

3D-ICE: A Compact Thermal Model for Early-Stage Design of Liquid-Cooled ICs

Arvind Sridhar, *Student Member, IEEE*, Alessandro Vincenzi, *Student Member, IEEE*, David Atienza, *Senior Member, IEEE*, and Thomas Brunschweiler, *Senior Member, IEEE*

Abstract—Liquid-cooling using microchannel heat sinks etched on silicon dies is seen as a promising solution to the rising heat fluxes in two-dimensional and stacked three-dimensional integrated circuits. Development of such devices requires accurate and fast thermal simulators suitable for early-stage design. To this end, we present 3D-ICE, a compact transient thermal model (CTTM), for liquid-cooled ICs. 3D-ICE was first advanced incorporating the 4-resistor model-based CTTM (4RM-based CTTM). Later, it was enhanced to speed up simulations and to include complex heat sink geometries such as pin fins using the new 2 resistor model (2RM-based CTTM). In this paper, we extend the 3D-ICE model to include liquid-cooled ICs with multi-port cavities, i.e., cavities with more than one inlet and one outlet ports, and non-straight microchannels. Simulation studies using a realistic 3D multiprocessor system-on-chip (MPSoC) with a 4-port microchannel cavity highlight the impact of using 4-port cavity on temperature and also demonstrate the superior performance of 2RM-based CTTM compared to 4RM-based CTTM. We also present an extensive review of existing literature and the derivation of the 3D-ICE model, creating a comprehensive study of liquid-cooled ICs and their thermal simulation from the perspective of computer systems design. Finally, the accuracy of 3D-ICE has been evaluated against measurements from a real liquid-cooled 3D-IC, which is the first such validation of a simulator of this genre. Results show strong agreement (average error < 10%), demonstrating that 3D-ICE is an effective tool for early-stage thermal-aware design of liquid-cooled 2D-/3D-ICs.

Index Terms—Liquid-cooling of ICs, thermal modeling, 3D-ICs

1 INTRODUCTION

THREE dimensional stacking of multiprocessor system-on-chips (3D MPSoCs), integrated using high-speed through-silicon vias (TSVs), possesses immense potential in accelerating the computational power of high performance servers and datacenters of the future [4]. But this vertical integration of CMOS circuits in the long-term is impeded by the inability of the current air-cooled heat sinks in handling rising heat fluxes [5]. On a larger-scale this thermal problem of ICs translates to the ever increasing cooling energy costs of today's datacenters [6]. Liquid cooling of integrated circuits using interlayer microchannel heat sinks etched directly on the IC has been advanced as one of the most promising solutions to this problem [7]-[9]. In addition to enabling further integration of CMOS circuits while maintaining safe operating temperatures, liquid cooling also increases cooling efficiency and enables energy harvesting in datacenters. Hence, it is touted as a long-term green energy solution for next-generation datacenters [10].

Prototypes of 2D and 3D stacked ICs with microchannel liquid cooling have been built by various research groups

around the world with promising results [8], [11], [12]. Further development of this technology and its large-scale use in the development of computers based on 3D ICs is strongly dependent upon sound early-stage design tools capable of 1) accurately predicting the thermal performance of these cooling technologies as well as 2) prescribing optimized architectural designs and dynamic run-time management tools that can maximize the electrical performance of these systems while maintaining safe operating temperatures. Research in both these directions have gained prominence in recent years [13]-[16] with emphasis on the development of an architectural thermal simulator and several thermal management approaches from a computer engineering perspective. However, a robust and low-complexity thermal modeling approach for 3D ICs is needed to evaluate the benefits of these works at the system-level for the future 3D MPSoCs.

In this context, the "3D Interlayer Cooling Emulator" (3D-ICE)- the first-ever compact model capable of performing transient thermal simulation of liquid-cooled ICs- was proposed in [1], [2]. This is accomplished using geometric compact thermal modeling, which raises the level of abstraction and enables the construction of an equivalent electrical circuit that simulates the conjugate thermal conduction-convection in complex silicon structures with microchannel heat sinks. This makes 3D-ICE ideally suited for early-stage thermal-aware design of 3D MPSoCs. We have recently released 3D-ICE as an open source thermal simulator/software library [3]. It has spawned other research efforts both in thermal management (for instance, [17]-[19]) and thermal simulation (for instance, [20], [21]) of liquid-cooled MPSoCs

• A. Sridhar, A. Vincenzi, and D. Atienza are with the Embedded Systems Laboratory (ESL), École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland.

E-mail: {arvind.sridhar, alessandro.vincenzi, david.atienza}@epfl.ch.

• T. Brunschweiler is with the Advanced Thermal Packing Group, IBM Research Laboratory, Rüschlikon, Switzerland. E-mail: tbr@zurich.ibm.com.

Manuscript received 07 Sept. 2012; revised 11 Mar. 2013; accepted 29 May 2013. Date of publication 13 June 2013; date of current version 12 Sep. 2014. Recommended for acceptance by E.-Y. Chung.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TC.2013.127

The contributions of this paper are:

- 1) to present a review of the existing literature on liquid-cooled ICs and thermal modeling from the perspective of computer architectural design.
- 2) to present a detailed derivation of 4RM-based CTTM and 2RM-based CTTM for 2-port cavities; and to extend the 3D-ICE model to liquid-cooled ICs that have multi-port microchannel cavities [22].
- 3) to present an evaluation of 3D-ICE simulation accuracy using transient temperature measurements from a real liquid-cooled 3D IC prototype, which is the first such validation for a thermal simulator of this class.
- 4) to use simulations of a realistic 3D MPSoC using the new enhanced 3D-ICE model to study and compare the impact of 2-port and 4-port cavities on IC temperature; and compare the simulation speeds of the 2RM-based and 4RM-based CTTMs, demonstrating the suitability of 3D-ICE in the early-stage design of 3D MPSoCs.

The rest of the paper is organized as follows. Section 2 reviews the previous work done in the area of thermal simulation and management of ICs. Section 3 presents a brief overview of 3D-ICE and its features. Section 4 describes the target 3D-IC architecture that is modeled in 3D-ICE. The compact 3D-ICE model for this test structure (both the 4RM- and the 2RM-based CTTM) is derived and presented in detail in Section 5. In addition, the 3D-ICE model is extended for multi-port cavities. The accuracy of 3D-ICE is validated using extensive experimental data from a real liquid-cooled 3D IC stack in Section 6. Simulation of a realistic 3D MPSoC using the 3D-ICE model is used to study the impact of 2-port and 4-port cavities on the temperature, and also the superior performance of the 2RM- over the 4RM-based CTTM in Section 7. Finally, the main conclusions of this work are summarized in Section 8.

2 SUMMARY OF RELATED WORK

A summary of the state-of-the-art in the manufacturing of liquid-cooled ICs and the related thermal simulation methodologies is presented in this section.

2.1 Liquid-Cooled IC Packages for 3D MPSoCs

The notion of liquid cooling of electronics has existed since the late 1970s [23]. In the light of the impending *thermal wall* in the scaling of bipolar integrated circuits, IBM and CRAY among others have advanced liquid-cooled cold plates and server blades as an effective and energy efficient alternative to the conventional air cooling systems. In 1981, Tuckerman and Pease [8] in their seminal work built and characterized the first IC cooled using water pumped through microchannels etched directly on the IC substrate. They demonstrated heat removal capabilities as high as $800\text{W}/\text{cm}^2$ while maintaining temperatures below 85°C . They also studied the relationship between the aspect ratio of microchannel cross-sections and the corresponding cooling efficiencies.

However, with the advent of CMOS circuits in the mid 1980s and the corresponding fall in the power consumption and heat dissipation in electronic devices, the interest in the development of liquid-cooled heat sinks dwindled. That interest has now been renewed with denser packaging of




Fig. 1. Trends of on-chip power density over a period of four decades [5].

electronic components in MPSoCs driving heat fluxes on conventional planar ICs beyond the $100\text{W}/\text{cm}^2$ mark (Fig. 1) [5], and the vertical integration of these devices into 3D ICs compounding the heat fluxes to alarmingly high levels (unto $250\text{W}/\text{cm}^2$ [9]).

While microchannel heat sinks and the heat transfer mechanisms in microscale structures have been well researched [24], there are only two known complete prototypes of 3D vertically stacked ICs with interlayer liquid cooling. The first, built by Dang et al. [12], uses microfluidic channels constructed inside the PCB to transport water to the IC package. Fluidic through-silicon vias (TSVs) are then used to transport water to the different cavities between the ICs. Transient measurements were undertaken to evaluate the cooling performance. One major challenge with this approach is the resulting complexity of PCB designs and the increased pressure drop (and hence the pumping energy spent) to pump water through these narrow and intricate fluid networks.

