

Interconnect and Thermal-aware Floorplanning for 3D Microprocessors *

W.-L. Hung, G.M. Link, Yuan Xie, N. Vijaykrishnan, and M. J. Irwin
The Pennsylvania State University, University Park, PA 16802, USA
{whung,link,yuanxie,vijay,mji}@cse.psu.edu

Abstract

Interconnects are becoming an increasing problem from both performance and power consumption perspective in future technology nodes. The introduction of 3D chip architectures, with their intrinsic capability of reducing wire length, is one of the promising solutions to mitigate the interconnect problem. While interconnect power consumption reduces due to the adoption of 3D designs, the stacking of multiple active layers leads to higher power densities. Thus, high peak temperatures are of major concern in 3D designs. Consequently, we present a thermal-aware floorplanner for 3D architectures. In contrast to most prior work, our floorplanner considers the interconnect power consumption in exploring a thermal-aware floorplan. Our results show that excluding interconnect power can result in peak temperatures being underestimated by as much as 15°C in 90nm technology. Finally, we demonstrate that our floorplanner is effective in lowering peak temperatures using a microprocessor design and four MCNC designs as benchmarks.

1. Introduction

Aggressive scaling of process technologies has enabled feature sizes to shrink continuously. While the performance of gates has been improving, there are concerns about the performance of wires in scaled technologies. As Ho et al. stated in [1], there are two kinds of wires. One category of wires is local wires inside logic modules that scale in length; another category is long wires that do not scale accordingly with technologies. In order to keep the delays of these long wires tractable, repeaters and flip-flops are inserted to prevent performance degradation. However, these additional components have detrimental impacts on interconnect power dissipation. Consequently, intermediate and global interconnects of current microprocessors contribute to a major portion of power consumption and also serve as impediments for better performance. Hence, many research efforts are devoted to seeking solutions which can overcome the limitation of wiring requirements for present and future chip designs.

There are various reasons why the interconnect has become the center of attention in terms of power consumption and performance of a chip. First, interconnects, unlike transistors, have not scaled down exponentially as we move to nano-meter era, and thus, a larger portion of total chip capacitance comes from the interconnect capacitance. Second, long interconnects, compared to the scaled transistors, are becoming exceptionally long. As a result, performance degradation is inevitable. Finally, as Kapur et al. states in [2] that with the introduction of repeaters and vias to compensate the performance lost, interconnect power consumption almost doubles.

One technique being actively researched to alleviate the problems of interconnects is the use of three dimensional (3D) integrated circuits. There are numerous novel 3D architectures under development. In this paper, we have considered one of the promising styles of 3D technologies: *wafer-bonding technology* [16]. Wafer-bonding technology holds active device layers after processing each active device layer separately and the connections are provided by 3D vertical vias. 3D architectures are effective in reducing wire length from the geometric point of view. If modules were carefully placed on a 3D chip, many of the long interconnects and the large power consumption associated with them can be reduced. A major concern in the adoption of 3D architecture is the increased power densities that can result from placing one computational block over another in the multi-layered 3D stack. Since power densities are already a major concern in 2D architectures, the move to 3D architectures could accentuate the thermal problem. However, 3D chips could offer some respite due to reduced interconnect power consumption due to the shortening of many long wires. Consequently, it is essential for one to investigate how the increased power densities of the stacked blocks and the reduced interconnect power consumption combine to influence the thermal behavior.

In this paper, we explore a floorplanning algorithm that attempts to reduce the peak temperatures of a 3D design. High temperature has adverse impacts on circuit performance. The interconnect resistance becomes larger while the driving strength of a transistor decreases with increasing temperature. In addition, leakage power has an exponential dependence on the temperature and can even result in ther-

*Acknowledgments: This work is supported in part by NSF Career 0093085, MARCO/DARPA-GSRC, and the TTC.

mal runaways. Finally, higher temperatures are known to accelerate failure mechanisms and reduce the lifetime of the device. While there have been prior proposals to do thermal aware floorplanning, the interconnect power consumption is often not considered. Our results in this work indicate that excluding interconnect power can result in peak temperatures to be under estimated by as much as 15°C in 90nm technology. In order to capture the actual influence of the interconnect, a detailed model of a microprocessor is used to model the interaction between the functional modules and relative interconnects. To the best of our knowledge, our approach is the first to use a real microprocessor design and to account for interconnect power consumption when performing 3D thermal-aware floorplanning.

