International Journal of Computer Science and Business Informatics

IJCSBI.ORG

Prospects of Thermal Management Techniques in Microprocessor Architecture

Ajaegbu Chigozirim
Department of Computer Science
Babcock University, Nigeria

Shodiya A.S
Department of Computer Science
Federal University of Agriculture, Abeokuta, Nigeria

Kuyoro Shade O.
Department of Computer Science
Babcock University, Nigeria

ABSTRACT
In all circuit networks, there is a need to balance the amount of heat generated by each circuit component with the power density. Managing the heat generated by electronic elements have always been one of the greatest tasks in any electronic design. Many thermal management techniques of microprocessor have been proposed and implemented for different systems. It was observed that the issue of microprocessor thermal-management needs to be addressed from the basis, that is, the application of the floor planning technique. Also, there is the need to balance the performance measure with the thermal technique in such a way that one needs not to be compromised for another. This paper identifies some of the thermal management techniques proposed by different authors, the gaps and the likely prospects.

Keywords
Thermal management, Microprocessor, Power, Electronic design, Floor planning

1. INTRODUCTION
The most recent, and arguably one of the most difficult obstacles to the exponential growth in transistor density predicted by Moore's Law is that of removing the large amount of heat generated within the tiny area of a microprocessor. The majority of the transistors on these chips exist to extract Instruction Level Parallelism and to reduce memory access times for a range of common applications; unfortunately, the performance benefits of many of these techniques are gradually becoming overshadowed by increased design complexity and power dissipation.

As we move towards billion transistor microprocessors, the growing power budgets of these chips must be addressed at all levels of the design cycle.

Thermal issues are fast becoming major design constraints in high-performance systems. Temperature variations adversely affect system reliability and prompt worst-case design.

This paper proceeds as follows. In the next section, some related works are discussed. Section 3 discusses the prospects of thermal management techniques and finally Section 4 gives the conclusion.

2. RELATED WORKS

The following are some of the recent related works on thermal management techniques:

2.1 Interconnect and Thermal-aware Floor planning for 3D Microprocessors [1]

This work centered on one of the promising technologies of 3D technology known as wafer-bonding technology. It addresses the problem of interconnect power consumption which was observed from their review of other thermal aware floor planning proposals to have been left uncovered. This according to their paper causes a high peak temperature in 3D designs. The objectives were as follows: To show the good effects of 3D wafer-bonding technology on interconnect issues; To Investigate how the increased power densities of the stacked blocks and reduced interconnect power consumption combine to influence thermal behavior and to study a floor planning algorithm that attempts to reduce the peak temperature in 3D designs. The work adopted a thermal-aware floor planning algorithm for 3D architecture methodology based on B*tree representation. The temperature estimation method used was an extended Hotspot model used for 2D temperature estimation. The extension comprises of a variable number of additional levels each composing of both silicon layer and inter-silicon glue material. The interconnect power consumption were extracted using Design Complier first encounter methodology and this was used in the floor planning algorithm.

The work showed that there is a net difference of 6˚C between peak temperature when the interconnect power consumption was not considered and when it was considered. This was achieved through the separation of high power density modules in different layers. The floor plan structure used was not stated that is if it is a slicing structure or a non-slicing structure. Thus B*tree floor plan model was not properly justified and the means of handling the issue of flexibility in B*tree floor plan model was also left out.

2.2 Thermal management for 3D processors via Task Scheduling [2]

Despite the good deeds of 3D technology such as reduction of interconnect delay across the dice, reduction of wiring both within and across disparate dice. 3D technology has the challenge of heat generation and this is more prone along its vertical direction. The paper seeks to provide an OS scheduling technique that will take care of heat generation in the vertical side of the 3D chip

multiprocessor. The main objective of their study was to propose an OS – level scheduling algorithm that performs thermal – aware task on 3D chip focusing on the heat generation on the vertical direction. Thus the specific objectives were: to review current works investigating the performance potential and the challenges attached to 3D Chip Microprocessor (CMP); to show that there is a strong thermal correlation among adjacent layers in a 3D processor; to establish the fact that cores farther from the heat sink are hotter than those nearer to it; to develop an algorithm that involves mostly sorting of the power and temperature. The paper adopted a comparative study of some task scheduling algorithms and the outcome of which was used to develop their own algorithm known as "Temperature balancing by stack" algorithm. Also, in order to establish the fact that in 3D chip, vertically adjacent cores have larger thermal impact among each other than horizontal adjacent cores, the researchers adopted a simple heat transfer model of a basic two-layer 3D chip structure that uses a face-to-back bonding technology because of its better scalability in layer count.

The proposed schedule algorithm when compared with the Baseline, Random, Round-Robin and temperature balancing by core, offers better performance in terms of thermal behavior, peak temperature reduction, thermal emergency reduction and performance improvement. In essence, their proposed algorithm removes 54% hardware DTM on average while resulting to 7.2% performance improvement over the base case.

