

Profile-Guided Microarchitectural Floorplanning for Deep Submicron Processor Design

Mongkol Ekpanyapong Jacob R. Minz Thaisiri Watewai[†] Hsien-Hsin S. Lee Sung Kyu Lim

School of Electrical
and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332

{pop, jrminz, lee, limsk} @ece.gatech.edu

[†]Dept. of Industrial Engineering
and Operations Research
University of California
Berkeley, CA 94720
thaisiri@uclink.berkeley.edu

Abstract

As process technology migrates to deep submicron with feature size less than 100nm, global wire delay is becoming a major hindrance in keeping the latency of intra-chip communication within a single cycle, thus decaying the performance scalability substantially. An effective floorplanning algorithm can no longer ignore the information of dynamic communication patterns of applications. In this paper, using the profile information acquired at the architecture/microarchitecture level, we propose a “profile-guided microarchitectural floorplanner” that considers both the impact of wire delay and the architectural behavior, namely the inter-module communication, to reduce the latency of frequent routes inside a processor and to maintain performance scalability. Based on our simulation results, the profile-guided method shows a 5% to 40% average IPC improvement when clock frequency is fixed. From the perspective of instruction throughput (in BIPS), our floorplanner is much more scalable than a conventional wire length based floorplanner.

Categories and Subject Descriptors:

B.7.2 [Integrated Circuits]: Design Aids

General Terms: Algorithms, Design, Performance

Keywords:

Microarchitectural Planning, Floorplanning, Computer Architecture.

1. INTRODUCTION

According to the projection of the International Technology Roadmap for Semiconductors (ITRS) [12], deep submicron process technology will soon be able to integrate more than one billion transistors onto a single monolithic die. Given the continuing and fast miniaturization of transistor feature size, global wire length is becoming a major hindrance in keeping performance scalable since its relative delay to the gate delay gradually worsens as

technology continues to shrink. Local wire length, on the other hand, will scale with a marginal impact in adding extra delay with respect to the same process technology [7]. Despite the use of different materials, device structures, circuit techniques or novel architectures including nanotechnology, this global interconnect limit still persists due to the nature of device physics and inflicts a substantial performance impact for chips manufactured with deep submicron technology, particularly, for microprocessors which keep pushing the envelope of high performance as the primary design objective.

For the last decade, due to the dramatic advancement of microelectronics and manufacturing technology, computer architects were able to improve processor performance simply by adding more computing capability and increasing resource capacity, for example, increasing cache sizes and hierarchy, enlarging reorder buffer, widening issue width, improving speculation by employing very complex branch predictors, to name a few. All of these architecture enhancement effectively resulted in higher processor performance in the past. On the other hand, CAD tool developers and circuit designers try to reduce the cycle time as much as possible, paying little attention of the entire design at the architectural level. With the increasing impact of global wire length, however, such design methods could lead to less optimal designs, if not totally ineffective, due to the intra-chip communication latency, and need to be largely changed by taking the wires into account. In this paper, we advocate the coalition of architecture design and physical design. By considering both simultaneously, we expect to achieve a much better overall performance improvement for microprocessors designed using deep submicron technology. Here we propose Microarchitectural Planning for circuit designers to consider location information during the design phase. After the design is completed, the detailed floorplanning should conform the location information given from our Microarchitectural Planning.

The rest of this paper is organized as follows. The next section discusses some related work. Section 3 overviews the implication of IPC and clock speed and outlines our approach. Section 4 introduces our profile-guided floorplanning and its mathematical foundation. The description of our architectural framework follows in the next section. Then we show our experimental results in Section 6. Finally, we conclude this work in Section 7.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DAC'04, June 7–11, 2004, San Diego, California, USA.
Copyright 2004 ACM 1-58113-828-8/04/0006 ...\$5.00.

2. RELATED WORK

With the growing concern in global wire delay, there exists a number of research efforts focusing on different aspects including circuits, microarchitectures, and the combination between logic synthesis and physical design. Agarwal et al. in [1] raised the issue of the wire length impact in designing conventional microarchitecture. They showed that reducing the feature size and increasing the clock rate do not necessarily imply an overall performance improvement for deep submicron processor designs. Cong et al. [4] confirmed their observation and showed that without considering clock speed, Instruction Per Cycle (IPC), widely used in architecture research, can be misleading in evaluating the performance of next generation processors.

