This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCPMT.2019.2930481, IEEE Transactions on Components, Packaging and Manufacturing Technology

Microfluidic Cooling of a 14nm 2.5D FPGA with 3D Printed Manifolds for High Density Computing: Design Considerations, Fabrication, and Electrical Characterization

Thomas E. Sarvey *Member, IEEE*, Ankit Kaul, Sreejith Kochupurackal Rajan, Aravind Dasu, Ravi Gutala, and Muhannad S. Bakir, *Senior Member, IEEE*

***Abstract*—2.5D integration is becoming a common method of tightly integrating heterogeneous dice with dense interconnects for efficient, high-bandwidth inter-die communication. While this tight integration improves performance, it also increases the challenge of heat extraction by increasing aggregate package powers and introducing thermal cross-talk between adjacent dice. In this work, a microfluidic heat sink is used to cool a 2.5D Stratix 10 GX FPGA consisting of an FGPA die surrounded by four transceiver dice. The heat sink utilizes a heterogeneous micropin-fin array with micropin-fin densities which are tailored to the local heat fluxes of the underlying dice. Enabled through a 3D-printed enclosure for fluid delivery, the assembled heat sink has a total height of 6.5 mm, including the tubes used for fluid delivery. The heat sink is tested in an open loop system with deionized water as a coolant and thermal performance is compared against a high-end air cooled heat sink. Improvements in die temperatures, computational density, and thermal coupling between dice are observed. The effect of the FPGA power on the surrounding transceiver die temperatures was reduced by a factor of 10× to over 100× when compared with the air cooled heat sink.**

***Index Terms*—microfluidic cooling, micropin-fin, FPGA, transceiver, EMIB, 2.5D**

I. Introduction

Interconnection has become a bottleneck in computing performance and efficiency at every level of integration. While the solution to this problem is decreasing interconnect length and increasing system density, the ability to remove heat already limits this approach. Computational density is primarily limited by the large volume of air that must be used to capture heat with a reasonable increase in temperature.

By switching to a liquid coolant, such as water, which has a volumetric heat capacity which is more than 3000× higher than that of air (close to standard temperature and pressure), the necessary volume for heat exchange and fluid delivery can be dramatically decreased. In addition to enabling higher density at the die-level in this fashion, microfluidic cooling can be applied at the package level to enable high density stacking of printed circuit boards (PCBs).

For example, high power computing accelerators are often integrated onto boards and connected to a main PCB through peripheral component interconnect express (PCIe) slots. The air cooled heat sink is the largest component on the board and limits the pitch at which the boards can be mounted, and therefore the number of accelerator boards that can be integrated into a single computer/server.

In addition to addressing the challenge of decreasing heat sink volume, these heat sinks must be able to address the needs of modern high power packages. These high power accelerator packages no longer include a single monolithic die, but several dice mounted in close proximity to one another, usually with an interposer or embedded bridge chips for high bandwidth interconnection [1]–[4]. These dice implement heterogeneous functionalities and may exhibit significant thermal coupling through their shared heat spreader [5].

As microelectronics transition from monolithic dice to large packages of heterogeneous chiplets, heat sinks can also be modified to match these heterogeneous designs. A low cost polymer manifold can be used to assemble microfluidic cooling chiplets to construct a large heterogeneous heat sink that can be attached to a 2.5D heterogeneous collection of chiplets. Alternatively, electrically functional chiplets could be assembled in a plastic (or similar) manifold. Bridge chips or interposers could be flip chip bonded to the chiplets, followed by the package substrate. Chiplets could either be etched, or jet impingement cooling could be used to obtain a substantial heat transfer coefficient across the backside of unetched dice. Alternatively, these cooling chiplets could be combined into a single heterogeneous micropin-fin heat sink which is attached to the backsides of the dice, the approach taken in the following work.

Prior work has successfully used heterogeneous micropin-fins as a means of normalizing heterogeneous power maps. By increasing the micropin-fin density proportionally to the underlying heat flux, the authors were able to reduce the temperature across a 500 µm × 500 µm hot spot to match the average surrounding temperature of a 1 cm × 1 cm test chip producing a heat flux of 250 W/cm^2 [6], [7]. This concept is extended in this work to address package-level heat flux and temperature heterogeneity across multiple dice in a 2.5D package.

In this work, a silicon micropin-fin heat sink is designed and fabricated to cool a Stratix 10 GX field-programmable gate array (FPGA) consisting of five heterogeneous dice. The micropin-fin geometry is locally varied to match the heat flux of the underlying dice. The silicon micropin-fin heat sink is embedded in a 3D-printed plastic piece which seals the top

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCPMT.2019.2930481, IEEE Transactions on Components, Packaging and Manufacturing Technology

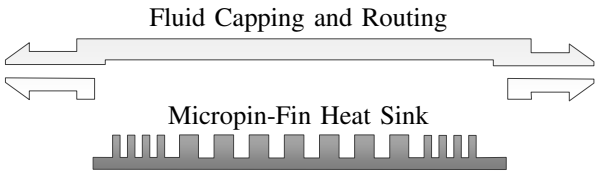

(a) Two parts of the microfluidic heat sink

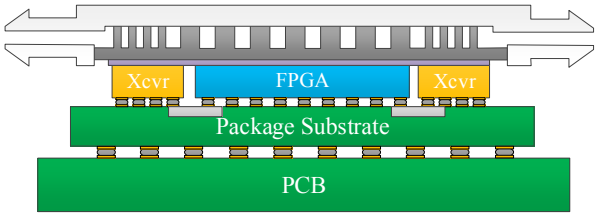

(b) Microfluidic heat sink assembled on dice

Fig. 1: Cross-sectional diagram of micropin-fin heat sink for 2.5D package

Fig. 2: Stratix 10 ES development board

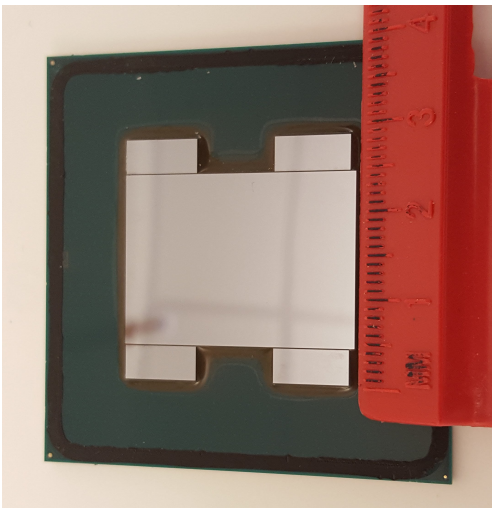

(a) Delidded Stratix 10 FPGA

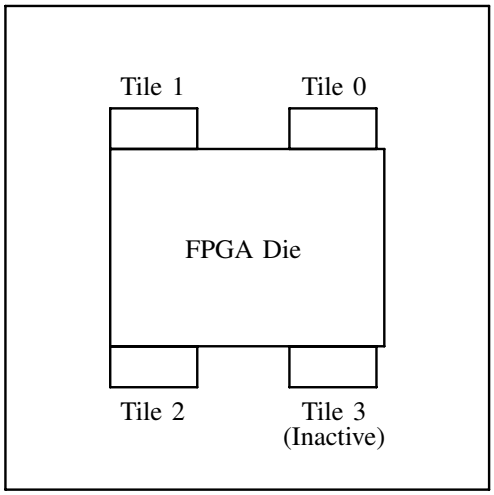

(b) Stratix 10 FPGA die layout

Fig. 3: Stratix 10 GX package

of the silicon micropin-fins and connects the heat sink to inlet and outlet tubing. A conceptual cross-sectional diagram of the heat sink concept can be seen in Fig. 1.

