Optimizing Pipelines for Power and Performance

Viji Srinivasan, David Brooks, Michael Gschwind, Pradip Bose,
Victor Zyuban, Philip N. Strenski, Philip G. Emma
IBM T.J. Watson Research Center
Yorktown Heights, NY 10598

Abstract

During the concept phase and definition of next generation high-end processors, power and performance will need to be weighted appropriately to deliver competitive cost/performance. It is not enough to adopt a CPI-centric view alone in early-stage definition studies. One of the fundamental issues confronting the architect at this stage is the choice of pipeline depth and target frequency. In this paper we present an optimization methodology that starts with an analytical power-performance model to derive optimal pipeline depth for a superscalar processor. The results are validated and further refined using detailed simulation based analysis. As part of the power-modeling methodology, we have developed equations that model the variation of energy as a function of pipeline depth. Our results using a set of SPEC2000 applications show that when both power and performance are considered for optimization, the optimal clock period is around 18 FO4. We also provide a detailed sensitivity analysis of the optimal pipeline depth against key assumptions of these energy models.

1 Introduction

Current generation high-end, server-class processors are performance-driven designs. These chips are still somewhat below the power and power density limits afforded by the package/cooling solution of choice in server markets targeted by such processors. In designing future processors, however, energy efficiency is known to have become one of the primary design constraints [7, 1].

In this paper, we analyze key constraints in choosing the "optimal" pipeline depth (which directly influences the frequency target) of a microprocessor. The choice of pipeline depth is one of the fundamental issues confronting the architect/designer during the very early stage microarchitecture definition phase of high performance, power-efficient processors. Even from a performance-only viewpoint, this issue has been important, if only to understand the limits to which pipelining can scale in the context of real workloads [15, 5, 9, 11, 21]. In certain segments of the market (typically desktops and low-end servers), there is often

a market-driven tendency to equate delivered end performance with the frequency of the processor. Enlightened customers do understand the value of net system performance; nonetheless, the instinctive urge of going primarily for the highest frequency processor in a given technology generation is a known weakness even among savvy end users and therefore processor design teams. Recent studies [9, 11, 21] seem to suggest that there is still room to grow in the pipeline depth game, with performance optima in the range of 8-11 FO4 inverter delays[1] per stage (consisting of 6-8 FO4 logic delay and 2-3 FO4 latch delay) for current out-of-order superscalar design paradigms. However, even in these performance-centric analysis papers, the authors do point out the practical difficulties of design complexity, verification and power that must be solved in attaining these idealized limits. Our goal in this paper is to examine the practical, achievable limits when power dissipation constraints are also factored in. We believe that such analysis is needed to realistically bound the scalability limit in the next few technology generations. In particular, power dissipation must be carefully minimized to avoid design points which on paper promise ever higher performance, yet under normal operating conditions, with commodity packaging and air cooling, only deliver a fraction of the theoretical peak performance.

In this paper, we first develop an analytical model to understand the power and performance tradeoffs for superscalar pipelines. From this model, we derive the optimal pipeline depth as a function of both power and performance. Subsequently, the results are validated and further refined using a detailed cycle-accurate simulator of a current generation superscalar processor. The energy model for the core pipeline is based on circuit-extracted power analysis for structures in a current, high performance PowerPC processor. We then derive a methodology for scaling these energy models to deeper and shallower pipelines. A key component of this methodology is the scaling of latch count growth as a function of pipeline depth. With these performance and power models we attempt to determine the optimal pipeline

---

[1]Fan-out-of-four (FO4) delay is defined as the delay of one inverter driving four copies of an equally sized inverter. The amount of logic and latch overhead per pipeline stage is often measured in terms of FO4 delay which implies that deeper pipelines have smaller FO4.

depth for a particular power-performance metric.

Our results based on an analysis of the TPC-C transaction processing benchmark and a large set of SPEC2000 programs indicate that a power-performance optimum is achieved at much shallower pipeline depths than a purely performance-focused evaluation would suggest. In addition, we find there is a range of pipeline depths for which performance increases can be achieved with a modest sacrifice in power-performance efficiency. Pipelining beyond that range leads to drastic reduction in power-performance efficiency.

The contributions of this paper are (1) energy models for both dynamic and leakage power that capture the scaling of different power components as a function of pipeline depth; (2) an analytical performance model that can predict optimal pipeline depth as well as the shift in the optimal point, when combined with the energy models; (3) cycle-accurate, detailed power-performance simulation with a thorough sensitivity analysis of the optimal pipeline depth against key energy model parameters.

This paper is structured as follows: We discuss the prior, related work in Section 2. In Section 3, we describe the proposed analytical model to study pipeline depth effects in superscalar processors. Section 4 presents the simulation-based validation methodology. In Section 5 we present results using both the analytical model and a detailed simulator. In Section 6 we present a detailed sensitivity analysis to understand the effect of variations in key parameters of the derived power models on the optimal pipeline depth. We conclude in Section 7, with pointers to future work.

2 Related Work

Previous work has studied the issue of "optimal" pipeline depth exclusively under the constraint of maximizing the performance delivered by the microprocessor.

An initial study of optimal pipeline depths was performed by Kunkel and Smith in the context of supercomputers [15]. The machine modeled in that study was based on a Cray-1S, with delays being expressed as ECL gate levels. The authors studied the achievable performance from scalar and vector codes as a function of gate levels per pipeline stage for the Livermore kernels. The study demonstrated that vector codes can achieve optimum performance by deep pipelining, while scalar (floating-point) workloads reach an optimum at shallower pipelines.

Subsequently, Dubey and Flynn [5] revisited the topic of optimal pipelining in a more general analytical framework. The authors showed the impact of various workload and design parameters. In particular, the optimal number of pipeline stages is shown to decrease with increasing overhead of partitioning logic into pipeline stages (i.e., clock skew, jitter, and latch delay). In this model, the authors considered only stalls due to branch mispredictions and did not consider data dependent stalls due to memory or register dependencies.

More recently, several authors have reexamined this topic in the context of modern superscalar processor microarchitectures. Hartstein and Puzak [9] treat this problem analytically and verify based on detailed simulation of a variety of benchmarks for a 4-issue out-of-order machine with a memory-execute pipeline. Simulation is also used to determine the values of several parameters of their mathematical model, since these cannot be formulated axiomatically. They report optimal logic delay per pipeline stage to be 7.7 FO4 for SPEC2000 and 5.5 FO4 for traditional and Java/C++ workloads. Assuming a latch insertion delay of 3 FO4, this would result in a total delay of about 10.7 FO4 and 8.5 FO4 per pipeline stage, respectively.

Hrishikesh et al. [11] treat the question of logic depth per pipeline stage empirically based on simulation of the SPEC2000 benchmarks for an Alpha 21264-like machine. Based on their assumed latch insertion delay of 1.8 FO4, they demonstrate that a performance-optimal point is at logic delay of 6.0 FO4. This would result in a total pipeline delay of about 8 FO4.

