

Microarchitecture Level Interconnect Modeling Considering Layout Optimization

Weiping Liao and Lei He

Dept. of Electrical Engineering, Univ. of California at Los Angeles, CA 90095, {wliao, [lhe](mailto:lhe@ee.ucla.edu)}@ee.ucla.edu

* corresponding author: Weiping Liao

Address:

University of California at Los Angeles
Department of Electrical Engineering,
53-135W, Engineering IV building
420 Westwood Plaza
Los Angeles, CA 90095

Office : (310) 267-5407

Fax : (310) 206-4685

Email : wliao@ee.ucla.edu

Date of Receiving: **to be completed by the Editor**

Date of Acceptance: **to be completed by the Editor**

Microarchitecture Level Interconnect Modeling Considering Layout Optimization¹

Weiping Liao and Lei He

Abstract — In this paper, we study microarchitecture-level interconnect modeling for power and performance. Considering structural interconnects, layer assignment, and concurrent repeater and Flip-Flop (FF) insertion, we develop cycle-accurate microarchitecture-level power and throughput simulation and obtain an accurate modeling of interconnects at the early design stage. Experiment show that the simulation reduces over-estimation by up to 2.24X compared to the conventional power estimation based on purely stochastic interconnects and fixed switching factor. Furthermore, we optimize throughput with consideration of FF insertion for interconnects and floorplanning optimization. We show that throughput is not always higher for an increased clock frequency, and there exists an optimal clock frequency to maximize throughput for a given microarchitecture and given floorplan. In addition, floorplan optimized for IPC (instructions per cycle)-critical interconnects has little on the total interconnect length but improves throughput by 23.49%. As FF insertion becomes necessary to achieve the clock frequency specified by ITRS, we conclude that the traditional design flow optimizing IPC and clock frequency separately is no longer valid, and coupled microarchitecture and layout optimization may improve both power efficiency and throughput.

¹ This paper is partially supported by NSF CAREER award CCR-0401682, SRC grant 1116, a UC MICRO grant sponsored by Mindspeed, and a Faculty Partner Award by IBM. We used computers donated by Intel and SUN Microsystems. Address comments to lhe@ee.ucla.edu.

Keywords — Microarchitecture, interconnect, repeater, flip-flop, floorplan

1 INTRODUCTION

The primary goal of processor design is to improve throughput within the power constraint. This goal is conventionally achieved by two separated design stages²: architects optimize IPC (Instruction Per Cycle) with microarchitecture innovations, and then VLSI circuit designers perform logic synthesis and layout design to retain IPC and maximize clock frequency. In most cases, interconnects are optimized at the second stage, but is not considered at the microarchitecture-level. As VLSI technology advances, the system delay has become dominated by the interconnect delay. A growing number of repeaters and Flip-Flops (FFs) are used to reduce the interconnect delay [1]. Because interconnects with inserted repeaters and FFs may greatly affect IPC and power, a microarchitecture is hardly optimized without considering interconnect and layout optimization.

However, most existing microarchitecture level simulation tools such as [2] - [5] do not explicitly characterize the impact of interconnects. At the layout and physical design level, there have been extensive studies on interconnect performance and power modeling considering repeater and FF insertion. Focusing on performance modeling in terms of interconnect delay and critical path estimation, [6], [7] studied the repeater insertion for optimal delay. Such studies are extended to consider the impact of process variation in the ultra deep submicron design era [8]. All these studies [6] - [8] only considered repeater insertion, assuming the clock period is longer than the delay of critical path. As technology keeps scaling, wire delay becomes dominated and easily exceeds the clock cycle time [9], making the insertion of FFs necessary. Targeting routing tree topology, [10] and [11] proposed concurrent FF and repeater insertion methodologies. However, no microarchitecture-level characteristics such as the structure interconnect in Section IV was considered in either [10] or [11].

² Note that in industry, there may exist *ad hoc* designs considering coupled optimization between IPC and clock frequency. However, those *ad hoc* designs do not present any general design methodology and are excluded from our discussion.

Concerning the power consumption by a large number of repeaters, [12] estimated the power for interconnect repeater insertion based on the stochastic wire length distribution [13], and studies delay-power trade-off for minimizing repeater power. [14] studied the trend of repeater power consumption for unit wire lengths for five technology generations from 180nm to 50nm. In both [12] and [14], an over-simplified repeater model (i.e., single-model to be defined in Section 2) is used and no FF insertion is considered. In addition, none of them considered structure interconnects, layer assignment or cycle-accurate interconnect simulation. Furthermore, targeting buffer trees, power-efficient repeater insertion considering dual- V_{dd} and dual- V_t technologies are studied in [15], [16]. Such methods are orthogonal to our study. With the accurate power estimation proposed in this paper, methods in [15], [16] can be conveniently extended to full-chip repeater power reduction.

At the microarchitecture level, [17] presents coupled system design and VLSI design for throughput optimization. However, [17] considers only buffer insertion but not FF insertion for interconnects. The initial study of this paper [18] studied the power and performance impact of concurrent repeater and FF insertion at microarchitecture level. Preliminary results showed in [18] that FF insertion has lower IPC but can improve the system throughput. [19] - [21] further developed efficient algorithms to consider the performance impact of FF insertion during floorplanning optimization. However, only IPC, but not the system throughput, was optimized in [19] - [21].

Considering interconnect layout optimization including floorplanning, layer assignment, and concurrent repeater and FF insertion, we develop in this paper a cycle-accurate microarchitecture-level power and throughput simulation and obtain an accurate modeling of interconnects at the early design stage. We also apply this simulation to optimize microprocessor throughput considering interconnect pipelining and floorplanning adjustment.

The rest of this paper is organized as follows. In Section 2, we study repeater and FF insertion for

individual wires. In Section 3, we study microarchitecture level interconnect power estimation and cycle-accurate power simulation with consideration of concurrent repeater and FF insertion. In Section 4, we optimize throughput considering interconnect pipelining and floorplanning optimization. We conclude in Section 5. An extended abstract about the preliminary results of this study was published in [18].

2 REPEATER AND FLIP-FLOP INSERTION

2.1 Interconnect and Device Models

In this paper, we model interconnects by the Π -type distributed RC circuit, and consider multiple interconnect layers. Top layers are used for wide and long global interconnects, and bottom layers are used for short local interconnects. Between them are the layers for intermediate interconnects. For the simplicity of presentation, we assume all wires are global wires in this section, and define the distinction of global and non-global wires in Sections 3. We assume that a unit length interconnect has resistance R_w and capacitance C_w , and model an inverter by its gate capacitance, drain capacitance and its effective resistance. We represent the gate, drain capacitances and effective output resistance for a minimum size inverter as C_0 , C_p and R_0 , respectively. A repeater can be a single inverter, or a cascaded inverter chain.

We use the Elmore delay to calculate interconnect delay, i.e.

$$T_d = \sum_i R_i \cdot C_{down} \quad (1)$$

where T_d is the total delay, R_i is the resistance of a wire segment and C_{down} is the sum of the downstream capacitances of R_i . We consider interconnect power including dynamic power and leakage power given by Equation (2) and (3), respectively:

$$P_{dynamic} = \frac{1}{2} \alpha V_{DD}^2 f_{clk} (S_g(C_0 + C_p) + l_g C_w + N_{FF} C_F) \quad (2)$$

$$P_{leakage} = V_{DD} I_{OFF} (S + N_F S_F) \quad (3)$$

where f_{clk} is the clock frequency, l is the wire length, α is the switching factor, I_{off} is the unit leakage current, and S is the total inverter size. Furthermore, N_F is the total number of FFs, C_F is the total capacitance of one FF, and S_F is the total gate size of one FF. We assume 100nm technology in this paper, with parameters in Table I, where the wire widths and heights are obtained from ITRS roadmap³, C_w and R_w are calculated by Berkeley Predictive Technology Model [22], the I_{off} is from [14], the α is 0.15 [23] and is fixed for logic and interconnects except the structure interconnects with cycle-accurate power simulation in Section 3.3. The other values are obtained from SPICE simulations.

