



7-9 October 2009, Leuven, Belgium

# Emulation-Based Transient Thermal Modeling of 2D/3D Systems-on-Chip with Active Cooling

David Atienza

Embedded Systems Laboratory (ESL), EPFL  
ESL-IEL-STI-EPFL, Station 11, CH1015, Lausanne, Switzerland  
E-mail: david.atienza@epfl.ch

**Abstract**—New tendencies envisage 2D/3D Multi-Processor System-On-Chip (MPSoC) as a promising solution for the consumer electronics market. MPSoCs are complex to design, as they must execute multiple applications (games, video), while meeting additional design constraints (energy consumption, time-to-market, etc.). Moreover, the rise of temperature in the die for MPSoCs, especially for forthcoming 3D chips, can seriously affect their final performance and reliability. In this context, transient thermal modeling is a key challenge to study the accelerated thermal problems of MPSoC designs, as well as to validate the benefits of active cooling techniques (e.g., liquid cooling), combined with other state-of-the-art methods (e.g., dynamic frequency and voltage scaling), as a solution to overcome run-time thermal runaway.

In this paper, I present a novel approach for fast transient thermal modeling and analysis of 2D/3D MPSoCs with active cooling, which relies on the exploitation of combined hardware-software emulation and linear thermal models for liquid flow. The proposed framework uses FPGA emulation as the key element to model the hardware components of 2D/3D MPSoC platforms at multi-megahertz speeds, while running real-life software multimedia applications. This framework automatically extracts detailed system statistics that are used as input to a scalable software thermal library, using different ordinary differential equation solvers, running in a host computer. This library calculates at run-time the temperature of on-chip components, based on the collected statistics from the emulated system and the final floorplan of the 2D/3D MPSoC. This approach creates a close-loop thermal emulation system that allows MPSoC designers to validate different hardware- and software-based thermal management approaches, including liquid cooling injection control, under transient and dynamic thermal maps. The experimental results with 2D/3D MPSoCs illustrate speed-ups of more than three orders of magnitude compared to cycle-accurate MPSoC thermal simulators, at the same time as preserving the accuracy of the estimated temperature within 3% of traditional approaches using finite-element simulations for 3D stacks and liquid cooling.

**Keywords** – Thermal modeling, transient analysis, FPGA emulation, 2D/3D MPSoC, active cooling, close-loop systems

## I. INTRODUCTION

The power density of high performance systems continues to increase with every process technology generation. Nowadays, several commercial multi-processor system-on-chip (MPSoC)

architectures are available several tens of cores, such as IBM's Cell [1], Sun's Niagara T1 [2] and Tiler's 64-core architecture [3]. However, in these new MPSoC architectures, power density increases the operating temperature and creates significant hot-spots on the die that need to be managed. Furthermore, 3D stacking is an emerging solution to increase the integration capabilities and frequency of forthcoming MPSoCs [4,5], but it substantially increases further power density due to the placement of computational units on top of each other. Therefore, temperature-induced problems exacerbate in 3D systems and are a major concern to be explored as early as possible in 3D MPSoC design and integration.

To explore the hardware/software (HW/SW) thermal interaction, cycle-accurate MPSoC simulators including SW thermal models exist, based on post-processing of run-time power consumption and floorplanning information [6, 7, 8]. However, these complex SW environments are very limited in performance (i.e., up to 100 KHz) due to signal management overhead and are not interactive with thermal control systems in real-time. Thus, they are not suitable for thermal control exploration in 2D/3D MPSoCs running complex real-life applications. Moreover, higher abstraction levels simulators attain faster simulation speeds, but lose significantly the accuracy for fine-grained thermal-aware architectural tuning or thermal modeling.

One alternative to cycle-accurate simulators is HW emulation. Various MPSoC emulation frameworks have been proposed [9, 10, 11]. Nevertheless, they are not designed for thermal exploration and are usually very expensive for (between \$100K and \$1M) and not flexible enough for MPSoC architecture exploration since their baseline architectures (e.g. processing cores or interconnections) are proprietary, not permitting internal changes. Furthermore, no flexible interconnection interfaces between HW emulation and also no fast thermal libraries that model active cooling behavior (e.g., liquid cooling [12, 13]) exist nowadays. Thus, thermal effects can only be verified in the last phases of the design process, typically when the final architecture and cooling components are available can be tested in the final system integration process, which can typically result in very expensive MPSoCs redesigns.

