is practically uniform in certain areas of the 3D chip. These areas with small temperature gradients are essentially due to the materials used and the power generation structure of the chip. We assume that the distribution of these areas in the system does not strongly depend on the executed applications [11], and that it can be estimated statically. On the other hand, the temperature of these areas depends on the application that is run on the system, and needs to be estimated dynamically. Using a non-uniform grid, the analysis speed is increased by removing nodes where those are not necessary: it is not useful to calculate precisely the temperature distribution in areas with small temperature gradients, but it is necessary to keep a higher node density for areas with high gradients.

In our methodology, we consider the 3D chip as a set of atomic blocks which perform elementary operations, identified using floorplan or layout information. A block can be contained in a single layer of Silicon, or includes multiple layers. As inputs for the methodology we require the power density map of each of these blocks and a set of process parameters (see Table I).

Overall, our approach consists of three phases:

1) Perform a fine-grain static thermal analysis on a small block
2) Use the collected data to generate a non-uniform mesh for each block
3) Use the new non-uniform mesh grid to perform an accurate *dynamic* thermal analysis of the full-chip

IV. Application to a Typical 3D MPSoC

A. A Typical 3D MPSoC

The architecture of the proposed example chip is based on a 3D Network-On-Chip (NoC) interconnecting 2D processor elements (PEs). The topology of the 3D NoC is a 3D Stacked Mesh topology, which has been described as a promising solution for 3D NoC [4], [5]. The PEs are conventional 2D RISC cores which have their own memories and communicate to each others through the NoC.

[1]Note that a set of boundary conditions has to be defined, in order to obtain a solution to the heat equation

The example chip is composed of four layers, and NoC connects the 4 layers with a bus spanning the vertical direction. This vertical bus is composed of Trough-Silicon-Vias (TSV), modeled on IBM's 3D process [12], with layers stacked in a face-to-back organization. Table I shows the main parameters of the 3D technology model.

TABLE I
3D PROCESS TECHNOLOGY PARAMETERS

Bonding Layer Thickness (μm)	10
Bulk Layer Thickness (μm)	48
Metal Layer Thickness (μm)	6
Signal TSV Diameter (μm)	5
Signal TSV Pitch (μm)	10
Signal TSV Aspect Ratio	5:1
Power/Ground TSV Diameter (μm)	40
Power/Ground TSV Pitch (μm)	400

First Phase: Fine-grained Static Thermal Analysis

To illustrate our approach, we show the detailed analysis of a 3D block that models a component of our platform's NoC, shown in Figure 1(a). Figure 1(b) shows the power density map of the block's layers.

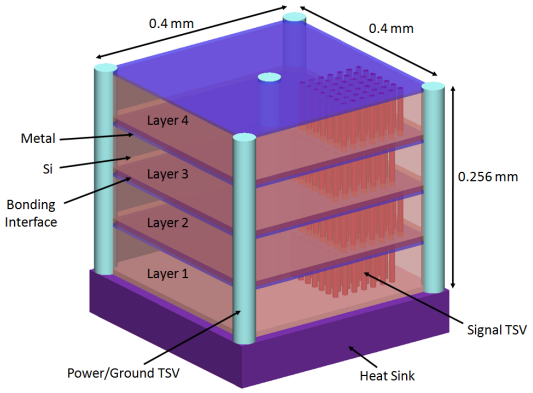

(a) 3D structure of the block

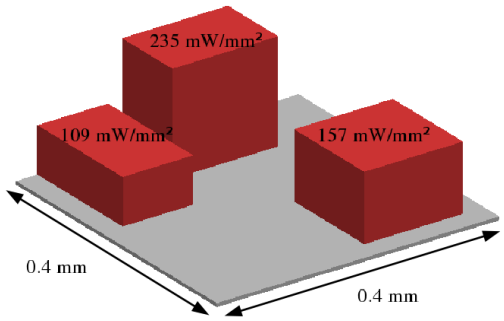

(b) Power density map of the block's layers

Fig. 1. The example 3D block

We perform a fine-grained analysis using a $256 \times 256 \times 256$ grid, which means that nodes have a distance of $1\mu m$ along the vertical axis, resulting in a detailed thermal model of the 3D structure. This first analysis is *static*, with the fixed power density map shown in Figure 1(b). It is performed only once to determine areas in which the temperature is homogeneous, meaning that its computational cost does not reduce the practical applicability of the methodology. We run the thermal simulation for $1ms$, as this is the characteristic time of the transient response [11]. Analyzing the thermal behavior for shorter periods (e.g. a clock cycle) would not be relevant, as the temperature varies slowly in the chip.

