Embedded Cooling Technologies For Densely Integrated Electronic Systems

Thomas E. Sarvey[1], Yang Zhang[1], Li Zheng[1], Paragkumar Thadesar[1],
Ravi Gutala[2], Colman Cheung[2], Arifur Rahman[2], and Muhannad S. Bakir[1]

[1]School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30318
[2]Altera Corporation, San Jose, CA 95134

***Abstract*—In modern integrated systems, interconnect and thermal management technologies have become two major limitations to system performance. In this paper, a number of technologies are presented to address these challenges. First, low-loss polymer-embedded vias are demonstrated in thick wafers compatible with microfluidics. Next, fluidic I/Os for delivery of fluid to microfluidic heat sinks are demonstrated in assembled 2.5D and 3D stacks. Then thermal coupling between dice in 2.5D and 3D systems is explored. Lastly, the utility of microfluidic cooling is demonstrated through an FPGA, built in a 28nm process, with a monolithically integrated microfluidic heat sink.**

I. INTRODUCTION

Modern computing systems are facing a number of issues that prohibit the types of performance increases that have previously been achieved through scaling of features, clock frequencies, and voltage rails. Two major barriers to continued performance improvement in current systems are system interconnects and heat dissipation techniques. In this paper, a number of technologies are presented to solve these challenges.

Fig. 1 shows maximum thermal design power (TDP), clock frequency, and number of cores for Intel desktop processors over the last 16 years [1]. A major shift can be noted around 2005, when TDP and clock frequency leveled off and the number of cores began to increase. This shift was largely brought on by the difficulty of dissipating more than 150 W from the CPU package with traditional cooling methods. This paradigm shift from single to multi-core processors allowed processor throughput to continue scaling with relatively constant TDP, albeit at the expense of single thread performance.

In fact, total processor throughput has continued to increase at a rate much larger than the rate at which memory bandwidth improves. This growing discrepancy between what the processor demands and what the memory can provide is known as the memory wall. The gap has largely been filled through consistently increasing cache sizes. By improving memory bandwidth, the optimal cache size can be shrunk, freeing die space for additional computing cores and increasing processor throughput [2].

Bandwidth, latency, and I/O power can be improved by decreasing interconnect lengths and increasing interconnect density. This can be accomplished through closer geometric integration of devices in the form of 2.5D or 3D systems.

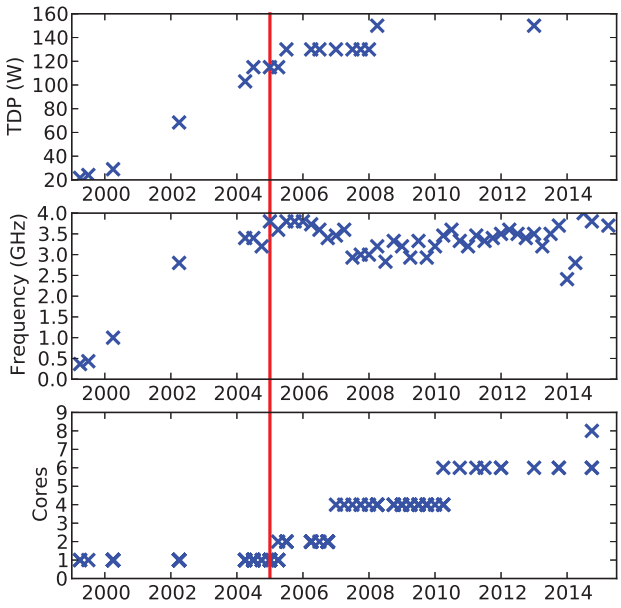

Fig. 1. Maximum thermal design power (TDP), clock frequency, and number of cores for Intel desktop CPUs

In 2.5D systems dice are mounted physically close to one another on an interposer with short, high density interconnects providing high bandwidth and low I/O power between devices on the interposer. Silicon interposers have been demonstrated with Field Programmable Gate Arrays (FPGAs) [3], [4]. Through the use of heterogeneous integration of FPGA slices and 28G transceivers, power was reduced by 50% and twice the capacity of what would have been achievable through scaling alone when going from the 40 nm to 28 nm node was gained.

Stacked 3D systems can improve integration density even further by stacking thin silicon dice on top of one another. Through silicon vias (TSVs) are a key enabler for stacked systems and can be used for inter-tier communication and power delivery.

However, while alleviating the interconnect problem, these

978-1-4799-8682-8/15/$31.00 ©2015 IEEE

two systems only worsen the thermal problem. Although these systems improve efficiency and therefore lower overall system power, they increase system power density. This dense integration leads to two major thermal issues: (1) higher heat fluxes that cannot be dissipated by traditional air cooled heat sinks and (2) thermal crosstalk between adjacent or stacked dice, which were once insulated from one another through their large physical distance, but now are in close physical, and thermal, proximity.

With respect to the problem of high heat fluxes, microfluidic cooling, which can provide far lower junction-to-ambient thermal resistances than traditional cooling techniques, can be used. The very small volumetric footprint of microfluidic heat sinks also makes them ideal for integration into high density computing systems. Three basic fashions in which microfluidic cooling can be utilized are shown in Fig. 2.