As a result, one major design challenge is the deployment of fast exploration methods of multiple HW and SW implementation alternatives for 2D and 3D MPSoCs with



7-9 October 2009, Leuven, Belgium

accurate estimations (e.g. performance, energy) that address the modeling of transient thermal behavior to tune the final 2D/3D MPSoC architectures.

In this paper I present a new HW/SW FPGA-based emulation framework of the 2D/3D MPSoC architectures, which enables realistic thermal studies in an early stage of system integration, including active (liquid) cooling modeling, as well as power, energy and performance constraints validation in real-time. First, the HW components of 2D/3D MPSoC components are mapped on an FPGA-based framework and statistics are extracted from three key MPSoC architectural levels (processors, memory subsystem and interconnections), while real-life applications are executed. Second, this run-time information is sent using a standard Ethernet connection to a dynamically adaptable SW thermal modeling tool running on a host PC. Third, this tool evaluates in real-time the thermal behaviour of the final MPSoC design, selecting different ordinary differential equation solvers according to the desired accuracy in thermal exploration and simulation time of 2D/3D chip stacks, and returns this information to the FPGA emulating the MPSoC design. This final step creates a closed-loop thermal simulation environment for 2D and 3D chips that enables testing temperature management strategies in real-time.

The experiments results with 2D/3D MPSoCs, using real-life case studies models of the UltraSPARC T1 [2] and other industrial platforms from Freescale [14], Philips [15], etc., show that this HW/SW emulation framework for transient thermal analysis can achieve speed-ups of more than three orders of magnitude compared to state-of-the-art cycle-accurate 2D/3D MPSoC thermal simulators, while keeping the accuracy of uncertainty levels of the simulated temperature obtained with the proposed method within 3% with respect to finite-element simulations.

The remainder of this paper is structured as follows. It starts in Section II with a detailed overview of prior art in thermal modeling and architectural simulation for 2D and 3D MPSoCs. Then, in Section III it is presented the proposed HW/SW thermal emulation flow for 2D/3D MPSoCs. Next, in Section IV it is described the 3D liquid cooling model. After that, in Section V we present the experimental setup and results with different 2D and 3D MPSoCs. Finally, in Section VI, I summarize the main conclusions of this work.

## II. RELATED WORK

It is widely accepted that 2D/3D MPSoCs represent a promising solution for forthcoming complex processing systems [18]. This has spurred research on modeling and prototyping MPSoC designs, using both HW and SW. From the SW viewpoint, solutions have been suggested at different abstraction levels, enabling tradeoffs between simulation speed and accuracy. First, fast analytical models have been proposed to prune very distinct design options using high-level languages (e.g., C or C++) [19]. Also, full system simulators, like Symics [20] and others [7, 8], have been developed for embedded SW debugging and can reach megahertz speeds, but are not able to accurately capture performance and power effects (e.g., at the interconnection level) depending on the cycle-accurate behavior of the HW. Second, transaction-level

modeling in SystemC, in academic [21] and industrial context [22, 23] has enabled more accuracy in system-level simulation at the cost of sacrificing simulation speed (about 100–200 KHz). Such speeds render unfeasible the transient testing of large systems due to overly long simulation times, conversely to the proposed 2D/3D thermal emulation framework. Moreover, in most cases SW simulations are only limited to a number of proprietary interfaces.

Finally, important research has been done to obtain cycle-accurate frameworks in low-level SystemC or HDL languages. Companies and universities have developed cycle-accurate simulators using post-synthesis libraries from HW vendors [27, 28]. However, their simulation speeds (10–120 KHz) are unsuitable for long MPSoC thermal exploration.