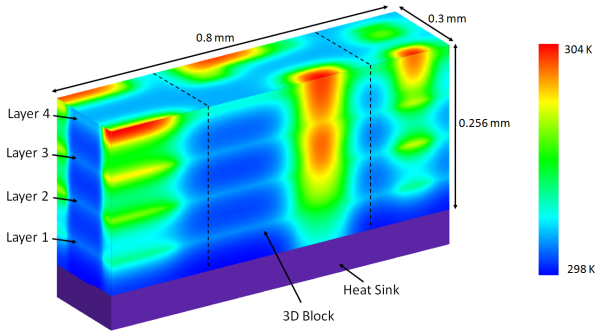

Fig. 2. Fine-grained thermal analysis (cut view after 1 ms)

It is worth noting that to compute the temperature distribution on our block ($0.4 \times 0.4 \times 0.256mm$) it is necessary to take an extended block to take into account the thermal exchanges with neighboring elements. The size of this extended block depends on the thermal conductivity of the rest of the chip. The distance traveled by a heat wave in $1ms$ can be approximated using Equation 1, and it is given by $L = \sqrt{\frac{\lambda_{th} \times t}{\rho c_p}}$, which gives $L \simeq 0.2mm$ in our case. It is not necessary to simulate a block with an extension larger than L, since the heat generated at its boundaries will not be able to reach the original block within the $1ms$ simulation time. Figure 2 shows the resulting temperature distribution inside the considered 3D block after $1ms$, which measures $0.8mm$ in width for this reason. The vertical and the top surfaces of the block are assumed to be adiabatic, while the bottom surface is assumed to be convective [10].

B. Second Phase: Non-Uniform Mesh Generation

Using a coarser mesh means increasing the simulation speed, but also losing accuracy. To generate our non-uniform grid we start by defining the maximum admissible error (when compared with the fine-grained analysis) for the temperature distribution. We use the standard Euclidean norm to compute the distance between two temperature distributions. To verify that the error introduced by a coarse-grained grid is sufficiently small we use the following equation:

$$\sum_{i \in Mesh} (T_i^{coarse\_grain} - T_i^{fine\_grain})^2 \leq N^2 \varepsilon \qquad (3)$$

where N is the number of nodes in the fine-grain grid. This equation allows the user to configure the desired level of accuracy.

The generation of the non-uniform mesh is achieved in two steps:

1) A coarse-grained uniform mesh (Figure 3(a)) is generated from the mesh used for the fine-grained analysis. All nodes in the fine-grained grid with small temperature

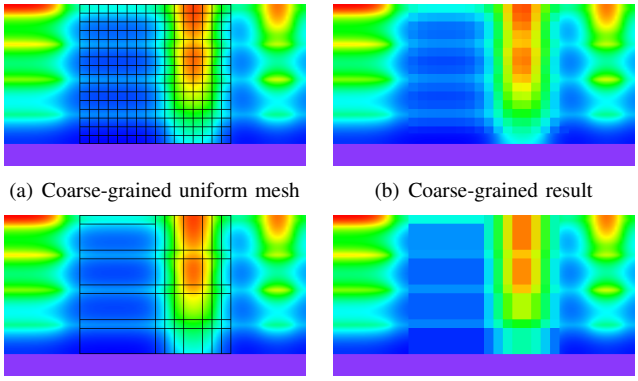

(a) Coarse-grained uniform mesh (b) Coarse-grained result

(c) Non-uniform mesh (d) Non-uniform result

Fig. 3. Generation of a non-uniform grid

gradients are merged into a single node in the coarse-grained mesh. Then, the temperature distributions of the fine- and coarse-grained meshes (Figure 3(a)) are compared to control verify the required error margin, and the coarsest uniform mesh for which Equation (3) holds is selected.

2) A non-uniform mesh (Figure 3(a)) is generated from the result of the first step. The error margin is controlled verifying that Equation (3) holds for the whole mesh.

C. Third Phase: System-Level Dynamic Thermal Analysis

In order to complete the dynamic thermal analysis for the overall MPSoC system, a global non-uniform grid is generated from the uniform grids corresponding to each block.

The proposed thermal analysis was implemented in a virtual platform (implementing Instruction-Set Simulators for the PEs' architecture) modeling our sample 3D MPSoC. An image processing application, Motion Compensated De-Interlacing (MCDI), was executed on the virtual platform. The MCDI application was programmed to be executed in parallel on 6 PEs in a pipelined architecture (task-level parallelism). The size of the processed images is 70×55 pixels and each PE execute 160K instructions to process an image. The results obtained with a non-uniform and an uniform grid are presented in Figure 4.

