

# GreenCool: An Energy-Efficient Liquid Cooling Design Technique for 3-D MPSoCs Via Channel Width Modulation

Mohamed M. Sabry, *Student Member, IEEE*, Arvind Sridhar, *Student Member, IEEE*, Jie Meng, *Student Member, IEEE*, Ayse K. Coskun, *Member, IEEE*, and David Atienza, *Member, IEEE*

**Abstract**—Liquid cooling using interlayer microchannels has appeared as a viable and scalable packaging technology for 3-D multiprocessor system-on-chips (MPSoCs). Microchannel-based liquid cooling, however, can substantially increase the on-chip thermal gradients, which are undesirable for reliability, performance, and cooling efficiency. In this paper, we present *GreenCool*, an optimal design methodology for liquid-cooled 3-D MPSoCs. *GreenCool* simultaneously minimizes the cooling energy for a given system while maintaining thermal gradients and peak temperatures under safe limits. This is accomplished by tuning the heat transfer characteristics of the microchannels using channel width modulation. Channel width modulation is compatible with the current process technologies and incurs minimal additional fabrication costs. Through an extensive set of experiments, we show that channel width modulation is capable of complementing and enhancing the benefits of temperature-aware floorplanning. We also experiment with a 16-core 3-D system with stacked dynamic random-access memory, for which *GreenCool* improves energy efficiency by up to 53% with respect to no channel modulation.

**Index Terms**—3-D ICs, energy efficiency, liquid cooling.

## I. INTRODUCTION

**3-D** STACKING technology enables building multiprocessor system-on-chips (MPSoCs) and integrated circuits (ICs) with higher transistor density per footprint, integrating heterogeneous technologies, and achieving more desirable tradeoffs between manufacturing cost and performance. Early 3-D stacked products in the market include stacked memory chips, package-on-package integration, and 2.5-D systems. Recently, research efforts for building 3-D MPSoCs and connecting layers in a 3-D stack using high-speed

Manuscript received March 29, 2012; revised August 1, 2012; accepted September 29, 2012. Date of current version March 15, 2013. This work was supported in part by the Swiss Confederation under Nano-Tera RTD Project CMOSAIK 123618 (scientifically evaluated by the SNSF), the EC in the 7th Framework Program PRO3D under STREP Project FP7-ICT-248776, NSF CAREER Grant 1149703, and the DAC Richard Newton Scholarship. This paper was recommended by Associate Editor G. Loh.

M. M. Sabry, A. Sridhar, and D. Atienza are with the Embedded Systems Laboratory, École Polytechnique Fédérale de Lausanne, Lausanne 1015, Switzerland (e-mail: mohamed.sabry@epfl.ch; arvind.sridhar@epfl.ch; david.atienza@epfl.ch).

J. Meng and A. K. Coskun are with the Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215 USA (e-mail: jiemeng@bu.edu; acoskun@bu.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCAD.2012.2226032




Fig. 1. Liquid-cooled 3-D IC with interlayer microchannels [6].

through-silicon vias (TSVs) have gained momentum [1]–[4]. However, the development of high-performance 3-D MPSoCs is strongly limited by the increase in on-chip temperatures. High-performance 3-D MPSoCs include a large number of processor cores in close proximity. Such high power densities, combined with the reduced cooling efficiency for the layers away from the cooling subsystem (e.g., heat sinks and fans), aggravate the existing temperature-induced problems [5].

Considering the significance of the thermal challenges in 3-D stacked design, interlayer liquid cooling using microchannels that are directly etched on the back of the substrates of individual layers (see Fig. 1) is a viable and scalable cooling solution for 3-D MPSoCs. Prototypes of 3-D interlayer water-based liquid-cooled packages have demonstrated the superior cooling capabilities of this technology, the corresponding reduction in the cooling energy, and also the relative ease with which the current CMOS processes can be modified for manufacturing liquid-cooled systems [7]–[9].

Despite these advances, interlayer liquid-cooling brings new thermal management challenges such as increased thermal gradients. Large gradients on chips cause reliability challenges as many failure mechanisms are accelerated by spatial thermal variations [10], [11]. In addition, device switching delay is affected by temperature and as a result, large gradients incur potential timing errors at run-time or higher design complexity to mitigate such errors. Large thermal gradients in liquid-cooled ICs appear mainly because of the rise in coolant temperature as liquid flows from inlet to outlet absorbing heat along the way [7], [12], [13]. Recent research has demonstrated that applying thermal management techniques designed for planar 2-D ICs to 3-D systems is suboptimal for meeting desired performance-temperature-energy requirements [14], [15]. While recent work has also developed thermal management policies for liquid-cooled 3-D systems [5], [6], [16], [17], these policies do not address the thermal gradients along the microchannels.

On the other hand, liquid cooling brings unique opportunities to design and optimize cooling. For example, microchannel heat transfer characteristics can be customized for specific target thermal loads [17], [18] in order to precisely control and minimize the cooling effort [6]. This paper builds on this customization concept and proposes gradient reduction and energy conservation using optimized liquid-cooling, or *GreenCool*, which is an optimal design methodology for liquid-cooled 3-D MPSoCs. *GreenCool* simultaneously reduces thermal gradients in liquid-cooled 3-D ICs and minimizes the energy consumed for pumping the coolant into the system. This is accomplished through channel width modulation [12], [13], which allows different segments of a microchannel to have different widths (as opposed to having a straight, fixed-width microchannel). In this way, it is possible to tune the heat transfer characteristics of the microchannel at a fine granularity. Hence, channel width can be optimally adjusted to cater to the thermal loads of local hot spots, as well as to compensate for the rise in coolant temperature as it flows from inlet to outlet. *GreenCool* accomplishes this by providing an optimal channel width profile for an entire 3-D MPSoC based on its unique heat flux footprint. *GreenCool* is also compatible with the current process technologies and incurs minimal additional fabrication costs. Our contributions in this paper are as follows.

- 1) We quantify the impact of varying the channel width on convective heat transfer and the overall thermal state of 3-D MPSoCs.
- 2) We develop *GreenCool* as an optimal design method for minimizing the coolant pumping power using channel modulation, subject to thermal constraints.
- 3) Using a 3-D MPSoC test vehicle, we demonstrate the effectiveness of *GreenCool*. We compare the energy efficiency of *GreenCool* with channel modulation against uniform channel width-based optimization (referred to as *GreenCool* without modulation). Our experiments, including floorplan exploration, show that *GreenCool* with modulation can adhere to given thermal constraints while saving up to 98% pumping power, regardless of how thermally suboptimal the floorplan is.
- 4) We also perform experiments using a 16-core 3-D MPSoC with stacked dynamic random-access memory (DRAM). We model the memory access latency and the overall application-level performance of this 3-D system in detail. On the 16-core system, our experiments illustrate that *GreenCool* with channel modulation improves energy efficiency by up to 53% compared to optimization without channel modulation.

The rest of this paper starts with an overview of related work. Section III describes the fundamentals of channel width modulation and presents the temperature simulation model for liquid-cooled 3-D ICs. Section IV formulates the optimization problem and describes the design space exploration. Section V demonstrates the impact of *GreenCool* under various floorplan optimizations. Section VI provides the performance modeling methodology for the 3-D MPSoC with stacked DRAM. We describe the experimental setup, 3-D MPSoC architecture, and the workload characteristics in

Section VII. Section VIII discusses the experimental results and Section IX concludes this paper.

## II. RELATED WORK

This section first provides an overview of the state-of-the-art in the design and thermal modeling of liquid-cooled ICs. We then discuss existing run-time and design-time thermal management techniques for liquid-cooled 3-D MPSoCs.

### A. Liquid Cooling Utilization in 3-D ICs

The seminal work done by Tuckerman and Pease [19] establishes the foundation of today's research efforts to build a complete 3-D stacked IC with interlayer water-based microchannel liquid cooling. They demonstrate that microchannel liquid-cooled heat sinks can remove heat fluxes on the order of 800 W/cm<sup>2</sup> while operating below 85 °C. They also present a preliminary theoretical and experimental study of the relationship between the aspect ratio of the microchannels and the heat transfer characteristics.

More recently, back-side water-based liquid cold plates (such as staggered microchannel and distributed return jet plates) are developed for IC cooling, which can handle up to 400 W/cm<sup>2</sup> in single-chip applications [20]. Prototypes of 3-D chips with interlayer microchannel heat sinks have been shown to handle up to 250 W/cm<sup>2</sup> of hot spot heat fluxes, demonstrating the scalability of interlayer liquid cooling [8], [9]. Enhanced heat transfer geometries, such as pin fins, have also been built as an alternative to microchannels, improving the heat transfer characteristics at the cost of higher unpredictability in coolant flow patterns [7], [21].

### B. Thermal Modeling of Liquid-Cooled ICs

While thermal simulation for conventional air-cooled ICs has been well investigated (see [22]), the thermal simulation of liquid-cooled ICs has recently garnered interest. A thermal model for 3-D systems with microchannel cooling [23] has been developed based on HotSpot [22].

Mizunuma *et al.* [24] propose a steady-state thermal simulation method, which is based on the extraction of thermal properties from numerical presimulations running on computational fluid dynamics simulators and their inclusion in a conventional thermal resistance circuit for microchannel-cooled 3-D ICs. Sridhar *et al.* [25] advance the first transient thermal simulation method for liquid-cooled ICs. Their simulator, 3D-ICE, uses a new compact model representation for forced convective cooling and creates an RC circuit for the microchannels. This RC circuit can then be integrated with the existing RC model for thermal conduction in the IC to perform transient conjugate conduction-convection simulations in liquid-cooled ICs.

A new porous medium approach is also proposed for 3D-ICE, which enables the simulations of enhanced heat transfer geometries such as pin fins [26]. Feng and Li [27] introduce a thermal simulation framework of 3-D stacks, where graphical processing units are utilized for accelerating temperature calculation.

### C. Run-Time Management for Liquid-Cooled 3-D MPSoCs

Several run-time thermal management techniques have been proposed to address the challenges presented by liquid-cooled

3-D MPSoCs. Coskun *et al.* [23] evaluate existing thermal management policies on a 3-D MPSoC with a fixed-flow rate value, and propose a run-time thermal management policy that dynamically adjusts the flow rate based on temperature measurements to balance the temperature on the 3-D MPSoC. Our recent work improves the energy efficiency of 3-D MPSoCs by using variable flow rate adjustment and thermally aware load balancing [15].