Sprangle and Carmean [21] extrapolate from the current performance of the Pentium 4 using IPC degradation factors for adding a cycle to critical processor loops, such as ALU, L1 and L2 cache latencies, and branch miss penalty for a variety of application types. The authors compute an optimal branch misprediction pipeline depth of 52 stages, corresponding to a pipeline stage total delay of 9.9 FO4 (based on a logic depth of 6.3 FO4 and a latch insertion delay of 3.6 of which 3 FO4 are due to latch delay and 0.6 FO4 represent skew and jitter overhead).

All of the above studies (as well as ours) assume that microarchitectural structures can be pipelined without limitation. Several authors have evaluated limits on the scalability and pipelining of these structures [6, 19, 22, 4, 11].

Collectively, the cited works on optimal pipelining have made a significant contribution to the understanding of workloads and their interaction with pipeline structures by studying the theoretical limits of deep pipelining. However, prior work does not address scalability with respect to increased power dissipation that is associated with deeper pipelines. In this work, we aim to build on this foundation by extending the existing analytical models and by proposing a power modeling methodology that allows us to estimate optimal pipeline depth as a function of *both* power and performance.

3 Analytical Pipeline Model

In the concept phase definition studies, the exact organization and parameters of the target processor are not known. As such, a custom, cycle-accurate power-performance simulator for the full machine is often not available or relevant. Therefore, the use of analytical reasoning models supplemented by workload characterization and limit studies (obtained from prior generation simulators or trace analysis programs) is common in real design groups. We

present such an analytical model to understand the power-performance optima and tradeoffs during the pre-simulation phase of a design project.

Figure 1 shows a high-level block diagram of the pipeline model used in our analysis. Our primary goal is to derive the optimum pipeline depth for the various execution units by estimating the various types of stalls in these pipes while using a perfect front-end for the processor. Although Figure 1 shows only one pipe for each unit (fixed point, floating point, load/store, and branch), the model can be used for a design with multiple pipes per unit as well.
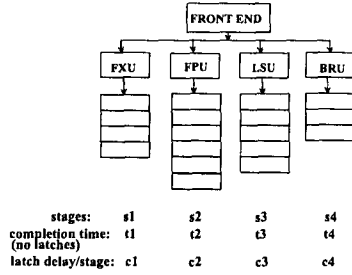

Figure 1. Pipeline Model

In Figure 1, let t_i be the latch-free logic time to complete an operation in pipe i, and s_i be the number of pipeline stages of pipe i. Assuming the same clock frequency for all the pipes, $t_i/s_i = t_j/s_j$, $\forall i, j$.

If c_i is the latch overhead per stage for pipe i, the total time per stage of pipe i is $T_i = ((t_i/s_i) + c_i)$, $\forall i$. As derived by Dubey and Flynn in [7], and Larson and Davidson (cited in Chapter 2 of [14]), the throughput of the above machine in the absence of stalls is given by $G = \sum(\frac{1}{T_i})$.

We now extend this baseline model to include the effect of data-dependent stalls. Workload analysis can be used to derive first-cut estimates of the probability that an instruction n depends on another instruction j for all instructions n, and can be used to estimate the frequency of pipeline stalls. This is illustrated in the example below, where an FXU instruction $(i + 1)$ depends on the immediately preceding instruction i, and will be stalled for $(s_1 - 1)$ stages assuming a register file bypass to forward the results. Similarly, another FXU instruction $(j+2)$ depends on instruction j, and will be stalled $(s_1 - 2)$ stages. Note that in the above workload analysis, if the source operands of an instruction i are produced by more than one instruction, the largest of the possible stalls is assigned to i.

```
inst (i)    add r1 = r2, r3
inst (i+1)  and r4 = r1, r5 --
stalled for (s1 - 1) stages

inst (j)    add r1 = r2, r3
inst (j+1)  or r6 = r7, r8
inst (j+2)  and r4 = r1, r5 --
stalled for (s1 - 2) stages
```

$$T_{fxu} = T_1 + Stall_{fxu-fxu} * T_1 + Stall_{fxu-fpu} * T_2 + Stall_{fxu-lsu} * T_3 + Stall_{fxu-bru} * T_4$$

where $Stall_{fxu-fxu} = f_1 * (s_1 - 1) + f_2 * (s_1 - 2) + \ldots$

The above equation represents the time to complete an FXU operation in the presence of stalls; f_i is the conditional probability that an FXU instruction m depends on another FXU instruction $(m - i)$ for all FXU instructions m, provided that instruction $(m - i)$ is the producer with the largest stall for instruction m. Similar expressions can be derived for T_{fpu}, T_{lsu}, and T_{bru}, the completion times of an FPU, LSU, and BRU operation, respectively. To account for superscalar ($>$ 1) issue widths, the workload analysis assumes a given issue width along with the number of execution pipes of various types (FXU, FPU, BRU, LSU), and issues independent instructions as an instruction bundle such that the bundle width $\leq$ issue width. Thus, the distance between dependent instructions is the number of instruction bundles issued between them. To account for the dependent instruction stalls due to L1 data cache misses we use a functional cache simulator to determine cache hits and misses. In addition, we split the load/store pipe into two, namely, load-hit and load-miss pipe; thereby, steering all data references that miss in the data cache to the load miss pipeline which results in longer stall times for the dependent instructions. Since the workload analysis is independent of the machine architecture details and uses only the superscalar issue width to determine the different stalls, it suffices to analyze each application once to derive the stalls.

The stalls modeled so far include only hazards in the execution stages of the different pipes. However, these pipes could also be waiting for instructions to arrive from the front-end of the machine. If u_i represents the fraction of time pipe i has instructions arriving from the front-end of the machine, the equation below gives the throughput of the pipeline in the presence of stalls. Note that $u_i = 0$ for unutilized pipelines, and $u_i = 1$ for fully utilized pipelines.

$$G = \frac{u_1}{T_{fxu}} + \frac{u_2}{T_{fpu}} + \frac{u_3}{T_{lsu}} + \frac{u_4}{T_{bru}}$$

Thus far, we focused on deriving the throughput of a pipeline as a function of the number of pipeline stages. In order to optimize the pipeline for both power and performance we use circuit-extracted power models described in Section 4.2.

If P_{s_i} is the total power of pipe i derived from the power models, the energy-delay product for pipe i is given by $ED = P_{s_i}/G^2$. Hence, the optimal value of s_i which minimizes the energy-delay product can be obtained using $d(ED)/ds_i = 0$. Note that depending on the class of processors the desired metric for optimization could be $d((BIPS)^\gamma/W)/ds_i = 0$ where $\gamma \geq 0$ is an exponent whose value can be fixed in specific circuit tradeoff contexts [26]; BIPS is billions of instructions per second.

335

4 Performance and Power Methodology

In this section, we describe the performance simulator used in this study as well as the details of our power modeling toolkit and the methodology that we use to estimate changes in power dissipation as we vary the pipeline depth of the machine.