Figure 4(a) shows the simulation performance for the evaluation of circuits of different sizes. Simulation times are given for the uniform and non-uniform meshes, at the same accuracy level. Figure 4(b) shows instead the accuracy obtained for different meshes. To evaluate the global accuracy, the temperature distribution obtained using the fine-grained mesh was used as reference.

V. Conclusion and Future Work

In this paper, we propose an efficent approach to perform the transient thermal analysis of 3D MPSoC with acceptable accuracy and acceptable runtime. Our approach is based on the static pre-characterization of a device, that can then be simulated at high speed while maintaining accuracy. Future work will focus on refining the technology model to take into account the the insertion of thermal vias and its effect

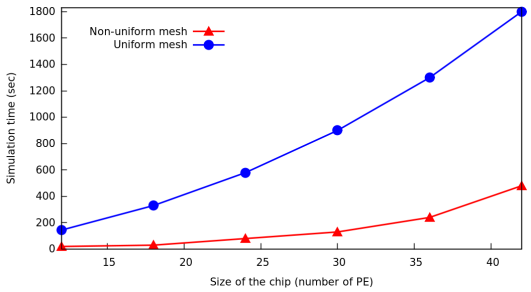

(a) Speed comparison at the same accuracy level

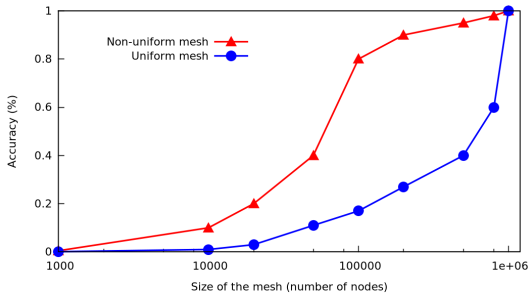

(b) Accuracy comparison at the same simulation speed

Fig. 4. Speed and Accuracy

on the temperature in the 3D SoC. We will evaluate the proposed methodology using more complete testbenches and we will investigate the use of parallel programming techniques to further accelerate the thermal analysis.

References

[1] G.H. Loh and Y. Xie, "3D Stacked Microprocessor: Are We There Yet?", *IEEE Micro.*, vol.30, pp.60-64, May-June 2010.

[2] T-Y. Chiang, K. Banerjee, and K. C. Saraswat, "Compact Modeling and SPICE-Based Simulation for Electrothermal Analysis of Multilevel ULSI Interconnects". In *Proc. ICCAD*, 2001, pp. 165-172.

[3] K. Puttaswamy and G.H. Loh, "Thermal Analysis of a 3D Die-Stacked High-Performance Microprocessor". In *Proc. of the ACM Great Lakes Symp. on VLSI*, 2006, pp. 19-24.

[4] V.F. Pavlidis and E.G. Friedman, "3-D Topologies for Networks-On-Chip", *IEEE Trans. Very Large Scale Integr. (VLSI) Systems*, vol. 15, pp. 1081-1090, Oct. 2007.

[5] B.S. Feero, and P.P. Pande, "Networks-on-Chip in a Three-Dimensional Environment: A Performance Evaluation", *IEEE Trans. Comput.*, vol. 58, pp. 32-45, Jan. 2009.

[6] J. Cong, A. Jagannathan, Y. Ma, G. Reinman, J. Wei and Y. Zhang "An Automated Design Flow for 3D Microarchitecture Evaluation". In *Proc. ASP-DAC*, 2006, pp. 384-389.

[7] B. Bailey and G. Martin, *ESL Models and Their Application: Electronic System Level Design and Verification in Practice (Embedded Systems)*, Springer, 2009

[8] C. Brandolese, F. Salice, W. Fornaciari and D. Sciuto, "Static power modeling of 32-bit microprocessors", *IEEE Trans. Comput.-Aided Des. Integr. Syst.*, vol. 21, pp. 1306-1316, Dec. 2002.

[9] A. Bona, V. Zaccaria, and R. Zafalon, "System Level Power Modeling and Simulation of High-End Industrial Network-on-Chip". In *Proc. DATE*, 2004, pp. 318-323.

[10] W. Huang, S. Ghosh, S. Velusamy and K. Sankaranarayanan, "HotSpot: A Compact Thermal Modeling Methodology for Early-Stage VLSI Design", *IEEE Trans. Very Large Scale Integr. (VLSI) Systems*, vol. 14, pp. 501-513, May 2006.

[11] S.S. Sapatnekar, "Addressing Thermal and Power Delivery Bottlenecks in 3D Circuits", In *Proc. ASP-DAC*, 2009, pp. 423-428.

[12] Anna W. Topol, Steven J. Koester, Douglas C. La Tulipe and Albert M. Young, "3D Fabrication Options for High-Performance CMOS Technology", Wafer Level 3-D ICs Process Technology, Springer US, 2008