

Thermal Management for 3D Processors via Task Scheduling

Xiuyi Zhou[†]

Jun Yang[†]

Yi Xu[†]

Yu Du[‡]

Youtao Zhang[‡]

[†]Electrical and Computer Engineering
University of Pittsburgh, Pittsburgh PA 15261

[‡]Computer Science
University of Pittsburgh, Pittsburgh PA 15261



Abstract—A rising horizon in chip fabrication is the 3D integration technology. It stacks two or more dies vertically with a dense, high-speed interface to increase the device density and reduce the delay of interconnects significantly across the dies. However, a major challenge in 3D technology is the increased power density which gives rise to the concern of heat dissipation within the processor. High temperatures trigger voltage and frequency throttlings in hardware which degrade the chip performance. Moreover, high temperatures impair the processor's reliability and reduce its lifetime.

To alleviate this problem, we propose in this paper an OS-level scheduling algorithm that performs thermal-aware task scheduling on a 3D chip. Our algorithm leverages the inherent thermal variations within and across different tasks, and schedules them to keep the chip temperature low. We observed that vertically adjacent dies have strong thermal correlations, and the scheduler should consider them jointly. Compared with other intuitive algorithms such as a Random and a Round-Robin algorithm, our proposed algorithm brings lower peak temperature and average temperature on-chip. Moreover, it can remove on average 46% of thermal emergency time and result in 5.11% (4.78%) performance improvement over the base case on thermally homogeneous (heterogeneous) floorplans.

Index Terms—3D processors, thermal-aware scheduling

1 INTRODUCTION

The 3D integration technology has gained significant attention recently. This is a technology that reduces wiring both within and across disparate dies, as wiring has become a major latency, area and power overhead in modern microprocessors. Studies have shown that wires can consume more than 30% of the power within a 2D CMP [3]. The 3D technology provides vertical stacking of two or more dies with a dense, high-speed interface, reducing the wire length by a factor of the square root of the number of layers used [16]. This significant reduction leads to improved performance and lower power dissipation on the interconnection.

One key challenge in 3D die stacking is the heat generation from the internal active layers because the power density per unit volume increases drastically in 3D. This exacerbates existing hotspots and can create new hotspots within the chip, especially when active logic circuits are vertically aligned. For example, the

peak temperature can increase by 17~20°C in a two-layer 3D implementation for an Alpha-like processor, compared to a 2D design [14], [22]. Other studies on logic-logic stacking 3D floorplans [1], [3], [23] also show similar thermal constraint.

There are existing dynamic thermal management (DTM) techniques such as dynamic voltage and frequency scaling (DVFS) at the architecture level to mitigate this problem. Hardware DTMs can respond to thermal crisis quickly and control the temperature efficiently by reducing the processor power, but inevitably leads to degraded performance. Recently, there has been an increasing interest in OS-assisted task scheduling on both single-core and 2D chip multiprocessors to alleviate the thermal condition on chip [5], [8], [17], [18], [21]. OS-assisted task scheduling can reduce the number of times DTMs are triggered while still meeting the thermal constraint. This technique not only improves the chip performance under the same thermal constraint, but also does not require any hardware modifications. Hardware DTMs are engaged only when task scheduling cannot keep the temperature below the thermal threshold.

In this paper, we propose a heuristic OS-level technique that performs thermal-aware task scheduling on a 3D chip multiprocessor (CMP). The proposed technique aims to improve task performance by keeping the temperature below the threshold to reduce the amount of DTMs. Unlike previous thermal-aware OS task scheduler for single core or 2D CMP, our scheduler for 3D chips must take into account the thermal conduction in the vertical direction. Early studies have shown that vertically adjacent dies have strong thermal correlations [2], [26]. For example, a core in one layer could become hot because of a high power task running in the same vertical column but at a different layer. Based on these observations, our proposed scheduler always considers the aggregated power of cores that are vertically aligned. Further, when a core is overheated, we choose to engage DTM on a vertically aligned core that generates the most power. Such an approach can greatly reduce the total power in one vertical column and quickly cool down the overheated core. Our experiments show that

the proposed scheduler outperforms a Random and a Round-Robin scheduler. On average, we can remove 46% of hardware DTMs and obtain a speedup of 5.11% over the baseline on the thermally homogeneous floorplan. With an enhanced version featured with dynamic tuning scheme, we can remove 46% hardware DTMs and result in 4.78% performance improvement over the base case on the thermally heterogeneous floorplan.

The remainder of this paper is organized as follows. Section 2 discusses previous related works. Section 3 elaborates the motivation of our thermal-aware heuristic algorithms. Section 4 compares our proposed scheduling algorithm with other alternatives. Section 5 introduces the experimental methodology. Section 6 reports the results and compares different algorithms. Section 7 concludes this paper.

2 PRIOR WORK

There have been many works recently investigating the performance potential and the challenges in 3D CMP designs. Mysore et al. [20] proposed to stack on top of a normal processor a profiling die that can identify memory leakage, perform diagnosis etc. to save the area and power on the main die. Black et al. [3] studied the performance advantages and thermal challenges for stacking a large DRAM and SRAM cache on a processor, as well as implementing a processor in two layers. Xie et al. [26] reported that the peak temperature in a 3D chip of 2 layers and one die per layer can be as high as 125°C. More importantly, there is only a difference of a couple of degrees, in the worst case, between the hotspots in the top die and the bottom die. This indicates a strong thermal correlation among adjacent layers in a 3D processor. To ensure better heat dissipation in a 3D chip, Puttaswamy et al. proposed a “Thermal Herding” design [23] which lowers the power of the chip by splitting individual function unit blocks across multiple layers, and places the most frequently switched part, or activity, closest to the heat sink. Alternatively, adding thermal vias can also alleviate the thermal conditions within a 3D chip. Goplen et al. [10] studied that proper placement of thermal vias in 3D IC design can obtain a maximum of 47.1% reduction in temperature. In the multicore domain, Loh et al. [19] introduced different approaches for implementing single-core and multicore 3D processors. Particularly, they pointed out that stacking separate cores (in multicore design) can significantly reuse the existing 2D designs, and the interface between the cores needs no more than a few thousand connections.

Compared to the previous work, this paper focuses mainly on software approaches to thermal management in 3D CMP. There have been proposals on OS-assisted thermal management for single core chip. The Hyb-DTM [17] technique controls temperature by limiting the execution of a hot job once it enters an alarm zone. This is achieved by lowering the priority of the hot

job so that the OS allocates fewer timeslices to it and gives cool jobs relatively more timeslices to execute. An ideal simulation study was performed in [18] to show the benefits of interleaving hot and cool job executions. However, neither performance study nor task switching overhead was considered. In the 2D multicore domain, Choi et al. [5] compared and implemented three different task schedulers, heat-balancing, deferred execution, and threading with cool-loops, to leverage temporal and spatial heat slacks among application threads. The proposed mechanisms are implemented in PowerPC5. Chong et al. [6] proposed a 3D MPSoC thermal optimization algorithm that conducts task assignment, scheduling, and voltage scaling for a set of real-time workloads. The goal was to slowdown the workloads as long as the deadlines are met. This is quite different from our approach which focuses on best performance and low thermal profile.