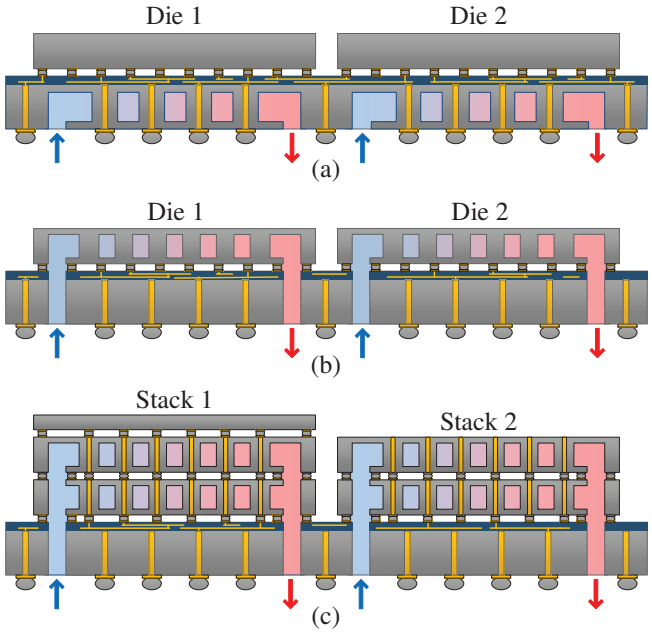

Fig. 2. Microfluidic cooling integrated in (a) the interposer, (b) the backside of the die, and (c) the backside of multiple dice in a 3D stack

For ICs with moderately high heat fluxes, interposer level cooling, show in Fig. 2(a), provides an approach with heat removal superior to an air cooled heat sink, without modification to the mounted ICs. In systems with higher heat fluxes, the monolithically integrated configuration shown in Fig. 2(b) may be preferred. By integrating the heat sinks into the backside of the dice, as close to the source of heat as possible, junction-to-ambient thermal resistance can be minimized. Lastly, in order to stack multiple high-power dice, each high-power tier in the 3D stack may need its own backside microfluidic heat sink, as shown in Fig. 2(c). By integrating a separate heat sink into each high-power tier, the number of high-power tiers can be scaled indefinitely with respect to thermal dissipation. This is in contrast to an air cooled heat sink mounted on the top of the 3D stack, where heat flux which must be dissipated by the heat sink is additive. In addition, with a top-mounted air cooled heat sink, the heat path from bottom dice to the heat sink includes all intermediate tiers. This makes microfluidic cooling a potential key enabling technology to 3D stacking of multiple high-power tiers.

With respect to the crosstalk problem, dedicated microfluidic heat sinks, as shown in Fig. 2(b) can provide significant decoupling in 2.5D systems. In 3D systems, the addition of a thermally resistive air gap layer between dice can be used for thermal decoupling (to be discussed in Section IV).

This paper is organized as follows: Section II describes polymer-embedded vias, a low loss TSV technology compatible with interposers containing microfluidics. In section III, fluidic I/Os are presented to address the issue of bringing fluid into and out of dice with microfluidic cooling. In section IV, the thermal cross talk problem is modeled in 2.5D and 3D systems and innovative solutions are proposed for each system. Lastly, in Section V, a monolithically integrated microfluidic heat sink is demonstrated and tested in a functional CMOS circuit for the first time. The microfluidic heat sink is integrated into the backside of a functional FPGA mounted to its native package substrate and development board.

II. Polymer-embedded Vias

Microfluidics can be integrated into the interposer in two ways. First, a microfluidic heat sink can be built directly into the interposer, as shown in Fig. 2(a). Moreover, the interposer is also useful for routing fluids laterally to monolithically integrated heat sinks, as shown in Fig. 3. In any of these configurations, the interposer may be several hundred micrometers thick. However, thick interposers result in longer TSVs with increased losses [5]. Glass and high-resistivity silicon have been explored in the literature to attain interposers with low-loss vias [6], [7]. However, the fabrication of vias in glass is challenging and high-resistivity silicon is relatively expensive.

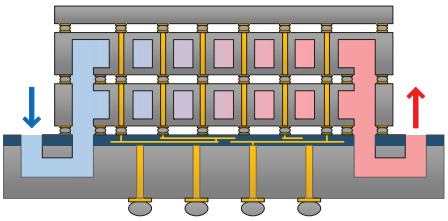

Fig. 3. Fluid can be routed through an interposer and delivered to monolithically integrated microfluidic heat sinks integrated into the backsides of dice.

To attain low-loss TSVs compatible with microfluidic cooling without complicated fabrication techniques, this paper demonstrates polymer-embedded vias with copper vias embedded within photodefined polymer wells in silicon [8]. The fabrication process for polymer-embedded vias is shown in Fig. 4. The process begins with etching of wells in a silicon wafer with a copper seed layer at the base and an SiO_2 etch stop layer between the silicon and copper. The SiO_2 is next

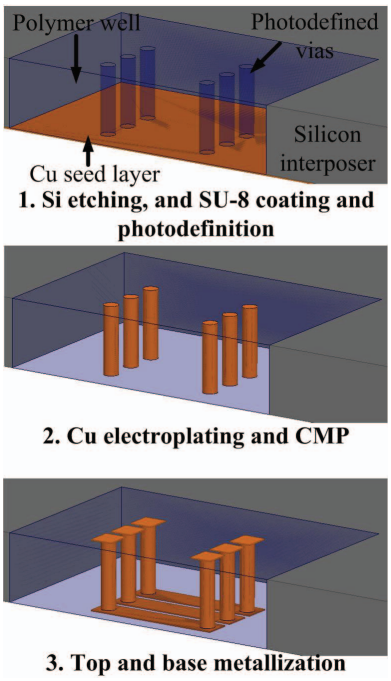

Fig. 4. Fabrication process for polymer-embedded vias.