Brunschwiler et al. at the IBM Research in Zurich independently built another set of prototypes of interlayer liquid cooled 2D and 3D ICs [9], [25] with various types of heat transfer geometries- microchannels, pin fins, 2-port cavities and 4-port cavities. Fluid is pumped using micropumps via a hermetically sealed manifold enclosing the entire IC. Resistive heaters that can be excited using controllable current input are fabricated on each die mimicking the various heat dissipating logic blocks in a real IC. Resistance temperature detectors (RTDs) with linear response to temperature enable accurate evaluation of the heat transfer performance of the microchannels. Extensive studies were performed in all these structures to validate the existing theory on heat transfer, develop correlations for new heat transfer geometries and finally to set benchmarks on fabrication and testing of future fully functional 3D MPSoCs with interlayer liquid-cooling. In this work, we used one of these prototypes- the *pyramid stack*- as shown in Fig. 2, for the detailed validation of our proposed thermal modeling approach in 3D-ICE. It consists of 4 dies and 4 microchannel cavities. A stepped pyramid-like structure constructed by stacking dies of slightly different sizes enable the electrical connections to the heaters and RTDs from each die to be wire bonded onto the PCB for easy measurements.




Fig. 2. The *pyramid stack* built by IBM Research in Zurich: a scanning electron microscope view of the fluid cavities and the electrical connections to the heater and RTDs (left); the final 3D IC package with the fluid manifold and the inlet/outlet connectors (right) [9], [25].

A detailed description of the experimental setup is given in Section 6.1.

2.2 Previous Thermal Modeling Methodologies

The concept of solving the partial differential equations governing transient heat conduction in solids by applying finite-difference approximation and building an equivalent electrical RC circuit is well-known in the heat transfer literature [26]. The rectangular structures typically encountered in ICs favor the application of finite-difference methods. Tools like HotSpot [27] make use of this methodology to create a compact thermal model for conventional air-cooled ICs with a very high abstraction level. These tools can then be integrated with power emulators such as Wattch [28], which provide the heat dissipation traces of individual floorplan elements in the IC, in order to obtain the time-domain temperature responses at different locations in the IC. Integrating this process with floorplanning tools enables efficient early-stage thermal-aware design of MPSoCs.

A similar compact modeling approach for liquid-cooled ICs did not exist in the literature before 3D-ICE. While the heat transfer mechanism of forced convective cooling in channels and even microchannels are well understood in the heat transfer literature [8], [29], [23], [30], this knowledge has not been effectively translated for the purposes of thermal-aware IC design. More recently, a steady-state thermal model for integrated microchannel cooling in 3D ICs was presented in [31]. In this method, the microchannels are discretized into small blocks along the direction of the flow and four new temperature nodes are added for each such block just inside each of the 4 channel walls. The temperature change in each of these nodes in the downstream direction is then calculated as a linear function of the heat flux entering a given upstream node using numerical presimulation. These linear functions, which model the “thermal wake” (the rise in temperature at a downstream location due to heating at an upstream location), are then incorporated in the heat conduction resistive model for the 3D stack. The main drawbacks of this approach are as follows. First, it is not a compact modeling methodology and focusses on many thermal effects that are irrelevant to the end user. Second, steady state conditions are assumed and no transient information can be obtained, rendering the method unsuitable for run-time thermal analysis and management. Third, the problem size is really large due to the dense coupling of nodes inside the microchannels. Finally, extensive numerical presimulations are needed for every thermal simulation executed.

3 3D-ICE IN EARLY-STAGE IC DESIGN

From an MPSoC engineering perspective, a compact thermal model capable of performing transient thermal analysis is necessary to enable the study of thermal-aware floorplanning [18] and dynamic run-time thermal management, such as DVFS. 3D-ICE addresses these needs as described next.

- **Compactness and Transient modeling:** 3D-ICE contains our new compact thermal modeling approach, called the *compact transient thermal model* or CTTM. We constructed it by the identification of the equivalent electrical representation of convective heat transport in fluid flow as a *voltage-controlled current source*. The CTTM raises the level of abstraction in this model is orders of magnitude higher than fine-grained commercial simulators, as was demonstrated in [1], [2].
- **Low complexity:** The 3D-ICE CTTM results in a significantly reduced problem size (hence, reduction in simulation time and memory consumption) compared to fine-grained simulations for liquid-cooled ICs; thus it has a complexity similar to conventional compact resistive thermal models for air-cooled ICs, as shown in Section 5.
- **Versatility and Accuracy:** By construction, 3D-ICE condenses all the information about the physics of forced convective cooling into a few parameters that can be easily modified. Hence, there is a high degree of flexibility in 3D-ICE to change not only the kind of heat transfer geometry (microchannels, pin fins etc), but also the methodology to compute heat transfer parameters for these structures [2]. Heat transfer coefficients (HTCs) computed using correlations for fully developed flows were incorporated in 3D-ICE in [1]. HTCs computed using numerical presimulations in commercial simulators such as ANSYS-CFX were included in [2]. In this work, we will include the HTCs from correlations for developing flows [32] and demonstrate, using the experimental validation in Section 6, that these correlations are sufficiently accurate for early-stage 3D MPSoC design, while being trivial in computational complexity compared to numerical presimulations.
- **Numerical Stability:** The circuit equations built using the 3D-ICE model are numerically integrated in time using the backward Euler method, which is unconditionally stable. Hence, stability of the simulations is guaranteed for any time-step.

4 TARGET ARCHITECTURE

A typical 3D IC with liquid cooling is built by etching the microchannels and TSVs on the back side of individual dies, aligning them and stacking them with a bonding process. There are various methods to accomplish each of the above steps, resulting in different final 3D IC architectures. In this work, and without loss of generality, we use the architecture shown in Fig. 3a. This structure, under development in IBM Research in Zurich, consists of vertically stacked dies with microchannel cavities glued together using bonding materials such as polyimide, with TSVs running through the microchannel walls. The structure is hermetically sealed using a polymer based manifold and the electrical connections from the dies are connected to the PCB via area-array TSVs through the substrate using area-array C4 bumps. The heat flux

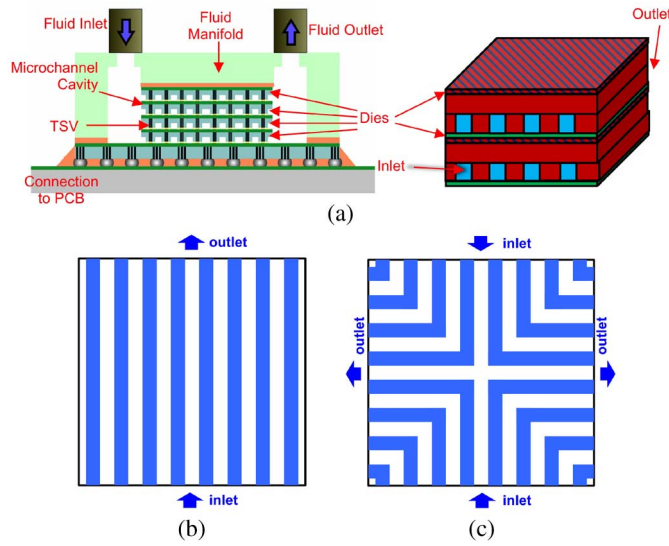


Fig. 3. (a) Target liquid-cooled 3D IC architecture (left) and the corresponding computational domain of 3D-ICE simulations (right). Cavity heat transfer geometries: (b) 2-port cavity and (c) 4-port cavity.

distribution, which is the input in the simulations, is defined for each die by means of a floorplan of the different logic blocks in the IC under study, and the corresponding power traces obtained from a power/performance emulation of the architecture.

The two types of cavity geometries are shown in Figs. 3b and 3c: Four-port microchannels cavities, unlike the traditional two-port devices, have two inlet and two outlet, where coolant enters the microchannels from the south and the north inlets, bends 90°, and exits via the east and west outlets. While increasing design complexity, there are various advantages to the four-port configuration when compared to the traditional two-port straight channel structures that make it interesting for the cooling of high-performance 3D MPSoCs [22]:

- 1) For the same die size, average length of the microchannels is reduced by half resulting in reduced pressure-drop (or pumping effort).
- 2) Alternatively, for the same average length of the microchannels, a die twice the size can be cooled using the same pumping effort.
- 3) For a given pressure-drop between the inlet and outlet ports, flow rates across the channels are non-uniform with very high fluid velocities in the corners resulting in non-uniform cooling- a property that can be utilized in hotspot minimization.

Since early-stage MPSoC designers are typically interested only in the temperatures inside the IC and not in the surrounding structures [17], the computational domain is limited to the volume occupied by the dies as shown in Fig. 3a. The polymer-based manifold used for its excellent workability and suitability to hermetic sealing is a poor conductor of heat. Hence, all the exposed surfaces of the 3D IC are assumed to be adiabatic with the microchannels being the sole heat sink for the 3D IC. Fluid is assumed to enter the inlet at a constant temperature. The outlet temperature depends upon the amount of heat absorbed by the fluid as it flows from inlet to the outlet. Typical footprints of such 3D ICs are 10mm × 10mm and the typical microchannel cross-sectional dimensions are 50 – 100μm × 50 – 100μm [9].

