

Co-Design of Multicore Architectures and Microfluidic Cooling for 3D Stacked ICs

Zhimin Wan, He Xiao, Yogendra Joshi*, Sudhakar Yalamanchili

Georgia Institute of Technology, Atlanta, USA

* Corresponding Author: yogendra.joshi@me.gatech.edu, +1 404 385 2810

Abstract

In this paper, we investigate the co-design of multicore architectures and microfluidic cooling for 3D stacked ICs. The architecture is a 16 core, x86 multicore die stacked with a second die hosting an L2 SRAM cache. First, a multicore x86 compatible cycle-level microarchitecture simulator was constructed and integrated with physical power models. The simulator executes benchmark programs to create power traces that drive thermal analysis. Second, the thermal characteristics under liquid cooling were investigated using a compact thermal model. Four alternative packaging organizations were studied and compared. Greatest overall temperature reduction is achieved under a given pumping power, with two tiers and two microgaps with the high power dissipation tier on the top. Third, an optimization of the pin fin parameters including the diameter, height, and longitudinal and transversal spacing was performed. This optimization is shown to achieve up to 40% improvement in energy/instruction and significant reductions in leakage power.

1 Introduction

Three-dimensional (3D) stacked ICs as an emerging technology have many advantages. Compared with planar ICs which place all the devices on the same plane, 3D stacked ICs integrate multiple devices in the vertical direction using through-silicon-via (TSV) technology. This vertical integration could reduce the global wire length by as much as 50% [1]. The wire limited clock frequency could be increased nearly four-fold [2]. Further, 3D stacked ICs enable heterogeneous layer integration.

However, as the level of integration continues to increase, one of the main challenges in architecting 3D systems is the heat dissipation capacity. Several contributing factors arise [3]: (1) the package heat flux increases due to reduced surface area per unit volume, (2) the interior chips in the 3D IC package could be overheated due to longer heat conduction paths incorporating insulating dielectric layers, and (3) heat dissipation non-uniformity on the chip resulting in hotspot heat fluxes 10 times or more larger than the average levels [4].

Liquid cooling using pin fin enhanced micro-gaps could be a viable solution to the increasing thermal challenges of 3D ICs due to its high heat transfer coefficients. Zhang et al. [5] fabricated an inter-tier pin fin enhanced microgap and showed that a staggered pin-fin heat sink is able to provide a thermal resistance as low as $0.27 \text{ K}\cdot\text{cm}^2/\text{W}$ with a flow rate of $70 \text{ mL}/\text{min}$ for a heat sink depth of $200 \mu\text{m}$. Jaspersen et al. [6] compared micro pin fin and micro channel heat sinks. Their results show that micro pin fin heat sink has a lower convection thermal resistance at liquid flow rates above approximately $60 \text{ g}/\text{min}$, with a higher pressure drop. Wan et al. [7, 8] built a compact model of 3D stacked ICs with inter-tier pin fin enhanced microgap under non-uniform heat flux. The maximum temperature was at the hotspot of the bottom tier. With the pin fin enhanced microgap cooling, the

maximum temperature could be maintained at $56 \text{ }^\circ\text{C}$. Ndao et al. [9] studied in-line and staggered circular pin fin-heat sinks, as well as offset strip fin heat sinks, and found that the latter outperforms the other cooling technologies. Bejan and Morega [10] reported the optimal geometry of an array of fins that minimizes the thermal resistance between the substrate and the flow forced through the fins, and found that the minimum thermal resistance of plate-fin arrays is approximately half of the minimum thermal resistance of heat sinks with continuous fins.

The architecture floor plan determines the non-uniform heat dissipation on the chip and the resulting thermal characteristics in turn affect the electrical performance. The above work either used uniform heat dissipation or did not consider the thermal-electric interaction of the 3D ICs. In this study, we investigated the co-design of architecture floorplans and inter-tier pin fin enhanced microgap for 3D stacked ICs using a compact 3D IC thermal model. First, an x86 compatible microarchitecture simulator was built to obtain execution profile information, which served as the input to a coupled power model to generate non-uniform power traces. Secondly, the power traces were input to the compact thermal model. Different pin fin enhanced microgaps were studied and the maximum temperature obtained. Two tiers and two microgaps with higher power tier on the top shows best thermal performance. Third, the compact thermal model was linked to an optimization tool constructed in Matlab. An optimization of the pin fin parameters including the diameter (D_p), height (H_p), and longitudinal (S_L) and transversal spacing (S_T) was performed using a genetic algorithm. It was found that the large pin fin dimensions are better for non-uniform heat dissipation, while smaller pin dimensions should be used for non-uniform heat dissipation with high heat flux hotspots. Finally, the improvement in electrical performance was analyzed and it was found that substantial savings in leakage power could be obtained after optimization.