Fetch Latencies		Decode Latencies	
Latency Parms	STD/INF	Latency Parms	STD/INF
NFA Predictor	1/0	Multiple Decode	2/0
L2 ICache	11/0	Millicode Decode	2/0
L3 (Instruction)	85/0	Expand String	2/0
I-TLB Miss	10/0	Mispredict Cycles	3/0
L2 I-TLB Miss	50/0	Register Read	1/1
Execution Pipe Latencies		Load/Store Latencies	
Latency Parms	STD/INF	Latency Parms	STD/INF
Fix Execute	1/1	L1 D-Load	3/3
Float Execute	4/4	L2 D-Load	9/9
Branch Execute	1/1	L3 (Data)	77/0
Float Divide	12/12	Load Float	2/2
Integer Multiply	7/7	D-TLB Miss	7/0
Integer Divide	35/35	L2 D-TLB Miss	50/0
Retire Delay	2/2	StoreQ Forward	4/0

Table 1. Latencies for 19 FO4 Design Point
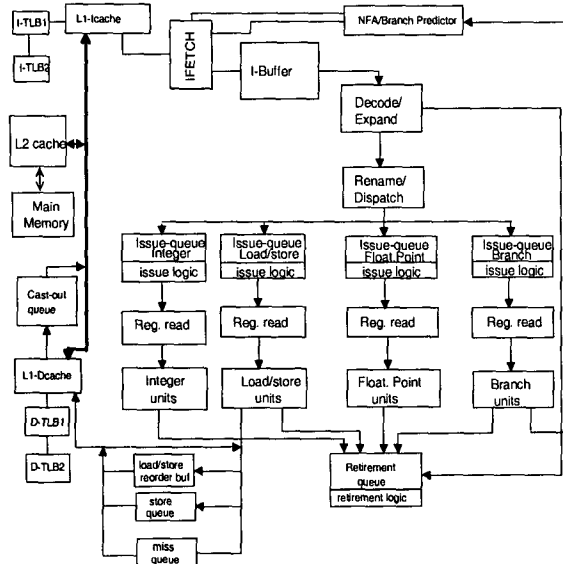
4.1 Performance Simulation Methodology

Figure 2. Modeled Processor Organization

We utilize a generic, parameterized, out-of-order 8-way superscalar processor model called Turandot [16, 17] with 32KB I and D-caches and a 2MB L2 cache. The overall pipeline structure (as reported in [16]), is repeated here in Figure 2. The modeled baseline microarchitecture is similar to a current generation microprocessor. As described in [16], this research simulator was calibrated against a pre-RTL, detailed, latch-accurate processor model. Turandot

supports a large number and parameters including configurable pipeline latencies discussed below.

Table 1 details the latency values in processor cycles for the 19 FO4 base design point of this study. We assume a 2 FO4 latch overhead and 1 FO4 clock skew and jitter overhead. The 19 FO4 latency values are then scaled with the FO4-depth (after accounting for latch and clock skew overhead). Each latency in Table 1 has two values: the first labeled STD, is for our detailed simulation model, and the second labeled INF, assumes infinite I-Cache, I-TLB, D-TLB, and a perfect front-end. The INF simulator model is used for validating the analytical model described in Section 3.

4.2 Power Simulation Methodology

To estimate power dissipation, we use the PowerTimer toolset developed at IBM T.J. Watson Research Center [3, 1] as the starting point for the simulator used in this work. PowerTimer is similar to power-performance simulators developed in academia [2, 23, 24], except for the methodology to build energy models.


Figure 3. PowerTimer Energy Models

Figure 3 above depicts the derivation of the energy models in PowerTimer. The energy models are based on circuit-level power analysis that has been performed on structures in a current, high performance PowerPC processor. The power analysis has been performed at the macro level using a circuit-level power analysis tool [18]. Generally, multiple macros combine to form one micro-architectural level structure which we will call a sub-unit. For example, the fixed-point issue queue (one sub-unit) might contain separate macros for storage memory, comparison logic, and control. Power analysis has been performed on each macro to determine the macro's unconstrained (no clock gating) power as a function of the input switching factor. In addition, the hold power, or power when no switching is occurring (SF = 0%), is also determined. Hold power primarily consists of power dissipated in latches, local clock buffers, the global clock network, and data-independent fraction of the arrays. The switching power, which is primarily combinatorial logic and data-dependent array power dissipation, is the additional power that is dissipated when switching factors are applied to the macro's primary inputs. These two pieces of data allow us to form simple linear equations for each macro's power. The energy model for a sub-unit is determined by summing the linear equations for each macro

within that sub-unit. We have generated these power models for all microarchitecture-level structures (sub-units) modeled in our research simulator [16, 17]. PowerTimer models over 60 microarchitectural structures which are defined by over 400 macro-level power equations.

PowerTimer uses microarchitectural activity information from the Turandot model to scale down the unconstrained hold and switching power on a per-cycle basis under a variety of clock gating assumptions. In this study, we use a realistic form of clock gating which considers the applicability of clock gating on a per-macro basis to scale down either the hold power or the combined hold and switching power depending on the microarchitectural event counts. We determine which macros can be clock gated in a fine-grained manner (per-entry or per-stage clock gating) and which can be clock gated in a coarse-grained manner (the entire unit must be idle to be clock gated). For some macros (in particular control logic) we do not apply any clock gating; this corresponds to about 20-25% of the unconstrained power dissipation. The overall savings due to clock gating relative to the unconstrained power is roughly 40-45%.

In order to quantify the power-performance efficiency of pipelines of a given FO4-depth, and to scale the power dissipation from the power models of our base FO4 design point across a range of FO4-depths, we have extended the PowerTimer methodology as discussed below.

Power dissipated by a processor consists of dynamic and leakage components, $P = P_{\text{dynamic}} + P_{\text{leakage}}$. The dynamic power data measured by PowerTimer (for a particular design point) can be expressed as $P_{\text{dynamic}}^{\text{base}} = CV^2 f *$ $(\alpha + \beta) * CGF$, where α is the average "true" switching factor in circuits; i.e., α represents transitions required for the functionality of the circuit and is measured as the switching factor by an RTL-level simulator run under the zero-delay mode. In contrast, β is the average glitching factor that accounts for spurious transitions in circuits due to race conditions. Thus, $(\alpha + \beta)$ is the average number of transitions actually seen inside circuits. Both α and β are averaged over the whole processor over non-gated cycles with appropriate energy weights (the higher the capacitance at a particular node the higher the corresponding energy weight). CGF is the clock gating factor which is defined as the fraction of cycles where the microarchitectural structures are *not* clock gated. The CGF is measured from our PowerTimer runs at each FO4 design point as described above. The remaining terms C, V, and f are effective switching capacitance, chip supply voltage, and clock frequency, respectively.

Next we analyze how each of these factors scales with FO4 pipeline depth. To facilitate the explanation we define the following variables: $FO4_{\text{logic}}$, $FO4_{\text{latch}}$ and $FO4_{\text{pipeline}}$, to designate the depth of the critical path through logic in one pipeline stage, the latch insertion delay including clock skew and jitter, and the sum of the two quantities, respectively, $FO4_{\text{pipeline}} = FO4_{\text{logic}} + FO4_{\text{latch}}$. We use FO4 and $FO4_{\text{pipeline}}$ interchangeably. In the remainder of the pa-

per, the qualifier 'base' in all quantities designates the value of the quantities measured for the base 19 FO4 design.

