A Silicon Interposer Platform Utilizing Microfluidic Cooling for High-Performance Computing Systems

Li Zheng, Yang Zhang, *Student Member, IEEE*, and Muhannad S. Bakir

***Abstract*— In this paper, a silicon interposer platform using microfluidic cooling is proposed for high-performance computing systems. The key advantage of the the silicon interposer is its very fine-pitch wiring, which enables high-bandwidth off-chip signaling for the chips assembled on the silicon interposer. Compared with conventional air cooling, embedded microfuidic cooling is used for better cooling and thermal isolation of chips on the silicon interposer. A test vehicle consisting of a silicon interposer and a silicon dice with microfluidic I/Os and embedded microfluidic heat sinks is fabricated and assembled for thermal measurements. At a flow rate of 50 mL/min, the measured temperature is 55.9 °C for a power density of 97 W/cm^2, which represents a normalized thermal resistance of 0.24 K · cm^2/W. The thermal simulations based on the measured thermal resistance show that 40.1% reduction in the silicon interposer temperature is achieved with microfluidic cooling compared to air cooling. Moreover, thermal coupling between the dice on the silicon interposer is significantly reduced with microfluidic cooling, which significantly benefits the integration density and the signaling performance by integrating chips more closely and reducing interconnect length.**

***Index Terms*— Bandwidth density, differential signaling, energy-per-bit, microfluidic cooling, silicon interposer, thermal resistance, thermal simulation.**

I. Introduction

SILICON interposer technology has been extensively explored in recent years due to its potential for high-bandwidth-density (BWD) signaling, form-factor reduction, and heterogeneous integration of logic, memory, microelectromechanical systems, and optoelectronics [1]–[3]. Dickson *et al.* [4] demonstrated the signaling of 10 Gb/s/channel using the silicon interposer interconnects of 2- to 6-μm wide and up to 6-cm long. With a large number of these signaling channels, very large off-chip BWD is achievable, which is critical to high-performance computing systems.

Cooling is a major challenge facing high-performance computing systems due to the ever increasing integration density. Microfluidic cooling, originally proposed in [4], has been demonstrated as a promising solution to large-power-density electronic systems. More recently, Zhang *et al.* [6] demonstrated the microfluidic cooling of a two-die stack operating at 100 W/cm^2 per tier with a maximum junction temperature of 47 °C using a staggered micropin-fin heat sink. Moreover, microfluidic I/Os consisting of solder-based fluidic microbumps and fluidic vias, which enable coolant delivery from a silicon interposer to an on-die microfluidic heat sink, have been demonstrated [7].

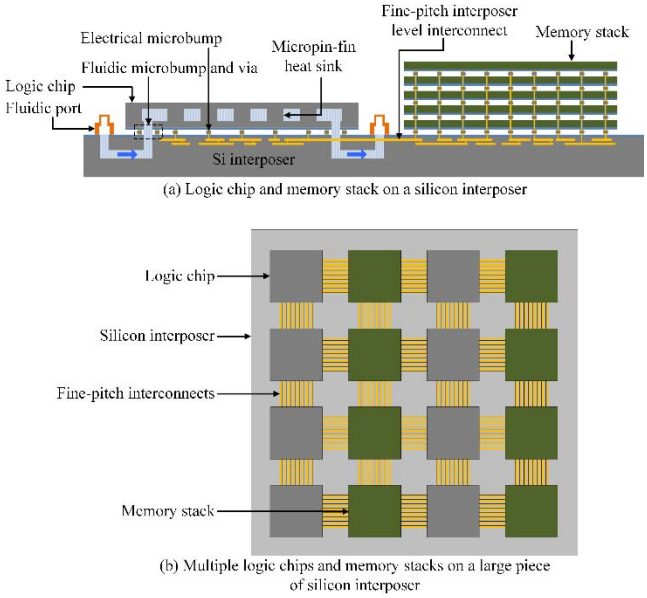

Fig. 1. High-performance computing system based on silicon interposer platform using microfluidic cooling. (a) Side view of a logic chip and a memory stack heterogeneously integrated on a silicon interposer. (b) Array of logic dice and memory stacks on a large silicon interposer with fine-pitch interconnects for high-bandwidth communication.

In this paper, we propose a silicon interposer platform using microfluidic cooling for high-performance computing systems; Fig. 1 illustrates our vision. A logic die with an embedded microfluidic heat sink and fluidic I/Os is assembled on a silicon interposer, as shown in Fig. 1(a), adjacent to a stack of memory dice. Coolant is pumped into the fluidic channels in the silicon interposer and distributed to the microfluidic heat sink through the fluidic microbumps and vias. Logic-to-memory high-bandwidth low-energy signaling is achieved using the short fine-pitch wires on the silicon interposer.

It is envisioned that a large array of such logic-memory pairs can be formed over a large silicon interposer, as shown in Fig. 1(b), perhaps using the silicon interposer bridging concept [8].

This paper is organized as follows. Section II presents the fabrication and the assembly of the test vehicle used to

Manuscript received September 26, 2014; revised March 25, 2015; accepted April 26, 2015. Date of publication September 15, 2015; date of current version October 2, 2015. This work was supported by the Defense Advanced Research Projects Agency (DARPA) through the DARPA ICECool Fundamentals and Applications under Contract W31P4Q-12-1-0014 and Contract HR0011-14-1-0002. Recommended for publication by Associate Editor R. Prasher upon evaluation of reviewers' comments.

The authors are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: lizheng@gatech.edu; steven.zhang@gatech.edu; muhannad.bakir@mirc.gatech.edu).

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

Digital Object Identifier 10.1109/TCPMT.2015.2470544

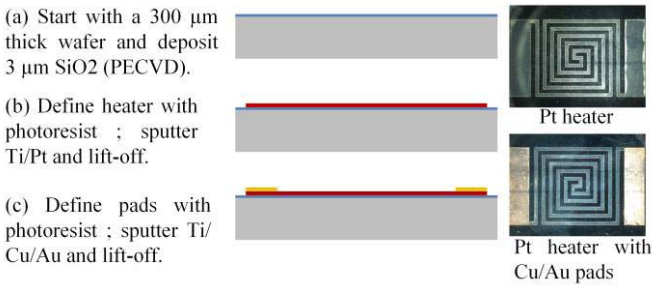

Fig. 2. Fabrication process for the Pt heater/RTD.

