Appears in the International Symposium on Low Power Electronics and Design 2009
San Francisco, California, August, 2009

Dynamic Power Gating with Quality Guarantees

Anita Lungu[1], Pradip Bose[2], Alper Buyuktosunoglu[2], and Daniel J. Sorin[3]

[1]Dept. of Computer Science
Duke University

[2]IBM T. J. Watson Research Center

[3]Dept. of Electrical and Compter Engineering
Duke University

ABSTRACT
Power gating is usually driven by a predictive control, and frequent mispredictions can counter-productively lead to a large increase in energy consumption. This energy vulnerability could be exploited by malicious applications such as a power virus, or it may be exposed by regular applications containing repetitive mispredictions patterns. We propose counteracting this vulnerability by using a guard mechanism to prevent power overruns.

Categories and Subject Descriptors
C.1.3 [**Processor Architectures**]: Other Architecture Styles - Pipeline processors

General Terms
Design, Performance, Experimentation

Keywords
power gating, power management, execution units, low power, microarchitecture

1. INTRODUCTION
One of the most promising mechanisms for reducing leakage energy is power gating, whereby leakage energy is saved by cutting the supply voltage to idle circuits. In prior work, Hu et. al [2] demonstrated how to effectively use microarchitectural information to predict when to power gate a microprocessor's units, such as its functional units. Recently, Intel has announced the adoption of power-gating in Nehalem [4], although the focus is core-level gating, and the exact dynamic algorithms have not been published.

Power gating can potentially save a significant amount of leakage energy but, because it is predictive, *it can also lead to significant energy penalties when mispredictions are frequent*. Every time a component is power gated off, some energy is needed for turning the component off and then on again. If the power gated unit needs to be brought back in use soon after being turned off, that energy cost is not offset by the leakage energy saved during power gating.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
ISLPED '09, August 19-21, 2009, San Francisco, California, USA.
Copyright 2009 ACM 978-1-60558-684-7/09/08...$10.00.

We show in this paper that the impact of power gating depends on the workload running on the processor, and we observe that both big wins *and big losses* can occur. Across the SPEC2006 benchmarks, we observe a large variability in the leakage energy savings of a well-known power gating algorithm [2] applied to functional units. On a given benchmark, power gating can save as much as 99% of a functional unit's leakage energy during some time frames, whereas during other time frames it can consume 70% *more* energy than a system without power gating.

When adding a feature such as power gating to a processor, we would like for it to both benefit the common case and do no harm in less common cases. To assess the common case benefit (e.g., average energy savings), it is generally sufficient to simulate the system across a wide range of benchmarks. To assess the potential for harm (e.g., consuming 70% extra energy), this simulation approach is often insufficient, because *a harmful situation may be unacceptable even if it is unlikely.* Consider, for example, a power virus application designed to attack the system by exploiting this vulnerability and increasing its energy consumption. When deciding whether to add a feature to a real design, we must consider the feature's behavior for applications that have not been simulated. Such an unsimulated behavior could be, for example, a scientific application running in a loop that consistently wastes a significant amount of energy through the addition of the power gating scheme. This consistent waste of energy is a design vulnerability of the power gating scheme that might prevent its adoption in a real processor.

When a proposed feature such as power gating can have both a large benefit and a large penalty, we would like to augment it with two mechanisms. First, we want a mechanism that enables the feature when that is likely to translate into a win and disables it when the result is likely to be a loss. Second, we want a mechanism that provides *a guaranteed bound on the worst-case behavior* of the feature. We propose adding both of these mechanisms. Specifically:

- We augment current dynamic power gating schemes with a Success Monitor that dynamically estimates whether a particular power gating scheme is successful at saving energy or not.
- Based on the Success Monitor, we add a feedback control mechanism, called the Success Monitor Switch, that enables/disables power gating depending on its success. For benchmarks initially exhibiting energy loss, the Success Monitor Switch saves energy and decreases performance penalties.
- We implement a Token Counting Guard mechanism that provably bounds the worst-case energy penalty to an acceptable threshold (2% of leakage energy) for a specified time interval.
- We combine the Success Monitor Switch with the Token Counting Guard to bound the energy loss while slightly increasing the average energy savings (by 2%) and decreasing the performance penalty by 63%.

