

Active Fluidic Cooling on Energy Constrained System-on-Chip Systems

Wen Yueh, *Member, IEEE*, Zhimin Wan, He Xiao, *Student Member, IEEE*,
Sudhakar Yalamanchili, *Fellow, IEEE*, Yogendra Joshi, *Fellow, IEEE*,
and Saibal Mukhopadhyay, *Senior Member, IEEE*

Abstract—This paper presents design, experimental characterization, and feasibility analysis of integrated in-package fluidic cooling for mobile systems-on-chips (SoCs). A pin fin interposer for fluidic cooling is designed and integrated with a commercial SoC. The demonstrated system integrates an active low-power piezoelectric pump controlled by the SoC itself and a metal/acrylic-based board-scale heat spreader and exchanger. Different software-based policies in the SoC for controlling the fluid flow based on SoC's temperature and performance are implemented and compared. The measurement results demonstrate that the in-package fluidic cooling improves the SoC's energy efficiency and reduces design footprint compared to external passive cooling.

Index Terms—Closed-loop control, microfluidic cooling, piezoelectric pump, system-on-chip (SoC).

I. INTRODUCTION

THERMAL management has emerged as a challenging problem in modern computer systems. The demand for efficient computation with more functionality has led the industry to adopt active fluidic solutions [1], [2]. Fluidic and liquid submerged cooling have been explored in high-performance servers and data centers [1], [2]. Traditionally, active cooling-based thermal management has not been explored for mobile systems-on-chips (SoCs). However, the computation capabilities of the SoCs used in smart phones, tablets, and various other applications such as robotics, autonomous avionics, and Internet of Things, have advanced significantly in recent years [3], [4]. With increasing computation demand and processing capabilities of these systems, thermal management is emerging as a key challenge for mobile/embedded SoC in various common use cases, for example, high-definition video processing [5]. An unmanaged high temperature reduces the performance, degrades

device energy efficiency, and accelerates device aging [6], [7]. The system's energy efficiency and user experience can be improved with a more aggressive cooling technique. However, the mobility needs impose stringent requirements on the form factor of cooling solutions.

This paper characterizes the effectiveness of an in-package microfluidic cooling technique using micropin fins for the thermal management of a commercial SoC running embedded and mobile applications. The cooling structure is embedded in a chip-scale silicon interposer that can be separately fabricated and integrated into the die during packaging, without interfering with the chip fabrication. The die attachment is more compact, less disruptive compared to etching the pin fins directly on the chip, and provides better heat extraction capacity compared to external cooling. The in-package cooling technology is fabricated in-house and attached to a commercial SoC (Snapdragon 600, Fig. 1). A low-power piezoelectric pump, controlled by the SoC, is integrated with the system. The experiments are performed using the single-phase cooling with deionized water, given its excellent thermal properties. The system-level temperature, power (processor and pump), and performance are measured considering benchmark applications running on the SoC. The results are compared against the external (on package) passive/active heat-removal technologies. A full system, designed with integrated platform-to-ambient heat spreader, demonstrates closed-loop fluidic cooling with SoC-based control policies that regulate fluid flow.

The experiment demonstrates that the in-package fluidic cooling can reduce the SoC energy consumption and integration footprint. The measurements over benchmark applications showed that the in-package cooling operated at 34 °C–35 °C lower temperature, 2%–47% lower energy, and 17%–89% better performance (completion time) than in a baseline (no cooling) SoC. Compared to the external passive cooling, in-package cooling reduced peak temperature by 23 °C–28 °C and peak energy by 3%–8%, including the pump power (peak 110 mW). The in-package cooling, compared to the passive cooling and external fluidic cooling, has 2.5× and 3× lower footprint, respectively, in SoC devices. By considering a realistic heat sink form factor and shape, closed-loop thermal performance of the system shows 19 °C–23 °C cooler thermal advantages and power advantage of 1%–6% for the benchmarks runs.

The rest of this paper is organized as follows. Section II discusses the related work, Section III presents the proposed

Manuscript received April 8, 2017; revised July 1, 2017; accepted July 30, 2017. Date of publication October 4, 2017; date of current version October 26, 2017. This work was supported in part by the National Science Foundation under CNS-1218745, in part by Sandia National Laboratories, in part by the Semiconductor Research Corporation, and in part by Qualcomm Inc. Recommended for publication by Associate Editor P. Dutta upon evaluation of reviewers' comments. (*Corresponding author: Wen Yueh.*)

W. Yueh, H. Xiao, S. Yalamanchili, and S. Mukhopadhyay are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30313 USA (e-mail: wyueh3@gatech.edu).

Z. Wan and Y. Joshi are with the George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA 30313 USA.

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/TCPMT.2017.2746667

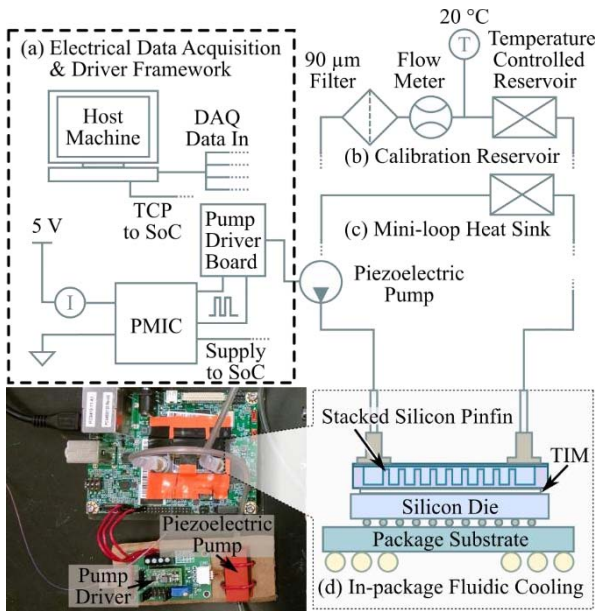


Fig. 1. Experimental characterization of the in-package fluidic cooling. (a) Schematic of the measurement setup. The in-package fluidic cooling is integrated with the SoC (Snapdragon 600). The piezoelectric pump's driver circuit directly draws current from the PMIC on the IFC6410 board. The hall-effect sensor measures the total current entering the board. The driver circuit takes PFM frequency control signal from the programmed GPIO pin 14. (b) Cartoon of the calibration setup with constant temperature bath. (c) Mini-loop setup to demonstrate the system can be lightweight and mobile. (d) Proposed micropin fin stack. The schematic and photographs of different cooling options: (a) IFC6410 board without cooling, (b) proposed in-package fluidic prototype mounted on the die, and (c) same board with passive air cooling solution.

system, Section IV discusses measurement results, and Section V concludes this paper.