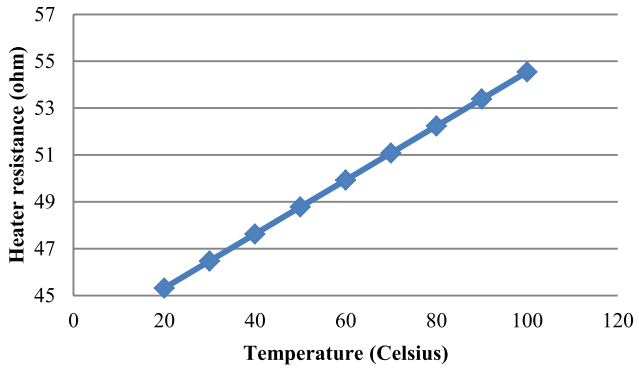

Fig. 3. Heater calibration.

demonstrate the vision, as shown in Fig. 1. In Section III, thermal measurements are reported for the fabricated test vehicle. The thermal benefits of the proposed platform using microfluidic cooling are also discussed. In Section IV, chip-to-chip electrical signaling benefits due to the lower interposer temperature and better thermal isolation compared to air cooling are evaluated. The conclusion is presented in Section V.

II. Test Vehicle and Thermal Measurements

A. Fabrication of Test Vehicle

The test vehicle consists of a pair of liquid-cooled silicon dice assembled on a silicon interposer with electrical and microfluidic interconnects. The number of electrical solder bumps (25-μm diameter) per die is 22 500. In order to facilitate thermal measurements, a thin-film platinum heater/resistance temperature detector (RTD) is integrated on the back-side of each of the assembled dice.

The fabrication process of the platinum heater is illustrated in Fig. 2. The process begins with depositing a 3-μm SiO_2 film on one side of a silicon wafer using plasma-enhanced chemical vapor deposition. Next, a Ti/Pt (25 nm/1 μm) film is sputtered on the SiO_2 layer followed by a liftoff process to form the heater. The last step is to sputter Ti/Cu/Au on the pads to enable electrical wire soldering to the electrical test equipment (power supply and data acquisition system). The area of the square heater is 0.5 cm^2.

Following fabrication, the heater is calibrated in an oven up to 100 °C. The calibration result indicates a good linear relationship between the temperature and the heater resistance, as shown in Fig. 3.

TABLE I
Resistance of the Electrical Microbumps (mΩ)

Microbump	Die #1	Die #2	Die #3
#1	11.6	12.8	13.8
#2	10.6	11.3	11.8
#3	12.9	13.2	12.6
Average	11.7	12.4	12.7

B. Assembly of Test Vehicle

SEM images of the electrical microbumps, staggered micropin fins, and fluidic microbumps and vias on the silicon die are shown in Fig. 4(a)–(c). The fabricated silicon dice were flip-chip bonded to a silicon interposer, which is also shown in Fig. 4.

The process used for the assembly is as follows: First, a pair of silicon dice with an integrated micropin-fin heat sink is sequentially bonded on a silicon interposer. Fig. 4(d) and (e) shows the X-ray images of the bonded dice. The resistance of the electrical microbumps was measured using a four-point measurement technique to quantify the bonding quality and the yield of the assembly process. The four-point resistance of the microbumps on the two assembled dice (Die #1 and Die #2) is listed in Table I; the results reflect the high yield of the assembly process.

For thermal measurements, another assembled die (Die #3) was capped with a silicon substrate containing the thin-film heater and the thermometer, as illustrated in Fig. 5; a thin thermal interface material (TIM) layer was applied between the interfacing surfaces to enhance the thermal contact. An epoxy film was applied to the edges for sealing. The last step is to attach the inlet/outlet port to the back side of the silicon interposer. The resistance of the electrical microbumps on Die #3 was also measured, as shown in Table I.

C. Thermal Measurements

Following assembly, microfluidic cooling experiments were conducted with the test vehicle. Deionized (DI) water at room temperature (~20 °C) was used as the coolant. The experimental setup is illustrated in Fig. 6. During the experiment, an adjustable digital gear pump drew the DI water from a reservoir. The DI water flowed through a mass flow meter and a polyester-based filter to remove particles (>20 μm) that could possibly clog the fluidic vias and the micropin-fin heat sink. A differential pressure gauge was used to measure the pressure at the input port. After flowing across the micropin-fin heat sink, the DI water exits the chip into another reservior. The temperature of the DI water was measured at both the inlet and output ports.

Once coolant flow commenced, the thin-film Pt heater was powered by an Agilent N6705B power analyzer to mimic the power dissipation of a functional die. The electrical resistance of the temperature sensor was recorded with an Agilent 34970A data logger. The electrical resistance values are used to calculate the heater/RTD temperature according to the calibration results shown in Fig. 3.