II. Experimental Setup

The experiments in this work were carried out using a Stratix 10 engineering silicon (ES) development kit from Intel. A photograph of the board can be seen in Fig. 2. The board carries a Stratix 10 GX FPGA which is a 2.5D device consisting of a 14 nm FPGA core die surrounded by four transceiver dice, connected through Intel's embedded multi-die interconnect bridge (EMIB). A photo of a delidded package can be seen in Fig. 3a.

Each transceiver tile (die) contains 24 transmitters and receivers, for a total of 96 in the package. Each receiver requires a dedicated reference clock connected directly through the package to the transceiver tile on which it resides. The Stratix 10 ES development board only has reference clocks connected to three of the four transceiver tiles, so only three of the four transceiver tiles are used in this work.

A. FPGA Benchmark Application

The benchmark program used in this work was designed to mimic a high power use case of the FPGA. It consists of a portion which dissipates power on the FPGA, and a portion which dissipates power on the transceivers. The core of the design consists of a streaming fast fourier transform (FFT) block followed by six first-in, first-out (FIFO) buffers operating on random inputs hard coded into the FPGA. This design is implemented on the FPGA through a combination of programmable logic blocks and digital signal processing (DSP) blocks. Much of the computational throughput as well as power dissipation comes from the FPGA's DSP blocks, which perform arithmetic operations on floating point operands. The FFT core was replicated 160 times across the FPGA and clocked at 475 MHz. The clock (and power) could be raised higher, but the voltage regulator modules (VRMs)

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCPMT.2019.2930481, IEEE Transactions on Components, Packaging and Manufacturing Technology

on the board are only designed to supply a maximum of 100 A of current on the VCC rail, causing instability when the FPGA was clocked at significantly higher frequencies. Even at 475 MHz, the FPGA used well over 100 A on the VCC rail and air was therefore blown over the VRMs to prevent them from overheating during experiments.

In addition to this FFT design, 72 transceiver channels were utilized to dissipate power on three of the four transceiver tiles. Each of the 72 transceiver channels were programmed to run in enhanced physical coding sublayer (PCS) mode with serial loopback enabled. This design was modified from a publicly available design from the Intel FPGA Wiki [8]. The maximum data rate within the Intel transceiver intellectual property (IP) for the GX series of Stratix 10 FPGAs is 16 Gbit/s per channel, but the transceiver clocks were increased during runtime to overclock the transceivers to a data rate of 22 Gbit/s. The future Stratix 10 TX FPGA will be available with up to 56 Gbit/s of bandwidth per channel and will likely dissipate even more power.

One temperature sensor on each die is read using a Nios II soft processor on the FPGA which feeds these numbers back to an attached computer for logging. Voltage and current on the power rails of the board were measured through on-board sensors which interface to a board test system (BTS) interface that comes with the development board.

The board has six power rails which can be monitored in the BTS interface. For power estimates, it was assumed that two of these rails represent power dissipated on the FPGA tile, while the sum of the remaining rails was assumed to be the power dissipated equally across the transceiver tiles. More granular, die specific power measurement is not possible because the individual dice within the package are not connected to independent power rails. Therefore, it was further assumed that the transceiver tiles dissipate identical amounts of power when running identical transceiver configurations. In actuality, there will be manufacturing variations across dice, but this assumption allowed for the estimation of power dissipation on each die for the purposes of heat sink design.

III. Heat Sink Design

Baseline power and temperature measurements were first made on the development board with the air cooled heat sink. First, power was measured before loading the benchmark design, with the clocks disabled to estimate leakage power and, in particular, the power dissipated on the unused transceiver die. Next, the benchmark application was loaded onto the FPGA and power and temperature were measured across the FPGA and three active transceiver dice. These initial measurements were made without waiting the full stabilization period used for the air cooled measurements reported later, which brought die temperatures closer to the temperatures observed with microfluidic cooling. As will be shown, this initial power estimate, which was used to design the microfluidic heat sink, was within 0.2 W of the actual total power measured with the microfluidic cooled heat sink at the highest flow rate.

The aforementioned assumption that power is equally distributed across the transceiver dice which run identical applications was used to estimate the power on each transceiver die for the purpose of designing the microfluidic heat sink. It was assumed that the leakage power determined in the previous section was equally split across all four dice, while the dynamic power was equally split across the three active transceiver dice. It was also assumed that the leakage power, measured without any design running, was equal to the leakage power when running the benchmark application. The power breakdown for each of the five physical dice used for simulations can be seen in Table I.

TABLE I: Stratix 10 die powers, areas, and power densities

Die	Power	Area	Power Density
FPGA	110 W	5.8 cm^2	19 W/cm^2
Transceiver Tile 0	21 W	0.38 cm^2	56 W/cm^2
Transceiver Tile 1	21 W	0.38 cm^2	56 W/cm^2
Transceiver Tile 2	21 W	0.38 cm^2	56 W/cm^2
Inactive Tile 3	0.72 W	0.38 cm^2	1.9 W/cm^2

An initial heterogeneous micropin-fin design was created by beginning with a high aspect ratio micropin-fin geometry which could be etched to the desired depth while maintaining nearly vertical sidewalls. This high aspect ratio micropin-fin design was used to populate the regions of the heat sink directly over the transceiver dice, while a second micropin-fin design was used to cool the region above the FPGA die. The pitch and diameter were scaled by the relative power density of the transceiver and FPGA dice to create the dimensions used to cool the center FPGA die.

IV. Heat Sink Simulation

Ansys FLUENT was used to simulate a silicon microfluidic heat sink with an applied power map matching the values shown in Table I, applied to the underside of the silicon micropin-fin heat sink. The micropin-fins were 460 µm tall, while the base silicon underneath the micropin-fins was 440 µm thick.

A top view of the simulated 3D model can be seen in Fig. 4a, with the fluid region highlighted in yellow. A symmetry plane down the center of the chip was used to reduce the model size. This means that the fourth transceiver tile was modeled as active and identical to the other transceivers, although the experimental results do not include this transceiver tile. The model assumes uniform inlet velocity and pressure outlet boundary conditions. This does do not include fluid delivery effects of the plastic enclosure used in experimentation. This model was meshed in ANSYS using proximity and curvature, yielding a mesh with 39.0 million elements. Though the simulation required nearly all available memory on the simulation computer node, which had 200 GB of usable memory, a denser second mesh with 12% more elements was created to observe the effect of mesh density on simulated pressures and temperature. The temperature rise above inlet and pressure drop differed from the original simulation by a maximum of 1.3% and 1.7%, respectively.

Simulation results were compared with simulations of a uniform micropin-fin geometry, where the micropin-fin geometry used over the FPGA was extended across the entire heatsink. This model can be seen in Fig. 4b. This model was meshed

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCPMT.2019.2930481, IEEE Transactions on Components, Packaging and Manufacturing Technology

using the same configuration as the non-uniform geometry, yielding a mesh with 31.8 million elements.