Frequency: FreqScale is the scaling factor that is used to account for the changes in the clock frequency with the pipeline depth. This factor applies to both hold power and switching power:

$$\text{FreqScale} = FO4_{\text{pipeline}}^{\text{base}} / FO4_{\text{pipeline}}$$

Latch: With fixed logic hardware for given logic functions, the primary change in the chip effective switching capacitance C with pipeline depth is due to changes in the latch count with the depth of the pipeline. LatchScale is a factor that appropriately adjusts the hold power dissipation, but does not affect the switching power dissipation.

$$\text{LatchScale} = \text{LatchRatio} * \left(\frac{FO4_{\text{logic}}^{\text{base}}}{FO4_{\text{logic}}} \right)^{\text{LGF}}$$

where *LatchRatio* defines the ratio of hold power to the total power and $LatchGrowthFactor(LGF)$ captures the growth of latch count due to the logic shape functions. The amount of additional power that is spent in latches and clock in a more deeply pipelined design depends on the *logic shape functions* of the structures that are being pipelined. Logic shape functions describe the number of latches that would need to be inserted at any cut point in a piece of combinatorial logic if it were to be pipelined. Values of $LGF > 1$ recognize the fact that for certain hardware structures the logic shape functions are not flat and hence the number of latches in the more deeply pipelined design increases super-linearly with pipeline depth. In our baseline model, we assume a LGF of 1.1 and study the sensitivity of the optimal pipeline depth to this parameter in Section 6.1.

Clock Gate Factor: In general, CGF decreases with deeper pipelines because the amount of clock gating potential increases with deeper pipes. This increased clock gating potential is primarily due to the increased number of cycles where pipeline stages are in stall conditions. This in turn leads to an increase in the clock gating potential on a per cycle basis. CGF is workload dependent and is measured directly from simulator runs.

Glitch: The final two factors that must be considered for dynamic power dissipation when migrating to deeper pipelines are α and β, the chip-wide activity and glitching factors.

The "true" switching factor α does not depend on the pipeline depth, since it is determined by the functionality of the circuits. The glitching factor at any net, on the other hand, is determined by the difference in delay of paths from the output of a latch that feeds the circuit to the gate that drives that net. Once a glitch is generated at some net, there

337

is a high probability that it will propagate down the circuit until the input of the next latch down the pipeline. Furthermore, the farther the distance from the latch output to the inputs of a gate the higher the probability of the existence of non-equal paths from the output of the latch to the inputs of this gate. Therefore the average number of spurious transitions grows with $FO4_{logic}$ – the higher the FO4 the higher the average glitching factor. Experimental data, collected by running a dynamic circuit-level simulator (PowerMill) on post-layout extracted netlists of sample functional units show that the average glitching factor β can be modeled as being linearly dependent on the logic depth:

$$\beta = \beta_{base} \frac{FO4_{logic}}{FO4_{logic}^{base}}$$

To account for the effect of the dependence of β on pipeline depth, we introduce the following factor which applies only to the switching power:

$$
\begin{aligned}
GlitchScale &= (1 - LatchRatio)\left(\frac{\alpha + \beta}{\alpha + \beta_{base}}\right) \\
&= \frac{1 - LatchRatio}{1 + \beta_{base}/\alpha}\left(1 + \frac{\beta_{base}}{\alpha}\frac{FO4_{logic}}{FO4_{logic}^{base}}\right)
\end{aligned}
$$

In this formula β_{base} is the actual glitching factor averaged over the baseline microprocessor for the base FO4 design point. Notice that β_{base} appears in the formula only in the ratio β_{base}/α. This is consistent with our experimental results showing that the glitching factor β is roughly proportional to the "true" switching factor α, for the range $0 < \alpha < 0.3$ (for higher values of α the growth of β typically saturates). For the set of six sample units that we simulated, with the logic depth ranging from 6 FO4 to 20 FO4, the ratio β/α was found to be roughly proportional to the logic depth of the simulated units, $FO4_{logic}$, with the coefficient equal to $0.3/FO4_{logic}^{base}$. Based on these simulation results we set $\beta_{base}/\alpha = 0.3$ for the whole microprocessor in the remainder of this section, and study the sensitivity of the results to variations in the glitching factor in Section 6.4.

Leakage Currents: As the technology feature size scales down and the power supply and transistor threshold voltages scale accordingly, the leakage power component becomes more and more significant. Since the magnitude of the leakage power component is affected by the pipeline depth, it is essential to include the effect of the leakage power in the analysis of the optimal pipeline depth. Assuming that the leakage power is proportional to the total channel width of all transistors in the microprocessor, we model the dependence of the leakage power on the depth of the pipeline as follows:

$$P_{leakage}^{FO4} = P_{leakage}^{base}\left(1 + \frac{w_{latch}}{w_{total}}\left[\left(\frac{FO4_{logic}^{base}}{FO4_{logic}}\right)^{LGF} - 1\right]\right)$$

where LGF is the $LatchGrowthFactor$ defined earlier, w_{latch}/w_{total} is the ratio of the total channel width of transistors in all pipeline latches (including local clock distribution circuitry) to the total transistor channel width in the base (19 FO4) microprocessor (excluding all low-leakage transistors that might be used in caches or other on-chip memories). If the technology supports multiple thresholds, or any of the recently introduced leakage reduction techniques are used on a unit-by-unit basis, such as MTCMOS, back biasing, power-down, or transistor stacking, then the above formula for the leakage power component needs to be modified accordingly and we leave the detailed study of these effects for future work. For the remainder of the study we set w_{latch}/w_{total} to 0.2 for the base 19 FO4 pipeline. Also, rather than giving the absolute value for the leakage current, $P_{leakage}^{base}$ in the base microprocessor, we will count it as a fraction of the dynamic power of the base design, $P_{leakage}^{base} = LeakageFactor^{base}P_{dynamic}^{base}$. We set the $LeakageFactor^{base}$ to the value of 0.1, typically quoted for state of the art microprocessors, and analyze the sensitivity of the results to $LeakageFactor$ in Section 6.5.

Total Power: The following equation expresses the relationship between the dynamic power for the base FO4 design, $P_{dynamic}^{base}$, leakage power and the scaled power for designs with different depths of the pipeline, P_{total}^{FO4}, considering all factors above.

$$P_{total}^{FO4} = CGF*FS*(LS+GS)*P_{dynamic}^{base} + P_{leakage}^{FO4}$$

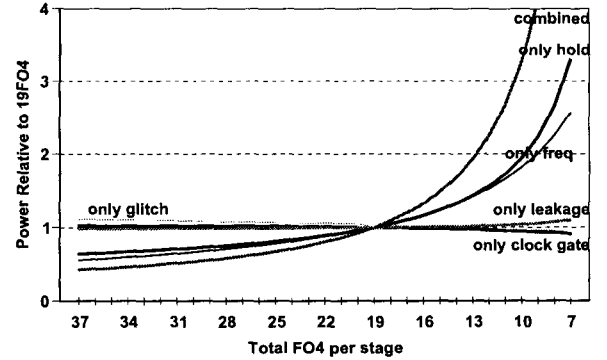where FS is $FreqScale$, LS is $LatchScale$, and GS is $GlitchScale$.

