

# THERMAL CONTROL IN 3D LIQUID COOLED PROCESSORS VIA HOTSPOT SEPARATION AND THERMOELECTRIC COOLING

Yue Hu, Shaoming Chen, Lu Peng, Edward Song, and Jin-Woo Choi

Division of Electrical & Computer Engineering, School of Electrical Engineering and Computer Science,  
Louisiana State University, Baton Rouge, LA, USA

## ABSTRACT

*Microchannel liquid cooling is a promising technique to handling the high temperature problem of three-dimensional (3D) processors. There have been a few works which made initial attempts to optimize liquid cooling by utilizing non-uniformly distributed channels, variable flow rate, wider channels, and Dynamic Voltage and Frequency Scaling (DVFS) combined with thread migration mechanisms. Although these optimizations could be better than a straightforward microchannel liquid cooling design, the cooling of 3D processors is limited due to design-time and run-time challenges. Moreover, in new technologies, the processor power density is continually increasing and this will bring more serious challenges to liquid cooling.*

*In this paper, we propose two thermal control techniques to control hotspots in liquid cooled 3D processors: 1) Core Vertically Placed (CVP) technique. According to the architecture of a processor core, two schemes are given for placing a core vertically onto multilayers. The 3D processor with the CVP technique can be better cooled since its separate hotspot blocks have a larger contact area with the cooler surroundings. 2) Thermoelectric cooling (TEC) technique. We propose to incorporate the TEC technique into the liquid-cooled 3D processor to enhance the cooling of hotspots. Our experiments show the CVP technique reduces the maximum temperature up to 29.58 °C, and 13.77 °C on average compared with the baseline design. Moreover, the TEC technique effectively cools down a hotspot from 96.86 °C to 78.60 °C. Furthermore, the CVP technique supports a 30% increase in processor frequency which results in a 1.27 times speedup of processor performance.*

## KEYWORDS

*3D processors, Core Vertical Placed (CVP), Liquid cooling, Thermoelectric Cooling (TEC).*

## 1. INTRODUCTION

Three-dimensional (3D) integration has been proposed to reduce the communication delay and interconnect power in processors [16]. Previous studies [7][26] demonstrate that significant performance speedup can be achieved and the total power consumption is reduced compared with a 2D design. However, the vertical integration of multiple circuit layers exacerbates existing high-temperature problems. It is well known that extended overheating in a processor damages its




Figure 1: Illustration of a microchannel liquid-cooled 3D processor

reliability and eventually shortens its lifetime. Therefore, eliminating hotspots remains one of the main challenges in 3D processor design.

Figure 1 illustrates the micro channel liquid cooling to a 3D processor. Heat mainly generates from active circuit layers, and the dissipated heat is predominantly removed through the coolant in micro channels. With regard to micro channel liquid cooling, IBM has developed a test vehicle on which up to 300 W/cm<sup>2</sup> of heat flux can be removed experimentally [12]. Although this result is inspiring, the cooling of actual 3D processors is limited to some extent due to four challenges:

- (1) The effect of liquid cooling itself is non-uniform. For instance, the cooling gets weaker along the direction of liquid flow since the liquid is gradually heating up.
- (2) The cooling demand varies significantly inside an actual processor. One example is that the cooling demand for the register file is much stronger than the one for the last level cache, since the power density of the register file is tens of times that of the last level cache.
- (3) The thermal resistance that exists between the active circuit layer and the coolant is one of the main bottlenecks for effective cooling.
- (4) The liquid flow rate, in practice, is limited for two reasons. First, the microchannel is relatively fragile, since the width of a channel wall is only about 100  $\mu\text{m}$ . Because a higher flow rate requires a comparatively higher pressure drop between the inlet and outlet, the microchannels may be destroyed. Second, pump power increases exponentially with the increase of flow rate. Therefore, an abnormally high flow rate could result in an unaffordable pump power.

In this paper, we propose two thermal control techniques to enhance the cooling of the liquid-cooled 3D processors:

- **CVP technique:** The CVP technique is to divide the circuits of each component (e.g., L1 data cache) in a processor core into several blocks and place them vertically onto multilayers. This technique is effective for challenges (3) and (4). First, the CVP technique reduces thermal resistance, since the divided and vertically placed component has a relatively large contact area with the horizontal surroundings. Second, a relatively low flow rate can achieve the same cooling effect, since the CVP technique helps cooling. To facilitate design, test, and verification, two schemes are given for different components in a processor core depending on the architecture of the components.

- **TEC technique:** A thermoelectric cooler works as a heatpump, absorbing heat from one side and dissipating it on the other using electric energy. The TEC technique is to incorporate thermoelectric coolers into the liquid-cooled 3D processor, so that they absorb heat from the active circuit layer (i.e. heat source) and release it to the microchannels. Since thermoelectric coolers can be activated to cool down hotspots, it is effective for challenges (1) and (2) by selectively cooling down high temperature areas.

We evaluated the effectiveness of the CVP and TEC techniques by running groups of benchmarks on 3D multi-core processors. According to power and thermal simulations, the CVP 3D processor reduces the maximum temperature by up to 29.58 °C for benchmark 256.BZIP2, and 13.77 °C on average compared with a conventional liquid-cooled 3D multi-core processor. After incorporating the TEC technique into the 3D processor, it effectively cools down a hotspot from 96.86 °C to 78.60 °C while running benchmark 256.BZIP2.

The remainder of this paper is organized as follows: Section 2 provides related work. Section 3 and section 4 describe the proposed CVP technique and the TEC technique individually. Section 5 explains the experimental methodology. We analyze experimental results in section 6 and present our conclusion in section 7.