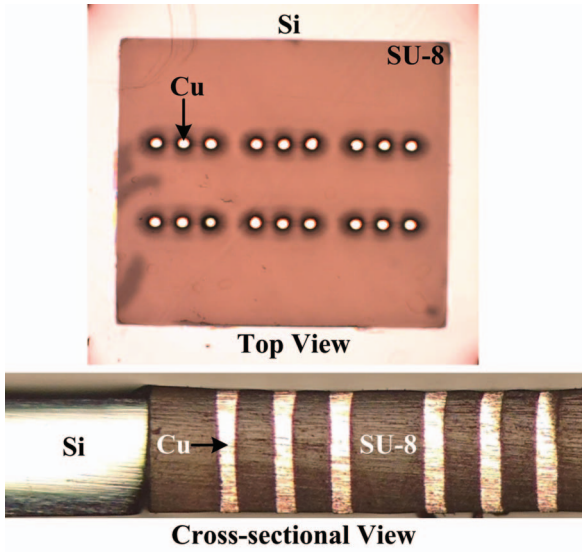

Fig. 5. Fabricated polymer-embedded vias.

etched followed by SU-8 coating and photodefinition yielding vias in polymer wells. Once the vias are formed, copper electroplating and chemical-mechanical polishing (CMP) are performed to attain copper vias in the polymer wells. Next, top and base metallization are fabricated to yield two-port RF measurement structures.

Fig. 5 illustrates the fabricated 370 µm tall polymer-embedded vias with 65 µm diameter and 150 µm pitch within a 1.8 mm x 1.8 mm well in silicon [8].

Using the fabricated two-port measurement structures with top and base metallization, RF measurements were performed

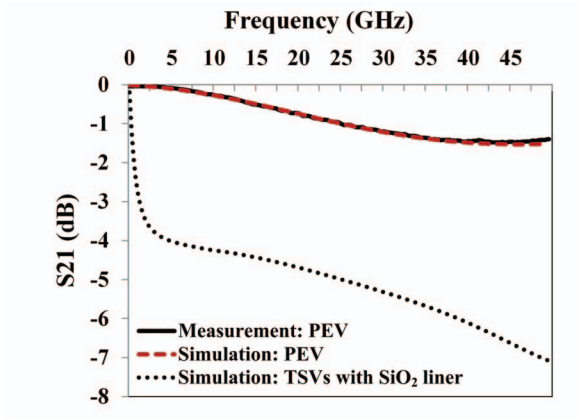

Fig. 6. Polymer-embedded via (PEV) measurement.

up to 50 GHz with an RF probe station, an Agilent N5245A PNA-X network analyzer and Cascade |Z| probes. The fabricated polymer-embedded vias yield a 78% reduction in insertion loss compared to similar-sized TSVs with 1 µm thick SiO_2 liner and 10 Ω–cm silicon resistivity at 50 GHz, as shown in Fig. 6 [8].

III. 2.5D and 3D Assembly with Fluid Delivery

Fluidic interconnections, which enable coolant delivery to embedded microfluidic heat sinks from package substrates or lower level dice, are essential to microfluidic cooled ICs. In this section, solder based annular shaped fluidic microbumps are presented as a novel fluidic I/O technology [9]. Fig. 7(a) shows the SEM images of the fabricated fluidic microbump (150 µm inner diameter and 210 µm outer diameter) and via (100 µm diameter), electrical microbumps (25 µm diameter and 50 µm pitch), and micropin-fin heat sink on a silicon die. Following fabrication, the silicon die is flip-chip bonded onto a silicon interposer [9] and the fluidic interconnections between the silicon die and interposer are formed by the fluidic microbumps, as shown in Fig. 7(a). The fluidic microbumps and electrical microbumps are very well aligned, as shown in X-ray image in Fig. 7(b).

Electrical and fluidic testing were conducted to verify the bonding. The resistance of the electrical microbumps was measured using the four point measurement technique. Fig. 8 lists the measurement results of three bonded samples. The average resistance of a single electrical microbump was 12.3 mΩ. Following resistance measurement, the micropin-fin heat sink was capped with a glass slide and inlet/outlet ports were attached to the backside of the interposer for fluidic testing. During the fluidic testing, deionized (DI) water was pumped into the heat sink continuously for 4 hours at two flow rates of 30 mL/min and 50 mL/min (each flow rate for 2 hours), and the pressure drop between the inlet and outlet ports was recorded. At 50 mL/min, the measured pressure drop reached 100 kPa. More importantly, the measured pressure drop was very stable during the testing and no leakage was observed, which indicates good fluidic interconnection.

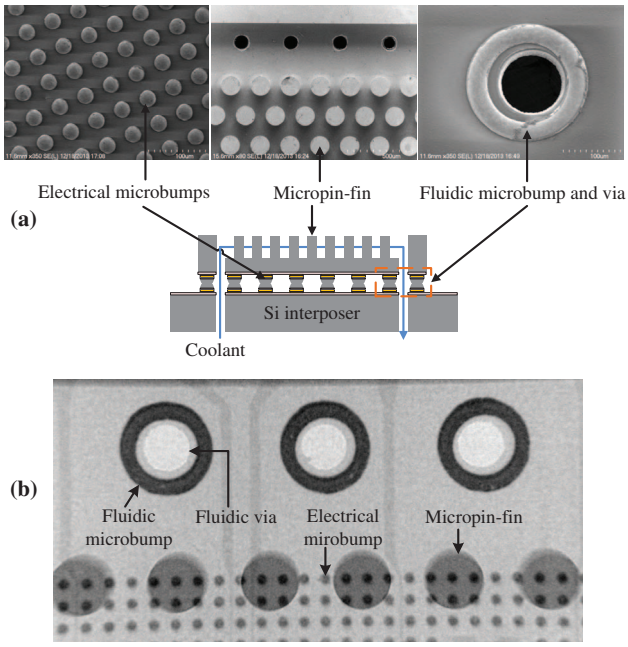

Fig. 7. a) SEM images of electrical microbumps, fluidic microbumps and vias, and micropin-fin heat sink on a silicon die; (b) X-ray images of the flip-chip bonded electrical microbumps, fluidic microbumps, and micropin-fins