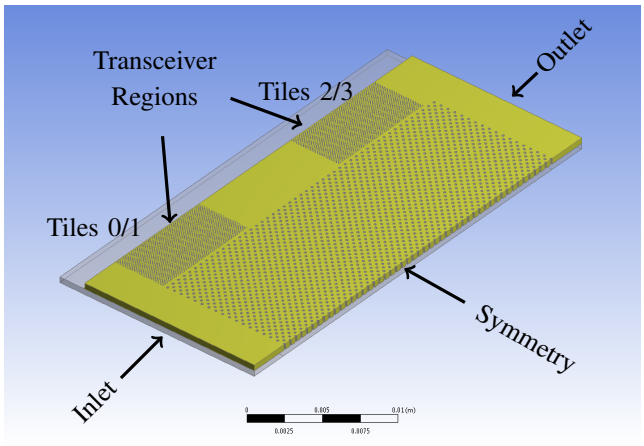

(a) Heterogeneous micropin-fin design

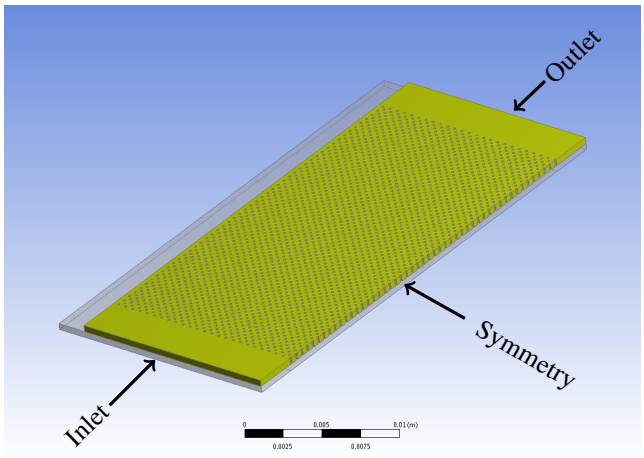

(b) Uniform micropin-fin heat sink

Fig. 4: Simulated micropin-fin heat sink models

Temperature dependent material properties were used within FLUENT for water viscosity and thermal conductivity as well as the thermal conductivity of silicon. The specific heat of water, which varies very little with temperature, was taken to be a constant 4182 J/(kg K).

Specific heat and viscosity values for water were taken from the NIST WebBook[9] and fit to polynomials over the temperature range 273.16 K to 370 K. Second and fourth order polynomial functions were fit to the thermal conductivity and viscosity values, respectively. The following polynomial function was generated for the thermal conductivity of water

$$k_w = -0.7491 + (7.430 \times 10^{-3})T - (9.658 \times 10^{-6})T^2 \quad (1)$$

where k_w is the thermal conductivity of water in W/(m K) and T is the temperature in K. The dynamic viscosity of the water is given by

$$\mu_w = 0.463482 - (5.383601 \times 10^{-3})T + (2.357235 \times 10^{-5})T^2 - (4.60188 \times 10^{-8})T^3 + (3.375844 \times 10^{-11})T^4 \quad (2)$$

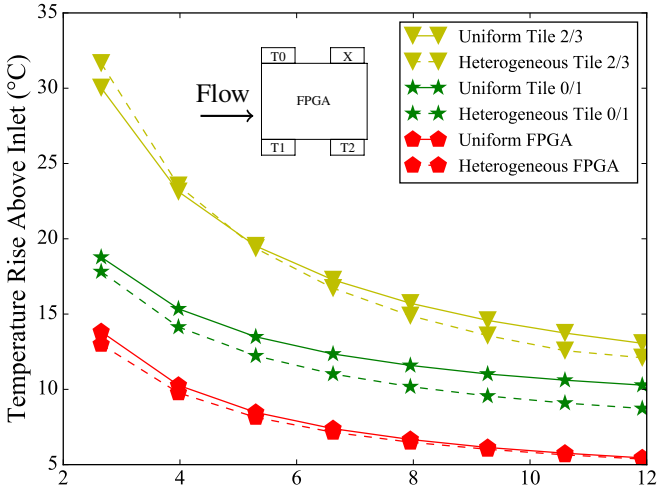

Fig. 5: Simulated temperatures of FPGA and transceiver regions with heterogeneous and uniform heat sinks

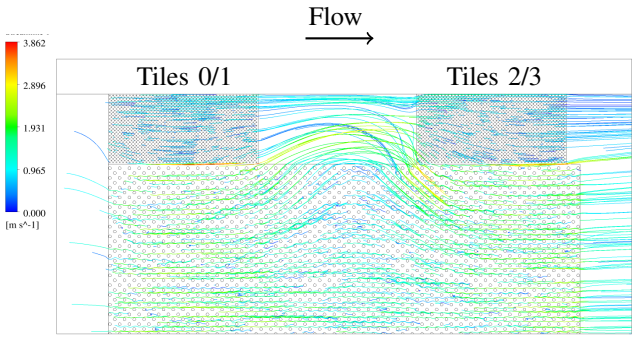

Fig. 6: Fluid flow streamlines with an inlet velocity of 0.9 m/s

where μ_w is the dynamic viscosity of water in units of Pa s. Across this temperature range, the two functions for water thermal conductivity and viscosity differ from the NIST data points by a maximum of 0.2% and 3.6%, respectively. The following correlation from [7] was used for the thermal conductivity of silicon

$$k_{Si} = 2122.1 - 16.765T + (4.818 \times 10^{-2})T^2 - (4.744 \times 10^{-5})T^3 \quad (3)$$

where k_{Si} is the thermal conductivity of silicon.

Simulations were performed in ANSYS FLUENT for both geometries with inlet flow velocities of 0.2 m/s to 0.9 m/s, corresponding to flow rates of 2.6 mL/s to 11.9 mL/s. Average temperatures across the FPGA and transceiver heat flux regions were extracted and can be seen in Fig. 5.

At low velocities, the temperatures between the uniform and non-uniform heat sinks were similar. However, as the flow rate increased the temperature of the transceiver dice dropped below the temperatures of the dice cooled with the uniform micropin-fin heat sink. At low fluid velocities, when the temperature rise is primarily due to heating of the fluid, the effect of this fluid bypass is more pronounced. This partial fluid redirection around the higher density regions can be seen

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCPMT.2019.2930481, IEEE Transactions on Components, Packaging and Manufacturing Technology

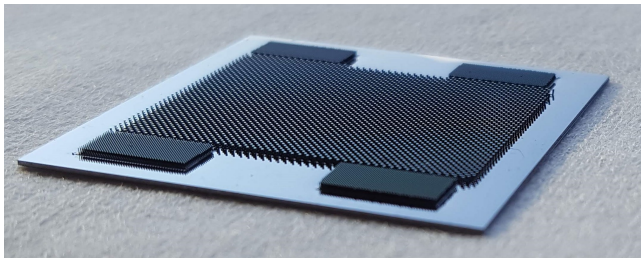

Fig. 7: Silicon micropin-fin heat sink for Stratix 10 FPGA

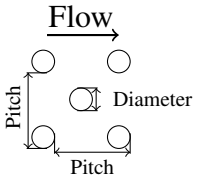

Transceiver Region Diameter	60 µm
Transceiver Region Pitch	200 µm
FPGA Region Diameter	175.6 µm
FPGA Region Pitch	585.1 µm
Height	460 µm

Fig. 8: Micropin-fin dimensions

through the streamlines in Fig. 6. Future heat sink versions may benefit from a barrier preventing lateral fluid movement between the regions.

V. Heat Sink Fabrication and Assembly

The complete microfluidic heat sink assembly consists of two parts: the etched silicon heat sink, through which heat is transferred to the fluid, and a 3D printed plastic enclosure which encases the silicon insert and routes fluid between the inlet/outlet tubes and the silicon micropin-fins.