II. RELATED WORK

Active cooling with microfluidics has been studied for high-performance computing [9]–[13]. The state-of-the-art methodology integrates fluidic systems on a silicon substrate or an interposer [13]. Applying the fluidic heat sink on the chip package improves heat-removal capacity. Furthermore, forming fluidic channels and micropin fins directly on the die may lead to even more effective cooling [14]. Tuckerman and Pease [9] and Sarvey *et al.* [10] reported impressively low thermal resistance measurements of 0.09 and 0.0389 K/W, respectively. However, fabricating channels/pin fins directly on the chip requires foundry support, which is difficult to adopt for relatively cost-sensitive systems. On the other hand, external “on-package” heat sinks are less effective and increase the SoC's footprint on the board. The ease of the integration, the system footprint, and the power dissipation associated with active cooling are additional criteria to determine a suitable cooling solution. The prior work on in-package fluidic cooling had shown integration feasibility with dummy silicon dies/heaters [9]–[13]. The simulation-based evaluation also has shown the fluidic cooling improves the system-level energy efficiency [15]–[17]. For example, Wan *et al.* [15] have considered the microarchitecture simulation of the 3-D ICs with fluidic cooling. Serafy *et al.* [16] have speculated that energy reduction may be achieved through chip stack cooling

simulation considering the pump energy. Xiao *et al.* [17] further modeled a high-performance system with a more advanced pin fin cooling channel. The prior works mainly focused on the high-power density systems, but the role of fluidic cooling in embedded SoC has been less studied.

There have been many studies on advanced heat spreading materials on cooling stack of small form factor [18], [19]. A recent high-performance mobile device has already integrated heat pipe for the thermal management [20]. To sustain the performance and maintain low-contact temperature, nanomaterials such as graphite sheets were utilized. Shaping the electromagnetic interference (EMI) shield on die that forms cavity pocket avoiding direct path from hot spot to enclosure surface has been reported in [21]. High-thermal conductivity materials may also be infused in the coolant to enhance the convection heat transfer [22]. Measurements independently reported by Wagner and Maltz [18] and Gurrum *et al.* [23] reported the average tablet surface temperatures are roughly 35 °C when the enclosure hotspots are at 41 °C or above in the ambient environment at 25 °C [18], [23].

We claim the novelty in the mobile form factor closed-loop cooling with software-based policies to control fluid flow employed in the SoC. We demonstrate a self-contained cooling loop for active fluidic cooling in the mobile form factor in Fig. 1. The experimental characterization of the in-package fluidic cooling contains electrical measurement setup in Fig. 1(a). A cartoon of the board and pump assembly used for the measurements is in Fig. 1(b) and (c). In parallel, fully custom 3-D stack system for high-performance computing has been explored in depth by Sarvey *et al.* [10]. However, the prior experimental characterizations considered a temperature-controlled fluid reservoir, which is infeasible in a mobile environment, and did not discuss the possibility of on-chip activity base fluid control. This paper significantly advances the state-of-the-art by demonstrating a closed-loop system in term of fluid control and in term of heat exchange to ambient. The proposed system considers the implication of using passive heat-sink-to-ambient-heat exchange to support mobile devices. Because the physical properties of the air convection are fixed with given form factor, the design focuses on heat spreading within the housing enclosure of the SoC platform and improving effective thermal capacitance. Compared to state-of-the-art solutions, the controllability of the active pump improves design simplicity and opens a degree of freedom for the choice of fluid/pressure, and allows economical heat sink material integration to the system.

III. SYSTEM INTEGRATION

A. Embedded SoC Platform

The platform for embedded SoC evaluation was the IFC6410 board from *Inforce Computing*. The embedded 28-nm SoC on the board is a Snapdragon 600 SoC with an Adreno graphics engine and an up to 1.9-GHz quad-core system. The SoC is flip-chip bonded, and backside of the die is exposed, which allows direct attachment of the silicon pin fin. The system runs on the operating system Linaro-Gnome, a Linux advanced




Fig. 2. Schematic and pictures of different cooling options. (a) IFC6410 board without cooling. (b) Proposed in-package fluidic prototype mounted on the die. (c) Same board with passive air cooling solution.

RISC machine distribution. The SoC's built-in temperature sensors collect thermal information of the chip during operation. Out of the 13 thermal sensors onboard, only four sensors for processors are considered for the analysis. The control framework on the SoC roughly consumes 2% core one's bandwidth in the background increasing core one's temperature by 2 °C higher than the remaining three cores. A watcher script written in bash reads the thermal register values and drives the pump on a preset thermal or frequency threshold. We implemented an additional debug, characterization, reporting background scripts to communicate with a host machine. For data logging purpose, the thermal information is averaged into one aggregated reading and transmitted to a desktop server at a fixed 1-s interval. A transmission control protocol (TCP) communication client updates the value to the connecting server on the network. Along with the thermal information, the core clock frequencies are also uploaded with the same framework.

B. Integration of Cooling Technology With SoC

The SoC with integrated in-package silicon-based active fluidic cooling and the baseline systems has been constructed. The in-package fluidic system is benchmarked against the bare die and natural convection heat sink.

1) *No Cooling*: The exposed die configuration came with the original system. The SoC relies on thermal throttling and frequency scaling to achieve thermal management. The stock system reaches 70 °C under a moderate workload without external intervention. The system configuration is shown in Fig. 2(a). The heat from the die will directly dissipate to the ambient air by natural convection. Although there is no thermal resistance of the thermal interface material and heat sink, the convection thermal resistance is very high due to the low natural heat transfer coefficient.

2) *Proposed Active In-Package Fluidic Cooling*: The micropin fin system is integrated with the Omegatherm 201 thermal paste between die and heat sink. The system was secured with electrical tape on top of the die. The microfluidic-pin fin region aligned with the active SoC die area. Note in Fig. 2(b) that the tubing and 20-mm extended connector areas on each side are not necessary for a commercially

integrated in-package fluidic cooler. The connection would be routed through the fluidic trace embedded in the printed circuit board [13]. Due to the high conductivity of the Si and its thin thickness, the conduction thermal resistance of the silicon pin fin channel is very small. Same for the microfluidic cooling, the heat transfer coefficient can be higher than the conventional air cooling by two orders of magnitude, and the convection thermal resistance is also very small. The thermal resistance of the thermal interface material can be dominant [14].

3) *Baseline Passive Air Cooling Heat Sink*: The natural air convection relies on the heat sink surface area to remove heat from the SoC. The attached aluminum heat sink has a dimension of 40 mm × 50 mm with fin height 30 mm. The heat sink was designed for the AMD RS780L chipset with 10-W peak power. The assembly is shown in Fig. 2(c). The heat sink and thermal interface material add additional thermal resistance between the die and ambient air. However, due to the improved heat spreading and larger surface area of the heat sink, the total thermal resistance can be reduced. The disadvantages are the limited cooling capability of the air-cooled heat sink and its increased total size of the device.