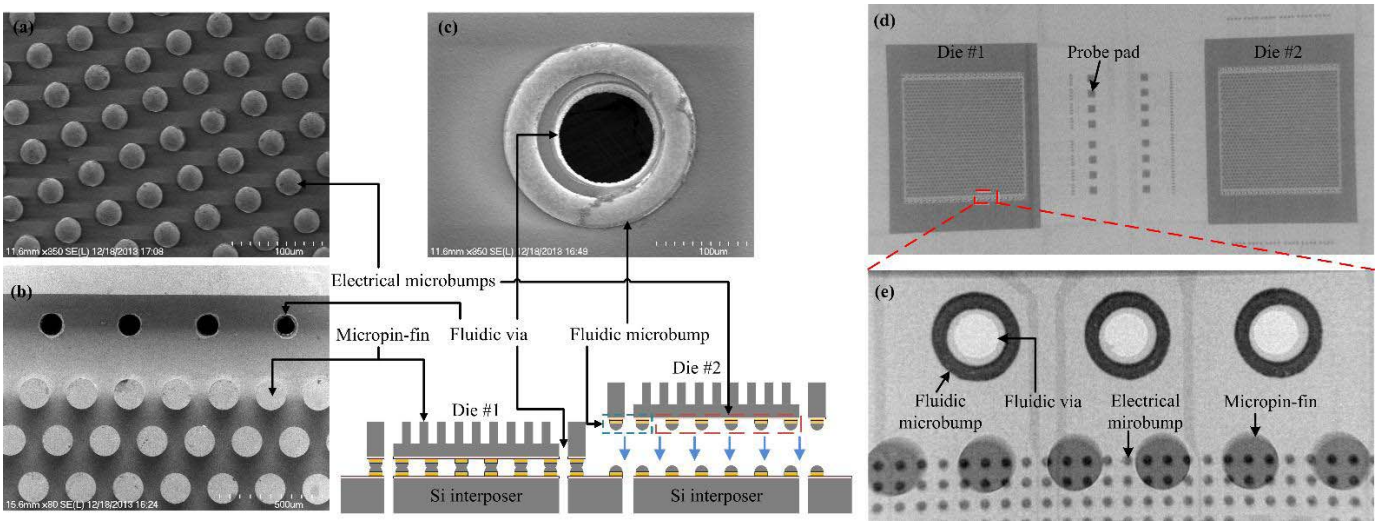

Fig. 4. SEM images of the fabricated silicon die and X-ray images of the bonded silicon dice and interposer. (a) Electrical microbumps. (b) Staggered micropin fins and fluidic vias. (c) Fluidic microbump and via. (d) Two bonded silicon dice on a silicon interposer. (e) Bonded fluidic and electrical microbumps.

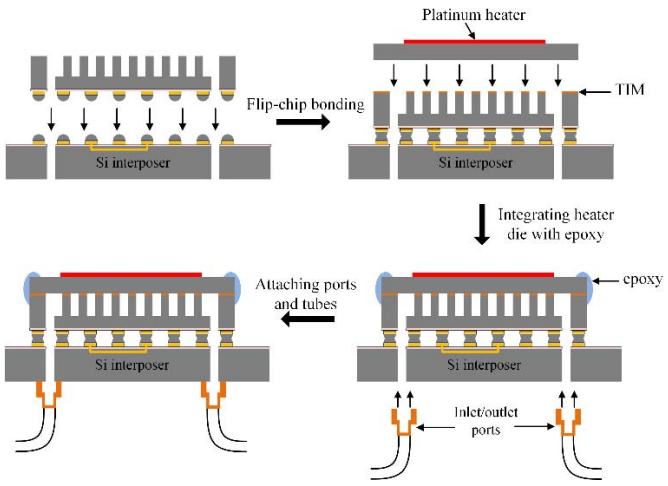

Fig. 5. Assembly process of the test vehicle for thermal measurement.

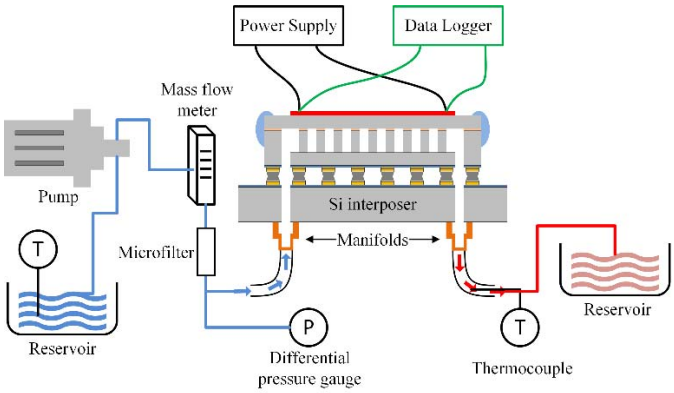

Fig. 6. Microfluidic cooling experiment setup.

III. Thermal Benefits of the Proposed Platform

A. Experimental Results

Flow rate is an important factor that affects the cooling performance of a micropin-fin heat sink as well as the pressure drop. Different flow rates, from 10 to 50 mL/min, were applied during the experiment. A power density of up to 100 W/cm^2 was applied to the heater. The inlet DI water was at room temperature (~20 °C).

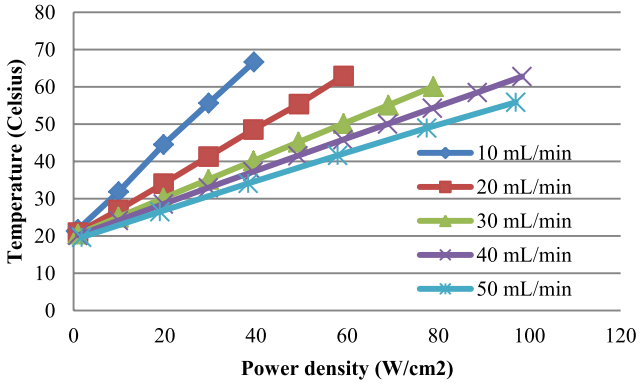

Fig. 7. Heater/RTD temperature versus power density for different flow rates.

The experimental results are shown in Fig. 7. As expected, the temperature of the heater increases linearly with the applied power density, and the temperature decreases as the flow rate increases for a given power density. The measured junction temperature is 55.9 °C at a power density of 97 W/cm^2 with a flow rate of 50 mL/min. Note that these results include the thermal resistance of the TIM layer between the capping layer and the micropin fins; the TIM was used to simplify the fabrication of the testbed.

B. Data Analysis Various

In this section, the thermal resistances of the test vehicle are calculated and analyzed to a first order using the measured temperatures. Equation (1) is used to calculate the total thermal resistance [5]

$$R_{\text{total}} = \frac{T_{\text{heater}} - T_{\text{inlet}}}{P} \tag{1}$$

where R_{total} is the total thermal resistance of the sample, T_{inlet} is the inlet DI water temperature (room temperature), T_{heater} is the heater temperature, and P is the power applied to the heater.

Based on the measured temperature, the summation of the conductive and convective thermal resistances can be calculated using [9]

$$R_{\text{cond}} + R_{\text{conv}} = \frac{T_{\text{heater}} - (T_{\text{inlet}} + T_{\text{outlet}})/2}{P} \tag{2}$$

where T_{outlet} is the DI water temperature at the outlet port, R_{cond} is the conductive thermal resistance, and R_{conv} is the convective thermal resistance.