Recently, Qian *et al.* [17] explore the use of a cyber-physical approach for 3-D MPSoCs thermal management with inter-tier liquid cooling. They construct their control mechanism with software-based thermal estimation and prediction. They use a nonuniform liquid flow in different microchannels to meet the cooling demands of different modules. Our previous work on run-time thermal management of liquid-cooled 3-D MPSoCs uses fuzzy-logic control to achieve energy-efficient thermal management, where a combined control of flow rate, dynamic voltage and frequency scaling, and task scheduling is used [6]. Finally, a recent work by Zanini *et al.* [16] explores the use of a convex optimization and hierarchical control to achieve low power thermal management. These run-time management methods, while effective, do not address the thermal gradients caused by the heat absorption in the microchannels.

#### D. Design-Time Optimization for Liquid-Cooled 3-D MPSoCs

Complementary to run-time thermal management, design-time solutions such as thermally-aware floorplanning have also been proposed for 3-D MPSoCs thermal management [28]–[31].

Qian *et al.* [17] propose a channel clustering methodology for this problem where microchannels are grouped into clusters of channels. Within a single cluster, the channels have the same flow rate in order to customize the cooling effort based on the demands of computing elements. Recently, Qian *et al.* [32] extend their previous work by proposing an energy-efficient microchannel clustering technique where the primary aim is to minimize the pumping energy consumed to achieve a given peak temperature constraint. Shi *et al.* [33] propose a customized channel allocation technique, where the density of the etched microchannels reflects the cooling demands of various regions of the IC. However, the channel allocation technique primarily targets the improvement of energy efficiency instead of the thermal distribution. Moreover, while these methods work when the hot spots line up perpendicular to the coolant flow, they do not address the case where multiple different hot spots lie along a channel.

Mizunuma *et al.* [34] use their thermal model to explore floorplanning solutions to homogenize temperature distributions in a 3-D IC. However, their work targets an unrealistic large number of identically sized functional blocks.

Brunschwiler *et al.* [13], [18] investigate channel width modulation, four-port fluid access, and the use of fluid guiding structures for thermal optimization. They show that by changing the channel width as we move from inlet to outlet, we can perform customized cooling and thus achieve thermal balancing. However, their channel modulation scheme relies on heuristics and without providing an optimality guarantee. In our recent work, we use optimal




Fig. 2. Test structure: a single microchannel cooling a strip of an IC with uniform heat flux distribution. 3-D and cross-sectional views.

control theory to find the best possible channel width profile and provide a mathematically precise solution to the problem of minimizing thermal gradients in the ICs [12]. Our approach finds the solution with the theoretical minimum possible thermal gradient for a given problem.

In this paper, we use the same principle of channel modulation as in [12] to design energy-efficient liquid cooling systems that keep thermal gradients under a user-defined limit. This paper differs from prior work in channel modulation or energy-efficient cooling design as follows.

- 1) Unlike previous design-time methods for energy-efficient liquid cooling, our proposal is applicable with the current process technologies and does not bring any complexity to the fluid delivery network.
- 2) While prior work mostly focuses on the peak temperatures, we also take the maximum spatial thermal variations into account in our optimization procedure. In addition, we show the sensitivity of our technique with respect to changes in the maximum thermal variation requirement.
- 3) While our proposal is a design-time technique, thus complementary to run-time thermal management strategies, it handles the thermal gradient reduction more efficiently compared to run-time techniques. As our technique optimizes the tradeoff between pumping power and thermal gradients [5], [6], it helps achieve more efficient thermal management.

### III. CHANNEL WIDTH MODULATION: CONCEPTS

This section first describes the causes for large thermal gradients in liquid-cooled ICs. We then present the concept of channel modulation, as a method used by GreenCool to reduce thermal gradients and to minimize the pumping effort. In addition, the mathematical formulation of the problem is described, which is later used to formulate and solve the proposed optimal design problem.

#### A. Thermal Gradients and Channel Modulation

Thermal gradients arise in liquid-cooled ICs for two reasons: 1) nonuniform heat fluxes resulting in nonuniform temperature distribution, and 2) sensible heat absorption by the coolant, creating uneven heat-sinking from inlet to outlet. The latter reason is far more dominant than the former one. This is illustrated using the example shown in Fig. 2. Here, a single microchannel heat sink cools a strip of silicon chip with




Fig. 3. Junction temperature distribution for the structure in Fig. 2 with (a) uniform nonmodulated channel width and (b) modulated channel width to compensate for sensible heat absorption.

uniform heat flux distribution  $\hat{q}_i$ . The direction of coolant flow is along the  $z$ -axis. The cross section of the structure is also shown in the figure. For this structure, the junction temperature (i.e., the temperature in the active layer of silicon) is plotted as a function of the longitudinal distance from the inlet in Fig. 3(a) [7]. As the figure shows, even in the case of uniform heat flux distribution, the junction temperature rises steadily with respect to the fluid inlet temperature from the inlet, reaching a maximum at the outlet. The junction temperature has three main contributors.

- 1) The temperature change due to conduction from the junction to the surface of the microchannel walls through the silicon substrate, represented by  $\Delta T_{\text{cond}}$ .
- 2) The temperature change due to the convective resistance at the surface of the microchannel walls, where the heat from silicon is carried away by flowing liquid. This is represented by  $\Delta T_{\text{conv}}$ . For simplicity of illustration, we assume the convective resistance, and hence  $\Delta T_{\text{conv}}$ , to be constant from inlet to outlet by neglecting the entry region effects.
- 3) The temperature change due to the sensible heat absorption (coolant rising in temperature due to the storage of thermal energy) from inlet to outlet. This is referred to as  $\Delta T_{\text{heat}}$ . As shown in Fig. 3, this is the only varying quantity in the plot and hence, the primary contributor of thermal gradients on the silicon junction.

The thermal gradient in Fig. 3(a) can be reduced by modifying one of the above three contributors. First, we consider  $\Delta T_{\text{cond}}$ .  $\Delta T_{\text{cond}}$  depends on the conductive thermal resistance  $R_{\text{conv}}$ , which in turn depends on the thermal conductivity of silicon and the thickness of the substrate—both of which are determined by the technology and the fabrication process, and are difficult to change. Next, we consider  $\Delta T_{\text{heat}}$ . The sensible heat absorption is a function that depends on the volumetric heat capacity of the coolant and the flow rate. The slope of this function (and hence the thermal gradient on the silicon junction) can be reduced by pumping the coolant at a higher flow rate (in other words, taking away heat from the channels more quickly, preventing the temperature from rising in the IC). However, this comes at the cost of higher pumping power.

The main idea behind GreenCool is to achieve lower thermal gradients without any rise in cooling energy costs or changes to the existing fabrication process. Hence, in this paper, we focus on  $\Delta T_{\text{conv}}$ . The convective temperature change depends on the convective resistance  $R_{\text{conv}}$ , which under steady-state conditions, depends on the channel aspect ratio [35]. Using the conventional CMOS fabrication process for etching the




Fig. 4.  $R_{\text{conv}}$  as a function of the channel width for the structure in Fig. 2.

channels, it is possible to modulate the width of the channel from inlet to outlet (and hence its aspect ratio) and create any kind of channel width profile, while keeping the height of the channels constant. Thus, channel width modulation requires only a change in the patterns on the masks used for etching channels amounting to minimal additional fabrication costs. To summarize, using careful design it is possible to modify the local channel aspect ratios so as to contain the pumping power while constraining the thermal gradients.

To understand how the channel width affects  $\Delta T_{\text{conv}}$  in detail, we explore the following equations governing the Nusselt number (a dimensionless form of heat transfer coefficient), and the product of friction factor and Reynold's number for microchannels, under fully developed conditions [35]:

$$\begin{aligned} \text{Nu} &= 8.235 \cdot (1 - 2.0421\text{AR} + 3.0853\text{AR}^2 \\ &\quad - 2.4765\text{AR}^3 + 1.0578\text{AR}^4 - 0.1861\text{AR}^5) \\ \text{fr} \cdot \text{Re} &= 24 \cdot (1 - 1.3553\text{AR} + 1.9467\text{AR}^2 \\ &\quad - 1.7012\text{AR}^3 + 0.9564\text{AR}^4 - 0.2537\text{AR}^5) \end{aligned} \quad (1)$$

where AR is the aspect ratio reciprocal (height/width) of the channel. Using the Nusselt number, the heat transfer coefficient (a measure of the amount of heat transferred per unit area for one Kelvin difference in temperature between the fluid and the microchannel wall surface, expressed in  $\text{W}/\text{m}^2\text{K}$ ) can be written as follows:

$$h = \frac{k_{\text{coolant}} \cdot \text{Nu}}{d_h} \quad (2)$$

where  $k_{\text{coolant}}$  is the thermal conductivity of the coolant and  $d_h$  is the hydraulic diameter of channel. The effective heat transfer coefficient as seen by the junction looking down the channel from the top can be written by projecting the heat transfer coefficient above from the side wall surfaces onto the top, as follows:

$$h_{\text{eff}} = h \frac{2H_C + w_C}{W} \quad (3)$$

where  $H_C$  is the height and  $w_C$  is the width of the channel, and  $W$  is the total width of the structure as shown in Fig. 2.

The convective resistance  $R_{\text{conv}}$  for this structure can be obtained as a reciprocal of this quantity. The  $R_{\text{conv}}$  for this structure is plotted as a function of  $w_C$  in Fig. 4, assuming water as the coolant,  $H_C = 100 \mu\text{m}$ ,  $W = 100 \mu\text{m}$ , and varying  $w_C$  from 10 to 50  $\mu\text{m}$ .

Fig. 4 shows that the convective resistance (and also  $\Delta T_{\text{conv}}$ ) drops quickly as the channel width is reduced. Since our goal is to modify the convective resistance to compensate for  $\Delta T_{\text{heat}}$ , we can postulate that the channel width must no longer be a constant, but instead, should be a function of the distance along the channel  $w_C(z)$ . The width must be larger near the inlet where the fluid temperature is low and smaller near the outlet where the fluid temperature is high. Hence, theoretically, for the case of uniform heat flux, it is possible to lower the final thermal gradient by steadily modulating the channel width from inlet to outlet, as shown in Fig. 3(b).

However, in reality, a variety of factors affect our ability to achieve an ideal thermal gradient. First, the heat flux distributions in a realistic IC is highly nonuniform and the resulting local rise and fall in incoming heat flux adds to the temperature gradients. This must also be compensated for in addition to the rise in the coolant temperature. Second, the manufacturing constraints and the placement of TSVs in the microchannel walls restrict our ability to increase or decrease the channel width. Third, reducing channel widths increases the frictional resistance to the flow of the coolant through the channels, and thus increases the pumping effort required (in other words, the energy required for cooling the IC). These factors, combined with our goal to modulate the channel width for reducing thermal gradients, represent conflicting optimization goals in the design. Hence, we define an optimal control design problem consisting of a cost function to be minimized, where the cost function and constraints are determined based on the goals and system specifications. In GreenCool, we use channel modulation to solve the optimal control problem of minimizing pumping power to cool a given 3-D MPSoC, while maintaining thermal gradients and peak temperatures of the IC below desired levels and respecting various manufacturing constraints.