C. In-Package Fluidic Cooling

Incorporating a heat-exchange layer in a direct contact to the SoC reduces the overall thermal resistance. In conventional external (on package) cooling, the thermal solution is integrated during component assembly. Traditional metal heat sink may not be directly attached to a silicon die due to potential EMI challenges and the mismatch in the thermal expansion coefficients. We present a silicon-based fluidic interposer designed to be permanently attached to the die and enclosed inside the EMI shield in Fig. 3(a). The microgap fabrication process started with a double-sided polished 4-in silicon wafer with a thickness 500 μm. In the first step, positive photoresist SPR-220 was spun and exposed to form a mask of the micro gap. Then, the wafer was etched in the deep reactive ion etching (DRIE) process. Using the standard Bosch process, which alternates between a plasma etching step and passivation step, the deep micro gap cavity with staggered micropin fin array was etched due to its high heat transfer coefficient [24].

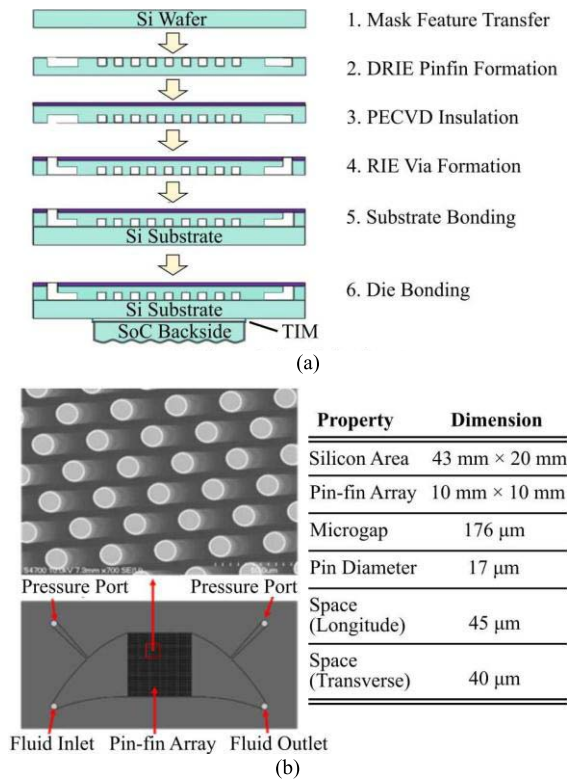


Fig. 3. Fabricated device and its corresponding features are highlighted. (a) Key steps for micro pin fin fabrication. (b) Parameters and the SEM image.

Tencor P15 profilometer was used to record the depth of the microgap. In the second step, the wafer was flipped and a $2\text{-}\mu\text{m}$ thickness silicon oxide layer was deposited by the plasma-enhanced chemical vapor deposition method as an insulation layer.

In the third step, the wafer was taken through a photolithography step and a reactive ion etching (RIE) process to remove the oxide and expose the silicon that was to be etched to form the fluid vias. After the RIE process, the wafer was put into DRIE to continue to etch the silicon and developed the fluid vias. Thereafter, the processed wafer was diced and the microgap samples were taken out of the wafer. In the last step, a $500\text{-}\mu\text{m}$ -thick silicon wafer was bonded to the diced microgap samples by epoxy to form a sealed device. The dimensions of the fabricated device are $43 \text{ mm} \times 20 \text{ mm}$. The microgap also includes pressure ports at the fluid inlet and outlet, which are not used in this paper. The area of the pin fin array is $1 \text{ cm} \times 1 \text{ cm}$. The depth of the microgap is $176 \mu\text{m}$, the diameter of the pins is $17 \mu\text{m}$, longitudinal spacing is $45 \mu\text{m}$, and the transversal spacing is $40 \mu\text{m}$. A fabrication flow is shown in Fig. 3(a). Fig. 3(b) shows an image of the fabricated device with the SEM image of the staggered pin fin array and the highlighted key parameters. The final pin fin assembly is $43 \text{ mm} \times 20 \text{ mm} \times 1 \text{ mm}$ without the fluid inlet and outlet tubing. This is the number we quoted comparing against the baseline heat sink volume.

D. Piezoelectric Pump

The choice of pump for active cooling in SoCs is limited by the motor's physical geometry and power consumption.

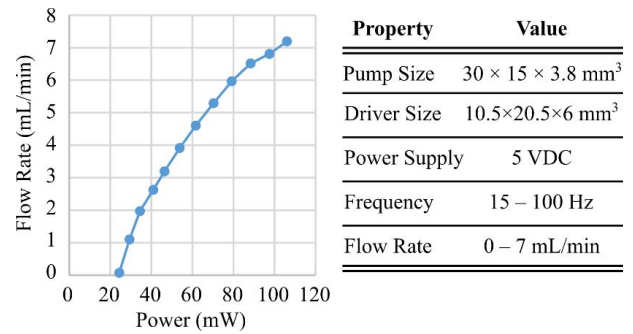


Fig. 4. Power characteristic and the corresponding parameters associated with the piezoelectric pump.

We considered compact high-flow-rate piezoelectric pumps for this investigation. As an example, the pump model MP6 manufactured by Bartels mikrotechnik is chosen for this paper. The pump dimension is $30 \text{ mm} \times 15 \text{ mm} \times 3.8 \text{ mm}$ and the driver board dimensions are $10.5 \text{ mm} \times 20.5 \text{ mm} \times 6 \text{ mm}$. The pump's peak power is 110 mW. The piezoelectric pump has peak pumping flow rate at 7 mL/min. The pump standby power is measured to be 10 mW. The piezoelectric pump's flow rate may be programmed with digitally controlled pulse-frequency modulation (PFM) driver. While the pump itself has a wider operating range, the operation point of interest for this work is tabulated in Fig. 4.

The PFM controllable micropump was powered through the 5-V rail from the power management-integrated circuit (PMIC) on the IFC6410 board. The onboard general-purpose input/output (GPIO) controlled the pump's PFM clock source. The total power of the system including the board (SoC + peripheral) and pump was measured with TCP202 hall-effect sensor on the 5-V line near the board's power socket.

E. Calibration Fluidic Loop

The configuration of the fluidic loop consists of the pump, controlled temperature reservoir, flowmeter, filter, and the SoC chip with the in-package fluidic cooling [see Fig. 1(b)]. A swappable Cole-Parmer digital gear pump that is capable of flow rates from 5.52 to 331.2 mL/min was also used during calibration besides the MP6 pump. A part of the flow loop was immersed into a controlled temperature bath. A McMillan S-114 flowmeter was calibrated to measure the volumetric flow rate. A $90\text{-}\mu\text{m}$ Swagelok filter was used to keep the inlet water clean and prevent clogging of the microgap.