5 COMPACT MODELING IN 3D-ICE

In this section, the proposed compact transient thermal model developed in 3D-ICE for 3D ICs with microchannel cooling is presented. In the ensuing subsections, we first briefly describe the conventional compact modeling of heat conduction in solids and how it relates to compact thermal simulators for conventional air-cooled ICs. Then, the compact transient thermal model (CTTM) for fluids is derived from the first principles and is incorporated in 3D-ICE.

5.1 Conventional Compact Modeling of Solids

The derivation of the conventional compact model for heat conduction in solids begins with the governing equation of heat transfer in solids [27], [29], which can be written in the differential form as:

$$C_v \frac{dT}{dt} + (-k\nabla^2 T) = \dot{q}, \quad (1)$$

where T is the temperature of the control volume, k is the thermal conductivity of the material, \dot{q} is the volumetric rate of generation of heat inside the volume and C_v is the volumetric specific heat of the material.

The above partial differential equation can be converted into an ordinary differential equation by applying the finite difference approximation to the spatial derivative (the second term on the left hand side) in the above equation [29], [26]. To this end, the given volume of solid is discretized along the 3 cartesian coordinates with discretization lengths Δx , Δy and Δz , respectively, to generate a thermal grid. If the temperature of each node in the grid is represented by its location as $T_{i,j,k}$, then the finite difference approximation for the above equation at the location (i, j, k) can be written as:

$$\Delta x \Delta y \Delta z \left(C_v \frac{dT}{dt} - k \frac{T_{i+1,j,k} - 2T_{i,j,k} + T_{i-1,j,k}}{\Delta x^2} - k \frac{T_{i,j+1,k} - 2T_{i,j,k} + T_{i,j-1,k}}{\Delta y^2} - k \frac{T_{i,j,k+1} - 2T_{i,j,k} + T_{i,j,k-1}}{\Delta z^2} \right) = \dot{q} \Delta x \Delta y \Delta z. \quad (2)$$

The well-known analogy between heat and electrical conduction is invoked here with the temperature represented as voltage, heat flow represented as electric current [27], the first term on the left hand side in the above equation represented as a capacitor and the rest of the terms on the left hand side represented as conductances, thus defining an RC circuit [29]. Then, for each tier in a 3D IC, the compact thermal model is generated as follows considering a single silicon layer of a die divided into 9 different *thermal cells*, each with length l , width w and height h , as shown in Fig. 5a. Each cell is then modeled as a node containing six resistances representing the conduction of heat in all the six directions (top → +z, bottom → -z, north → +y, south → -y, east → +x and west → -x), and a capacitance representing the heat storage inside the cell, as shown in Fig. 4a.

The conductance of each resistor and the capacitance of the thermal cell are calculated using the material properties of the solid and the cell dimensions. Current sources, representing




Fig. 4. A single thermal cell: (a) solid and (b) liquid.

the sources of heat, are connected to the cells wherever there is heat dissipation. Next, the nodes of these thermal cells are connected to the nodes of their neighboring cells through the interfaces by computing the equivalent conductances between them. This creates the following system of ordinary differential equations:

$$\mathbf{G}\mathbf{T}(t) + \mathbf{C}\dot{\mathbf{T}}(t) = \mathbf{U}(t), \quad (3)$$

where $\mathbf{T}(t)$ is the vector of all node temperatures (as a function of time) ordered according to their numbering in Fig. 5a, \mathbf{C} is a diagonal matrix of all cell capacitances, $\mathbf{U}(t)$ is a vector of inputs (heat sources as a function of time) wherever they exist. \mathbf{G} is the conductance matrix of the form shown in Fig. 6a. The diagonal term $\sum g_i$ represents the sum of all conductances between node i and its neighbors. The formulation of equations, as described above, can be extended to structures containing multiple layers of thermal cells and a *thermal grid* for an entire IC can be generated.

5.2 3D-ICE Compact Modeling for Liquid-Cooling

The energy conservation equation for heat transfer in flowing liquids can be written (similar to the case of solids) in the differential form as follows [29]:

$$C_v \frac{dT}{dt} + \nabla \cdot (-k\nabla T) + C_v \vec{u} \cdot \nabla T = \dot{q}. \quad (4)$$

Compared to Eq. (1), the above equation contains an added term on the left hand side. Here, \vec{u} is the velocity of outflow of the fluid at the surface of the control volume. This term indeed represents the net outflow of heat from the control volume due to convection. This convection term can be calculated for each surface of a small cuboidal thermal cell as a product of the velocity of the fluid flowing out, the surface temperature, the area of the surface and the volumetric heat capacity of the fluid. Keeping this in mind, we apply finite difference approximation, similar to the case of solids, for a given liquid cell with unidirectional fluid flow (towards the north

direction as shown in Fig. 4b). As a result, we obtain the following equation: shown in (5) at the bottom of the page.

In the above equation, the terms k_{xx} , k_{yy} and k_{zz} are the conductivity of the fluid in the x , y and z directions, respectively. $u_{avg,y}$ is the average velocity of the fluid through the cell in the y direction, namely, the only non-zero component of the velocity, with fluid entering from the front end and exiting via the rear end of the fluid cell as indicated in Fig. 4b. The terms T_{S2} and T_{S1} represent the surface temperatures of the rear and the front faces, respectively, and ΔA_y is the front face area.

By invoking the electrical analogy, the above term can be translated into a *voltage-controlled current source* in the equivalent RC circuit calculated as follows:

$$J_{conv} = c_{conv}(T_{S2} - T_{S1}), \quad (6)$$

where $c_{conv} = C_v u_{avg,y} \Delta A_y$ and the surface temperatures T_{S2} and T_{S1} can be calculated using central differencing. These *voltage controlled current sources* model the transport of heat from the inlet to the outlet of the microchannel and, hence, account for the rise in temperature of the coolant as it flows through the microchannel. The inlet temperature serves as the boundary condition in this model.

The conductance terms in this cell would represent heat transfer from the walls of the channel into the fluid. In [1] we used a 4-resistor model (4RM-based CTTM) to represent heat transfer from all the four sides of a microchannel into the fluid as shown in Fig. 7a. These resistances can be calculated using the heat transfer coefficients h_{conv} obtained using empirical correlations or numerical presimulations. Conduction of heat along the fluid flow direction was neglected in comparison to the convective heat transport which dominates the heat transfer in this direction. Hence, for a silicon layer with microchannels, as shown in Fig. 5b, the application of 4RM-based CTTM would result in a conductance matrix as shown in Fig. 6b. As can be seen, the structure and the sparsity (i.e the locations of the non-zero values in the matrix) of both a solid-only model and the 3D-ICE model are identical resulting in similar CPU time and memory performance during 3D MPSoC thermal simulations, as described in Section 3. The complete derivation of this model can be found in [1].

$$\Delta x \Delta y \Delta z \left(C_v \frac{dT}{dt} - k_{xx} \frac{T_{i+1,j,k} - 2T_{i,j,k} + T_{i-1,j,k}}{\Delta x^2} - k_{yy} \frac{T_{i,j+1,k} - 2T_{i,j,k} + T_{i,j-1,k}}{\Delta y^2} - k_{zz} \frac{T_{i,j,k+1} - 2T_{i,j,k} + T_{i,j,k-1}}{\Delta z^2} \right) + C_v u_{avg,y} \Delta A_y (T_{S2} - T_{S1}) = \dot{q} \Delta x \Delta y \Delta z. \quad (5)$$




Fig. 5. Discretization of a silicon layer in an IC: (a) without microchannel and (b) with microchannel.

The 4RM-based CTTM requires the boundaries of the thermal cells to conform to the solid-liquid interfaces. Hence, the discretization in the x direction is constrained by the microchannel dimensions as illustrated in Fig. 7a (a small section of the microchannel layer is shown here). Since the typical microchannel widths are at least 2 orders of magnitude smaller than the lateral dimensions of a typical IC, this may result in very fine discretization and consequently prohibitively large simulation times for early-stage designs (See Section 7).

We overcome this with the introduction of the 2-resistor model based CTTM or 2RM-based CTTM [2]. This is accomplished by homogenizing the entire microchannel cavity layer in the 3D IC into a single *porous medium* based on the theory described in [22]. Here, we computed an effective heat transfer coefficient by projecting the heat entering from the side walls onto the top and the bottom walls as follows:

$$h_{eff,porous} = \frac{\int h_{conv} dA_{wetted}}{A_{projected}}, \quad (7)$$

where A_{wetted} is the actual wetted surface area and $A_{projected}$ is the area on which the heat transferred from wetted surface is projected.