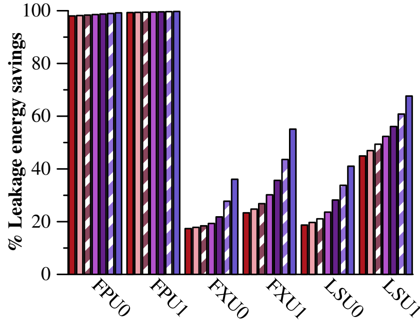

Figure 1. Best-case potential to power gate (SPECint)

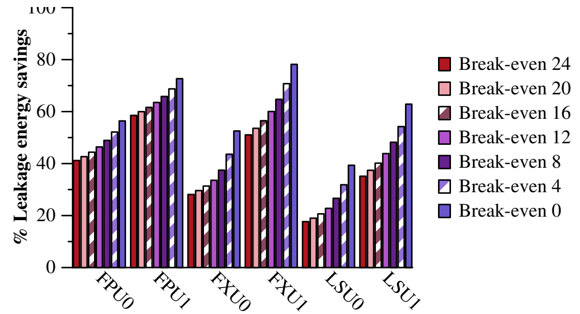

Figure 2. Best-case potential to power gate (SPECfp)

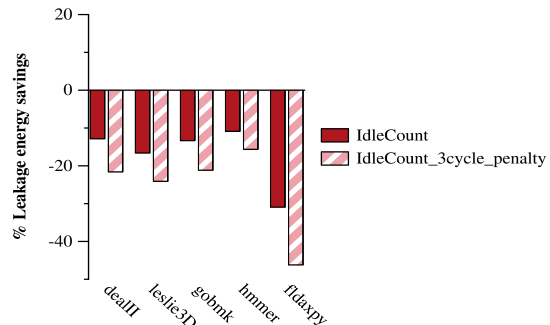

Figure 3. Benchmarks wasting energy with IdleCount

2. Power Gating: Potential and Pitfalls

Power gating requires, for each circuit that can be turned off, the presence of a header (or footer) "sleep" transistor that can set the supply voltage of the circuit to ground level (or V_{DD} level for footer) during idle times. Power gating also requires control logic to predict when would be a good time to power gate the circuit.

Every time the control logic decides to power gate the circuit, an Energy Overhead cost is incurred. This energy overhead is due to 1) distributing the sleep signal to the header transistor before the circuit is actually turned off and 2) turning off the sleep signal and driving the Virtual V_{DD} when the circuit is powered-on again. Power gating's *break-even point* is the time when the cumulative leakage energy savings equals the energy overhead. Adopting prior terminology [2], the time between the decision to power gate and when the unit has reached the break-even point is "uncompensated," and the time after the break-even point is "compensated."

Power gating can impact performance due to two different sources: the addition of the "sleep" transistor and supplementary cycles spent waiting for functional units to wake up. The addition of the "sleep" transistor can imply a small hit on cycle time. This can be avoided by a small increase in V_{DD} to counteract the voltage droop across the "sleep" transistor. The second source of performance impact can be largely mitigated by using information from the decode stage or issue queue to proactively wake up needed units.

2.1 An Example Power Gating Scheme

The baseline power gating scheme that we will augment in this paper is a microarchitectural technique for power gating functional units that was developed by Hu et al. [2]. They evaluate the potential for power gating the functional units and propose an IdleCount algorithm for deciding when to assert the power gate signal. The IdleCount algorithm counts the number of cycles a unit has been idle and decides to shut the unit off when a fixed threshold (called the *idle_detect*) has been reached. We will illustrate two of our proposed mechanisms, the Success Monitor Switch and Token Counting Guard, by adding them on top of the IdleCount algorithm. However, as we discuss in Section 3, our mechanisms are not restricted to the IdleCount algorithm.

2.2 Power Gating Potential

Workloads must have a significant amount of idleness for power gating to be effective. We show that this is indeed the case by exploring the power gating potential present for functional units in the SPEC 2006 benchmarks. The data presented was obtained using the methodology presented in detail in Section 4.

Figure 1 (SPECint) and Figure 2 (SPECfp) show the energy saving potential through power gating of each unit, as a percentage of the total leakage energy of that unit, as a function of the break-even point. Considering the analytical model developed by Hu et al. [2], the break-even point for functional units developed in current technology parameters is between 9 and 24. The energy savings values are calculated for an oracle algorithm that knows the future workload behavior and turns a unit off immediately when that decision saves energy. The oracle also wakes-up a power gated unit when required without incurring any performance penalty at wake-up.