Figure 4. Power Growth Breakdown

Figure 4 shows contributions of different factors in the above formula, depending on the pipeline depth of the design $FO4_{pipeline}$. The 19 FO4 design was chosen as a base pipeline.

The line labeled "combined" shows the cumulative increase or decrease in power dissipation. The line labeled "only clock gate" quantifies the amount of additional clock gating power savings for deeper pipelines. The relative effect of scaling in clock gating is fairly minor with slightly

more than 10% additional power reduction when going from the 19 FO4 to 7 FO4 design points. There are several reasons why the effect of clock gating is not larger. First, the fraction of power dissipation that is not eligible to be clock gated becomes larger with more clock gating leading to diminishing returns. Second, some of the structures are clock gated in a coarse-grained fashion and while the average utilization of the structure may decrease it must become idle in all stages before any additional savings can be realized. Finally, we observe that clock gating is more difficult in deeper pipelined machines, because it is harder to deliver cycle-accurate gating signals at lower FO4.

The two lines labeled "only freq" and "only hold" show the power factors due to only frequency and hold power scaling, respectively.[2] Overall, dynamic power increases more than quadratically with increased pipeline depth.

Figure 4 shows that the leakage component grows much less rapidly than the dynamic component with the increasing pipeline depth. There are two primary reasons for this. First, the leakage power does not scale with frequency. Second, the leakage power growth is proportional to the fraction of channel width of transistors in pipeline latches, whereas the latch dynamic hold power growth is proportional to the fraction of the dynamic power dissipated in pipeline latches. Obviously, the former quantity is much smaller than the latter.

4.3 Workloads and Metrics Used in the Study

In this paper, we report experimental results based on PowerPC traces of a set of 21 SPEC2000 benchmarks, namely, *ammp, applu, apsi, art, bzip2, crafty, equake, facerec, gap, gcc, gzip, lucas, mcf, mesa, mgrid, perl, sixtrack, swim, twolf, vpr,* and *wupwise*. We have also used a 172M instruction trace of the TPC-C transaction processing benchmark. The SPEC2000 traces were generated using the tracing facility called *Aria* within the MET toolkit [17]. The particular SPEC2000 trace repository used in this study was created by using the full reference input set. However, sampling was used to reduce the total trace length to 100 million instructions per benchmark program. A systematic validation study to compare the sampled traces against the full traces was done in finalizing the choice of exact sampling parameters [12].

We use BIPS^3/W (*energy * delay^2*) as a basic energy-efficiency metric for comparing different FO4 designs in the power-performance space. The choice of this metric is based on the observation that dynamic power is roughly proportional to the square of supply voltage (V) multiplied by clock frequency and clock frequency is roughly proportional to V. Hence, power is roughly proportional to V^3 assuming a fixed logic/circuit design. Thus, delay cubed multiplied by power provides a voltage-invariant power-performance characterization metric which we feel is most

---

[2]Although these two factors increase linearly with clock frequency, we are plotting against FO4-depth which is $1/clock\ frequency$.

appropriate for server-class microprocessors (see discussion in [1, 7]). In fact, it was formally shown in [26] that optimizing performance subject to a constant power constraint leads to the BIPS^3/W metric in processors operating at a supply voltage near the maximum allowed value in state of the art CMOS technologies. As a comparative measure, we also consider BIPS/W (energy) and BIPS^2/W (energy-delay [8]) in the baseline power-performance study.

5 Analytical Model and Simulation Results

5.1 Analytical Model Validation

Figure 5. BIPS: Analytical Model vs. Simulator

We now present the performance results using the simple analytical model (see Section 3) and compare it with the results using our detailed cycle-accurate simulator. The results in this section are presented for the average of the SPEC2000 benchmarks described in Section 4.3. For fair comparison, we have modified the simulator to include a perfect front-end of the machine. The model and the simulator use latencies shown in the column labeled "INF" in Table 1.

Figure 5 shows BIPS as a function of the FO4 delay per stage of the pipeline. The BIPS for the analytical model was computed after determining the stalls for the different pipelines using independent workload analysis as explained in Section 3.

From Figure 5 we observe that the performance optimal pipeline depth is roughly 10 FO4 delay per pipeline stage for both the model and the simulator. The BIPS estimated using the model correlates well with the BIPS estimated using the simulator except for very large and very small FO4-depth machines. Since the stalls are determined in the model independent of the pipeline and only once for a workload, it is possible that for shallow pipelines with large FO4 delay per stage, the model underestimates the total stall cycles, and hence the BIPS computed by the model is higher than the BIPS obtained from the simulator. The analytical model currently only uses data dependent stalls to derive the optimal pipeline depth. Hence, for the purpose of the validation experiment, the model and the simulator assume a perfect front-end. However, to estimate the effect of resource stalls we modeled a large but finite register renamer,

339

instruction buffer, and miss queue in the simulator keeping the rest of the front-end resources infinite and perfect. For deeper pipelines ($\leq$ 10 FO4), we observe from the simulator that the stalls due to the resources become appreciable and the analytical model begins to overestimate the BIPS relative to the simulator.


Figure 6. BIPS³/W: Model vs. Simulator

Figure 6 shows that the optimal FO4-depth that maximizes the BIPS³/W is around 19-24 FO4 for the simulator. We observe that the analytical model tracks the simulator results more closely while optimizing performance alone as seen in Figure 5. Since the analytical model overestimates the performance for shallow pipelines, the cubing of BIPS in Figure 6 compounds these errors.

5.2 Detailed Power-Performance Simulation

Figure 7. Simulation Results for SPEC2000

In the remainder of this work we consider the power and performance results using the detailed performance simulator with parameters corresponding to the STD column in Table 1. Figure 7 shows the results for five metrics: BIPS, IPC, BIPS/W, BIPS²/W, and BIPS³/W.

Figure 7 shows that the optimal FO4-depth for performance (defined by BIPS) is 10 FO4, although pipelines of 8 FO4 to 15 FO4 are within 5% of the optimal. Because of the super-linear increase in power dissipation and sublinear increases in overall performance, the BIPS/W always decreases with deeper pipelines. BIPS³/W shows an optimum point at 18 FO4. BIPS³/W decreases sharply after the optimum and at the performance-optimal pipeline depth of 10 FO4, the BIPS³/W metric is reduced by 50% over

the 18 FO4 depth. For metrics that have less emphasis on performance, such as BIPS²/W and BIPS/W, the optimal point shifts towards shallower pipelines, as expected. For the BIPS²/W, the optimal pipeline depth is achieved at 23 FO4.


Figure 8. Simulation Results for TPC-C

Figure 8 presents similar results for the TPC-C trace. The optimal BIPS for TPC-C is very flat from 10-14 FO4 (within 1-2% for all 5 design points which is much flatter than SPEC2000). Using BIPS³/W, the optimal pipeline depth shifts over to 25-28 FO4. The main reason that the optimal point is shallower for TPC-C is that BIPS decreases less dramatically with decrease in pipeline length (relative to SPEC2000). This is slightly counterbalanced because power increases at a slower rate for TPC-C with deeper pipes, because the additional amount of clock gating is more pronounced due to large increases in the number of stall cycles relative to the SPEC2000 suite.