A photograph of the silicon micropin-fin heat sink can be seen in Fig. 7. The heat sink was fabricated through Bosch-process etching of a 900 µm thick silicon wafer to a depth of approximately 460 µm.

Since both sets of micropin-fin dimensions were etched to approximately the same depth, the micropin-fins over the transceivers were much higher aspect ratio than those over the FPGA (approximately 7.7:1). Therefore, the micropin-fin dimensions used for the transceivers were chosen first with a pitch-to-diameter ratio which yielded minimal sidewall taper when etched to approximately 460 µm. The micropin-fin dimensions over the FPGA region were then chosen by scaling both the pitch and diameter by the respective heat flux densities of the two regions. The dimensions are summarized in Fig. 8.

A scanning electron microscope (SEM) image of both micropin-fin regions can be seen in Fig. 9. The micropin-fins over the transceiver regions were fabricated with a relatively high aspect ratio of approximately 7.7:1. Since the profile of the etched sidewalls depends on the surrounding features, the geometry and fabrication recipe which yielded minimal taper across the bulk of the micropin-fin arrays produced significant taper along the outermost edge of the micropin-fins. These outermost rows represent a small fraction of the total number of micropin-fins.

The height of the micropin-fins were measured using an Olympus LEXT optical profilometer. Since etch depth can locally vary based on local features, measurement windows

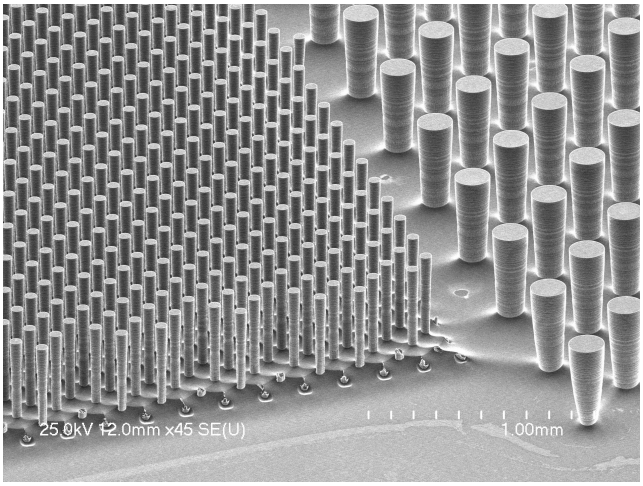

Fig. 9: SEM image of high and low density micropin-fins

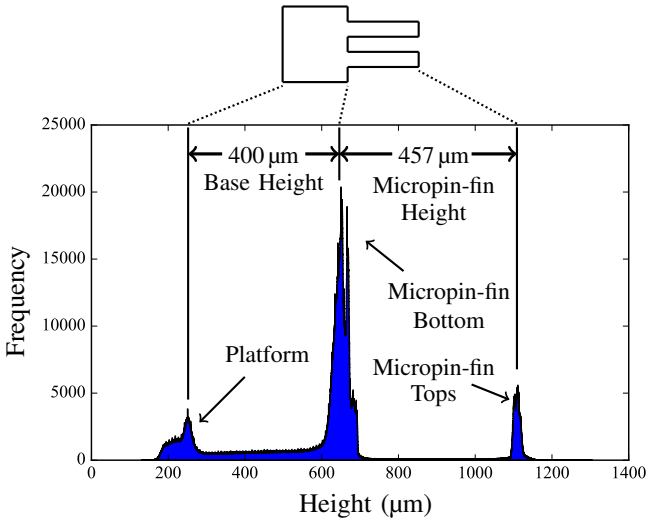

Fig. 10: Histogram of micropin-fin heat sink heights

were stitched together to collect height data across approximately one quarter of the chip, capturing one full transceiver micropin-fin region and approximately one quarter of the FPGA micropin-fin region. Since the entire micropin-fin heat sink is nearly symmetrical across two planes, this quarter of the chip is assumed to be representative of the entire chip. 13.3 million height measurements across this region are aggregated into a histogram in Fig. 10 with a bin width of 0.1 µm. From this histogram it can be estimated that the base silicon had a mode height of 400 µm and the micropin-fins had a mode height of 457 µm.

The 3D printed enclosure can be seen in Fig. 11. A recess with a nominal depth of 450 µm exists for the silicon heat sink. A further recess with an additional nominal depth of 450 µm exists to enclose the micropin-fins. The height of the micropin-fins was made to be approximately 10 µm taller than this recess to reduce the likelihood of a gap existing between the tops of the micropin-fins and the plastic. Since the micropin-fins were designed to be taller than the cavity, the edges of the die were not expected to touch the edges of the plastic, but this gap

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCPMT.2019.2930481, IEEE Transactions on Components, Packaging and Manufacturing Technology

was filled with epoxy.

Epoxy was dispensed with a syringe along the edges of the cavity before inserting the silicon die. A groove was added between the edges where epoxy was applied and the micropin-fin region, so that epoxy which was pushed out during assembly would fill this groove before clogging the micropin-fins. A close-up photograph of the two levels of recess, the groove, and an inlet slit can be seen in Fig. 11b.

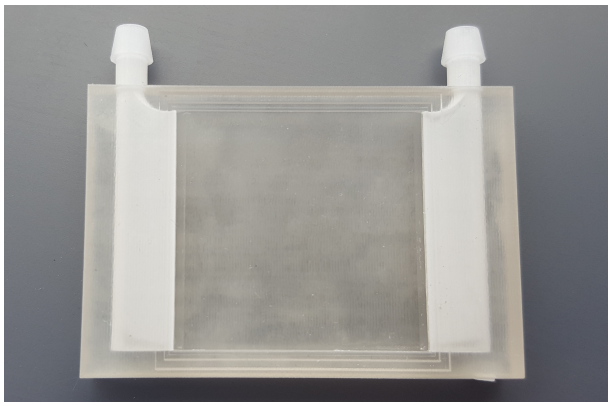

(a) Top view of 3D printed enclosure

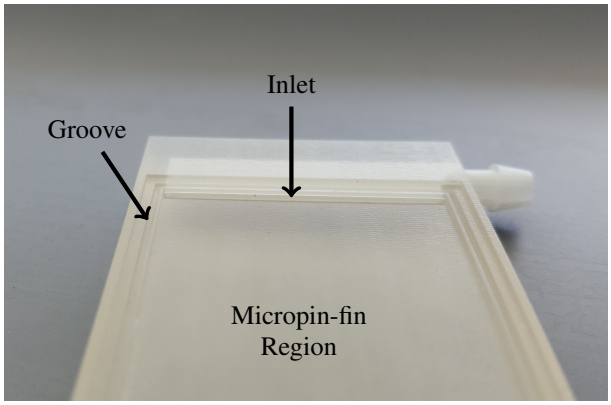

(b) Close up of enclosure features

Fig. 11: 3D printed enclosure and fluid routing device

The height profile across the enclosure was taken using an Olympus LEXT optical profilometer (from left to right in Fig. 11b) and can be seen in Fig. 12. As the plots shows, the region of the enclosure which sits atop the micropin-fins is not completely uniform. Therefore, the micropin-fins are likely to make contact at the edges of the cover, leaving up to approximately 100 µm above some regions of the heat sink where fluid can flow outside of the micropin-fin arrays.

After applying epoxy to the edges of the plastic cover, the silicon heat sink was inserted into the cavity and excess epoxy was wiped from the outside of the assembly. A photograph of the assembled heat sink can be seen in Fig. 13. The entire heat sink was designed to be nearly planar so that it could make contact with the tops of the silicon dice without touching any surrounding components or the board.

