Adaptive Dynamic Frequency Scaling for Thermal-Aware 3D Multi-core Processors

Hong Jun Choi[1], Young Jin Park[1], Hsien-Hsin Lee[2], and Cheol Hong Kim[1,*]

[1] School of Electronics and Computer Engineering,
Chonnam National University, Gwangju, Korea
[2] School of Electrical and Computer Engineering, Georgia Institute of Technology,
Atlanta, Georgia, U.S.A.
chj6083@gmail.com, blueboy24@nate.com,
leehs@gatech.edu, chkim22@jnu.ac.kr

Abstract. 3D integration technology can provide significant benefits of reduced interconnection delay and low power consumption in designing multi-core processors. However, the 3D integration technology magnifies the thermal challenges in multi-core processors due to high power density caused by stacking multiple layers vertically. For this reason, the 3D multi-core architecture cannot be practical without proper solutions to the thermal problems such as Dynamic Frequency Scaling(DFS). This paper investigates how the DFS handles the thermal problems in 3D multi-core processors from the perspective of the function-unit level. We also propose an adaptive DFS technique to mitigate the thermal problems in 3D multi-core processors by assigning different DFS levels to each core based on the corresponding cooling efficiency. Experimental results show that the proposed adaptive DFS technique reduces the peak temperature of 3D multi-core processors by up to 10.35°C compared to the conventional DFS technique, leading to the improved reliability.

Keywords: Processor architecture, 3D integration technology, Thermal management, Dynamic frequency scaling, Multi-core processor.

1 Introduction

As process technology scales down and integration densities continue to increase, interconnection delay has become one of the major constraints in improving the performance of microprocessors[1-2]. As one of the most promising solutions to reduce the interconnection delay of multi-core processors, three dimensional(3D) integration technology has drawn significant attentions[3]. In 3D multi-core processors, multiple cores are stacked vertically in the same package and components on different layers are connected through direct through-silicon vias(TSVs)[4]. Therefore, the 3D integration technology reduces the interconnection delay of multi-core processors by

---

[*] Corresponding author.

B. Murgante et al. (Eds.): ICCSA 2012, Part IV, LNCS 7336, pp. 602–612, 2012.
© Springer-Verlag Berlin Heidelberg 2012

decreasing the global wire length significantly compared to the 2D technology[5-7]. Reduced wire length also leads to the decreased power consumption [8-9]. For this reason, the 3D integration technology has drawn considerable attentions in designing recent multi-core processors.

Despite the benefits mentioned above, the 3D integration technology cannot be practical without proper solutions to the thermal problems in the processors. As mentioned in [10], the thermal problems are exacerbated in the 3D multi-core processors compared to the 2D multi-core processors mainly for two reasons. One is that the vertically stacked silicon layers cause rapid increase of power sources. The other is that the thermal conductivity of the thermal interface material(TIM) is lower than that of silicon and metal, resulting in a higher power density as more stacked TIM layers. Unfortunately, high temperature in the processor causes increased cooling costs, negative impact on the reliability and performance degradation. Therefore, more advanced cooling methods are required in designing upcoming 3D multi-core processors[2].

Dynamic thermal management(DTM) techniques using dynamic frequency scaling(DFS), dynamic voltage scaling(DVS), clock gating or computation migration have been proposed to relieve the thermal stress in 2D chips[11]. These techniques solve the thermal problems to a great extent by lowering the average temperature or by keeping the temperature under a given threshold[12]. However, conventional DTM techniques might not be sufficient to mitigate the thermal problems of 3D multi-core processors, because the peak temperature of 3D chips is much higher than that of 2D chips. Therefore, more aggressive DTM techniques for 3D chips should be investigated[13]. In our previous research, we analyzed the detailed thermal behavior of 3D multi-core processors where the conventional DFS technique is applied[2]. Based on the results in [2], in this paper, we propose an adaptive DFS technique to reduce the peak temperature of 3D multi-core processors. In the processor with the proposed adaptive DFS technique, the cores run with different DFS levels based on the corresponding cooling efficiency of the core, resulting in the reduced peak temperature.

The rest of this paper is organized as follows. Section 2 provides an overview of 3D integration technology and dynamic thermal management techniques. In Section 3, the proposed adaptive DFS technique is described. Section 4 provides the simulation methodology and the detailed results. Finally, Section 5 discusses our conclusions and future work.