In this paper, we assume all interconnects are two-pin nets. This assumption has been used widely in the literature for high-level estimation [12], [13]. Specifically, as shown in Figure 1, we assume every interconnect has one driver and one load. Both the driver and load are inverters with the 4X minimum inverter size. We study the repeater and FF insertion for two objective functions: one is to meet the delay target with minimum number of FFs, or min-FF; and the other is to meet the delay target with minimum total interconnect power consumption, or min-power.

2.2 Min-FF Solution

It has been assumed in [12], [14] that for repeater insertion, the input capacitance C_{in} and effective resistance for each repeater are equal to $S \square C_0$ and $\frac{R_0}{S}$ respectively, where S is the size of the repeater.

Under this assumption, each repeater is a single inverter, named single model. To drive a large load, a repeater may contain a chain of cascaded inverters, where C_{in} of a repeater is equal to C_0 times the size of first inverter in the inverter chain. The formulas to determined S and the location of each inverter along the interconnect are presented in [12] and [14]. We call this type of repeater cascaded

repeater. An inverter in a cascaded repeater is a stage, and the size ratio between two consecutive inverters is the stage ratio. In addition, we also consider a hybrid model where the first stage is a chain of cascaded inverters, but the rest are single inverters. In the hybrid model, the cascaded repeater is put at the beginning of the interconnect, and the location of other single inverters can be calculated based on the formulas used in the single model. The hybrid model may lead to a good solution when the inverter in the last stage of the first repeater is large enough to drive the rest of the single repeaters. We illustrate the three repeater insertion models in Figure 2.

We study the power optimization problem under a given delay target for interconnects. The existing analytical repeater insertion methods [12], [14] can only be used for the single model. We find the solution by the following enumeration. For a cascaded model, we enumerate the number of repeaters, the first inverter size, the uniform stage ratio and the stage number for each repeater. Again, we assume that all repeaters are identical. For the hybrid model, we enumerate the number of repeaters, the design of the first cascaded repeater, and the uniform design of the rest of the repeaters using the single model. For each combination, we calculate the delay and power. If the delay is smaller than our delay target, we call this combination a valid solution. We choose the valid solution with the smallest number of FFs. If there is more than one valid solution, we choose the one with the lowest power consumption. We also do pruning during enumeration. If we have obtained a valid solution with repeater size S , all solutions with repeater size greater than S should be skipped because they definitely consume more power. If a wire is too long to meet the delay target, we insert FFs to break the wires into shorter wires. If we have obtained valid solutions for one wire, we can reuse the solutions for wires of the same length.

Table II shows our experiment results from all three models as discussed above. We use the wire lengths 4mm, 8mm, and 1cm, and clock frequencies 1GHz, 2GHz and 3GHz. We assume that the

³ Note the width and height of global wires are from 130nm technology as we assume the global interconnects do not scale [1]

delay target is 80% clock period. No FF insertion is needed for wires up to 10mm and 4mm for 1GHz and 2GHz clock frequencies (see **highlights** in the table), respectively. Among these cases, the hybrid model achieves up to 15.09% power reduction compared with the single model. The hybrid model also has the smallest number of FFs for the same wire and delay target. This is further illustrated in Table III. For target delay, the longest wire without FF insertion in the hybrid model can be $1.5X$ of that in the single model.

2.3 *Min-Power Solution*

Although the hybrid model provides better power consumption for the same wire length, FF number and clock frequency, we also observe from the Table II that the single model with more FFs actually has lower power consumption than the hybrid model with fewer FFs. The reason is that for all repeater insertion models, the resulting power consumption is super-linear with respect to the wire length as shown in Figure 3, where the wire length increases by $4X$ from 1mm to 4mm, the power consumption increases by more than $7X$. It is easy to see that instead of inserting FFs merely to meet the delay target, we can reduce power by aggressively inserting more FFs. Figure 4 shows the power for different wire lengths for same target delay but different numbers of FFs. According to Figure 4, when enough FFs are inserted, the power curve becomes nearly linear with respect to wire length. On the other hand, FF insertion is not always beneficial. The more FFs inserted, the more power is consumed by the FFs. There exists a point where the extra power consumed by FFs outweighs the power saving by FF insertion, i.e. there is an optimal number of FF to be inserted for minimal power consumption.

The min-power solution finds the concurrent repeater and FF insertion method with minimum power and less delay than the delay target. Again, we use enumeration to find the min-power solution. We enumerate a range of reasonable FF numbers. For each number, we find the repeater insertion solution as discussed before. Finally, we choose the solution with the minimum total power. We

present the results under min-power FF insertion and hybrid repeater model in Table II. The min-power method can reduce the interconnect power by up to 40.39% compared with the min-FF method. However, the effectiveness of min-power method may not be over-emphasized because it depends on specific interconnect length distribution in individual design. For some specific design, the power reduction of the min-power method over the min-FF method may be small, as in our example in Section 3.2.

2.4 Runtime Reduction

In our implementations, we use a lookup table for concurrent repeater and FF insertion solution since there is no closed-form solution. Tables are built for each interconnect length and clock frequency. Each table entry contains the concurrent repeater and FF insertion solution and the optimal power. With lookup table we can greatly reduce runtime and speed up our interconnect power estimation in Section 3.

3 MICROARCHITECTURE LEVEL INTERCONNECT POWER ESTIMATION AND CYCLE-ACCURATE SIMULATION

In this section we refine interconnect power modeling as follows. We first assume purely stochastic interconnects and fixed switching factor, perform layer assignment and develop type I interconnect power estimation. Then we introduce the concepts of random interconnects and structural interconnects, and develop type II interconnect power estimation. Finally we consider accurate activity rate for interconnects based on cycle-accurate simulation, and develop type III interconnect power estimation, which is also called power simulation. As we have already seen in Section 2, the hybrid model achieves the lowest power and least number of FFs compared with the other two models. In this section and the rest of this paper, we only use the hybrid model for interconnect

power estimation unless specified otherwise.

3.1 Power Estimation with Stochastic Interconnects

Interconnects are routed in different metal layers for routability and performance optimization, and layer assignment has a significant impact on power estimation. In our layer assignment, we assume the top two layers are used for global interconnects. We further assume that on these two layers, 50% of the area is used by power/ground and clock routing. Therefore, the total area occupied by all global interconnects are $2 \times 50\% \times \text{Chip_size}$, and the minimum length of global interconnects l_{gmin} satisfies Equation (4):

$$2 \times 50\% \times \text{Chip_size} = \int_{l_{min}}^{l_{max}} \text{Global_pitch_width} g(l) dl \quad (4)$$

where l_{max} is the maximum length of interconnects and it is $2\sqrt{N}$ with N being the total number of gates on the chip. $i(l)$ is the length density function. l_{gmin} can be used as the length boundary between global and intermediate interconnects.

Similarly, we find the length boundary between intermediate and local interconnects l_{mmin} by Equation (5):

$$\text{Layer_number} \times \text{Chip_size} = \int_{l_{mmin}}^{l_{gmin}} \text{Intermediate_pitch_width} g(l) dl \quad (5)$$

where Layer_number is the number of intermediate layers, and the area utilization rate is 100% for the intermediate layer. We assume the Layer_number is an even number, and keep increasing Layer_number until the interconnects with the length of l_{mmin} can meet the delay target without repeater insertion. Interconnects with length less than l_{mmin} are local interconnects and are assigned to local layers.