The lid of the Stratix 10 package on the development board was removed and the backsides of the five dice were cleaned

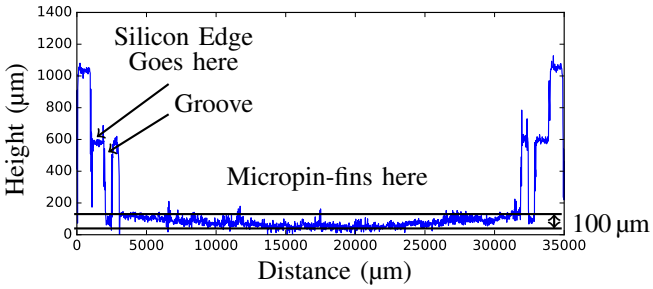

Fig. 12: Profile of micropin-fin heat sink enclosure

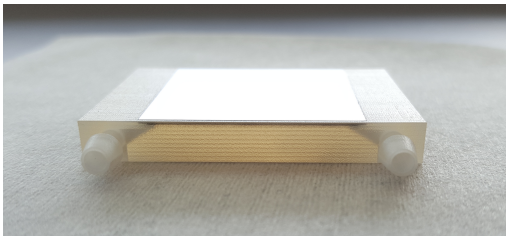

Fig. 13: Assembled microfluidic heat sink

with isopropyl alcohol. The board with the delidded die can be seen in Fig. 14.

After assembling the complete microfluidic heat sink, MasterGel Pro thermal interface material (TIM), which has a thermal conductivity of 8 W/(m K), was applied to the back sides of the five Stratix 10 dice and the heat sink was mounted on top. Pressure was applied using a custom 3D printed mounting bracket. The edges of the heat sink nearest to the inlet and outlet were visually aligned to the edges of the package, while the other two edges were aligned to the edges of the mounting bracket, which was designed to be situated in the center of the package and have the same width as the heat sink enclosure. Images of the mounted heat sink can be seen in Fig. 15.

VI. Air Cooled Heat Sink

A Cooler Master MA621P TR4 air cooled heat sink with one fan was used as a baseline for comparison[10]. The heat sink is designed for a large Advanced Micro Devices (AMD) Threadripper central processing unit (CPU) package and covers the entire 5 cm × 5 cm FPGA package. The heat sink was mounted using custom 3D printed mounting brackets as well

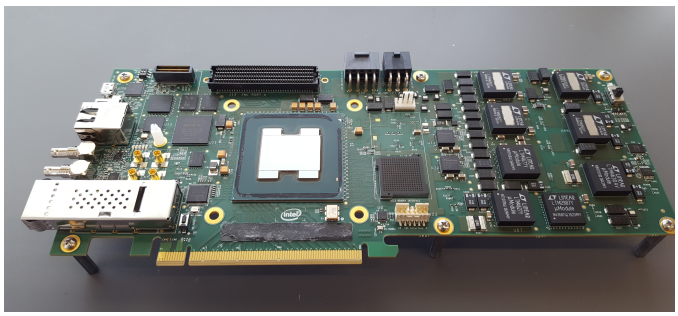

Fig. 14: Delidded Stratix 10 FPGA on development board

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCPMT.2019.2930481, IEEE Transactions on Components, Packaging and Manufacturing Technology

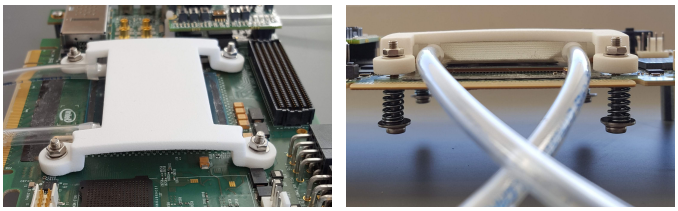

(a) 3D printed mount (b) Screws and springs applying pressure

Fig. 15: Custom mounting for microfluidic heat sink on Stratix 10 development board

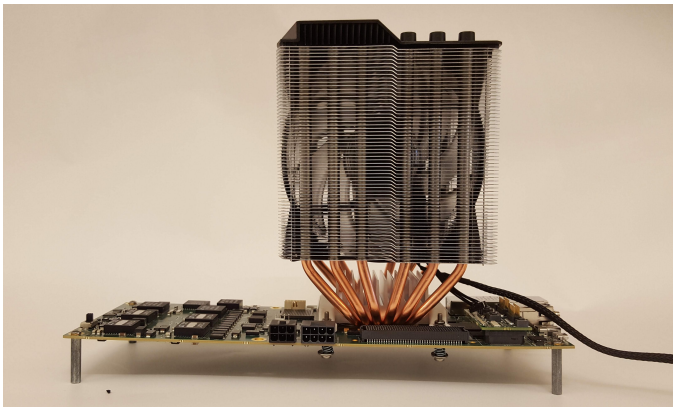

Fig. 16: Stratix 10 ES development board with air cooled heat sink

as the same screws and springs used to mount the microfluidic heat sink. MasterGel Pro TIM was applied between the base of the heat sink and the FPGA heat spreader. The assembled heat sink and FPGA development board can be seen in Fig. 16. The 3D printed mounting brackets and springs can be seen in Fig. 17.

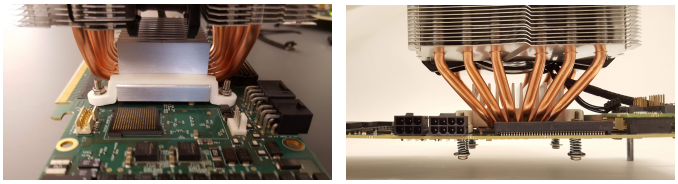

(a) 3D printed mount (b) Screws and springs applying pressure

Fig. 17: Custom mounting for air cooled heat sink on Stratix 10 development board

VII. Experimental Results

The assembled heat sink and board were tested in an open loop system with deionized water as a coolant. Temperature measurements were made at the inlet, outlet, and in the surrounding ambient air using k-type thermocouples. Flow rates were measured with a Kobold rotameter and Ohmega electronic flow meter, both calibrated through repeated filling of a known volume of fluid. Calibration took place with

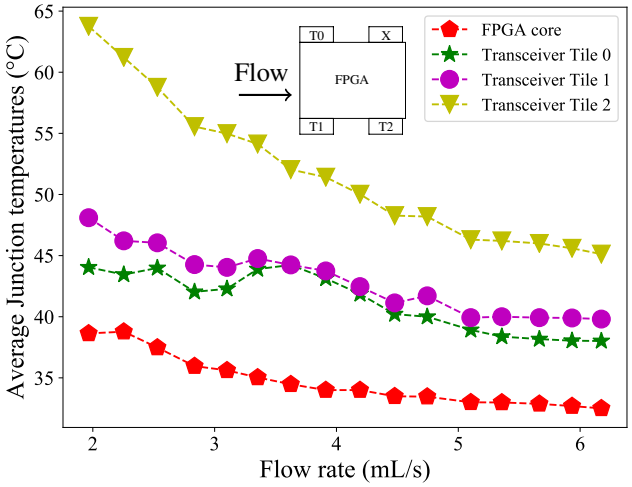

Fig. 18: Die temperatures vs. flow rate with micropin-fin heat sink

deionized water at $21.2 \pm 0.3\,^{\circ}\mathrm{C}$, which was within $2.4\,^{\circ}\mathrm{C}$ of all temperatures used for testing. Die temperatures were measured using on-die temperature diodes with temperature measurement IP integrated into the FPGA design.