3 MOTIVATION AND RATIONALE

3.1 A representative floorplan

There have been a number of 3D CMP floorplans, as shown in Figure 1 (a)-(c), proposed in literature [1], [3], [19]. In these figures, cores are stacked on each other, with extended cache or memory in between. We observed that for a 3D stacked chip to be scalable in layer count, it is inevitable to encounter more than one active cores in one vertical core column, no matter how the active cores and cache banks are placed in the floorplan. Further, if we look at the distance of each core stack to the heatsink (on either the top or bottom of the chip), we can classify these floorplans into two categories.

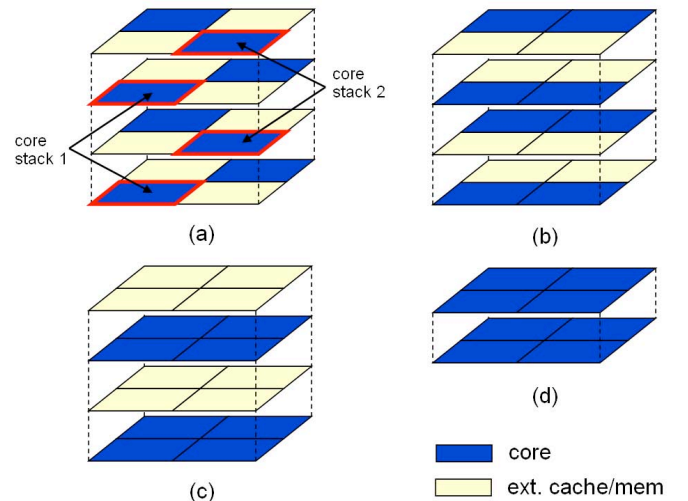


Fig. 1. 3D chip multiprocessor floorplan options.

Figure 1(a) and (b) represent the first category in which the distance of some core stacks, e.g. core stack 1 in (a), to the heatsink is different from others such as core stack 2. These floorplans are *thermally heterogeneous*, meaning that the heat dissipation property of different core stacks is different. For example, if the heatsink is on the bottom

of the stacked chip (as illustrated in Figure 2), core stack 2 is further away from the heat sink than core stack 1. Thus, heat dissipation for cores in stack 2 will be more difficult than those in stack 1. In contrast, Figure 1(c) has a rather homogeneous thermal property because all cores are equally distant from the heat sink. Our preliminary work [28] focused only on homogeneous floorplans while this paper considers both.

Despite these distinctions among different floorplans, they still share some important property. The heat from any core can quickly propagate vertically to other cores above and below. For all these floorplans, the cache layers almost serve as heat conductance between the core layers. Considering this commonality among various 3D floorplans, we choose to use the floorplan in Figure 1(d) as a representative to first introduce the general rationales behind our scheduling algorithms. Then, we will discuss details of our algorithms for homogeneous and heterogeneous floorplans respectively. In Figure 1(d), there are two layers, and each layer contains four cores. The cache banks are subsumed within each core.

3.2 Vertically adjacent layers have strong thermal correlations

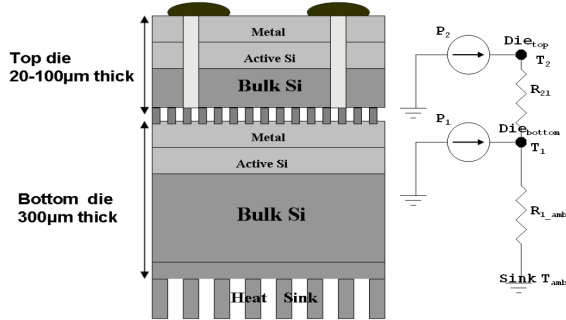


Fig. 2. A face-to-back 3D die stacking structure (adapted from [3]), and the corresponding thermal model.

Similar to a regular 2D processor where heat dissipates mostly in the vertical direction [13], 3D chips also have better heat conductivity in vertical than horizontal direction. This implies that vertically adjacent cores have larger thermal impact among each other than horizontally adjacent cores. We will use a simple heat transfer model to explain this phenomenon. Figure 2 shows a basic two-layer 3D chip structure (adapted from [3]). We use a face-to-back bonding technology for better scalability in layer count. The top layer is thinned for better electrical characteristics and improved physical construction of the through silicon vias for power delivery and I/O. A thin die also has better heat conductivity than a thick die such as the bottom die. As we can see, the distance between the two active silicon dies are very small ($< 100\mu m$). This directly determines the high heat conductivity between the two adjacent dies. The heat transfer model for this 3D chip is shown on the right of the figure. Here one die is modeled using

one node. Its temperature and power are denoted as T and P respectively. R_{21} represents the thermal resistance between the two nodes. R_{1_amb} represents the thermal resistance between the bottom node and the ambient air. We omit the thermal capacitance here to model only the steady state temperature (In our experiments later, both thermal resistance and capacitance are modeled.). Let T_1 and T_2 be the temperature (relative to the ambient air) in the bottom and top node respectively. Then,

$$T_1 = R_{1_amb}(P_1 + P_2) \quad (1)$$

$$T_2 = R_{1_amb}(P_1 + P_2) + R_{21}P_2 \quad (2)$$

Hence, the temperature difference between the two nodes is $R_{21}P_2$. From the parameter used in literature [3], [7], [13], [24], R_{21} is $0.0108 - 0.0159K/W$. P_2 represents the power generated by the entire die. This value is in the range of $40 - 70W$ for a typical single-core processor. Therefore, the temperature difference between the top and bottom die is merely a $0.43 - 1.11K$.

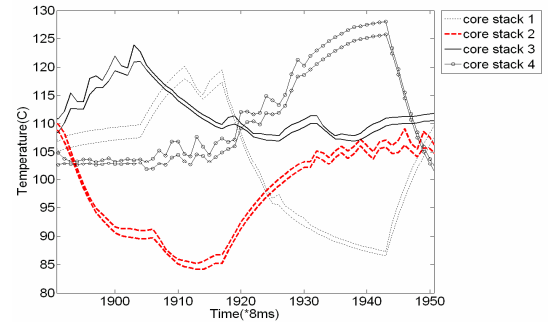


Fig. 3. Thermal correlation between adjacent dies.

Such a strong thermal correlation between the two adjacent dies can also be demonstrated from our simulation. Figure 3 shows a typical thermal profile of running eight threads concurrently on eight cores as floorplaned in Figure 1 (the experimental setup will be introduced in Section 5). Here eight threads are eight different benchmarks chosen from the benchmark suite we use. We refer to vertically aligned two cores as a core stack. We can see from Figure 3 that there are four distinct clusters of temperature curves. Each cluster has drastically different variations from others. However, each cluster has two lines that are very close to each other. Their variations are almost always synchronized. The four clusters correspond to the four core stacks in the floorplan. And the two lines in each cluster correspond to the temperature variation of the two cores per stack. This experiment shows clearly the strong correlation between adjacent dies, as the temperatures for different core stacks hardly have any dependencies among them, but within each core stack, the temperatures of the two cores are strongly correlated. Such correlation can still be observed for a 4-layer floorplan in our experiments, as the intermediate thin cache layers serve as good heat conductors among their vertical core neighbors.