The most important alternative nowadays to MPSoC simulation is HW emulation. In industry, one of the most complete sets of statistics is provided by Palladium II [9], which can accommodate very complex systems (i.e., up to 256 Mgate). However, its main disadvantages are its operation frequency (approximately 1.6 MHz) and cost (around \$1 million). Then, ASIC integrator [10] is much faster for MPSoC architectural exploration. Nevertheless, its major drawback is its limitation to only up to few ARM-based cores and only AMBA interconnects. The same exploration limitation of proprietary cores occurs with Heron SoC emulation [24] and Zebu-XL [11], both based on multi-FPGA emulation in the order of MHz. They can be used to validate intellectual property blocks, but are not flexible enough for fast MPSoC design exploration or detailed statistics extraction. In the academic world, a recent emulation platform for exploring performance of MPSoC alternatives is TC4SOC [25]. It uses a proprietary 32-bit VLIW core and enables exploration of interconnects by using an FPGA, but, it does not enable detailed extraction of statistics and performing thermal modeling at the other three architectural levels proposed in this work, i.e., memory hierarchy, interconnects and processing cores. Finally, an interesting approach that uses FPGA prototyping to speed-up co-verification of pure SW simulators is described in [26], which uses a cycle-by-cycle synchronization basis with the C/C++ SW part by using an array of shared registers in the FPGA that can be accessed by both sides at a speed of 1 MHz, outlining the potential benefits of combined HW-SW frameworks which it is exploited in the proposed approach to reach an MPSoC emulation speed of hundreds of MHzs.

Turning our attention to thermal modeling, [6] presented a thermal/power model for 2D super-scalar architectures. It can predict the temperature variations between the different components of a processor and show the expected influence in performance. Additionally, [14, 17] have investigated the impact of temperature and voltage variation across the die of 2D and 3D MPSoCs. Their results show that the temperature can vary by more than 25 degrees across the die and tiers. In all these works, a “1D” approximation is often assumed to evaluate the thermal behavior [29, 30]. This means that the power is uniformly produced on active levels (or on parts of them), one per stratum. This assumption may lead to strongly underestimated maximum temperature. Thus, several authors [31] use this simplification but perform detailed simulation of

3D thermal effects due to the presence and localization of supervias, and analyze the local (3D) and global (1D) modeling contribution to the maximum temperature, showing that thermal resistance can be higher than 1D thermal resistance due to local 3D effects and even more fine-grain transient analysis need to be performed to avoid thermal overestimations.

Finally, numerical thermal simulations have been carried out to convert power dissipation distribution into a temperature distribution in a 3D IC [32]. Based on the past work, the development of a fundamental analytical model for heat transport in 3D integrated circuits is highly desirable. Such an analytical model provides a framework in which to analyze the general problem of heat dissipation in 3D ICs, and will enable simple thermal design guidelines.

A key component of the 3D technology is the through-silicon via (TSV) that enables communication between the two dies as well as with the package. Several works have analyzed the optimization of placement of TSVs for heat dissipation in 3D ICs [31, 32]. Other works [33] propose analytical and finite-element models of heat transfer in 3D electronic circuits and use this model to analyze the impact of various geometric parameters and thermo-physical properties (through silicon vias, inter-die bonding layers, etc.) on thermal performance of a 3D IC, but at the cost of very time-consuming simulations. Overall, all of these works prove the importance of hot spots in 2D high-performance multi-core systems (and even more in 3D structures), as well as the need of accurate and fast transient temperature analysis tools for the different architectural components of MPSoCs (cores, TSVs, etc.). Thus, the proposed emulation method aims at estimating accurately the transient temperature of integrated circuits implementing 2D/3D MPSoCs, including active cooling mechanisms (e.g., liquid cooling).

### III. HW/SW THERMAL EMULATION FLOW FOR MPSoCs

In Figure 1 it is depicted an overview of the instantiation of the proposed HW/SW thermal emulation environment for a Freescale-based 3-core MPSoC [14], implemented onto a Xilinx Virtex-V FPGA, modeling the transient thermal behavior of the system while executing multiple multimedia applications (e.g., SW-defined radio, video streaming, etc.) and a multi-processor *operating system* (OS). The system can be scaled to any number of cores sub-systems by using appropriate FPGAs.