Equations (3)–(5) are used to calculate the convective thermal resistance [10]

$$R_{\text{conv}} = \frac{1}{h \cdot A_t} \tag{3}$$

$$A_t = A_b + \eta \cdot A_{\text{fin}} \tag{4}$$

$$\eta = \frac{\tanh\left(2H_{\text{fin}}\sqrt{h/k_{\text{si}}D}\right)}{2H_{\text{fin}}\sqrt{h/k_{\text{si}}D}} \tag{5}$$

where h is the heat transfer coefficient, which is determined by the dimensions and the placement of the micropin fins and coolant velocity, A_t is the total effective heat transfer area, A_b is the base area exposed to coolant, η is the fin efficieny, A_{fin} is the total surface area of the micropin fins exposed to coolant, H_{fin} is the height of the micropin fins, D is the diameter of the micropin fins, and k_{si} is the thermal conductivity of silicon. Given that the dimensions and the layout of the micropin-fins heat sink and the DI water velocity are the same as in [10], the heat transfer coefficient h is approximated as the reported value (18 235 W/m^{2}K for a flow rate 40 mL/min in this experiment) from [10].

After obtaining R_{conv}, R_{cond} is simply calculated with (2). R_{cond} has three components: 1) $R_{\text{cond\_TIM}}$; 2) $R_{\text{cond\_si}}$; and 3) $R_{\text{cond\_sio}_2}$, which are related to the TIM layer, the bulk silicon (300-μm thick), and the silicon dioxide film (3-μm thick) beneath the Pt heater, respectively. $R_{\text{cond\_si}}$ and $R_{\text{cond\_sio}_2}$ are calculated using

$$R_{\text{cond\_si}} = \frac{T_{\text{si}}}{k_{\text{si}} \cdot A} \tag{6}$$

$$R_{\text{cond\_sio}_2} = \frac{T_{\text{sio}_2}}{k_{\text{sio}_2} \cdot A} \tag{7}$$

where T_{si} is the thickness of the bulk silicon, T_{sio_2} is the thickness of the silicon dioxide film, A is the total heating area, and k_{sio_2} is the thermal conductivity of silicon dioxide. After calculating $R_{\text{cond\_si}}$ and $R_{\text{cond\_sio}_2}$, $R_{\text{cond\_TIM}}$ can be obtained.

Table II lists the normalized thermal resistances for 1-cm^2 heating area, where R_{total} and $R_{\text{cond}} + R_{\text{conv}}$ are calculated from the measured temperature using (1) and (2). The remaining thermal resistances are derived.

In a functional die, the generated heat would be directly beneath the heat sink without a TIM layer. Thus, we adjust the thermal resistance by subtracting the thermal resistance of the TIM layer. Since the thermal resistance of the TIM layer is a part of R_{cond}, which is not a function of flow rate, it can be subtracted from the total thermal resistance for the different

TABLE II
THERMAL RESISTANCES (K · CM^2/W) AT 40 mL/min

Thermal resistance	Value	Derivation
R_{total}	0.43	Measured temperature, equation (1)
$R_{cond} + R_{conv}$	0.34	Measured temperature, equation (2)
R_{conv}	0.165	Heat transfer coefficient [10], equation (3)
R_{cond}	0.175	$(R_{cond} + R_{conv}) - R_{conv}$
R_{cond_si}	0.02	Equation (6)
R_{cond_sio2}	0.02	Equation (7)
R_{cond_TIM}	0.135	$R_{cond} - R_{cond_{si}} - R_{cond\_sio2}$
$R_{total} - R_{cond\_TIM}$	0.295	

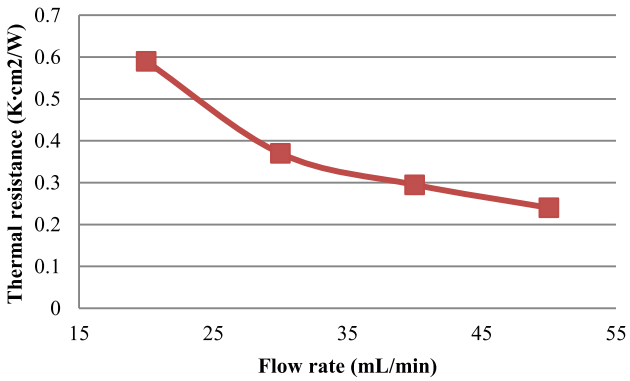

Fig. 8. Adjusted thermal resistance versus flow rate.

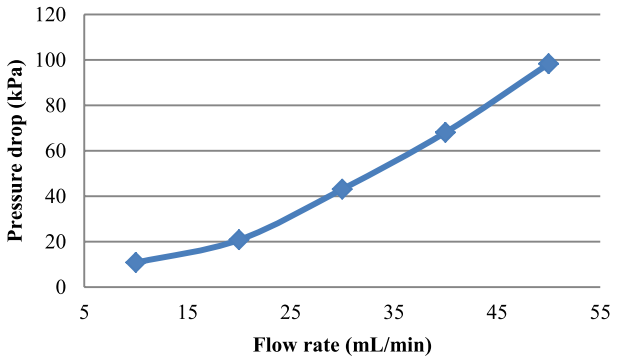

Fig. 9. Pressure drop versus flow rate.

flow rates. Fig. 8 shows the approximated thermal resistance (without TIM) for various flow rates. The adjusted thermal resistance is 0.24 K · cm^2/W at a flow rate of 50 mL/min.

Fig. 9 shows the measured pressure drop across the assembled die. As expected, the pressure drop increases with the increasing flow rate. At 50 mL/min, it reaches 98.3 kPa. One reason for this high pressure drop is the large opening area difference between the input port (1.25 cm × 1.25 cm) and fluidic vias. Optimizing the transition between the input port and the fluidic vias would help reduce the pressure drop.

C. Thermal Simulation

In this section, we use a finite-difference-method-based thermal simulator [11], [12] to compare air and microfluidic-cooled silicon-interposer-based systems, as shown in Fig. 10. The simulator uses nonconformal meshing, which enables us

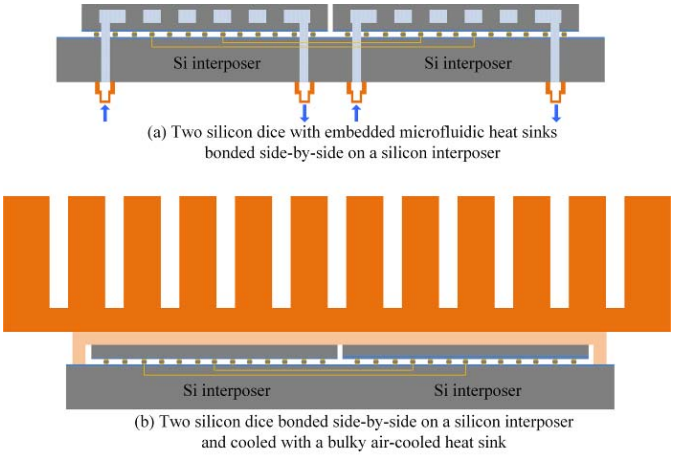

Fig. 10. Microfluidic cooling and air cooling for silicon-interposer-based systems.