In each thermal cell of this homogenous medium in the cavity, both the solid and the liquid cell parameters exist side by side (Fig. 7b) with their individual contributions to the thermal grid determined by a parameter called *porosity*, ϵ defined as the volumetric fraction of the cavity occupied by the liquid. Hence, thermal cells are no longer constrained by microchannel dimensions and even *multiple* microchannels can be covered by a single thermal cell in the x direction in this homogenous medium, as illustrated in Fig. 7. We compute these parameters, representing convective heat transfer from the walls into the coolant, convective heat transport downstream, heat storage in both the coolant and the silicon wall, vertical heat conduction through the microchannel walls and heat conduction along the walls parallel to coolant flow, as follows:

$$\begin{aligned} g_{conv,top/bottom} &= \frac{1}{R_{conv}} = h_{eff,porous} \cdot (l \cdot w), \\ J_{conv} &= c_{conv}(T_{S2} - T_{S1}) \cdot \epsilon, \\ c_{coolant} &= C_{v,coolant} \cdot (l \cdot w \cdot h) \cdot \epsilon, \\ g_{cond,top/bottom} &= \frac{1}{R_{cond}} = k_{Si} \cdot \frac{l \cdot w}{h/2} \cdot (1 - \epsilon), \\ g_{downstream} &= \frac{1}{R_{downstream}} = k_{Si} \cdot \frac{l \cdot h}{w/2} \cdot (1 - \epsilon), \\ c_{Si} &= C_{v,Si} \cdot (l \cdot w \cdot h) \cdot (1 - \epsilon). \end{aligned} \quad (8)$$

Detailed derivation of these parameters can be found in [2]. A comparison between our 4RM-CTTM and 2RM-CTTM representations is illustrated in Fig. 7. The 2RM-based CTTM not only frees the user from the restrictions of the microchannel dimensions in discretizing the 3D IC structure for simulation, but also allows for the inclusion of virtually any kind of




Fig. 6. Structure of the G matrix for (a) test silicon layer-Figs. 5a and 5b test silicon layer with microchannel-Fig. 5b.




Fig. 7. The 3D-ICE models for a small section of the microchannel layer: (a) 4RM-based CTTM and (b) 2RM-based CTTM.




Fig. 8. Heat transfer geometries: (a) microchannel, (b) pinfin inline, and (c) pinfin staggered.




Fig. 9. Comparison of heat transfer coefficients calculated using different methods as a function of distance along the channel.

heat transfer geometry such as pin fins (Fig. 8) as long as its porosity and the effective heat transfer coefficient (Eq. (7)) is known, making 3D-ICE an extremely versatile tool for 3D MPSoC thermal simulation [2].

5.2.1 Heat Transfer Coefficients Used in 3D-ICE

For microchannels in our experiments, we calculate the surface heat transfer coefficients as follows:

$$h_{conv} = \frac{k_{coolant} \cdot Nu}{d_h}, \quad (9)$$

where $k_{coolant}$ is the thermal conductivity of the coolant, Nu is the Nusselt number of the flow and d_h is the hydraulic diameter of channel, defined as $\frac{2h}{h+1}$ for rectangular channels. The Nusselt number in the current implementation was derived from correlations proposed by Shah and London [32] for *developing flows* with the assumption of isothermal channel perimeter. Here, the heat transfer coefficient is expressed as a function of a dimensionless distance y^* along the channel as follows:

$$Nu = 3.04 + \frac{0.0244}{y^*} + \frac{0.448}{AR} - \frac{0.0000269}{y^{*2}} + \frac{0.02}{AR^2} - \frac{0.000678}{y^* \cdot AR},$$

$$y^* = \frac{\pi}{4} \left(\frac{y}{Re \cdot Pr \cdot d_h} \right), \quad (10)$$

where AR is the aspect ratio of the channel, Re is the Reynolds number of the flow and Pr is the Prandtl number. Note that the above formula for the Nusselt number is different from the correlations for *fully developed flows* (again from [32]) used in [1], [2]. This new correlation is incorporated in 3D-ICE to increase the accuracy of the simulations.

This is illustrated using comparisons against numerical simulations in the fine-grained simulator ANSYS CFX [33] in Fig. 9. Here, a 1cm long channel of cross-sectional dimensions $50\mu\text{m} \times 100\mu\text{m}$ subjected to uniform heat flux from the top

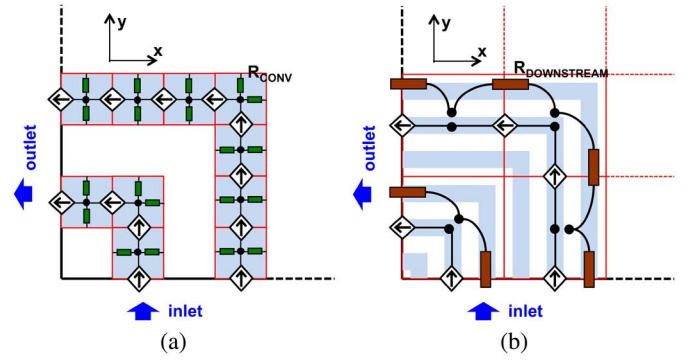


Fig. 10. 3D-ICE model for the 4-port cavity: (a) 4RM-based CTTM and (b) 2RM-based CTTM.


and bottom directions was simulated using ANSYS CFX. From the simulation results, the local heat transfer coefficients were computed, and then compared to the HTC obtained using correlations for both fully developed flows and developing flows. The HTCs as a function of distance along the channel are plotted in Fig. 9. As can be seen, HTCs computed using the developing flow correlations (maximum error 9%) match those from CFX much better than HTCs computed using the fully developed flow correlations (maximum error 33%), with no additional computational cost. Hence, the new HTCs implemented in 3D-ICE are expected to provide greater accuracy in thermal simulations, as demonstrated in Section 6.

5.3 New CTTM for Multi-Port Microchannel Cavities

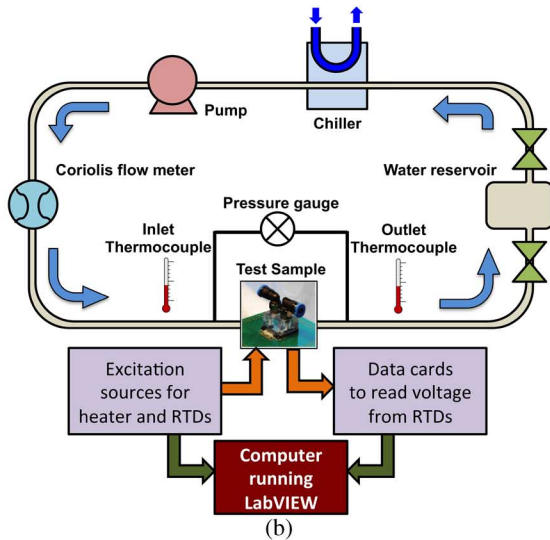
Based on the theory described in the previous sections, the 3D-ICE model can be extended for the case of microchannel cavities with four-ports [22], [25] as illustrated in Fig. 3c.

The 4RM- and the 2RM-based CTTM can be extended for the four-port cavity, by changing the direction of the voltage-controlled current sources in the liquid thermal cells representing convective heat transport based on the location of the thermal cell. This is illustrated for the 4RM-based CTTM in Fig. 10a (only a small section of one quadrant of the cavity is shown here). In this top-view, only the liquid thermal cell components are shown for the sake of clarity (also, the connections in the z direction are omitted). The connections representing the convective resistance to the microchannel walls remain the same as in Fig. 7a. Note that in this case, the cell size in both x and y directions (Δx and Δy) are determined by the cross-sectional dimensions of the microchannel (in the two-port scenario, only Δx was fixed by microchannel dimensions).

The 2RM-based CTTM for the four-port cavity is shown in Fig. 10b. Here, thermal cells of the porous medium (in red) encompassing multiple channels are superimposed on the cavity structure. Note that the direction of the voltage-controlled current sources again depends upon the location of the thermal cell vis a vis the underlying microchannel geometry. The conductive resistance for the microchannel walls $R_{downstream}$ also bends along with the flow direction in these cells. Connections to the top and bottom microchannel walls (and hence, to the other layers in the model), again remain the same as in Fig. 7b and the individual circuit components are scaled by the porosity as in Eq. (8). Note that both the 4RM- and the 2RM-based CTTMs described



(a)



(b)

Fig. 11. (a) Experimental setup and (b) schematic.

above can be generalized to any n -port microchannel cavity. In this work, we focus primarily on the 4-port case.

5.3.1 Flow Rate in Individual Channels

As mentioned earlier, since the pressure drop across all channels are kept uniform, the flow rate varies along the channels. Hence, the flow rate for individual channels must be computed separately. The Darcy-Weisbach equation for laminar flows [23] is employed here to compute the volumetric flow rate \dot{V}_i in the i^{th} channel, based on its length l_{Ci} and pressure drop ΔP :

$$\dot{V}_i = \frac{\Delta P}{8\mu l_{Ci}} \frac{(H_C \cdot w_C)^3}{(H_C + w_C)^2}, \quad (11)$$

where w_C and H_C are the cross-sectional width and the height of the channels, and μ is the dynamic viscosity of the coolant. Using this relationship, the flow rate in each channel can be computed and substituted in the J_{conv} terms in Eqs. (6) and (8) for the 4RM- and 2RM-based CTTM respectively.