The rest of the paper is organized as follows. Section 2 reviews previous work related to this research. Section 3 explains how to estimate temperature. Section 4 presents our approach in evaluating the impacts of interconnect during floorplanning process. Section 5 presents and discusses experimental results. We conclude this paper in the last section.

2. Related Work

Recent studies on microarchitecture focused on the performance issue. Full chip power modeling including functional units and interconnects is rarely considered. In [4], the performance-meeting requirement was involved during the microarchitecture evaluation. However, the power factor was not considered and an unsuitable set of benchmarks for representing actual microprocessor architecture were used. Ekpanyapong et al. [5] proposed a profile-driven microarchitectural floorplanning in evaluating placement of functional units of the microprocessor, with SimpleScalar incorporated. Unfortunately, the only concern of their work is to satisfy performance demand during the floorplanning process.

Interconnect and interconnect buffers are now first-order timing and power considerations in VLSI design [6]. This change has imposed challenges across all design levels. It is no longer possible to accurately produce the power consumption and performance of a design without prior knowledge about its floorplan to predict the structure of its interconnect. A number of researchers have considered the impacts of chip-level interconnect in power and performance aspects [2, 3, 7].

Tools for modeling thermal effects on chip-level placement have been developed. Recently, Cong et al. [8] proposed a thermal-driven floorplanning algorithm for 3D ICs. Chu and Wong [9] proposed using a matrix synthesis problem (MSP) to model the thermal placement problem. A standard cell placement tool to even thermal distribution has been introduced by Tsai and Kang [10] with their proposed compact finite difference method (FDM) based temperature

modeling. In [17], thermal effect was formulated as another force in a force-directed approach to direct the placement procedure for a thermally even standard cell placement. Another design metric, reliability, was taken care of in [11] when doing a multi-layer System-on-Package floorplanning, thermal issue was neglected, though. Nevertheless, interconnect power factor is never the center of attention in these floorplanning/placement techniques.

Some researchers have looked at the thermal problem in 2D microarchitectural floorplanning. For example, Han et al. in [19] used an Alpha floorplan; however, they use hypothetical interconnects as opposed to our interconnect information extracted from a real processor design. The tradeoff between performance and temperature is explored in [18]; nevertheless, a 5% of extra space is required to get a lowering temperature solution with their HotFloorplan tool. Both the above work neglected the interconnect power consumption and they do not consider floorplanning of 3D architectures. Although [20] targets on the same problem as our work, their approach of 3D partition followed by mixed ILP floorplanning approach may generate sub-optimal solution with prohibitively high runtime overhead.

3. Temperature Estimation

Skadron et al. [12] proposed a thermal modeling tool called HotSpot, which provides temperature estimation of a microprocessor at the functional module level by employing the principle of thermal-electrical duality. An RC network of thermal capacitances and resistances of functional modules are constructed and then temperatures at the center of functional modules are calculated by using circuit-solving techniques. The inputs to HotSpot are the floorplan and the power consumption number of individual modules, and the specifications of heat spreader and heat sink are also provided to define the heat-removing ability.

While the HotSpot tool was originally intended to be a fast means of modeling temperatures of 2D architectures, as we move from 2D to 3D, HotSpot is no longer capable of generating an accurate temperature profile. As such, we extend this original model by including a variable number of additional levels, each composed of both a silicon layer and an inter-silicon "glue" material in our new tool called *HS3D*. This tool was validated for accuracy by first comparing to the original HotSpot tool, which showed identical temperature estimates. To validate the multi-layer modeling, HS3D was compared to a commercial FEM tool, Flotherm, which showed an average temperature mis-estimation of 3° and a maximum deviation of 5°. This compares well to the validation of many other thermal estimation libraries. In addition, we note that the maximum discrepancies occurred 'downstream' of the airflow modeled in Flotherm. As HS3D uses a simple resistive model for heat transfer to ambient, heat transfer from airflow back to the heatsink is not modeled,

resulting in a significant portion of the discrepancy.