2 System model

2.1 Power model

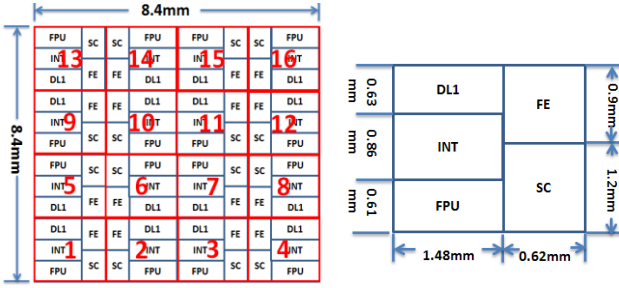


Figure 1: Floor plan for logic tier

We model a 16 core, x86 processor, each with its own L1 cache and all cores sharing a banked, coherent L2 cache interconnected by a 2D mesh interconnect. The simulation model is a cycle-level timing model that is driven by a multicore emulator front-end that boots a linux operating system and executes compiled 32-bit x86 binaries. The goal of this infrastructure is to generate timing, energy, and power behaviors that are as close as possible to commodity processors. The floor plan used in this study is shown in Figure 1. The 16 cores are placed on the 8.4mm x 8.4mm chip. Every core consists of five modules: Frontend (FE), scheduler (SC), integer unit (INT), floating point unit (FPU) and memory (DL1). The L2 cache consists of 16 equal sized L2 cache banks arrayed on a 8.4mm x 8.4 mm die. Each L2 bank has a 1 Mbyte capacity. This floor plan was generated using the McPAT [11] modeling library using publicly available information about commodity x86 processors.

Simulations are run for 500M clock cycles to warm up the processor state and reach a “region of interest” in the benchmark program. This is a region wherein the computational characteristics are representative, since they primarily avoid operating system boot code and application startup and initialization code. Once execution has reached the region of interest, the power at each block in the processor floor plan is sampled every 10 microseconds to produce a power trace. Such traces are used to drive the thermal models. In general, we draw upon benchmark programs from the SPLASH and PARSEC benchmark suites. In the specific results reported here, the power traces were generated from the Canneal benchmark in the PARSEC benchmark suite.

The physical model employs various configurations of the logic tier and memory tier with microfluidic cooling. While there are many more configurations and packaging options that could have been explored, we emphasize two main points– i) the methodology for co-design, and ii) demonstration that co-design matters.

A typical power distribution of each module in logic tier is listed in Table 1. The power consumption of core 3 and core 4 is lower than other 14 cores. The power consumption on memory tier is uniform and 30% of that of logic tier based on our test.

Table 1: Power dissipation for each module

	FPU (W)	INT (W)	DL1 (W)	SC (W)	FE (W)
Core 3	0.11	0.25	0.17	0.12	0.11
Core 4	0.11	0.27	0.25	0.13	0.11
Other Cores	1.69	2.06	2.81	2.31	1.63

2.2 3D stacked ICs structure model

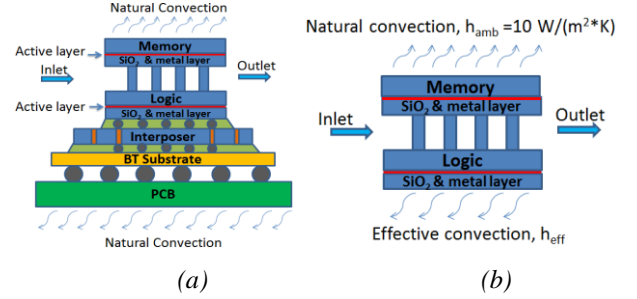


Figure 2: (a) 3D stacked ICs structure (b) Simplified structure

Figure 2(a) shows the notional packaged 3D stacked IC structure considered in the present study. Two tiers, logic and memory, are enclosed in the system. The red region is the active layer in which most of the heat is generated. Below the active layer is the SiO₂ & metal layer used for bonding. Between the two tiers is the pin fin enhanced microgap, incorporating fluid flow. The two tiers are placed on a bismaleimide triazine (BT) substrate through a silicon interposer. The BT substrate is attached to the printed circuit board using solder ball array. The top and bottom of the system are assumed to be natural convection cooled, with heat transfer coefficient of 10 W/(m²*K). Table 2 shows the dimensions and properties of materials.