3.3 The die layers further from the heat sink are usually hotter

Not only are the cores in a stack strongly correlated in their temperatures, but also the ones on the top are usually hotter than those near the bottom. This has also been noted in the literature for steady state temperatures [2], [19]. For clarity, we refer to the cores further from the heat sink as “top” cores, as illustrated in Figure 2. The intuition is that the bottom cores are closer to the heat sink, therefore, their heat can be removed more quickly. Here we give a more analytical analysis taking into account the thermal capacitance as well.

Suppose in the thermal model depicted in Figure 2, the thermal capacitance between the top die and ambient air is C_2 . Then,

$$\frac{T_2 - T_1}{R_{21}} = P_2 - C_2 \frac{dT_2}{dt}, \quad (3)$$

As mentioned earlier, P_2 , which represents the power of a modern processor, has a typical value range of 40 – 70W. C_2 represents how quickly temperature changes from the top die. For a thin die within $100\mu m$ in a 2-layer 3D chip, the thermal capacitance is reported as $23.6 - 37.4 mW \cdot s / K$ [3], [24]. dT_2/dt is the temperature change rate within a short time. From our experimental experience, and many other results in the literature, temperature varies slowly with time. For example, we observed a less than $6^\circ C$ increase in temperature in a $8ms$ window using Hotspot 3.0.2 for 3D chips. Hence, the right hand side of equation 3 is usually positive with a range of 12 – 52.3W. Therefore, T_2 is usually higher than T_1 .

We also performed simulations to testify the above observation. We intentionally put the coolest job (lowest average temperature in a 2D chip) in our benchmark suite on the top die, and the hottest job on the bottom die in a 2-core stacked 3D chip setting. The temperatures of the two cores are shown in Figure 4. We can see that the top core has higher temperatures than the bottom layer almost always. Such an observation serves as a guideline to the development of our heuristic scheduling algorithm.

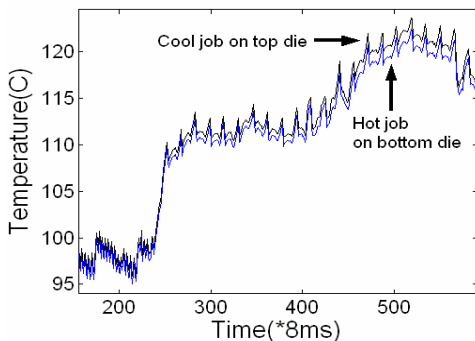


Fig. 4. Demonstration of the top die being hotter than the bottom die.

4 SCHEDULING ALGORITHMS

The strong correlations among the cores in one stack leads to a scheduling that considers the entire stack as a whole. The fact that top cores are hotter than the bottom cores suggests that threads within a core stack should be placed with care. Furthermore, we take advantage of this observation and introduce a new voltage/frequency scaling mechanism that results in the fastest temperature drop within the shortest amount of time, once the peak temperature within a stack reaches the thermal threshold. In this section, we present a sequence of thread scheduling algorithms.

Since we have two categories of floorplans, we will select a representative homogeneous floorplan as shown in Figure 1(c), and a representative heterogeneous floorplan as shown in Figure 1(a). Both the homogeneous and heterogeneous floorplan will be applied with five algorithms: Baseline, Random, Round-robin, Balancing-by-core, and our proposed Balancing-by-stack algorithm.

4.1 The baseline

We use the Linux 2.6 scheduler [4] as our baseline algorithm. In this scheduler, each core has a task queue that keeps track of all running tasks on that core. Each queue contains two priority lists: active and expired list. At runtime, the core selects to execute the tasks in the active list, according to some policy. Once a task uses up its time quota, it is moved to the expired list. If all tasks are in the expired list, an epoch has finished, and the scheduler iterates the process by swapping the two lists. Each task in the active list has 10 – 200ms of CPU cycle quota, depending on its own priority. By default, the core switches to a different task every 100ms. Thus, in our 8-core 3D chip, upon the scheduling interval of every 100ms, the scheduler selects a task from each core’s active list according to its original policy, and then assigns it to a different randomly selected core.

This algorithm is simple, and has low context switch overhead compared to other algorithms introduced later. However, it may run into the risk of putting two hot tasks into the same core stack, which may lead to extremely high temperature that results in long and harsh voltage/frequency scaling penalty to both tasks. Moreover, once a poor scheduling has been made, it stays in that condition for a long period of time (100ms until the next scheduling time), exacerbating the already serious thermal condition within the chip.

4.2 Random (Baseline+)

A quick fix of the baseline scheduler is to increase the scheduling frequency. In the normal Linux OS, any context switch interval between 10 – 200ms may be used [4]. A minimum of 10ms is recommended to avoid unnecessary context switch overhead. We used 8ms as our scheduling interval mainly due to experimental restriction on collecting the power traces. Also, 8ms is

close to the thermal constant of the core under testing. However, the algorithm can be directly applied to any scheduling interval recommended in Linux such as $10ms$ if those restrictions don't apply. Further, we take into account the extra context switch overhead using an $8ms$ scheduling interval during our experiments. We performed a real machine measurement on the time required to perform a single context switch. For an $8ms$ interval, it is $\sim 0.44\%$, a mild penalty that can be easily offset by the performance gain from a better scheduling.

With the improved baseline scheduling algorithm (termed Random to reflect the scheduling decision), the chip can exit a poor thermal condition due to an unwise scheduling more quickly, resulting in less harmful impact.

4.3 Round-Robin

The Random scheduler may result in uneven distribution of power and temperature as tasks are assigned randomly to any core. A Round-Robin scheduler (RR) can overcome this by rotating tasks among cores in a fixed order periodically. Therefore, after N iterations where N is the number of cores, each task has executed on every core for one scheduling interval, e.g. $8ms$. This can help balancing the power and temperature distribution in the long run.

4.4 Temperature balancing by core

An alternative way to balance the heat among the cores is to explicitly arrange the tasks according to their power consumption and the core temperatures. Essentially, a high power task should be assigned to a low temperature core. At each scheduling point, the scheduler sorts the power consumption of all tasks and the current temperature of each core. It then assigns the task with the highest power to the coolest core, the 2^{nd} highest power to the 2^{nd} coolest core, and so forth.