3. WIRE DELAY ISSUES

Ho, Mai and Horowitz in [7] classified wires into three categories based on their delay impact: 1) wires that scale in length, such as local interconnected wire within logic blocks; 2) wires that do not scale in length and is super-linear when feature size is reduced; 3) repeated wire, i.e. long wire with repeated buffers along the wire. In next generation deep submicron processor design, it is likely that repeaters will be inserted frequently for global wires to prevent the wire delays from becoming non-linear. In this paper, we assume repeated wires to be the dominant wire and examine their performance impact from the perspective of floorplanning. Based on predicted value of resistor, capacitance and other parameters from [12, 7], repeated wire delay is approximated to be 80pS/mm for 30nm technology, which we use as the baseline for our discussion in this paper. Note that a FO4 gate delay for 30 nm is approximately 17pS.

Flip-flop insertion is a technique to alleviate the impact of wire delay for achieving target clock frequency. As discussed in [9], by inserting flip-flops between modules, clock frequency can be increased through deeper pipelining and leads to a higher billion instructions per second (BIPS). Nevertheless, the improvement cannot always be anticipated especially for designs with a small feature size for flip-flop insertion will cause IPC degradation from increasing the latency as shown in Figure 1, in which if the module 2 is to be moved toward the righthand side from the lefthand side, more flip-flops need to be inserted, leading to a longer latency. Therefore, inserting flipflops without a meticulous measure, even though we can accelerate the clock rate, it does not guarantee an overall performance improvement since the latency of communication between modules might now take longer.

4. PROFILE-GUIDED FLOORPLANNING

As shown in Figure 2, we propose a unified framework that combines technology scaling parameters and the execution profiling information of applications to guide the floorplanning of a given processor architecture design. First of all, a machine description is provided as the input for the microarchitecture simulator in which profiling counters were instrumented for bookkeeping module-to-module communication. Then a cycle-accurate simulation is performed to collect and extract the amount of interconnection traffic between modules for given benchmark programs. For cache-like or buffer-like structures, the area size and




Figure 1: Impact of Wire Delay on Latency

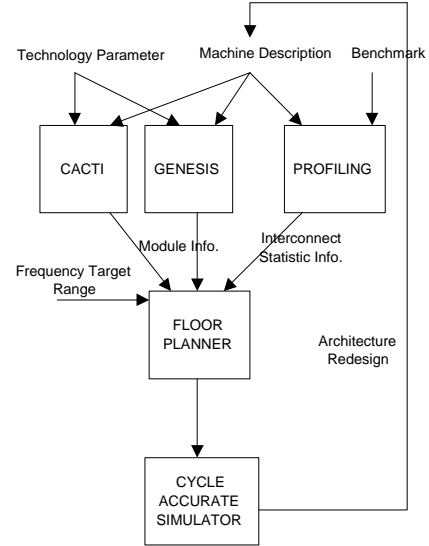


Figure 2: Profile-Guided Floorplanning

module delay are estimated using a industry tool from HP Western Research Labs called CACTI [11]. For scaling the other structures such as ALUs, we use GENESYS [5] developed at Georgia Tech which has a baseline using a .35 μm Verilog model. After the timing and area information of each module is collected, the statistical interconnection traffic and the processor target frequency range are fed into our profile-guided floorplanner to generate a floorplan for deriving the timing model (i.e. inter-module latency) of the microarchitecture as a result of the generated floorplan. With the new latencies, the architecture performance simulation is performed for simulating the IPC numbers, and later the actual performance in BIPS can be calculated together with the cycle time. We keep iterating the entire methodology for exploring alternate architecture designs by increasing or decreasing the module size, or even splitting the module geometry for achieving the performance goal. Results from our profile-guided microarchitectural floorplanning can be used to guide the final global floorplanning such that performance is not degraded.

Now we formulate the mathematical programming models for our floorplanner. First, we present a mixed integer nonlinear programming (MINP) model which minimizes the weighted wirelength. Since to find the optimal solution of the MINP model is NP-hard, we then linearize and relax integral constraints in such a way as to stay close to the optimal solution within a reasonable runtime. Finally, we demonstrate the complete floorplanning algorithm.