5.4 3D-ICE Implementation

This subsection presents a brief note on the implementation of the 3D-ICE thermal simulator. In order to maintain the numerical stability of simulations, the ordinary differential equations constructed in 3D-ICE (Eq. (3)) are numerically




Fig. 12. Layout of heaters and RTDs in the bottom die.

integrated using backward Euler method, and solved using the SuperLU sparse matrix solver. Both steady state and transient simulations can be performed. Power traces for inputs to the simulation can be provided off-line using text files along with project *stack files* that describe the structure of the IC. Alternatively, it is also possible to feed dynamic input power traces and retrieve the output thermal data via a network socket, thus, creating an efficient HW/SW cosimulation platform for MPSoCs [34]. Further implementation details can be found in [3].

6 EVALUATION OF 3D-ICE ACCURACY USING MEASUREMENTS

In this section, we validate 3D-ICE against experimental data from an actual liquid-cooled 3D IC stack, which is the first such validation of this type of simulator.

6.1 Experimental Setup

We built an experimental setup for the accuracy evaluation of 3D-ICE at the IBM Research Laboratory in Zurich using the water-cooled 2-port *pyramid stack* (Fig. 2) as the test structure. A complete fluid loop with a controllable micropump was constructed. A Coriolis flow meter measured the flow rate and pressure sensors were used to maintain a safe pressure drop between the inlet and the outlet of the package. A chiller was used to maintain the inlet temperature at a constant 20°C . Thermocouples were used to measure the temperatures at the inlet and outlet, and of the ambient. The test setup and its schematic are shown in Fig. 11.

6.1.1 Heaters and Resistance Temperature Detectors

Three out of the four dies in the pyramid stacks have heaters and resistance temperature detectors (RTDs) fabricated on them. Each of these dies contain four heaters of size $2\text{mm} \times 5\text{mm}$ laid out from inlet to outlet. There is one RTD in the middle of each heater measuring the temperatures at these locations. The heaters and RTDs are numbered for identification. The layout of these heaters is shown in Fig. 12 [25]. Since the purpose of this validation is to ascertain the capability of 3D-ICE to accurately simulate heat flow in liquid-cooled microchannel heat sinks, it is desirable to suppress any other route for heat to escape into the ambient.




Fig. 13. Differences between the pyramid stack and the model built in 3D-ICE (drawing not to scale). The active die in the experiments is highlighted using the red arrow.

To this end, the inlet temperature of the coolant was always kept lower than the ambient (26°C).

However, this means that some heat might enter from the ambient into the test sample. This form of heating is fairly uniform, but because water in the microchannels accumulate heat as it flows from inlet to outlet, it is observed that even with all heaters on the IC turned off, there is a slight increase of temperature from inlet to outlet creating a shift in the baseline temperature as we move from RTD 1 to RTD 4. Hence, to identify and eliminate this component from the measurements, all the measurement experiments were performed again with the direction of the coolant reversed as shown in Fig. 12 (because the resulting heat absorption from ambient would remain the same in both cases). In addition, to further suppress the effects of ambient on the measurements, only the bottommost die (marked with a horizontal arrow in Fig. 13) in the stack was considered for excitation and measurements in our experiments.

6.1.2 Excitation Sources and Measurement Devices

Precision voltage and current sources were used to excite the heaters and the control inputs to the RTDs. Voltage data acquisition from the heaters and the RTDs was performed at a sampling frequency of 5 kHz- sufficient to resolve the observed time constants- and fed into a computer running the LabVIEW program [35], which dumped the voltage data into a database. This was then collected and processed by a Matlab program to be then used by 3D-ICE to run the test simulations.

6.1.3 Simulation Parameters

While constructing the model, the assumptions discussed in Section 4 were applied. The differences between the actual test stack and the model built in 3D-ICE for validation is illustrated in Fig. 13.

Firstly, the model doesn't have the pyramid structure to contain complexity. The extra area introduced in the pyramid

stack for the electrical connections to the heaters and RTDs contributed to less than 1% increase in the area, thus having too small an effect on heat spreading to be considered in the model. Secondly, the thermal conductivity of silicon dioxide was uniformly applied to all the metallization layers since the contribution of aluminum in these layers is minimal due to low wiring density. Thirdly, while the fabrication specification for the low conductivity polyimide bonding layer thickness was $4\mu\text{m}$, scanning electron microscope studies revealed the actual thickness to be closer to $3\mu\text{m}$, which was used in simulations. The list of material properties used in the model are tabulated in Table 1. The 2RM-CTTM was applied in all the simulations with a thermal cell size of $200\mu\text{m} \times 200\mu\text{m}$.

6.2 Characterization of Heaters and RTDs

Before proceeding with the validation, the temperature response of the heaters and RTDs must be properly studied and calibrated for accurate measurements. For this, we performed a first set of benchmarking experiments to characterize the thermal responses of the heaters and RTDs. First, coolant was pumped in at a known inlet temperature without firing any heater, and the temperature of the IC was allowed to reach steady state. Then the resistance of each heater and RTD was measured. This procedure was repeated for two other inlet temperatures within the expected range of the final experiments. We plotted the data for each heater and RTD and we used the resulting linear fit as the benchmark to compute temperatures from voltage responses during the final measurements.

TABLE 1
Thermal Properties of Materials Used in the Experiments

| Material | k (W/m · K) | C_v (MJ/m ³ · K) |
|-----------------|---------------|-------------------------------|
| Silicon | 151 | 1.642 |
| Silicon dioxide | 1.38 | 1.654 |
| Alumina | 36 | 3.037 |
| Polyimide | 0.14 | 1.645 |
| Water | 0.6069 | 4.173 |




Fig. 14. Sample transient results for 5 Hz excitation in Heater 3: (a) RTD 2, (b) RTD 3, and (c) RTD 4.

This initial benchmarking confirmed the manufacturer's specification of a sensitivity of $2000\text{ppm}/^{\circ}\text{C}$ for the resistance of all the heaters and RTDs to temperature. The measurements from RTD 1 was found to be unreliable and hence, data from this sensor was discarded during the final experiments. Once the benchmarking experiments were completed, the exact heat input in each heater during the final experiments was computed as follows. For each experiment, the excitation voltages driving the heaters was measured along with the voltages from the RTDs at each time point. Next, using the temperatures computed from the RTDs, the exact resistances of the corresponding heaters were calculated based on their benchmarking for each time instant during the experiments. Finally, using the knowledge of the resistances and the voltages of the heaters, the heat being dissipated in all the heaters at each time point during the experiments was accurately computed and fed into 3D-ICE as inputs.

6.3 Comparison of Transient Thermal Response

After the benchmarking of heaters and RTDs, we performed the final experiments to measure the temperature response to various forms of heater excitations. In each experiment, one heater was activated and the temperature responses in RTD 2-4 were measured. Out of each of these experiments data lasting a period of 1 second was used for study. We ensured that during at least 75% of this time period the heater was active and switching in order to visualize and study temperature transients. Heaters were fired with a square-wave voltage waveform at three different frequencies during the experiments: 10 Hz, 5 Hz and 1 Hz, with a 50% duty-cycle in each case. In addition, we also performed experiments with step rise and step fall in input voltage followed by allowing the temperature to settle to the steady state.

Different voltage levels were chosen to perform experiments with both low heat flux ($35\text{W}/\text{cm}^2$) and high heat flux ($150\text{W}/\text{cm}^2$) levels in the individual heaters. To evaluate the effects of coolant flow rate on the temperatures, all the above experiments were repeated for multiple coolant flow rates- 88 ml/min, 140 ml/min and 175 ml/min (measured for all the four cavities combined, resulting in per cavity flow rates of 22 ml/min, 35 ml/min and 43.7 ml/min, respectively)- all while ensuring that the pressure drop between the inlet and outlet does not exceed 1 bar. A total of 70 experiments were performed (including forward and reversed coolant flows, as shown in Fig. 12).

Sample transient result comparisons for the three RTDs by exciting Heater 3 with 5 Hz voltage waveforms are shown in Fig. 14. The coolant flow was in the reverse direction at a rate of 175 ml/min for these plots. Results from 3D-ICE are always plotted using solid lines and the corresponding results from measurements are plotted using dotted lines. A smoothing filter was applied to the measurement data to remove noise. Each of the subfigures corresponds to results from one RTD. The coloring scheme for all the results in this section is: red for RTD 2, green for RTD 3 and blue for RTD4. In Fig. 14b the heat input in the heater is also shown in a secondary axis. As can be seen, the temperature transients in each case are captured well using simulations. A detailed analysis of the error incurred during the validation experiments is discussed in the next subsection.

6.4 Analysis of Error with Respect to Measurements

To visualize the errors from all the 70 experiments and the patterns therein, the following plotting scheme was used. In each experiment, we noted ϵ_{avg} the maximum absolute error (ϵ_{max}) and the average absolute error (ϵ_{avg}) between 3D-ICE and the data during the entire simulation interval. We also noted the maximum rise in temperature from the baseline (ΔT_{max}) during the measurement. Next, the errors were added to ΔT_{max} and plotted against ΔT_{max} as a scatter plot. The resulting plots for ϵ_{max} and ϵ_{avg} are shown in Fig. 15. The different colors in the data points correspond to data from the different RTDs. The various lines from origin provide a measure of the extent of deviation of the 3D-ICE results from the measured data (i.e. solid line if there is an exact match, dashed line for an error of 10% and so on).