2 Related Work

2.1 3D Integration Technology

3D die stacking is a new technology that increases transistor density by integrating two or more dies vertically[15]. Figure 1 shows the structural comparison between a 2D two-core processor and a 3D two-core processor. The 3D die stacking enables a significant reduction of wire length both within a die and across dies in the microprocessor. In the 3D microprocessor, blocks can be placed vertically on multiple dies to

reduce the wire distance, latency and power. For this reason, the 3D integration technology has the potential to change the processor-design constraints by providing substantial power and performance benefits compared to the 2D design technology. Despite the promising advantages, the 3D integration technology brings us some considerable issues. Thermal problem is one of the critical issues for the 3D microprocessor. The vertically stacked silicon dies causes rapid increase of power density, and the increased power density causes increased cooling costs, negative impact on the reliability and performance degradation. For this reason, more effective cooling techniques are required in designing upcoming 3D multi-core processors.

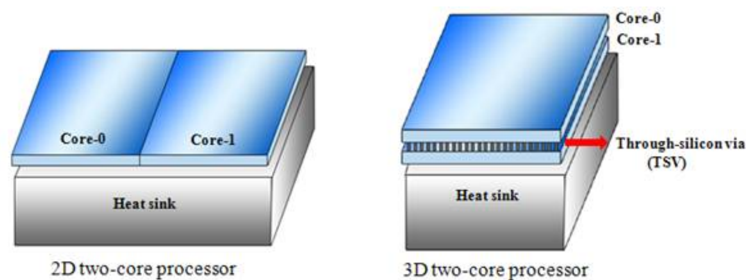

Fig. 1. Structural comparison between 2D two-core processor and 3D two-core processor

2.2 Dynamic Thermal Management Techniques

For the conventional 2D chips, dynamic thermal management(DTM) techniques have been widely used to address the thermal problems when the chip temperature exceeds the thermal limit supported by the cooling solutions[16]. A lot of researches have provided a large number of DTM techniques. They can be categorized into two different groups: One is hardware-based techniques and the other is software-based techniques. The hardware-based DTM techniques such as dynamic frequency scaling(DFS), dynamic voltage scaling(DVS) and clock gating[11] are more aggressive and effective in managing the temperature than the software-based techniques. However, the hardware-based techniques incur more execution-time overhead compared to the software-based techniques such as energy-aware process scheduling[17] and OS-level task scheduling[18].

The thermal problems are expected to be more severe in 3D multi-core processors compared to 2D multi-core processors. Therefore, more effective DTM techniques are required for addressing the thermal problems of 3D multi-core processors. The DFS technique has been regarded as one of the most efficient DTM techniques for 2D multi-core processors[19]. In this work, we analyze the detailed thermal behavior of 3D multi-core processors with the conventional DFS technique varying application features, cooling characteristics and frequency levels. We also propose an adaptive DFS technique to reduce the peak temperature of the 3D multi-core processor by assigning different DFS levels to each core depending on the corresponding cooling efficiency.

3 Adaptive Dynamic Frequency Scaling Technique

In the multi-core processor using the conventional DFS technique, the DFS levels applied to the cores have no difference because the target processor of the conventional DFS technique is the 2D multi-core processor where all the cores have comparable cooling efficiency. However, there is a significant difference in the cooling efficiency of the cores in the 3D multi-core processor depending on the vertical location of the core. For this reason, so as to mitigate the thermal problems in the 3D multi-core processor based on the fact that the cores with better cooling efficiency (the cores located nearer to the heat sink) can be clocked at a higher frequency compared to the cores with worse cooling efficiency[2], we propose an adaptive DFS technique which assigns different DFS levels to the cores by considering the corresponding cooling efficiency. We determine the DFS levels to the cores based on the theory in [20] to find the optimal frequency according to the temperature[21-22].