### B. Thermal Model

For solving the optimal control problem defined above, it is essential to find an analytical formulation for the heat transfer problem in 3-D ICs with microchannel liquid cooling. This analytical formulation must be in the form of an ordinary differential equation providing the mathematical platform on which an optimal control algorithm can work. The goal of our optimization is to compute the channel width profile as a function of the distance from the inlet, which minimizes the intended cost function. Hence, the steady-state temperatures of the 3-D IC must be written as a function of this distance in the analytical formulation, with the channel widths as the input parameters. If the distance from the inlet is measured along the coordinate axis  $z$ , then we need to find an equation of the form

$$\frac{d}{dz} \mathbf{T}(z) = \Phi(z, \mathbf{w}_C(z), \mathbf{T}(z)) \quad (4)$$

where  $\mathbf{T}(z)$  is the vector of temperatures on the IC that we are interested in and  $\mathbf{w}_C(z)$  is a vector of width functions of




Fig. 5. Test structure: a single microchannel cooling two active silicon layers.

different microchannels written as a function of  $z$ . Our goal is to find the  $w_C(z)$  that minimizes the gradients in  $\mathbf{T}(z)$ .

In order to find the steady-state analytical model for heat transfer in an IC cooled by a microchannel, we consider a single microchannel structure of length  $d$  shown in Fig. 5 between two silicon layers, and two silicon side walls. The width of the entire structure is  $W$ . The width of the channel is a function of distance  $w_C(z)$ . The height of the channel and the height of the silicon walls above and below the channel are constants, with values of  $H_C$  and  $H_{Si}$ , respectively. Heat flux distributions  $\hat{q}_{i1}(z)$  and  $\hat{q}_{i2}(z)$  (measured here as the heat per unit length along  $z$  axis, W/m) are applied to the top and the bottom layers of the silicon—referred to as the active layers. Coolant enters the microchannel at the inlet ( $z = 0$ ) with a constant temperature  $T_{Cin}$ , absorbs heat from silicon along the way, and exits at the outlet ( $z = d$ ). All the exposed surfaces of silicon are assumed to be adiabatic, hence the microchannel heat sink is the only way for the heat to exit the system. Using the electrical analogy for heat transfer in ICs with microchannel heat sinks, where temperature is represented by voltage and heat flow is represented by current, circuit parameters can be written for a cell of this structure (see Fig. 6), representing the following five types of heat transfer occurring in the structure [25], [26]:

- 1) longitudinal heat conduction inside the two active silicon layers, parallel to the microchannel ( $\hat{g}_l$ );
- 2) vertical heat conduction from the active silicon layers to surface of the top and bottom walls of the microchannel ( $\hat{g}_w(z)$ );
- 3) vertical heat conduction between the active silicon layers through the silicon side walls enclosing the microchannel ( $\hat{g}_{v,si}$ );
- 4) convective heat transfer from the surface of the microchannel walls into the bulk of the flowing coolant ( $\hat{h}(z)$ );
- 5) convective heat transport downstream along the channel due to the mass transfer (flow) of the coolant ( $q_C(z)$ ).

The state variables in this formulation are the temperatures in the two active silicon layers,  $T_1(z)$  and  $T_2(z)$ , and the heat flowing in these layers parallel to the channel  $q_1(z)$  and  $q_2(z)$ .  $T_C(z)$  represents the temperature of the coolant as a function of the distance from the inlet. Assuming the silicon thermal




Fig. 6. Cell of the test structure with length  $\Delta z$  and at a distance  $z$  from the inlet, and the equivalent electrical circuit.

conductivity is  $k_{Si}$ , the volumetric heat capacity of the coolant is  $c_v$ , and the volumetric flow rate is  $\dot{V}$ , we can write the following parameters for this circuit [25], [26]:

$$\begin{aligned} \hat{g}_l &= k_{Si} \times W \times H_{Si} \text{ (W} \times \text{m)} \\ \hat{g}_w(z) &= k_{Si} \times \frac{(W - w_c(z))}{(2H_{Si} + H_C)} \text{ (W/m)} \\ \hat{g}_{v,Si} &= k_{Si} \times \frac{W}{H_{Si}} \text{ (W/m)} \\ \hat{h}(z) &= h(z, w_c(z)) \text{ (W/m)} \\ \hat{g}_v(z) &= (\hat{g}_{v,Si}^{-1} + \hat{h}(z)^{-1})^{-1} \text{ (W/m)} \\ q_C(z) &= c_v \times \dot{V} \times T_C(z) \text{ (W)}. \end{aligned} \quad (5)$$

The heat transfer coefficient  $\hat{h}(z)$  is a function of various parameters, namely Reynold's number of the flow, the coolant thermal conductivity, its viscosity, the distance from the inlet, and the width of the channel  $w_c(z)$ . Our model is independent of the method used to estimate heat transfer coefficients, the correlation studies based on experiments, or numerical techniques. In this study, we adopt the heat transfer coefficient calculated using the Nusselt number correlations (as a function of channel aspect ratio) presented by Shah and London for isothermal channel perimeters as shown in (1) [35].

Using the parameters in (5), the state-space analytical model for heat transfer in the test structure can be derived as follows:

$$\frac{d}{dz} \mathbf{X}(z) = \mathbf{F}(z, w_c(z), \mathbf{X}(z)) + \mathbf{G}(\hat{\mathbf{q}}_i(z), T_{Cin}) \quad (6)$$

where

$$\mathbf{X}(z) = \begin{bmatrix} T_1(z) \\ T_2(z) \\ q_1(z) \\ q_2(z) \end{bmatrix} \quad \hat{\mathbf{q}}_i(z) = \begin{bmatrix} \hat{q}_{i1}(z) \\ \hat{q}_{i2}(z) \end{bmatrix}. \quad (7)$$

$\mathbf{F}(z, w_c(z), \mathbf{X}(z))$  is a nonlinear function of distance, channel width, and the states of the system.  $\mathbf{G}(\hat{\mathbf{q}}_i(z), T_{Cin})$  is another vector function that is independent of the states and dependent solely on the input heat flux distributions and the inlet temperature of the coolant. Since all the exposed surfaces of the IC are assumed to be adiabatic, the boundary conditions for the above analytical model can be defined as follows:

$$q_i(0) = q_i(d) = 0, \quad i = 1, 2. \quad (8)$$

The derivation of the functions in (7) is beyond the scope of this paper. The proposed model has been validated against the numerical simulator 3D-ICE [25]. Our model can be extended for the case of multiple channels adjacent to each other, by taking into account of the additional heat spreading in the lateral ( $y$ ) direction. Each added channel brings two additional nodes: one for the top layer and one for the bottom layer of silicon (each node constitutes a temperature variable and a heat flow variable). It is also possible to combine two or more channels under a single set of top and bottom nodes to reduce the model complexity, by scaling the per-unit-length parameters in (5) suitably.

#### IV. GREENCOOL OPTIMIZATION METHODOLOGY

This section describes the GreenCool methodology for optimizing the energy efficiency of 3-D MPSoCs. Our method minimizes the required pumping energy while maintaining given peak temperature and thermal gradient constraints. For the sake of simplicity, in the optimization procedure described below, we assume that the fluid is always under fully developed conditions. We also assume that fluid parameters such as viscosity and density are constant and temperature independent for the computation of convective resistances [35]. It is important to note that the proposed approach is able to incorporate a more complex relationship between the channel width and the flow profiles or heat transfer characteristics. However, the study of such relations is beyond the scope of this paper.

##### A. Cost Function Derivation

Since our primary optimization metric is the pumping energy, we use Bernoulli's equation that relates the required input power ( $Q_{\text{pump}}$ ) with the required pressure drop ( $\Delta \mathbf{P}$ ) and the total flow rate ( $\dot{\mathbf{V}}$ )

$$Q_{\text{pump}} = \frac{\Delta \mathbf{P} \cdot \dot{\mathbf{V}}^T}{\eta} \quad (9)$$

where  $\eta$  is the pump efficiency,  $\Delta \mathbf{P}$  is the pressure drop vector among the existing  $N$  microchannels in the 3-D MPSoC (i.e.,  $\Delta \mathbf{P} = [\Delta P_0, \Delta P_1, \dots, \Delta P_{N-1}]$ ), and  $\dot{\mathbf{V}}$  is the flow rate vector of all the flow rates in the different channels (i.e.,  $\dot{\mathbf{V}} = [\dot{V}_0, \dot{V}_1, \dots, \dot{V}_{N-1}]$ ). Since the pump efficiency ( $\eta$ ) is mainly dependent on the pump characteristics and not on the modulated channel or the 3-D MPSoC thermal state, it does not have any impact on the optimization procedure. Hence, we remove its impact from our upcoming derivations and calculations. We use the Darcy Weisbach equation to calculate the pressure drop. For the  $i^{\text{th}}$  microchannel ( $i \in [0, 1, \dots, N-1]$ ) of length  $d$  and a channel width function  $w_{C,i}(z)$ , the pressure drop across it ( $\Delta P_i$ ) is calculated as follows:

$$\Delta P_i = \int_0^d 2 \frac{f_r \cdot \text{Re}_i(z) \cdot \rho \cdot v_i(z)^2}{\text{Re}_i(z) \cdot Dh_i(z)} dz \quad (10)$$

where

$$v_i(z) = \frac{\dot{V}_i}{w_{C,i}(z) \cdot H_C} \quad \text{(flow velocity)} \quad (11)$$

$$\text{Re}_i(z) = \frac{Dh_i(z) \cdot v_i(z)}{\nu} \quad \text{(Reynold's number)} \quad (12)$$

$$Dh_i(z) = 2 \frac{w_{C,i}(z) \cdot H_C}{w_{C,i}(z) + H_C} \quad (\text{hydraulic diameter}) \quad (13)$$

$$\text{AR}_i(z) = \frac{H_C}{w_{C,i}(z)} \quad (\text{aspect ratio reciprocal}). \quad (14)$$

The product of the friction factor and Reynold's number  $f_r \cdot \text{Re}(z)$  is given by (1) where AR is replaced by  $\text{AR}_i(z)$ . The equations above demonstrate that there is a nonlinear dependency of the pressure drop on the channel width function. Our goal is to find a channel width profile ( $w_C(z) = [w_{C,0}(z), w_{C,1}(z), \dots, w_{C,N-1}(z)]$ ) and a total flow rate value ( $\dot{V}$ ) such that the used pumping energy is minimized. Hence, we can formulate the problem of optimum channel width and flow rate selection as follows:

$$\begin{aligned} \min_{w_C(z), \dot{V}} J &= \Delta P \cdot \dot{V}^T \quad (15) \\ \text{subject to} & \quad 1) \text{ (6) and (8)} \\ & \quad 2) \text{ Design constraints.} \end{aligned}$$

### B. Design Constraints

In our channel modulation formulation for maximizing energy efficiency, there are several design constraints that should not be violated.