6 Sensitivity Analysis

The derived equations that model the dependence of the power dissipation on the pipeline depth depend on several parameters. Some of these parameters, although accurately measured for the baseline microprocessor, are likely to change from one design to another, whereas others are difficult to measure accurately. In this section, we perform sensitivity analysis of the optimal pipeline depth to key parameters of the derived power models such as $LatchGrowthFactor$, $LatchRatio$, latch insertion delay ($FO4_{latch}$), $GlitchRatio$ and $LeakageFactor$.

6.1 Latch Growth Factor

$LatchGrowthFactor$ (LGF) is determined by the intrinsic logic shape functions of the structures that are being pipelined. We have analyzed many of the major microarchitectural structures to identify ones that are likely to have $LatchGrowthFactor$ greater than 1. One structure that we will highlight is the Booth recoder and Wallace tree which is common in high-performance floating point multipliers [13, 20], as shown in Figure 9. Figure 9 shows the exponential reduction in the number of result bits as the data passes through the 3-2 and 4-2 compressors. We have estimated

340

the amount of logic that can be inserted between latch cut points for 7, 10, 13, 16, and 19 FO4 designs by assuming 3-4 FO4 delay for 3-2 and 4-2 compressor blocks. As the 7 and 10 FO4 design points require latch insertions just after the Booth multiplexor (where there are 27 partial products), there would be a large increase in the number of latches required for these designs. We note that the 7 FO4 design also requires a latch insertion after the booth recode stage.
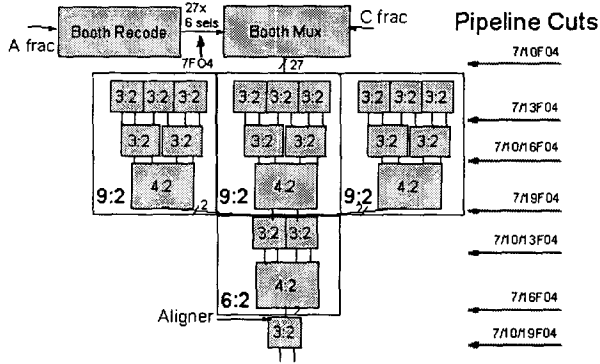
**Figure 9. Wallace Tree Diagram and Latch Cut points for 7/10/13/16/19 FO4**

Figure 10 gives estimates for the cumulative number of latches in the FPU design as a function of the FO4 depth of the FPU. For example, the first stage of the 10 FO4 FPU requires 3x as many latches as the first stage of the 19 FO4 FPU because the first latch cut point of the 19 FO4 FPU is beyond the initial 9:2 compressor tree. Overall, the 10 FO4 FPU requires nearly 3x more latches than the 19 FO4 FPU.
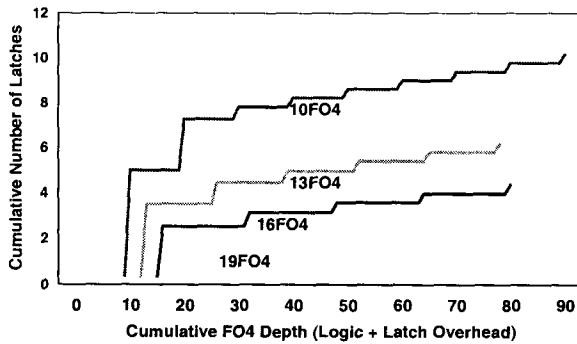
Figure 10. Cumulative latch count for FPU

There are many other areas of parallel logic that are likely to see super-linear increases in the number of latches such as structures with decoders, priority encoders, carry look ahead logic, etc. Beyond the pipelining of logic structures, deeper pipelines may require more pre-decode information to meet aggressive cycle times, which would require more bits to be latched in intermediate stages. On the other hand, the number of latches that comprise storage bits in various on-chip memory arrays (such as register files and

queues) does not grow at all with the pipeline depth, meaning that the $LGF = 0$ for those latches. Thus, designs with overall $LGF < 1$ are also possible.
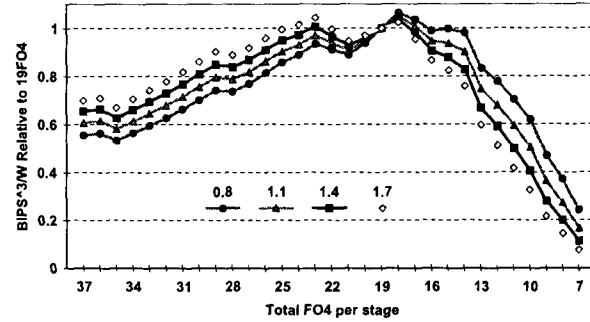

Figure 11. BIPS^3/W varying $LatchGrowthFactor$

Figure 11 quantifies the dependence of the optimal pipeline depth on LGF using the BIPS^3/W metric. It shows that the optimal pipeline FO4 tends to increase as LGF increases above the value of 1.1, assumed in the baseline model. As a point of reference, our estimate for the latch growth factor for the 10 FO4 vs. 19 FO4 Booth recoder and Wallace tree is $LGF = 1.9$, while for the entire FPU LGF is slightly less than 1.7.

6.2 Latch Power Ratio

Latch, clock, and array power are the primary components of power dissipation in current generation CPUs. This is especially true in high-performance, superscalar processors with speculative execution which require the CPU to maintain an enormous amount of architectural and non-architectural state. One possible reason why the $LatchRatio$ could be smaller than the base value of 0.7 chosen in Section 4, is if more energy-efficient SRAM arrays are used in high-power memory structures instead of latches to reduce the data independent array power (which we include as part of hold power).

Figure 12. BIPS^3/W varying $LatchRatio$

Figure 12 shows the optimal FO4 design point while varying the $LatchRatio$ of the machine from 80% to 40%. We see that while the optimal FO4 point remains 18 FO4, it is less prohibitive to move to deeper pipelines with smaller

latch-to-logic ratios. For example, with a *LatchRatio* of 0.4, the 13 FO4 design point is only 19% worse than the optimal one while it is 27% worse than optimal with a *LatchRatio* of 0.6.

6.3 Latch Insertion Delay

| Latch FO4	2	3	4	5
Relative Latch Energy	1.0	0.53	0.36	0.29
Relative Clocking Energy	1.0	0.62	0.49	0.43

Table 2. Latch Insertion Delay (excluding skew and jitter) vs. Relative Latch Energy

With the large amount of power spent in latches and clocking, designers may consider the tradeoff between latch delay and power dissipation as a means to design more energy-efficient CPUs. Researchers have investigated latch power vs. delay tradeoff curves both within a given latch family and across latch families [25, 10]. Table 2, derived from [25], shows latch FO4-delay vs. latch energy across several latch styles. The first row of Table 2 shows the latch insertion delay, excluding clock skew and jitter overhead which is assumed to be constant for all latches. The second row shows the relative latch energy, excluding energy dissipated in the clock distribution. The third row shows the relative energy of the clock system, including both energy of latches (70% of the total) and clock distribution system (30%). It is assumed that the clock distribution energy cannot be completely scaled with reduced latch load. There is significant overhead simply from driving the wires necessary to distribute the clock over a large area.