From the plots we observed that the global average error was 8.5%, while the maximum error was contained within 20% for most cases. However, in some cases the maximum errors can be as high as 32%. This suggests that the points in time when these huge errors occur must be rare. Moreover, in almost all cases the sign of the errors were positive, i.e., 3D-ICE overestimated the temperatures, potentially leading to safe but conservative thermal-aware 3D MPSoC designs. In order to spatially isolate the source of this error, we combined the data from all the RTDs into a single set, and then split this set into two parts: errors measured at the location of the active heater and errors measured in locations downstream of the heater in order to study the effect of thermodynamic boundary layer near the active heaters on the error. These plots are shown in Fig. 16.

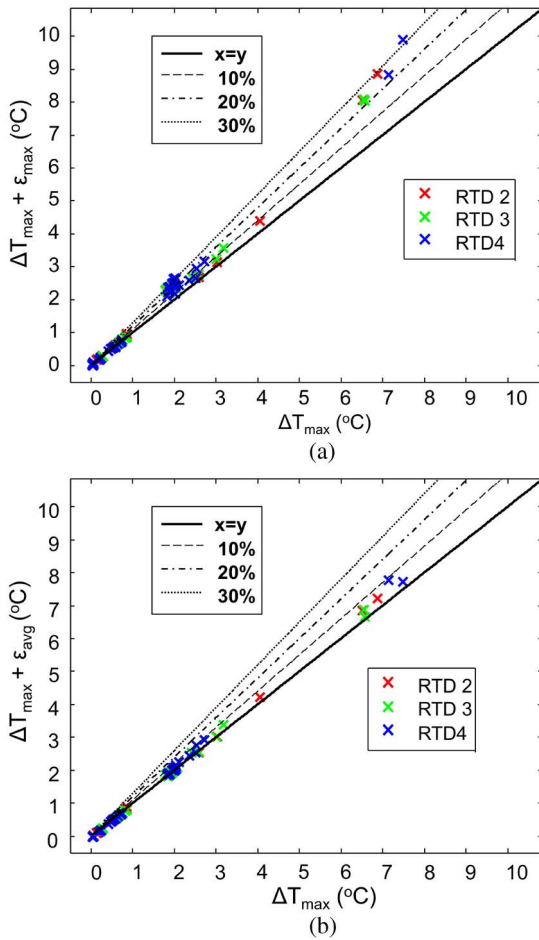


Fig. 15. Scatter plot for (a) maximum error ϵ_{max} and (b) average error ϵ_{avg} .

As can be seen from this figure, the maximum errors in downstream locations are lower ($< 20\%$) compared to the errors near the heater (32%), confirming our hypothesis that the source of the error is the changing heat transfer coefficients in the entrance region. However, we found that average errors in both cases were still small ($< 8.5\%$). This again suggests that even at the entrance regions, errors occur only at few time points of simulation.

In order to temporally isolate the location of these errors, only those results from the above two sets were selected, where the input voltage waveform to the heater was a step rise or a step fall. From these results, we first extracted the error in the steady state temperatures. Next, we normalized all the transient waveforms between 0 and 1 (corresponding to the baseline and the steady state temperatures) and measured the differences between 3D-ICE and data in the rising and falling time constants (measured as the time between the switching of the input and time to reach $1 - e^{-1} = 67\%$ of the final value). A sample plot comparing the normalized waveforms from 3D-ICE and data is shown in Fig. 17. We found that the steady state temperatures from 3D-ICE differed only slightly from the measurement data (10%) irrespective of the location of measurement. However, the time constants showed some interesting trends at locations near the heaters and downstream of the heaters, as illustrated in Fig. 18.

Note that the signs of these lines are negative in both Figs. 18a and 18b indicating that responses in 3D-ICE rise

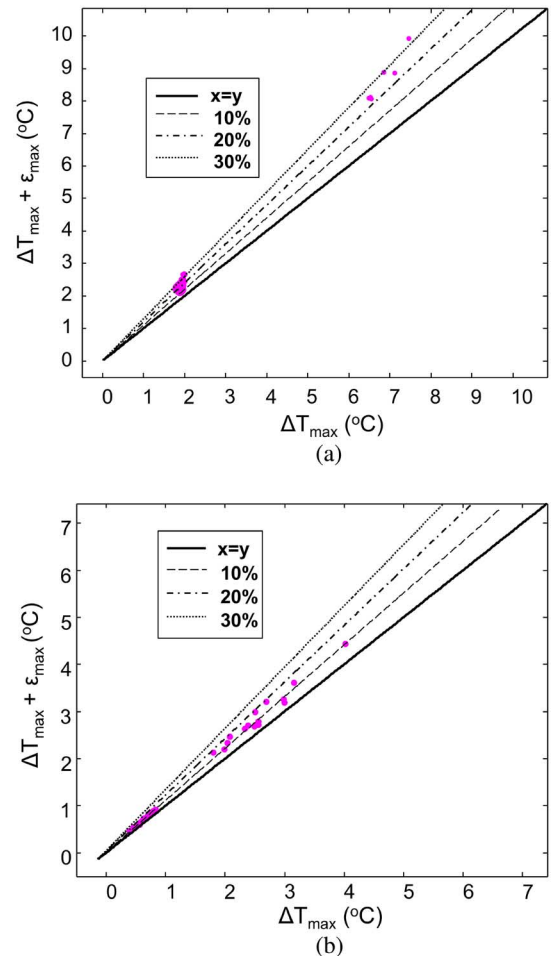


Fig. 16. Scatter plot for maximum errors ϵ_{max} measured (a) at the location of the active heater and (b) at the locations downstream of the active heater.

and fall faster than data. However, while the time constants in 3D-ICE are within 20% of the data at downstream locations, the deviation is fairly large near the heater (35%). Note that the lines indicating the extent of deviation in Fig. 18b correspond to only 5%, 10% and 20%. Given that steady state errors are low, this observation can account for the fact that larger errors in 3D-ICE can occur only at few points in time and space of the simulation, namely, at the rising and falling of the temperature near the heat source. The possible sources of these errors are: increased heat capacitance of the real stack due to the heat spreading into regions outside of computational domain and the inlet/outlet fluid cavities, impurities in the materials, process variations and noise in measurements.

In summary, errors are generally low (i.e., the global average error of is 8.5%) when compared to transient thermal measurements from a real liquid-cooled 3D IC. Larger errors occur very rarely in the simulation interval, are limited to small regions in the IC, and are still less than 30%. Hence, our experiments demonstrate that the accuracy of 3D-ICE is sufficient for the purposes of early-stage 3D MPSoC design. Moreover, this accuracy was obtained while incorporating heat transfer coefficients from correlation studies in the 3D-ICE model, demonstrating that extensive numerical pre-simulations are not needed for this purpose.

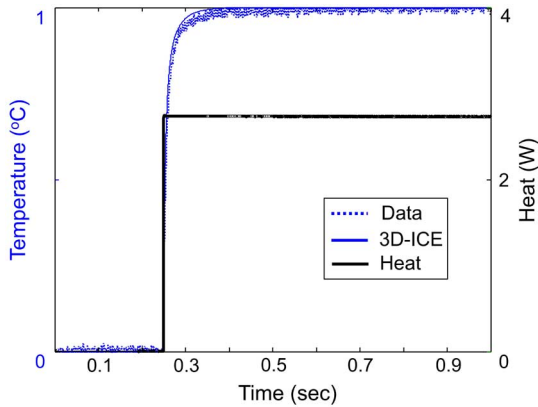


Fig. 17. Sample normalized transient response for computation of time constant.

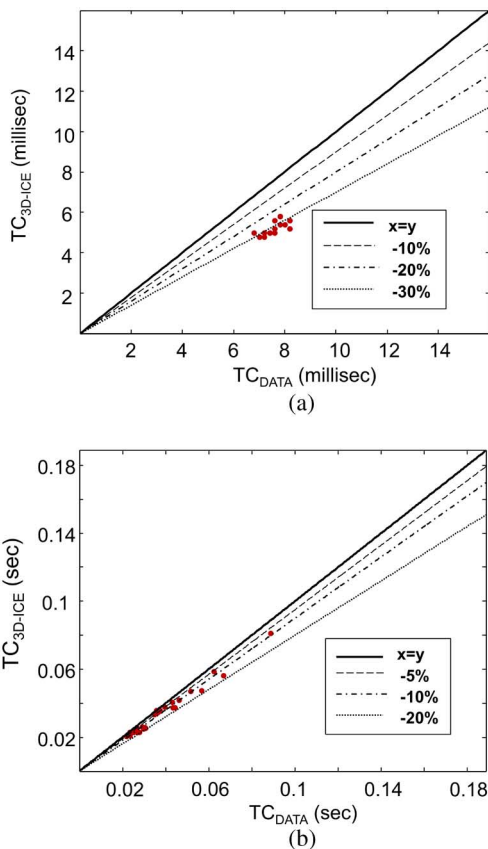


Fig. 18. 3D-ICE vs data scatter plots for time constant (a) at the location of the active heater and (b) at locations downstream of the active heater.