4.1 Mixed Integer Non-linear Programming Model

The parameters used in the model MINP are defined as follows. Let N denote the set of all flexible modules, and E denote the set of directed edges where a directed edge (i, j) represents a wire from module i to module j . Also, let α be the repeated delay per mm , and λ_{ij} be the statistical traffic on wire (i, j) . g_i is the gate delay of module i . $w_{min,i}$ and $w_{max,i}$ denote the minimum half width and the maximum half width of module i respectively. The area of module i is denoted by a_i . Finally, f_{ij} is the minimum number of flip-flops on wire (i, j) that is the number of flipflops from the circuit design before considering the impact of wire delay.

Let L ($1/\text{clock_frequency}$) denote the target cycle time of the system and is the input to our floorplanner. In the model, we need to determine the values of the following decision variables: Let (x_i, y_i) denote the location of the center of module i on \mathbb{R}_+^2 space. X_{ij} and Y_{ij} represent $|x_i - x_j|$ and $|y_i - y_j|$ respectively. z_{ij} is the number of flip-flops on wire (i, j) . w_i denotes the half width of module i . To avoid overlapping between any two modules i and j , where $i < j$, we also need a set of binary variables so that at least one of the followings holds.

$$\begin{aligned} x_i + w_i &\leq x_j - w_j && i \text{ is on the left of } j \\ x_i - w_i &\geq x_j + w_j && i \text{ is on the right of } j \\ y_i + \frac{a_i}{4w_i} &\leq y_j - \frac{a_i}{4w_i} && i \text{ is below } j \\ y_i - \frac{a_i}{4w_i} &\geq y_j + \frac{a_i}{4w_i} && i \text{ is above } j \end{aligned}$$

We thus let c_{ij} and d_{ij} be the binary variables such that

$$(c_{ij}, d_{ij}) = \begin{cases} (0, 0) & \text{if } i \text{ is on the left of } j \\ (0, 1) & \text{if } i \text{ is on the right of } j \\ (1, 0) & \text{if } i \text{ is below } j \\ (1, 1) & \text{if } i \text{ is above } j \end{cases}$$

The constraints for non-overlapping among modules are given from inequality (8) to (11) in the MINP formulation illustrated in Figure 3 wherein the objective function (1) is to minimize the net weighted delay. Constraint (2) is obtained by the definition of latency. (3) to (6) define the absolute values. Constraint (7) represents the minimum number of flip-flops required. (12) states the possible range of each module's half width. (13) are non-negative constraints of module's location. (14) manifests that (c_{ij}, d_{ij}) are binary. Finally, (15) specifies that the number of flip-flops must be integer. Also note that M is a sufficiently large number.

4.2 Mixed Integer Linear Programming Model

To avoid non-linearity of (10) and (11), we linearize the half height of module i (h_i) using linear approximation as shown in Figure 4. Note that this approximation will guarantee that the solution in the approximated model is feasible (due to over approximation). The resulting mixed integer linear programming formulation, MILP, is shown in Figure 5.

4.3 Linear Programming Model

Despite the fact that we can convert MINP into MILP, the problem still remains NP-hard for two dimensional bin-packing. Hence, it will limit our search space by MILP run-

MINP

Minimize $\sum_{(i,j) \in E} \lambda_{ij} z_{ij}$ (1)

Subject to

$$z_{ij} \geq \frac{g_i + \alpha(X_{ij} + Y_{ij})}{L} \quad (i, j) \in E \quad (2)$$

$$X_{ij} \geq x_i - x_j \quad (i, j) \in E \quad (3)$$

$$X_{ij} \geq x_j - x_i \quad (i, j) \in E \quad (4)$$

$$Y_{ij} \geq y_i - y_j \quad (i, j) \in E \quad (5)$$

$$Y_{ij} \geq y_j - y_i \quad (i, j) \in E \quad (6)$$

$$z_{ij} \geq f_{ij} \quad (i, j) \in E \quad (7)$$

$$x_i + w_i \leq x_j - w_j + M(c_{ij} + d_{ij}) \quad i < j \in N \quad (8)$$

$$x_i - w_i \geq x_j + w_j - M(1 + c_{ij} - d_{ij}) \quad i < j \in N \quad (9)$$