4. Experimental Methodology

In this section, we describe our experimental methodology. First, we introduce the Alpha-like, detailed microprocessor architecture. This architecture model is then mapped to OKI 160nm and TSMC 90nm libraries by Design Compiler and then placed and routed through First Encounter to extract power consumption numbers of functional modules, interconnects among modules, and area of each module as well. This information will next be used in guiding the algorithm to generate a solution with the lowered hot spot temperatures from our 2D/3D floorplanner.

4.1 Processor model

In order to effectively explore the architecture-level interconnect power consumption of a modern microprocessor, we need a detailed model which can act for the current-generation high-performance microprocessor designs. Based upon this requirement, we have used IVM [14], a Verilog implementation of an Alpha-like architecture (denoted as Alpha in the rest of the paper) at register-transfer-level, to evaluate the impacts of both interconnect and module power consumptions at the granularity of functional module level. A diagram of the processor was shown in Figure 1. Each functional block in Figure 1 represents a module used in our floorplanner. The registers between pipeline stages are also modeled but not shown in the figure.

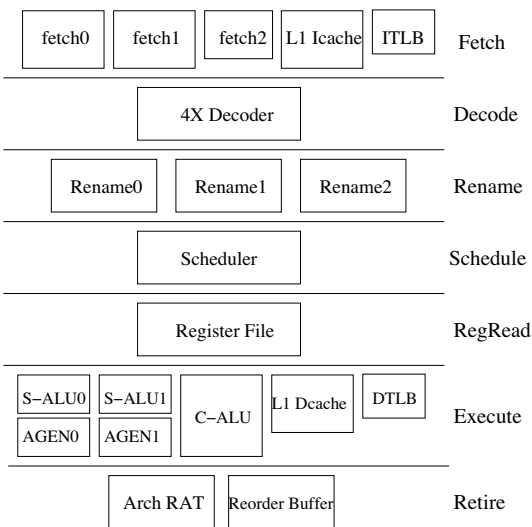


Figure 1. Processor model diagram.

4.2 B*-tree floorplan model

The floorplanner used in this work is based on the B*-tree representation. The B*-tree is proposed by Chang et al. in [13]. While the original B*-tree structure was developed

and used for the 2D floorplanning problem, we modify the perturbation function to handle 3D floorplans in this work. There are six perturbation operations used in our algorithm and they are listed below:

- (1) Node swap, which swaps two modules.
- (2) Rotation, which rotates a module.
- (3) Move, which moves a module.
- (4) Resize, which adjusts the aspect-ratio of a soft module.
- (5) Interlayer swap, which swaps two modules at different layers.
- (6) Interlayer move, which moves a module to a different layer.

The first three perturbations are the original moves defined in [13]. Since these moves only have influence on the floorplan in single layer, more interlayer moves, (5) and (6) are needed to explore the 3D floorplan solution space.

4.3 Simulated annealing engine

A simulated annealing engine is used to perturb floorplanning solutions. The inputs to our floorplanning algorithm are the area of all functional modules as shown in Figure 1 and interconnects among modules. However, the actual dimension of each module is unknown *a priori* except for its area before placement. That is, we have to treat them as *soft* modules. Thus, we provide the choice of adjusting aspect-ratio as one perturbation operation. During simulated process, each module dynamically adjusts its aspect-ratio to fit closely with the adjacent modules, that is, with no dead space between two modules. A traditional weighted cost representing optimization costs (area and wire length) is generated after each perturbation.

Different from 2D floorplanning, our 3D floorplanner uses a two-stage approach. The first stage tries to partition the blocks into the appropriate layers, and tries to minimize the packed area difference between layers and total wire length using all perturbation operations (one through six listed in previous subsection). However, due to the fact the first stage was trying to balance the packed areas of the different layers, the floorplan of some of the layers may not be compactly packed. The second stage is intended to overcome this problem. Thus, in the second stage, we start with the partitioning solution generated by the first stage and focus on adjusting the floorplan of each layer simultaneously with the first four operations. At this point, there are no interlayer operations to disturb module partition of each layer obtained from stage one.