We obtain the chip size from ITRS and assume the chip area for random logic by subtracting cache

area from the total chip area. For type I interconnect power estimation, we use the length density function $i(l)$ from the stochastic length distribution methodology [13] to calculate the boundaries between local, intermediate and global interconnects in layer assignment. We set the length of one gate pitch as the square root of the logic gate area obtained from ITRS. The typical Rent's exponent of 0.55 is used. The gate count, gate area, and gate pitch are shown in Table IV. For min-FF and min-power methods, the system clock frequency is 3GHz and we assume the interconnect delay target is about 80% of the clock period. There is no delay target for min-delay method as the minimum interconnect delay depends on the interconnect length.

Figure 5 shows the type I interconnect power calculated by the three different repeater and FF insertion solutions. In the first solution, repeaters are inserted for minimum delay, or min-delay, i.e., we insert repeaters as long as it can reduce delay and we do not insert any FFs. The power reduction from the min-power method mainly comes from the reduced repeater area. We define one equivalent repeater as one minimum size inverter. A repeater with total size S can be mapped to S equivalent repeaters. For any repeater and FF insertion solution, the total power is decided by total wire capacitances, the number of equivalent repeaters and FFs. Table V shows the total number of equivalent repeaters and FFs for all three solutions. Note that in the min-delay method we do not insert any FF, and there is no guarantee that the delay target for min-FF and min-power can be satisfied in the min-delay method. From Table V we can see that the min-FF and min-power solutions reduce the number of equivalent repeaters by $3.40X$ and $8.23X$, respectively. Although the number of FFs in min-power solution is almost $8X$ of that in min-FF solution, the min-power solution still saves 14.03% power as it reduces the number of equivalent repeaters by 58.76%.

3.2 Power Estimation with Structural Interconnects

Stochastic interconnect distribution is assumed in [11], [12], [14] and in our type I interconnect power estimation. However, major components in a system-on-a-chip are often connected by

varieties of busses that can be modeled accurately. To capture this, we introduce the concepts of random interconnects and structural interconnects. The random interconnects are interconnects inside each module and can be calculated by the same stochastic model as in type I interconnect power estimation. The structural interconnects are address and data busses between related modules, and their lengths are decided by the floorplan of the layout.

We consider high-performance SuperScalar processors, and summarize the configuration of processors under study in Table VI. Based on the die photo of the MIPS R10000 microprocessor [24], we first design the floorplanning without a L2 cache, and then incorporate a L2 cache into the floorplanning according to appropriated area ratio between L2 cache and other modules, as shown in Figure 6. We measure the lengths of busses according to the Manhattan distances between the centers of modules connected by the busses. Table VIII shows the bit-width and lengths for all busses.

The number of long interconnects are reduced with the introduction of structural interconnects. Therefore in type II interconnect power estimation, we need to re-calculate the overall wire length distribution and layer assignment. The interconnect density function $i(l)$ for a system is now the sum of all interconnect density functions among all modules and busses, given by Equation (6):

$$i(l) = \sum_k i_k(l) \quad (6)$$

where subscript k iterates over all modules and busses. Using the same number of layers as in Table IV, the new length boundaries with consideration of structural interconnects are shown in table IX. Compared to Table IV, the boundaries for both global/intermediate and intermediate/local are reduced due to the reduced number of long interconnects. In other words, a higher portion of random interconnects can be assigned to the global and intermediate layers for reduced delay and in turn reduced buffer numbers. This may help to reduce interconnect power.

Considering the new layer assignment, we apply the power estimation method based on the

stochastic length distribution to each module independently and obtain the interconnect power for each module (see Table VII). We also apply concurrent repeater and FF insertion to obtain the interconnect power for busses (see Table VIII). In Table X, adding power for all modules and busses, we obtain the total type II interconnect power at the microarchitecture level and compare it with the type I interconnect power estimation. Based on this table, type I interconnect power estimation overestimates the interconnect power by $1.31X$ and $1.16X$ for the min-FF and min-power solutions, respectively. Part of the power reduction is due to the reduced number of long interconnects, which in turn reduces the number of equivalent repeaters and FFs. The equivalent repeaters are reduced by $3.08X$ and $1.74X$ for min-FF and min-power solutions, respectively. Compared to the min-FF solution, the min-power solution uses slightly fewer repeaters. With consideration of power used by FFs, the min-power solution reduces the full-chip interconnect power by 3.2% compared to the min-FF solution. On one hand, the min-power solution actually provides us the lower bound of full-chip interconnect power; on the other hand, power reduction by the min-power method compared to the min-FF method depends on specific interconnect length distribution and may not always be substantial. Furthermore, as a min-power solution may greatly reduce IPC, it is not necessarily used in practice.

3.3 Cycle-accurate Power Simulation

To obtain accurate activity for interconnects, we further incorporate our interconnect power models with concurrent repeater and FF insertion into the sim-outorder simulator of SimpleScalar toolset [2]. We perform the following cycle by cycle simulation: if a module is accessed, we count its active (dynamic + leakage) interconnect power, otherwise we only count its leakage power. On the other hand, for each bus, we count the number of bit-line transitions every cycle. The dynamic power in that cycle equals the number of transitions times the dynamic switching power per bus bit-line. Note

that the dynamic switching power is the full switching power ($\frac{1}{2}CV^2$) without the empirical fixed switching factor. The leakage power for each bus is always equal to the total number of bit lines times the leakage power per bus bit-line. By counting only leakage power for idle modules we implicitly consider clock gating.

We run simulations for total seven SPEC 2000 benchmarks: *bzip2*, *gcc*, *gzip*, *mcf*, *parser*, *mesa*, *equake*. Among them, *mesa* and *equake* are floating-point benchmarks, while the rest are integer benchmarks. During each simulation, the benchmark is first fast forwarded by 10 million instructions to avoid the startup effect, and is then simulated for 10 million instructions. Table XI reports the total type III interconnect power. By applying cycle-accurate simulations and clock gating, the average interconnect power by arithmetic mean⁴ for all benchmarks can be reduced by *1.71X* and *1.74X* for the min-FF and min-power solution, respectively, compared with type II interconnection power estimation. Compared with type I interconnect power estimation, the overall reduction of over-estimation is *2.24X* and *2.02X* for min-FF and min-power solution, respectively. Given such big differences in power, type III interconnect power estimation with cycle-accurate simulation is needed to obtain accurate interconnect power and validate power reduction innovations.

4 THROUGHPUT MAXIMIZATION CONSIDERING INTERCONNECT PIPELINING AND FLOORPLANNING OPTIMIZATION

In this section, we optimize throughput using BIPS (billion instructions per second) as the metric. We call interconnects with FFs inserted as pipelined interconnects, and compare pipelined interconnects and logic gates with voltage scaling. Then, we introduce throughput maximization by optimizing the clock frequency and floorplaning, respectively.

⁴ We assume each benchmark runs equally often.