Table 2: Material dimensions and properties

	Thickn ess (μm)	Length (mm)	Width (mm)	Thermal Conductivity (W/(m ² *K))	
				k_{xy}	k_z
PCB	1600	100	100	56.9	0.36
BT Substrate	950	20	20	13.4	0.21
Interposer	1.69	2.06	2.81	2.31	1.63
Logic/Memory	100	8.4	8.4	149	149
SiO ₂	10	8.4	8.4	1.4	1.4
Solder Ball	D=600 μm , pitch =1000 μm . Between PCB and BT substrate.			0.05	14.1
Micro Bump	D=12 μm , pitch =250 μm . Between interposer and BT substrate.			0.63	0.63
TSV	D=25 μm , pitch=150 μm .			401	401

A simplified structure was proposed in Figure 2(b). An effective heat transfer coefficient is applied on the bottom of SiO₂ & metal layer. To obtain the effective heat transfer coefficient, a finite element (FE) heat conduction model, including the materials from the underfill below the SiO₂ & metal layer to PCB was built (Figure 3). The effective heat transfer coefficient was obtained by:

$$R = \frac{1}{h_{\text{eff}}A_o} = \frac{T_{o,\text{avg}} - T_{\text{amb}}}{Q} \quad (1)$$

Where R is the thermal resistance between the heating surface and the ambient. Also, Q is the total power applied on the surface of underfill (2 W), A_o is the surface area of the oxide layer, T_{o,avg} is the average temperature of the heating surface, and T_{amb} is the ambient temperature (20 °C in the present study).

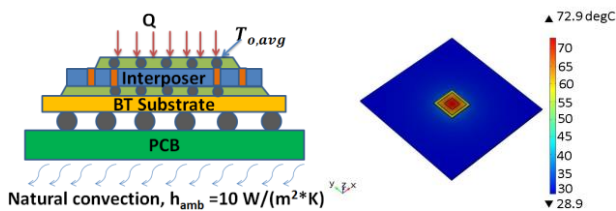


Figure 3: Conduction FE model and temperature result

The average temperature of the heating surface is 70.4 °C. So the effective heat transfer coefficient is 562.4 W/(m²*K).

2.3 Compact thermal model

A compact thermal model [7] was used which discretized the 3D stacked ICs model into multiple control volumes, each around one pin (Figure 4). The metal layer was not included for simplicity. The arrows show the energy flows in the vertical direction within one control volume. The temperature of the active layer was assumed uniform within the control volume. The resulting set of equations was solved simultaneously by iteration.

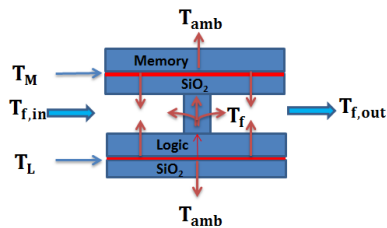


Figure 4: Control volume around one pin

For the prediction of heat transfer coefficient and pressure drop (ΔP), the following correlations [12] were used:

$$f = C^* (H_p/D_p)^{\alpha_1} (S_L/D_p - 1)^{\alpha_2} (S_T/D_p - 1)^{\alpha_3} Re^m \quad (2)$$

Table 3. Coefficient for friction factor

	C	α_1	α_2	α_3	m
Re<100	3.1335	0.4485	0.4965	0.5553	0.6292
Re>100	1.246	0.3362	0.4478	0.4615	0.4393

$$j = C^* (H_p/D_p)^{\alpha_1} (S_L/D_p - 1)^{\alpha_2} (S_T/D_p - 1)^{\alpha_3} Re^m \quad (3)$$

$$j = h / (\rho V_{\text{max}} C_p) * Pr^{2/3} \quad (4)$$

Table 4. Coefficients for colburn j factor correlation

	C	α_1	α_2	α_3	m
Re<100	0.5885	0.0072	0.1432	0.1289	0.5697
Re>100	0.4481	0.1285	0.1707	0.0804	0.4864