One problem of 3D floorplanning is the final packed area of each layer must match to avoid penalties of chip area. For example, assuming two layers, L1 and L2, if the final width of packed modules of L1 is larger than the final width of packed modules of L2 and the height of L1 is smaller than that of L2, a significant portion of chip area is wasted due

to the need for the layer dimensions to match for manufacturing. Thus, care must also be taken in both stages of our algorithm so that the dimension of each layer will be *compatible*. Thus, we adopt the concept of dimension deviation $dev(F)$ in [11]. The goal is to minimize $dev(F)$, which tells the deviation of the upper-right corner of a floorplan from the average Ave_x, Ave_y values. The value, Ave_x can be calculated by $\sum ux(f_i)/k$, where $ux(f_i)$ is the x-coordinate of upper-right corner of floorplan i and k indicates the number of layers. The value Ave_y can be obtained in similar manner. Thus, $dev(F)$ is formulated as $\sum_i^{\#layers} |Ave_x - ux(f_i)| + |Ave_y - uy(f_i)|$. The modified cost function for 3D floorplanner can be written as

$$cost = \alpha * area + \beta * wl + \gamma * dev(F) \quad (1)$$

where $area, wl$ are chip area and wire length, respectively.

4.4 Interconnect power distribution

With the detailed model of the Alpha-like microprocessor in Verilog, the actual power consumptions of functional modules and interconnects are extracted by Design Compiler and First Encounter, and used in our floorplanning algorithm. After running our 2D floorplanner, we have the total interconnect length for 2D architecture. Since there is still no available tool can be used to accurately model the power dissipation of all interconnects in 3D architecture, one simple way to obtain this number is through scaling. The approximation is done by scaling 3D interconnect length with the interconnect length from 2D architecture. Thereby, the total interconnect power consumption in 3D is formulated as equation (2),

$$3DP_{int} = 2DP_{int} * (3DL_{int}/2DL_{int}) \quad (2)$$

where $3DP_{int}$ represents the power consumption of all nets in 3D and $3DL_{int}$ indicates the total net length accumulated through 3D floorplanner. Since HS3D only models thermal effects at the per-module level, we need a mechanism to account for those interconnect-induced power consumptions. That is, an approach to distribute the power consumed by each net to the modules is required. We accomplish this goal by taking the intuition that power consumption is relative to capacitance and capacitance is proportional to the module area. Thus, from equation (2), the power value of each net, n_i , is the ratio of the net length of n_i to the total net length multiplied by total net power either in 2D or 3D. Finally, the amount of net power contributes to the connecting module b_j can be stated as follows:

$$Net_{ij} = (n_i/3DL_{int}) * 3DP_{int} * (A_{b_j}/TBA_{n_i}) \quad (3)$$

where Net_{ij} indicates the amount of power from net n_i contributing to module b_j , A_{b_j} represents the area of functional module b_j , and TBA_{n_i} tells the total area of connected modules of n_i .

4.5 Temperature approximation

Although HS3D can be used for providing temperature feedbacks, when evaluating a large number of solutions during simulated procedure, it is not wise to involve the time-consuming temperature calculation every time. Other than using the actual temperature values, we have adopted the power density metric as a thermal-conscious mechanism in our floorplanner. Since the temperature is heavily dependent on power density based on a general temperature-power equation: $T = P * R = P * (t/k * A) = (P/A) * (t/k) = d * (t/k)$, where t is the thickness of the chip, k is the thermal conductivity of the material, R is the thermal resistance, and d is the power density. Thus, we can substitute the temperature and adopt the power density, according to the equation above, to approximate the 3-tie temperature function, $C_T = (T - T_o)/T_o$, proposed in [8] to reflect the thermal effect on a chip. As such, the 3-tie power density function is defined as $P = (P_{max} - P_{avg})/P_{avg}$, where P_{max} is the module with the maximum power density while P_{avg} is the average power density of all modules. The cost function for 2D architecture used in simulated annealing can be written as

$$cost = \alpha * area + \beta * wl + \gamma * P \quad (4)$$

For 3D architectures, we also adopt the same temperature approximation for each layer as horizontal thermal consideration. However, since there are multiple layers in 3D architecture, the horizontal consideration alone is not enough to capture the coupling effect of heat. The vertical relation among modules also needs to be involved and is defined as: $OP(TPm) = \sum (Pm + Pm_i) * overlap_area$, where $OP(TPm)$ stands for the summation of the power density of module, Pm , and all overlapping module m_i with module m and their relative power densities multiplying their corresponding overlapped area.