4.1 *Throughput Metric and Voltage Scaling*

Our metric for throughput optimization is BIPS (Billion Instruction Per Second) defined as:

$$BIPS = \frac{IPC \times \text{clock_frequency}}{10^9} \quad (7)$$

It can be maximized by increasing either IPC or clock frequency. Raising supply voltage (V_{dd}) is often applied to obtain a higher clock frequency, and such technique is called voltage scaling. Similar to the first order approximation between gate delay and V_{dd} , a proper V_{dd} can be decided by to obtain a desired clock frequency freq for logic module⁵, where V_t is the threshold voltage. In this case, the inverter effective output resistance R_d used by repeater and FF insertion also varies with respect to V_{dd} . In our experiments, we assume that V_t is 20% of V_{dd} and $V_{dd} = 1\text{Volt}$ leads to 3GHz clock frequency as specified by the ITRS. We obtain R_d under different V_{dd} by SPICE simulation, and summarize the values for clock frequency, V_{dd} and R_d used in our experiments in Table XII.

However, the delay of pipelined interconnects behaves differently with respect to V_{dd} scaling. Figure 7 plots the normalized delays for logic gates and pipelined interconnects. It is easy to see that when V_{dd} increases, gates reduce delay much faster than pipelined interconnects and there is an increasing gap between them. Since we scale V_{dd} according to the gate delay, the pipelined interconnect can not sustain the same clock frequency increase as the logic gates and modules. Therefore, we have to re-design repeater and FF insertion for interconnects in order to obtain the increase of clock frequency decided by the logic modules.

4.2 *Throughput Maximization by Clock Frequency Scaling*

With FF insertion, IPC and clock frequency are no longer independent to each other. For a given microarchitecture and floorplan, the increased clock frequencies require that more FFs be inserted. This degrades IPC however. Therefore, there may exist an optimal clock frequency to maximize

⁵ We assume that local interconnects within the logic module and logic gates have the same performance scaling characteristics [1].

throughput when the clock frequency and IPC are well balanced.

We study throughput maximization for the same microarchitecture and floorplan as in Section 3. We evaluate BIPS with respect to clock frequencies between 2GHz and 4.5GHz. For each clock frequency, we first obtain the min-FF solution for concurrent repeater and FF insertion, then modify SimpleScalar according to the resulting FF insertion, and report the simulated IPC and BIPS in Figure 8. For all benchmarks, when clock frequency increases from 2GHz to 3GHz, although IPC slightly decreases, BIPS keeps increasing due to the increased clock frequency. When clock frequency exceeds 3GHz, IPC decreases severely due to FFs inserted on a critical path, such as data busses between LSQ and L1 d-cache. As a result, BIPS does not improve even when clock frequency is increased up to 4.5GHz. Figure 8 clearly shows that there does exist an optimal clock frequency for BIPS maximization for a given microarchitecture and floorplan, and this clock frequency is 3GHz in our example.

4.3 Throughput Maximization by Floorplanning Optimization

Floorplanning directly affects the lengths of structural interconnects, and in turn, the interconnect pipelining solution. By adjusting the floorplan, we may reduce interconnect pipeline stages for better IPC and BIPS. Because Figure 8 has shown a severe IPC degradation when the clock frequency increases from 3GHz to 3.5GHz, we target the 3.5GHz clock frequency for adjusting the microprocessor floorplan.

Figure 9 presents two floorplans, A and B, for the SuperScalar processors we study. Floorplan A is same as that in Figure 6 and floorplan B is the new floorplan optimized for IPC-critical interconnects. The differences between them are highlighted in the figure and include: (i) we move LSQ closer to L1 d-cache, and eliminate one FF between them. (ii) we distribute the four integer function units, remove one FF between RUU and IALU1/IALU2, but introduce one extra FF between RUU and

IMULT. Because multiplication and division take much longer than addition, IMULT has a much larger latency than IALU. Intuitively, the IPC gain of IALU1/IALU2 outweighs the IPC loss of IMULT. For a similar reason, we exchange the locations of FALU and FPMULT such that FALU is closer to RUU but FPMULT is further away.

As shown in Figure 10, floorplan B optimized for IPC-critical interconnects increases IPC (as well as BIPS) by 23.49%⁶. Although the IPC improvement is significant, floorplan B only reduces 5% of the total structural interconnect length, the objective to minimize in the conventional floorplan. Therefore, in the presence of interconnect pipelining, the floorplan should consider both the conventional objective of minimizing total interconnect length and the new objective of maximizing IPC.

5 CONCLUSIONS AND DISCUSSIONS

Considering structural interconnects, layer assignment, and concurrent repeater and FF (flip-flop) insertion, we have developed cycle-accurate microarchitecture-level power and throughput simulations and obtained an accurate modeling of interconnects at the early design stage. Experiments have shown that the simulation reduces over-estimation by up to 2.24X compared to the conventional power estimation based on purely stochastic interconnects and fixed switching factor. Given such a difference, cycle-accurate simulation becomes a necessity to validate microarchitecture innovations for power optimization.

With the presence of pipelined interconnects, we have shown that throughput is not always higher for an increased clock frequency, and there exists an optimal clock frequency to maximize throughput for a given microarchitecture and floorplan. We have illustrated that floorplanning optimized for IPC (instructions per cycle)-critical interconnects has little effect on the total

⁶ Although the performance improvement are specific to the floorplan we study, our method can be applied to study general cases

interconnect length but it improves throughput by 23.49%. Therefore, future floorplanning optimization should consider both the conventional objective of minimizing total interconnect length and the new objective of maximizing IPC.

As FF insertion becomes necessary to achieve the clock rate specified by ITRS, we conclude that the traditional design flow of optimizing IPC and clock rate separately is no longer valid, and coupled microarchitecture and layout optimization may improve both power efficiency and throughput. Such co-optimization has been further studied in recent work on automatic floorplanning optimization with interconnect pipelining [25-27].

In this paper, we assume two-pin interconnects. Similar assumption has been used extensively for early-stage estimation [12, 14]. In the future, we will extend our study to consider multi-pin interconnects in a fashion similar to [11].

REFERENCES

- [1] D. Sylvester and K. Keutzer, "Impact of small process geometries on microarchitectures in systems on a chip," Proceedings of the IEEE (2001), vol. 89, no. 4, pp. 467-489, 2001.
- [2] D. Burger and T. Austin, "The simplescalar tool set version 2.0". Technical Report, Computer Science Department, University of Wisconsin-Madison (1997).
- [3] W. Ye, N. Vijaykrishnan, M. Kandemir, and M. J. Irwin, "The design and use of simplepower: a cycle-accurate energy estimation tool," Proceedings of Design Automation Conference (2000).
- [4] D. Brooks, V. Tiwari, and M. Martonosi, "Wattch: A framework for architectural-level power analysis optimization," Proceedings of the International Symposium of Computer Architecture (2000).
- [5] W. Liao, J. M. Basile, and L. He, "Leakage power modeling and reduction with data retention", Proceedings of the International Conference on Computer-Aided Design (2002).
- [6] Lukas P. P. van Ginneken, "Buffer placement in distributed rc-tree networks for minimal elmore delay", Proceedings of the IEEE International Symposium on Circuits and Systems (1990), pp. 865-868.
- [7] C. J. Alpert and A. Devgan, "Wire segmenting for improved buffer insertion," Proceedings of the Design Automation Conference (1997), pp. 588-593.
- [8] J. Xiong, K. H. Tam, and L. He, "Buffer insertion considering process variation," Proceedings of the Design, Automation and Test in Europe, (2005), pp. 970-975.
- [9] J. Cong, Y. Fan, X. Yang, and Z. Zhang, "Architecture and synthesis for multi-cycle communication," Proceedings of the International Symposium on Physical Design, (2003), pp. 190-196.