The instantiation flow of the proposed HW/SW transient thermal emulation environment is created in four steps. First, the HW part of the MPSoC emulator is defined. It implies synthesizing the MPSoC architecture onto a certain FPGA target technology (a multi-FPGA environment exists, if it is necessary). Second, specific hardware sniffers are included in the FPGA that monitor particular signals of each component of the target 2D/3D MPSoC architecture [15]. The purpose of this step is to define a very fast method to extract the switching activity of MPSoC components, while being transparent to the normal MPSoC operation; thus, the power extraction method does not interfere or alter the actual run-time power consumption, conversely to including SW profiling to extract the execution statistics of the target MPSoC.

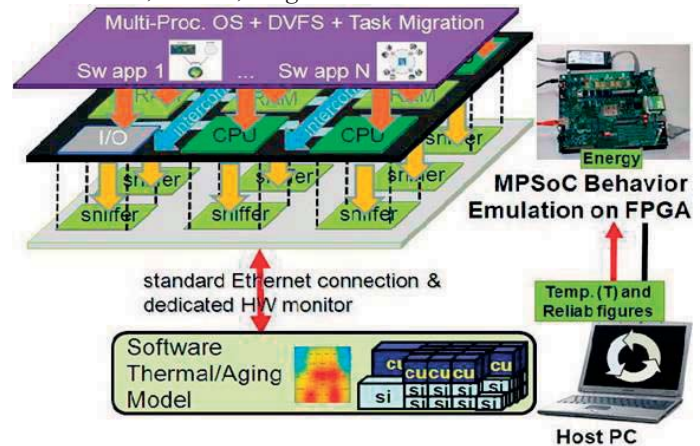


Fig. 1. Overview of the 2D/3D MPSoC thermal emulation framework

In the third step, the run-time power information is sent using a standard Ethernet connection to a dynamically adaptable SW thermal modeling tool running on a host PC. This tool evaluates in real-time the thermal behaviour of the final 2D/3D MPSoC design using a thermal model developed for bulk silicon chip systems [15], and calculates the temperature of each cell according to the floorplan of the emulated MPSoC, the frequency/voltage of each MPSoC component at run-time, as well as the specific leakage power at run-time for each component in the MPSoC. It includes different types of ordinary differential equation solvers [16] (Forward Euler 1<sup>st</sup> order, Crank-Nicholson 2<sup>nd</sup> order method, Runge-Kutta 4<sup>th</sup> order method, etc.), which enables multiple trade-offs between accuracy and thermal modeling time in 2D/3D chip stacks.

Finally, in the fourth step of the thermal emulation flow, the temperatures calculated by the SW thermal library are sent back to the FPGA emulating the MPSoC system (see link between the host PC and the FPGA on the right side of Figure 1) and are stored in registers of the FPGA that emulate the presence of thermal sensors in the target MPSoC in certain positions of the floorplan.

This final mechanism provides real-time temperature information visible by the running multi-processor OS on the modeled 2D/3D MPSoC, as the registers storing the predicted temperature are memory-mapped in a restricted position of the memory hierarchy visible only by the OS of the MPSoC. Then, the emulated temperature sensors are updated by the thermal monitoring subsystem in regular intervals, typically configurable in the range of 10 ms to 1s, according to the system designer interest. Thus, thanks to a handshake mechanism between the thermal model and the multi-processor OS middleware to synchronize the upload/download of temperatures, our extended framework implements a closed-loop thermal monitoring system, which enables exploring the impact of thermal control mechanisms in the transient thermal behavior of 2D/3D MPSoCs at multi-megahertz speeds.

### IV. ACTIVE COOLING MODEL FOR 2D/3D MPSoC

Modeling of the 3D stacked architecture with liquid cooling can be accomplished in three steps: (i) defining a grid-level thermal RC network of 2D/3D chip stacks, (ii) adding models



for the interlayer material and TSVs distribution and (iii) modeling water flowing in independent thermal cell layers, which represent microchannels in the stacks. These three steps are detailed in this section.