## **2. RELATED WORK**

### **2.1. 3D Stacked Chips**

Three dimensional (3D) integration is a recently developed chip manufacturing technology which stacks multiple circuit layers into the same chip and connect different layers through vertical interconnect tunneling [21][24][25]. Compared with conventional two dimensional (2D) layouts, the 3D chips have a few advantages: (1) better performance due to the reduced interconnect length and latency; (2) lower power consumption for the interconnect because of the reduction of wiring length; (3) higher packing density due to the third dimension. Combined with the current trend of chip-multiprocessing, 3D integration proved significant performance advantage can be achieved over 2D designs. The multiple circuit layers in a 3D chip include processor cores, caches, interconnects and routers. Recently, researchers also integrate DRAM [23] and new generation of memories such as MRAM [32] and STT-RAM [25] into 3D chips.

One of the main challenges for a 3D chip is its high temperature resulting from accumulative heat dissipated from circuits which easily generates hotspots in the processors, especially in the middle layers where it is difficult for conventional cooling air to reach. To avoid a high peak temperature, processor designers utilizes mitigating methods such as lowering executing frequency [14][15] and throttling a hotspot core's execution. These technologies inevitably reduce a 3D chip's performance. On the other hand, extended overheating for a chip will hurt its reliability and eventually shorten its lifetime. Conventional thermal management in 3D chips include DVFS [20], job scheduling/migration [27][1], thermal aware floorplan placement [13][17][19][22]. Although these techniques mitigate 3D chip hotspots in a certain degree, the effectiveness with a conventional air cooling system will be limited for increasingly complicated and overheating 3D processors.

## 2.2. Microchannel liquid cooling

Very recently, there have been a few works that make initial attempts to optimize microchannel liquid cooling. Their optimizations can be divided into static and dynamic optimizations.

- **Static optimizations:** Static optimizations are applied at the time of design, such as non-uniform distributed channels [29], thermal balancing using channel modulation [28], and channel dimension studies [9][10] that aim to provide optimal cooling. Different from these works, the proposed CVP technique optimizes the liquid cooling from the circuit level of 3D processors. This also means the former static optimizations can be combined with our proposed CVP technique to further enhance microchannel liquid cooling.

- **Dynamic optimizations:** Dynamic optimizations consist of liquid flow rate adjustment [14], dynamic voltage and frequency scaling (DVFS) [20], thread migration [1], and the combinations [15][27] of these techniques. Compared with former optimizations, the proposed TEC technique has significant advantages in terms of performance overhead, energy efficiency of liquid cooling, and response time. First, the DVFS and the thread migration have obvious performance overhead. This is prohibitive in a 3D processor where performance is still the top priority in most cases. Second, after flow rate is increased over a certain value, the cooling effect slowly increases; in contrast, pump power increases exponentially. Moreover, increasing flow rate provides enhanced cooling to the whole 3D processor, while the hotspot is usually a very small localized area in an actual 3D processor. This decreases the energy efficiency of flow rate adjustment. On the other hand, thermoelectric coolers can be selectively activated to cool the localized hotspots. Third, the response time of thermoelectric coolers is only about 5  $\mu$ s [5], which is much shorter than the hundreds of milliseconds in flow rate adjustment. In fact, the proposed TEC design can be combined with the flow rate adjustment, to provide a fine-tuned control aims at optimal cooling in terms of both the cooling effect and the energy efficiency.

## 3. CORE VERTICALLY PLACED (CVP) TECHNIQUE

In this section, we will discuss the CVP technique in terms of both the circuit level placement schemes and thermal analysis.

### 3.1. Placement Schemes

Figure 2(a) shows the layout of a conventional 3D eight-core processor that is cooled using four microchannel layers. We compared this layout with the one in which core layers and cache layers overlap each other, and found that this layout, which was also used in [27], is the optimal layout for liquid cooling for all processor cores placed near the cooler liquid inlet. Therefore, we used this layout as a baseline for later comparisons. Figure 2(c) shows the floorplan of a processor core in Figure 2(a).

The layout of the CVP 3D processor is shown in Figure 2(b). As can be observed, a single processor core is divided into nearly equal parts and placed vertically onto four layers, in contrast with the baseline in Figure 2(a). A processor core is further divided into several components, as shown in Figure 2(c). Circuits of each component, such as the register file and the multiplication unit, are divided into nearly equal parts and placed vertically onto four layers. Depending on the

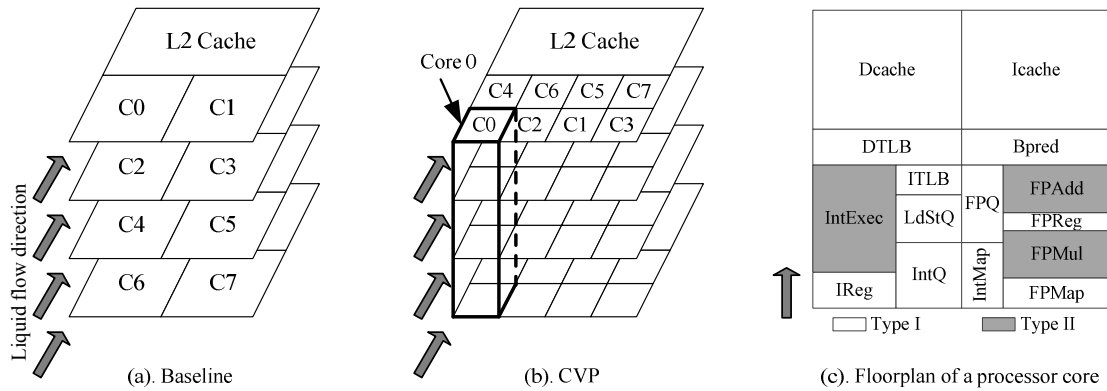


Figure 2 : Layouts of 3D eight-core processors.