- [10] R. Lu, G. Zhong, C.-K. Koh, and K.-Y. Chao, "Flip-flop and repeater insertion for early interconnect planning," Proceedings of Design, Automation and Test in Europe, (2002), pp. 690-695.
- [11] P. Cocchini, "Concurrent flip-flop and repeater insertion for high performance integrated circuits," Proceedings of the International Conference on Computer-Aided Design, (2002).
- [12] G. C. Pawan Kapur and K. C. Saraswat, "Power estimation in global interconnects and its reduction using a novel repeater optimization methodology," Proceedings of the Design Automation Conference, (2002).
- [13] V. D. J.A. Davis and J. Meindl, "A stochastic wire-length distribution for gigascale integration (GSI)-part 1: derivation and validation," IEEE Transaction on Electron Devices (1998), vol. 45, no. 3.
- [14] K. Banerjee and A. Mehrotra, "Power dissipation issues in interconnect performance optimization for sub-180 nm designs," Proceedings of Symposium on VLSI Circuits, (2002).
- [15] K. H. Tam and L. He, "Power-optimal dual-vdd buffered tree considering buffer stations and blockages," Proceedings of the Design Automation Conference, (2005).
- [16] Y. Chang, K. Tam, and L. He, "Power-optimal repeater insertion considering V_{dd} and V_t as design freedoms," Proceedings of the International Symposium on Low Power Electronics and Design, (2005).
- [17] J. Cong, A. Jagannathan, G. Reinman, and M. Romesis, "Microarchitecture evaluation with physical planning," Proceedings of the Design Automation Conference, (2003).
- [18] W. Liao and L. He, "Full-chip interconnect power estimation and simulation considering concurrent repeater and flip-flop insertion," Proceedings of the International Conference on Computer-Aided Design, (2003).

- [19] M. Ekpanyapong, J. R. Minz, T. Watwai, H.-H. S. Lee, and S. K. Lim, "Profile-guided microarchitectural floorplanning for deep submicron processor design," Proceedings of the Design Automation Conference, (2004), pp. 634-639.
- [20] C. Long, L. J. Simonson, W. Liao, and L. He, "Floorplanning optimization with trajectory piecewise-linear model for pipelined interconnects," Proceedings of the Design Automation Conference, (2004), pp. 640-645.
- [21] V. Nookala, Y. Chen, D. J. Lilja, and S. S. Sapatnekar, "Microarchitecture-aware floorplanning using a statistical design of experiments approach," Proceedings of the Design Automation Conference, (2005), pp. 579-584.
- [22] Berkeley Predictive Technology Model (BPTM) 0.10um SPICE Model Cards, July 2000.
- [23] D. Sylvester and K. Keutzer, "Getting to the bottom of deep submicron," Proceedings of the International Conference on Computer Aided Design, (1998), pp. 203-211.
- [24] N. Vasseghi and et al, "200MHz superscalar RISC processor circuit design issues," Proceedings of the IEEE International Solid-State Circuits Conference, (1996), pp. 356-357.
- [25] M. Ekpanyapong, J. R. Minz, T. Watwai, H. S. Lee, S. K. Lim, "Profile-guided microarchitectural floorplanning for deep submicron processor design," Proceedings of the ACM/IEEE Design Automation Conference (2004), pp. 634-639.
- [26] C. Long, L. J. Simonson, W. Liao, and L. He, "Floorplanning optimization with trajectory piecewise-linear model for pipelined interconnects," Proceedings of the ACM/IEEE Design Automation Conference (2004), pp. 640-645.
- [27] V. Nookala, Y. Chen, D. J. Lilja, and S. S. Sapatnekar, "Microarchitecture-aware floorplanning using a statistical design of experiments approach," Proceedings of the ACM/IEEE Design Automation Conference (2005), pp. 579-584.




Figure 1. The repeater and FF insertion problem in two-pin nets.

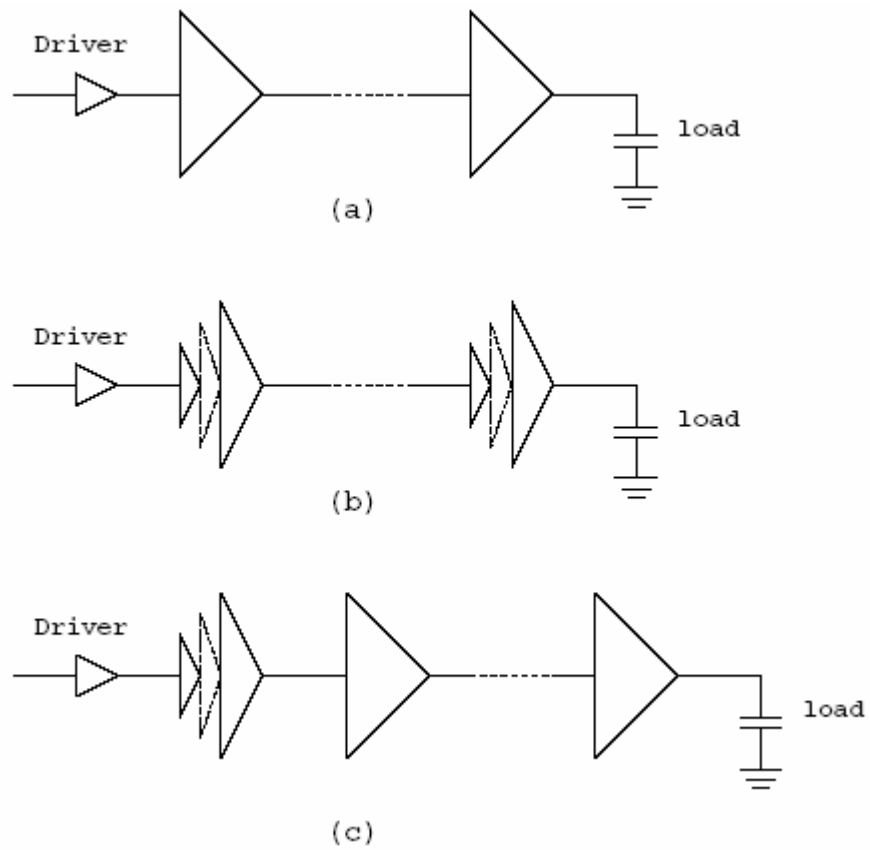


Figure 2. The three modes for repeater insertion: (a) single model; (b) cascaded model; and (c) hybrid model.




Figure 3. The power consumption with repeater insertion vs. different wire lengths. We choose 100nm technology and the delay target as 80% of clock period for a 3GHz system clock frequency. Only the results from the hybrid model are shown. No FF is inserted.

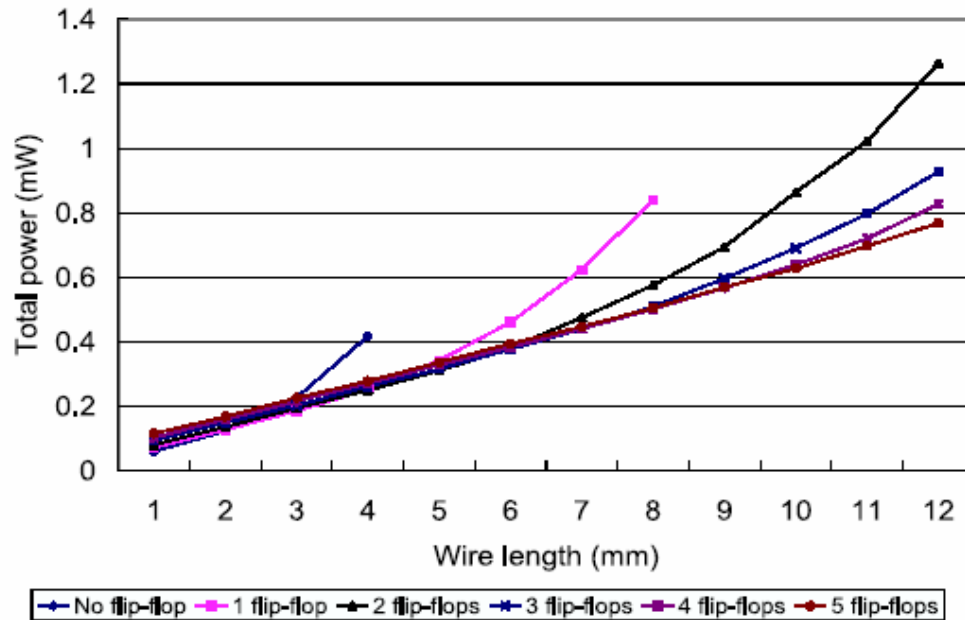


Figure 4. The power for different wire lengths, under different FF insertion. The experiment settings are the same as those in Figure 3.