An initial experiment was conducted at the lowest flow rate to find the time necessary for the temperatures to reach steady state. Little systematic variation was observed after the first measurement. Nonetheless, temperatures were allowed to stabilize for one minute prior to taking all measurements with the microfluidic heat sink. A similar experiment was conducted with the air cooled heat sink, which took significantly longer to stabilize. Therefore, all of the air cooled measurements in this section were recorded after a wait time of 15 minutes. Data was then taken for approximately one minute at a rate of approximately four data points per second. The averages of die temperatures over these one-minute periods are reported in this section. The maximum standard deviation of any die temperatures across any of these periods was $0.87\,^{\circ}\mathrm{C}$.

The temperature of the FPGA die as well as the three active transceiver dice can be seen as a function of flow rate in Fig. 18. The FPGA die temperature measurement is the lowest of the four dice, while Tile 2 has the highest temperature due to its close proximity to the outlet. The temperature difference between Tile 2 and the other tiles decreases as flow rate is increased and the fluid temperature at the outlet drops. These results agree reasonably well with those from simulations. The measured FPGA temperatures were approximately $3\,^{\circ}\mathrm{C}$ to $4\,^{\circ}\mathrm{C}$ higher than the simulated temperatures of the region above the FPGA and the measured transceiver temperatures were approximately $6\,^{\circ}\mathrm{C}$ to $8\,^{\circ}\mathrm{C}$ higher than the simulated temperatures of the regions above the transceivers. It is expected that the die temperatures are higher than those at the base of the heat sink (where temperature was simulated) due to the heat conduction through the TIM. Reducing the TIM bond line thickness may be an area for further temperature improvement.

Pressure drop, measured across the heat sink with no power

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCPMT.2019.2930481, IEEE Transactions on Components, Packaging and Manufacturing Technology

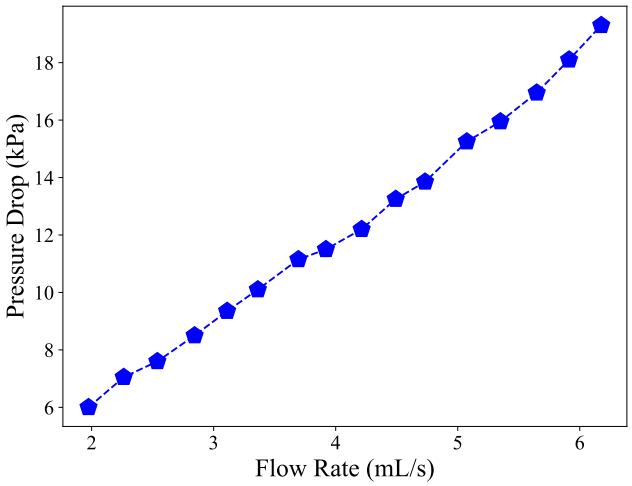

Fig. 19: Heterogeneous micropin-fin heat sink pressure drop vs. flow rate

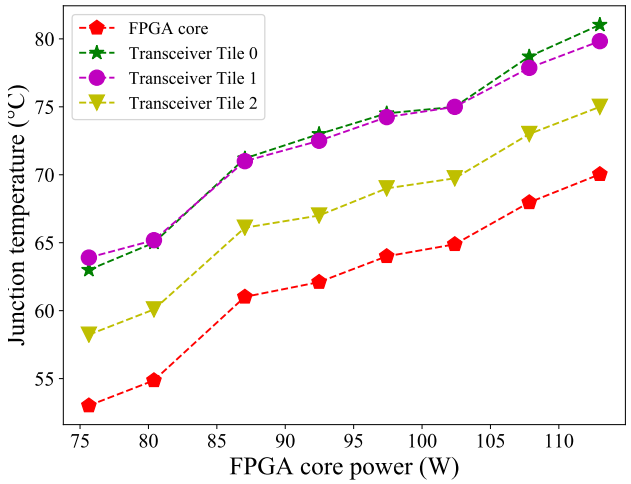

Fig. 20: Air cooled heat sink die temperatures vs. FPGA core power

applied to the FPGA, can be seen in Fig. 19. The measurement includes the inlet/outlet ports and short lengths of tubing in the fluid path.

An additional experiment was performed in which the power on the FPGA die was modulated by varying the clock frequency to the FFT computational blocks. The results of the experiment with air cooling can be seen in Fig. 20. As expected, the temperature of the FPGA die increases with increasing power, with a slope of approximately 0.46 °C/W, computed from the first and last points in the plot. However, it can also be seen that the slopes of the transceiver die temperatures are all approximately equal to the slope of the FPGA die temperature. From the lowest FPGA die power of 75.6 W, corresponding to an FFT clock frequency of 300 MHz, to the highest FPGA die power of 113 W, corresponding to 475 MHz, the temperature rose 17.0 °C degrees on the FPGA die while the temperatures of transceiver tiles 0, 1, and 2 rose by 18.0 °C, 15.9 °C, and 16.8 °C, respectively. This indicates strong thermal coupling between adjacent dice, where the conditions on the FPGA die have a large effect on the temperatures of the surrounding dice.

Similarly, the transceiver power was varied by modulating the data rates of the transceivers while the FPGA core power was kept constant. Due to the details of the FPGA design, varying the transceiver data rates also affected the FPGA die power. Therefore, the FFT clock rate was modulated to keep FPGA die power within 1.7% of 111 W for all three transceiver powers shown in Fig. 21, thus isolating the effect of transceiver power on die temperatures. The transceiver temperatures fell by 17.5 °C to 25 °C when the power was reduced from 69.37 W to 25.55 W. This change in transceiver power, while holding FPGA die power nearly constant, led to a change in FPGA die temperature of 10 °C. Hence there is strong thermal coupling from the FPGA die to surrounding dice (Fig. 20) and from surrounding dice to the FPGA (Fig. 21). This strong thermal coupling is likely due to the very thick heat spreaders, both integrated into the lid of the die, and the base of the air cooled

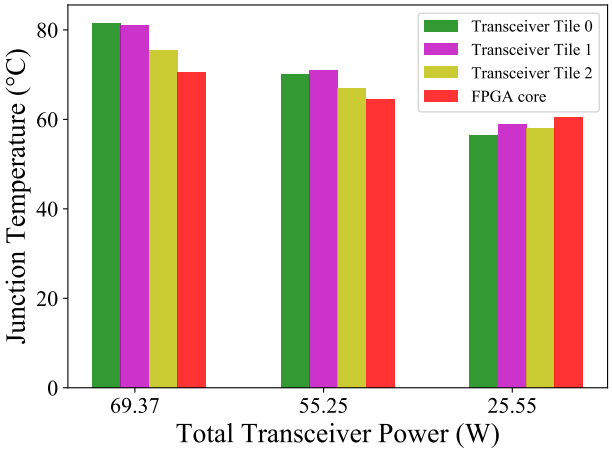

Fig. 21: Air cooled heat sink die temperatures vs. transceiver power

heat sink.

The power of the FPGA die was similarly varied with the microfluidic cooled heat sink with a flow rate of 6.18 mL/s and an inlet temperature of 19.5 ± 0.3 °C. The results can be seen in Fig. 22. The FPGA die temperature increased with increasing power, with a slope of approximately 0.057 °C/W, which is significantly lower than the slope seen with the air cooled heat sink because of the comparatively lower thermal resistance of the microfluidic heat sink. While the temperature of the transceiver tiles with air cooling closely followed the temperature of the FPGA die, the temperatures of transceiver tiles 0 and 1 are nearly constant with microfluidic cooling as the FPGA die power changes from 77.5 W to 117.4 W. The average Tile 2 temperature measurement increases by 1.7 °C because it is located close to the outlet and likely receives more fluid which has been warmed by the FPGA die.