Overall, we observe a large potential for power gating across all units. Furthermore, there is a significant increase in power gating potential associated with a decrease in the break-even point. This result is intuitive since a lower break-even point signifies that more idle intervals can be effectively used for power gating by the oracle algorithm. If circuit-level techniques become available to reduce the energy overhead associated with power gating, then the additional energy savings are significant. The large potential for saving leakage energy by power gating processor units justifies microarchitectural level techniques for harvesting it. However, power gating can have its own pitfalls which we present next.

2.3 Power Gating Pitfalls

Due to the inherently speculative characteristic of power gating algorithms, it is possible for power gating to result in significant energy penalties. This possibility represents a vulnerability of the power gating algorithm that could result in a design decision against implementing the scheme in a real processor design.

To investigate this possibility, we implemented the IdleCount algorithm described in Section 2.1. Using the methodology described in Section 4, we ran the IdleCount algorithm on the SPEC2006 benchmarks and computed the energy savings. Several of the SPEC2006

benchmarks (*dealII, leslie3d, gobmk,* and *hmmer*) exhibited an energy penalty due to the addition of the IdleCount power gating feature. We also tested a microbenchmark called *fldaxpy* that exhibits a type of behavior that is difficult for the power gating scheme to predict correctly. For these benchmarks, Figure 3 presents the percentage of leakage energy saved (with respect to a system with no power gating) by using IdleCount to power gate the FXU0 unit for a value of the idle_detect threshold of 5. Solid bars represent energy savings assuming zero performance penalty at wake-up from power gating while the striped bars assume a three cycle wake-up penalty. We see that the speculative power gating algorithm can indeed incur significant energy penalties (negative savings).

We conclude that a feature added to save power can at runtime end up costing extra energy. Moreover this type of runtime behavior can be encountered in regular benchmarks. To reduce the worst-case impact of power gating (and predictive schemes, in general), we add two mechanisms to the policy, which we discuss next.

3. MECHANISMS TO IMPROVE COMMON CASE AND BOUND WORST CASE

In this section, we describe two mechanisms that we propose adding to predictive schemes. The first is a Success Monitor that assesses the dynamic benefit of the predictive scheme. We can use the Success Monitor to make better predictions about when to power gate. By not using power gating when it is not saving energy, we decrease the performance penalty associated with power gating. Our second mechanism is a Token Counting Guard that provides a provable worst-case bound on the possible penalty associated with mispredictions. For both mechanisms, we discuss them in general but describe details for their specific application to power gating.

3.1 Improving Common Case: Success Monitor

We propose adding a Success Monitor to predictive mechanisms (e.g., power gating) whose success or loss depends on the dynamic behavior of the application. The goal of the Success Monitor is to estimate whether a particular policy is successful or harmful during a certain time interval that we call a *monitoring interval*. Based on the estimate from the Success Monitor, the control logic can better predict when to power gate (discussed in Section 3.1.2).

3.1.1 Success Monitor Structure and Behavior

The Success Monitor requires four values to dynamically estimate the success or loss of a policy, the first two of which are obtained at runtime (and called *efficiency counters*) and the latter two of which can be approximated to a constant for a given system:

- Number of successful instances per monitoring interval,
- Number of harmful instances per monitoring interval,
- Reward of a successful instance, and
- Cost of a harmful instance.

In the context of power gating, a successful instance is any compensated cycle (i.e., a cycle when a power gated unit remains idle after reaching the break-even point). We keep track of energy savings or penalties by using tokens, one token corresponding to the leakage energy used by the unit during one cycle. The reward of a successful instance is thus one token. A harmful instance is represented by any case when the unit needs to be woken up before reaching the break-even point. We pessimistically assign a cost equal to the Energy Overhead for that unit for any harmful state. The unit might, in fact, have been idle for a significant number of cycles before being woken up, so using the pessimistic estimate might disable power gating even when it was marginally saving energy. However, it allows us to guarantee, by using the Token Counting Guard described in Section 3.2, that the energy penalty is below the bound set by the user.