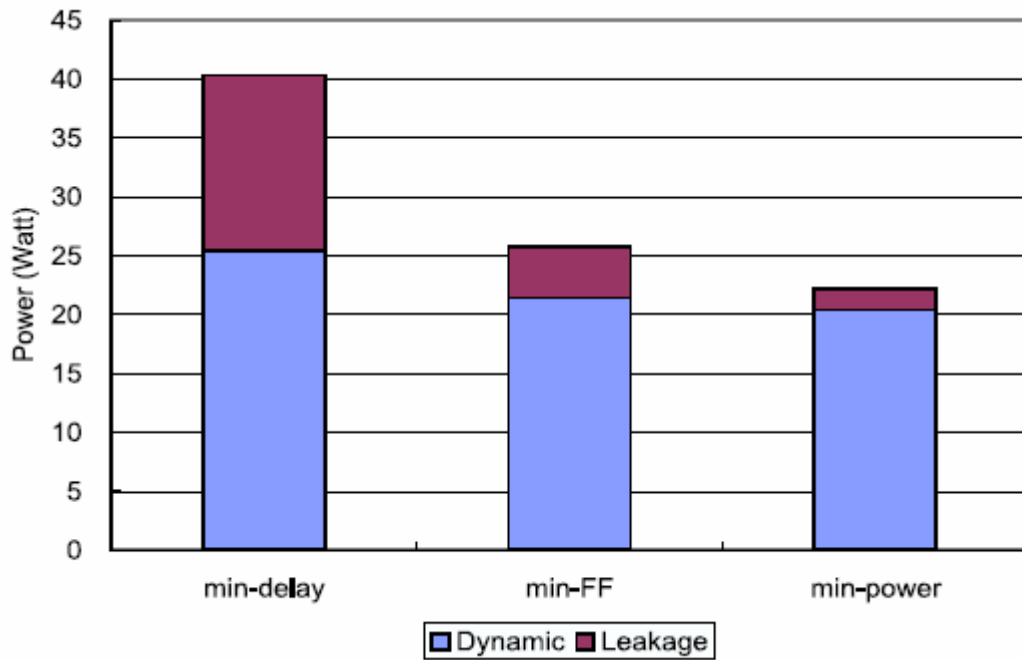


Figure 5. Type I interconnect power estimation for three repeater and FF insertion solutions.

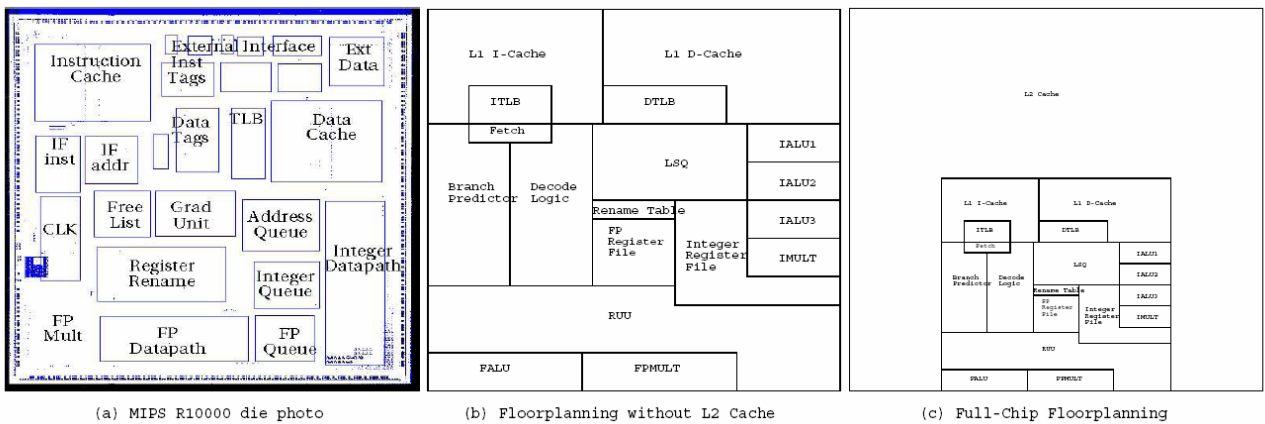


Figure 6. The die photo of the MIPS R10000 and the floorplanning similar to the MIPS R10000 used in our experiments.






Figure 7. The Delay for logic gate and pipelined interconnects with respect to different V_{dd} . Both gate and pipelined interconnect delay at 0.667V V_{dd} are chosen as a reference for normalization for each kind of delay. The gate delay is characterized by the delay of a minimum size inverter with a load of FO4. The pipelined interconnect delay is the delay of the bus between L1 i-cache and L2 cache with one FF inserted.



(a) IPC



(b) BIPS

Figure 8. The IPC and BIPS results for different clock frequencies with FF insertion.

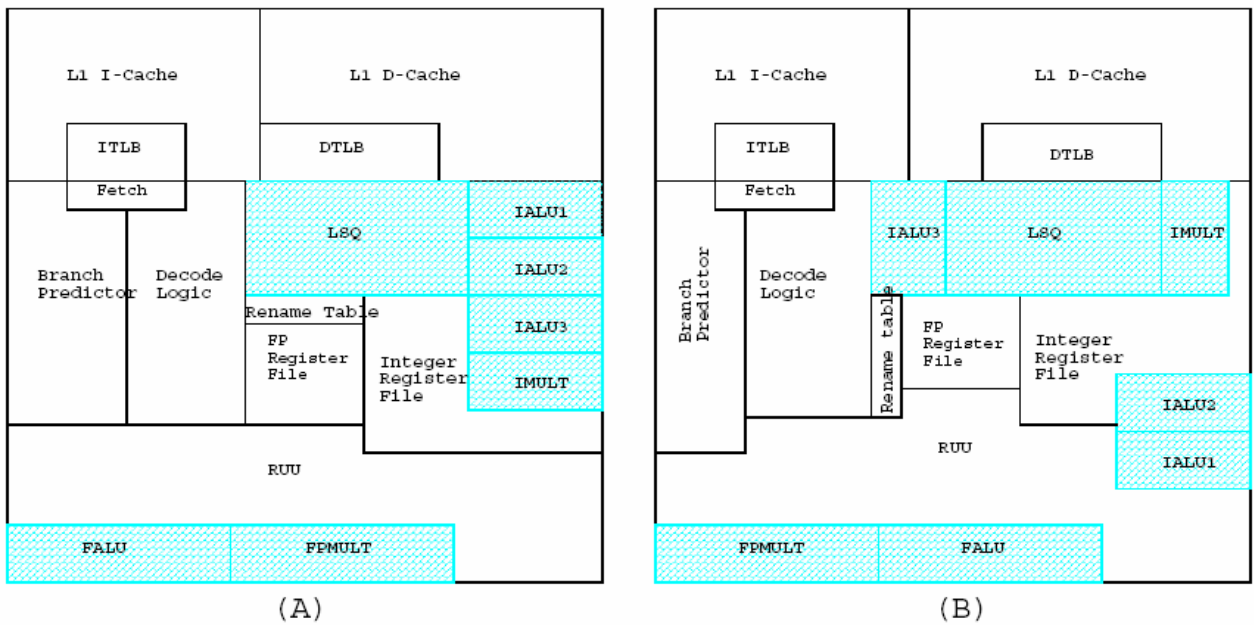


Figure 9. Two different floorplans with difference between them shadowed: (A) the floorplan in Figure 6; and (B) the new floorplan after adjustment. L2 cache is omitted.