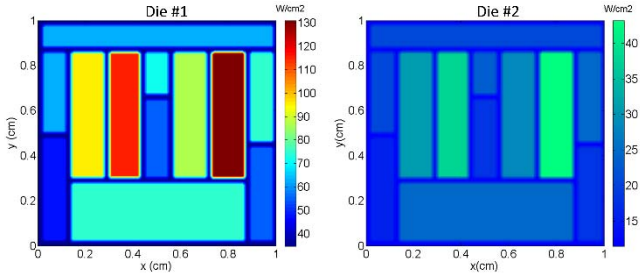

Fig. 11. Power maps of the two dice (74.63 W for die #1 and 24.88 W for die #2).

to evaluate large geometries that are consistent with typical interposer and heat spreader dimensions. The simulator is validated by ANSYS with an error less than 3%.

Fig. 10(a) shows two logic dice with an embedded microfluidic heat sink assembled side by side on a silicon interposer. In the air-cooled case, a bulky air-cooled heat sink and a heat spreader are attached on top of the two dice, as shown in Fig. 10(b). For the microfluidic-cooled case, a thermal resistance of 0.24 K · cm^2/W (from the reported experiments in Section III-B) was used in the thermal models. The air-cooled heat sink design and attributes are similar to that used for the Intel i7 microprocessor; the heat spreader is 5 cm × 4.5 cm, and the total thermal resistance from the heat spreader to the ambient is 0.218 K/W [13], which can be converted into a resistance of 0.8918 K · cm^2/W from the die surface to the ambient according to the spreading resistance model in [14].

The power maps of the two logic dice are based on the Intel i7 microprocessor, which have a maximum thermal design power of 84 W [15], [16], as shown in Fig. 11. In our thermal models, we use a silicon die of 1 cm × 1 cm. Since this area is smaller than the Intel i7 microprocessor, we scale the maximum power proportionally to 74.63 W. Referring to Fig. 11, we assume that the left die (die #1) is operating at maximum power, and the right die (die #2) is operating at one third of the maximum power. The size of the silicon interposer is 2 cm × 3 cm. There is a 1-mm gap between the edges of the two dice.

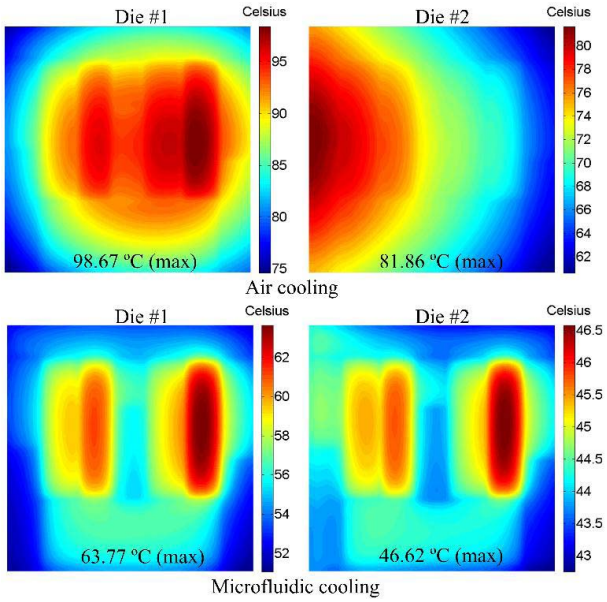

Fig. 12. Simulated temperature maps of dice with microfluidic cooling and air cooling.

The simulated temperature maps of the air- and microfluidic-cooled dice are shown in Fig. 12. It is obvious that microfluidic cooling significantly reduces the die temperature; the temperature reductions in Die #1 and Die #2 are approximately 34.90 °C (35.4%) and 35.24 °C (43.1%), respectively. Another observation is that with air cooling, the temperature of the high-power die is severely coupled to the low-power die due to thermal conduction within the heat spreader. In contrast, the temperature map of the low-power die using microfluidic cooling does not reflect thermal coupling. In addition to the lower junction temperature, microfluidic cooling provides better thermal isolation between the high-power and low-power dice.

Fig. 13 shows the power maps of the silicon interposers using air and microfluidic cooling. There is approximately a 33.4 °C (40.1%) reduction in temperature (average temperature of the space between the two dice), and a better thermal isolation is achieved with microfluidic cooling.

In order to reduce the temperature coupling of the two dice in the air cooling configuration, we could increase the space between the high-power and low-power dice (albeit at the cost of reduced system integration and increased interconnect lengths, which is discussed in Section IV). Fig. 14 shows the temperatures of the two dice (maximum temperature) and the silicon interposer (average temperature of the space between the two dice) as a function of spacing between the two dice.

Since the die size and the interposer size are assumed to be 1 cm × 1 cm and 2 × 3 cm, respectively, the maximum space between the two dice is 10 mm. Although 10-mm space is not practical and opposes the purpose of high-density integration on silicon interposer, we swept the spacing distance from 0.1 to 10 mm for wide range exploration. As shown in Fig. 14, for air cooling, there is ∼27.8 °C (25.4%), 6.3 °C (6.4%), and 13.1 °C (15.3%) reduction in temperature for the low-power die, high-power die, and silicon interposer,

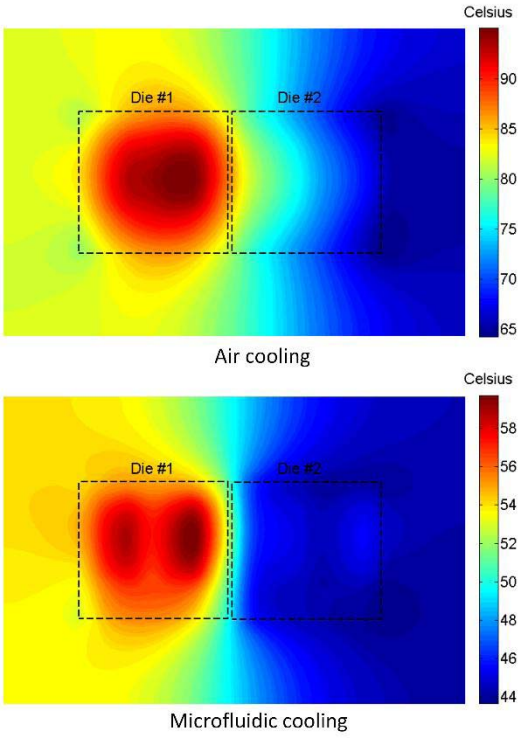

Fig. 13. Simulated temperature maps of silicon interposers with microfluidic cooling and air cooling.