Though, the proposed algorithm showed better improvement in task scheduling thermal management, it however did not take care of likely increase in the number of context switches. Also, the work did not put into consideration heat generation components like integrated heat spreader (HIS), die package, socket and motherboard which might cause variation in the heat conductivity between the two adjacent dice of the adapted face – to – back bonding technology.

2.3 System-Level Dynamic Thermal Management for High Performance Microprocessors [3]

The Independent use of hardware dynamic thermal management (HDTM) and software dynamic thermal management (SDTM) in managing the overall temperature has resulted to interaction gap between different techniques. This led to high performance impact that is, the chip temperature having high impact on the system performance. The objectives of this work were: to develop a novel regression – based thermal model for providing fast and accurate prediction of temperature covering system – wide thermal effects; to present a comparative analysis of the pros and cons of the SDTM, HDTM and HybDTM. The methodology used a performance counter of P4 processor along with their proposed regression based thermal model. The regression based thermal model was integrated inside the Linux 2.6.9 kernel which uses hardware performance counter as input to gather the thermal profiles of the individual applications running on the system and to estimate the overall temperature at runtime. Also,

in order to account for the system wide effect temperature including memory, they developed a simple power and thermal model for a memory using an adopted approach. They also designed micro benchmarks targeting different processor functional units.

At the end of the work, a hybrid DTM model was developed. The model was able to account for the thermal impact of individual applications at runtime while at the same time exercising control over thermal emergencies. When evaluated with SPEC2000 benchmark, the proposed HybDTM outperforms SDTM and HDTM in terms of effectiveness in eliminating all thermal emergencies and execution time overhead. The work failed to put into consideration the following: imprecision due to sensor placement; avoidance of the memory clock enabled getting too low when the mechanism is applied because CLE can only be in its activation mode when it is high; it did not consider the HDTM part from the architectural level.

2.4 Thermal-aware Task scheduling at the system software level [4]

Despite the existing approaches on distribution of power consumption across a multi-core chip, in order to mitigate the thermal dissipation problem without compromising performance, the researchers observed that an improved method for dynamic redistribution of power are needed to meet power and thermal envelops without giving up performance in order to deal with and exploit workload variability. The main objective was to develop a prototype Linux – based implementation of a thermal aware scheduler. This led to the various specific objectives as follows: to examine the thermal time constants on the POWER5 chip; to demonstrate that an OS scheduler can mitigate on-chip hotspots through leveraging spatial heat-slack; to demonstrate that an OS scheduler can mitigate hot spots by leveraging temporal heat slack through intelligent workload scheduling; to demonstrate that the OS scheduler can have a temperature reducing effect by co-scheduling System Management Technique (SMT) – mode tasks intelligently. A 1.2GHz POWER5 system running Bare Metal Linux (IBM) was used for the experiment. The thermal sensors were calibrated using the spatially – resolved imaging of microprocessor power (SIMP) methodology. SIMP methodology calibrates the sensor by capturing infra – red images of the POWER5 chip when it reaches stable temperatures and make comparison of image temperature with the reading of digital temperature sensor value. The infrared imaging was made possible through the replacement of the mental heat sink of the POWER5 system with a liquid heat sink that is transparent. From their experiment, they observed the following:

- For thermal time constant studies, the change in workload can cause changes in on-chip temperatures up to 16˚Celsius.
- In the case of core hopping (for leveraging spatial heat slack), it was observed that core hopping can cause temperature changes of 5.5˚Celsius while causing less than 1% average slowdown.

- In task scheduling, it was also observed that mixing a cold task with a hot one reduces on – chip temperature as much as 5˚C.
- Again, it was observed that a cold task with a hot task on SMT reduces on-chip temperature by up to 3˚Celsius.

In essence, their schemes enabled lowering of on-chip unit temperature by changing the workload in a timely manner with OS and existing hardware support. In all, their work failed to look at the unfairly penalization of heterogeneous workload that may arise.

2.5 Dimetrodon: Processor- level Preventive Thermal management via Idle Cycle Injection [5]

The paper centered on the challenge attached to traditional dynamic thermal management techniques. These techniques focused on worse-case thermal emergencies but are not designed to operate under normal thermal conditions. The general objective was to control the CPU temperature while ensuring that additional energy is not consumed by CPU as a result of that. Thus the specific objectives were as follows: to propose the use of idle cycle injection for per thread preventive thermal control; to implement and evaluate an idle cycle injection in a preventive thermal management context on real time hardware in a commodity OS and comparison with a hardware – based approach; to characterize the application- level impact across worst case thermal load and real world workloads. The method that was adopted in the paper was the design of an OS scheduler known as Dimetrodon. The scheduler uses an idle cycle injection as a preventive thermal management mechanism targeting the average case scenario.

Their proposed model was validated using CPUburn package (Specifically burn6). Their implementation gave a throughput of average 1.0% lower than expected. Also when Dimetrodon was compared with race – to – ldle technique over an average of five trials (using some selected benchmarks), Dimetrodon consumed between 97.6% as against 103.7% of the energy of race – to – idle. In all, their work centered much on temperature reduction of running applications thereby giving less attention to performance enhancement. Again, the effect of Dimetrodon was not discussed in a situation where concurrent processes are involved and the effect of ambient temperature was not considered.