Transceiver powers were similarly varied with the microfluidic heat sink with a flow rate of 6.19 mL/s and inlet temper-

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCPMT.2019.2930481, IEEE Transactions on Components, Packaging and Manufacturing Technology

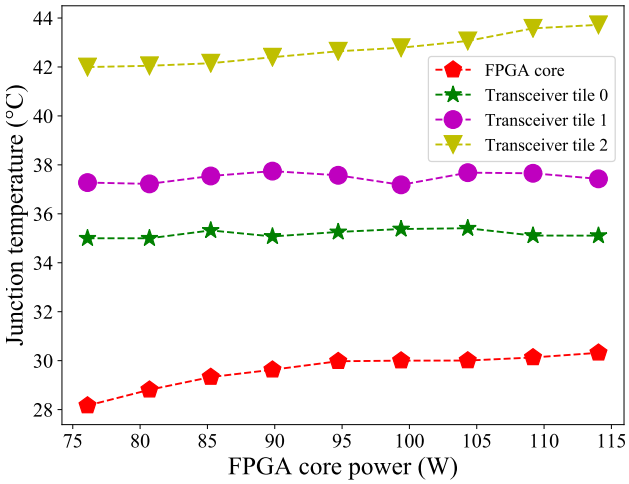

Fig. 22: Heterogeneous micropin-fin heat sink die temperatures vs. FPGA core power

TABLE II: Thermal coupling between dice: change in die temperatures as a function of FPGA die power

Die	Air	Microfluidic
FPGA	0.46 °C/W	0.057 °C/W
Transciever Tile 0	0.48 °C/W	0.003 °C/W
Transciever Tile 1	0.43 °C/W	0.004 °C/W
Transciever Tile 2	0.45 °C/W	0.045 °C/W

ature of 19.2 ± 0.2 °C. The results can be seen in Fig. 23. While transceiver temperatures dropped by 10.1 °C to 14.6 °C when total transceiver power was changed from 63.34 W to 22.61 W, the average temperature measurement on the FPGA die only changed by 0.13 °C. The effects of FPGA die power on the steady state temperatures of all four dice with both the air cooled heat sink and microfluidic cooled heat sink are summarized in Table II. The effect of the FPGA power on the surrounding transceiver die temperatures was reduced by a factor of $10\times$ to over $100\times$ when compared with the air cooled heat sink.

This significant reduction in thermal coupling between adjacent dice is likely a result of the significantly reduced thickness of the microfluidic heat sink. Heat is rapidly transferred to the fluid and extracted, whereas heat must conduct through a large shared thermal mass before being extracted by the air cooled heat sink. This also has an effect on the thermal time constant of the two systems. It was observed that the time necessary for temperature to reach a steady state with the air cooled heat sink was approximately 10 min to 15 min, but only 1 s to 3 s with the microfluidic heat sink.

The temperature measurements corresponding to the baseline design with 475 MHz FFT clock and transceivers at 22 Gbit/s are summarized in Table III for both the air cooled heat sink and the microfluidic heat sink at the highest flow rate. The temperatures of all four dice with the microfluidic cooled heat sink were lower than those with the air cooled heat sink. For reference, the results of an additional test with

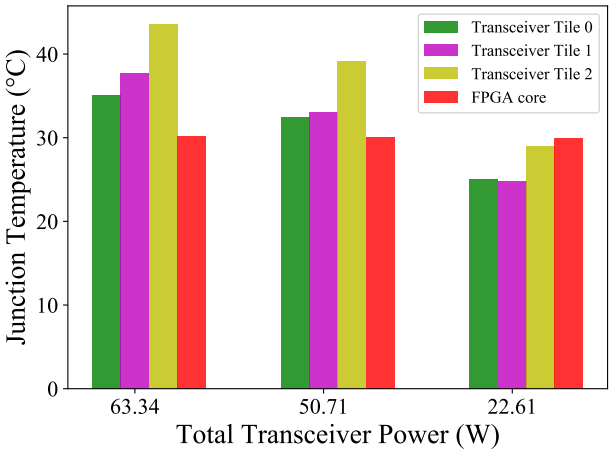

Fig. 23: Heterogeneous micropin-fin heat sink die temperatures vs. transceiver power

TABLE III: Stratix 10 temperature measurements with air cooled heat sink and microfluidic cooled heat sink

Die	Air	Liquid	Air with no Compute
FPGA	70.5 °C	30.1 °C	38.5 °C
Transciever Tile 0	81.5 °C	35.1 °C	48.5 °C
Transciever Tile 1	81 °C	37.7 °C	52 °C
Transciever Tile 2	75.5 °C	43.6 °C	46 °C

the air cooled heat sink can also be seen in Table III where all 160 FFT compute cores were deactivated, with the transceivers remaining on. It can be seen that to reach temperatures similar to those of the microfluidic-cooled dice, all FFT computation needed to be shut down.

In addition to the benefits of low temperatures and low thermal coupling, the microfluidic heat sink offers the potential for very high density computing. The height of the air cooled heat sink used for baseline measurements was 172.2 mm, while the height of the microfluidic cooled heat sink is 6.5 mm. If the boards were to be stacked at high density, the pitch of the boards with the microfluidic heat sink would not be limited by the height of the heat sink, but rather by other connectors on the board, which give the board a total height of approximately 17.4 mm. This would still allow the boards to be mounted at approximately $10\times$ the density of boards with the air cooled heat sink. Lower profile air cooled heat sinks are available, but using them would likely require further sacrifices in other areas, such as die temperature.

Improvements in total power draw were also realized through reduced temperatures provided by the microfluidic cooled heat sink. When running the baseline design with the air cooled heat sink, a total power of 182.4 W was measured. With the microfluidic cooled heat sink, this power was 172.5 W. Power dissipation on the FPGA die was 3.5% higher with air cooling while power on the transceiver dice was 9.6% higher, likely because the highest temperatures (and therefore leakage powers) were experienced on the transceiver dice with air cooling.

VIII. Conclusion

In this work, a microfluidic heat sink was designed, fabricated, and tested for the cooling of a 2.5D Stratix 10 FPGA. Novel aspects of the heat sink include 1) The 3D-printed enclosure used for fluid delivery, 2) the heterogeneous silicon micropin-fin heat sink, and 3) the low profile of the entire heat sink assembly, including fluid delivery. These features enabled improvements in 1) die temperatures, 2) compute density, and 3) thermal coupling between adjacent dice in a 2.5D package.

Heat transfer was achieved with silicon micropin-fins with a height of approximately 460 µm and non-uniform density tailored to the non-uniform heat fluxes of the underlying dice. The micropin-fins were enclosed in a 3D-printed enclosure giving the entire heat sink a height of 6.5 mm. Results were compared with a high performance air cooled heat sink and a temperature reduction ranging from 31.9 °C to 46.4 °C was observed across the four active dice under test. It was also found that thermal coupling between the FPGA die and surrounding transceiver dice was very strong with the air cooled configuration. With the microfluidic cooled heat sink this effect was reduced on the transceiver tile near the outlet and nearly eliminated on the other transceiver tiles. In addition to the reduction in temperature and thermal coupling between dice, a large increase in compute density was also enabled through the use of the low profile microfluidic heat sink.