Table 1. DFS levels for the conventional DFS vs. DFS levels for the adaptive DFS

	Level-1	**Level-2**	**Level-3**	**Level-4**
core-0 (conventional DFS)	1GHz	2GHz	3GHz	4GHz
core-1 (conventional DFS)	1GHz	2GHz	3GHz	4GHz
core-0 (adaptive DFS)	0.5GHz	1.5GHz	2.5GHz	3.5GHz
core-1 (adaptive DFS)	1.5GHz	2.5GHz	3.5GHz	4.5GHz

The proposed adaptive DFS technique assigns different DFS levels to each core in the 3D multi-core processor by considering the cooling efficiency of the core. Therefore, in the 3D multi-core processor using the adaptive DFS technique, a higher baseline frequency is assigned to the core with better cooling efficiency while a lower baseline frequency is assigned to the core with worse cooling efficiency. In the 3D two-core processor depicted in Figure 1, the cooling efficiency of core-1 is better than that of core-0 owing to the shorter distance from the heat sink. For an example, as shown in Table 1, the proposed adaptive DFS technique assigns different DFS levels to the cores while the conventional DFS technique assigns same DFS levels to the cores. In the table, core-1 represents the core near the heat sink while core-0 means the core far from the heat sink. Based on the assumption that the conventional DFS has four DFS levels(1GHz, 2GHz, 3GHz, 4GHz), the higher four DFS levels(1.5GHz, 2.5GHz, 3.5GHz, 4.5GHz) are assigned to the core-1 in the adaptive DFS technique. To make the sum of the frequency to the same value, the lower four DFS levels(0.5GHz, 1.5GHz, 2.5GHz, 3.5GHz) are assigned to the core-0 in the adaptive DFS technique. As shown in Table 1, the sum of all the frequencies for the adaptive DFS is equal to that for the conventional DFS.

The baseline frequency of the core with relatively better cooling efficiency for the adaptive DFS is higher than that for the conventional DFS, while the baseline frequency of the core with relatively worse cooling efficiency for the adaptive DFS is lower than that for the conventional DFS. The frequency gap between the DFS levels for the adaptive DFS is equal to that for the conventional DFS.

We expect that the proposed adaptive DFS technique can reduce the peak temperature of the 3D multi-core processor compared to the conventional DFS, since the cooling efficiency of each core depending on the distance from the heat sink is considered in the proposed adaptive DFS technique.

4 Experiments

4.1 Experimental Methodology

Our baseline processor for the simulated core is Alpha21264(EV6)[23] without L2 cache, as shown in Figure 2. We extend it to a dual-core configuration for multi-core simulations. We use SimpleScalar [24] as our system simulator, which provides a detailed cycle-level modeling of processors. In addition, Wattch [25] is used to obtain a detailed power trace. We select two applications(mcf, gcc) from SPEC CPU2000 benchmark[26], because the chosen applications show a significant difference in the perspective of the peak temperature of the processor.

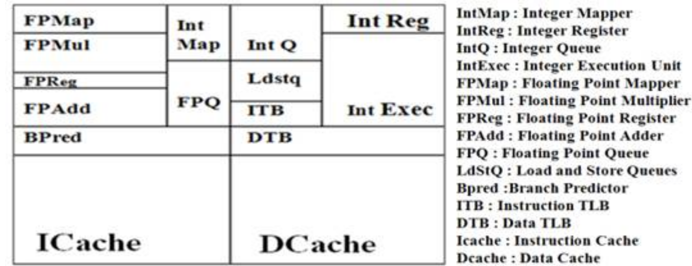

Fig. 2. Floorplan of Alpha21264(EV6)

HotSpot version 5.0 is used as our thermal modeling tool[27] and modeled parameters of the silicon layer and thermal interface material are shown in Table 2. Thermal modeling parameters can be obtained by considering material properties described in CRC handbook[28]. In the table, core-0 represents the core far from the heat sink, while core-1 is near the heat sink.

Table 2. Thermal modeling parameters

Parameter	Value			
	TIM-0	**CORE-0**	**TIM-1**	**CORE-1**
Specific heat capacity (J/m^{3}K)	4.00e6	1.75e6	4.00e6	1.75e6
Thickness(m)	2.00e-5	1.50e-4	2.00e-5	1.50e-4
Resistivity (mK/W)	0.25	0.01	0.25	0.01

In the experiments, the DFS technique is engaged continuously instead of reacting to thermal emergencies. The conventional DFS technique is applied with four

levels(1GHZ, 2GHz, 3GHz and 4GHz) and the sum of the frequencies of two cores is 5GHz for all simulations.