1) *Constraints on Channel Widths:* TSV arrays are the driving factors of 3-D stacked integration, and the microchannel heat sinks implemented in 3-D ICs must be compatible with them. Hence, it must be ensured that the maximum channel width is bounded to give clearance for the etching processes involved in the fabrication of TSVs. This bound depends on the TSV pitch and diameter. On the other hand, channels cannot be arbitrarily thin. Thin channels are not only difficult to fabricate, but also result in excessive resistance to coolant flow requiring larger pumping effort [13]. These considerations require a minimum channel width to be defined. Thus, our optimal design problem is constrained with the following inequality for  $N$  channels:

$$\begin{aligned} w_{C\min} \leq w_{C,i}(z) \leq w_{C\max} \\ \forall z \in [0, d], \quad i = 0, 1, \dots, N-1. \quad (16) \end{aligned}$$

2) *Equality of Pressure Drop in Microchannels:* Liquid-cooled 3-D MPSoCs are typically manufactured using a hermetically sealed manifold which forms the basic structure that helps the transfer of the coolant from an external source into the various cavities in the target system (see Fig. 1). As Fig. 1 shows at the inlet and outlet there is a single reservoir cavity from which fluid enters different channels on various layers. Hence, a reasonable assumption is that the pressure drop across the channels is the same in the IC. This condition can be enforced in the proposed method by using the following additional design constraint:

$$\Delta P_i = \Delta P_{i+1}, \quad i = 0, 1, \dots, N-2. \quad (17)$$

3) *3-D MPSoC Thermal Variation Constraints:* One of the aims of our cooling system is to ensure that the 3-D MPSoC operates under safe conditions. Safe operating conditions are defined by maximum limits on both the peak temperature and the spatial thermal variations (i.e., thermal gradients). In

our prior work [12], we have used optimal design techniques to minimize thermal gradients in an IC. Also, in the same work [12], we have limited the pressure drop (and hence the power supplied to the cooling systems) while solving the problem, to the maximal operational safe limits shown in liquid-cooled 3-D IC prototypes [7]. In this paper, our main objective is to minimize the cooling effort while maintaining safe operating conditions. Hence, in (15), the peak temperature and thermal gradients must appear as constraints.

The peak temperature constraint  $\max(\mathbf{T})$  is usually determined based on the critical thresholds provided by the manufacturer of the chip. However, thermal gradients have different definitions based on the designer's motivations and desirable operating conditions. A simple definition is the difference between the maximum and minimum temperature in IC. A more sophisticated definition would entail us to minimize the spatial gradient slope of the temperature on the IC surfaces. In this paper, we choose the former to simplify the overall optimization process. The thermal design constraints used in our problem formulation are

$$\max(\mathbf{T}) \leq T_{\max} \quad (18)$$

$$\max(\mathbf{T}) - \min(\mathbf{T}) \leq \Delta T_{\max} \quad (19)$$

where  $\mathbf{T}$  is the vector of temperatures computed for a given IC.  $T_{\max}$  and  $\Delta T_{\max}$  are the peak temperature and the thermal gradient constraints.

## V. ILLUSTRATION OF GREENCOOL USING A 3-D TEST VEHICLE

In this section, we illustrate with a simple example the potential of combining thermally aware floorplanning and GreenCool in a single optimization loop for global energy efficiency. In particular, we perform this analysis to show some insights that whether thermally optimized, energy-efficient 3-D MPSoC design is obtained with equal contribution of the aforementioned techniques, or a single technique has the major contribution on the overall efficiency.

### A. Test 3-D MPSoC

In this analysis, we use a simple illustrative two-tier 3-D MPSoC, where the top and bottom tiers are identical, and water is the coolant that flows through the microchannels. Within a single tier, the overall area ( $1 \times 1 \text{ cm}^2$ ) is split into four identical sections, each section is of an area ( $0.5 \times 0.5 \text{ cm}^2$ ). In addition, there are four power sources that have the power dissipation values [30, 50, 70, 90]W/cm<sup>2</sup>. These power sources are allocated to each tier. Thus, there are 24 different combinations (floorplans) to the target 3-D MPSoC. Since floorplans that are mirror images lateral to the direction of coolant flow are identical, we consider only 12 of the combinations in Fig. 7. In these figures, the direction of the coolant in the microchannels is from the bottom to the top of the figure. The structural parameters of this 3-D MPSoC are shown in Table I (see Fig. 5 for more details about these parameters). Obviously, some of these floorplans are better with respect to thermal considerations than the others. Our objective in this experiment is to study how channel width modulation affects

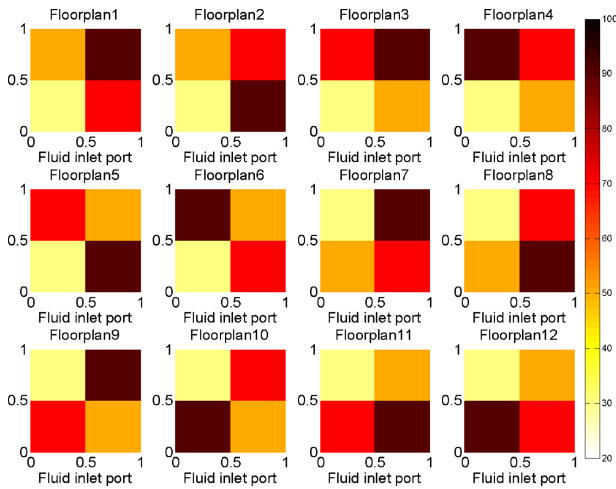


Fig. 7. Planar view of the 12 possible floorplans that can be generated when allocating the different power sources to the target 3-D MPSoC.

TABLE I  
SYSTEM PARAMETERS USED IN THE ANALYSIS

Symbol	Parameter	Value ( $\mu\text{m}$ )
$H_C$	Channel height	200
$H_{Si}$	Silicon tiers thickness	50
$w_{C,min}$	Minimum channel width	50
$w_{C,max}$	Maximum channel width	130
$W$	Channel pitch	200

thermally good and thermally inefficient designs within a joint architectural exploration.

**B. Channel Width Variation Impact**

First, we show in this section the thermal gradient and peak temperature impact when the channel width is varied. In this set of experiments, we fix the applied pressure drop to  $\Delta P = 4$  bar, and we use a uniform (nonmodulated) channel width. In addition, we simulate each floorplan with different channel widths  $W_C \in \{40, 50, \dots, 130\} \mu\text{m}$ , and we plot, for all the cases, the scatter graph between the peak temperature and the thermal gradient ( $\max(T) - \min(T)$ ) in Fig. 8.

We observe from Fig. 8 that inappropriate (suboptimal) selection of the channel width leads to undesired thermal response presented by high peak temperature and high thermal gradient. Irrespective of the floorplan, Fig. 8 shows a trend in high peak temperature associated with high thermal gradient when peak temperature is higher than  $60^\circ\text{C}$ . Thus, it is crucial that the channel widths are optimized for each floorplan to satisfy the thermal constraints.

In addition, we observe that floorplanning has an impact on thermal efficiency. Fig. 8 shows that each floorplan that achieves a low thermal gradient implies a low peak temperature such as floorplans 11 and 12. However, there are other floorplans that achieve low peak temperature (e.g., floorplans 2, 5, and 8) but with a relatively high thermal gradient. This is a clear indication that floorplans can be designed to be thermally efficient and when combined with optimized channel width, or even modulation, more optimized designs can be achieved. Moreover, this paper shows that a low thermal gradient is a more crucial and tighter constraint than the peak temperature constraint. Thus, in the following analysis, we

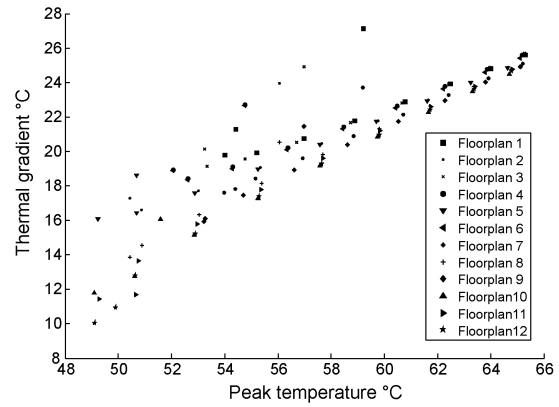


Fig. 8. Thermal gradient and peak temperature scatter graph of various floorplans in Fig. 7 with varying channel widths and applied pressure drop  $\Delta P = 4$  bar.

limit the search space for an optimized channel profile with a maximal peak temperature of  $60^\circ\text{C}$ .

**C. Channel Modulation Impact**

To study the impact and effectiveness for channel modulation, we run the optimized design method explained in Section IV under two scenarios.

- 1) We solve the problem in (15) by enforcing a uniform channel width. That is, instead of finding the channel width as a function of the distance  $z$  from inlet, we find a single value of channel width for all channels and the corresponding pumping power that minimizes the cost function in (15). This approach corresponds to a conventional microchannel design. We denote this scenario as *GreenCool without modulation* throughout the rest of this paper and serves as the reference for comparison with our proposed optimized channel modulation.
- 2) We solve the problem in (15) to perform our proposed optimized channel width modulation. This scenario will be referred to as the *GreenCool with modulation* throughout the rest of this paper. The thermal constraints in each case are defined as  $T_{\max} = 60^\circ\text{C}$  and  $\Delta T_{\max} = 10^\circ\text{C}$ .

To reduce the simulation time of computation of channel width profiles, we exploit the simplicity of the given floorplan and divide the area of the IC in Fig. 7 into two halves—the left and the right half, and group the microchannels under each half as a single unit. Each of this group lies directly below a unique set of two floorplan elements from inlet to outlet and hence, the channel width modulation method applied to them is affected mainly by that set of floorplan elements. Therefore, a single optimized channel width profiling is applied to all the microchannels in a group. Each microchannel can be treated individually at the cost of a much higher simulation time.