The rationale behind this is that for a module with relatively high power density in one layer, we want to minimize its accumulated power density from overlapping modules located in different layers. We can define the set of modules to be inspected, so the total overlap power density is $TOP = \sum OP(TPi)$, for all modules in this set. The cost function for 3D architecture is thus modified as follows:

$$cost = \alpha * area + \beta * wl + \phi * dev(F) + \gamma * P + \delta * TOP \quad (5)$$

At the end of algorithm execution, the actual temperature profile is reported by our HS3D tool.

5 Experimental Results

We implemented the proposed floorplanning algorithm in C++. The thermal model is based on the HS3D. The implementation of microprocessor has been mapped to OKI 160nm library by Design Compiler and placed and routed by

Circuit	2D						2D(thermal)					
	wire (um)	area (mm ²)	dead space(%)	peakT	peakT (Int)	run time(s)	wire (um)	area (mm ²)	dead space(%)	peakT	peakT (Int)	run time(s)
Alpha	339672	29.43	3.21	108.13	114.50	81	381302	29.68	4.02	104.13	106.64	213
xerox	542926	19.69	1.76	114.54	123.75	8.24	543855	19.84	2.49	104.71	110.45	13
hp	133202	8.95	1.38	109.58	119.34	10.33	192512	8.98	1.75	106.66	116.91	25.78
ami33	44441	1.21	4.60	127.37	128.21	81	51735	1.22	5.55	116.07	116.97	101
ami49	846817	37.43	5.31	115.79	119.42	98	974286	37.66	5.90	103.46	108.86	240
Average	1	1	1	1	1	1	1.18	1.01	1.25	0.93	0.93	2.08

Table 1. Floorplanning results of 2D architecture.

Circuit	3D					3D(thermal)				
	wire (um)	area (mm ²)	peakT	peakT (Int)	run time(s)	wire (um)	area (mm ²)	peakT	peakT (Int)	run time(s)
Alpha	210749	15.49	126.01	135.11	363	240820	15.94	117.48	125.47	519
xerox	297440	9.76	127.21	137.51	17	294203	9.87	119.64	127.31	47
hp	124819	4.45	125.67	137.90	16	110489	4.50	120.39	134.39	47
ami33	27911	0.613	164.60	165.61	113	27410	0.645	154.65	155.57	347
ami49	547491	18.55	130.22	137.71	386	56209	18.71	124.19	132.69	892
Average	0.68	0.50	1.17	1.18	2.69	0.55	0.52	1.10	1.12	6.01

Table 2. Floorplanning results of 3D architecture.

First Encounter under 1GHz performance requirement. We have also synthesized the microprocessor design to TSMC 90nm low power library in order to exhibit the significance of interconnect power when technologies scaled. There are totally 34 functional modules and 168 netlists extracted from the processor design. The area and power consumptions from the actual layout served as inputs to our algorithm. The experiment was run on a dual Intel Xeon (3.2 GHz, 2GB RAM) machine running Linux. Other than Alpha processor, we have also used MCNC benchmarks to verify our approach. A similar approach in [10] is used to assign the average power density for each module in the range of 2.2×10^4 (W/m²) and 2.4×10^6 (W/m²). The total net power is assumed to be 30% of total power of modules due to lack of information for the MCNC benchmarks and the total wire length used to be scaled during floorplanning is the average number from 100 test runs with the consideration of area factor alone. The widely used method of half-perimeter bounding box model is adopted to estimate the wire length. Throughout the experiments, two-layer 3D architecture was assumed due to a limited number of functional modules and excessively high power density beyond two layers; however, our approach is capable of dealing with multiple-layer architecture.