Figure 13. BIPS varying latch power-delay

Replacing the baseline fast 2 FO4 latches with slower, lower power latches increases the latch insertion delay overhead, which impacts both the performance-optimal and power-performance-optimal pipeline depth. Figure 13 shows the processor performance versus pipeline depth for four latches from Table 2. We see that the performance maxima shift towards shallower pipelines as the latch insertion delay increases. For example, with a 3 FO4 latch the performance-optimal FO4-depth is 11 FO4 and with a 4 FO4 latch it becomes 16 FO4. ·
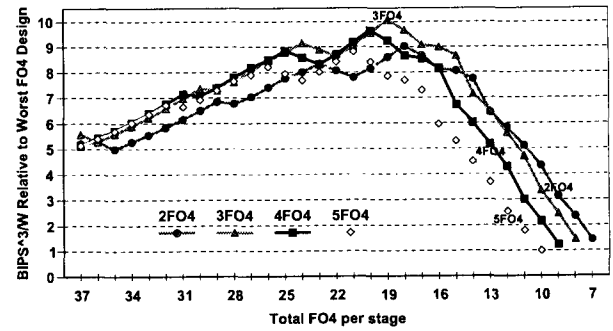
Figure 14. BIPS³/W varying latch power-delay

Figure 14 shows the impact of the latch insertion delay on the BIPS³/W rating of the processor for the same range of pipeline depths. In this figure, all of the data points are shown relative to the 10 FO4 design with the base 5 FO4 latch. Unlike curves on all previous sensitivity graphs, curves in Figure 14 do not intersect at the base design point, because different curves represent different designs, with different power and performance levels.

Figure 14 shows that using the fastest 2 FO4 latch results in the best BIPS³/W rating for processors with stages less than 14 FO4. For processors with pipelines ranging from 15 FO4 to 24 FO4 a lower power 3 FO4 latch is the most energy efficient, whereas for shallower pipelines (25 FO4 or more) the highest BIPS³/W is achieved with even slower 4 FO4 latches.

The use of the 3 FO4 latch, combined with the choice of the pipeline depth in the range from 15 FO4 to 24 FO4 improves the BIPS³/W rating of the processor by more than 10%, compared to the base case of 2 FO4 latches. The graph also shows that the optimal BIPS³/W design point shifts towards shallower pipelines as high-performance latches are replaced with lower power ones. For example, the 18 FO4 design point is optimal for a processor using 2 FO4 latches, whereas the 19 FO4, 20 FO4, and 21 FO4 design points are optimal for processors using 3 FO4 latches, 4 FO4, and 5 FO4 latches, respectively.

6.4 Glitch Factor

In this subsection we quantify the sensitivity of the optimal pipeline depth to the glitching factor. There are no practical means for accurately measuring the actual value of β_{base}/α, averaged over the whole microprocessor. Instead we measured the glitching factor for a selected set of functional units and used the averaged value of $\beta_{base}/\alpha = 0.3$ throughout the analysis. In this section we analyze the sensitivity of the optimal pipeline depth to the value of β_{base}/α.

Figure 15 shows the dependence of the BIPS³/W rating of the processor on the pipeline depth for three values of β_{base}/α. From this figure we see that higher glitching factors favor deeper pipelines. However, in the base design the decrease in power dissipation related to reduced glitching in deeper pipelines did not have a substantial impact on the
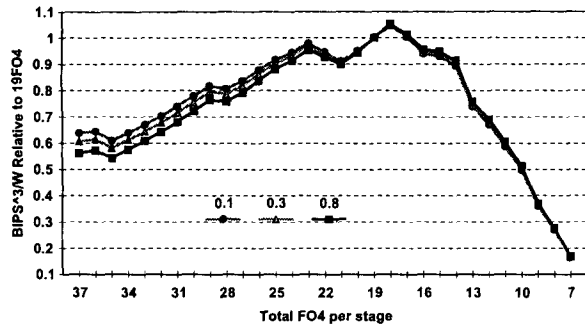
Figure 15. BIPS^3/W varying $\beta_{\text{base}}/\alpha$

optimal pipeline depth, primarily because of the relatively small fraction of power dissipated in combinatorial switching. For designs which have smaller *LatchRatio* values, this effect could be more significant.

6.5 Leakage Factor

As explained earlier, the leakage power component grows more slowly with the pipeline depth than the dynamic component. Therefore, the optimum pipeline depth depends on the *LeakageFactor*. Throughout the analysis we assumed that for the base 19 FO4 microprocessor the *LeakageFactor* ($P^{\text{base}}_{\text{leakage}}/P^{\text{base}}_{\text{dynamic}}$) to be 0.1. However, as the technology feature size scales down and the power supply and transistor threshold voltages scale accordingly, the leakage power component becomes more and more significant. To study the effect of the growing fraction of the leakage power component we measured the sensitivity of the optimal pipeline depth to the value of the *LeakageFactor*.


Figure 16. BIPS^3/W varying *LeakageFactor*

Figure 16 shows the BIPS^3/W rating of the processor versus pipeline depth for three values of the *LeakageFactor*: a value of 0 that represents older CMOS technologies, a value of 0.1, assumed in the current model, and values of 0.5 and 1.0, projected for future generation CMOS technologies (arguably, extreme values). The results in Figure 16 show that unless leakage reduction techniques become the standard practice in the design of high-end microprocessors, the high values of the *LeakageFactor* projected for future generations of CMOS technologies

may tend to shift the optimum pipeline depth towards slightly deeper pipelines. For current generation technologies the result for the optimal pipeline depth is sufficiently stable with respect to reasonable variations in the *LeakageFactor*.

Summary of the Sensitivity Analysis: In this section we considered the sensitivity of optimal pipeline length to five key parameters in the power models using the BIPS^3/W metric. We did not observe a strong dependence of the results on the assumptions and choices of any of these parameters, which demonstrates the stability of the model, and its applicability to a wide range of designs. To summarize the results, higher values of the *LatchGrowthFactor* favor shallower pipelines, lower values of the *LatchRatio* favor deeper pipelines, the use of lower-power latches favors shallower pipelines, higher values of the *GlitchFactor* favors deeper pipelines, and, finally, higher leakage currents favor deeper pipelines.

7 Conclusions

In this paper, we have demonstrated that it is important to consider both power and performance while optimizing pipelines. For this purpose, we derived detailed energy models using circuit-extracted power analysis for microarchitectural structures. We also developed detailed equations for how the energy functions scale with pipeline depth. Based on the combination of power and performance modeling performed, our results show that a purely performance-driven, power-unaware design may lead to the selection of an overly deep pipelined microprocessor operating at an inherently power-inefficient design point.