For each floorplan in Fig. 7, the energy spent on pumping the coolant for both the GreenCool without modulation and with modulation cases is shown in Table II. We limit the maximum pumping pressure applied to 10 bars, as indicated by our industry partners [8]. In Table II, we mark the operating points that exceed the maximum possible pressure drop with XX. Without modulation, we observe that there are significant differences between the thermal performances of the floor-



TABLE II  
POWER SAVINGS OBSERVED WHEN GRADIENT CONSTRAINT IS 10 °C AND PEAK TEMPERATURE CONSTRAINT IS 60 °C

Floorplan number	GreenCool Without Modulation			GreenCool With Modulation			Power Savings (% Value)
	$\Delta P$ (bars)	$V$ ml/min per cavity	Pumping power required (W)	$\Delta P$ (bars)	$V$ ml/min per cavity	Pumping power required (W)	
1	6.84	609	6.945	2.85	140	0.67	90.4
2	3.57	284	1.687	3.29	114	0.62	63.1
3	XX	XX	XX	2.83	155	0.73	98.4
4	XX	XX	XX	9.66	204	3.29	97.2
5	XX	XX	XX	2.83	155	0.73	96.9
6	XX	XX	XX	9.66	204	3.29	97.2
7	6.84	609	6.945	2.85	140	0.67	90.4
8	2.79	168	0.781	3.30	114	0.62	20.1
9	XX	XX	XX	2.83	155	0.73	98.4
10	4.74	652	5.513	2.49	91	0.38	92.6
11	2.58	95	0.408	2.51	81	0.34	17.4
12	2.62	106	0.463	2.44	80	0.32	29.9

Pressure drops that exceed the maximum allowable drop are marked by XX.

plans, i.e., some require very little pumping power owing to better thermal design and some require very large or even unfeasible pumping power due to poor thermal design. For example, floorplans 4 and 6 require very high pumping power, which is unfeasible. In these two floorplans, the high heat flux elements are placed near the outlet of the microchannels where the coolant is already hot. This aggravates the thermal response at these locations, raising the peak temperatures and thermal gradients considerably. Hence, coolant must be pumped at a very large flow rate to remove heat and maintain the temperatures below the constraints. Therefore, the pumping power increases significantly.

In contrast, GreenCool with modulation performs better compared to without modulation for every single floorplan. The savings in pumping power reach 98%. Moreover, GreenCool reduces the pumping pressure of all the cases to the allowable range (i.e., ten bars or less). It must be noted that the savings are higher for the poor thermal designs, such as floorplans 4 and 6. This is because channel modulation compensates for the rise in coolant temperatures near the outlets by customizing the cooling properties. Even for the case of better designs, such as floorplans 8, 11, and 12, there is still about 17–30% savings in the pumping power. In addition, the range of the required pumping power for all the floorplans is considerably reduced compared to the uniform channel width case. Thus, irrespective of how thermally optimal the floorplan is, optimized channel modulation manages to reduce cooling costs across the design space while meeting the thermal constraints. This implies that by utilizing GreenCool, thermal constraints during floorplanning can be relaxed and the layout can be optimized solely with other major design constraints such as performance and area.

## VI. TARGET 3-D MPSOC ARCHITECTURE AND PERFORMANCE MODELING METHODOLOGY

3-D MPSoCs with stacked DRAM are expected to be among the early commercial 3-D products to appear in the market. Performance improvements achieved by reducing the memory access latency through DRAM stacking are likely to cause corresponding increases in power and temperature, resulting in interesting performance-temperature tradeoffs. This section introduces our target 3-D MPSoC with stacked DRAM and discusses the modeling of memory access latency in detail.

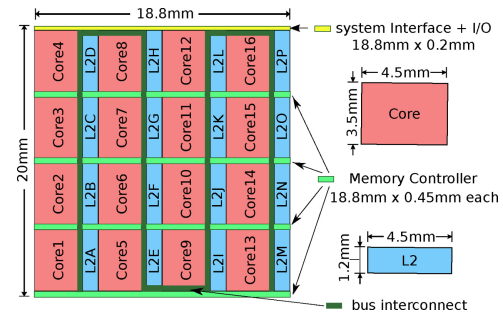


Fig. 9. Layout of the logic layer of the target 3-D system.




Fig. 10. Illustration of the 3-D system with DRAM stacking that has (a) single-bus memory access and (b) 4-way parallel memory access, respectively.

### A. Target System

Our target 3-D MPSoC is a DRAM-on-multicore architecture. All the processing cores and caches in our target 3-D MPSoC are on a logic layer, and the DRAM layer is stacked below it. We use TSVs for vertically connecting the logic and DRAM layers. We assume face-to-back, wafer-to-wafer bonding. We explore the target 3-D MPSoC with and without L2 caches. For memory-bounded benchmarks running on 3-D systems with high-bandwidth memory access, we expect the 3-D MPSoC without L2 cache to provide smaller area and lower cost without sacrificing performance.

We illustrate the floorplan of the logic layer of the 16-core 3-D MPSoC in Fig. 9. We model the core architecture based on the cores in AMD Magny-Cours processors, as described in our prior work [36]. We assume that the processor is manufactured using 45 nm technology, and that the 3-D MPSoC with L2 caches has a total die area of 376 mm<sup>2</sup>.

To simulate the data transfer between the logic layer and the DRAM layer in the 3-D MPSoC, we consider three different schemes: single-bus memory access, 4-way parallel memory access, and 8-way parallel memory access. As illustrated in Fig. 10, 4-way parallel scheme allows four on-chip memory controllers accessing the four DRAM ranks simultaneously.




Fig. 11. Six layouts for modeling on-chip wire delay of target 3-D systems. The green blocks represent the memory controllers.

TABLE III  
ON-CHIP WIRE DELAY

Floorplan	Average Distance (mm)	Wire Delay (ns)
CaseI_X, CaseI_Y	12.5	4.58
CaseII_X, CaseII_Y	13.75	5.03
CaseIII_X, CaseIII_Y	14.375	5.26

*B. Modeling Memory Access Latency*

For MPSoCs, the access latency from the last-level caches (LLC) to main memory is composed of three components: 1) the propagation delay between the LLC to the memory controller (LLC-to-controller delay); 2) the data request time spent at the memory controller (memory controller processing latency); and 3) the data retrieval time spent at the DRAM.

In order to model the LLC-to-controller delay, we assume that all the last-level L1 or L2 caches are connected to the memory controllers through a shared bus. The global bus interconnect is routed around the chip in a serpentine fashion, as illustrated in Fig. 9. We use energy-optimized repeater-inserted pipelined channels to model the on-chip bus interconnect [37]. The wire propagation delay is linear with respect to the wire length. We assume each pipeline stage has a propagation delay of 183 ps per mm [38].

We explore the average LLC to memory controller distance for the six different layouts as shown in Fig. 11. The memory controllers in CaseI, CaseII, and CaseIII have one, two, and four memory controller blocks, respectively. The average LLC to memory controller distance and the corresponding round-trip wire delay for all the six layouts are listed in Table III. The wire delay in Table III is obtained by multiplying the round-trip average distance by the propagation delay per mm.

The memory controller processing latency is strongly affected by the memory request queuing delay [39], that is, the time spent by a memory request waiting to get scheduled. We apply queuing theory to model the memory controller queuing delay, where the memory request queue is modeled as a M/D/1 queuing system. In M/D/1 queuing formula, the queuing delay depends on two parameters: arrival rate and service rate. We estimate the service rate by considering the DRAM access time ( $t_{RAS}$  and  $t_{RP}$ ) and the availability of parallel memory access in the 3-D MPSoC. We set the row active time  $t_{RAS} = 36$  ns and row precharge time  $t_{RP} = 15$  ns as reported by MICRON’s

TABLE IV  
MAIN MEMORY ACCESS LATENCY FOR THE 3-D MPSOC WITH SINGLE-BUS MEMORY ACCESS

LLC-to-MC	5 ns LLC-to-controller delay
Memory controller	25 ns memory controller processing time
Main Memory	On-chip 1 GB DRAM $t_{RAS} = 36$ ns, $t_{RP} = 15$ ns
Total delay	Total = 81 ns
Memory bus	On-chip memory bus, 2 GHz, 64 Byte bus width




Fig. 12. Memory request queuing delay in different memory access schemes. Average access rates of 0.0035, 0.012, and 0.025 are obtained by simulating single-bus, 4-way parallel, and 8-way parallel access schemes, respectively.

DDR3 SDRAM [40]. Thus, we model the memory request queue service rate for the 3-D MPSoC with single-bus access as 0.02 per cycle. We assume that the service rate is four times and eight times of the service rate of the single bus access with 4- and 8-way parallel memory access, respectively. Fig. 12 presents the queuing delay of the memory requests.

DRAM access latency consists of address decoding time, column and row active time, and data transfer time. In our 3-D MPSoC, we consider a 1 GB DRAM which consists of four ranks, each of which has four banks. We apply the same timing parameters for the DRAM layer of the target 3-D MPSoC as in MICRON’s SDRAM [40], which is consistent with the assumptions used in earlier studies [41], [42]. The wire delay on the DRAM layer is not explicitly modeled in this study, since the intrarank delay is already taken into account in the DRAM row active time. Table IV summarizes the access times for the 3-D MPSoC with single-bus memory access.

Table V lists the bus width, number of TSVs, and the total area of TSVs for the 3-D MPSoCs with the single-bus, 4-way, and 8-way parallel memory access schemes. In our experiments, we use TSVs with a diameter of 50  $\mu$ m and a center-to-center pitch of 100  $\mu$ m. We select these TSV parameters as they are verified to be compatible with interlayer liquid cooling in manufactured prototypes [6], [8], [18].

VII. GREENCOOL ON 3-D MPSOC: EXPERIMENTAL SETUP

*A. Performance Simulation*

In our simulation framework, M5 full-system simulator [43] with Alpha instruction set architecture is used for simulating the performance of our target systems. We model the target 3-D MPSoC in M5 by configuring the main memory access latency and the bus width/speed between the last level caches and main memory to mimic the high data transfer bandwidth provided by the TSVs. Table IV lists the memory access latencies used in the performance simulations.