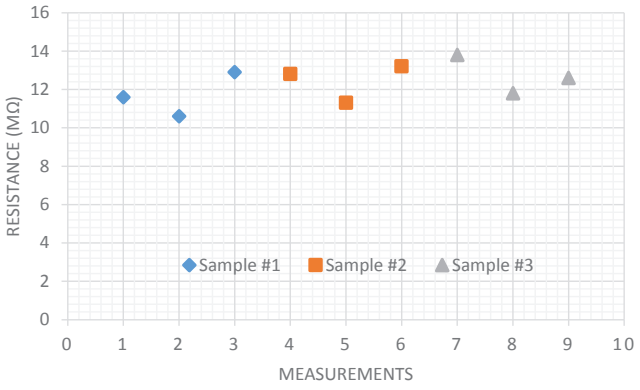

Fig. 8. Resistance of electrical microbumps measured on three bonded samples

Besides the structure in Fig. 7(a), the fluidic microbumps can also be utilized for other 2.5D and 3D microfluidic cooled structures. Fig. 9(a) illustrates the fabricated 2.5D sample with embedded micropin-fin heat sink and fluidic I/Os [8]. Two silicon dice with fluidic I/Os and heat sinks are assembled side-by-side on a silicon interposer. Moreover, a 3D structure with fluidic I/Os is demonstrated [10], as shown in Fig. 9(b). The left X-ray image shows the two silicon dice stacked on a silicon interposer; the right X-ray image shows the fluidic I/Os on the two dice at a tilted angle.

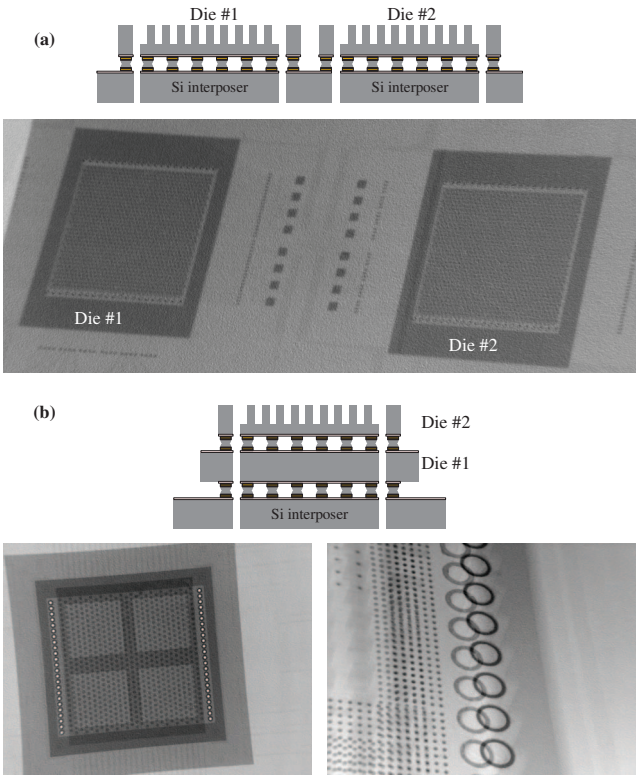

Fig. 9. Diagrams and X-ray images of (a) two silicon dice with microfluidic I/Os and embedded micropin-fin heat sink (2.5D) and (b) two silicon dice with microfluidic I/Os and embedded micropin-fin heat sink stacked on a silicon interposer (3D); top view X-ray image, and angled view of fluidic I/Os.

IV. Thermal Coupling in 2.5D and 3D Systems

Thermal coupling forms a fundamental challenge for 2.5D and 3D ICs, where each die within the stack experiences unwanted thermal crosstalk from the others, particularly between high-power dice and low-power, temperature sensitive, components. In this section, the thermal coupling issue and proposed solutions for 2.5D and 3D heterogeneous stacks are discussed. For 2.5D systems, microfluidic cooling integrated at the die level can provide better thermal isolation than a traditional air cooled heat sink. For 3D heterogeneous stacks, an additional air gap isolation is proposed to deal with the strong vertical thermal coupling phenomenon.

A. Thermal Coupling in 2.5D

In an interposer based 2.5D stack with an air cooled heat sink (Fig. 10(a)), thermal coupling results from heat transfer through the heat spreader as well as the heat sink base and the interposer. Therefore, to avoid thermal coupling, the low-power logic die has to be put further away from the other logic dice, resulting in a communication penalty due to longer interposer-level interconnects [11]. To solve this problem, die-specific microfluidic cooling is proposed, as shown in Fig. 10(b). Due to the separate heat paths for each component,

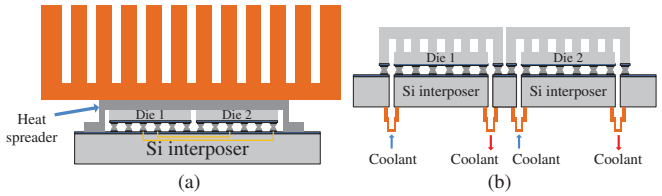

Fig. 10. Silicon interposer based 2.5D stacks using (a) conventional air cooled heat sink and (b) embedded die-specific microfluidic heat sink