7 SIMULATION OF A REALISTIC 3D MPSoC USING 3D-ICE

In this section, the new 3D-ICE model is used to simulate a realistic liquid-cooled 3D MPSoC to show the impact of 4-port cavities on the temperatures of an IC with respect to the conventional 2-port cavities. For this, we choose the 2-die ULTRASPARC T1 stacked architecture of footprint size $1.1\text{cm} \times 1.1\text{cm}$ (one die contains the cores and the other die contains the memory caches). Benchmark power traces were obtained from the execution of a real-life application on this platform, as discussed in [14]. A liquid-cooled microchannel cavity is sandwiched between the 2 dies. All simulations were




Fig. 19. Temperature map of the ULTRASPARC T1 MPSoC architecture in a 2-die liquid cooled IC stack with (a) 2-port and (b) 4-port cavity.

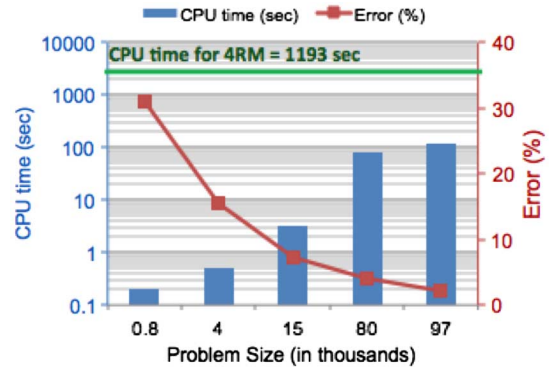


Fig. 20. CPU time comparison and error incurred by the 2RM- w.r.t. the 4RM-based CTTM for the 4-port cooled 3D MPSoC, for various discretization sizes.

performed using the 3D-ICE models implemented in Matlab, running on a Corei7-3770 3.40 GHz CPU with 32 GB RAM.

First, the conventional 2-port cavity is assumed with a pressure drop of 1 bar between the inlet and outlet. The steady state temperatures are simulated using the conventional 3D-ICE 4RM-based CTTM. The resulting thermal map of the core layer is shown in Fig. 19a (all the temperatures are in Celcius). Next, the same 2-die structure is simulated assuming a 4-port microchannel cavity, but with the same pressure drop between the inlet and the outlet ports. For this, the new 4RM-based CTTM for 4-port cavities as described in Section 5.3 is used. The resulting temperatures are shown in Fig. 19b. The following observations can be made based on these results:

- 1) There is a slight reduction in maximum temperature from 2-port to 4-port cavity ($\sim 4^\circ\text{C}$).
- 2) Due to the nature of the flow and the associated accumulation of heat in the coolant, the hotspots are shifted from the north-end of the die in 2-port cavity to the central part of the die in 4-port cavity.

The above observations have significant implications for the thermal-aware floorplanning of 3D MPSoCs when using 4-port cavities. The corners in the 4-port cavity are cooled more efficiently due to the shorter channels and the higher flow rates. Thus, to fully incur the advantages of the 4-port cavity, high power dissipating elements (such as crossbar in this case) must be moved to the corners away from the center of the die.

7.1 Simulation Speed of 2RM- versus 4RM-Based CTTM

We had earlier demonstrated that the 4RM-based CTTM is orders of magnitude faster than commercial simulators such

as ANSYS CFX in [1]. In [2], we also demonstrated the additional speed up of 2RM- over the 4RM-based CTTM for the 2-port microchannels. Now, we compare the performances of the new 2RM- and 4RM-based CTTMs for 4-port microchannel cavities by simulating the same test 3D MPSoC stack. For this, we consider the temperature results from 4RM-based CTTM above as the benchmark, and measure the errors incurred by the 2RM-based CTTM against it. The thermal cell size for the 4RM-based CTTM is, as described in Section 5.3, fixed to $50\mu\text{m} \times 50\mu\text{m}$ and results in a problem size of 244k nodes.

Next, we simulate the same structure using the 2RM-based CTTM by varying the thermal cell size: $100\mu\text{m} \times 100\mu\text{m}$, $110\mu\text{m} \times 110\mu\text{m}$, $250\mu\text{m} \times 250\mu\text{m}$, $500\mu\text{m} \times 500\mu\text{m}$ and $1100\mu\text{m} \times 1100\mu\text{m}$ resulting in various problem sizes. Maximum error with respect to the 4RM-based CTTM is noted in each case. The simulation times for the 2RM-based CTTM are plotted against the problem size for various discretizations in Fig. 20 on a log scale. This plot also shows the error incurred by the 2RM-based CTTM in the secondary axis. The simulation time for the 4RM-based CTTM (about 1200 seconds) is indicated using the green line. As can be seen from the graph, an optimal point can be found (for the case of $250\mu\text{m} \times 250\mu\text{m}$ cell size with a problem size of 15k nodes) where the error is only 7% with a speed up of 375x against the 4RM-based CTTM. Hence, the 2RM-based CTTM in 3D-ICE serves as an efficient tool for early-stage thermal-aware design of 3D MPSoCs while not sacrificing accuracy.

8 CONCLUSIONS

In this work we have reviewed the state-of-the-art in the manufacturing and thermal modeling of liquid-cooled ICs. We have presented a detailed derivation of the 3D-ICE thermal model for liquid-cooled ICs, highlighted the main novelties in it and extended it for the case of 4-port microchannel cavities. We performed extensive measurements from a real liquid-cooled 3D IC stack to evaluate the accuracy of 3D-ICE, which is the first such experimental validation of a simulator of this genre. Finally, we provided a case study with a realistic liquid-cooled 3D MPSoC to both study the effect of 2-port and 4-port cavities on the die temperatures, and to study the simulation speed of 3D-ICE, to demonstrate its suitability for early-stage thermal-aware design of 3D MPSoC architectures.

ACKNOWLEDGMENTS

The authors would like to thank the members of the Advanced Thermal Packaging group at the IBM Research-Zürich, for their support that made this work possible. The authors would also like to thank the users of 3D-ICE for their feedback and input over the last 2 years, which helped in fine-tuning and enhancing the 3D-ICE model. This research has been partially funded by the Nano-Tera RTD project CMOSAIIC (ref. 123618), which is financed by the Swiss Confederation and scientifically evaluated by SNSF, as well as by the PRO3D STREP project (ref. FP7-ICT-248776) financed by the EC in the 7th Framework Programme.