Figure 10. The IPC comparison between two floorplans in Figure 9 with FF insertion. The clock frequency is 3.5GHz.

TABLE I
TECHNOLOGY PARAMETERS

<i>Technology</i>		100nm		
<i>I_{off}($\mu A/u$)</i>		6.33		
<i>A</i>		0.15		
<i>Minimum size inverter</i>	<i>R0 (KΩ)</i>	9.79		
	<i>C0 (fF)</i>	0.91		
	<i>Cp (fF)</i>	0.75		
<i>FF</i>	<i>CF (fF)</i>	16.6		
	<i>SF</i>	10		
	Interconnects			
		Global	Intermediate	Local
<i>Width (nm)</i>		335	160	122.5
<i>Height (nm)</i>		670	272	196
<i>RW (KΩ/m)</i>		89.106	459.559	832.986
<i>CW (pF/m)</i>		204.802	180.068	176.188

TABLE II

THE POWER CONSUMPTION FOR DIFFERENT WIRE LENGTHS AND DIFFERENT CLOCK FREQUENCIES, UNDER THREE MODELS FOR REPEATER INSERTION. THE SYMBOL “R#” MEANS THE NUMBER OF REPEATERS. THE COLUMN OF “POWER REDUCTION” IS THE POWER REDUCTION BY MIN-POWER SOLUTION COMPARED WITH MIN-FF SOLUTION UNDER HYBRID MODEL.

GHz	Wire length (mm)	Min-FF									Min-Power			Power reduction
		Single Model			Cascaded Model			Hybrid Model			Power (mW)	FF	R#	
		Power (mW)	FF	R #	Power (mW)	F	R #	Power (mW)	F	R #				
	4	0.0751	0	1	0.0751	0	1	0.0751	0	1	0.0751	0	1	0.00%
1	8	0.1956	0	2	0.1956	0	2	0.1903	0	2	0.1524	1	1	19.93%
	10	0.3415	0	4	0.3415	0	4	0.29	0	3	0.1921	2	1	33.74%
	4	0.2005	0	2	0.1793	0	1	0.1793	0	1	0.1573	1	1	12.29%
2	8	0.4054	1	2	0.3631	1	1	0.3631	1	1	0.3181	2	1	12.38%
	10	0.4113	2	1	0.6785	1	2	0.5268	1	2	0.3981	3	1	24.44%
	4	0.2514	1	1	0.2514	1	1	0.4168	0	2	0.2489	2	1	40.28%

3	8	0.5094	3	1	0.5752	2	1	0.8403	1	2	0.5009	4	1	40.39%
10	0.6383	4	1	1.1142	2	1	0.8636	2	2	0.6269	5	1	27.41%	

TABLE III
THE LONGEST WIRE THAT REPEATER INSERTION ALONE IS ABLE TO MEET THE DELAY TARGET WITHOUT FF INSERTION. THE DELAY TARGET IS 80% OF THE CLOCK PERIOD.

<i>Clock (GHz)</i>	Longest wire under delay constrain (mm)		
	Single model	Cascaded model	Hybrid model
1	11.41	12.07	14.92
2	4.41	5.40	6.91
3	2.38	3.38	4.34

TABLE IV
THE LENGTH BOUNDARIES DECIDE BY LAYER ASSIGNMENT WITH THE GATE PITCH AND SYSTEM CLOCK FREQUENCY

<i>Technology</i>		100nm
Total gate count		8615384
Gate area		6.5 μm^2
Gate pitch		2.55 μm
Global and intermediate interconnect boundary	In gate pitch	1389
	In mm	2.6163
Intermediate and local interconnect boundary	In gate pitch	85
	In mm	0.1479

TABLE V
THE NUMBER OF REPEATERS AND FFs INSERTION IN ALL THREE SOLUTIONS IN FIGURE 5. THE MIN-DELAY SOLUTION DOES NOT USE ANY FF INSERTION

Solution	Number of equivalent repeaters	Number of FFs
Min-delay	23559861	N/A
Min-FF	6938919	11850
Min-power	2861590	84378

TABLE VI

THE CONFIGURATION OF THE SUPERSCALAR PROCESSORS WE SIMULATE.

Parameter	Value
Processor Core	
RUU size	64 instructions
LSQ size	32 instructions
Fetch Queue size	8 instructions
Fetch width	4 instructions/cycle
Decode width	4 instructions/cycle
Issue width	4 instructions/cycle
Commit width	4 instructions/cycle
Functional Units	3 integer addition, 1 integer multiplication/division, 1 FP addition, 1 FP multiplication/division
Branch Predictor	Combined, Bimodal 4K table 2-Level 1K table, 10-bit history 4K chooser
Memory Hierarchy	
L1 instruction-cache	1 read port and 1 write port, 32K, 4-way (LRU) 32B blocks, 1-cycle load-use penalty
L1 data-cache	2 read/write ports, 32K, 4-way (LRU) 32B blocks, 1-cycle load-use penalty
L2 unified cache	512K, 8-way (LRU) 64B blocks, 12-cycle latency
TLB	128 entry, fully associative 30-cycle miss latency
Main memory	256-cycle latency

TABLE VII
MODULES AND THEIR CORRESPONDENT MICROARCHITECTURE, GATE COUNT, AND POWER UNDER DIFFERENT REPEATER AND FLIP-FLOP INSERTION MODELS. THE CACHES AND REGISTER FILES ARE NOT CONSIDERED BECAUSE THEY ARE PURELY MEMORY ARRAYS.

Module	Microarchitecture	Gate count	Min-FF (W)			Min-Power (W)		
			Total	Dynamic	Leakage	Total	Dynamic	Leakage
Fetch	Fetch queue	30154	0.0221	0.022	0	0.0221	0.022	0
Decode	Decode queue	1270769	2.6948	2.544	0.1508	2.6659	2.5385	0.1274
Branch	Branch predictor	1193231	2.4921	2.3567	0.1354	2.4686	2.3525	0.1161
Rename	Renaming table	280000	0.4074	0.3963	0.0111	0.4074	0.3963	0.011
RUU	Register Update Unit	2373538	5.8504	5.4163	0.4341	5.6859	5.3754	0.3106

LSQ	Load/Store queue	1300923	2.7744	2.6174	0.157	2.7433	2.6114	0.1318
IALU[1-3]	One integer ALU	318769	0.4799	0.4659	0.014	0.4799	0.4659	0.0139
IMULT	Integer multiplier	318769	0.4799	0.4659	0.014	0.4799	0.4659	0.0139
FALU	Floating-point ALU	598769	1.0573	1.0155	0.0418	1.0559	1.0156	0.0403
FPMULT	Floating-point multiplier	598769	1.0573	1.0155	0.0418	1.0559	1.0156	0.0403
Sum:		8921231	18.2753	17.2474	1.0279	18.0243	17.1909	0.8334

TABLE IX
NEW INTERCONNECT LENGTH BOUNDARIES BETWEEN LOCAL, INTERMEDIATE AND GLOBAL WIRES, AFTER WE DISTINGUISH THE STRUCTURAL INTERCONNECTS AND RANDOM INTERCONNECTS.