As this work is the first quantitative evaluation of power and performance optimal pipelines, we also performed a detailed sensitivity analysis of the optimal pipeline depth against key parameters such as latch growth factor, latch ratio, latch insertion delay, glitch, and leakage currents. Our analysis shows that there is a range of pipeline depth for which performance increases can be achieved at a modest sacrifice in power-performance efficiency. Pipelining beyond that range leads to drastic reduction in power-performance efficiency with little or no further performance improvement.

Our results show that for a current generation, out-of-order superscalar processor, the optimal delay per stage is about 18 FO4 (consisting of a logic delay of 15 FO4 and 3 FO4 latch insertion delay) when the objective function is a power-performance efficiency metric like BIPS^3/W; this is in contrast to an optimal delay of 10 FO4/stage when considering the BIPS metric alone. We used a broad suite of SPEC2000 benchmarks to arrive at this conclusion.

The optimal pipeline depth depends on a number of parameters in the power models which we have derived from current state-of-the-art microprocessor design methodologies. Also, as already established through recent prior work,

such optimal design points generally depend on the input workload characteristics. Our simulation-based experiments on a typical commercial application (TPC-C) shows that although the optimal pipeline depth is around 10-14 FO4 for performance-only optimization, it increases to 24-28 FO4 when we consider power and performance optimizations.

In future work, we would like to consider other architectural options than those available in the base model, for example, single- versus multi-threading, wide versus narrow issue machines, and inorder versus out-of-order designs. In addition, we would also like to consider more circuit-level power-performance tradeoffs opportunities and factor that into the analysis of optimal pipelines. Finally, there are numerous circuit and technology techniques which affect leakage and therefore power-performance optimality; we hope to consider all of these factors in our future work.

References

[1] D. Brooks et al. Power-aware Microarchitecture: Design and Modeling Challenges for the next-generation microprocessors. *IEEE Micro*, 20(6):26–44, Nov./Dec. 2000.

[2] D. Brooks, V. Tiwari, and M. Martonosi. Wattch: A framework for architectural-level power analysis and optimizations. In *Proceedings of the 27th Annual International Symposium on Computer Architecture (ISCA-27)*, June 2000.

[3] D. Brooks, J.-D. Wellman, P. Bose, and M. Martonosi. Power-Performance Modeling and Tradeoff Analysis for a High-End Microprocessor. In *Power Aware Computing Systems Workshop at ASPLOS-IX*, Nov. 2000.

[4] M. Brown, J. Stark, and Y. Patt. Select-free instruction scheduling logic. In *Proceedings of the 34th International Symposium on Microarchitecture (MICRO-34)*, pages 204–213, December 2001.

[5] P. Dubey and M. Flynn. Optimal pipelining. *J. Parallel and Distributed Computing*, 8:10–19, 1990.

[6] P. G. Emma and E. S. Davidson. Characterization of branch and data dependencies in programs for evaluating pipeline performance. *IEEE Transactions on Computers*, C-36(7):859–875, 1987.

[7] M. J. Flynn, P. Hung, and K. Rudd. Deep-Submicron Microprocessor Design Issues. *IEEE Micro*, 19(4):11–22, July/Aug. 1999.

[8] R. Gonzalez and M. Horowitz. Energy dissipation in general purpose microprocessors. *IEEE Journal of Solid-State Circuits*, 31(9):1277–84, Sept. 1996.

[9] A. Hartstein and T. R. Puzak. The optimum pipeline depth for a microprocessor. In *Proceedings of the 29th International Symposium on Computer Architecture (ISCA-29)*, May 2002.

[10] S. Heo, R. Krashinsky, and K. Asanovic. Activity-sensitive flip-flop and latch selection for reduce energy. In *19th Conference on Advanced Research in VLSI*, March 2001.

[11] M. Hrishikesh, K. Farkas, N. Jouppi, D. Burger, S. Keckler, and P. Sivakumar. The optimal logic depth per pipeline stage is 6 to 8 FO4 inverter delays. In *Proceedings of the 29th International Symposium on Computer Architecture (ISCA-29)*, pages 14–24, May 2002.

[12] V. Iyengar, L. H. Trevillyan, and P. Bose. Representative traces for processor models with infinite cache. In *Proc. 2nd. Symposium on High Performance Computer Architecture (HPCA-2)*, Feb. 1996.

[13] R. Jessani and C. Olson. The floating-point unit of the PowerPC 603e microprocessor. *IBM J. of Research and Development*, 40(5):559–566, Sept. 1996.

[14] P. Kogge. *The Architecture of Pipelined Computers*. Hemisphere Publishing Corporation, 1981.

[15] S. R. Kunkel and J. E. Smith. Optimal pipelining in supercomputers. In *Proceedings of the 13th International Symposium on Computer Architecture (ISCA-13)*, pages 404–411, June 1986.

[16] M. Moudgill, P. Bose, and J. Moreno. Validation of Turandot, a fast processor model for microarchitecture exploration. In *Proceedings of the IEEE International Performance, Computing, and Communications Conference (IPCCC)*, pages 451–457, Feb. 1999.

[17] M. Moudgill, J. Wellman, and J. Moreno. Environment for PowerPC microarchitecture exploration. *IEEE Micro*, 19(3):9–14, May/June 1999.

[18] J. S. Neely, H. H. Chen, S. G. Walker, J. Venuto, and T. Bucelot. CPAM: A common power analysis methodology for high-performance VLSI design. In *Proc. of the 9th Topical Meeting on the Electrical Performance of Electronic Packaging*, pages 303–306, 2000.

[19] S. Palacharla, N. Jouppi, and J. Smith. Complexity-Effective Superscalar Processors. In *Proceedings of the 24th International Symposium on Computer Architecture (ISCA-24)*, 1997.

[20] P. Song and G. D. Micheli. Circuit and architecture tradeoffs for high-speed multiplication. *IEEE Journal of Solid-State Circuits*, 26(9):1184–1198, Sept. 1991.

[21] E. Sprangle and D. Carmean. Increasing processor performance by implementing deeper pipelines. In *Proceedings of the 29th International Symposium on Computer Architecture (ISCA-29)*, May 2002.

[22] J. Stark, M. Brown, and Y. Patt. On pipelining dynamic instruction scheduling logic. In *Proceedings of the 33rd International Symposium on Microarchitecture (MICRO-33)*, pages 57–66, Dec. 2000.

[23] N. Vijaykrishnan, M. Kandemir, M. Irwin, H. Kim, and W. Ye. Energy-driven integrated hardware-software optimizations using SimplePower. In *Proceedings of the 27th Annual International Symposium on Computer Architecture*, June 2000.

[24] V. Zyuban. *Inherently Lower Power High Performance Superscalar Architectures*. PhD thesis, University of Notre Dame, March 2000.

[25] V. Zyuban and D. Meltzer. Clocking strategies and scannable latches for low power applications. In *Proc. of Int'l Symposium on Low-Power Electronics and Design*, 2001.

[26] V. Zyuban and P. Strenski. Unified Methodology for Resolving Power-Performance Tradeoffs of the Microarchitectural and Circuit Levels. In *Proc. of Int'l Symposium on Low-Power Electronics and Design*, pages 166–171, 2002.