Connections between the circuits, all components fall into two types. The circuits of type I components are relatively loosely connected, such as the register file. On the other hand, in type II components, which consist of complex arithmetic and logic units, such as the multiplication unit, circuits are relatively tightly connected with wires. The type I and type II components are displayed individually with white and gray backgrounds in Figure 2(c).

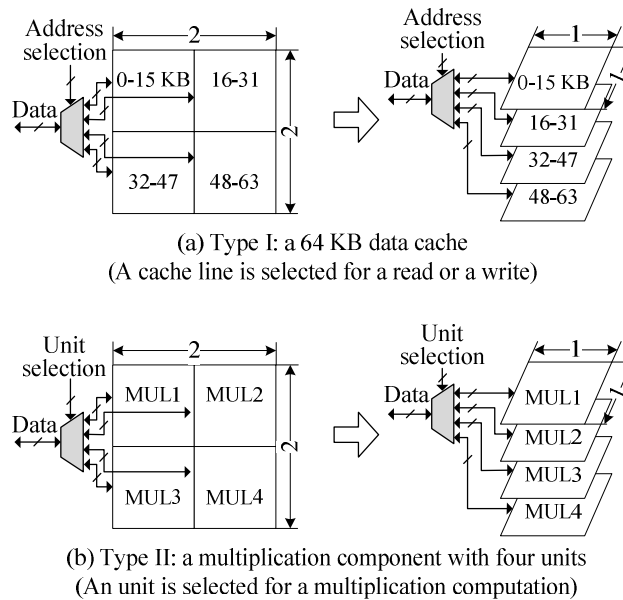


Figure 3: Two schemes for the CVP technique.  
(The left part is for the baseline, and the right one is for the CVP technique)

Figure 3 illustrates the two vertical placement schemes for the CVP technique at the circuit level. The type I components are loosely connected, because they are mainly made up of static random-access memory (SRAM). These components can be divided according to the SRAM size. Figure 3(a) gives such an example for a 64 KB data cache. As can be observed, a quarter of the data cache (i.e. 16 KB) is placed vertically onto one of the four layers. For each operation, a cache line is selected for a read or a write depending on the address selection. On the other hand, the type II

components are tightly connected. Therefore, to facilitate the test and verification of the integrated circuit (IC), these components are divided by units. Figure 3(b) shows such an example for a multiplication component with four units. As can be observed, each multiplication unit is placed vertically onto a separate layer in the CVP design. Depending on the unit selection, a multiplication unit is selected for a multiplication computation.

As we can see, for both type I and type II components, the divided circuits are still functional. For example, the divided 16 KB data cache can be correctly read or written with an input and output port; the multiplication unit can still execute a multiplication computation after being placed vertically. Therefore, the CVP 3D processor will not bring much complexity to the test and verification steps of the IC design. In fact, to further facilitate test and verification, we can design an extra input and output port for each divided component, so that each divided small core on a single layer can be tested and verified independently.

### 3.2. Thermal Analysis




Figure 4: Thermal illustration of the CVP technique.

Figure 4 shows two 3D processors with each circuit layer (i.e. heat source) being cooled by the liquid layer underneath. Figure 4(a) shows a baseline design in which a hotspot block (i.e. a component) is located in the center of the first circuit layer. Heat that dissipates from the hotspot component flows horizontally to the comparatively cooler surroundings and vertically to the liquid layers. Note that the heat flow to the liquid layers is not shown in the figure. Besides, the thermal maps shown in Figure 4 are just for illustration, instead of being generated from actual thermal simulations.

Figure 4(b) shows a 3D processor using the CVP technique. This scheme is the same as that illustrated in Figure 2(b). Compared with the baseline, the whole hotspot block is equally divided into four smaller blocks, and each of them is placed vertically onto a separate circuit layer. In the vertical direction, each hotspot block is equally cooled by one liquid layer underneath in both the baseline and CVP design. However, horizontally, the vertically placed smaller blocks have a larger total contact area with the cooler surroundings than that of the baseline. Therefore, 3D processors that use the CVP technique can be better cooled.




Figure 5: Illustration of the TEC technique.

#### 4. THERMOELETRIC COOLING (TEC) TECHNIQUE

The TEC technique is based on the thermoelectric effect which creates a temperature difference using electric energy. Similar to a heat pump, it absorbs heat from one side and dissipates it on the other. It has been reported that the TEC device is able to sustain a heat flux up to 1250 W/cm<sup>2</sup> [11] with a response time of about 5 μs [5]. This makes it suitable for the effective and rapid cooling of hotspots.

A recent study demonstrates that using the TEC device as an active cooling method is an effective approach to mitigating the hotspots of an integrated circuit [6]. Moreover, the feasibility of incorporating the TEC device into a 3D IC has already been demonstrated [18]. It can be especially useful to the 3D processors when the heat is difficult to remove. Specifically, in microchannel liquid cooling, the TEC device can be used to absorb heat from the active circuit layers and release it to the channel layers.

The typical structure of a thermoelectric cooler is shown in Figure 5(a). When a voltage is applied to the thermoelectric cooler, two electric fields are generated along the N-type and the P-type semiconductors in opposite directions. Therefore, thermal energy will flow downward in both semiconductors. Thus, heat is absorbed from one side and dissipated on the other using electric energy.