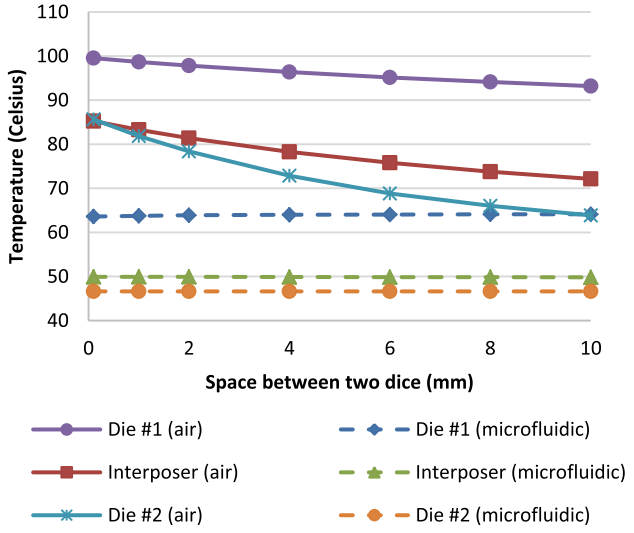

Fig. 14. Temperature of silicon dice and interposers as a function of the space between the two dice.

respectively, when the space increases from 0.1 to 10 mm. For microfluidic cooling, separating the dice does not impact the temperature. Again, although increasing the space helps reduce temperature coupling in the air cooling configuration, it significantly reduces the integration density (less number of dice on an interposer) and the signaling performance (discussed in Section IV).

IV. Signaling Benefits of the Proposed Platform

In this section, the fine-pitch interconnects on the proposed silicon interposer are modeled, and the signaling benefits of the proposed platform are analyzed.

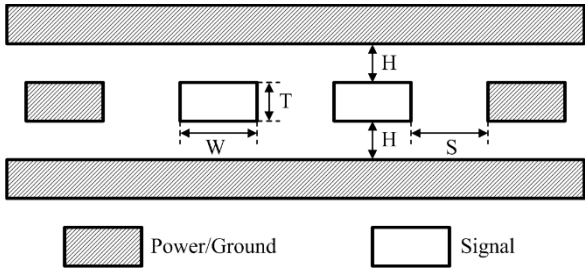

Fig. 15. Stripline differential signaling on silicon interposer.

A. Signaling Model

In this paper, low-swing current mode, bipolar, and unidirectional differential signaling scheme is chosen for high bandwidth and good noise immunity signaling. We follow the modeling methodology presented in [17] and [18]. The stripline structure illustrated in Fig. 15 is used.

The two major parameters are channel attenuation and noise, as shown in

$$I_{\min} \geq \frac{V_{\text{margin}}}{\text{A} \cdot Z_{\text{in}}} \tag{8}$$

where $I_{\min}$ is the minimum required current swing, V_{margin} is the noise margin, A is the attenuation of the channel, and Z_{in} is the interconnect impedance. Basically, (8) requires the signal magnitude after attenuation to dominate the noise.

Resistance, inductance, capacitance, and conductance of the interposer interconnects are extracted to calculate the attenuation at a given data rate (or frequency) [19]. Resistance is calculated using the frequency-dependent model developed for the interposer interconnects [20]. Capacitance and conductance are calculated using the models from [21]. Inductance is derived from its relationship with the propagation velocity and the capacitance.

The noise margin is the sum of the noise margin required at a given bit error rate and other fixed noise sources, such as receiver offset and sensitivity [17]. Channel crosstalk is neglected because the signaling channels are well shielded by ground wires and planes, as shown in Fig. 15.

After determining $I_{\min}$, the power consumption of the signaling channel is calculated as follows:

$$P = V_{\text{dd}} \cdot I_{\min} \tag{9}$$

where P is the signaling channel power and V_{dd} is the power supply voltage [18].

B. Silicon Interposer Interconnect Analysis

Fine-pitch wiring is a key advantage of the silicon interposer technology. Using a larger number of interconnects, very large signaling bandwidth can be achieved. However, reducing the pitch of the interconnects (and thus, interconnect width and height) increases the interconnect resistance, which leads to higher energy-per-bit (EPB). Thus, there is a tradeoff between signaling bandwidth and EPB. To capture this tradeoff, we use a composite metric BWD/EPB as defined in [18], [20]. BWD is defined as the aggregate bandwidth within a unit width.

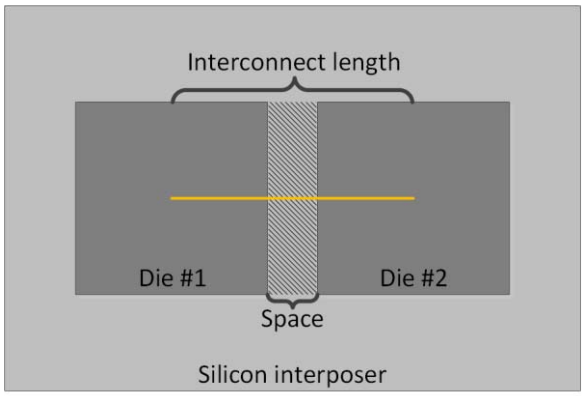

Fig. 16. Center-to-center distance for interconnect length.

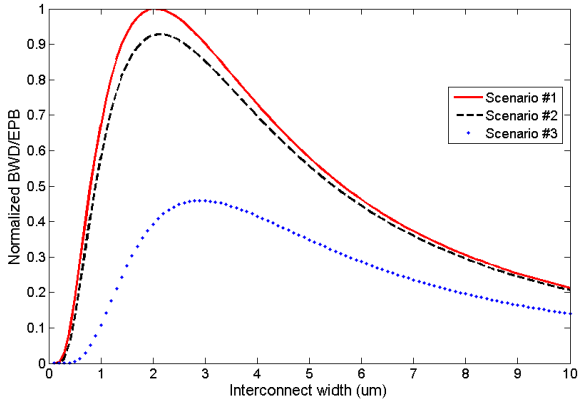

Fig. 17. Normalized BWD/EPB as a function of interconnect width for the three scenarios.