Such a mechanism should perform a better job in balancing the temperature distribution among cores than RR. However, recall that there is a strong thermal correlation between two adjacent layers, and the cores in one stack have only a small difference in temperatures. This implies that if a core stack contains the hottest core, it probably also contains the 2^{nd} hottest core. When the temperature Balancing-by-core algorithm is applied, the tasks with the lowest and 2^{nd} lowest power are scheduled to this hot core stack. Similarly, the tasks with the highest and 2^{nd} highest power will be scheduled to the coolest core stack. After that, the hottest/coolest core stack will have the largest temperature drop/rise, which may lead to temperature oscillations and task thrashing between those two stacks, potentially leading to more thermal emergencies. In that case, a RR, or a Random algorithm may be a better solution.

Another issue with this mechanism is how the power consumption of each task is obtained. Recently, there has been proposals on obtaining the runtime power

consumption of an application through probing the *performance counters* in a processor [15]. We also adopt this approach and assume that each core is equipped with such counters that can be used for power estimation. Note that our power estimation need not be very accurate, as we only need the sorted order of the power, not the absolute values.

4.5 Temperature balancing by stack

The core-based temperature balancing algorithm can create thrashing of tasks between the hottest core stack and the coolest core stack, as we analyzed earlier. This is because the algorithm, while trying to balance the temperatures among all cores, treats each core independently. However, as adjacent dies have strong temperature correlations, cores in the same stack should indeed be considered together. Intuitively, we can assume that each stack is a "super" core that has cores with similar temperatures. Hence, scheduling of the tasks within three dimensions can be reduced to scheduling of "super" tasks within two dimensions. Apparently, a super task is a set of tasks that are assigned to a super core, i.e., a core stack.

4.5.1 Algorithm for homogeneous floorplans

We treat homogeneous and heterogeneous floorplans differently in this algorithm as their super cores have different thermal property. We first elaborate on the algorithm for homogeneous floorplan.

Super tasks. Let L be the number of layers in a 3D chip, and N be the number of cores per layer. As a super core contains L cores, a super task should also contain L tasks and there are N super tasks. The scheduling of N super tasks among N super cores is now simply a 2D problem, where a balanced temperature distribution is desired. Hence, we first balance the power among super tasks, i.e., let each super task have about the same power, and then balance the temperatures among super cores by scheduling a relatively high power super task onto a relatively cool super core.

To balance the power among super tasks, we first sort the powers of all $N \times L$ tasks. Let B_{1-N} be N initially empty bins. We will fill powers into these bins such that each bin will contain L tasks, and the total powers of each bin are about the same. In descending order of powers, we put each power value into a bin that has the smallest current total power among all bins. This policy attempts to reduce the gap between the smallest and the largest total power in each step, in order to generate a relatively balanced total power across N bins. Finally, all powers within a bin form a super task. We remark that our policy is only a heuristic as an optimum solution may require an exhaustive search. We aim for a simple, effective, yet low-complexity heuristic because the scheduler makes the decision at runtime.

Task distribution among and within super cores. The goal of producing super tasks is to generate relatively balanced power distribution across super cores. Once the super tasks are formed, we sum up the temperatures of all L cores in a super core, and sort them. Similar to the previous procedure, we assign the hottest super core with the super task of the lowest power, and so on. Figure 5 shows an example of scheduling 8 tasks onto a 2-layer, 4-core-per-layer, 3D chip. Step (a)-(c) depict the procedure except for how tasks within a super task are allocated onto different cores within a stack.

As discussed earlier, the top cores are usually hotter than the bottom cores in a core stack. Hence, we should allocate tasks of higher powers onto the bottom cores for better heat removal, and tasks of lower powers onto the top cores. For example, if the temperatures of the cores from bottom up are strictly increasing, then the tasks allocated to them should have strictly decreasing powers from bottom up. Figure 5's last step illustrates this policy in a two-layer floorplan.

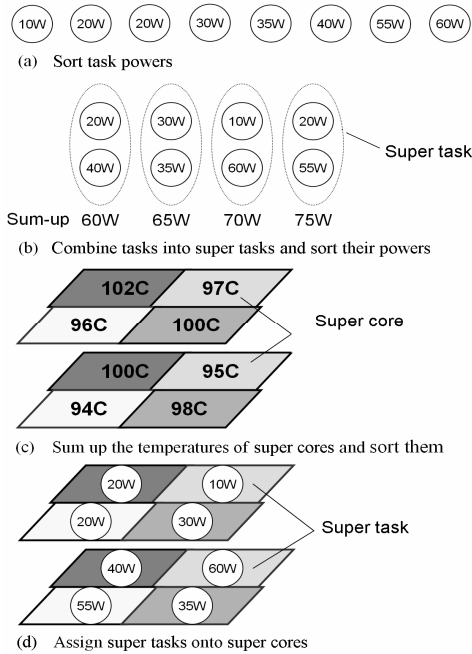


Fig. 5. The temperature balancing-by-stack algorithm.

Scheduling procedure. To sum up, on every scheduling interval (8ms in our case), the scheduler performs the following steps:

- 1) Sort the powers of all tasks. Form super tasks. Sort the power sums of the super tasks from *low to high*.
- 2) For each super core, sum up the temperatures for all cores. Sort the temperature sums for all super cores from *high to low*.
- 3) Create a sequential one-one mapping between the sorted super tasks and sorted super cores.
- 4) In each super core, allocate the tasks in their increasing power order onto the cores with decreasing temperature order.

Our algorithm involves mostly sorting of the powers and temperatures. Therefore, its time complexity is $O(NL \log(NL))$.

4.5.2 Algorithm for heterogeneous floorplans

The major difference in heterogeneous floorplans is that different super cores have different heat dissipation capability due to their varying distances to the heatsink. For this reason, even if two super cores are of the same present temperature, the same super task assigned to them will result in different future temperatures. Therefore, unlike the algorithm for homogeneous floorplans where the total power among super tasks should be well balanced, the task bundling in heterogeneous floorplan should intentionally create power imbalance to generate balanced temperature distribution among super cores. However, it is difficult to estimate how much power difference we should create among super tasks because the future temperature depends on not only power but also the present temperature and duration of the power. Therefore, for a given set of power values, our algorithm forms the super tasks with minimum, moderate and maximum total power difference (denoted as Min-diff, Mod-diff and Max-diff respectively), and dynamically make the selection of super tasks.

Let $P_1 \cdots P_n$ be n powers in ascending order. Super tasks with Min-diff, Mod-diff and Max-diff are formed as follows, assuming each super task contains L tasks:

- Max_diff: $\{P_1, P_2 \cdots P_L\}, \{P_{L+1}, \cdots P_{2L}\}, \cdots$
- Mod_diff: $\{P_1, P_{L+1}, P_{2L+1} \cdots\}, \{P_2, P_{L+2}, \cdots\}, \cdots$
- Min_diff: The principle is to balance the total powers of super tasks. This is identical to the algorithm for homogeneous floorplans (Section 4.5.1).