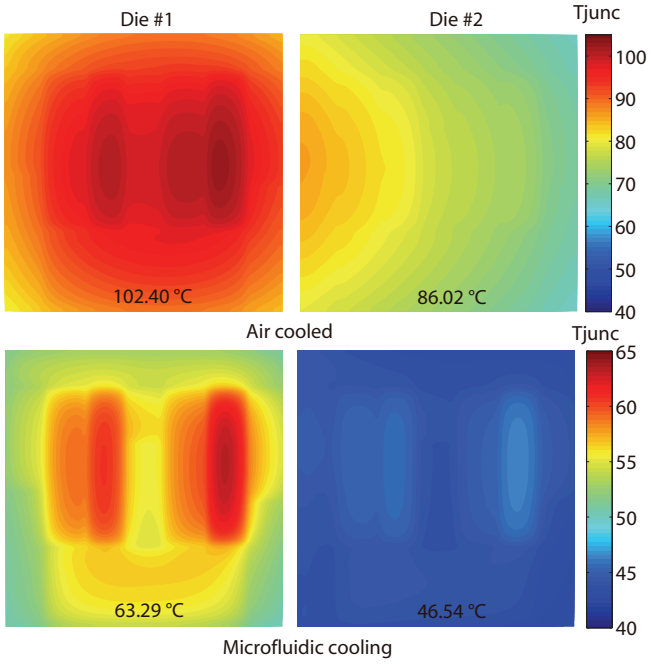

Fig. 11. Simulated temperature maps of the dice with microfluidic cooling and air cooling

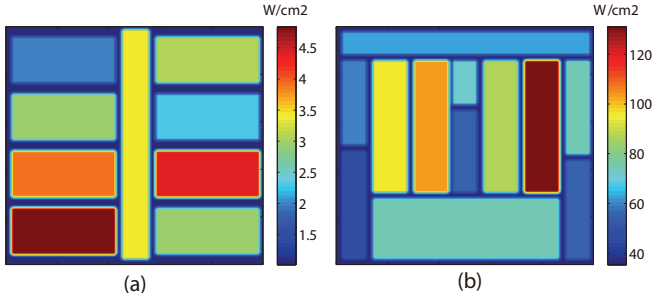

Fig. 12. Power maps (a)memory die (b)processor die

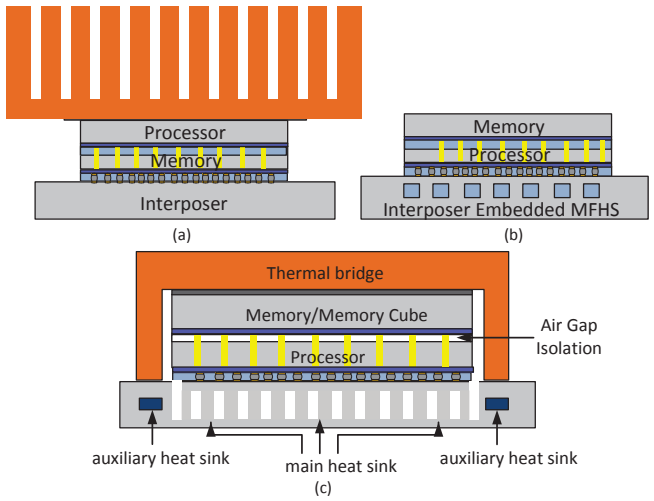

Fig. 13. (a)3D stack with air cooled heat sink (b)3D stack with interposer microfluidic cooling (c)Proposed stack with separate microfluidic cooling and thermal bridge

most of the heat generated by the chips is expected to dissipate through the microfluidic heat sinks and there will be less thermal coupling between the two dice [12].

To compare the two stacks shown in Fig. 10, we use a finite volume method based thermal simulator to quantify the temperature of both stacks [12]. The simulator utilizes non-conformal meshing, which enables us to evaluate large geometries that are consistent with typical interposer and heat spreader dimensions. The simulator was validated against ANSYS with an error of less than 3%. For the microfluidic cooled case, a thermal resistance of 0.24 K cm^2/W [11] was used in the thermal models. The air-cooled heat sink design and attributes are similar to those used for the Intel i7 microprocessor [13]; the heat spreader is 5 cm by 4.5 cm and the total thermal resistance from the heat spreader to ambient is 0.218 K/W.

The power map of the logic die is based on the Intel i7 microprocessor and shown in Fig. 12 (b). The processor power is assumed to be 74.49 W [12]. 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 x 3 cm. There is a 1 mm gap between the edges of the two dice.

The simulated temperature maps of the dice with microfluidic cooling and air cooling are shown in Fig. 11. As expected, microfluidic cooling significantly reduces the die temperature. The temperature reductions in Die #1 and Die #2 are approximately 39.11 °C and 39.48 °C, 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 through the heat spreader. In contrast, the temperature map of the low-power die using microfluidic cooling does not reflect this thermal coupling. Thus, in addition to the lower junction temperature, microfluidic cooling provides better thermal isolation between high-power and low-power dice.

B. Heterogeneous 3D stack and thermal isolation

For heterogeneous 3D ICs, thermal coupling occurs in the vertical direction. As a result, even a low-power die will exhibit a high temperature when stacked with a high-power logic die. Moreover, the temperature variation over time of the high-power die will be reflected in the low-power die temperature due to this strong thermal coupling. This may be

TABLE I
COMPARISON OF DIFFERENT STACKS

Unit: (°C)	T_{max} (Memory)	T_{max} (Processor)
Air cooled stack	75.06	76.44
Microfluidic cooled stack	65.38	66.05
Proposed stack w/o TSVs	39.63	64.64
Proposed stack with TSVs	51.76	61.44

harmful in heterogeneous ICs where the low-power dice are temperature sensitive, such as memory, photonics, and MEMS. Therefore, a conventional air cooled heat sink (Fig. 13 (a)) and even the interposer microfluidic heat sink (Fig. 13 (b)) will not sufficiently decouple the two dice in a heterogeneous stack.