$$y_i + \frac{a_i}{4w_i} \leq y_j - \frac{a_i}{4w_i} + M(1 - c_{ij} + d_{ij}) \quad i < j \in N \quad (10)$$

$$y_i - \frac{a_i}{4w_i} \geq y_j + \frac{a_i}{4w_i} - M(2 - c_{ij} - d_{ij}) \quad i < j \in N \quad (11)$$

$$w_{min,i} \leq w_i \leq w_{max,i} \quad i \in N \quad (12)$$

$$x_i, y_i \geq 0 \quad i \in N \quad (13)$$

$$c_{ij}, d_{ij} \in \{0, 1\} \quad i < j \in N \quad (14)$$

$$z_{ij} \in \mathbb{N} \quad (i, j) \in E \quad (15)$$

Figure 3: MINP Formulation




Figure 4: Block Area Constraint Approximation
 ning time for a large number of modules. To remedy this problem, we further relax MILP into Linear Programming (LP) as shown in Figure 6. We derive a method similar to those described in [8, 3] by recursively bi-partitioning a space into smaller subregions.

To relax the integrality while maintaining the feasibility and staying close to the optimal solution, we first relax the integrality of z_{ij} to be a real number. We also solve several linear programming problems to determine the non-overlapping relationship $-(c_{ij}, d_{ij})$.

We start the algorithm by creating a large block containing all modules. In each iteration we divide each current block into two smaller subblocks by specifying the centers of each subblock as well as their corresponding boundaries. Then we randomly assign modules currently in the block into one of the centers of these two subblocks. For the next iteration, the modules currently in the block must remain in the same block ((20)-(23)), and all the modules assigned to the same subblock have center of gravity of area at the center of that subblock ((24)-(25)). Note that even if a module assigned to a particular subblock, it can be moved across the subblock boundary as long as the preceding two conditions are satisfied. Once the locations of all modules are determined at the current iteration (by LP(u) as shown in Figure 6), each subblock will become a block in the next iteration. We terminate the algorithm when each subblock contains exactly one module. The following notations are defined before we detail the whole algorithm. Let $B(u)$ denote the set of all blocks at iteration u , and $M_j(u)$ denote

<p>MILP</p> $\text{Minimize } \sum_{(i,j) \in E} \lambda_{ij} z_{ij} \quad (16)$ <p>Subject to (2)—(9), (12)—(15) and the following</p> $y_i + m_i w_i + k_i \leq y_j - m_j w_j - k_j + M(1 - c_{ij} + d_{ij}) \quad i < j \in N \quad (17)$ $y_i - m_i w_i - k_i \geq y_j + m_j w_j + k_j - M(2 - c_{ij} - d_{ij}) \quad i < j \in N \quad (18)$
--

Figure 5: MILP Formulation

<p>LP(u): Minimizing net weighted delay</p> $\text{Minimize } \sum_{(i,j) \in E} \lambda_{ij} z_{ij} \quad (19)$ <p>Subject to (2)—(7), (12)—(13) and the following</p> $x_i + w_i \leq r_j \quad i \in M_j(u), j \in B(u) \quad (20)$ $x_i - w_i \geq l_j \quad i \in M_j(u), j \in B(u) \quad (21)$ $y_i + m_i w_i + k_i \leq t_j \quad i \in M_j(u), j \in B(u) \quad (22)$ $y_i - m_i w_i - k_i \geq b_j \quad i \in M_j(u), j \in B(u) \quad (23)$ $\sum_{i \in S_{jk}(u)} a_i x_i = \sum_{i \in S_{jk}(u)} a_i \times \bar{x}_{jk} \quad k \in \{1, 2\}, j \in B(u) \quad (24)$ $\sum_{i \in S_{jk}(u)} a_i y_i = \sum_{i \in S_{jk}(u)} a_i \times \bar{y}_{jk} \quad k \in \{1, 2\}, j \in B(u) \quad (25)$

Figure 6: LP Formulation

the set of modules currently in block j at iteration u . Also, let $S_{jk}(u)$ be the set of modules assigned to the center of subblock k ($k \in \{1, 2\}$) contained in block j at iteration u . We denote the center of subblock k contained in block j by $(\bar{x}_{jk}, \bar{y}_{jk})$. Finally, let r_j, l_j, t_j, b_j denote the right, left, top, and bottom boundary of block j . The algorithm is summarized in Figure 7.