Global and intermediate interconnect boundary	In gate pitch	499
	In mm	1.272
Intermediate and local interconnect boundary	In gate pitch	24
	In mm	0.061

TABLE VIII
THE CONFIGURATIONS OF ALL BUSSES, AND THEIR PER-BIT-LINE POWER PER SWITCH AFTER DIFFERENT REPEATER AND FLIP-FLOP INSERTION MODELS. THE INT_CDB AND FP_CDB ARE THE RESULT BUSSES FOR INTEGER AND FLOATING-POINT UNITS, RESPECTIVELY. NOTE THAT THESE VALUES ARE FIRST-ORDER ESTIMATION AND ARCHITECTURE-SPECIFIC.

Tow modules the bus connects	Bus width with ECC	Bus length (mm)	Power per bit line (mW)				
			Hybrid + min-FF		Hybrid + min-power		
			Dynamic	Leakage	Dynamic	Leakage	
Address Bus							
Fetch	ITLB	39	0.6693	0.0346	0.0063	0.0346	0.0063
Fetch	L1 i-cache	39	1.9661	0.1014	0.0184	0.1014	0.0184
Fetch	Branch	39	2.5936	0.1412	0.0367	0.1412	0.0367
LSQ	DTLB	39	1.50597	0.0769	0.0127	0.0769	0.0127
DTLB	L1 d-cache	39	2.5099	0.1358	0.0342	0.1358	0.0342
L1 d-cache	L2 cache	39	6.8605	0.4309	0.1696	0.4309	0.1696
L1 i-cache	L2 cache	39	7.8645	0.534	0.2659	0.534	0.2659
Data Bus							
Fetch	L1 i-cache	288	1.9661	0.1014	0.0184	0.1014	0.0184
Fetch	Decode	288	2.5936	0.1412	0.0367	0.1412	0.0367
Decode	Rename	288	2.1753	0.1137	0.0228	0.1137	0.0228
Rename	RUU	228	2.5936	0.1412	0.0367	0.1412	0.0367

Rename	LSQ	24	1.8406	0.0945	0.0165	0.0945	0.0165
LSQ	L1 d-cache	72	4.0159	0.2754	0.1462	0.2754	0.1462
RUU	IALU1	144	7.9482	0.5453	0.2785	0.5453	0.2785
RUU	IALU2	144	7.1115	0.453	0.1874	0.453	0.1874
RUU	IALU3	144	6.1912	0.3711	0.1342	0.3711	0.1342
RUU	IMULT	144	5.2709	0.2936	0.0747	0.2936	0.0747
RUU	FALU	144	4.6016	0.2523	0.057	0.2523	0.057
RUU	FPMULT	144	2.0079	0.1037	0.019	0.1037	0.019
RUU	Int reg	576	3.4303	0.2121	0.0848	0.2121	0.0848
RUU	FP reg	288	1.5478	0.0792	0.0133	0.0792	0.0133
Int reg	LSQ	576	3.1374	0.1864	0.0709	0.1864	0.0709
FP reg	LSQ	288	3.0538	0.1758	0.0595	0.1758	0.0595
L1 d-cache	L2 cache	72	6.8605	0.4309	0.1696	0.4309	0.1696
L1 i-cache	L2 cache	72	7.8645	0.534	0.2659	0.534	0.2659
	INT_CDB	288	7.9482	0.5453	0.2785	0.5453	0.2785
	FP_CDB	288	4.6016	0.2523	0.057	0.2523	0.057

TABLE X
TOTAL INTERCONNECT POWER FOR BOTH TYPE I AND II INTERCONNECT POWER ESTIMATION. THE “REP” AND “FF” REPRESENT THE NUMBER OF EQUIVALENT REPEATERS AND FF, RESPECTIVELY. THE POWER IS IN THE UNIT OF WATT.

	Min-FF					Min-power				
	Total power	Dynamic	Leakage	Rep	FF	Total power	Dynamic	Leakage	Rep	FF
Type I	25.82	21.43	4.39	6938919	11850	22.2	20.39	1.81	2861590	84378
Type II	19.77	18.33	1.43	2249406	1690	19.14	18.15	0.99	1648925	15574
Reduction	1.31X	1.17X	3.06X	3.08X	7.01X	1.16X	1.12X	1.84X	1.74X	5.42X

TABLE XI
THE TYPE III INTERCONNECT POWER WITH BOTH RANDOM AND STRUCTURAL INTERCONNECTS. LEAKAGE POWER IS OMITTED BECAUSE IT IS NOT AFFECTED BY CLOCK GATING. THE UNIT OF POWER IS WATT.

Power		Min-FF power		Min-power power	
		Total	Dynamic	Total	Dynamic
	<i>bzip2</i>	13.03	11.6	12.46	11.47
	<i>gcc</i>	7.26	5.83	6.76	5.77
	<i>gzip</i>	11.15	9.72	10.6	9.62
	<i>mcf</i>	14.36	12.93	13.8	12.82
	<i>parser</i>	12.2	10.77	11.64	10.66
	<i>mesa</i>	11.22	9.79	10.67	9.68

<i>equake</i>	11.5	10.07	10.95	9.97
Average	11.53	10.07	10.98	10.00
Type II	19.77	18.33	19.14	18.15
Difference	1.71X	1.81X	1.74X	1.82X

TABLE XII
SUPPLY VOLTAGE V_{dd} AND EFFECTIVE OUTPUT RESISTANCE FOR A MINIMUM SIZE INVERTER (R_d) UNDER DIFFERENT CLOCK FREQUENCIES.

Clock frequency (GHz)	2	2.5	3	3.5	4	4.5
V_{dd} (V)	0.667	0.833	1.0	1.167	1.333	1.5
R_d (k Ω)	11.96	9.18	9.79	6.98	6.46	6.09

BIOGRAPHIES

Weiping Liao received B.S. degree and M.S. degrees, both in physics, from the University of Science and Technology of China, Hefei, China, in 1996 and 1999, respectively, the M.S. degree in computer engineering from the University of Wisconsin, Madison in 2002, and the Ph.D. degree in electrical engineering from the University of California at Los Angeles in 2005. He is currently a senior architect at NVIDIA Corporation, Santa Clara, CA, focused on graphics processor design.

Lei He obtained Ph.D. degree in computer science from UCLA in 1999. He joined the faculty of electrical engineering, UCLA in 2002. He was a faculty member at University of Wisconsin, Madison between 1999 and 2001. His research interests include (i) modeling and design considering signal integrity and process variation for integrated circuits and packages, (ii) programmable logic fabrics and reconfigurable embedded systems, and (iii) power-efficient circuits and systems. He has published over 110 technical papers and is technical program committee member for a number of conferences, including Design Automation Conference, International Symposium on Low Power Electronics and Design, and International Symposium on Field Programmable Gate Array.

He was granted the Dimitris N. Chorafas Foundation Prize for Engineering and Technology in 1997, Distinguished Ph.D. Award from the UCLA Henry Samueli School of Engineering and Applied Science in 2000, US National Science Foundation CAREER award in 2000, UCLA Chancellor's faculty career development award (highest class) in 2003, IBM faculty partner award in 2003, and Northrop Grumman Excellence in Teaching award in 2005.