TABLE III
THREE SCENARIOS FOR SIGNALING ANALYSIS

	Scenario #1	Scenario #2	Scenario #3
Interconnect length (cm)	1.1 cm	1.1 cm	2.0 cm
Spacing (mm)	1 mm	1 mm	10 mm
Temperature (Celsius)	49.9	83.3	72.1
Cooling	Microfluidic cooling	Air cooling	Air cooling
Copper resistivity	$1.88 \times 10^{-8}\,\Omega$·m	$2.09 \times 10^{-8}\,\Omega$·m	$2.02 \times 10^{-8}\,\Omega$·m

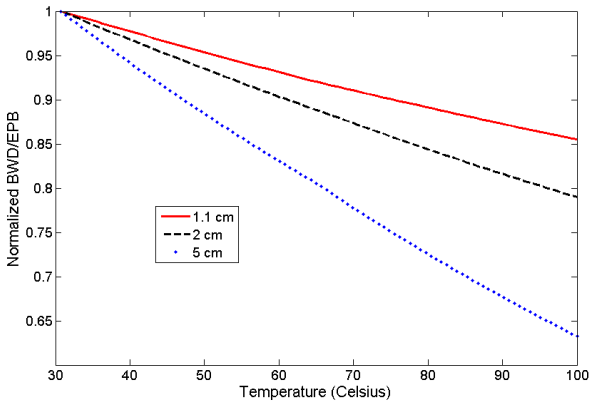

Fig. 18. Normalized BWD/EPB as a function of temperature for three interconnects of different lengths.

We assume interconnect thickness, T, and dielectric layer thickness, H, are constant ($T = H = 2\ \mu$m); space S is two-third of the interconnect width W, as shown in Fig. 14. We also assume a signaling frequency of 5 GHz (or data rate 10 Gb/s). The impact of interconnect width, length, and operating temperature, which impacts interconnect resistance, on BWD/EPB is investigated.

The interconnect temperature is assumed to be the average temperature of the space between the two dice (shaded region in Fig. 16), as shown in Fig. 14. The temperature impacts signaling through resistivity, as indicated by

$$\rho(T) = \rho_0[1 + \alpha(T - T_0)] \tag{10}$$

where ρ is resistivity, ρ_0 is resistivity at temperature T_0, where T_0 is reference temperature, and α is temperature coefficient of resistivity.

For a given interconnect length and temperature, there is an optimal width that maximizes BWD/EPB, since the increasing width improves EPB but reduces BWD, and vice versa. In the following analysis, we use the die center-to-center distance as interconnect length, as shown in Fig. 16. Table III lists the three scenarios for analysis.

Fig. 17 shows the normalized BWD/EPB for the three scenarios. For scenario #1 and scenario #2, BWD/EPB is maximized with an optimal interconnect width of 2 and 2.1 μm, respectively. There is approximately a modest 7.76 % improvement in BWD/EPB for microfluidic cooling (scenario #1) compared with air cooling (scenario #2). In scenario #3, the interconnect length is increased to 2 cm for temperature and thermal coupling reduction. As noted in Section III-C, a 10 mm space is not practical and does not enable the high-density integration on silicon interposer. We use this scenario for the comparison and the exploration of the impact of interconnect length on signaling performance. As shown in Fig. 17, the optimal width for the 2-cm-long interconnect increases to 2.9 μm; however, the BWD/EPB (at optimal width) becomes less than half of that of scenario #1 and scenario #2.

Fig. 18 shows the normalized BWD/EPB for interconnects of 1.1-, 2-, and 5-cm length as a function of temperature; the results are normalized to each interconnects' BWD/EPB at 30 °C. We can see that the BWD/EPB decreases rapidly for longer interconnects. When interposer temperature increases from 30 °C to 80 °C, there is ∼11.7%, 16.33%, and 28.5% reduction in BWD/EPB for interconnects of 1.1-, 2-, and 5-cm length, repectively. Thus, for large silicon interposers with long interconnect length, microfluidic cooling could potentially provide some benefits for chip-to-chip singling.

V. CONCLUSION

A silicon interposer platform using microfluidic cooling is proposed for high-performance computing systems. The thermal benefits of the proposed platform are analyzed based on microfluidic cooling experimentation and thermal simulation.

Low temperature and better thermal isolation may be achieved with microfluidic cooling. Moreover, the platform provides means for the improvement in signaling principally due to tighter integration (shorter interconnects between dice) while maintaining lower temperatures.

ACKNOWLEDGMENT

The authors would like to thank the cleanroom staff at Georgia Tech for their help with the fabrication processes.