F. Compact Fluidic Loop

Once calibrated, the compact fluidic loop only contains the micropump and enclosure heat sink that is designed to be small enough to be a cell phone back plate [see Fig. 1(c)]. Limited by the system-mobility criteria, the heat spreader to ambient area is confined to lesser or equal to the device housing enclosure. The on-die fluidic cold plate may serve as both sprint computing heat buffer and hot-spot spreading layer. The active cold plate on the SoC carries heat away quickly through convection to the heat sink, and away from the area directly on top of the SoC. This prevents uneven thermal conduction




Fig. 5. Fluid to ambient coldplate’s mechanical drawing and machined assembly is shown.

to the stack on top of the SoC and produce hot spot on mobile device surface.

In a handheld system, the SoC thermal design point (T_{SoC}) has an upper limit of 90 °C. The peak SoC power (P_{SoC}) is roughly at 8 W. The enclosure surface temperature (T_S) is limited by touch temperature at 41 °C and is a function of the overall thermal resistance and lateral heat spreading. In order to satisfy enclosure contact temperature and the SoC’s thermal design point (TDP), the former is bounded by where $T_{SoC} < 90$ °C, $T_S < 41$ °C, and $P_{SoC} < 8$ W. The constraint limits the vertical $R_{TH} < 6.1$ K/W, which is not difficult to achieve even for mobile form factor. Low-cost housing material such as acrylic glass or plastic may still be used at millimeter thicknesses. The more stringent limitation for handheld systems is the spreading resistance from the SoC to the edge of the enclosure. The limitation may be significantly improved by forced convection. In this paper, an acrylic-based cold plate and an aluminum-based cold plate are designed in Fig. 5. The plate area is 60 mm × 132 mm and the fluid area is 50 mm × 111 mm × 0.5 mm. The pin fins are 2 mm × 2 mm squares. The longitudinal spacing is 4 mm and transversal spacing is 3 mm. The fins are designed for structural support. The acrylic-based design demonstrates the possibility of the display side cooling. The aluminum-based design resembles the back-side cooling. Both plates are covered with plain 1-mm acrylic sheet.

The choice of nonmetal heat spreader drastically improves the material cost of the overall design (see Fig. 6). A self-contained fluidic cooled SoC can adopt the proposed methodology without board redesign. This approach converts application-dependent and cost-conscious system to fluidic cooled design with ease.

IV. EXPERIMENTAL OBSERVATION ON FLUIDIC COOLED SOC

A subset of *Splash-2* benchmark suite was used to demonstrate the thermal behavior of the SoC with different cooling technologies considering the workload [25]. The Linux




Fig. 6. Integration configuration with metal/clear surfaces, the configuration was used for data collection for closed-loop control. We showcase different material for heat spreading; modeling handheld/mobile form factor designs.

workload generator tool, Stress, was used to bring the thermal condition to the steady state [26]. In all benchmarks, the stress tool spun $\text{sqrt}()$ on all four cores. The benchmarks ran on all four cores when the compiled code supports multiprocessing.

A. Pump Power and SoC Power Tradeoff

If the leakage power at a higher die temperature exceeds the power required to run the active cooling loop, an SoC platform with an active cooling system may consume lower power compared to an existing baseline platform with passive cooling. The preceding hypothesis is validated in Fig. 7. We first ran the stress tool in all the four cores over a period that was long enough to reach a steady-state temperature. Next, we terminated the application to reduce power and allowed the temperature to cool to steady state. The power dissipation of the board and the temperature of the SoC were measured by enabling and disabling the fluidic loop. A measured flow rate of 7 mL/min was considered in the experiment. Fig. 7(a) shows the power and temperature when stress was running. We observed the SoC heats up to 55 °C within 1 min of operation, and 70 °C for a continuous full load of 10 min. When the fluidic loop was disabled, the system consumes additional 174 mW of power at the 1-min mark and additional 490 mW of power at the 10-min mark. We believe the additional power was due to the increased temperature induced leakage current. Fig. 7(b) shows the power and temperature during the low-workload (idle) condition. With the fluidic loop “turned-OFF,” although the SoC temperature remains higher, the system power becomes lower than the case when pump was ON. We believe this was because, the aggressive circuit/microarchitecture level idle power management techniques in commercial mobile SoCs significantly reduce the leakage current during low-workload conditions. Consequently, the overhead associated with the pumping power made the system power larger with the fluidic loop turned ON. This observation leads to an optimization

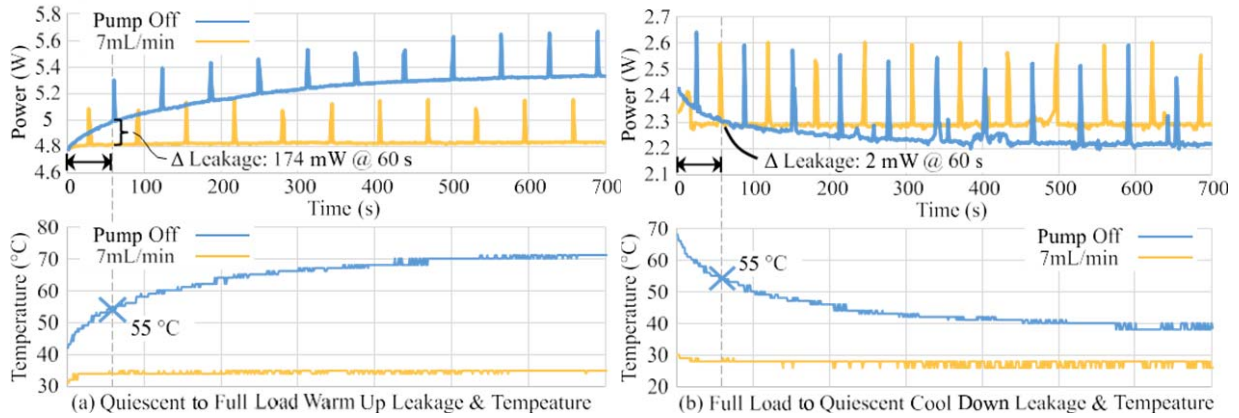


Fig. 7. Measurements show the system power and the SoC temperature following enabling/disabling of the fluidic loop in the in-package cooling technology. (a) At a high workload condition with full utilization of the cores, the system without active cooling operates at a higher temperature and sustains a higher leakage. The active cooling reduces temperature, and hence leakage, to reduce the total system power even after accounting for the pumping power. (b) On the other hand, in the idle or low utilization condition, the SoC employs aggressive idle power management to electrically minimize leakage power; consequently, the temperature reduction with the active cooling does not translate to power saving. The pumping power overhead makes the fluidic cooling less efficient. The measurement shows the need to couple electrical power management techniques with active fluidic cooling for an optimal power management system targeting low-power SoCs.