#### A. RC Network for 2D/3D Stacks

2D/3D thermal modeling can be accomplished using an automated model that forms the RC circuit for certain grid dimensions. In this work, it is used the model proposed in [15], which has been extended to include 3D modeling capabilities as discussed in [17]. The extension for the existing multi-layered thermal modeling provides a new interlayer material model to include the TSVs (cf. Section IV-B) and the microchannels (cf. Section IV-C). Then, in a typical automated thermal model, the thermal resistance and capacitance values of the blocks or grid cells are computed initially at the start of the simulation, considering that the system properties do not vary at runtime. To model the heterogeneous characteristics of the interlayer material including the TSVs and microchannels, I introduce two major differences to other works: (1) as opposed to having a uniform thermal resistivity value of the layer, our infrastructure enables having various resistivity values for each grid, (2) the resistivity value of the cell can vary at runtime. The interlayer material is divided into a grid, where each grid cell except for the cells of the microchannels has a fixed thermal resistance value depending on the characteristics of the interface material and TSVs. The thermal resistivity of the microchannel cells is computed based on the liquid flow rate through the cell, and the characteristics of the liquid at runtime.




Fig. 2. Manufactured 5-tier stack chip for 3D thermal library validation

The proposed RC thermal model has been calibrated for the manufacturing technologies of 2D MPSoCs using experimental data based on the technologies used by industrial partners (Sun, Freescale, IBM, etc.). Then, the tuning of the version of the thermal library for 3D MPSoCs has been performed by manufacturing a 5-tier 3D chip stack with resistors and thermal sensors, as shown in Figure 2.

Exhaustive experiments have been performed in the 5-tier stack to characterize the possible inaccuracy of the proposed RC thermal network for 2D/3D chip stacks. One of the measured sets of experiments is shown in Figure 3, with heat sources modeling cores in the first tier and measurements in the last tier (to create the largest possible temperature variation between

measurements and heat source injection). This figure shows that the average error in the worst thermal propagation part in 3D stack (inter-tier thermal model) is very small, i.e., 2.7%, and the maximum error is less than 5%.

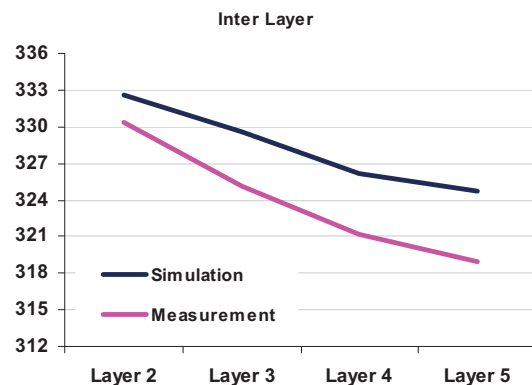


Fig. 3. Thermal measurements of inter-layer heat propagation vs. RC-network simulations for the 5-tier stack chip

#### B. Through-Silicon-Vias Modeling

In order to model the effect of TSVs on the thermal behavior of 3D MPSoCs, it is necessary to first perform a study to determine which modeling granularity is required. In the TSV model it is required to provide a TSV density for each unit (i.e., core, cache, interconnect line, etc.). Therefore, it is assumed that the effect of the TSV insertion to the heat capacity of the interface material is negligible, which is reasonable as the total area of TSVs constitutes a very small percentage of the total area of the material. Then, it is differentiated among the different block functionalities to adjust the TSV density. For example, a crossbar structure requires a high TSV density, while a processing core does not require any modeling of TSV interference in its thermal spreading properties. As a result, we assign a TSV density to each unit based on its functionality and system design choices. The TSV dimensions are set to  $10\mu\text{m} \times 10\mu\text{m}$ , and a minimum spacing of  $10\mu\text{m}$  from each side of the TSV is employed. In fact, the experiments developed in the calibrated 3D stack thermal simulation model of the 5-tier stack indicates that a block-level granularity provides very similar results to providing the exact locations of TSVs, while it has a very important complexity reduction in transient thermal analysis.

#### C. Active (Liquid) Cooling Modeling

Next, active cooling properties (i.e., liquid cooling) have been modeled using additional layers of thermal cells with different cooling thermal conductance and resistance properties than silicon and metal layers, using IBM's technology [12, 13]. In fact, in a 3D system with liquid cooling, the local junction temperature can be computed using a resistive network, as shown in Figure 4.