Table 1 and 2 show the experiment results of our approach when considering traditional metrics (area and wire) and thermal effect. The dead space columns in Table 1 demonstrate the effectiveness of our tightly compact floorplanner with the ability of adjusting aspect-ratio for soft modules. The peak temperature results from these tables also reiterate the importance of including the interconnect power consumption which most prior work ignore. The difference between peak temperatures not considering interconnect power (peakT) and considering it (peakT(Int)) is 6°C on the average. When taking thermal effect into account, our thermal-aware floorplanner can reduce the peak temperature by 7% on the average while increasing wire-length by 18% and providing a comparable chip area as com-

pared to the floorplan generated using traditional metrics. Note that the temperature estimation doubles the floorplanner execution time for the thermal-aware case.

When we move to 3D architectures, the peak temperatures increased by 18% (on the average) as compared to the 2D floorplan due to the increased power density. However, the wire length and chip area reduced by 32% and 50%, respectively. The adverse effect of the increased power density in the 3D design can be mitigated by our thermal-aware 3D floorplanner. We lower the peak temperature by 6% (peakT(Int)) with little area increase as compared to the 3D floorplanner which does not account for thermal behavior. Figure 2 shows the temperature profile of the peak and average temperatures with (peakT(Int), avgT(Int)) and without (peakT, avgT) interconnect power for both 2D and 3D architecture. of the Alpha design. As expected, the chip temperature is higher when we step from 2D to 3D architecture without thermal consideration. Although the wire length is reduced when moving to 3D and thus accordingly reduces interconnect power consumption, the temperature for 3D architecture is still relatively high due to the accumulated power densities from different layers and the diminished chip area. After applying our thermal-aware floorplanner, the peak temperature is lowered to a moderate level through the separation of high power density modules in different layers.

Table 3 shows the ratio of interconnect power to total power consumption. It can be seen that in the 3D case the interconnect power becomes less dominant part of the overall power consumption, only accounting 15% of overall power in both cases as compared to around 21% and 24% (Ther) for the 2D cases. Note that these numbers are lowered than 30% because the total wire length used for scaling is obtained prior including wire factor in the cost function.

In the next experiment, we demonstrate the importance of the impact of interconnect when moving to a more advanced technology. Figure 3 shows the temperature difference between 160nm and 90nm technologies under 2D and

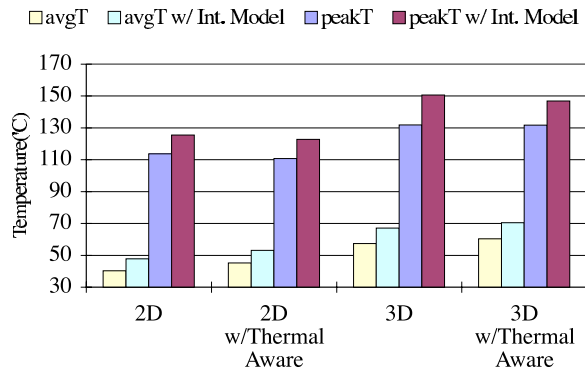


Figure 2. Thermal profiles of Alpha processor for 2D and 3D architectures.

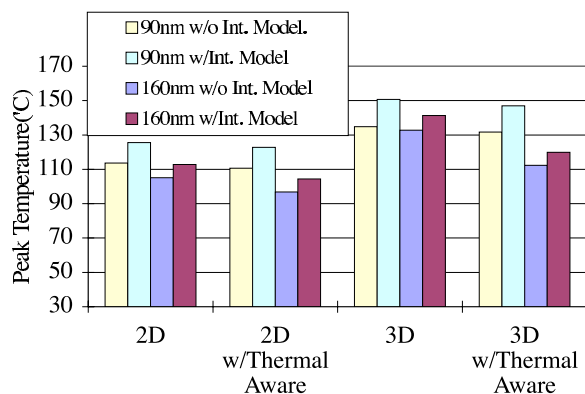


Figure 3. Temperature profiles of Alpha processor for 160nm and 90nm technologies.

3D architecture configurations. As can be seen from the figure, the temperature difference not considering interconnect power can be as high as 15°C in 90nm, which is higher than those (6~8°C) in 160nm technology. Based on this result, we state that more attention should be drawn to interconnect in future technologies and it is imperative to include interconnect power estimates in guiding any thermal-aware floorplanning.