Note that the LP(u) will be feasible if we start from an initial block that is large enough, and at each iteration each subblock is created according to the total area of all the modules assigned to that subblock.

Once the algorithm is terminated with locations of all modules, we investigate the non-overlapping relationship (8) to (11). If for each pair (i, j) exactly one of four possible pairs (c_{ij}, d_{ij}) is satisfied, we fix such (c_{ij}, d_{ij}) , otherwise we randomly fix one of the two possible relationships - horizontal (left or right) and vertical (above or below). By fixing all (c_{ij}, d_{ij}) , we solve MILP without (15) so that it is a linear programming problem. Note that each time we perform a random procedure, we repeat many times with different random values and select the best result.

We also have another LP model in which we minimize the total wirelength. We apply the similar approach as our model. The LP'(u) for wirelength minimization is described in Figure 8.

5. SIMULATION INFRASTRUCTURE

SimpleScalar 3.0 tool suite [2] is used as our architecture simulator for both profile collection and performance simulation. The detailed microarchitecture evaluated in Section 6 is illustrated in Figure 9 and their varied configurations experimented are listed in Table 1. Each functional block in Figure 9 represents a module used by our floorplanner. In order to facilitate physical-design driven micro-architectural exploration we added some new fea-

<p style="text-align: center;">Bi-Partitioning LP Algorithm</p> <p>For a given L</p> <p>Step 0: Initialize $B(1) = \{1\}, M_1(1) = N$.</p> <p>Step u: for <i>count</i> = 1 to <i>run</i></p> <p style="padding-left: 20px;">Specify $S_{jk}(u), (\bar{x}_{jk}, \bar{y}_{jk})$. Solve LP($u$).</p> <p style="padding-left: 20px;">Update $B(u+1), M_j(u+1), r_j, l_j, t_j$, and b_j.</p> <p style="padding-left: 20px;">Goto final step if each $M_j(u+1)$ contains exactly 1 module</p> <p style="padding-left: 20px;">end for</p> <p style="padding-left: 20px;">Project best obj. into next iteration $u+1$</p> <p>Final step: Solve MILP by fixing c_{ij}, d_{ij} obtained from Recursive Bipartitioning.</p> <p style="padding-left: 20px;">Project best z_{ij}</p>

Figure 7: Recursive Bi-partitioning LP Algorithm

<p>LP'(u): Minimizing total wirelength</p> $\text{Minimize } \sum_{(i,j) \in E} (X_{ij} + Y_{ij}) \quad (26)$ <p>subject to (3)—(6), (12)—(13), (20)—(23), (24)—(25)</p>
--

Figure 8: Minimize Total Wirelength

tures. First, we extended the SimpleScalar to include a new feature of configurable pipeline depth. The fetch unit of the simulator was efficiently modified to accommodate desired pipeline depths providing a handle to study the effects of lengthening the processor pipeline. In terms of performance, pipeline depth has a direct impact on the operational frequency of the processor. The number of pipeline stages is also affected by our extension of the functional units. In the modified simulator, the functional units or execution resources are completely made configurable in terms of operation and issue latency and can be specified as configuration input.

Provisions were also made to consider wire delays in the simulator. The existing simulator assumes that the forwarding latency between blocks is always one cycle which is less reasonable while operating at extremely high frequency given the increased wire delays and larger die areas. In our modified simulator, the forwarding latencies of the function units to other units or pipeline stages were made configurable, enabling a more realistic IPC projection.