In this figure, the thermal resistance of the wiring layers ( $R_b$ ), the thermal resistance of the silicon ( $R_{Si}$ ) and the convective thermal resistance are combined to model the 3D stack. Considering the heat flux ( $q$ ) as the source and the chip back-side temperature ( $T_{\text{fluid}}$ ) as the ground, the electrical circuit can

be solved to get the junction temperature ( $T_{\text{junction}}$ ). Thus, the total thermal resistance ( $R_{\text{tot}}$ ) of the junction is computed as in Equation 2 [13]. The parameters of the equations are listed in Table I, and their values are fixed according to [12] and [13].

$$R_{\text{tot}} = R_{\text{cond}} + R_{\text{conv}} + R_{\text{heat}} \quad (1)$$

$$R_{\text{tot}} = 1/(G_{\text{si}}/t + 1/R_{\text{b}}) + A/(h \cdot A_{\text{t}}) + A/(V \cdot P \cdot c_{\text{p}}) \quad (2)$$




Fig. 4. Equivalent 3D resistive network including liquid cooling

According to [13], it is considered a base flow rate between 15ml and 150ml/min. Then, each channel has a width of 500 $\mu\text{m}$  and a depth of 300 $\mu\text{m}$ .

TABLE I  
2D/3D THERMAL AND LIQUID COOLING CHARACTERISTICS

Parameter name	Definition
$R_{\text{tot}}$	Total thermal resistance
$R_{\text{cond}}$	Conductive thermal resistance
$R_{\text{conv}}$	Convective thermal resistance
$R_{\text{heat}}$	Thermal resistance of passive material layers
$G_{\text{si}}$	Variable thermal conductivity of Si (dependent on T)
$t$	Si baseline thickness
$R_{\text{b}}$	Thermal resistance of interconnection layers
$A$	Area of high power dissipation intensity
$h$	Heat transfer coefficient of fluid
$A_{\text{t}}$	Total surface area
$V$	Volumetric flow rate
$P$	Density
$c_{\text{p}}$	Heat capacity

Then, Figure 5 outlines the emulated microchannels liquid cooling and TSVs layout. Thus, the microchannels are distributed uniformly on each tier and the fluid flows through each channel with the same flow rate, which can be modified at run-time by the OS, and do not intersect with TSVs.

## V. EXPERIMENTAL SETUP AND RESULTS

The proposed HW/SW thermal emulation framework for 2D/3D MPSoCs has been compared with different SW thermal libraries for 2D/3D MPSoCs [6, 7, 17], while running intensive MPSoCs processing kernels. The obtained results are depicted in Figure 6 and show significant speed-ups with respect to state-of-the-art temperature estimation frameworks [14,15, 17].

In particular, these results outline that the proposed modeling approach for MPSoC HW/SW thermal emulation scales significantly better than state-of-the-art SW simulators for

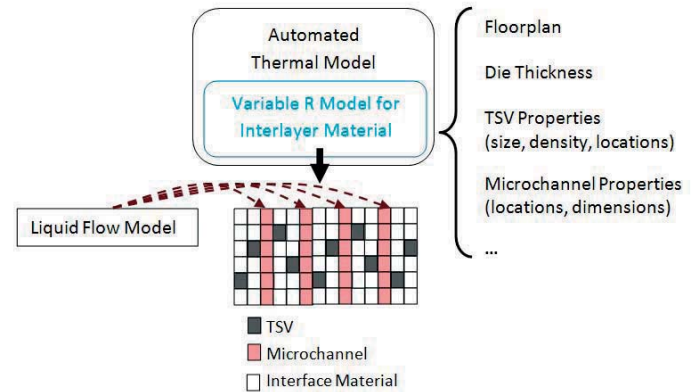


Fig. 5. Representation of microchannels/TSVs layout in emulated 3D MPSoCs

transient thermal analysis. In fact, the results of the exploration of 2D thermal behavior on a commercial 8-core MPSoC [2] has shown that the proposed thermal emulation can achieve speed-ups of more than to 800 $\times$  with respect to thermal simulators.