Figure 5(b) shows the cross-section of a liquid-cooled 3D processor using the TEC technique. The thermoelectric cooler layer is located in between the active circuit layer, which is the main heat source, and the channel layer. For the safety of the integrated circuit, an electrical insulator layer is applied to separate the thermoelectric cooler layer from the circuits and channels as shown in the figure. One suitable candidate for the electrical insulator is chemical vapor deposition (CVD) diamond [2].

Thermoelectric coolers are activated to improve cooling when the localized hotspot is unable to be cooled down by liquid cooling. On the other hand, they may worsen cooling slightly when inactive. This is because the thermoelectric cooler layer is located between the active circuit layer and the channel layer, and it has a relatively low thermal conductivity, as listed in Table 2.

## 5. EXPERIMENTAL METHODOLOGY

Our experimental analysis includes performance, power consumption, and overall thermal impact.

We used SESC [4], a cycle accurate architectural simulator to build the performance model for 3D eight-core processors running with SPLASH2, SPEC2000 and SPEC2006 benchmarks. For SPLASH2 benchmarks, we ran them on an eight-thread version; while for SPEC2000 and SPEC2006 benchmarks, we ran multi-programmed benchmarks. Each processor core, whose floorplan is shown in Figure 2(c), is based on the Alpha 21264 microarchitecture. Table 1 lists the main microarchitecture parameters.

The power model consists of leakage power and dynamic power. The leakage power, whose density is assumed to be 10 W/cm<sup>2</sup> at a normal temperature 25 °C, is proportional to temperature [6]. On the other hand, the dynamic power is modeled from two input sources. One is the access rate of each component, such as the intensity of the reads and writes to the register files, which is obtained from the performance model. The other is the energy it consumes for each operation, such as an integer addition operation. In addition, the energy parameters of the core-related and the L2 Cache-related were obtained individually from Wattch [8] and CACTI [3].

For our thermal analysis, we used the two-resistance model of 3D-ICE 2.0 [30][31]. 3D-ICE, which stands for 3D Interlayer Cooling Emulator, is a Linux based open source Thermal Emulator Library that can perform both steady and transient thermal analysis to the microchannel liquid cooled 3D integrated circuits. We incorporated the thermoelectric coolers (TECs) model into the 3D-ICE. The TECs' power, which is equal to QTEC per unit time, is modeled as:

$$Q_{TEC} = \frac{Q_{pump}}{COP} \quad (1)$$

Parameters	Values
Technology	32 nm
Voltage	1.1 V
Frequency	3.0 GHz
Fetch / Issue/ Commit Width	4/ 4/ 5
INT/ FP Window Size	96/ 64
LoadStore/ INT/ FP Units	2/ 2/ 3
Load/ Store Queue Size	80/ 80
Latency of INT ALU/ Mult/ Div	1/ 4/ 12 cycles
Latency of FP ALU/ Mult/ Div	1/ 2/ 10 cycles
L1 Instruction/ Data Cache Size	64/ 64 KB
L1 Instruction/ Data Cache	8/ 8
L1 Instruction/ Data Block Size	64/ 64 B
L2 Cache Size	16 MB
L2 Cache Associativity	16
L2 Cache Block Size	64 B

Table 1: Performance and power parameters



Thermal Parameters			
Material	Conductivity(W/(m·K))		Heat Capacity(m <sup>3</sup> ·K)
Silicon	130		1635660
BEOL	2.25		2174502
CVD Diamond	1800		1764530
TEC	1.2		1153750
Channel Walls	160		1641101
Water	0.6		4172000
Coolant Heat Transfer Coefficient: 37132.3W/( m <sup>2</sup> ·K)			
Dimension parameters (μm)			
Active layer thickness	2	BEOL layer	12
CVD Diamond layer thickness	5	TEC layer thickness	10
Channel layer thickness	10	Silicon slab thickness	50
Channel width	50	Channel wall width	100
Others			
Length/Width of 3D processor at each layer (mm)			10/ 10
Liquid Flow Rate Per Channel Layer (mL/min)			56.371
Pressure Drop of The Pump (Pa)			2.0e5
Inlet water temperature (K)			300.00

Table 2:Thermal simulation parameters.

QTEC denotes the TECs energy it consumes to pump an amount of heat  $Q_{pump}$  from one side on the other. The efficiency of TECs is termed as coefficient of performance (COP). For example, for COP of 3, it takes 1 W power to pump 3 W of heat. We made the conservative assumption that the maximum heat flux that TECs can pump is 1000 W/cm<sup>2</sup> and COP is equal to 3. In addition, we designed a dynamic and smooth algorithm to control the working intensity of TECs. TECs are disabled when a hotspot is cooler than a certain threshold temperature. Otherwise, TECs are enabled and their working intensity smoothly varies depending on the difference between the hotspot temperature and the threshold. The reliability and lifetime of processors will degrade exponentially when the hotspot temperature is higher than a safe temperature. We conservatively assumed the safe temperature to be 80 °C. Moreover, to ensure efficient hotspot control, the threshold temperature was set to be 2 °C lower than the safe temperature. The thermal parameters are listed in Table 2.

## 6. EXPERIMENTAL RESULTS

### 6.1. Core vertical placed technique

Figure 6 shows the maximum temperature comparisons between the baseline and the CVP design. The maximum temperature of the baseline design varies from 35.63 °C (470.LBM) to 96.86 °C (256.BZIP2), which indicates the selected benchmarks are diverse in their thermal behaviors. In the figure, the CVP design reduces the maximum temperature up to 29.58 °C (256.BZIP2), and 13.77 °C on average compared with the baseline. The significant temperature reduction is due to two factors. First, as illustrated in Figure 4, the divided hotspot units have a larger total contact area with the cooler surroundings; thus heat can dissipate easily. Second, such cooling improvement is stronger when the maximum temperature is high in the baseline. On the other hand, benchmark 181.MCF, whose IPC per core is 0.24, has a negligible temperature reduction of

3.61 °C due to its comparatively lower temperature in the baseline. Other benchmarks are in between these two extremes. Overall, the CVP technique effectively mitigates hotspot temperatures for all benchmarks. In addition, such cooling improvement is stronger when the maximum temperature is comparatively higher in the baseline.