To solve this thermal coupling challenge, we propose a new 3D stack shown in Fig. 13 (c). The proposed architecture has three novel features: 1) a microfluidic heat sink (MFHS) is integrated in the interposer, and consists of two separate parts. The main MFHS is directly under the high-power chip, i.e. the processor. It serves as the main thermal path for the stack. The auxiliary MFHS is located at the periphery of the interposer and is used to cool the "thermal bridge". 2) An air gap thermal isolation is integrated between the high-power and low-power dice to decouple the thermal crosstalk, and 3) a "thermal bridge" is attached on the top of the isolated low-power die to provide a separate cooling path for it.

Next we used the thermal model mentioned previously to evaluate the three stacks. The power map of the memory and processor dice are shown in Fig. 12 (a) and (b) respectively. The layout is based an 8 Gb 3D DDR3 DRAM design from Samsung [14] and the total power is estimated from the *Micron* DDR3 DRAM datasheet [15], which gives a value of 2.82 W. The thermal specifications are the same as the previous section.

Table I illustrates the maximum temperature of each die in all 3D stack scenarios. From the results, we find our proposed architecture has the lowest temperature. Moreover, our proposed architecture decouples the heat from the processor and the memory. In the first two scenarios, the memory exhibits a temperature value that is similar to the processor. For the proposed stack with thermal isolation, the memory temperature is only 39.63 °C even though the processor temperature is as high as 64.64 °C (assuming no TSVs). However, when TSVs passing through the air gap are inserted in our proposed stack, the thermal coupling increases, as expected. From the last row of Table I, we find the temperature difference of the two dice is only 9.68 °C when 4,900 uniformly spaced TSVs are inserted, compared to 25.01 °C in the case without TSVs. Thus, the TSVs clearly impact the thermal isolation.

However, when the TSVs are clustered in a certain area, thermal coupling is expected to occur only in that area. In this way, the heat from the processor die will be localized. For the memory die, the clustered TSVs usually act as the I/O pins and are outside of the memory cell circuits (labeled by a dashed-line box in Fig. 14 (a)). Hence the memory cell circuits will become relatively free from the impact of the processor because there are no TSVs in this area. Inspired by the above

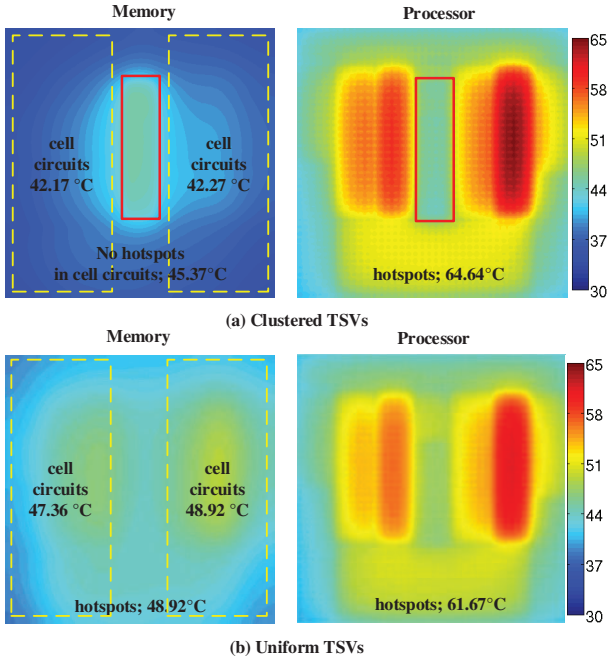

Fig. 14. Thermal maps of clustered TSVs and uniform TSVs (a) TSVs are clustered in the solid-line box (b) The same number of TSVs are uniformly distributed

analysis, we cluster the TSVs only in the center, labeled by the solid-line rectangle in Fig. 14 (a). The results with clustered TSVs and uniformly distributed TSVs are shown in Fig. 14.

In the clustered TSV case, the maximum temperature of the whole DRAM die drops by 3.55 °C compared to the uniform TSV case. However, the maximum temperature of the cell array circuits is only 42.27 °C, which is a drop of 6.65 °C and is much closer to the temperature in the ideal case without TSVs, 39.63 °C. By clustering the TSVs far from the memory cells, the most thermally-sensitive portion of the die is effectively isolated from the high-power die. Thus, we conclude that, in conjunction with air gap isolation, clustering TSVs can localize thermal coupling in 3D stacks.

V. A MONOLITHICALLY INTEGRATED MICROPIN-FIN HEAT SINK IN AN ACTIVE FPGA

FPGAs serve a wide range of applications with diverse thermal and packaging requirements. There are two mainstream packages for today's FPGA solutions: wire bond BGA (WBBGA) and flip chip BGA (FCBGA). The WBBGA package has a lower thermal performance, with devices dissipating 5–10 W, and they are used in low-end applications. The FCBGA package has a higher thermal performance, with devices dissipating 10–150 W, and they are used in mid-range and high-performance applications.

As FPGA system performance continues to increase in future generations due to innovations in architecture, device, and packaging technologies, power density will likely increase and drive the need for advanced cooling solutions. FPGA

companies are some of the early adopters of die stacking technology in the semiconductor industry. The need for 3D integration of FPGAs is compelling for high performance applications. The ability to efficiently remove heat from the 3D stack is a critical component in ensuring an optimal solution. However, one of the potential show stoppers of deploying 3D IC technology in high-performance FPGA products is the ability to remove heat from 3D stacks. The realization of a 3D FPGA will require innovative cooling solutions and thermally aware design optimization at various levels of design abstraction.

To date, microfluidic cooling has focused on "dummy" silicon dice with metallic heaters representing the heat producing circuitry. We present here, for the first time, a functional microfluidic-cooled CMOS circuit.