3. PROSPECTS

In the case of thermal – aware floor planning in 3D Microprocessors, the B*tree floor plan algorithm is mainly limited to non-slicing structure. Since research has shown that B*tree has the challenge of flexibility, the suggested future work will be on how a slicing structure (i.e. a binary tree whose leaves denote modules and internal nodes specify horizontal or vertical cut lines) binary tree algorithm and B*tree can be incorporated together in a 3D floor plan design. (See Table 1)

Table 1. List of different 3D floorplan structures [7]

Representation	Advantages (a) and Disadvantages (d)
Binary Tree	a1. Efficient a2. Flexible to deal with hard, pre-placed, soft, and rectilinear modules, etc a3. Smaller encoding cost a4. Can operate on the tree directly, no need to do transformation during processing a5. Can evaluate area cost incrementally a6. Transformation between representation and placement takes linear time
	d1. Can handle only the slicing structure
Sequence pair/BSG	a1. Can handle only the slicing structure a2. Very flexible in representation
	d1. Time – consuming d2. The solution space is large d3. Sequence encoding cost is high d4. Harder to transform between a sequence pair and a placement d5. Sequence pair cannot handle soft modules directly d6. BSG incurs redundancies
O-tree	a1.Can handle non-slicing structure a2. The solution space is smaller a3. Transformation between representation and placement takes only linear time a4. Encoding by fewer bits than sequence pair and BSG
	d1. Less flexible than BSG/sequence pair in representation d2. Tree structure is irregular, harder for implementation d3. Need to encode and operative on module sequence d4. Need to transform between the tree and its placement during processing d5. Inserting positions are limited, might deviate from the optimal during solution perturbation
B*tree	a1. Can handle non-slicing structure a2. Binary-tree based, efficient a3. Flexible to deal with hard, pre-place, soft, and rectilinear modules, etc a4. Smaller encoding cost a5. Except for handling soft modules, only need to transform from a tree to its placement during processing, which takes only linear time a6. Can evaluate area cost incrementally a7. The solution space is smaller
	d1. Less flexible than BSG/sequence pair in representation

In case of hybrid technique, the HDTM should be modeled in such a way that it will be able to accommodate the motherboard heat and ambient temperature. In other words the number of hardware event counters should be made to go to infinity while due consideration is also given to the memory clock enabled temperature. This will help to build more formidable hybrid DTM model.

4. CONCLUSION

The quest for computer performance, has led to increment in the number of transistors per microprocessor. This increment has also resulted to compact space between components which in turn affects the voltage scaling of the components. Power dissipation in microprocessors affects the temperature of the system. Increment in system's temperature can result to system failure thus the need for an effective thermal management technique capable of controlling the temperature scaling of a particular system.

REFERENCES

[1] Hung, W. L., Xie, Y., Vijaykrishnan, N., Addo-Quaye, C., Theocharides, T., and Irwin, M. J. Interconnect and Thermal-aware Floorplanning for 3D Microprocessors. In Proceedings of the 7th International Symposium on Quality Electronic Design (ISQED'06), (2006), *IEEE Computer Society* http://www.cse.psu.edu/~yuanxie/Papers/ISQED05-Hung.pdf

[2] Zhou, X., Yang, J., Xu, Y., Du, Y. and Zhang, Y. Thermal Management for 3D Processors via Task Scheduling. In Proceedings of the 37th International Conference on Parallel Processing, (ICPP'08) *IEEE Computer Society*, (2008), 115-122.

[3] Kumar, A., Shang, L., Peh, S., Jha, N. System-Level Dynamic Thermal Management for High-Performance Microprocessors. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*. 27, 1 (2008), 96-108.

[4] Jeonghwan, C.; Chen-Yong, C.; Hubertus, F.; Hendrik, H.; Alan, W. and Pradip, B. Thermal-aware Task Scheduling at the system software level. In *Proceedings of the international symposium on Low power electronics and design* (ISLPED '07), (2007), 213-218.

[5] Bailis, P., Vijay, J. R., Sanjay, G., David, B., and Margot, S. Dimetrodon: processor-level preventive thermal management via idle cycle injection. In *Proceedings of the 48th Design Automation Conference* (DAC), San Diego, California, (2011), 89-94.

[6] Yun-Chih, C., Yao-Wen, C., Guang-Ming, W., and Shu-Wei W. B*-Trees: a new representation for non-slicing floorplans Proceeding DAC '00. In *Proceedings of the 37th Annual Design Automation Conference* (DAC'00) (2000), 458-463.

[7] Schmidt, R. Challenges in Electronic Cooling- Opportunities for Enhanced Thermal management Techniques- Microprocessor Liquid Cooled Minichannel Heat Sink. *Heat Transfer Engineering*. 25, 3 (2004), 3-12.