4.2 Thermal Impact of DFS on 3D Multi-core Processors

In order to analyze the thermal impact of DFS levels, the simulations are performed by running the same application(mcf) on both cores, only varying the DFS levels. Therefore, four schemes can be simulated as follows: 1GHz/4GHz, 2GHz/3GHz, 3GHz/2GHz, 4GHz/1GHz, where the slash separates the frequencies of both cores. The former represents the frequency of core-0 which is far from the heat sink and the latter refers to the frequency of core-1 which is near the heat sink.

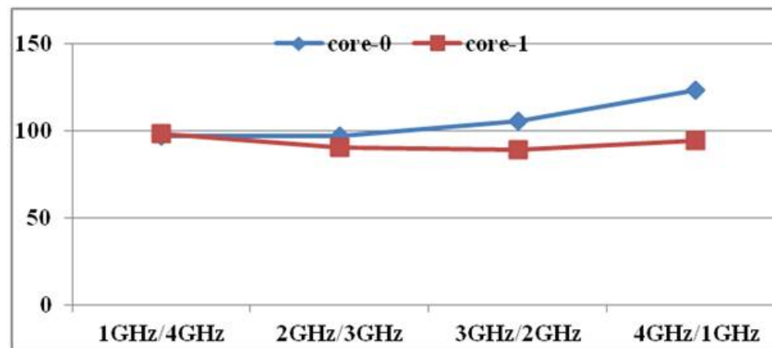

Fig. 3. Peak temperature(°C) according to DFS levels

The peak temperature according to DFS levels is shown in Figure 3. The peak temperature for the compared four schemes are 98.09°C(1GHz/4GHz), 97.17°C(2GHz/3GHz), 105.59°C(3GHz/2GHz) and 123.1°C(4GHz/1GHz). For the first scheme(1GHz/4GHz) and the last scheme(4GHz/1GHz), the same frequency values are used, just swapping the frequencies on the two cores, while the peak temperature changes from 98.09°C to 123.1°C. The same pattern can be seen in the second and the third scheme, the peak temperature varies from 97.17°C to 105.59°C. As shown in the graph, same frequency values yield a totally different thermal profile, because the cooling efficiency of each core in the 3D multi-core processor is not comparable. In general, the core near the heat sink shows better cooling efficiency than the core far from the heat sink. Therefore, to reduce the temperature of the hot unit in the core, the core with relatively better cooling efficiency should be clocked at a higher frequency than the core with relatively worse cooling efficiency.

4.3 Thermal Impact of Workload Distribution on 3D Multi-core Processors

In order to identify the thermal pattern according to the workload distribution for 3D multi-core processors, mcf and gcc applications are chosen to be run because these two benchmark applications show a great difference in generating thermal impact on

the processors. In the 2D single-core processor, mcf shows very high peak temperature(123.5°C), while gcc gets low peak temperature(89.9°C). In the experiments, two different applications are run on two cores with four DFS levels, resulting in eight different schemes as follows:

mcf1GHz/gcc4GHz, mcf2GHz/gcc3GHz, mcf3GHz/gcc2GHz, mcf4GHz/gcc1GHz, gcc1GHz/mcf4GHz, gcc2GHz/mcf3GHz, gcc3GHz/mcf1GHz, gcc4GHz/mcf1GHz

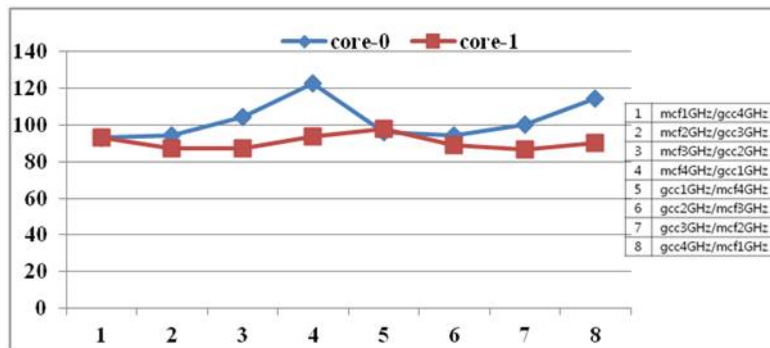

Fig. 4. Peak temperature(°C) according to workload distribution

The running application and DFS level of the core-0 are shown before the slash and those of the core-1 are shown after the slash. For the first four schemes, gcc is run on the core with better cooling efficiency and mcf is run on the other core. Then, applications are swapped for the last four schemes to observe the thermal impact of the workload distribution on 3D multi-core processors. As shown in Figure 4, the peak temperatures of eight simulated schemes are 93.08°C, 94.43°C, 104.1°C, 122.53°C, 97.51°C, 94.12°C, 100.01°C and 114.49°C, respectively. The core-1 yields lower temperature than the core-0 despite the type of workload because of its near-heat-sink location advantage(better cooling efficiency) except for fifth scheme(gcc1GHz/mcf4GHz). For the fifth scheme, the frequency of the core-1 is too higher than that of the core-0.