TABLE I
SYSTEM STEADY-STATE POWER AND TEMPERATURE
COMPARISONS WITH ASSEMBLY CLEARANCE

Cooling Structure	Symbol	Peak Temp.	Peak Power	Footprint	Height
No Cooling	nc	84 °C	6.28 W	N.A.	N.A.
Passive Air Cooling	ac	64 °C	5.83 W	2000 mm ²	30 mm
In-package Fluidic Cooling	ifc	40 °C	5.39 W	800 mm ²	1 mm

problem between pump power and leakage power. An efficient heat exchange can lead to system-level power reduction by balancing leakage temperature and pump power.

B. Steady-State Temperature at Full Utilization

We next consider temperature at the full utilization scenario. The comparisons of system power, temperature, and footprint/height of the cooling solution are shown in Table I. All measurements were performed considering the same “stress” workload. The fluidic channels were driven from the same-pump at fixed 7 mL/min for both external and in-package cooling. The system power includes the pump power. The measurement results show that the in-package cooling can reduce temperature by 24 °C and, hence, the leakage current, resulting in 450-mW power saving over passive cooling during the peak workload condition, even after accounting for the pumping power. The reduced temperature may have an additional benefit of slowing down the aging process of the devices and may improve the lifetime of the SoC.

C. Application-Dependent Power

The active fluidic cooling improves the performance through avoiding overheating beyond the TDP. The transient temperature measurements were performed considering the *Splash-2* benchmark. The application, *raytrace*, which is known to be unfriendly to dynamic voltage and frequency scaling (DVFS), was applied [27]. Fig. 8 shows the *raytrace* application on both the bare system and that with in-package fluidic cooling. The system without fluidic cooling had limited thermal headroom. Once the hot cores exhausted their sprinting budget

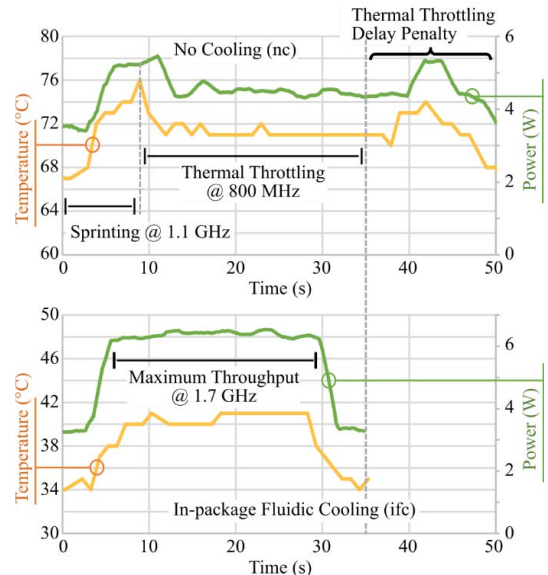


Fig. 8. Measurement results of the temperature and power characteristics with the bare-die (no cooling) case and the in-package-active-cooling case. The traces were collected from the benchmark “Raytrace.” Without any thermal management, the higher temperature limited operating time in high-performance (high power) mode and induced throttling, thereby increased the computation time. The higher computation time led to higher energy dissipation. The system with the active in-package cooling ran at a higher power mode without throttling resulting lower computation time and, hence, lesser energy dissipation.

in a short period, the internal power controller of the SoC forced the system to operate at a reduced power state to stay within the thermal threshold. Moreover, the thermal throttling also incurred a delay penalty. When active in-package cooling is used, we observed that the SoC stayed at a higher power state during the entire operation, without violating thermal constraints. Consequently, the application runtime and the computation energy were reduced. Overall, we observed the system without cooling consumed 31.5% more energy than the in-package fluidic cooling for

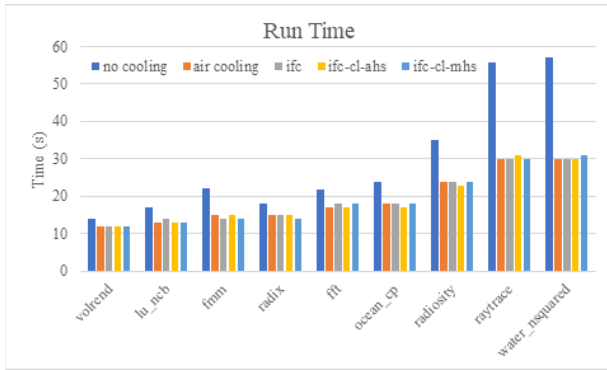


Fig. 9. Measurement results for various Splash-2 benchmarks running on the SoC for cumulative execution time.

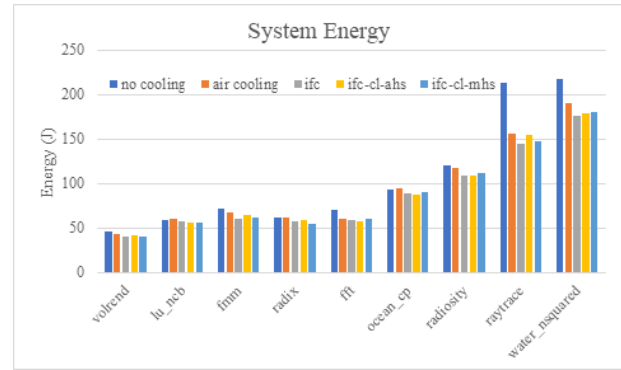


Fig. 10. Measurement results for various Splash-2 benchmarks running on the SoC for the benchmark computation energy.

the same workload. The chip peak temperature was also 32 °C higher.

D. Mobile Heat Sink Performance

The measured temperature differences for aluminum plate cooler between the inlet and outlet are 30.6 °C and 28.3 °C, respectively, at the peak 4.8-W board power. The acrylic-based fluidic sink temperatures are 31.3 °C and 29.6 °C, respectively. The temperature percentage difference between metal and acrylic material is less than 3%. This demonstrated the passive heat exchange for the low-power system dominated by air and the high thermal conductive material is not of a major concern. Furthermore, the forced convection shows lesser than 8% surface temperature gradient across the inlet and outlet for the aluminum cold plate and 6% for the acrylic cold plate.