Intuitively, when the temperature difference among super cores is large, a super task with Max-diff is desired. However, if the power difference among tasks is also large, using the Max-diff may be an overkill. A Mod-diff combination may be sufficient. Therefore, our decision relies on both the temperature gradient (denoted as ΔT) among the super cores and the power range (denoted as ΔP) of the tasks. Let

$$\theta = \frac{\Delta T}{\Delta P}. \quad (4)$$

When θ is small, the temperature gradient (ΔT) is relatively small compared with the power range (ΔP) of the tasks. Super tasks of Min-diff are more appropriate in this situation because we need only to perform mild temperature adjustment. On the other hand, when θ is large, a more aggressive task bundling to create power difference is necessary, hence the selection will favor Max-diff.

During our experiments, we use two heuristic θ values: $\theta_1 = 0.5$ and $\theta_2 = 1$ as the thresholds for choosing different algorithms. The choice of these two values is based on our experimental settings, and may vary with thermal properties of the floorplan. If θ falls in the range

of $[0, \theta_1)$, `Min_diff` will be chosen. If θ is in the range of $[\theta_1, \theta_2]$, `Mod_diff` is selected. If θ is greater than θ_2 , `Max_diff` will be selected.

4.5.3 A new thermal management scheme

A critical component in company with our proposed scheduling algorithm is how to handle thermal emergencies once a core temperature increases above the hardware threshold. Conventionally, such a core will be put to a low power state through DVFS. In a 3D chip, since the top cores are usually hotter, thermal emergencies usually occur in the top layers. Moreover, our scheduler puts cooler tasks on the top layers, which means that those tasks are more likely to undergo DVFS, leading to their degraded performance.

The problems of such conventional thermal management are twofold. First, the cooler tasks could be penalized more often than the hotter tasks, which brings a fairness issue among different tasks. Intuitively, hotter tasks should be restrained by the system due to their potential harmful impact to the chip. Second, applying DVFS to the cooler tasks on the top layers does not yield the same efficiency as in a 2D chip. This is because it takes longer time to cool down the top cores due to their high power neighbors at the bottom. In fact, it is because of those hot bottom tasks are top cores overly heated. Therefore, a more rational thermal management should employ the scalings to the source of an overheating — the bottom cores that are running high power tasks.

More formally, when core A of a super core S is overheated, the thermal management will select core B with the highest power in S to engage DVFS. B may or may not be identical to A . Such a thermal management strategy solves the two above problems effectively. First, cool tasks are not penalized more often than hot tasks because if a cool task becomes a temperature victim, the hot task that caused the problem is penalized. Second, all cores in S , including A and B , are quickly cooled because the total power of S is reduced with the maximum strength. For example in Figure 5, if the super core containing the 20W-40W super task tripped a thermal emergency on the 20W core, and suppose the DVFS reduces the power of a core by half, then our scheme will reduce the total power of this super core to $20 + 40/2 = 40W$, while the conventional thermal management will only reduce it to $20/2 + 40 = 50W$. As we can see, if DVFS is applied to a relatively low power task, the result is inferior because a task is being penalized, but the total power in the chip is not reduced as much. This is often the case for the temperature Balancing-by-core scheduler as it tends to allocate cool tasks on the top layer (since it is usually hotter).

As a result, our mechanism brings down the temperature of the hotspot at the highest speed, resulting in minimum penalty to the overall performance of this super core. We will show later that our proposed temperature balancing-by-stack scheduling algorithm with improved thermal management results in the much less amount of

thermal emergencies and the much better performance among all previous schemes.

5 EXPERIMENTAL METHODOLOGY

5.1 Floorplan setup

Our detailed experiments are conducted on floorplans as depicted in Figure 1(a) and (c). Each floorplan has four layers and a total of eight cores. We simulated 8 P4 Northwood cores in 3.0GHz clock frequency. Each core has a size of $1.144 \times 1.144 \text{ cm}^2$. The remaining space is left for extended cache or memory. So the die size is $2.289 \times 2.289 \text{ cm}^2$. Other physical parameters of the floorplan are similar to [3].

5.2 Simulation tool and power trace collection

We used Hotspot [13] version 3.0.2 as our simulation tool. We chose the grid model to experiment our 3D floorplan. We substituted the 4th-order Runge-Kutta method with TILTS [12] to generate accurate temperatures at high speed.

Hotspot takes power traces as inputs, and temperature variation within a die is a slower process compared to other metrics such as IPC. Hence, we need to collect extended power traces to model realistic temperature variations such as warming up and cooling down due to task scheduling. As mentioned earlier, we adopt the recently proposed performance counter based method [15], [25] to collect runtime hardware activities of a program on a real machine. We obtained the power model (calibrated) from [15], [25] to produce long power traces for programs from a Linux machine with a Pentium 4 core. The traces contain powers for each functional unit, and all traces are complete execution of the programs in SPEC2K.

For scheduling algorithms that require power information (Balancing-by-core and Balancing-by-stack), we use the power in the last $8ms$ interval to predict the power in the next interval. That is, the scheduling decisions are based on local power information. The scheduler needs not to know whether a program is globally hot or cool. Also, we use the last power predictor in the scheduler due to its simplicity. We experimented with more complex power predictors and found that their overhead, both in time and space, is not appropriate for on-line scheduling [27]. Most of the benchmarks exhibit $\sim 5\%$ power mis-prediction rate. Our experiments show that an error within 5% makes last power prediction accurate enough for the scheduler.

5.3 Benchmark classification

We first ran the power traces of each benchmark to obtain its temperature profile as shown in Figure 6. We next classified these benchmarks as hot(power-intensive), cool(power non-intensive), and mild(between hot and cool). After that, we created 9 workload combinations, as listed in Table 1, each with one or more

hot tasks. The workload mixes without hot tasks are less thermally critical and thus, are not considered here. In Table 1, when the number of benchmarks in one combination is less than 8, copies of the benchmarks will be created to ensure that every core in the floorplan has one task to run. This resembles the situation of running parallel threads of the same program in multicore processors.

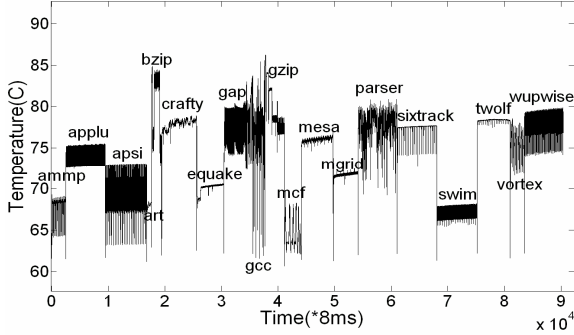


Fig. 6. Temperatures of the benchmark in SPEC2000

TABLE 1
the combination of benchmarks in simulation

HC	crafty mcf
HC	sixtrack swim
HHCC	bzip twolf art ampp
HMMC	wupwise equake applu ampp
HM	gzip mgrid
HM	parser equake
HHMM	crafty gzip mgrid apsi
HHMMCCC	gap twolf equake mgrid vortex ampp art swim
HHHHCCCC	bzip gzip sixtrack wupwise ampp art mcf swim

5.4 DVFS implementation and context switching overhead

We modified Hotspot to incorporate the hardware DVFS. Every $80\mu s$, $1/100$ of a scheduling interval, Hotspot checks if the temperature has trespassed the threshold. If so, the voltage is lowered from $1.3V$ to $1.1V$ and the frequency is reduced by $4/5$. We charge $30\mu s$ of overhead on every voltage/frequency transition. During a DVFS scaling, if the temperature persists above the threshold after one $80\mu s$, the scaling continues and no additional DVFS switch overhead is charged. We do not choose multi-level DVFS scheme to avoid unnecessary switch overhead in every level transition.