For the fourth scheme(mcf4GHz/gcc1GHz) and the eighth scheme(gcc4GHz/mcf1GHz), the frequency of each core is same, but the temperature of the eighth scheme is lower than that of the fourth scheme. Same pattern can be seen in the third and seventh scheme, reflecting the thermal impact of assigning different workloads to the cores with different cooling efficiency. Therefore, to reduce the temperature of the cores, the relatively heavier workload should be assigned to the core with relatively better cooling efficiency.

4.4 Conventional DFS Technique vs. Adaptive DFS Technique for 3D Multi-core Processors

In the processor using the conventional DFS technique, the DFS levels applied to each core are identical. However, our simulation results show that the core with better

cooling efficiency can be clocked at a higher frequency to mitigate the thermal problems in 3D multi-core processors. For this reason, the proposed adaptive DFS technique assigns different DFS levels to the cores depending on the corresponding cooling efficiency.

For a dual-core processor using the adaptive DFS technique, the DFS levels applied to the core with better cooling efficiency are 1.5GHZ, 2.5GHz, 3.5GHz and 4.5GHz, while the DFS levels applied to the core with worse cooling efficiency are 0.5GHz, 1.5GHz, 2.5GHz and 3.5GHz. In the simulations, the sum of the frequencies of two cores in the adaptive DFS is set to 5GHz, because the sum of the frequencies of two cores is 5GHz in the conventional DFS. Therefore, for an example, one core is set to operate at 3.5GHz in the adaptive DFS, the frequency of the other core is set to 1.5GHz. Moreover, in order to make the results comparable with previous simulation results, same applications(mcf and gcc) are used. Simulated eight schemes for the adaptive DFS technique can be listed as follows:

mcf0.5GHz/gcc4.5GHz, mcf1.5GHz/gcc3.5GHz, mcf2.5GHz/gcc2.5GHz, mcf3.5GHz/gcc1.5GHz, gcc0.5GHz/mcf4.5GHz, gcc1.5GHz/mcf3.5GHz, gcc2.5GHz/mcf2.5GHz, gcc3.5GHz/mcf1.5GHz

The application and DFS level of the core-0 are shown before the slash and those of core-1 are shown after the slash. The peak temperature of the core in each scheme is shown in Figure 5.

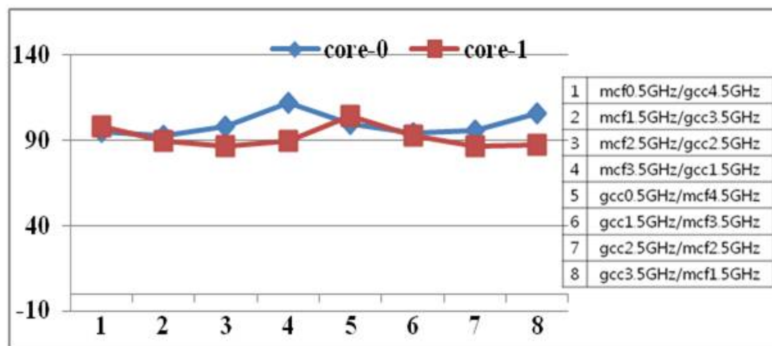

Fig. 5. Peak temperature(°C) with adaptive DFS

As shown in Figure 4 and Figure 5, the peak temperature with the adaptive DFS is lower than that with the conventional DFS except for the first scheme(mcf0.5GHz/ gcc4.5GHz) and the fifth scheme(gcc0.5GHz/ mcf4.5GHz). The purpose of the adaptive DFS is reducing the frequency of the core with worse cooling efficiency while the frequency of the core with better cooling efficiency increases. In the simulated dual-core processor, the core with better cooling efficiency is core-1 while the core-0 has worse cooling efficiency. For the first scheme with the conventional DFS (shown in Figure 4), the peak temperature of core-1 is similar to that of core-0. For the fifth scheme with the conventional DFS, the peak temperature of core-1 is higher than that of core-0. For the first and fifth scheme with the adaptive DFS, the frequency of the