In order to validate the compact thermal model, a full computational fluid dynamics and heat transfer (CFD/HT) model was built. D_p was 100 μm , S_L and S_T were 200 μm , and H_p was 300 μm . The chip dimension was 8.4 mm x 8.4 mm. Figure 5 shows the model and boundary conditions used. A symmetric boundary condition was used to simplify the model. The uniform heat dissipation of the active layer of the logic tier was 160 W. The uniform dissipation of active layer of the memory tier was 80 W. The inlet boundary condition was water at 20 °C with inlet velocity 0.58 m/s. The outlet boundary condition was atmospheric pressure. The fluid properties were evaluated at mean fluid temperature.

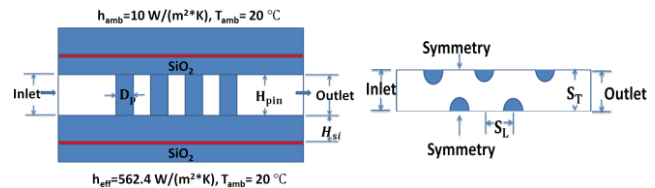


Figure 5: Full CFD/HT model

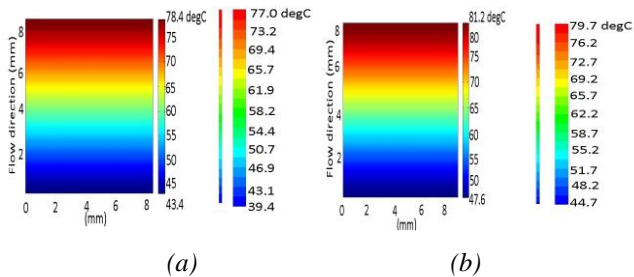
Table 5 shows the mesh independence study of the full CFD/HT model. It shows that when the number of elements increases from 3745k to 4133k, the pressure drop changes 0.09% and the maximum temperature changes less than 0.2%. So 3745k elements were used subsequently.

Table 5: Mesh independence study

Number of elements	ΔP (Pa)	T _{max,logic} (°C)	T _{max,memory} (°C)
2777k	20949	76.87	79.98
3745k	20934	76.96	79.74
4133k	20953	76.95	79.89

Figure 6(a) shows the comparison of the temperature distribution of logic tier between compact thermal model and detailed CFD/HT model. The temperature increases almost linearly along the flow direction due to the uniform heating. The difference in the maximum temperature between the two models was 1.8%, while that in minimum temperature was 9.2%. In the compact thermal model, an average heat transfer coefficient was used for every column of the pin fins. However the heat transfer coefficient at the

inlet of the full CFD/HT model was much higher than the average heat transfer coefficient. This resulted in the detailed CFD/HT model prediction being lower than that of compact model.



(a) Logic tier, compact thermal model/ detailed CFD/HT model

(b) Memory tier, compact thermal model/ detailed CFD/HT

Figure 6: Comparison of temperature distribution of logic tier, memory tier between compact thermal model and detailed CFD/HT model

Figure 6(b) is the comparison of temperature distribution of memory tier between compact thermal model and detailed CFD/HT model. The difference in the maximum temperatures between the two models was 1.8%, while that in minimum temperature was 6.1%. This confirmed the validity of the compact thermal model. The compact thermal model took about 45 seconds to compute, while the detailed CFD/HT model took about 3 hours and 20 min on a Win 7 machine with 3.4 GHz CPU and 8.0 GB memory.

3 3D stacked ICs under realistic power map

The compact thermal model was used to analyze the thermal characteristics of 3D stacked ICs under realistic power map obtained in Section 1. The total pumping power was 0.03 W, water inlet was at 20 °C, and the outlet boundary condition was atmospheric pressure.

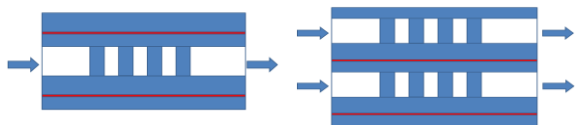


Figure 7: Two types of microgap configuration

Two types of microgap configuration were studied as in Figure 7. The first configuration has only one microgap while the other one has two microgaps. The D_p was 100 μm , S_L and S_T 200 μm , and H_p 200 μm [13].