Although the temperatures of all of the dice were reduced relative to experiments with the air cooled heat sink, the transceiver tiles remained warmer than the FPGA die. Fluid may have bypassed these higher density regions, either through the gap between the tops of the micropin-fins and the cover, or through the lower density micropin-fins in the center region of the heat sink. Future work could more precisely control the temperatures of each die by separately delivering fluid to each region of the heat sink through the plastic manifold.

IX. Acknowledgements

We would like to thank DARPA for their support of this work under grant FA8650-16-1-7674. We would also like to thank Dan Pritsker, Deepika Ganju, Dale Elson, and Rajsaktish Sankaranarayanan from Intel for their assistance with the Stratix 10 platform.

This research was supported in part through research cyber infrastructure resources and services provided by the Partnership for an Advanced Computing Environment (PACE) at the Georgia Institute of Technology, Atlanta, Georgia, USA. This work was performed in part at the Georgia Tech Institute for Electronics and Nanotechnology, a member of the National Nanotechnology Coordinated Infrastructure, which is supported by the National Science Foundation (Grant ECCS-1542174).

References

[1] B. Banijamali, S. Ramalingam, H. Liu, *et al.*, "Outstanding and innovative reliability study of 3d TSV interposer and fine pitch solder micro-bumps," in *2012 IEEE 62nd Electronic Components and Technology Conference*, May 2012, pp. 309–314. DOI: 10.1109/ECTC.2012.6248847.

[2] D. Greenhill, R. Ho, D. Lewis, *et al.*, "3.3 a 14nm 1ghz fpga with 2.5d transceiver integration," in *2017 IEEE International Solid-State Circuits Conference (ISSCC)*, 2017, pp. 54–55. DOI: 10.1109/ISSCC.2017.7870257.

[3] D. Foley and J. Danskin, "Ultra-performance pascal gpu and nvlink interconnect," *IEEE Micro*, vol. 37, no. 2, pp. 7–17, 2017, ISSN: 0272-1732. DOI: 10.1109/MM.2017.37.

[4] J. Macri, "Amd's next generation gpu and high bandwidth memory architecture: Fury," in *2015 IEEE Hot Chips 27 Symposium (HCS)*, 2015, pp. 1–26. DOI: 10.1109/HOTCHIPS.2015.7477461.

[5] Y. Zhang, T. E. Sarvey, and M. S. Bakir, "Thermal evaluation of 2.5-d integration using bridge-chip technology: Challenges and opportunities," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 7, no. 7, pp. 1101–1110, 2017, ISSN: 2156-3950. DOI: 10.1109/TCPMT.2017.2710042.

[6] T. E. Sarvey, Y. Hu, C. E. Green, *et al.*, "Integrated circuit cooling using heterogeneous micropin-fin arrays for nonuniform power maps," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 7, no. 9, pp. 1465–1475, 2017, ISSN: 2156-3950. DOI: 10.1109/TCPMT.2017.2704525.

[7] D. Lorenzini, C. Green, T. E. Sarvey, *et al.*, "Embedded single phase microfluidic thermal management for nonuniform heating and hotspots using microgaps with variable pin fin clustering," *Int. J. Heat and Mass Transfer*, vol. 103, pp. 1359 –1370, 2016, ISSN: 0017-9310.

[8] P. Schepers, *Stratix10 devkit 72ch nios embedded prbs*, Apr. 5, 2017. [Online]. Available: https://fpgawiki.intel.com/wiki/File:Stratix10_Devkit_72Ch_Nios_Embedded_Prbs.zip.

[9] E. Lemmon, M. McLinden, and D. Friend, *Thermophysical Properties of Fluid Systems*, ser. NIST Chemistry WebBook, NIST Standard Reference Database Number 69. Gaithersburg MD, 20899: National Institute of Standards and Technology. DOI: 10.18434/T4D303. (visited on 06/2018).

[10] Cooler Master, *Masterair ma621p tr4 edition*. [Online]. Available: https://coolermaster.egnyte.com/dl/mVCD4IBr0j/ (visited on 06/29/2019).

Thomas E. Sarvey received B.S. degrees in physics and computer engineering from the University of Maryland, College Park in 2012. He received a Ph.D. degree in electrical and computer engineering from the Georgia Institute of Technology in Atlanta, GA in 2018, where his research focused on densely integrated, 2.5D and 3D electronic systems and their enabling thermal technologies. He is currently a staff CPU performance engineer with Futurewei Technologies in Austin Texas.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCPMT.2019.2930481, IEEE Transactions on Components, Packaging and Manufacturing Technology

Ankit Kaul received the B.E. degree in electrical and electronics engineering from Rashtreeya Vidyalaya College of Engineering (RVCE), Bengaluru, India, in 2013. He received the M.S. degree in electrical and computer engineering from Georgia Institute of Technology, Atlanta, GA, USA (2018), where he is currently pursuing the Ph.D. degree in electrical and computer engineering. His current research interests include modeling and experimental evaluation of thermal challenges in novel heterogeneous integration technologies.

Sreejith Kochupurackal Rajan (S 18) received the B. Tech degree in electronics and communication engineering from Cochin University of Science and Technology, Kerala, India in 2013 and M.S. degree in electrical and computer engineering from Georgia Institute of Technology, Atlanta, GA, U.S.A in 2016, where he is currently pursuing the Ph.D. degree in electrical and computer engineering.

His current research interests include the study of interconnection and thermal challenges in advanced heterogenous integration architectures.

Aravind Dasu is a Principal Engineer and the Director of Strategic Planning for Advanced Technology and Innovation at the Programmable Solutions Group of Intel Corporation. He has a PhD from Arizona State University (2004) in Electrical Engineering. Over the past twenty years, his research has spanned a range of technologies in FPGAs and reconfigurable computing.

Ravi Gutala is a senior member of technical staff at Intel in the Programmable Solutions Group. He has over 25 years of experience in the semiconductor industry and performed pioneering work in high endurance flash memory design. He currently has 18 granted patents in the area of circuit/memory design.

Muhannad S. Bakir (SM'12) is a Professor in the School of Electrical and Computer Engineering at Georgia Tech. Dr. Bakir and his research group have received thirty paper awards including six from the IEEE Electronic Components and Technology Conference (ECTC), four from the IEEE International Interconnect Technology Conference (IITC), and one from the IEEE Custom Integrated Circuits Conference (CICC). Dr. Bakirs group was awarded 2014 and 2017 Best Papers of the IEEE Transactions on Components Packaging and Manufacturing Technology (TCPMT). Dr. Bakir is the recipient of the 2013 Intel Early Career Faculty Honor Award, 2012 DARPA Young Faculty Award, 2011 IEEE CPMT Society Outstanding Young Engineer Award, and was an Invited Participant in the 2012 National Academy of Engineering Frontiers of Engineering Symposium. He is also a recipient of the 2018 IEEE Electronics Packaging Society (EPS) Exceptional Technical Achievement Award "for contributions to 2.5D and 3D IC heterogeneous integration, with focus on interconnect technologies." Dr. Bakir is also the co-recipient of the 2018 McKnight Foundation Technological Innovations in Neuroscience Awards.

Dr. Bakir serves on the editorial board of IEEE Transactions on Components, Packaging and Manufacturing Technology (TCPMT) and IEEE Transactions on Electron Devices (TED). Dr. Bakir serves as a Distinguished Lecturer for IEEE EPS.