core-1 is higher than that with the conventional DFS. Therefore, the peak temperature of two schemes with the adaptive DFS scheme increases. However, the peak temperature with the adaptive DFS is lower than that with the conventional DFS by 5.01°C on the average. Especially, for the fourth scheme(gcc3.5GHz/ mcf1.5GHz) with the adaptive DFS, the peak temperature is lower than that with the conventional DFS by up to 10.35°C. Simulation results prove that assigning different DFS levels to each core depending on the cooling efficiency can reduce the peak temperature of the 3D multi-core processor, resulting in the improved reliability and better performance.

5 Conclusions

In this paper, we analyzed the thermal behavior of 3D multi-core processors varying DFS levels and workload distribution. We demonstrated that the core near the heat sink can be clocked at a higher frequency than the core far from the heat sink to keep the performance without thermal emergencies. The workload with bigger thermal influence also can be assigned to the core with better cooling efficiency to reduce the overall temperature of the 3D multi-core processor. We also proposed an adaptive DFS technique to mitigate the thermal problems in the 3D multi-core processor by assigning different DFS levels to each core based on the corresponding cooling efficiency. According to our simulation results, the proposed adaptive DFS technique reduces the peak temperature of the 3D multi-core processor by 5.01°C on the average compared to the conventional DFS technique. Therefore, we expect that the proposed adaptive DFS can be a good solution to reduce the peak temperature of the upcoming 3D multi-core processors.

Acknowledgements. This work was supported by the National Research Foundation of Korea Grant funded by the Korean Government (NRF-2011-013-D00105).

References

1. Joyner, J.W., Zarkesh-Ha, P., Davis, J.A., Meindl, J.D.: A Three-Dimensional Stochastic Wire-Length Distribution for Variable Separation of Strata. In: Proc. of IEEE International Interconnect Technology Conference, SanFrancisco, USA, pp. 132–134 (June 2000)
2. Park, Y.J., Zeng, M., Lee, B.S., Lee, J.A., Kang, S.G., Kim, C.H.: Thermal Analysis for 3D Multi-core Processors with Dynamic Frequency Scaling. In: Proc. of IEEE/ACIS International Conference on Computer and Information Science, Kaminoyama, Japan, pp. 69–74 (August 2010)
3. Jang, H.B., Yoon, I., Kim, C.H., Shin, S., Chung, S.W.: The impact of liquid cooling on 3D multi-core processors. In: Proc. of the 2009 IEEE International Conference on Computer Design, California, USA, pp. 472–478 (October 2009)
4. Zhu, C., Gu, Z., Shang, L., Dick, R.P., Joseph, R.: Three-dimensional chip-multiprocessor run-time thermal management. IEEE Transactions on Computer-Aided Design of Lntegrated Circuits and Systems 27(8), 1479–1492 (2008)