REFERENCES

- [1] A. Sridhar, A. Vincenzi, M. Ruggiero, T. Brunswiler, and D. Atienza, "3D-ICE: Fast Compact Transient Thermal Modeling for 3D-ICs with Inter-Tier Liquid Cooling," *Proc. Int'l Conf. Computer-Aided Design (ICCAD)*, pp. 463-470, 2010.
- [2] A. Sridhar, A. Vincenzi, M. Ruggiero, T. Brunswiler, and D. Atienza, "Compact Transient Thermal Model for 3D-ICs with Liquid Cooling Via Enhanced Heat Transfer Cavity Geometries," *Proc. Thermal Investigations of ICs and Systems (THERMINIC)*, pp. 1-6, 2010.
- [3] A. Sridhar et al., Available: <http://esl.epfl.ch/3D-ICE>. Accessed on Sep. 2010.
- [4] P. Ruch, T. Brunswiler, W. Escher, S. Paredes, and B. Michel, "Toward Five-Dimensional Scaling: How Density Improves Efficiency in Future Computers," *IBM J. Research and Development*, vol. 55, pp. 15:1-15:13, Sept./Oct. 2011.
- [5] I. Canturk, "Workload Adaptive Power Management with Live Phase Monitoring and Prediction," PhD thesis, Princeton University, June 2007.
- [6] D. Abts, M. Marty, P. Wells, P. Klausler, and H. Liu, "Energy Proportional Datacenter Networks," *Proc. Int'l Symp. Computer Architecture (ISCA'10)*, pp. 338-347, June 19-23, 2010.
- [7] B. Agostini et al., "State of the Art of High Heat Flux Cooling Technologies," *Heat Transfer Eng.*, vol. 28, no. 4, pp. 258-281, 2007.
- [8] D.B. Tuckerman and R.F.W. Pease, "High-Performance Heat Sinking for VLSI," *IEEE Electron Device Letters*, vol. 5, no. 5, pp. 126-129, May 1981.
- [9] T. Brunswiler et al., "Interlayer Cooling Potential in Vertically Integrated Packages," *Microsystem Technologies*, vol. 15, no. 1, pp. 57-74, 2009.
- [10] T. Brunswiler, B. Smith, E. Ruetsche, and B. Michel, "Toward Zero-Emission Data Centers through Direct Reuse of Thermal Energy," *IBM J. Research and Development*, vol. 53, pp. 11:1-11:13, May 2009.
- [11] F. Alfieri et al., "3D Integrated Water Cooling of a Composite Multilayer Stack of Chips," *J. Heat Transfer*, vol. 132, no. 12, pp. 565-574, 2010.
- [12] B. Dang, M. Bakir, D. Sekar, C. King, and J. Meindl, "Integrated Microfluidic Cooling and Interconnects for 2d and 3d Chips," *IEEE Trans. Advanced Packaging*, vol. 33, no. 1, pp. 79-87, Feb. 2010.
- [13] H. Mizunuma, C.L. Yang, and Y.C. Lu, "Thermal Modeling for 3D-ICs with Integrated Microchannel Cooling," *Proc. Int'l Conf. Computer-Aided Design (ICCAD)*, pp. 256-263, 2009.
- [14] M.M. Sabry et al., "Fuzzy Control for Enforcing Energy Efficiency in High-Performance 3D Systems," *Proc. Int'l Conf. Computer-Aided Design (ICCAD)*, pp. 642-648, 2010.
- [15] A.K. Coskun, J. Meng, D. Atienza, and M.M. Sabry, "Attaining Single-Chip, High-Performance Computing through 3D Systems with Active Cooling," *IEEE Micro*, vol. 31, no. 4, pp. 63-75, July/Aug. 2011.
- [16] J. Xie and M. Swaminathan, "Electrical-Thermal Co-Simulation of 3d Integrated Systems with Micro-Fluidic Cooling and Joule Heating Effects," *IEEE Trans. Components Packaging and Manufacturing Technology*, vol. 1, no. 2, pp. 234-246, Feb. 2011.
- [17] M.M. Sabry, A. Coskun, D. Atienza, T. Rosing, and T. Brunswiler, "Energy-Efficient Multi-Objective Thermal Control for Liquid-Cooled 3D Stacked Architectures," *IEEE Trans. Computer-Aided Design*, vol. 30, no. 12, pp. 1883-1896, Dec. 2011.
- [18] I. Arnaldo, A. Vincenzi, J. Ayala, J. Risco, J. Hidalgo, M. Ruggiero, and D. Atienza, "Fast and Scalable Temperature-Driven Floor Plan 3d Design in MPSoCs," *Proc. IEEE Latin Am. Test Workshop (LATW2012)*, pp. 98-103, Apr. 2012.
- [19] P. Kumar and L. Thiele, "System-Level Power and Timing Variability Characterization to Compute Thermal Guarantees," *Proc. 9th IEEE/ACM/IFIP Int'l Conf. Hardware/Software Codesign and System Synthesis (CODES+ISSS '11)*. ACM, pp. 179-188, 2011.
- [20] X.-X. Liu, Z. Liu, S.-D. Tan, and J. Gordon, "Full-Chip Thermal Analysis of 3d ICs with Liquid Cooling by GPU-Accelerated GMRES Method," *Proc. 13th Int'l Symp. Quality Electronic Design (ISQED)*, pp. 123-128, Mar. 2012.
- [21] A. Vincenzi, A. Sridhar, M. Ruggiero, and D. Atienza, "Accelerating Thermal Simulations of 3d ICs with Liquid Cooling Using Neural Networks," *Proc. Great Lakes Symp. Very Large Scale Integration (GLSVLSI)*, UT, pp. 15-20, 2012.

- [22] T. Brunswiler, S. Paredes, U. Drechsler, B. Michel, W. Cesar, G. Toral, Y. Temiz, and Y. Leblebici, "Validation of the Porous-Medium Approach to Model Interlayer-Cooled 3d-Chip Stacks," *Proc. Int'l Conf. 3D System Integration (3DIC)*, pp. 1-10, 2009.
- [23] F. Incropera, *Liquid Cooling of Electronic Devices by Single-Phase Convection*. Wiley, 1999.
- [24] S. Garimella, V. Singhal, and D. Liu, "On-Chip Thermal Management with Microchannel Heat Sinks and Integrated Micropumps," *Proc. IEEE*, vol. 94, no. 8, pp. 1534-1548, Aug. 2006.
- [25] T. Brunswiler, S. Paredes, U. Drechsler, B. Michel, W. Cesar, Y. Leblebici, B. Wunderle, and H. Reichl, "Heat-Removal Performance Scaling of Interlayer Cooled Chip Stacks," *Proc. 12th IEEE Intersociety Conf. Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, pp. 1-12, June 2010.
- [26] F. Incropera, D. Dewitt, T. Bergman, and A. Lavine, *Fundamentals of Heat and Mass Transfer*. Wiley, 2007.
- [27] W. Huang, S. Ghosh, S. Velusamy, K. Sankaranarayanan, K. Skadron, and M. Stan, "Hotspot: A Compact Thermal Modeling Methodology for Early-Stage VLSI Design," *IEEE Trans. Very Large Scale Integration (VLSI) Systems*, vol. 14, no. 5, pp. 501-513, May 2006.
- [28] D. Brooks, V. Tiwari, and M. Martonosi, "Wattch: A Framework for Architectural-Level Power Analysis and Optimizations," *Proc. Indian Science Congress Assoc. (ISCA)*, pp. 83-94, 2000.
- [29] J. Lienhard-IV and J. Lienhard-V, *A Heat Transfer Textbook*. Phlogiston Press, 2006.
- [30] J. Koo, S. Im, L. Jiang, and K. Goodson, "Integrated Microchannel Cooling for Three-Dimensional Electronic Circuit Architectures," *ASME J. Heat Transfer*, vol. 127, pp. 49-58, 2005.
- [31] H. Mizunuma et al., "Thermal Modeling and Analysis for 3D-ICs with Integrated Microchannel Cooling," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 30, no. 9, pp. 1293-1306, Sept. 2011.
- [32] R. Shah and A. London, *Laminar Flow Forced Convection in Ducts*. Academic Press, 1978.
- [33] ANSYS, <http://www.ansys.com/products/fluid-dynamics/cfx/>. Accessed on 2011.
- [34] D. Atienza, P. Valle, G. Paci, F. Poletti, L. Benini, G.D. Michelli, J. Mendias, and R. Hermida, "HW-SW Emulation Framework for Temperature-Aware Design in MPSoCs," *ACM Trans. Design Automation for Embedded Systems (TODAES)*, vol. 12, pp. 1-26, 2007.
- [35] National-Instruments, <http://www.ni.com/labview/>. Accessed on 2011.



Arvind Sridhar (S'07) received the BEng degree in electronics and communication engineering from the College of Engineering Guindy, Anna University, Chennai, India, in 2006, and the MAsc degree in electronics from Carleton University, Ottawa, Canada, in 2009. He is currently pursuing doctoral studies in the Embedded Systems Laboratory at École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland. He was a research scholar in the CAD laboratory at Carleton University between 2006 and 2009 and interned at IBM Research, Zurich, Switzerland, in 2011.



Alessandro Vincenzi (S'11) received the BS (in 2007) and MS (in 2010, summa cum laude) degrees in computer science from the University of Parma, Parma, Italy and the University of Verona, Verona, Italy, respectively. In 2010, he joined the Embedded Systems Laboratory (ESL) group at the Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, where he is working toward the PhD degree in electrical engineering. His research interests include thermal modeling of electronic devices as well as program-

ming on parallel and high performances architectures. He received the Best Student Award from the University of Parma (2004), at the end of his first year of studies.



David Atienza (M'05-SM'13) received the MS degree from the Complutense University of Madrid, Madrid, Spain, and the PhD 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 coauthor of more than 180 publications in peer-reviewed international journals and conferences, several book chapters, and six U.S. patents in these fields. He received the ACM SIGDA Outstanding New Faculty Award in 2012, a Best Paper Award at the Very Large-Scale Integration and System-on-Chips (VLSI-SoC) 2009, and six Best Paper Award Nominations at Design Automation Conference (DAC) 2013 and 2004, Design Automation & Test in Europe (DATE) 2013, High Performance Computing & Simulation (HPCS) 2012, Workshop on Exploitation of Hardware Accelerators at HPCS (WEHA-HPCS) 2010, and International Conference on Computer-Aided Design (ICCAD) 2006. He is an associate editor of *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, *IEEE Design and Test*, and Elsevier *Integration*. He is a member of the Executive Committee of the IEEE Council on Electronic Design Automation (CEDA), since 2008, and was a GOLD member of the Board of Governors of the IEEE Circuits and Systems Society (CASS), from 2010 to 2012.



Thomas Brunswiler (M'03-SM'13) received the master's degree in microsystems technology from the Interstate University of Applied Science, Buchs, Switzerland, in 2001, and the PhD degree in electrical engineering from the Technical University of Berlin, Germany, in 2012. He joined the Advanced Thermal Packaging Group, IBM Research, Zurich, Switzerland, in 2001, and worked on and coordinated numerous governmental and joint projects on photonic devices, organic LEDs, and more recently, 3D integration

and the related aspects of heat removal and power delivery. Currently, he supports Dr. Matthias Kaiserswerth, the director of the IBM Research, Zurich, as the technical assistant in strategic matters. He has authored over 50 publications, one book chapter, and over 25 patents. He received three Best Paper Awards at ITherm and SEMI-THERM and was honored in 2009 with the Harvey Rosten Award for Excellence in the Physical Design of Electronics. He is currently on the technical committee of the IBM internal packaging community, ITherm, and InterPACK.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.