Figure 6: Maximum temperature comparisons between the baseline and the CVP design.

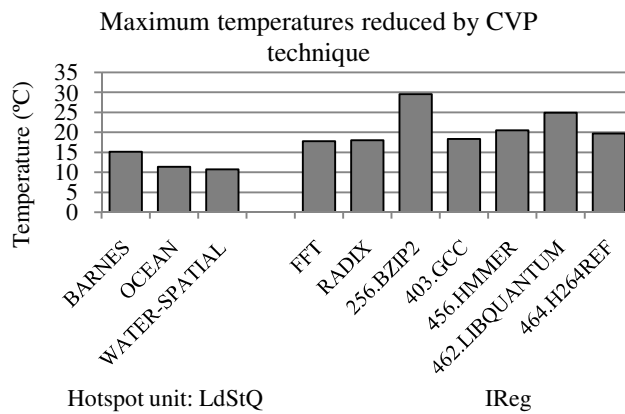


Figure 7: Cooling effectiveness of the CVP technique for different hotspot units

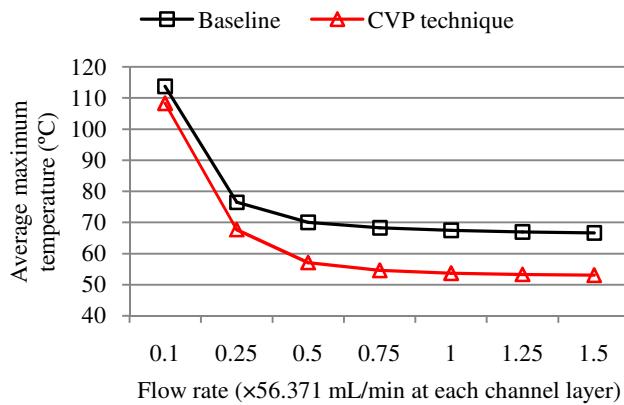


Figure 8: Cooling effectiveness of the CVP technique under different liquid flow rates

From Figure 6, we also observe that the CVP technique has different cooling effects on different benchmarks. One of the reasons is the cooling effectiveness also depends on the locations of hotspots in a processor. Figure 7 shows the cooling effectiveness, i.e. the maximum temperatures reduced by the CVP technique. To highlight the effectiveness of the CVP technique, only a portion of benchmarks whose maximum temperature is higher than 70 °C in the baseline design are shown. Obviously, the CVP technique is more effective in cooling on benchmarks whose hotspot units are IReg compared with benchmarks whose hotspot units are LdStQ. This is because the IReg component has a relatively worse heat conduction to the horizontal surroundings. Specifically, in Figure 2(c), the IReg component is located at the corner of a core where heat can only be dissipated to the up and right directions, while the LdStQ component has a horizontal contact in all four directions. The CVP technique improves cooling mainly through increasing a hotspot's horizontal contact with the cooler surroundings; thus, a higher temperature reduction is observed for the hotspot components (i.e. IReg in this study) which have a comparatively worse heat conduction with the horizontal surroundings.

In addition, the cooling effectiveness of the CVP technique is also observed to depend on the liquid flow rate of microchannel liquid cooling. Figure 8 shows the average maximum temperatures of all the benchmarks with varying liquid flow rates. Overall, the average maximum temperature decreases with the increase of liquid flow rate for both the baseline and the CVP technique. Furthermore, the temperature reduced by the CVP technique (i.e. the temperature difference between the baseline and the CVP technique) under the same liquid flow rate is obviously proportional to the flow rate until it is higher than  $(0.75 \times 56.371 \text{ mL/min at each channel layer})$ . This is because a relatively low flow rate results in relatively low vertical heat


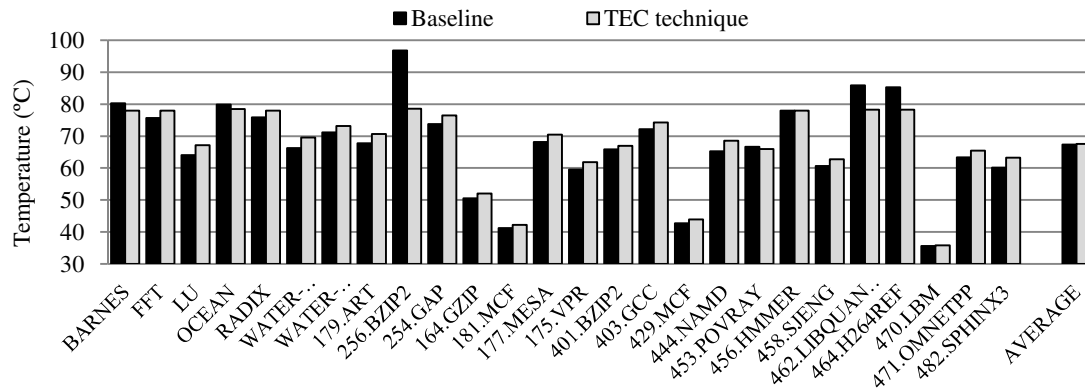



Figure 9: Example (256.BZIP2) that illustrates the cooling effect for the CVP and the TEC technique.

conduction from the hotspot units to the micro channels. This makes the vertical instead of the horizontal heat conduction become a bottleneck in heat dissipation. Thus, the CVP technique which mainly increases the horizontal heat conduction does not help much in lowering the hotspots' temperatures when liquid flow rate is relatively low. On the other hand, when flow rate is relatively high, such as the flow rates higher than  $(0.75 \times 56.371 \text{ mL/min at each channel layer})$ , the horizontal instead of the vertical heat conduction is the bottleneck. In this case, a relatively large temperature reduction can be observed with the CVP technique.