Other overheads in our proposed scheduler is mainly the increased number of context switches. We measured this time in a Linux machine by enforcing a large number of context switches between two tasks, and calculating the average switch time from the increased execution time of these two tasks. This quantity in our test machine is $\sim 35\mu s$. Later we will see that our proposed scheduler can still outperform Linux baseline scheduler even with much higher context switch frequency.

6 RESULTS AND ANALYSIS

The metrics we use to evaluate different scheduling algorithms are peak temperature of all cores, the reduction in time that a task stays above the thermal threshold (termed “thermal emergency reduction” in later discussion), and performance improvement in terms of total execution time reduction of all tasks. The peak temperature indicates how well a scheduler can alleviate the worst cases of the thermal condition on-chip. The thermal emergency reduction indicates the capability of a scheduler to control the temperature below the hardware threshold. The performance improvement is the result of both the thermal emergency reduction and the efficiency of lowering the temperature during an emergency. Next, we present the results for homogeneous and heterogeneous floorplans separately.

6.1 Homogeneous floorplan

In the following we will introduce the experiment results on the thermally homogeneous floorplan. Five schedulers, Baseline, Random, Roundrobin, Balancing-by-core, Balancing-by-stack, are tested in the experiments.

6.1.1 Thermal behavior comparison of different schedulers

First, let us see a qualitative comparison among different schedulers on the homogeneous floorplan. Figure 7 shows a close-up of temperature traces for 8 cores running the HMMC workload under different scheduling algorithms. Here, we did not enforce DVFS at the threshold because otherwise, many high temperature curves would be capped at the threshold. As we can see, the baseline algorithm can result in a large temperature gradient across different core stacks. A $\sim 13^\circ C$ difference between the hottest and the coolest core stack is observed in this figure. For Random and RR scheduler, the temperature gradient within the 3D chip gradually reduces because their scheduling interval is $8ms$, much smaller than that in the baseline. Temperature gradient is between $4-6^\circ C$ in these schedulers. Finally, both the Balancing-by-core and our proposed Balancing-by-stack schedulers create the smallest temperature gradient among all cores. The temperature curves of all cores almost overlap entirely. The width of the temperature band is $1-2^\circ C$ only, indicating an excellent balance of temperature among the cores. However, the Balancing-by-core scheduler generates more fluctuation. Note that an ideal temperature balancer would create a $0^\circ C$ among all cores. Hence, our proposed Balancing-by-stack algorithm is only a couple of degrees from the ideal case.

6.1.2 Peak temperature reduction

Balancing the temperatures across the chip can help to reduce the peak temperatures among all cores. Figure 8 shows the peak temperature generated from each

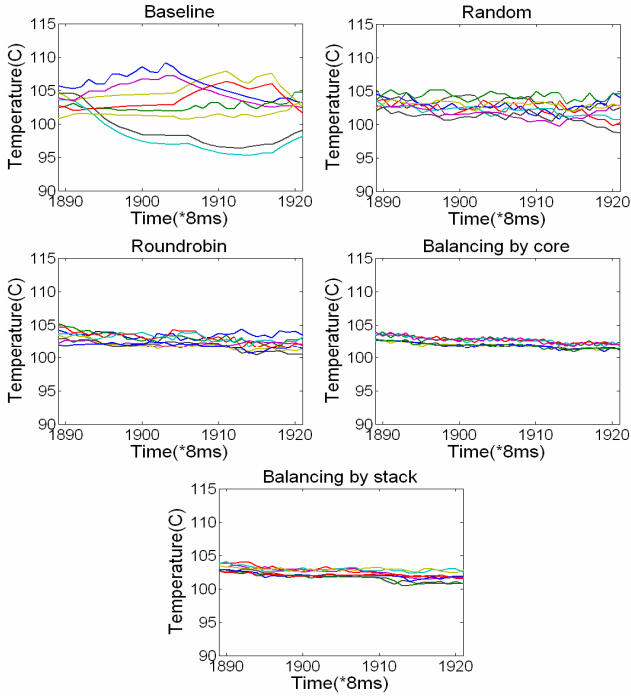


Fig. 7. A zoom-in of temperature variation over time under different scheduling algorithms.

scheduling algorithm assuming there are no DVFS employed (otherwise, the peak temperature is just the thermal threshold). We can see from the figures that baseline algorithm can generate the highest peak temperature of 118.31°C . The Random, RR, Balancing-by-core and Balancing-by-stack can reduce the peak temperature better and better. Our proposed Balancing-by-stack scheduling generates the second lowest peak temperature of 113.71°C , a 4.6°C lower than the baseline and a mere 0.03°C higher than that of Balancing-by-core.

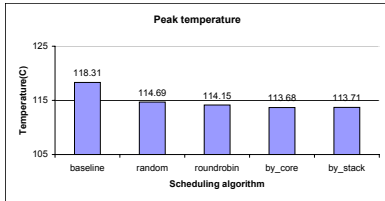


Fig. 8. Peak temperatures of different scheduling algorithms.

6.1.3 Thermal emergency reduction

A direct benefit from scheduling the tasks is the reduced thermal emergency time, i.e. the time a core temperature stays above the hardware thermal threshold. Note that this metric does not necessarily correlate with the peak temperatures reported in Figure 8, which are collected under *no DVFS*. For example, a relatively low peak temperature may still trip DVFS if temperature oscillates around the threshold often. Figure 9 shows thermal emergency time reductions from different algorithms,

normalized to the baseline case. As we can see, the Random, RR and Balancing-by-core can reduce the emergency time by 30.9%, 37.41% and 36.4% on average respectively. Our Balancing-by-stack algorithm removes the most emergency time in 8 cases of 9 benchmarks. An average of 46.23% reduction is observed, with a range of 6.06%-96.04%. Also, the Balancing-by-core algorithm turns out to introduce as much emergency time as RR algorithms even with lower peak temperature. This is because (1) it tends to create temperature oscillations among core stacks as discussed in Section 4.4; and (2) it tends to allocate cooler tasks on the top layer where DVFS is usually engaged for a long time. The consequence is that the overall power in the entire chip is not reduced as much as in other schedulers where high power tasks can be scaled during emergencies. Therefore, a Balancing-by-core scheduler may not be a good scheduling candidate in practice.