We select four parallel applications from the PARSEC benchmarks suite [44] as our target workloads: bodytrack, canneal, ferret, and streamcluster. We run the PARSEC benchmarks in M5 with sim-large input data sets in the region of interest (ROI). We execute each benchmark with the detailed

TABLE V  
BUS AND TSV CONFIGURATIONS

Memory access	Bus width (byte)	TSV Number	TSV Area (mm <sup>2</sup> )
Single-bus	64	512	5.78
4-way parallel	256	2048	21.2
8-way parallel	512	4096	41.8

TABLE VI  
SYSTEM PARAMETERS USED IN THE EXPERIMENTS

Symbol	Parameter	Value
$T_{Cin}$	Fluid inlet temperature	27 °C
$T_{max}$	Peak temperature constraint	60 °C
$H_C$	Channel height	100 $\mu$ m
$H_{Si}$	Silicon tiers thickness	50 $\mu$ m
$w_{C,min}$	Minimum channel width	30 $\mu$ m
$w_{C,max}$	Maximum channel width	80 $\mu$ m
$W$	Channel pitch	100 $\mu$ m

out-of-order CPUs with accurate memory access simulations, and collect statistics at every 10ms for 100 sampling steps or until the end of ROI. The performance statistics collected from M5 simulations are used as inputs to the processor power model. At every sampling interval, we also track the number of L2 cache accesses and the memory accesses to each DRAM bank. The access statistics are then fed into cache and DRAM power models to generate the power traces for the L2 caches and DRAM banks.

### B. Power Model

In order to obtain the run-time dynamic power for the processing cores, we use McPAT 0.7 [45] for 45 nm process. We set  $V_{dd}$  to 1.1 V and operating frequency to 2.1 GHz for our target systems in the McPAT simulations. After getting the run-time dynamic core power from McPAT, we calibrate them using the average core power from measurements on an AMD Magny-Cours processor. We derive a calibration factor by correlating the average dynamic power computed by McPAT with the measured power. We then use this factor to scale the dynamic core power consumption of each benchmark. The average dynamic core power across the benchmarks in the 3-D MPSoC with single bus access is 6.18 W, while the average core power is 7.7 W in the 3-D MPSoC with 8-way parallel memory access without L2 caches. The core leakage power is 2.63 W at nominal temperature of 343 K. As we modulate the on-chip temperature through liquid cooling, range of thermal variations is limited. Thus, we use a fixed leakage value without explicitly modeling the temperature dependence.

The L2 cache power is calculated using Cacti 5.3 [46]. The dynamic L2 power is scaled using L2 cache access rates from the performance simulations. The average L2 cache power is 0.62W and the leakage power for L2 cache is 0.45 W. The DRAM power in the 3-D system is calculated using MICRON's DRAM power calculator [47], which takes the memory read and write access rates as inputs. We obtain detailed DRAM power traces for each of the DRAM banks at every sampling interval. The average on-chip DRAM bank power across all the benchmarks in 3-D MPSoC with single-bus access is 1.19 W, while in 3-D MPSoC with 8-way parallel memory access without L2 cache is 3.33 W.

The total on-chip memory controller power for the target 3-D MPSoCs is estimated based on existing memory controllers

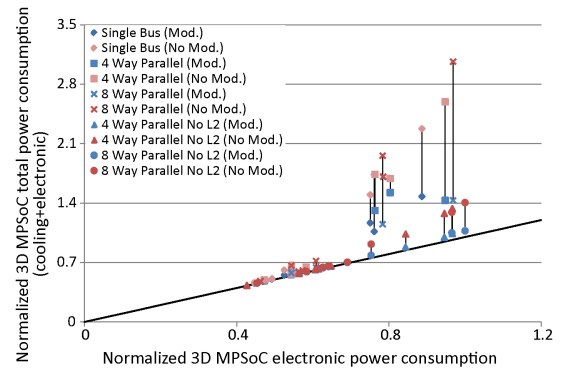


Fig. 13. 3-D MPSoC power efficiency for GreenCool with and without modulation cases when  $\Delta T_{max} = 10$  °C.

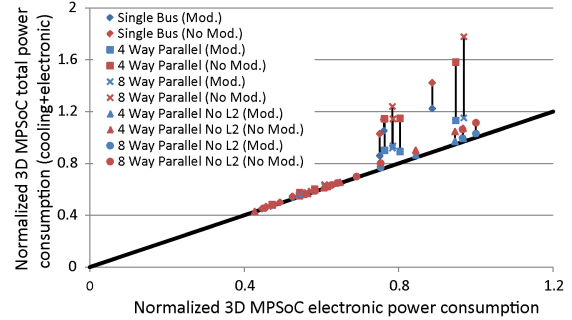


Fig. 14. 3-D MPSoC power efficiency for GreenCool with and without modulation cases when  $\Delta T_{max} = 12$  °C.

for many-core systems [48] as 5.9 W. Thus, the system interface and I/O power as well as the on-chip bus power are negligible with respect to the total chip power. It has been shown that the total on-chip bus power for running PARSEC and NAS workloads is less than 2.0 W even for a 64-core system [37].

## VIII. GREENCOOL ON 3-D MPSOC: RESULTS

In this section, we apply GreenCool with and without modulation to the aforementioned five architectures. Each architecture is simulated with eight different power traces (using peak and average power consumption for each of the 4 benchmarks), which makes a total of 40 case studies. First, we simulate air cooling for the same 3-D MPSoC architectures in HotSpot 5.0 [22] using default package parameters of HotSpot. The resulting peak temperatures are between 77 °C and 99 °C. In our experiments, we use (single-phase) water as the coolant.

Next, we explore the impact of thermal gradients on GreenCool performance for the 40 cases using four different thermal gradient constraints in (19), i.e.,  $\Delta T_{max} \in \{10, 12, 15, 20\}$  °C. We show the system parameters and the constraints in Table VI. When we apply GreenCool to the target 3-D MPSoC, we limit the peak temperature to  $T_{max} = 60$  °C. This choice is based on our analysis in Section V-B, where we find that the feasible region for GreenCool for meeting the given thermal gradient constraints is when the peak temperature is below 60 °C.

Figs. 13–15 use scatter plots to illustrate the 3-D MPSoC power efficiency obtained for thermal gradient constraints of 10, 12, and 15 °C. In these figures, we plot the 3-D MPSoC total power consumption (computational power and cooling) against solely the computational power consumption (cores,




Fig. 15. 3-D MPSoC power efficiency for GreenCool with and without modulation cases when  $\Delta T_{\max} = 15^\circ\text{C}$ .

L2 cache, DRAM, and interconnect). For better visualization, all the power values in this plot are normalized to the maximum computational power consumption observed in all the power simulations in Section VII-B ( $P_{\max} = 463\text{ W}$ ). In each of these scatter plots, a different marker is used for each of the five target architectures. In addition, marker colors, blue, and red, are used to distinguish between the results obtained from GreenCool with modulation and without modulation.

To quantify the energy efficiency based on these figures, we use a metric called the power usage effectiveness (PUE), which is a metric often used for quantifying the efficiency of data centers. In this paper, we define PUE as follows:

$$\text{PUE} = \frac{\text{3-D MPSoC total power consumption}}{\text{3-D MPSoC computational power consumption}}. \quad (20)$$

PUE can be visualized in Figs. 13–15 as the ratio of the y-coordinate to the x-coordinate for each data point. Ideally, PUE must be equal to 1, representing a case where no cooling effort is needed at all, and all the energy spent is purely for computational performance. This ideal power efficiency is represented by the solid black straight line in Figs. 13–15. Based on (20) and Figs. 13–15, we make the following observations.

First, for each simulation, GreenCool with modulation exhibits higher energy efficiency than without modulation. GreenCool manages to enhance PUE by up to 54% and on average by 11.4% when  $\Delta T_{\max} = 10^\circ\text{C}$ . Similar energy efficiency improvement trends are observed when  $\Delta T_{\max} = 12^\circ\text{C}$  (6% average 35% peak savings) and when  $\Delta T_{\max} = 15^\circ\text{C}$  (2.3% average and 14% peak savings). The peak temperatures and gradients satisfy the constraints as follows:

$$\begin{aligned} \max(\mathbf{T}) &\in [39, 49]^\circ\text{C} \text{ when } \Delta T_{\max} = 10^\circ\text{C} \\ \max(\mathbf{T}) &\in [42, 51]^\circ\text{C} \text{ when } \Delta T_{\max} = 12^\circ\text{C} \\ \max(\mathbf{T}) &\in [45, 54]^\circ\text{C} \text{ when } \Delta T_{\max} = 15^\circ\text{C} \\ \max(\mathbf{T}) &\in [50, 60]^\circ\text{C} \text{ when } \Delta T_{\max} = 20^\circ\text{C} \end{aligned}$$

Second, changing the thermal gradient constraint has a very significant impact on the PUE. Fig. 13 shows the poor efficiency of the target 3-D MPSoC, especially for high power consumption cases, when we constrain the maximum thermal gradient to only  $10^\circ\text{C}$ . This is because a significant portion of the total energy is used in the cooling system to satisfy this constraint, which indicates either a suboptimal thermal floorplan design or the difficulty of meeting this constraint on

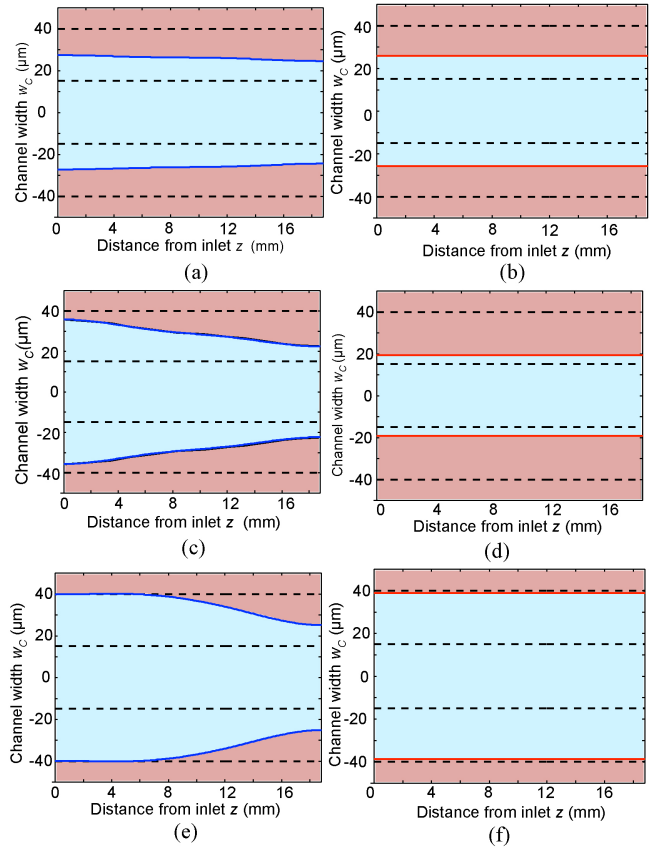


Fig. 16. Channel width profile obtained, by applying GreenCool (a), (c), (e) with modulation, and (b), (d), (f) without modulation for 4-way parallel access architecture running ferret. (a) Modulation  $\Delta T_{\max} = 10^\circ\text{C}$ . (b) No modulation  $\Delta T_{\max} = 10^\circ\text{C}$ . (c) Modulation  $\Delta T_{\max} = 12^\circ\text{C}$ . (d) No modulation  $\Delta T_{\max} = 12^\circ\text{C}$ . (e) Modulation  $\Delta T_{\max} = 15^\circ\text{C}$ . (f) No modulation  $\Delta T_{\max} = 15^\circ\text{C}$ .

the target system. However, Fig. 15 shows significant PUE gain for these high power consumption cases, due to the relaxed thermal gradient constraint. Obviously, relaxed constraint results in larger thermal gradients, and hence may not represent an effective use of GreenCool. Fig. 14 represents an optimum tradeoff point between the two cases, as we maintain small gradients and enhance PUE by reducing the cooling cost. In these experiments, we see that GreenCool improves the energy efficiency even under tight thermal constraints.