6 Conclusion

In this paper, we have presented a thermal-aware floorplanner for a 3D IC. Our floorplanner is unique in that it accounts for the effects of the interconnect power consumption in estimating the peak temperatures. We have demonstrated the importance of accounting the interconnect power in guiding the floorplanner as the peak temperatures can be underestimated by as much as 15°C in 90nm technology node without including interconnect power. Finally, we have shown the effectiveness of the thermal-aware tool in reducing peak temperatures using one microprocessor design and four MCNC benchmarks.

in %	2D	2D (Ther)	3D	3D (Ther)
Alpha	20.92	26.00	15.15	15.37
xerox	23.16	23.19	14.17	14.04
hp	27.96	27.79	19.97	18.08
ami33	13.44	15.63	8.5	8.3
ami49	22.46	27.96	19.32	19.55
Average	21.58	24.11	15.42	15.06

Table 3. Interconnect power profile.

References

- [1] R. Ho, K. W. Mai, and M. A. Horowitz. The Future of Wires. In *Proc. of the IEEE* 2001.
- [2] P. Kapur, G. Chandra, and K. C. Saraswat. Power estimation in global interconnects and its reduction using a novel repeater optimization methodology. In *DAC 2002*.
- [3] D. Sylvester and K. Keutzer. Getting to the bottom of deep submicron. In *ICCAD 1998*.
- [4] J. Cong, A. Jagannathan, G. Reinman, and M. Romesis. Microarchitecture evaluation with physical planning. In *DAC 2003*.
- [5] M. Ekpanyapong, J. R. Minz, T. Watwai, H.-H. S. Lee, and S. K. Lim. Profile-guided microarchitectural floorplanning for deep submicron processor design. In *DAC 2004*.
- [6] J. Cong and D. Z. Pan. Interconnect performance estimation models for design planning. In *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, pages 739–752, June 2001.
- [7] W. Liao and L. He. Full-chip interconnect power estimation and simulation considering concurrent repeater and flip-flop insertion. In *ICCAD 2003*.
- [8] J. Cong, J. Wei, and Y. Zhang. A Thermal-Driven Floorplanning Algorithm for 3D ICs. In *ICCAD 2004*.
- [9] C. Chu and D. F. Wong. A matrix synthesis approach to thermal placement. In *ISPD 1997*.
- [10] C.-H. Tsai and S.-M. S. Kang. Cell-level placement for improving substrate thermal distribution. In *IEEE Trans. on Computer-Aided Design of Integrated Circuits and System*, vol. 19, pages 253–266, Feb. 2000.
- [11] P. H. Shiu, and S. K. Lim. Multi-layer Floorplanning for Reliable System-on-Package. In *ISCAS 2004*.
- [12] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan. Temperature-aware microarchitecture. In *Proc. of international symposium on Computer architecture*, pages 2–13, 2003.
- [13] Y.-C. Chang, Y.-W. Chang, G.-M. Wu, and S.-W. Wu. B*-trees: a new representation for non-slicing floorplans. In *DAC 2000*.
- [14] IVM. <http://www.crhc.uiuc.edu/ACS/tools/ivm/about.html>.
- [15] N. J. Wang, J. Quek, T. M. Rafacz, and S. J. Patel. Characterizing Effects of Transient Faults on a High-Performance Process Pipeline. In *Proc. of the 2004 International Conference on Dependable Systems and Networks*.
- [16] R. Reif, et al. Fabrication Technologies for Three-Dimensional Integrated Circuits. In *ISQED 2002*.
- [17] Brent Goplen and Sachin Sapatnekar. Efficient Thermal Placement of Standard Cells in 3D ICs using a Force Directed Approach. In *ICCAD 2003*.
- [18] Thermal-aware 3D Microarchitectural Floorplanning. A Technical Report, May 20, 2005, <http://www.cs.virginia.edu/techrep/CS-2005-08.pdf>
- [19] Y. Han, I Koren and C. A. Moritz. Temperature Aware Floorplanning. *Second Workshop on Temperature-Aware Computer Systems*, 2005.
- [20] Microarchitectural Floorplanning for Thermal Management. A Technical Report, GIT-CERCS-04-37, <http://www.cercs.gatech.edu/tech-reports/tr2004/git-cercs-04-37.pdf>