The first case we studied was one microgap with logic tier at bottom and memory tier on the top. Figure 8 shows the temperature distribution of logic and memory tier for this case. The non-uniform temperature distribution was due to the non-uniform heat dissipation. The maximum temperature of the logic tier was 93.1 °C at the outlet. Every core has the same power distribution, except cores 3 and 4. Due to the bulk fluid temperature rise, the maximum temperature was at the outlet. Although the uniform heat

dissipation of memory tier was only 30% of that of the logic tier, the temperature distribution was non-uniform and the maximum temperature of the memory tier reached 82.2 °C because of the cross-tier heat conduction. The pressure drop ΔP was 27.3 kPa and the mass flow rate \dot{m} was 1.1 g/s for this case.

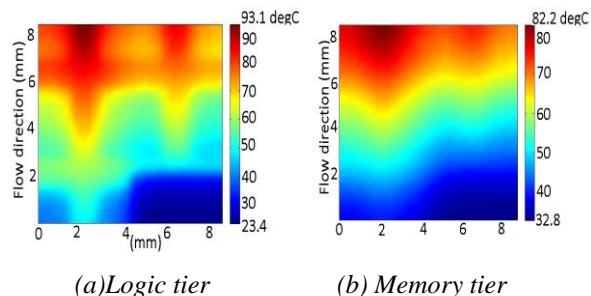


Figure 8: Temperature distribution of logic and memory tier for case 1

Table 6: Results for 4 cases

	Configuration	$T_{\text{max,logic}}$ (°C)	$T_{\text{max,memory}}$ (°C)
Case 1	One microgap with logic tier at bottom and memory tier on the top	93.1	82.2
Case 2	One microgap with logic tier on the top and memory tier at the bottom	114.9	77.1
Case 3	Two microgaps with logic tier at bottom and memory tier on the top	87.7	54.8
Case 4	Two microgaps with logic tier on the top and memory tier at the bottom	72.7	58.3

Table 6 shows the maximum temperature of the logic and memory tier for 4 different cases. Compared with case 1, the temperature of logic tier in case 2 was 114.9 °C, which was higher than that of case 1. The temperature of memory tier was 77.1 °C, which was lower than that of case 1. This is because the thermal conductivity of silicon is higher than that of silicon oxide. So for two tiers with one microgap configuration, the logic tier with high heat dissipation should be placed at the bottom. For case 3, the maximum temperature of logic tier was reduced by about 5.4 °C, and that of the memory tier was reduced by about 27.4 °C compared with case 1. The pumping power is determined as the product of the pressure drop and volume flow rate. When the pumping power of each microgap was reduced to half of case 1, both the pressure drop and volume flow rate were reduced. Thus, volume flow rate of each microgap should be larger than half of case 1. The total mass flow rate for case 3 is 1.6 g/s - higher than case 1. Therefore, the bulk fluid temperature rise was lower than case 1 and the

maximum temperature of case 3 should be lower than case 1. Further, the pressure drop was reduced to 18.3 kPa. Thus the two tiers with two microgaps was superior, both in thermal and hydraulic performance. Compared with previous cases, the maximum temperature of the logic tier in case 4 was further reduced to 72.7 °C, since the logic tier has microgaps below and above. The temperature of the memory tier was slightly increased, because it had only one sided microgap cooling compared with case 3.

The above four cases show that the two tiers with two microgaps, and high heat dissipation tier with double side microgap cooling has the best thermal performance.

4 Optimization

In the above cases, the pin fin dimensions were fixed. In this section, the compact thermal model was linked to the Matlab optimization tool box. A genetic algorithm was used to find an optimized pin fin structure, which minimized the maximum temperature of the logic tier for the configuration of case 4. The heat dissipation for logic and memory tiers is the same as before, with pumping power fixed at 0.03 W. The optimization range of pin fin diameters was 100 μm ~ 200 μm ; The range of ratio of longitudinal spacing to pin diameter was 1.5 ~ 2.25; The range of ratio of transversal spacing to pin diameter was 1.5 ~ 2.25; and the range of ratio of pin height to pin diameter is 1 ~ 3. Water was used as coolant, with inlet temperature at 20 °C.