E. System Benchmarks

Five systems were compared: the baseline case with no cooling, the passive heat sink cooling, and different scenarios of the proposed system with the microfluidic cooling. On the radiator side, the constant temperature bath (in-package fluidic cooling—*ifc*), the acrylic cold plate (in-package fluidic cooling and acrylic heat-spread—*ifc-cl-ahs*), and the aluminum cold-plate (in-package fluidic cooling with closed-loop and metal heat spreader—*ifc-cl-mhs*) are benchmarked. The bare system without the additional thermal management was thermally throttled to a lower power state during the execution and took more time to complete, as shown in Fig. 9. Longer benchmarks such as *raytrace* and *water_nsquared* show the most slowdown due to exhausting thermal headroom. All the systems with cooling (passive and in-package fluidics) sustained comparable runtime without significant throttling. The energy dissipation was calculated considering the total power (board+pump) and the completion time, in Fig. 10. The cumulative leakage power was more apparent in longer benchmarks. In all the applications, the energy advantage of the in-package fluidic cooling over the no cooling was up to 2%–47% (20%–47% for top 3). In a throughput-based application, like *raytrace*, or applications with significant higher floating point operations per second (e.g., *fmm*, *water_nsquared*), the in-package fluidic cooling showed energy advantages over the passive cooling by

3%–8% (7%–8% for top 3). Furthermore, the DVFS-friendly benchmark *ocean cp* (because of its process synchronization barriers) still consumed more power in the passive cooling compared to the in-package fluidic cooling by 6%. Compared to closed-loop systems, the nonmetal solution consistently beats passive cooling power consumption by 6%–7% for benchmarks longer than 20 s, except for *raytrace*, which shows <1% advantage. Note that all these measurements are of system energy usage, which includes the pump, executing realistic CPU load, on an SoC with advanced DVFS power reduction techniques, built with off-the-shelf parts except custom pin fins, and still shows 1%–6% energy advantage over the traditional approach. We claim such a technique as net-zero-energy cooling that the operating temperature of SoC, compared with that in traditional techniques, reduces by 19 °C–23 °C at no additional operating cost in the system; see Fig. 11. For the same T_j max limit, the power dissipation for the in-package fluidic cooling can be further increased beyond the capability of the air cooling.

F. Closed-Loop Control Methodology

Fluidic cooling is an effective methodology for transmitting heat to the enclosure surface and reducing in-package temperature. The active cooling system takes advantage of the on-chip power reduction to power the pump, and still achieve lower system-level power. However, the workload and power on the die can be transient, depending on the running content. As we can see in Section IV-A, the power dissipation at idle state is very low and thus the temperature is very low. There is no need for the active liquid cooling. Therefore, further control analysis has been performed to balance the pump power and on-chip temperature. Few simplistic control mechanisms are benchmarked: always on (ao), maximum temperature threshold (mt), and maximum frequency threshold (mt). A watcher script written in bash reads the thermal register values and enables the pump on a preset thermal or frequency threshold. The SoC drives the GPIO pins to enable the pump based on these thresholds. Since the control loop is self-regulated onboard, the manage power and sampling overhead are also considered by construction. Note that all our closed-loop benchmarks consider the power saving during benchmark code execution. We are not benchmarking with idle cores similar to

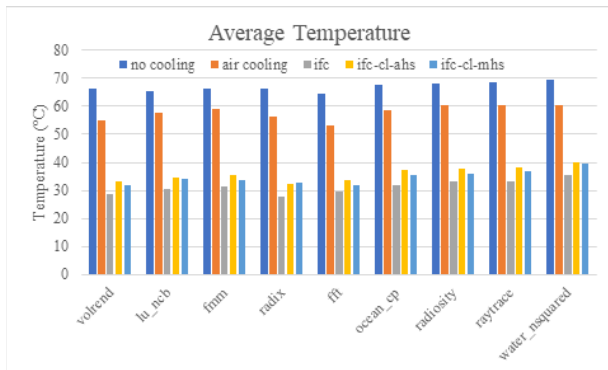


Fig. 11. Measurement results for various Splash-2 benchmarks running on the SoC for the average temperature.

experiments conducted in Fig. 7. The experiment is to show improvement on runs with at least one core active at any given time.

1) *Always On*: The always on control scheme is equivalent to the in-package fluidic cooling experiments where the pump is constantly driven to the maximum flow rate. Whether the workload and thus the power dissipation of the die changes or not, the pump flow rate will be kept constant. In this case, the pump flow rate should be high enough to account for the worst power scenario. Thus, overcooling can happen, which lowers the system energy efficiency.

2) *Maximum Temperature Threshold*: A bang–bang controller is implemented for the fluidic pump in software on the SoC. The script completely shut off the pump’s pulse frequency modulator signal and the controller power gate pump’s boost regulator for minimal sleep power consumption. The temperature threshold for each processor is polled every second and when any processor’s temperature exceeds the temperature threshold, the enable signal turns on the pump and the boost regulator. The temperature threshold is selected to be 55 °C. In this case, if the temperature is high, the pump flow rate will increase to reduce the die temperature until the temperature goes below the threshold. If the die temperature is low, the pump flow rate can be reduced. This algorithm can increase the energy efficiency.

3) *Maximum Frequency*: A similar bang–bang controller is implemented for the fluidic pump in software. The frequency threshold for each processor is polled every second, and when any processor’s frequency exceeds the frequency threshold, the enable signal turns on the pump and the boost regulator. And frequency threshold is selected to be 1 GHz.

The steady-state power consumption shows less than 1% of difference between each cooling methodology at full load. The time to completion is shown in Fig. 12. The computation energy is shown in Fig. 13. The temperature is shown in Fig. 14. The energy benchmark shows the energy saving applying different techniques is not significant. The major reason is the temperature in all the cooling objective is not producing significant temperature difference. The pump power itself also is relatively small compared to the board power, roughly 2% of the total power. Any improvement in the pump power will be less prominent during active computing. During quiescent state, the frequency-based activation achieves

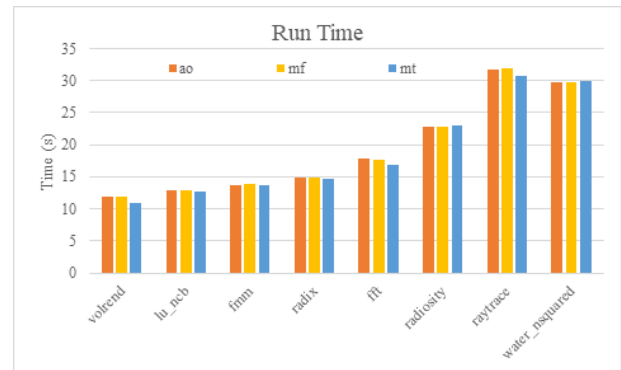


Fig. 12. Measurement results for various Splash-2 benchmarks running on the SoC for the completion time. The results highlight the pump enabling policy.

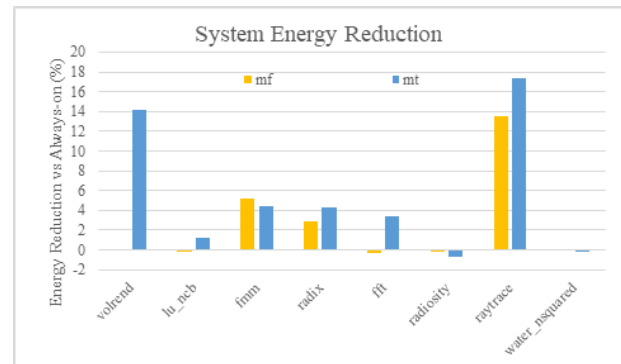


Fig. 13. Measurement results for various Splash-2 benchmarks running on the SoC for the total computation energy. The results highlight the pump enabling policy.