To further understand the variations in hotspot mitigation, we investigated the benchmarks' thermal behaviors and determined the exact time when and the exact hotspot component where the maximum temperature occurred. Figure 9(a) shows the instructions per cycle (IPC) and the heat flux of a hotspot component in real time for benchmark 256.BZIP2, in which the hotspot component is the integer register file during the whole run-time. As can be observed, the curves of the IPC and the heat flux have strong coherence. The comparatively higher heat flux between



(a) Maximum temperature comparisons



(b) Average TEC power

Figure 10: Cooling effect of the TEC technique.

56 and 77 milliseconds is caused by the frequent accesses to the integer register file. Figure 9(b) shows the real-time maximum temperature comparison between the baseline and the CVP design for benchmark 256.BZIP2. Although the curve of the CVP design is similar to that of the baseline, the temperature of the CVP design is obviously lower. For example, the CVP technique reduces the maximum temperature from 96.86 °C to 67.28 °C at the peak temperature point.

## 6.2. Thermoelectric cooling technique

In the TEC technique, thermoelectric coolers are applied between the active circuit layer, which is the main heat source, and the channel layer to remove the localized hotspot as illustrated in Figure 5(b). We conservatively assume the safe temperature of processors to be 80 °C. The TEC technique aims to control processors' temperature to be lower than the safe temperature so that the processor reliability and lifetime can be guaranteed.

As can be observed in Figure 10(a), overall, the TEC technique has different effect in cooling compared with the cooling of CVP shown in Figure 6. For instance, for benchmark 256.BZIP2, the TEC technique decreases temperature from 96.86 °C to 78.60 °C compared with the baseline. On the other hand, the TEC technique increases the temperature from 73.77 °C in the baseline, to 76.48 °C for benchmark 254.GAP. To further study the different cooling effect caused by the TEC technique, we divide all benchmarks into two types. In the baseline, the benchmarks whose maximum temperature is higher than the safe temperature are denoted as HITEMP benchmarks. In contrast, the benchmarks whose maximum temperatures are lower than the safe temperature are denoted as LOTEMP benchmarks. As can be observed, the TEC technique decreases the maximum temperature for all HITEMP benchmarks compared with the baseline. For all LOTEMP benchmarks, the TEC technique increases the maximum temperatures by about 2 °C, but the temperatures are still lower than the safe temperature. Overall, with the TEC technique, the maximum temperatures are lower than the safe temperature for all benchmarks.

For all HITEMP benchmarks, the TEC technique helps cooling. This is because thermoelectric coolers are allowed to enhance the cooling to hotspots once hotspot temperature is detected to be higher than the threshold (78 °C) that is 2 °C lower than the safe temperature in the experiments. The thermoelectric coolers will remain active until the hotspot temperature is cooled to be lower than the threshold temperature. The TEC power consumed is shown in Figure 10(b). The TEC technique significantly improves the reliability and lifetime of processors which decrease exponentially with the increasing of the maximum temperature.

On the other hand, for all LOTEMP benchmarks, the TEC technique increases the maximum temperatures by about 2 °C. This is because a thermoelectric cooler layer is located between the active circuit layer and the channel layer as shown in Figure 5(b). The thermoelectric cooler layer, which has a relatively low thermal conductivity as listed in Table 2, increases the thermal resistance between the active circuit layer and the channel layer. Therefore, the comparatively higher thermal resistance in the TEC technique increases temperatures for all LOTEMP benchmarks. However, the effect on the reliability and lifetime is negligible since the temperature is always maintained lower than the safe temperature for all LOTEMP benchmarks.

In addition, as can be observed in Figure 9(b), during the time interval of 59-78 milliseconds, the TEC technique reduces the maximum temperature to lower than 80 °C compared with the baseline design. A temperature reduction, as high as 18.26 °C, is observed at 72 milliseconds. In contrast, the TEC technique slightly increases the temperature when the temperature in the baseline is lower than 80 °C.

Figure 11 shows the average TEC power of all the benchmarks when the maximum temperature is cooled to be lower than different safe temperatures. As can be observed, TEC power is higher for a lower safe temperature. This is because for a comparatively lower safe temperature, thermoelectric coolers need to be enabled in more benchmarks, in longer time for a benchmark, and the thermoelectric coolers in a larger area needs to be enabled at a certain time.

To further demonstrate the effectiveness of the proposed CVP and TEC techniques, Figure 12 shows the thermal maps of the top layer of the 3D processors, which are the hottest layers, when the maximum temperature occurs for benchmark 256.BZIP2. As can be seen, the CVP technique significantly reduces the maximum temperature from 96.86 °C to 67.28 °C by dispersing hotspot blocks. Moreover, the overall temperature of the CVP design is lower compared with the

baseline. The thermal map of the TEC design is quite similar to that of the baseline design, but the hotspots are effectively removed by the TEC technique, resulting in an 18.26 °C decrease in the maximum temperature.




Figure 12: Thermal maps of 3D processors.

(The numbers inside brackets are the maximum temperatures in the thermal maps. This figure is more readable in color.)




Figure 13: Performance speedup.