REFERENCES

[1] J. U. Knickerbocker *et al.*, "3D silicon integration," in *Proc. 58th IEEE Electron. Compon. Technol. Conf.*, Lake Buena Vista, FL, USA, May 2008, pp. 538–543.
[2] M. S. Parekh, P. A. Thadesar, and M. S. Bakir, "Electrical, optical and fluidic through-silicon vias for silicon interposer applications," in *Proc. 61st IEEE Electron. Compon. Technol. Conf.*, Lake Buena Vista, FL, USA, May/Jun. 2011, pp. 1992–1998.
[3] R. Ho *et al.*, "Silicon photonic interconnects for large-scale computer systems," *IEEE Micro*, vol. 33, no. 1, pp. 68–78, Jan./Feb. 2013.
[4] T. O. Dickson *et al.*, "An 8×10-Gb/s source-synchronous I/O system based on high-density silicon carrier interconnects," *IEEE J. Solid-State Circuits*, vol. 47, no. 4, pp. 884–896, Apr. 2012.
[5] D. B. Tuckerman and R. F. W. Pease, "High-performance heat sinking for VLSI," *IEEE Electron Device Lett.*, vol. 2, no. 5, pp. 126–129, May 1981.
[6] Y. Zhang, A. Dembla, Y. Joshi, and M. S. Bakir, "3D stacked microfluidic cooling for high-performance 3D ICs," in *Proc. IEEE 62nd Electron. Compon. Technol. Conf.*, May/Jun. 2012, pp. 1644–1650.
[7] L. Zheng, Y. Zhang, G. Huang, and M. S. Bakir, "Novel electrical and fluidic microbumps for silicon interposer and 3-D ICs," *IEEE Trans. Compon., Packag., Manuf. Technol.*, vol. 4, no. 5, pp. 777–785, May 2014.
[8] H. S. Yang, C. Zhang, M. Zia, L. Zheng, and M. S. Bakir, "Interposer-to-interposer electrical and silicon photonic interconnection platform using silicon bridge," in *Proc. IEEE Opt. Interconnects Conf.*, Coronado, CA, USA, May 2014, pp. 71–72.
[9] T. Brunschwiler *et al.*, "Interlayer cooling potential in vertically integrated packages," *Microsyst. Technol.*, vol. 15, no. 1, pp. 57–74, Aug. 2008.
[10] Y. Zhang, L. Zheng, and M. S. Bakir, "3-D stacked tier-specific microfluidic cooling for heterogeneous 3-D ICs," *IEEE Trans. Compon., Packag., Manuf. Technol.*, vol. 3, no. 11, pp. 1811–1819, Nov. 2013.
[11] T. E. Sarvey, Y. Zhang, Y. Zhang, H. Oh, and M. S. Bakir, "Thermal and electrical effects of staggered micropin-fin dimensions for cooling of 3D microsystems," in *Proc. IEEE Intersoc. Conf. Thermal Thermomech. Phenomena Electron. Syst. (ITherm)*, May 2014, pp. 205–212.
[12] J. Xie and M. Swaminathan, "Fast electrical-thermal co-simulation using multigrid method for 3D integration," in *Proc. IEEE 62nd Electron. Compon. Technol. Conf.*, May/Jun. 2012, pp. 651–657.
[13] Intel. *Intel Core i7 Processor Families for the LGA2011-0 Socket: Thermal Mechanical Specification and Design Guide*. [Online]. Available: http://www.intel.com/content/dam/www/public/us/en/documents/design-guides/core-i7-lga-2011-guide.pdf, accessed Sep. 2015.
[14] S. Lee, S. Song, V. Au, and K. Moran, "Constricting/spreading resistance model for electronics packaging," in *Proc. ASME/JSME Thermal Eng. Conf.*, Mar. 1995, pp. 199–206.
[15] Intel. *3rd Generation Intel Core Processor Family Quad Core Launch Product Information*. [Online]. Available: http://download.intel.com/newsroom/kits/core/3rdgen/pdfs/3rd_Generation_Intel_Core_Product_Information.pdf, accessed Sep. 2015.
[16] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi, "McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Proc. 42nd Annu. IEEE/ACM Int. Symp. Microarchitecture*, Dec. 2009, pp. 469–480.
[17] H. Cho, P. Kapur, and K. C. Saraswat, "Power comparison between high-speed electrical and optical interconnects for interchip communication," *J. Lightw. Technol.*, vol. 22, no. 9, pp. 2021–2033, Sep. 2004.
[18] V. Kumar *et al.*, "Airgap interconnects: Modeling, optimization, and benchmarking for backplane, PCB, and interposer applications," *IEEE Trans. Compon., Packag., Manuf. Technol.*, vol. 4, no. 8, pp. 1335–1346, Aug. 2014.
[19] D. M. Pozar, *Microwave Engineering*. New York, NY, USA: Wiley, 2011.
[20] V. Kumar, L. Zheng, M. S. Bakir, and A. Naeemi, "Compact modeling and optimization of fine-pitch interconnects for silicon interposers," in *Proc. IEEE Int. Interconnect Technol. Conf. (IITC)*, Jun. 2013, pp. 1–3.
[21] R. Venkatesan, "Multilevel interconnect architectures for gigascale integration (GSI)," Ph.D. dissertation, Dept. Elect. Comput. Eng., Georgia Inst. Technol., Atlanta, GA, USA, 2003.

Li Zheng received the B.S. degree from Zhejiang University, Hangzhou, China, in 2006, and the M.S. degrees in electrical and computer engineering from Shanghai Jiao Tong University, Shanghai, China, and the Georgia Institute of Technology, Atlanta, GA, USA, in 2009, where he is currently pursuing the Ph.D. degree in electrical and computer engineering.

His current research interests include microfluidic cooling, off-chip signaling, and power delivery for high-performance silicon interposer and 3-D integrated systems.

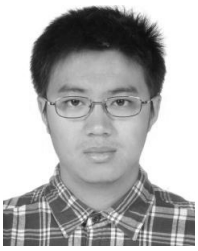

Yang Zhang (S'13) received the B.S. degree in microelectronics and mathematics from Peking University, Beijing, China, in 2012. He is currently pursuing the Ph.D. degree in electrical engineering with the Georgia Institute of Technology, Atlanta, GA, USA.

Muhannad S. Bakir received the B.E.E. degree from Auburn University, Auburn, AL, USA, in 1999, and the M.S. and Ph.D. degrees in electrical and computer engineering from the Georgia Institute of Technology (Georgia Tech), Atlanta, GA, USA, in 2000 and 2003, respectively.

He is currently an Associate Professor and the ON Semiconductor Junior Professor with the School of Electrical and Computer Engineering, Georgia Tech. His current research interests include 3-D electronic system integration, advanced cooling and power delivery for 3-D systems, biosensors and their integration with CMOS circuitry, and nanofabrication technology.

Dr. Bakir is a member of the International Technology Roadmap for Semiconductors Technical Working Group for Assembly and Packaging. He and his research group have received 14 conference and student paper awards, including five from the IEEE Electronic Components and Technology Conference, four from the IEEE International Interconnect Technology Conference, and one from the IEEE Custom Integrated Circuits Conference. He was a recipient of the 2011 IEEE CPMT Society Outstanding Young Engineer Award, the 2012 DARPA Young Faculty Award, and the 2013 Intel Early Career Faculty Honor Award, and was an Invited Participant in the 2012 National Academy of Engineering Frontiers of Engineering Symposium. His group received the 2014 Best Paper Award of the IEEE TRANSACTIONS ON COMPONENTS PACKAGING AND MANUFACTURING TECHNOLOGY Award in the area of advanced packaging. He is an Editor of the IEEE TRANSACTIONS ON ELECTRON DEVICES and an Associate Editor of the IEEE TRANSACTIONS ON COMPONENTS, PACKAGING AND MANUFACTURING TECHNOLOGY. In 2015, he was elected by the IEEE CPMT Society to serve as a Distinguished Lecturer for a four-year term.