For accurate prediction and optimization of performance, the wires can no longer be isolated from architecture-level evaluation but must be modeled as *units* that consume power and have delays. For example, the Pentium 4 processor design has dedicated 2 pipeline stages for moving signal across the chip due to wire delay [6]. Wires, as much as other logical, storage, and control units of a processor must be "designed" appropriately. In order to have a simultaneous view of the frequency and the IPC, we have to explore the floorplan and processor organization together. The wire parameters can be captured for architectural optimization through floorplanning, to a reasonable degree of accuracy. The wires directly affect the forwarding latency and the pipeline depth of the processor. For performance evaluation, we use the information provided by the floorplanner to derive essential simulation parameters such as pipeline depth and forwarding latency. The forwarding latency is a function of the distance between units in the floorplan and the number of flip-flops between modules. If the floorplan has been optimized for clock speed, the pipeline depth of the processor reflects it. In our experiments, we expect an improvement in performance (in archi-

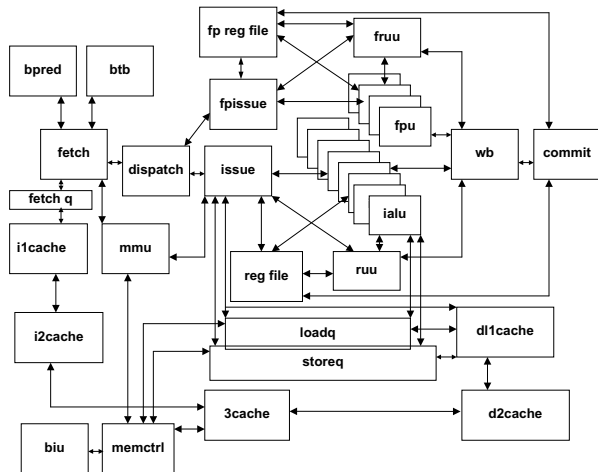


Figure 9: Processor Microarchitecture Model

Micro-arch Structure	Size				# of Bits
	Cfg 1	Cfg 2	Cfg 3	Cfg 4	
M. width	2	4	8	8	N/A
Bpred	128	512	512	512	2
BTB	128	512	512	512	96
RUU	64	128	512	512	168
L1 Icache	8KB	64KB	8KB	8KB	512
L1 Dcache	8KB	64KB	8KB	8KB	512
L2 Ucache	64KB	512KB	128KB	128KB	1024
L3 Ucache	None	None	2MB	2MB	1024
ITLB	32	128	128	128	112
DTLB	32	128	128	128	112
ALU	2	4	4	8	N/A
FPU	1	2	2	4	N/A
LSQ	16	64	128	128	84
Mem port	1	4	8	8	N/A

Table 1: Microarchitecture Configurations

tektural simulation) if the frequency of forwarding traffic between units are included in our floorplan formulation. The forwarding frequency driven floorplanning tries to put heavy traffic units closer together, minimizing their latencies as a function of their distance. Our floorplanning algorithms described in Section 4 were implemented in C, compiled with gcc with -O3, integrated with Simplescalar tool set, and executed on Pentium III 750MHz machines. We use lp_solve [10] to solve our linear programs.

6. EXPERIMENTAL RESULTS

We performed experiments on 10 SPEC2000 benchmarks, 6 from the integer suite and 4 from the floating point one. The training input set was used for profile collection while the IPC performance results were gathered using the reference input set. Each simulation was fast-forwarded by 200 million instructions and then simulated for 100 million instructions. The maximum execution time spent in both floorplanning and simulation was less than two hours for each benchmark.

Figure 10 shows the IPC improvement based on the same clock frequency with different microarchitecture configurations for the baseline wirelength algorithm (WL) and our profile-guided floorplanning algorithm (PGF). Note that the baseline floorplan is generated by solely minimizing the total wirelength. In minimizing wirelength objective, some control path modules are tied together to make it practical. For example instruction fetch and branch predictor always stay close together, otherwise it is possible that branch predictor can take more cycles than the ac-

tual branch address calculation. The results show that the baseline algorithm ceases to deliver higher IPCs as the number of the modules and/or the size of the modules are increased due to the limitation of wirelength. On the other hand, our profile-guided floorplanner continues to increase the IPCs since the dynamic information of interconnection traffic between modules is taken into account during floorplanning, thus shortening the cycle latency among heavily communicated functional units. For the most complex microarchitecture (Cfg 4), the processor design based on our technique outperforms the baseline by 40%.