An Altera Stratix® V FPGA, built in a 28nm process, was post processed to integrate a micropin-fin heat sink directly into the back of the flip chip bonded silicon die. This microfluidic-cooled FPGA was then tested and benchmarked against a Stratix® V with a stock air cooled heat sink. All testing was performed with Altera Stratix® V DSP development boards.

The FPGA was processed while bonded to its package substrate. The packaged FPGA was first desoldered from the development board before processing. The metal heat spreader was then removed revealing the silicon die and the Bosch process was used to etch the micropin-fins into the back of the FPGA die. The etched micropin-fin dimensions can be seen in Fig. 15.

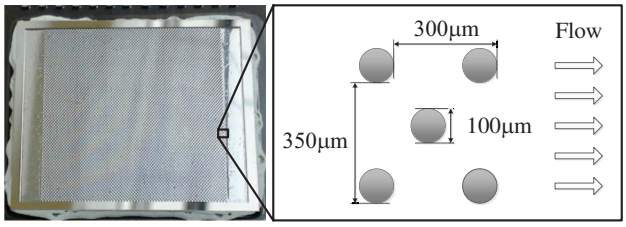

Fig. 15. Micropin-fin dimensions

After the micropin-fins were etched in the silicon die, a capping layer with fluid inlet and outlet ports was attached with high temperature epoxy. The FPGA was then resoldered to the development board. Finally, the edges of the cap-die interface were sealed with epoxy and Nanoports were attached for fluid delivery. The re-assembled FPGA on the development board can be seen in Fig. 16.

The microfluidic-cooled FPGA was tested using an open loop testbed with deionized water. It was benchmarked against a second Stratix® V DSP development board with the stock air cooled heat sink provided with the development kit. A pulse compression algorithm was used to benchmark the two FPGAs. The design consists of 9 processing cores, which can be turned on and off during run time. FPGA temperatures and powers can be seen in Table II for 0–9 active processing

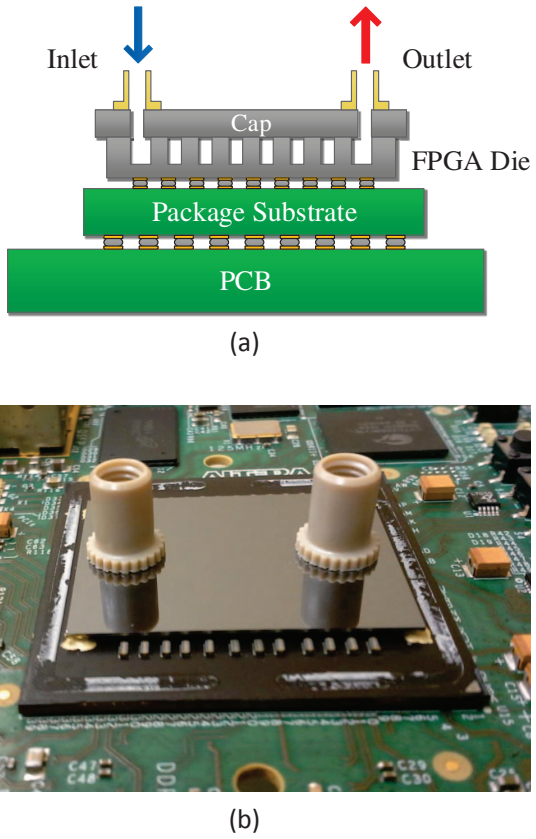

Fig. 16. (a) Cross sectional diagram and (b) photograph of FPGA assembled on development board with Nanoports for fluid delivery

cores for both the FPGA with integrated microfluidic heat sink (MFHS) and with an air cooled heat sink (ACHS).

The ambient air temperature was 19.4 °C and the deionized water inlet temperature for the microfluidic cooled FPGA was 20.1–20.6 °C. The deionized water was pumped through the FPGA at a a rate of 147 mL/min at a pressure drop of 70.2–72.5 kPa (including inlet/outlet port pressure drop). Air was blown over the voltage regulator on the development board to prevent it from overheating. Had the testing been done without cooling of the voltage regulator, the FPGA temperatures would have been higher, particularly with the air cooled heat sink.

FPGA temperature and power measurements were taken using the Altera Power Monitor tool. Temperature measurements are read from a temperature diode located at the corner of the FPGA die. Therefore, in order to account for the temperature gradient caused by heating of the fluid, data was taken with water flowing in both directions, thus producing the range of microfluidic cooled FPGA temperatures shown in Table II.

The FGPA with integrated micropin-fin heat sink showed marked improvements in die temperature, maximum throughput, and leakage current. A maximum die temperature of 60 °C was set based on the built in temperature warning on the Stratix® V development board. The stock air cooled FPGA ran

TABLE II
FPGA THERMAL AND POWER MEASUREMENTS WITH MICROFLUIDIC HEAT SINK (MFHS) AND AIR COOLED HEAT SINK (ACHS)

Cores	MFHS FPGA Power (W)	ACHS FPGA Power (W)	MFHS FPGA Temp (°C)	ACHS FPGA Temp (°C)
0	13.2	13.7	21–22	43
1	15.4	16.0	21–23	46
2	17.6	18.3	22–23	49
3	19.8	20.5	22–23	51
4	21.9	22.8	22–23	53
5	24.0	25.1	22–23	56
6	26.2	27.5	22–23	59
7	28.3	29.8	22–24	61[a]
8	30.4	—	22–24	—
9	32.4	—	22–24	—

[a]Temperature warning on board illuminated.

6 out of the available 9 cores before reaching this temperature while the microfluidic-cooled FPGA ran all 9 cores with a steady state die temperature of less than 24 °C. In addition to this increase in usable silicon, the microfluidic cooled FPGA ran at 5% lower total power when running 7 cores due to a reduction in leakage current.