5. Puttaswamy, K., Loh, G.H.: Dynamic Instruction Schedulers in a 3-Dimensional Integration Technology. In: Proc. of the ACM Great Lake Symposium On VLSI, Philadelphia, USA, pp. 153–158 (May 2006)
6. Puttaswamy, K., Loh, G.H.: Implementing Caches in a 3D Technology for High Performance Processors. In: Proc. of the International Conference on Computer Design, San Jose, USA, pp. 525–532 (October 2005)
7. Reed, P., Yeung, G., Black, B.: Design Aspects of a Microprocessor Data Cache using 3D Die Interconnect Technology. In: Proc. of the International Conference on Integrated Circuit Design and Technology, pp. 15–18 (May 2005)
8. Puttaswamy, K., Loh, G.H.: Thermal Analysis of a 3D Die Stacked High Performance Microprocessor. In: Proc. of ACM Great Lakes Symposium on VLSI, Philadelphia, USA, pp. 19–24 (May 2006)
9. Davis, W.R., Wilson, J., Mick, S., Xu, J., Hua, H., Mineo, C., Sule, A.M., Steer, M., Franzon, P.D.: Demystifying 3D ICs: The Pros and Cons of Going Vertical. IEEE Design Test Computers 22, 498–510 (2005)
10. Cong, J., Luo, G.J., Wei, J., Zhang, Y.: Thermal-Aware 3D IC Placement Via Transformation. In: Proc. of ASP-DAC (2007), Yokohama, Japan, pp. 780–785 (January 2007)
11. Brooks, D., Martonosi, M.: Dynamic Thermal Management for High-performance Microprocessors. In: Proc. of the 7th International Symposium on High-Performance Computer Architecture, Monterrey, Mexico, pp. 171–182 (January 2001)
12. Coskun, A.K., Ayala, J.L., Atienza, D., Rosing, T.S., Leblebici, Y.: Dynamic Thermal Management in 3D Multicore Architectures. In: Proc. of Design, Automation & Test in Europe Conference & Exhibition, Nice, France, pp. 1410–1415 (April 2005)
13. Kumar, R., Zyuban, V., Tullsen, D.M.: Interconnections in multi-core architectures: Understanding mechanisms, overheads and scaling. In: Proc. of the 32th Annual International Symposium on Computer Architecture, Madison, USA, pp. 408–419 (June 2005)
14. Takahashi, K., Sekiguchi, M.: Through Silicon Via and 3-D Wafer/Chip Stacking Technology. In: Proc. of the 2006 Symposium on VLSI Circuit Digest of Technical Papers, Honolulu, USA, pp. 89–92 (June 2006)
15. Black, B., Annavaram, M., Brekelbaum, N., DeVale, J., Jiang, L., Loh, G.H., McCauley, D., Morrow, P., Nelson, D.W., Pantuso, D., Reed, P., Rupley, J., Shankar, S., Shen, J., Webb, C.: Die Stacking (3D) Microarchitecture. In: Proc. the 39th Annual IEEE/ACM International Symposium on Microarchitecture, Florida, USA, pp. 469–479 (December 2006)
16. Kumar, A., Shang, L., Peh, L.S., Jha, N.K.: System-level dynamic thermal management for high-performance microprocessors. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems 27(1) (January 2008)
17. Mishra, R., Rastogi, N., Zhu, D., Mosse, D., Melhem, R.: Energy Aware Scheduling for Distributed Real-Time Systems. In: The Proc. International Parallel and Distributed Processing Symposium, Nice, France, pp. 21–30 (April 2003)
18. Zhou, X., Xu, Y., Du, Y., Zhang, Y., Yang, J.: Thermal Management for 3D Processors via Task Scheduling. In: Proc. of the 2008 37th International Conference on Parallel Processing, Portland, USA, pp. 115–122 (September 2008)
19. Skadron, K., Stan, M.R., Sankaranarayanan, K., Huang, W., Velusamy, S., Tarjan, D.: Temperature-aware microarchitecture: modeling and implementation. ACM Transactions on Architecture and Code Optimization 1(1), 94–125 (2004)

20. Shi, B., Zhang, Y., Srivastava, A.: Dynamic Thermal Management for Single and Multicore Processors Under Soft Thermal Constraints. In: Proc. of the 16th ACM/IEEE International Symposium on Low Power Electronics and Design, Poster, Austin, USA (August 2010)
21. Li, M., Zhao, W.: Visiting power laws in cyber-physical networking systems. Mathematical Problems in Engineering 2012, Article ID 302786, 13 pages (2012)
22. Li, M., Zhao, W.: Asymptotic identity in min-plus algebra: a report on CPNS. Computational and Mathematical Methods in Medicine 2012, Article ID 154038, 11 pages (2012)
23. Kessler, P.E.: The Alpha 21264 Microprocessor. IEEE Micro 19(2), 24–36 (1999)
24. Burger, D.C., Austin, T.M.: The SimpleScalar tool set, version 2.0. ACM SIGARCH CAN 25(3), 13–25 (1997)
25. Brooks, D., Tiwari, V., Martonosi, M.: Wattch: a framework for architectural-level power analysis and optimizations. In: Proc. of the 27th International Symposium on Computer Architecture, pp. 83–94 (June 2000)
26. Henning, J.L.: SPEC CPU 2000: Measuring CPU Performance in the New Millennium. IEEE Computer 33(7), 28–35 (2000)
27. Hotspot, `http://lava.cs.virginia.edu/HotSpot`
28. CRC Press, CRC Handbook of Chemistry, `http://www.hbcpnetbase.com`