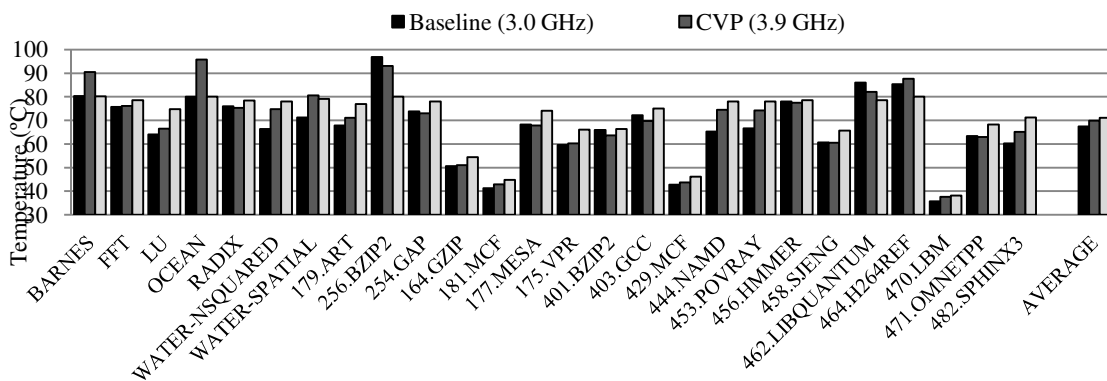


Figure 14: Maximum temperature comparisons.

### 6.3. CVP+TEC

To explore the cooling potential of the combination of the proposed two techniques, the frequency of processor cores and L2 Cache has been increased from 3.0 GHz to 3.9 GHz through DVFS. Figure 13 presents performance speedup after DVFS. As this figure shows, 30% of frequency increase (i.e. 3.0 GHz to 3.9 GHz) results in a 1.27 times speedup in geometric mean. The reason that some benchmarks have a relatively low speedup is DVFS only increases the frequency of processor cores and L2 Cache, instead of main memory. For example, benchmark 429.MCF has a relatively low speed up, i.e. 1.10 is because it is a memory-intensive benchmark whose memory access rate is 6.10 GB/s. Figure 14 shows the maximum temperature comparisons among the baseline design, the CVP design, and the (CVP+TEC) design. The processor frequency in the baseline design is 3.0 GHz, while it is increased to 3.9 GHz in both the CVP and the (CVP+TEC) design. The maximum temperature of the CVP design (i.e. 95.80 °C) is lower than that of the baseline design (i.e. 96.86 °C). In other words, the CVP technique can support at least a 30% increase in frequency with the same cooling effect. Moreover, the (CVP+TEC) design effectively controls all temperatures to be lower than 80 °C.

## 7. CONCLUSION

In this paper, we proposed two thermal control techniques for 3D processors with microchannel liquid cooling. First, the CVP technique, which is applied during the design phase, is evaluated to be able to significantly improve the cooling effect by dividing hotspot blocks and placing them vertically onto multilayers. To facilitate design, test, and verification of the integrated circuit design, two schemes are given for the detailed vertical placement. Second, the TEC technique is a run-time cooling control technique. We proposed to incorporate the TEC technique into the liquid-cooled 3D processor and make use of it to absorb heat from the active circuit layer and dissipate it on the channel layer, so that the cooling to the localized hotspot can be enhanced. Our experiments demonstrate that the CVP technique reduces the maximum temperature up to 29.58 °C, and 13.77 °C on average compared with the baseline. Moreover, the TEC technique effectively cools down a hotspot from 96.86 °C to 78.60 °C. In the future, we will consider to test more benchmarks including emerging big data applications for the proposed thermal control techniques.

## REFERENCES

- [1] <http://www.ansys.com/products/fluid-dynamics/cfx>.
- [2] [http://www.diamond-materials.com/downloads/cvd\\_diamond\\_booklet.pdf](http://www.diamond-materials.com/downloads/cvd_diamond_booklet.pdf).
- [3] <http://www.hpl.hp.com/research/cacti>.
- [4] <http://sesc.sourceforge.net>.
- [5] A. Bar-Cohen and P. Wang. On-Chip Thermal Management and Hot-Spot Remediation. In Nano-Bio-Electronic, Photonic and MEMS Packaging 2010.
- [6] S. Biswas, M. Tiwari, T. Sherwood, L. Theogarajan, and F. T. Chong. Fighting Fire with Fire: Modeling the Datacenter-Scale Effects of Targeted Superlattice Thermal Management. In International Symposium on Computer Architecture (ISCA) 2011.
- [7] B. Black, M. Annavaram, N. Brekelbaum, J. DeVale, L. Jiang, G. H. Loh, D. McCauley, P. Morrow, D. W. Nelson, D. Pantuso, P. Reed, J. Rupley, S. Shankar, J. Shen, and C. Webb. Die Stacking (3D) Microarchitecture. In International Symposium on Microarchitecture (MICRO) 2006.
- [8] D. Brooks, V. Tiwari, and M. Martonosi. Wattch: A Framework for Architectural-Level Power Analysis and Optimizations. In ISCA 2000.