We also examine the impact of GreenCool with large thermal gradients, i.e.,  $\Delta T_{\max} = 20^\circ\text{C}$ . In this case, GreenCool manages to reduce the pumping power by up to 41% with respect to GreenCool without modulation. In fact, the optimum channel width in this case is determined as the maximum width constraint we define in Table VI. However, we observe that the enhancement in the PUE in this case is 2%. This slight enhancement indicates that both GreenCool with and without modulation achieve low PUE values (PUE = 1.02 with modulation); thus, both techniques are energy-efficient.

Third, we observe performance differences among the various architectures. While running benchmarks with low or average power consumption, all the explored architectures are energy-efficient. However, architectures shared, 4-way, and 8-way parallel show extremely poor PUE at high power consumption values, whereas architectures 4-way and 8-way parallel No-L2 do not deviate much from PUE = 1.0 for any of

the power profiles. The observed results are counter-intuitive as due to the presence of L2 cache, the overall power density is reduced. Hence, the required pumping power for the cases with L2 cache should be lower than those without L2 cache. However, these results are primarily dominated by the effects of the die areas of the explored architectures. The larger die area, hence longer channel, in architectures shared, 4-way, and 8-way parallel causes the coolant to travel a longer distance along the IC, leading to more sensible heat accumulation compared to that of the smaller die area in architectures 4-way and 8-way parallel No-L2. To compensate for this increased sensible heat accumulation, coolant flow rate and the pumping power must be significantly increased. Moreover, the pumping power for the fluid to flow with the same flow rate in longer channels is larger than the case of smaller channels. Compared to 4-way parallel No-L2, 4-way parallel architecture requires, on average, 250% increase in pressure drop (from 3.62 to 12.9 bars) for GreenCool with modulation and 161% without modulation (from 4.63 to 13.23 bars) to satisfy the thermal gradient constraint  $\Delta T_{\max} = 10^\circ\text{C}$ . Similarly, compared to 8-way parallel No-L2, 8-way parallel requires about 450% higher pressure drop to satisfy the thermal constraints. 8-way parallel No-L2 has 5.2% higher throughput (instructions per second) compared to 8-way parallel, owing to the high DRAM access bandwidth and low latency in the 3-D architecture. Thus, removing the L2 caches in 3-D MPSoCs with stacked DRAM reduces the cooling cost in addition to potential benefits in design cost and complexity, while maintaining high performance.

In Fig. 16, we show the modulated channel width profile obtained after running GreenCool with and without modulation on the 4-way parallel access architecture running ferret (peak power). As our target 3-D MPSoC has a homogeneous architecture, GreenCool provides the same channel profile for all the microchannels. However, our previous work shows that GreenCool can also handle heterogeneous architectures with large power dissipation variations [12]. Our results in Section V also illustrate that GreenCool achieves outstanding results for heterogeneous floorplans. Overall, GreenCool provides the optimum channel profile regardless of the specific thermal gradient constraint. The resulting channel shape, however, varies depending on the gradient constraints. For high power benchmarks (as in Fig. 16), we observe lower variation in channel widths for the low thermal gradient  $\Delta T_{\max} = 10^\circ\text{C}$ . This is primarily due to the inclusion of the flow rate  $\dot{V}$  as a optimization parameter and the fact that we primarily target minimizing the pumping energy. In some experiments the search space is narrower due to the stringent thermal constraints ( $\Delta T_{\max} = 10^\circ\text{C}$ ), hence, limiting the feasible channel width profiles to a smaller set. When constraints are relaxed, the feasible channel width profile is augmented leading to a more optimum operating point.

## IX. CONCLUSION

In this paper, we presented GreenCool, an energy-efficient liquid cooling optimization method for thermal balancing and energy reduction. We showed the impact of channel width modulation on convective heat transfer coefficients. In addition, we formulated an optimal design problem to

minimize the pumping energy by utilizing channel modulation, subject to peak temperature and maximum thermal gradient constraints. We compared GreenCool against applying the same optimization procedure, but using fixed-width channels (i.e., no modulation). When exploring GreenCool under various floorplanning approaches, we showed that GreenCool reduces cooling power by up to 98% with respect to no modulation. In addition, GreenCool augments the benefits of thermally-aware floorplanning, facilitating better floorplan optimization for other design parameters. We also conducted detailed experiments on a 3-D MPSoC with stacked DRAM, for which GreenCool achieves up to 53% energy-efficiency improvement compared to the no modulation case.

## REFERENCES

- [1] Y. Xie, G. H. Loh, B. Black, and K. Bernstein, "Design space exploration for 3-D architectures," *ACM J. Emerging Technol. Comput. Syst.*, vol. 2, no. 2, pp. 65–103, 2006.
- [2] K. Bernstein *et al.*, "Interconnects in the third dimension: Design challenges for 3D-ICs," in *Proc. IEEE/ACM Design Automat. Conf.*, Jun. 2007, pp. 562–567.
- [3] B. Black, M. Annavaram, N. Brekelbaum, J. DeVale, L. Jiang, G. H. Loh, D. McCaule, P. Morrow, D. W. Nelson, D. Pantoso, P. Reed, J. Rupley, S. Shankar, J. Shen, and C. Webb, "Die stacking (3-D) microarchitecture," in *Proc. IEEE/ACM Int. Symp. Microarchit.*, Dec. 2006, pp. 469–479.
- [4] D. Sacchetto, M. Zervas, Y. Temiz, G. D. Micheli, and Y. Leblebici, "Resistive programmable through silicon vias for reconfigurable 3-D fabrics," *IEEE Trans. Nanotechnol.*, vol. 11, no. 1, pp. 8–11, Jan. 2012.
- [5] A. K. Coskun, J. Meng, D. Atienza, and M. M. Sabry, "Attaining single-chip, high-performance computing through 3-D systems with active cooling," *IEEE Micro*, vol. 31, no. 4, pp. 63–75, Jul.–Aug. 2011.
- [6] M. M. Sabry *et al.*, "Energy-efficient multiobjective thermal control for liquid-cooled 3-D stacked architectures," *IEEE Trans. Comput.-Aided Design*, vol. 30, no. 12, pp. 1883–1896, Dec. 2011.
- [7] T. Brunschweiler *et al.*, "Interlayer cooling potential in vertically integrated packages," *Microsyst. Technol.*, vol. 15, no. 1, pp. 57–74, 2009.
- [8] F. Alfieri *et al.*, "3-D integrated water cooling of a composite multilayer stack of chips," *J. Heat Transfer*, vol. 132, no. 12, pp. 121402-1–121402-9, 2010.
- [9] B. Dang *et al.*, "Integrated microfluidic cooling and interconnects for 2-D and 3-D chips," *IEEE Trans. Adv. Packag.*, vol. 33, no. 1, pp. 79–87, Feb. 2010.
- [10] JEDEC: jep122E Failure Mechanisms and Models for Semiconductor Devices (2011, Oct.) [Online]. Available: <http://www.jedec.org/standards-documents/docs/jep-122e>
- [11] A. K. Coskun, T. S. Rosing, and K. C. Gross, "Utilizing predictors for efficient thermal management in multiprocessor SoCs," *IEEE Trans. Comput.-Aided Design*, vol. 28, no. 10, pp. 1503–1516, Oct. 2009.
- [12] M. M. Sabry, A. Sridhar, and D. Atienza, "Thermal balancing of liquid-cooled 3-D MPSoCs using channel modulation," in *Proc. Des. Autom. Test Eur.*, 2012, pp. 599–604.
- [13] T. Brunschweiler *et al.*, "Hotspot-optimized interlayer cooling in vertically integrated packages," in *Proc. MRS Fall Meeting*, Dec. 2009, pp. 223–234.
- [14] C. Zhu *et al.*, "Three-dimensional chip-multiprocessor run-time thermal management," *IEEE Trans. Comput.-Aided Design*, vol. 27, no. 8, pp. 1479–1492, Aug. 2008.
- [15] A. K. Coskun *et al.*, "Energy-efficient variable-flow liquid cooling in 3-D stacked architectures," in *Proc. Des. Autom. Test Eur.*, Mar. 2010, pp. 111–116.
- [16] F. Zanini, M. M. Sabry, D. Atienza, and G. D. Micheli, "Hierarchical thermal management policy for high-performance 3d systems with liquid cooling," *IEEE J. Emerging Sel. Top. Circuits Syst.*, vol. 1, no. 2, pp. 88–101, Jun. 2011.
- [17] H. Qian *et al.*, "Cyber-physical thermal management of 3-D multicore cache-processor system with microfluidic cooling," *ASP J. Low Power Electron.*, vol. 7, no. 1, pp. 1–12, 2011.
- [18] T. Brunschweiler *et al.*, "Angle-of-attack investigation of pin-fin arrays in nonuniform heat-removal cavities for interlayer cooled chip stacks," in *Proc. IEEE Symp. Semicond. Thermal Measurement Manag.*, Mar. 2011, pp. 116–124.
- [19] D. B. Tuckerman and R. F. W. Pease, "High-performance heat sinking for VLSI," *IEEE Electron Device Lett.*, vol. 2, no. 5, pp. 126–129, May 1981.
- [20] T. Brunschweiler *et al.*, "Direct liquid-jet impingement cooling with micron-sized nozzle array and distributed return architecture," in *Proc. IEEE ITherm*, May–Jun. 2006, pp. 196–203.
- [21] Y. J. Kim, Y. K. Joshi, A. G. Fedorov, Y. J. Lee, and S. K. Lim, "Thermal characterization of interlayer microfluidic cooling of three-dimensional integrated circuits with nonuniform heat flux," *J. Heat Transfer*, vol. 132, no. 4, pp. 041009-1–041009-9, 2010.
- [22] K. Skadron, M. R. Stan, W. Huang, V. Sivakumar, S. Karthik, and D. Tarjan, "Temperature-aware microarchitecture," in *Proc. Int. Symp. Comput. Archit.*, Jun. 2003, pp. 2–13.
- [23] A. K. Coskun, J. Ayala, D. Atienza, and T. S. Rosing, "Modeling and dynamic management of 3-D multicore systems with liquid cooling," in *Proc. Int. Conf. Very Large Scale Integr.-SoC*, Oct. 2009, pp. 60–65.