Fig. 14. Measurement results for various Splash-2 benchmarks running on the SoC for the average temperature. The results highlight the pump enabling policy.

79 mW of power reduction because pump is mostly offline. The temperature-based activation achieves 41 mW of power reduction due to constant temperature regulation. In active states, the techniques reduce 0%–16% of the overall power under DVFS (Fig. 13). As an energy-conscious cooling policy, the frequency-based cooling should be employed because of the benefit of lower quiescent power and advantage from the DVFS leakage reduction. The temperature threshold-based control does not seem necessary for fluid control especially if the thermal throttling is built in the SoC. For the three techniques, the frequency threshold-based control for pump is observed to be the most effective approach. The benefit of in-package active cooling can be more significant if the

power/frequency of the device increases. Dynamic fluid control methodologies have been implemented on a real device and show higher energy efficiency. In this paper, the pump flow rate is controlled after the temperature or frequency threshold is reached. This might cause some delay in response. A more efficient control can be that the pump flow rate is controlled based on the projection of the workload and temperature.

V. CONCLUSION

This paper experimentally demonstrated closed-loop in-package fluidic cooling for a commercial SoC platform. A chip-scale in-package microfluidic cooling technique based on micropin fins in a silicon interposer was fabricated and attached to the commercial SoC. The onboard thermal management was demonstrated using a low-power piezoelectric pump controlled by the SoC. Compared to a baseline (no cooling) and an external passive cooling, measurements using the *Splash-2* benchmark showed that in-package cooling achieved 2%–47% and 3%–8% less energy consumption, respectively. In the closed-loop measurement, in-package cooling with simple acrylic heat sink achieved thermal advantages of 19 °C–23 °C and power advantages of 1%–6% over external passive cooling. Moreover, in-package cooling had a reduced assembly footprint and height compared to external passive and fluidic cooling. Our observations suggested that in mission-critical operations when cores must operate at the maximum load to deliver the required throughput, the in-package cooling solution can successfully complement the electrical techniques (e.g., power gating, voltage–frequency scaling) that manage temperature and reduce total system power. Although the power in this paper was not very high, the work demonstrated the feasibility of integrating compact chip-scale fluidic cooling structures into SoC without the need of fabricating channels/pin fins directly on a silicon die. The successful in-package fluidic cooling integration with a commercial SoC will motivate future work on a co-design between innovative cooling structures and advanced SoC power management. The active in-package cooling may bridge efficient fluidic control and workload management that conducts online thermal-power co-optimization.

REFERENCES

- [1] M. J. Ellsworth, L. A. Campbell, R. E. Simons, M. K. Iyengar, R. R. Schmidt, and R. C. Chu, "The evolution of water cooling for IBM large server systems: Back to the future," in *Proc. Thermal Thermomech. Phenomena Electron. Syst.*, 2008, pp. 266–274.
- [2] T. J. Chainer, M. D. Schultz, P. R. Parida, and M. A. Gaynes, "Improving data center energy efficiency with advanced thermal management," *IEEE Trans. Compon., Packag., Manuf. Technol.*, vol. 7, no. 8, pp. 1228–1239, Aug. 2017.
- [3] B. C. Ward, J. L. Herman, C. J. Kenna, and J. H. Anderson, "Outstanding paper award: Making shared caches more predictable on multicore platforms," in *Proc. Real-Time Syst. (ECRTS)*, Jul. 2013, pp. 157–167.
- [4] J. Gubbi, R. Buyya, S. Marusic, and M. Palaniswami, "Internet of Things (IoT): A vision, architectural elements, and future directions," *Future Generat. Comput. Syst.*, vol. 29, no. 7, pp. 1645–1660, 2013.
- [5] L. Johnson. (2014). *Sony Answers 4K Overheating Concerns*. [Online]. Available: <http://www.trustedreviews.com/news/sony-answers-4k-overheating-concerns-says-only-shoot-in-small-bursts>
- [6] W. Song, S. Mukhopadhyay, and S. Yalamanchili, "Architectural reliability: Lifetime reliability characterization and management of many-core processors," *IEEE Comput. Archit. Lett.*, vol. 14, no. 2, pp. 103–106, Jul./Dec. 2014.

- [7] C.-C. Chen and L. Milor, "Microprocessor aging analysis and reliability modeling due to back-end wearout mechanisms," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 23, no. 10, pp. 2065–2076, Oct. 2014.
- [8] Z. Wan, W. Yueh, Y. Joshi, and S. Mukhopadhyay, "Enhancement in CMOS chip performance through microfluidic cooling," in *Proc. Thermal Invest. ICs Syst. (THERMINIC)*, 2014, pp. 1–5.
- [9] D. B. Tuckerman and R. F. W. Pease, "High-performance heat sinking for VLSI," *IEEE Electron Device Lett.*, vol. EDL-2, no. 5, pp. 126–129, May 1981.
- [10] T. E. Sarvey *et al.*, "Embedded cooling technologies for densely integrated electronic systems," in *Proc. Custom Integr. Circuits Conf. (CICC)*, Sep. 2015, pp. 1–8.
- [11] Z. Wan *et al.*, "Thermal analysis and improvement of high power electronic packages," in *Proc. 12th Int. Conf. Electron. Packag. Technol. High Density Packag.*, Aug. 2011, pp. 1–5.
- [12] W. R. Hamburg and J. S. Fitch, "Packaging a 150-W bipolar ECL microprocessor," *IEEE Trans. Compon., Hybrids, Manuf. Technol.*, vol. 16, no. 1, pp. 28–38, Feb. 1993.
- [13] J. H. Lau, "TSV interposers: The most cost-effective integrator for 3D IC integration," *Chip Scale Rev.*, vol. 15, no. 5, pp. 23–27, 2011.
- [14] Z. Wan and Y. Joshi, "Pressure drop and heat transfer characteristics of pin fin enhanced microgaps in single phase microfluidic cooling," in *Proc. ASME Int. Mech. Eng. Congr. Expo.*, Nov. 2013, p. V010T11A086.
- [15] Z. Wan, H. Xiao, Y. Joshi, and S. Yalamanchili, "Co-design of multi-core architectures and microfluidic cooling for 3D stacked ICs," *Microelectron. J.*, vol. 45, no. 12, pp. 1814–1821, 2014.
- [16] C. Serafy, A. Srivastava, and D. Yeung, "Continued frequency scaling in 3D ICs through micro-fluidic cooling," in *Proc. Thermal Thermomech. Phenomena Electron. Syst. (ITherm)*, 2014, pp. 79–85.
- [17] H. Xiao, Z. Wan, S. Yalamanchili, and Y. Joshi, "Leakage power characterization and minimization in 3D stacked multi-core chips with microfluidic cooling," in *Proc. Semiconductor Thermal Meas. Manage. Symp. (SEMI-THERM)*, 2014, pp. 207–212.
- [18] G. R. Wagner and W. Maltz, "On the thermal management challenges in next generation handheld devices," in *Proc. ASME Int. Tech. Conf. Exhib. Packag. Integr. Electron. Photon. Microsyst.*, 2013, p. V002T08A046.
- [19] J. Lee, D. W. Gerlach, and Y. K. Joshi, "Parametric thermal modeling of heat transfer in handheld electronic devices," in *Proc. Thermal Thermomech. Phenomena Electron. Syst.*, 2008, pp. 604–609.
- [20] H. Uchida, T. Shioga, S. Aoki, S. Ogata, and H. Nagaoka, "Loop heat pipe," U.S. Patent 13 591 397, Aug. 22, 2012.
- [21] L. Shao *et al.*, "On-chip phase change heat sinks designed for computational sprinting," in *Proc. Semiconductor Thermal Meas. Manage. Symp. (SEMI-THERM)*, 2014, pp. 29–34.
- [22] S. S. Gupta *et al.*, "Thermal conductivity enhancement of nanofluids containing graphene nanosheets," *J. Appl. Phys.*, vol. 110, no. 8, p. 084302, 2011.
- [23] S. P. Gurrum *et al.*, "Generic thermal analysis for phone and tablet systems," in *Proc. Electron. Compon. Technol. Conf. (ECTC)*, 2012, pp. 1488–1492.
- [24] T. Brunschweiler *et al.*, "Interlayer cooling potential in vertically integrated packages," *Microsyst. Technol.*, vol. 15, no. 1, pp. 57–74, Aug. 2008.
- [25] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta, "The SPLASH-2 programs: Characterization and methodological considerations," *ACM SIGARCH Comput. Archit. News*, vol. 23, no. 2, pp. 24–36, 1995.
- [26] A. Waterland. (2014). *Stress*. [Online]. Available: <http://people.seas.harvard.edu/apw/stress/>
- [27] W. Kim, M. S. Gupta, G.-Y. Wei, and D. Brooks, "System level analysis of fast, per-core DVFS using on-chip switching regulators," in *Proc. High Perform. Comput. Archit.*, Feb. 2008, pp. 123–134.