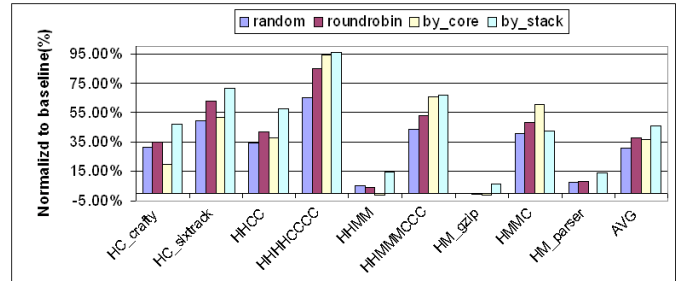


Fig. 9. Thermal emergency time reductions in homogeneous floorplans.

6.1.4 Performance improvement

Corresponding to the thermal emergencies removed, our proposed Balancing-by-stack algorithm achieves the best performance speedup among all algorithms discussed. This is shown in Figure 10. The performance is the total execution time of all 8 tasks in a workload. The results are normalized to the baseline performance. On average, the Balancing-by-stack achieves a 5.11% speedup, while the Random, RR, and Balancing-by-core algorithm achieve 1.45%, 1.72% and 1.65% improvement respectively. This is primarily due to the amount of thermal emergencies our algorithm removed, as well as the high efficiency in handling them with our new thermal management mechanism.

We also notice that in some occasions, the performance may not improve even if the thermal emergency time is reduced. This could happen when the temperature floats around the thermal threshold, but does not increase overly high. In such a scenario, there could be many DVFS triggered, which introduce high transition penalty and overkills the gains from scheduling. For example in the HMMC workload, the Balancing-by-core removed 18.01% of thermal emergency time in the Balancing-by-stack, but its performance is 0.54% worse

than the Balancing-by-stack. Our Balancing-by-stack removes more thermal emergency time than other schedulers in th 8 cases of 9 benchmark combinations, and therefore, achieves the most performance improvement.

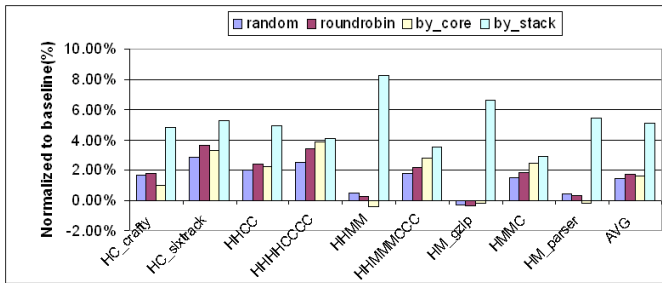


Fig. 10. Performance speedups for homogeneous floorplans.

6.2 Heterogeneous floorplan

In addition to the five algorithms applied to the homogeneous floorplan, two additional algorithms are also tested for heterogeneous floorplans. The first is the revised Balancing-by-stack algorithm with dynamic super task forming mechanisms. The algorithm is designed to tackle the thermal heterogeneity of the floorplan as discussed in Section 4.5.2. The second is a pseudo-optimal algorithm that tests the quality of each discussed algorithm. We term this algorithm a “1-step-optimal” since it tries all task bundling mechanisms and chooses the one that triggers the fewest DTMs in *one next step*. Notice that this is not a true optimal algorithm which would go beyond one-step to enumerate all possible schedules and pick the optimum one (and so is termed “1-step” only). Although it is not realistic to adopt “1-step-optimal” algorithm online due to its complexity, it does indicate the potential for improvement of the discussed algorithms.

6.2.1 Thermal emergency reduction

Figure 11 shows the thermal emergency time reduction for different algorithms normalized to the baseline case. As we can see, Random and RR perform relatively poorly compared with other algorithms because of the heterogeneity in the floorplan. They achieve 12.41% and 12.35% of thermal emergency time reduction respectively. Our proposed dynamic Balancing-by-stack algorithm achieves a total of 46.37% reduction, only 1.92% away from the 1-step-optimal on average, and is better than the remaining algorithms. For example, it removes 9.22% more emergency time than the original Balancing-by-stack algorithm. This indicates that dynamically tuning of the task bundling is very helpful to a thermally heterogeneous floorplan. The Balancing-by-core algorithm is slightly better than dynamic Balancing-by-stack in three cases: HHHHCCCC, HHMMCCCC and HM_gzip. This is because when ΔT and ΔP do

not change, our dynamic Balancing-by-stack algorithm will select the same one from the fixed power bundling schemes, while a slight re-ordering of core temperature will cause Balancing-by-core to form a different and better power bundle more flexibly. Also, Balancing-by-core slightly surpasses 1-step-optimal in HHMMCCCC and HM_gzip workloads. This is because the 1-step-optimal does not generate a global optimal schedule.

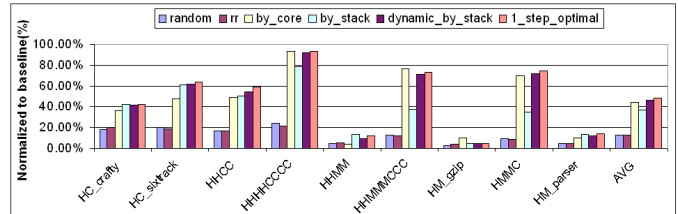


Fig. 11. Thermal emergency time reductions in heterogeneous floorplans.

6.2.2 Performance improvement

Compared with the thermal emergencies removed, our dynamic Balancing-by-stack algorithm achieves the best performance speedups on average among all the algorithms except 1-step-optimal.

Figure 12 shows that Random and RR achieve 0.39% and 0.31% speedup respectively, which is notably lower than 1.45% and 1.72% speedup shown in Figure 10, indicating that Random and RR are not as helpful in heterogeneous floorplans as in homogeneous ones. Balancing-by-stack achieves 2.46% speedup more than Balancing-by-core, though Figure 11 shows it removes 6.74% less thermal emergency time than Balancing-by-core. The reason behind this is that Balancing-by-core tends to generate many overhead in DTM mode switches though the total *time* above the emergency threshold is low which was reported in Figure 11. Finally, the dynamic Balancing-by-stack algorithm achieves the best performance speedup of 4.78% with negligible gap from the 1-step-optimal.

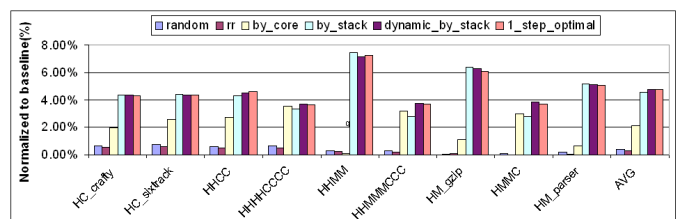


Fig. 12. Performance speedups for heterogeneous floorplan.

7 CONCLUSIONS

We have demonstrated in this paper that OS task scheduling is an effective approach to improve the thermal conditions in 3D chip multiprocessors. It can reduce

the peak temperature within the chip, reduce the thermal emergency triggering amount, and improve the overall performance of the executing tasks.