Figure 11 evaluates the impact with respect to the total wirelength using our technique. Instead of minimizing wire length, our algorithm attempts to improve overall performance and could potentially result in a longer total wirelength. As the results shown in the figure, our technique does not actually incur any significant increase in wires for less complexity processors, in some cases it is even better than the baseline¹. For the outlier case shown in Config3 of bzip2, data are mostly forwarded from RUU to ALUs with very little inter-ALU communication, hence the only constraint imposed is from the RUU. Whereas in Config4 case, there are more data forwarding between ALUs, hence the imposed constraint do not allow each ALU to move freely, inhibiting potential wire length increase. For complex machines such as the configuration 4, we increase the total wire by 68% in average. Even though the total wirelength is increased, it only affects the inter-module latencies of those less frequently used routes.

We present the overall performance in normalized BIPS for different clock frequencies in Figure 12, in which the clock period is decreased from 0.2ns (5GHz), 0.18ns (5.5GHz) down to 50ps per cycle (20GHz). The results, averaged across all benchmarks, are normalized to the BIPS delivered by 5GHz baseline machine using the floorplanner that minimizes the wirelength. As shown, the advantage of using a profile-guided method significantly stands out when the clock speed is increased. When clock frequency is quadFor a 20GHz processor, our technique shows almost 3 times improvement in BIPS over the baseline, whereas minimizing wirelength can gain only twice. Hence our floorplanner is more scalable than a conventional approach. Event hough we didn't show the comparison between timing driven floorplanning here, it can be approximately considered as minimizing wirelength on higher clock frequency. As shown in Figure 12, our approach can be as good as next generation processor and it is even better for high clock frequency.

7. CONCLUSIONS

Due to the miniaturization of feature size, global wire delay needs to be addressed and considered as an architectural entity for microprocessor designs. In this work, we proposed a profile-guided floorplanning algorithm which uses both the technology scaling parameters and the information of dynamic interconnection traffic between microarchitectural modules to guide the floorplanning for performance optimization. The results show that our proposed method can effectively improve the overall applications' performance by up to 40% over a conventional wirelength based floorplanner.

¹Note that the baseline itself has some randomness involved and is also a heuristics, not an optimal solution.




Figure 10: Performance Improvement for a 10GHz Processor (WL: Wirelength, PGF: Profile-Guided Floorplan)

8. REFERENCES

- [1] V. Agarwal, M. S. Hrishikesh, S. W. Keckler, and D. Burger. Clock Rate versus IPC: The End of the Road for Conventional Microarchitectures. In *ISCA*, 2000.
- [2] T. M. Austin. SimpleScalar tool suite. <http://www.simplecalar.com>.
- [3] N. Sehgal B. Halpin, C.Y.R. Chen. Timing Driven Placement using Physical Net Constraints. In *DAC01*, 2001.
- [4] J. Cong, A. Jagannathan, G. Reinman, and M. Romesis. Microarchitecture Evaluation with Physical Planning. In *DAC*, 2003.
- [5] J. C. Eble, V. K. De, D. S. Wills, and J. D. Meindl. A Generic System Simulator (GENESYS) for ASIC Technology and Architecture Beyond 2001. In *Int'l ASIC Conference*, 1996.
- [6] P. N. Glaskowsky. Pentium 4 (Partially) Previewed. *Microprocessor Report*, 14(8):11-13, August 2000.
- [7] R. Ho, K. W. Mai, and M. A. Horowitz. The Future of Wires. *Proceedings of the IEEE*, 2001.
- [8] J. Kleinhaus and et al. GORDIAN: VLSI Placement by Quadratic Programming and Slicing Optimization. *IEEE Tran. on CAD*, 1991.
- [9] W. Liao and L. He. Full-Chip Interconnect Power Estimation and Simulation Considering Concurrent Repeater and Flip-flop Insertion. In *ICCAD*, 2003.
- [10] Eindhoven University of Technology. LP_solve. ftp://ftp.es.ele.tue.nl/pub/lp_solve/.
- [11] P. Shivakumar and N. P. Jouppi. CACTI 3.0: An Integrated Cache Timing, Power, and Area Model. Technical Report 2001.2, HP Western Research Labs, 2001.
- [12] SIA. National Technonology Roadmap for Semiconductors, 2001.




Figure 11: Impact on Total Wire Length




Figure 12: Performance versus Frequency Scaling (Cfg3)