Wen Yueh (S'08–M'15) received the B.S. and M.S. degrees in electrical and computer engineering from Rutgers University, New Brunswick, NJ, USA, in 2009, and the Ph.D. degree from the Georgia Institute of Technology, Atlanta, GA, USA, in 2015.

His current research interests include system level multiphysics simulator, self-adaptive circuit design for many-core processor thermal management, and energy aware low-power memory architecture.



Zhimin Wan received the B.S. and M.S. degrees from the Huazhong University of Science and Technology, Wuhan, China, in 2009 and 2011, respectively, and the Ph.D. degree in mechanical engineering from the Georgia Institute of Technology, Atlanta, GA, USA, in 2016.

He is involved in on-chip microfluidic cooling.



He Xiao (S'14) is currently pursuing the Ph.D. degree with the Georgia Institute of Technology, Atlanta, GA, USA.

His current research interests include simulating and characterizing physical effects on multicore microarchitectures using the 3-D CMOS technology, as well as exploring adaptive architectures based on thermal analysis, and also include computer architecture, low-power design, programming models, and compiler optimization.



Sudhakar Yalamanchili (S'79–M'82–SM'91–F'14) received the B.E. degree in electronics from Bangalore University, Bengaluru, India, and the Ph.D. degree in electrical and computer engineering from the University of Texas at Austin, Austin, TX, USA, in 1984.

He was a Senior Research Scientist with the Honeywell Systems and Research Center, Minneapolis, MN, USA, where he then became the Principal Research Scientist. He was the Principal Investigator for projects in the design

and analysis of multiprocessor architectures for embedded applications. He served as a member of Honeywell's Program Technical Advisory Board to Microelectronics and Computer Technology Corporation, Austin, and was an Adjunct Faculty with the Department of Electrical Engineering, University of Minnesota, Minneapolis, MN, USA. He joined the Georgia Institute of Technology, Atlanta, GA, USA, in 1989, where he is currently a Regents Professor and a Joseph M. Pettit Professor of computer engineering with the School of Electrical and Computer Engineering.

Dr. Yalamanchili is a member of the ACM. He contributes professionally with regular service on editorial boards and conference and workshop program committees. He served on the Editorial Board of *Computer Architecture Letters* from 2011 to 2015, as the Program Co-Chair for the IEEE/ACM International Symposium on Networks on Chip in 2014, and the IEEE International Symposium on Workload Characterization in 2015, and on program committees for the IEEE/ACM Supercomputing in 2017, the IEEE/ACM International Symposium on Microarchitecture in 2015, the IEEE International Symposium on High-Performance Computer Architecture in 2014, and the IEEE/ACM International Symposium on Computer Architecture in 2014. He currently serves as a Co-Director Center for Experimental Research in Computer Systems.



Yogendra Joshi (SM'03–F'12) received the Ph.D. degree in mechanical engineering and applied mechanics from the University of Pennsylvania, Philadelphia, PA, USA, in 1984.

He is currently a Professor and John M. McKenney and Warren D. Shiver Distinguished Chair with the G.W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA, USA, where he is the Principal Investigator with the Office of Naval Research Consortium for Resource-Secure Outposts, and the Site Director for the National Science Foundation Industry/University Cooperative Research Center on Energy Efficient Electronic Systems. His current research interests include multiscale thermal management.

Dr. Joshi is an elected Fellow of the ASME and the American Association for the Advancement of Science. He was a co-recipient of the ASME Curriculum Innovation Award in 1999, the Inventor Recognition Award from the Semiconductor Research Corporation in 2001, the ASME Electronic and Photonic Packaging Division Outstanding Contribution Award in Thermal Management in 2006, the ASME J. of Electronics Packaging Best Paper of the Year Award in 2008, the IBM Faculty Award in 2008, the IEEE SemiTherm Significant Contributor Award in 2009, the IIT Kanpur Distinguished Alumnus Award in 2011, the ASME InterPack Achievement Award in 2011, the ITherm Achievement Award in 2012, and the ASME Heat Transfer Memorial Award in 2013.



Saibal Mukhopadhyay (S'97–M'06–SM'12) received the B.E. degree in electronics and telecommunication engineering from Jadavpur University, Calcutta, India, in 2000, and the Ph.D. degree in electrical and computer engineering from Purdue University, West Lafayette, IN, USA, in 2006.

He was a Research Staff Member with the IBM T. J. Watson Research Center, Yorktown Heights, NY, USA, from 2006 to 2007, where he was an Intern in 2003, 2004, and 2005. He joined the faculty of the Georgia Institute of Technology, Atlanta, GA, USA, in 2007. He has authored or co authored over 60 papers in reputed conferences and journals and holds four U.S. patents. His current research interests include technology-circuit co-design methodologies for low-power and variation-tolerant static random access memory in sub-65-nm silicon technologies.