Table 7: Optimization results for non-uniform heat dissipation without hotspot.

	D_p (μm)	S_L (μm)	S_T (μm)	H_p (μm)	\dot{m} (g/s)	ΔP (kPa)	$T_{\text{max,logic}}$ (°C)
1	100	200	200	200	1.6	18.3	72.7
2	100	150	150	243	1.2	25.5	81.6
3	100	200	200	100	1.0	29.7	92.3
4(optimized)	194	290	420	400	3.4	8.9	57.6

Table 7 shows the optimization results. The pin fin dimensions 1, 2, 3 were selected from literature [13, 14]. Although there are more pin fins for the smaller pin fin dimensions, the mass flow rate was smaller and pressure drop higher due to the higher flow resistance, which lead to a larger bulk fluid temperature rise. Therefore, the maximum temperature of logic tier was higher. The low flow resistance for larger pin fin dimensions increases the mass flow rate significantly, and the bulk fluid temperature rise was reduced. So the maximum temperature of logic tier was smaller.

The previous optimization was for non-uniform heat dissipation without hotspot. Next, the power of DL1 module in the 7th core was increased to 28.1 W, while other modules remained the same. So, the maximum temperature was now

at this hotspot. The optimized pin fin dimensions in Table 8 show that for the non-uniform heat dissipation with hotspot, smaller pin fin dimensions achieved better thermal performance. Compared with large pin fin dimensions, which produced large mass flow rate and small bulk fluid temperature rise, smaller pin fin dimensions resulted in more convection surface area. While this overall result is expected, the present analysis provides a quantitative definition of the optimized design.

Table 8: Optimization results for non-uniform heat dissipation with hotspot.

	D_p (μm)	S_L (μm)	S_T (μm)	H_p (μm)	\dot{m} (g/s)	ΔP (kPa)	$T_{\text{max,logic}}$ (°C)
1	100	200	200	200	1.6	18.2	149.5
2	100	150	150	243	1.2	25.2	136.1
3	100	200	200	100	1.0	29.4	158.5
4	194	290	420	400	3.4	8.8	151.9
5(optimized)	116	175	175	349	1.6	18.3	131.1

5 Electrical performance analysis

To calculate the impact on energy efficiency gained by optimizing the pin fin structure, we measure the energy per instruction (EPI) of the microprocessor under the previous 4 pin fin dimensions in Table 7, which records the average energy spent to process one single instruction during execution. Specifically, we tracked the energy consumption of our 16 core microarchitecture for 200 million cycles under the 4 different pin fin configurations, and computed the EPI based on the execution information and the McPAT power and energy models [11]. The 16 core microarchitectures is configured in a 16 nm process, with each core running at 3GHz under a supply voltage 1.0 V. The test case is *barnes* from SPLASH2 benchmark. Figure 9 gives the normalized EPI (relative to the worst case).

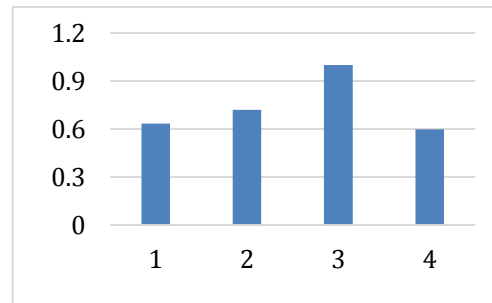


Figure 9: Energy per Instruction Comparison among all 4 pin fin structures

Given that optimized pin fin structure provides best thermal dissipation, its EPI has the lowest value, indicating that the energy saved throughout execution is 40% over the worst case as depicted in Figure 9. As a result of temperature drop due to the optimized pin fin structure, the leakage power of

each component in both the logic and memory tiers thus decreases accordingly. The data for leakage power is shown in Table 9.

Table 9: Predicted Leakage power of the 16 core microarchitecture under different pin fin organizations

	Logic Leakage (W)	Memory Leakage (W)	Total Leakage (W)
1	5.85	6.09	11.94
2	9.22	9.64	18.86
3	20.01	20.90	40.91
4(optimized)	4.51	4.70	9.21

As listed in Table 9, it is always benefit to replace the microarchitecture with the optimized pin fin structure, as it will dramatically reduce the leakage power and keep system EPI small.