Moreover, the thermal exploration of 3D MPSoCs with active cooling (liquid) modeling shows even larger speed-ups (more than 1000 $\times$ ) due to power extraction and thermal synchronization overhead in thermal simulators [6, 7, 17].

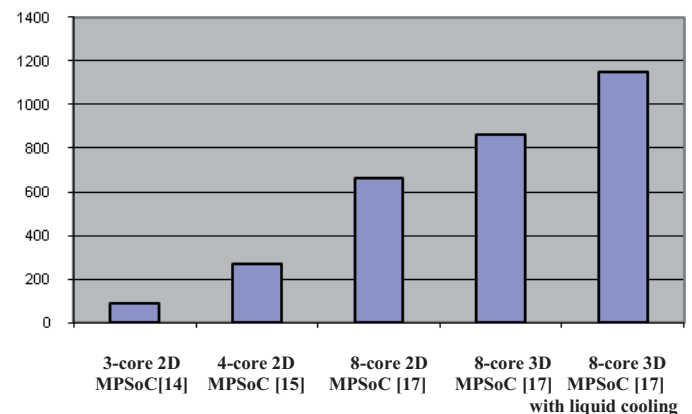


Fig. 6. Simulation speed-ups of the proposed HWSW thermal emulation framework for transient thermal analysis with respect to state-of-the-art 2D/3D thermal simulators

In the second set of experiments it has been evaluated the accuracy of the proposed thermal model with respect to transient thermal analysis of 3D liquid cooling-based MPSoCs. To this end, it has been compared the temperature evolution at the junction (cf. Equation 2) using finite-element simulations [13] (red straight line) with the estimated temperature of the linear model of Figure 4 (yellow dashed line). The results are shown in Figure 7, which indicates that the variations between both types of simulations are less than 1.5% on average (encircled area). Furthermore, while the proposed emulation framework using the simple liquid cooling model for straight channels can calculate the junction temperature evolution in the order of few milliseconds, the detailed finite-element simulation can take few hours. Thus, it illustrates the potential of linear thermal estimation methods for simple geometries of liquid microchannels using a laminar flow regime.




Fig. 7. Temperature evolution at the junction in liquid-cooling based systems using finite-element simulation and the proposed model for straight liquid cooling channels

All in all, these experiments outline the potential benefits of the proposed HW/SW thermal emulation framework to explore the design space of complex thermal management policies in 2D/3D MPSoCs, compared to existing methods based on SW cycle-accurate simulators or finite-elements simulation, which suffer from very important speed limits for long simulations, necessary to achieve representative transient thermal analysis of real systems.

## VI. CONCLUSIONS

2D and emerging 3D MPSoC architectures have been proposed as a promising solution to exploit the available area in forthcoming computing systems. In this paper I have presented a new HW/SW FPGA-based emulation framework that enables the rapid analysis of run-time thermal behavior in 2D/3D MPSoCs with active liquid cooling. The experimental results have shown that the proposed framework obtains detailed transient thermal exploration with a speed-up of more than 1000 $\times$  with respect to cycle-accurate MPSoC simulators, even more when active (liquid) cooling effects are considered in the overall thermal system analysis. Furthermore, almost no loss in thermal estimation accuracy (less than 3%) is experienced with respect to classical (and very time-consuming) finite-element simulations. Overall, this HW/SW thermal emulation approach is a promising mechanism to perform long-time transient behavior characterization in 2D and 3D MPSoC stacks.

## ACKNOWLEDGMENT

The author would like to thank Ayse K. Coskun and Prof. Tajana Simunic Rosing from UCSD, Prof. Luca Benini at Bologna University, and the group of Advanced Packaging Technologies at IBM Zürich for their useful feedback and inputs in the validation of the 3D thermal modeling and liquid cooling technology. This work has been supported in part by the Swiss NanoTera NTF Project - CMOSAIC, and a FPGA donation of the OpenSPARC University Program of Sun Microsystems.