Although the large difference in performance between the two FPGAs is partially a function of the air cooled heat sink which comes with the board, the thermal resistance of the microfluidic heat sink is approximately 0.08 °C/W (using the average of the recorded FPGA temperatures at maximum power), which is far better than the best air cooled heat sinks. Improving the air cooled heat sink would also likely require increasing its size, making dense integration impossible, whereas this microfluidic heat sink is integrated at the die-scale, which enables 3D stacking.

VI. CONCLUSION

In this paper, a number of technologies to solve the thermal problems of modern integrated circuits were presented. Polymer embedded vias were demonstrated for low loss interposer communication. Fluidic I/Os were presented as a means of delivering fluid to and from microfluidic heat sinks in 2.5D and 3D stacks. The challenge of thermal coupling in these dense systems was then addressed through the use of microfluidic cooling and an air isolation gap. Lastly, an Altera FPGA with monolithically integrated micropin-fin heat sink was presented, showing that this cooling technology can be integrated into real systems and deliver the performance expected from isolated heat sink demonstrations.

Current integrated circuits, including the Altera FPGA used in this paper, are built with current thermal dissipation limitations taken into account as can be seen in the very low temperatures achieved with the microfluidic heat sink in this paper. Eventually, circuitry can be designed taking these new, improved, thermal constraints into account, allowing designers to exploit higher clock speeds, or maintain lower temperatures and leakage powers. Eventually this will enable 3D stacking of multiple functional, high-power logic dice in a 3D stack for maximum computing density, throughput, and efficiency.

ACKNOWLEDGMENTS

We would like to thank Dale Ibbotson, Karthik Chandrasekar, and Vadali Mahadev of Altera Corporation for their help with this work. We would also like to thank DARPA for their support of this work under ICECool Applications (HR0011-14-1-0002) and the DARPA Young Faculty Award (N66001-12-1-4240).

REFERENCES

[1] "Intel® product information," Intel® Corporation, accessed 4/23/15. [Online]. Available: http://ark.intel.com

[2] K. Bowman, A. Alameldeen, S. Srinivasan, and C. Wilkerson, "Impact of die-to-die and within-die parameter variations on the clock frequency and throughput of multi-core processors," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 17, no. 12, pp. 1679–1690, Dec 2009.

[3] P. Quinn, "Fpga based silicon innovation exploiting "more than moore" technology," in *Ph.D. Research in Microelectronics and Electronics (PRIME), 2013 9th Conference on*, June 2013, pp. 11–12.

[4] A. Rahman, J. Schulz, R. Grenier, K. Chanda, M. Lee, D. Ratakonda, H. Shi, Z. Li, K. Chandrasekar, J. Xie, and D. Ibbotson, "Interconnection requirements and multi-die integration for fpgas," in *Interconnect Technology Conference (IITC), 2013 IEEE International*, June 2013, pp. 1–3.

[5] J. S. Pak, C. Ryu, and J. Kim, "Electrical characterization of through silicon via (TSV) depending on structural and material parameters based on 3D full wave simulation," in *Proc. International Conference on Electronic Materials and Packaging (EMAP)*, Nov. 2007, pp. 1–6.

[6] V. Sukumaran, G. Kumar, K. Ramachandran, Y. Suzuki, K. Demir, Y. Sato *et al.*, "Design, fabrication, and characterization of ultrathin 3-D glass interposers with through-package-vias at same pitch as TSVs in silicon," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 4, no. 5, pp. 786–795, May 2014.

[7] I. Ndip, B. Curran, K. Lobbicke, S. Guttowski, H. Reichl, K.-D. Lang, and H. Henke, "High-frequency modeling of TSVs for 3-d chip integration and silicon interposers considering skin-effect, dielectric quasi-TEM and slow-wave modes," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 1, no. 10, pp. 1627–1641, Oct. 2011.

[8] P. A. Thadesar, L. Zheng, and M. S. Bakir, "Low-loss silicon interposer for three-dimensional system integration with embedded microfluidic cooling," in *Proc. IEEE Symposium on VLSI Technology: Digest of Technical Papers*, Jun. 2014, pp. 1–2.

[9] L. Zheng, Y. Zhang, G. Huang, and M. Bakir, "Novel electrical and fluidic microbumps for silicon interposer and 3-d ics," *Components, Packaging and Manufacturing Technology, IEEE Transactions on*, vol. 4, no. 5, pp. 777–785, May 2014.

[10] L. Zheng, Y. Zhang, and M. S. Bakir, "Silicon interposer with embedded microfluidic cooling for high-performance 3-D ICs," in *Proc. IEEE Electronic Components and Technology Conference (ECTC)*, Accepted, May 2015.

[11] L. Zheng, Y. Zhang, and M. Bakir, "A silicon interposer platform utilizing microfluidic cooling for close integration of high-performance ics," *IEEE Trans. Compo. Packag. Manuf. Technol.*, 2014.

[12] Y. Zhang, Y. Zhang, and M. Bakir, "Thermal design and constraints for heterogeneous integrated chip stacks and isolation technology using air gap and thermal bridge," *IEEE Trans. Compo. Packag. Manuf. Technol.*, 2014.

[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

[14] U. Kang, H.-J. Chung, S. Heo, S.-H. Ahn, H. Lee, S.-H. Cha *et al.*, "8gb 3d ddr3 dram using through-silicon-via technology," in *Proc. IEEE Solid State Circuits Conf.*, 2009.

[15] Micron, "Technical Note Calculating Memory System Power for DDR3." [Online]. Available: http://www.micron.com/support/power-calc