6 Conclusions

In this paper a co-design of architecture floorplans and inter-tier microgap with pin fin for 3D stacked IC were studied.

- The configuration of two tiers and two microgaps with the high power dissipation tier under double side cooling shows the best thermal performance.
- For non-uniform power dissipation without hotspot, large pin fin dimensions which produce larger mass flow rate are better than small pin fin dimensions.
- For non-uniform power dissipation with hotspot, small pin fin dimensions which produce larger convection surface area are better than large pin fin dimensions.
- The optimized pin fin dimensions could save energy dramatically in terms of leakage power.

Acknowledgments

The authors gratefully acknowledge the support of Sandia National Laboratories and the National Science Foundation under grant CNS-855110. Discussions with Dr. Muhannad Bakir, Yue Zhang and Xuefei Han are acknowledged.

Literature

- [1] N. H. Khan, S. M. Alam, and S. Hassoun, "System-level comparison of power delivery design for 2D and 3D ICs," Proc. 3DIC 2009, San Francisco, CA, USA, pp. 1-7, 2009.
- [2] M. Bamal et al., "Performance comparison of interconnect technology and architecture options for deep submicron technology nodes," 2006 International Interconnect Technology Conference, San Francisco, CA, USA, pp. 202-204, 2006.
- [3] H. C. Chien et al., "Thermal evaluation and analyses of 3D IC integration SiP with TSVs for network system applications," 2012 62nd Electronic Components and Technology Conference, San Diego, CA, USA pp. 1866-1873, 2012.
- [4] S. M. Sri-Jayantha, G. Mcvickerm K. Bernstein, J. U. Knickerbocker, "Thermomechanical modeling of 3D electronic packages," IBM Journal of Research & Development, vol. 52, no. 6, pp. 623-634, 2008.
- [5] Y. Zhang et al., "Coupled electrical thermal 3D IC centric microfluidic heat sink design and technology," 2011 61st Electronic Components and Technology Conference, Lake Buena Vista, FL. USA, pp. 2037-2044, 2011.
- [6] B. A. Jaspersen, Y. Jeon, K. T. Turner, F. E. Pfefferkorn, and W. L. Qu, "Comparison of micro-pin-fin and microchannel heat sinks considering thermal-hydraulic performance and manufacturability," IEEE Transactions on Components and Packaging Technology, vol. 33, no. 1, 2010.
- [7] Z. M. Wan, Y. J. Kim, Y. Joshi, "Compact modelling of 3D stacked die inter-tier microfluidic cooling under non-uniform heat flux," Proc. ASME-IMECE 2012, Houston, TX, USA, 2012.
- [8] Z. M. Wan, Y. Joshi, "Transient Compact modelling of 3D stacked die inter-tier microfluidic cooling under non-uniform heat flux," Proc. ASME-IPACK 2013, San Francisco, CA, USA.
- [9] S. Ndao, Y. Peles, M. K. Jensen, "Multi-objective thermal design optimization and comparative analysis of electronics cooling technologies," International Journal of Heat and Mass Transfer, vol. 52, pp. 4317-4326. 2009.
- [10] A. Bejan, A. M. Morega, "Optimal arrays of pin fins and plate fins in laminar forced convection," Journal of Heat Transfer, Vol. 115, pp. 75-81, 1993.
- [11] S. Li et al., "McPAT: An Integrated Power, Area, and Timing Modeling Framework for Multicore and Manycore Architectures", 42nd Annula IEEE/ACM International Symposium on Microarchitecture, MICRO-42, New York, NY, USA, pp. 469-480, 2009.
- [12] Z. M. Wan, Y. Joshi, "Pressure drop and heat transfer characteristics of pin fin enhanced microgaps in single phase microfluidic cooling," Proc. ASME-IMECE 2013, San Diego, CA, in press.
- [13] T. Brunschwiler et al., "Interlayer cooling potential in vertically integrated packages," Microsystem Technologies, vol. 15, no.1, pp. 57-74, 2009.
- [14] A. Kosar, Y. Peles, "Thermal hydraulic performance of MEMS-based pin fin heat sink," Journal of Heat Transfer, vol. 128, no. 2, pp. 121-131, 2006.