In particular, we have shown that our proposed scheduling mechanism, Balancing-by-stack, outperforms other intuitive algorithms in the thermally homogeneous floorplan because of the following three properties. First, our scheduler takes into account the high thermal correlations among the layers in one core stack, and schedules tasks in bundles. Second, within every stack of cores, hot tasks are allocated to the layers that are closest to the heat sink for best heat dissipation. Third, upon a thermal emergency, power scaling is engaged in a core stack whose temperature exceeds the threshold, and to the core that generates the largest power in this stack. This can quickly cool down the core stack, reducing the performance penalty imposed to the task.

We also presented a revised Balancing-by-stack algorithm for heterogeneous 3D floorplans. This algorithm refined task bundling mechanism to adapt to the heterogeneity in thermal distribution among the cores. Such a refinement shows its strength over other algorithms in both thermal emergency reduction and performance speedup. Moreover, it achieves nearly the full potential suggested by a local optimal algorithm.

REFERENCES

- [1] Mans Swarthy, Rajeev Balasubramonian, "Exploring the Design Space for 3D Clustered Architectures", *3rd IBM Watson Conference on Interaction between Architecture, Circuits, and Compilers (P=ac2)*, Yorktown Heights, October 2006.
- [2] K. Banerjee, S. Souri, P. Kapur, and K. Saraswat, "3-D ICs: A Novel Chip Design for Improving Deep-Submicrometer Interconnect Performance and Systems-on-Chip Integration," *Proceedings of the IEEE*, vol. 89, pp. 602–633, May 2001.
- [3] Bryan Black, et al., "Die Stacking (3D) Microarchitecture," *MICRO 2006*, 469-479
- [4] D. Bovet, M. Cesati, "Understanding the Linux kernel, 3rd Edition," *O'Reilly Publisher*, November, 2005.
- [5] Jeonghwan Choi, "Thermal-aware Task Scheduling at the System Software Level," *ISLPED'07*, August 27-29, 2007, Portland, Page 213-218.
- [6] S. Chong, L. Shang, and R. P. Dick, "Three-dimensional multi-processor system-on-chip thermal optimization," *Proc. International Conference Hardware/Software Codesign and System Synthesis*, Sep. 2007
- [7] Yangdong Deng, W. P. Maly, "2.5-dimensional VLSI system integration," *IEEE Trans. VLSI Syst.*, 13(6): 668-677 (2005).
- [8] D. Donald, M. Martonosi, "Techniques for multicore thermal management: classification and new exploration," *ISCA*, pp 78-88, 2006.
- [9] Brent Goplen, Sachin S. Sapatnekar "Thermal via placement in 3D ICs," *ISPD 2005*, pp.167-174.
- [10] Brent Goplen, Sachin S. Sapatnekar, "Placement of Thermal Vias in 3-D ICs Using Various Thermal Objectives," *IEEE Trans. on CAD of Integrated Circuits and Systems* 25(4), pp.692-709 (2006).
- [11] Y. Han, I. Koren, C. M. Krishna, "Temptor: A lightweight runtime temperature monitoring tool using performance counters," *the 3rd Workshop on Temperature-Aware Computer Systems, Held in conjunction with ISCA-33*, 2006.
- [12] Yongkui Han, Israel Koren, C. Mani Krishna, "TILTS: A Fast Architectural-Level Transient Thermal Simulation Method," *Journal of Low Power Electronics*, 3(1): 13-21 (2007).
- [13] W. Huang, M. R. Stan, K. Skadron, K. Sankaranarayanan, and S. Ghosh, "HotSpot: A Compact Thermal Modeling Method for CMOS VLSI Systems," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 14(5):501-513, May 2006.
- [14] W. Hung, G. M. Link, Y. Xie, V. Narayanan, and M. J. Irwin, "Interconnect and Thermal-aware Floorplanning for 3D Microprocessors," *the 7th ISQED*, pp. 98-104, 2006.
- [15] C. Isci, M. Martonosi, "Runtime power monitoring in highend processors: methodology and empirical data," *MICRO*, pp. 93-104, 2003.
- [16] J. Joyner, P. Zarkesh-Ha, and J. Meindl, "A Stochastic Global Net-length Distribution for a Three-Dimensional System on Chip (3D-SoC)," *the 14th IEEE International ASIC/SOC Conference*, 2001.
- [17] A. Kumar, L. Shang, L.-S. Peh, N. Jha, "HybDTM: A coordinated hardware-software approach for dynamic thermal management," *DAC*, pp. 548-553, 2006.
- [18] E. Kursun, C.-Y. Cher, A. Buyuktosunoglu, P. Bose, "Investigating the effects of task scheduling on thermal behavior", *the 3rd Workshop on Temperature-Aware Computer Systems, Held in conjunction with ISCA-33*, 2006.
- [19] Gabriel H. Loh, Yuan Xie, Bryan Black, "Processor Design in 3D Die-Stacking Technologies," *IEEE Micro*, 27(3): 31-48 (2007).
- [20] Shashidhar Mysore, Banit Agrawal, Navin Srivastava, Sheng-Chih Lin, Kaustav Banerjee, Timothy Sherwood, "Introspective 3D chips," *ASPLOS 2006*, pp. 264-273.
- [21] M. D. Powell, M. Goma, T. N. Vijaykumar, "Heat-and-run: leveraging SMT and CMP to manage power density through the operating system," *ASPLOS 2004*, pp. 260-270.
- [22] Kiran Puttaswamy, Gabriel H. Loh "Thermal analysis of a 3D die-stacked high-performance microprocessor," *ACM Great Lakes Symposium on VLSI*, 19-24, 2006.
- [23] Kiran Puttaswamy, Gabriel H. Loh, "Thermal Herding: Microarchitecture Techniques for Controlling HotSpots in High-Performance 3D-Integrated Processors," *HPCA*, pp. 193-204, 2007.
- [24] K. Skadron, K. Sankaranarayanan, S. Velusamy, D. Tarjan, M.R. Stan, and W. Huang, "Temperature-Aware Microarchitecture: Modeling and Implementation," *ACM TACO*, 1(1):94-125, Mar. 2004.
- [25] Wei Wu, Lingling Jin, Jun Yang, Pu Liu, Sheldon X.-D. Tan "A systematic method for functional unit power estimation in microprocessors," *DAC 2006*, 554-557.
- [26] Yuan Xie, Gabriel Loh, Bryan Black, and Kerry Bernstein, "Design Space Exploration for 3D Architecture," *ACM Journal of Emerging Technologies for Computer Systems*, Vol. 2. No. 2, pp.65-103, April 2006.
- [27] Jun Yang, Xiuyi Zhou, Marek Chrobak, Youtao Zhang, Lingling Jin, "Dynamic Thermal Management through Task Scheduling," *ISPASS*, pp.191-201, 2008.
- [28] Xiuyi Zhou, Yi Xu, Yu Du, Youtao Zhang, Jun Yang, "Thermal Management for 3D Processor via Task Scheduling," *International Conference on Parallel Processing*, Sept. 2008.