- [24] H. Mizunuma, C. L. Yang, and Y. C. Lu, "Thermal modeling for 3D-ICs with integrated microchannel cooling," in *Proc. Int. Conf. Comput.-Aided Des.*, Nov. 2009, pp. 256–263.
- [25] A. Sridhar, A. Vincenzi, M. Ruggiero, T. Brunschwiler, and D. Atienza. (2010). *3D-ICE: Fast Compact Transient Thermal Modeling for 3D-ICs With Inter-Tier Liquid Cooling* [Online]. Available: <http://esl.epfl.ch/3d-ice.html>
- [26] A. Sridhar, A. Vincenzi, M. Ruggiero, T. Brunschwiler, and D. Atienza, "Compact transient thermal model for 3D-ICs with liquid cooling via enhanced heat transfer cavity geometries," in *Proc. THERMINIC*, Oct. 2010, pp. 1–6.
- [27] Z. Feng and P. Li, "Fast thermal analysis on GPU for 3D-ICs with integrated microchannel cooling," in *Proc. Int. Conf. Comput.-Aided Des.*, Nov. 2010, pp. 551–555.
- [28] J. Cong, J. Wei, and Y. Zhang, "A thermal-driven floorplanning algorithm for 3D-ICs," in *Proc. Int. Conf. Comput.-Aided Des.*, Nov. 2004, pp. 306–313.
- [29] K. Sankaranarayanan, S. Velusamy, M. Stan, and K. Skadron, "A case for thermal-aware floorplanning at the microarchitectural level," *J. Instruction-Level Parallelism*, vol. 8, pp. 1–16, Oct. 2005.
- [30] M. Healy *et al.*, "Multiobjective microarchitectural floorplanning for 2-D and 3-D ICs," *IEEE Trans. Comput.-Aided Des.*, vol. 26, no. 1, pp. 38–52, Jan. 2007.
- [31] W.-L. Hung *et al.*, "Interconnect and thermal-aware floorplanning for 3-D microprocessors," in *Proc. Int. Symp. Quality Electronic Design*, Mar. 2006, pp. 98–104.
- [32] H. Qian, C. Chang, and H. Yu, "An efficient channel clustering and flow rate allocation algorithm for non-uniform microfluidic cooling of 3-D integrated circuits," *Integr. VLSI J.*, vol. 46, no. 1, pp. 57–68, 2013.
- [33] B. Shi *et al.*, "Non-uniform micro-channel design for stacked 3D-ICs," in *Proc. Design Autom. Conf.*, Jun. 2011, pp. 658–663.
- [34] H. Mizunuma *et al.*, "Thermal modeling and analysis for 3D-ICs with integrated microchannel cooling," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 30, no. 9, pp. 1293–1306, Sep. 2011.
- [35] R. Shah and A. London, *Laminar Flow Forced Convection in Ducts*. New York: Academic Press, 1978.
- [36] J. Meng, K. Kawakami, and A. K. Coskun, "Optimizing energy efficiency of 3-D multicore systems with stacked dram under power and thermal constraints," in *Proc. Design Autom. Conf.*, Jun. 2012, pp. 648–655.
- [37] J. Meng, C. Chen, A. K. Coskun, and A. Joshi, "Run-time energy management of manycore systems through reconfigurable interconnects," in *Proc. ACM/IEEE Great Lakes Symp. VLSI*, May 2011, pp. 43–48.
- [38] Y. Jin, K. H. Yum, and E. J. Kim, "Adaptive data compression for high-performance low-power on-chip networks," in *Proc. IEEE/ACM Int. Symp. Microarchit.*, Nov. 2008, pp. 354–363.
- [39] M. Awasthi *et al.*, "Handling the problems and opportunities posed by multiple on-chip memory controllers," in *Proc. PACT*, Sep. 2010, pp. 319–330.
- [40] Micron Technology, Inc. (2011). *DRAM Component Datasheet* [Online]. Available: <http://www.micron.com/http://www.micron.com>
- [41] G. H. Loh, "3D-stacked memory architectures for multicore processors," in *Proc. Int. Symp. Comput. Archit.*, Jun. 2008, pp. 453–464.
- [42] G. L. Loi, B. Agrawal, N. Srivastava, S.-C. Lin, T. Sherwood, and K. Banerjee, "A thermally-aware performance analysis of vertically integrated (3-D) processor-memory hierarchy," in *Proc. Design Autom. Conf.*, July 2006, pp. 991–996.
- [43] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. Shoib, N. Vaish, M. D. Hill, and D. A. Wood, "The gem5 simulator," *SIGARCH Comput. Archit. News*, vol. 39, no. 2, pp. 1–7, 2011.
- [44] C. Bienia, "Benchmarking modern multiprocessors," Ph.D. dissertation, Dept. Comput. Sci., Princeton Univ., Princeton, NJ, Jan. 2011.
- [45] S. Li *et al.*, "Mcpat: An integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Proc. IEEE Symp. Microarchit.*, Dec. 2009, pp. 469–480.
- [46] D. Tarjan, S. Thozioyoor, and N. P. Jouppi, "CACTI 4.0," HP Labs., Palo Alto, CA, Tech. Rep. HPL-2006-86, 2006.
- [47] Micron Technology, Inc. (2011). *DRAM Power Calculations* [Online]. Available: <http://www.micron.com/http://www.micron.com/>
- [48] J. Howard, S. Dighe, Y. Hoskote, S. Vangal, D. Finan, G. Ruhl, D. Jenkins, H. Wilson, N. Borkar, G. Schrom, F. Paillet, S. Jain, T. Jacob, S. Yada, S. Marella, P. Salihundam, V. Erraguntla, M. Konow, M. Riepen, G. Droege, J. Lindemann, M. Gries, T. Apel, K. Henriss, T. Lund-Larsen, S. Steibl, S. Borkar, V. De, R. Van Der Wijngaart, and T. Mattson, "A 48-core 1A-32 message-passing processor with DVFS in 45 nm CMOS," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2010, pp. 108–109.



**Mohamed M. Sabry** (S'12) received the B.S. (Hons.) and M.S. degrees in electrical engineering from Ain Shams University, Cairo, Egypt, in 2005 and 2008, respectively. He is currently pursuing the Ph.D. degree in electrical engineering with the Department of Electrical Engineering, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland.

He is currently a member of the Embedded Systems Laboratory, EPFL. His current research interests include system design and resource management methodologies in embedded systems, and multiprocessor system-on-chips


(MPSoCs), especially temperature and reliability management of 2-D and 3-D MPSoCs.

Mr. Sabry was a recipient of the Best Student Award when he was pursuing the B.S. degree.



**Arvind Sridhar** (S'07) received the B.Eng. degree in electronics and communication engineering from the College of Engineering Guindy, Anna University, Chennai, India, in 2006, and the M.A.Sc. degree in electronics from Carleton University, Ottawa, ON, Canada, in 2009. He is currently pursuing the Doctoral degree with the Embedded Systems Laboratory, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland.

He was a Research Scholar with the Computer-Aided Design Laboratory, Carleton University, from 2006 to 2009, and an Intern with the Advanced Thermal Packaging Group, IBM Research, Zurich, Switzerland, in 2011. He is an author of 3-D-ICE, the first compact transient thermal simulator for 2-D/3-D ICs with liquid cooling, which is currently being used by researchers in more than 100 universities and laboratories worldwide.



**Jie Meng** (S'12) received the M.S. degree in electrical engineering from McMaster University, Hamilton, ON, Canada. She is currently pursuing the Ph.D. degree with the Department of Electrical and Computer Engineering, Boston University, Boston, MA.

Her current research interests include multicore and 3-D stacked architectures, focusing on energy awareness and performance improvement for future systems.



**Ayse K. Coskun** (M'06) received the M.S. and Ph.D. degrees in computer science and engineering from the University of California, San Diego.

She is currently an Assistant Professor with the Department of Electrical and Computer Engineering, Boston University (BU), Boston, MA. She was with Sun Microsystems (now Oracle), San Diego, prior to her current position with BU. Her current research interests include energy-efficient computing, multicore systems, 3-D stack architectures, computer architectures, and embedded systems and software.

Dr. Coskun received the Best Paper Award at the IFIP/IEEE VLSI-SoC Conference in 2009 and at the HPEC Workshop in 2011. She currently serves on program committees of many design automation conferences, including DAC, DATE, GLSVLSI, and VLSI-SoC. She has served as a Guest Editor for the ACM TODAES journal. She is a member of ACM.



**David Atienza** (M'05) received the M.S. degree from the Complutense University of Madrid, Madrid, Spain, and the Ph.D. degree from the Inter-University Microelectronics Center, Leuven, Belgium, in 2001 and 2005, respectively, both in computer science and engineering.

He is currently a Professor of electrical engineering and the Director of the Embedded Systems Laboratory, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland. His current research interests include system-level design methodologies for high-performance multiprocessor system-on-chip (MPSoC) and embedded systems, including new 2-D/3-D thermal-aware design for MPSoCs, wireless body sensor networks, HW/SW reconfigurable systems, dynamic memory optimizations, and network-on-chip design. He is a co-author of more than 180 publications in peer-reviewed international journals and conferences, several book chapters, and six U.S. patents in these fields.

Dr. Atienza was the recipient of the ACM Special Interest Group on Design Automation Outstanding New Faculty Award in 2012, the Best Paper Award at the Very-Large-Scale Integration-System on Chip in 2009, and four Best Paper Award Nominations at the HPCS 2012, WEHA-HPCS 2010, ICCAD 2006, and DAC 2004 conferences. He is an Associate Editor of the IEEE TRANSACTION ON COMPUTER-AIDED DESIGN OF CIRCUITS AND SYSTEMS, and Elsevier *Integration*. He has been a member of the Executive Committee of the IEEE Council on EDA since 2008. He was a GOLD Member of the Board of Governors of the IEEE Circuits and Systems Society from 2010 to 2012.