## REFERENCES

- [1] D. Pham et al., Design and Implementation of a First-Generation Cell Processor., *Proc. ISSCC*, 2005.
- [2] P. Kongetira et al., Niagara: A 32-way multithreaded SPARC processor., *IEEE Micro*, 2005.
- [3] Tiler Corporation, Tiler's 64-core architecture, 2008, [www.tiler.com/products/processors.php](http://www.tiler.com/products/processors.php)
- [4] W. Davis, et al., Demystifying 3D ICs: The Pros and Cons of Going Vertical. *IEEE Des&Test*, 2005.
- [5] M. Healy, et al., Multiobjective microarchitectural floorplanning for 2D and 3D ICs. *IEEE Transactions on CAD*, 2007
- [6] K. Skadron et al., Temperature-aware microarchitecture: Modeling and implementation (Hot-spot simulator), *IEEE TACO*, 2004
- [7] G. Paci et al., Exploring temperature-aware design in low-power MPSoCs, *Proc. DATE*, 2006.
- [8] M.-N. Sabry, High-precision thermal models, *IEEE TCPT*, 2005.
- [9] Cadence Palladium II, 2005. <http://www.cadence.com>.
- [10] ARM integrator AP, 2004. <http://www.arm.com>.
- [11] Emulation Engineering. Zebu models, <http://www.eve-team.com>.
- [12] T. Brunschwiler et al., Direct liquid-jet impingement cooling with micron-sized nozzle array and distributed return architecture. *Proc. JIHERM*, 2006.
- [13] T. Brunschwiler, et al., Interlayer cooling potential in vertically integrated packages, *Microsyst. Technologies*, 2008.
- [14] F. Mulas et al., Thermal Balancing Policy for Streaming Computing on Multiprocessor Architectures, *Proc. DATE*, 2008.
- [15] D. Atienza, et al., A fast HW/SW FPGA-based thermal emulation framework for multi-processor system-on-chip. *Proc. DAC*, 2006.
- [16] B. Richard, et al., *Numerical Analysis*, Brooks Cole, 2000.
- [17] A.K. Coskun, et al., Dynamic Thermal Management in 3D Multicore Architectures, *Proc. DATE*, 2009.
- [18] A. Jerraya, et al. Multiprocessor SoCs. *Morgan Kaufmann*, 2005.
- [19] G. Braun, et al. Processor/memory co-exploration on multiple abstraction levels. *Proc. DATE*, 2003.
- [20] P. S. Magnusson, et al. Simics: A full system simulation platform. *IEEE Computer*, 2002.
- [21] P. Paulin, et al. Steppn: A system-level exploration platform for network processors, *IEEE Des&Test*, 2002.
- [22] Coware, Convergence and Lisatek product lines, 2006. <http://www.coware.com>
- [23] ARM, PrimeXSys platform architecture and methodologies, white paper, 2005, <http://www.arm.com/>
- [24] Heron Engineering, SoCemulation, 2004, <http://www.hunteng.co.uk>
- [25] M. D. Nava, et al., An open platform for developing MPSoC, *IEEE Computer*, 2005.
- [26] Y. Nakamura, A fast HW/SW co-verification method for SoC by using a C/C++ simulator and FPGA emulator with shared register communication, *Proc. of DAC*, 2004.
- [27] Mentor Graphics, Platform express and primecell, 2005, <http://www.mentor.com/>
- [28] Synopsys, Realview Max-sim ESL Exploration Framework, 2004, <http://www.synopsys.com/>
- [29] T.-Y. Chiang, et al, Thermal analysis of heterogeneous 3-D ICs with various integration scenarios, *Proc. of IEDM*, 2001.
- [30] A. Rahman, et al., Thermal analysis of three-dimensional integrated circuits, *Proc. of IITC*, 2001.
- [31] P. Leduca et al., Challenges for 3d IC integration: bonding quality and thermal management, *Proc. of IITC*, 2007.
- [32] K. Puttaswamy, et al., Thermal analysis of a 3d die-stacked high-performance microprocessor, *Proc. of GLSVLSI*, 2006.
- [33] A. Jain, et al., Thermal modeling and design of 3D integrated circuits, *Proc. of ICTTPES*, 2008.