- [9] T. Brunswiler, B. Michel, H. Rothuizen, U. Kloter, B. Wunderle, H. Oppermann, and H. Reichl. Interlayer cooling potential in vertically integrated packages. In *Microsystem Technology 2009*.
- [10] T. Brunswiler, S. Paredes, U. Drechsler, and B. Michel. Heat-removal performance scaling of interlayer cooled chip stacks. In *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm) 2010*.
- [11] I. Chowdhury, R. Prasher, K. Lofgreen, G. Chrysler, S. Narasimhan, R. Mahajan, D. Koester, R. Alley, and R. Venkatasubramanian. On-chip cooling by superlattice-based thin-film thermoelectrics. In *Nature Nanotechnology 2009*.
- [12] E. G. Colgan, B. Furman, M. Gaynes, W. S. Graham, N. C. LaBianca, J. H. Magerlein, R. J. Polastre, M. B. Rothwell, R. J. Bezama, R. Choudhary, K. C. Marston, H. Toy, J. Wakil, J. A. Zitz, and R. R. Schmidt. A Practical Implementation of Silicon Microchannel Coolers for High Power Chips. In *IEEE Transaction on Components and Packing Technologies 2007*.
- [13] J. Cong, J. Wei, Y. Zhang. A Thermal-Driven Floorplanning Algorithm for 3D ICs. In *IEEE/ACM International conference on Computer-aided design (ICCAD) 2004*.
- [14] A. K. Coskun, D. Atienza, T. S. Rosing, T. Brunswiler and B. Michel. Energy-Efficient Variable-Flow Liquid Cooling in 3D Stacked Architectures. In *Design, Automation & Test in Europe Conference & Exhibition (DATE) 2010*.
- [15] A. K. Coskun, J. Meng, D. Atienza, and M. M. Sabry. Attaining Single-Chip, High-Performance Computing through 3D Systems with Active Cooling. In *IEEE Micro 2011*.
- [16] S. Das, A. Fan, K. Chen, C. S. Tan, N. Checka, and R. Reif. Technology, Performance, and Computer-Aided Design of Three-Dimensional Integrated Circuits. In *International Symposium on Physical Design (ISPD) 2004*.
- [17] M. B. Healy, M. Vittes, M. Ekpanyapong, C. Ballapuram, S. K. Lim, H. S. Lee and G. H. Loh. Multi-Objective Microarchitectural Floorplanning For 2D and 3D ICs. In *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, Vol. 26, No. 1, 2007*.
- [18] L. L. Hsu, P. Wang, X. Wei, and H. Zhu. Thermoelectric 3D Cooling. US Patent. Patent Number: US 8,030,113 B2. Date of Patent: 10/04/2011.
- [19] W. Hung, G. Link, Y. Xie, N. Vijaykrishnan and M. J. Irwin. Interconnect and Thermal-aware Floorplanning for 3D Microprocessors. In *IEEE International Symposium on Quality Electronic Design (ISQED) 2006*.
- [20] W. Kim, M. S. Gupta, G. Wei, and D. Brooks. System Level Analysis of Fast, Per-Core DVFS using On-Chip Switching Regulators. In *International Symposium on High-Performance Computer Architecture (HPCA) 2008*.
- [21] F. Li, C. Nicopoulos, T. Richardson, Y. Xie, N. Vijaykrishnan, M. Kandemir. Design and Management of 3D Chip Multiprocessors using Network-in-memory. In *ISCA 2006*.
- [22] X. Li, Y. Ma and X. Hong. A Novel Thermal Optimization Flow Using Incremental Floorplanning For 3D ICs. In *Asia and South Pacific Design Automation Conference (ASP-DAC) 2009*.
- [23] G. H. Loh. 3D-Stacked Memory Architectures for Multi-core Processors. In *ISCA 2008*.
- [24] G. L. Loi, B. Agrawal, N. Srivastava, S. Lin, T. Sherwood and K. Banerjee. A Thermal-Aware Performance Analysis of Vertically Integrated (3D) Processor-Memory Hierarchy. In *Design Automation Conference (DAC) 2006*.
- [25] A. K. Mishra, X. Dong, G. Sun, Y. Xie, N. Vijaykrishnan and C. R. Das. Architecting On-Chip Interconnects for Stacked 3D STT-RAM caches in CMPs. In *ISCA 2011*.
- [26] K. Puttaswamy and G. H. Loh. Implementing Caches in a 3D Technology for High Performance Processors. In *International Conference on Computer Design (ICCD) 2005*.
- [27] M. M. Sabry, A. K. Coskun, D. Atienza, T. S. Rosing, and T. Brunswiler. Energy-Efficient Multiobjective Thermal Control for Liquid-Cooled 3-D Stacked Architectures. In *IEEE Transaction on Computer-Aided Design of Integrated Circuits and Systems, 2011*.
- [28] M. M. Sabry, A. Sridhar, and D. Atienza. Thermal Balancing of Liquid-Cooled 3D-MPSoCs Using Channel Modulation. In *DATE 2012*.
- [29] B. Shi, A. Srivastava, and P. Wang. Non-Uniform Micro-Channel Design for Stacked 3D-ICs. In *DAC 2011*.
- [30] A. Sridhar, A. Vincenzi, M. Ruggiero, T. Brunswiler, and D. Atienza. Compact Transient Thermal



- Model for 3D ICs with Liquid Cooling via Enhanced Heat Transfer Cavity Geometries. In Thermal Investigations of ICs and Systems (THERMINIC) 2010.
- [31] A. Sridhar, A. Vincenzi, M. Ruggiero, T. Brunswiler, and D. Atienza. 3D-ICE: Fast Compact Transient Thermal Modeling for 3D ICs with Inter-tier Liquid Cooling. In ICCAD 2010.
- [32] G. Sun, X. Dong, Y. Xie, J. Li and Y. Chen. A Novel Architecture of the 3D Stacked MRAM L2 Cache for CMPs. In HPCA 2009.
- [33] X. Zhou, J. Yang, Y. Xu, Y. Zhang, and J. Zhao. Thermal-Aware Task Scheduling for 3D Multicore Processors. In IEEE Transaction on Parallel and Distributed Systems, 2010.