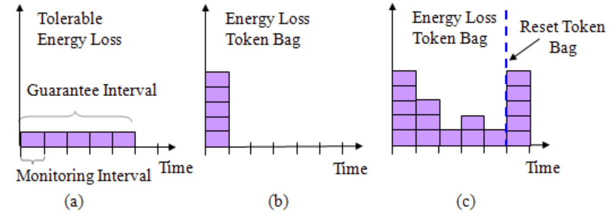

Figure 5. Token bag concept

3.1.2 Using Success Monitor to Improve Prediction

The information from the Success Monitor can be used by a hardware mechanism or by a high level software entity to dynamically change the power gating policy. In this work, we use the Success Monitor to drive an enable/disable signal for the power gating control logic. When the monitor estimates that the power gating policy has been harmful over the previous monitoring interval, we disable the policy during the next monitoring interval. Otherwise, the policy remains enabled. We call this solution the Success Monitor Switch. The efficiency counters are incremented and made available to the monitor regardless of whether the power gating policy is enabled. This permits the Success Monitor Switch to re-enable power gating when the monitor expects it to be beneficial.

Figure 4 shows the monitoring mechanism and this particular scenario.The lower part of the figure depicts the baseline IdleCount algorithm in which the system can be in one of three states: on (or active), power gated off but uncompensated (Off_U in figure), and power gated off and compensated.

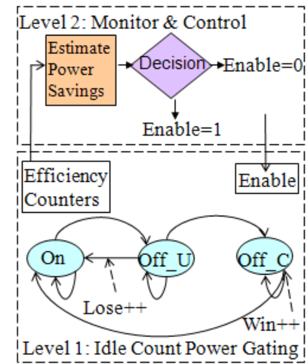

Figure 4. Efficiency counters

The success efficiency counter is incremented each time the unit remains in a power gated compensated state (Win++ in the figure). The harmful efficiency counter is incremented each time the unit goes from a power gated uncompensated state to being active again (Lose++).

3.2 Bounding Worst Case: Token Counting Guard

We propose adding a Token Counting Guard mechanism that provides a guarantee on the worst-case behavior of a policy. The guarantee is given over a time interval, called the *guarantee interval,* which is an integer multiple of the monitoring interval (Figure 5a). We associate tokens with the quantities we wish to limit for the power gating scheme. One token equals the leakage power of the unit over one cycle. A token bag holds the tokens that a unit can consume over the course of one guarantee interval, as illustrated in Figure 5b. Figure 5c shows how the number of tokens in the token bag varies over time. The token bag is updated as follows. At the beginning of a guarantee interval, the token bag is reset to a fixed, non-zero value that represents the entire amount of energy penalty that can be tolerated over the current guarantee interval. For example if we wish to guarantee a maximum leakage energy penalty of

2% over 100 monitoring intervals each 50 cycles long, then the token bag is initiated to 100 tokens.[1] The choice of interval lengths depends on the technology (e.g., thermal time constants, tolerable power overshoots, etc.) and the guarantees we wish to provide.

At the end of each monitoring interval, the token bag is updated depending on the energy savings or penalty estimated by the Success Monitor over this interval. The token bag will be increased if energy was saved or it will be decreased if energy was wasted. The quantity by which the token bag is updated corresponds to the token equivalent of the energy saved or wasted.

At the beginning of each monitoring interval, a decision is made, based on the number of tokens in the bag, whether to enable power gating for the next monitoring interval. If there are enough tokens to tolerate the worst possible behavior for the next monitoring interval, then the power gating is enabled. Otherwise it is disabled. Once power gating is disabled, it remains disabled until the end of the guarantee interval when tokens become available again. The benefit of the token bound mechanism is that it limits the penalty incurred by power gating in the worst-case scenario. However, we wish to achieve this bound without disabling power gating when it could save energy. *The key to achieving this goal is that there is a significant amount of energy savings slack across one guarantee interval for most workloads.* The power gating scheme is only disabled when all tokens have been consumed for that guarantee interval. By disabling power gating only in instances when it probably wastes energy, we see slightly greater energy savings for a system with Token Counting Guard compared to a system without it.

The application of the Success Monitor and the Token Counting Guard is not restricted to power gating schemes nor to power management in general. Any feature that, depending on runtime behavior, can succeed or not can benefit from these mechanisms.

3.3 Summary of Added Hardware

Success Monitor Switch. One counter, which we call the *power gating counter*, is required to count the number of cycles a unit is idle after being power gated, up to the break-even point. The counter is updated by an incrementer. Its size is $\lceil \log(\text{break_even point}) \rceil$, 5 bits for a break-even point of 19. Two additional counters are the efficiency counters. The size of the *success efficiency counter* is $\lceil \log(\text{monitoring interval}) \rceil$, which is 6 bits for a 50-cycle monitoring interval. Its value is incremented each time the power gating counter is at the break-even point value and the unit remains power gated. The size of the *harmful efficiency counter* is $\lceil \log(\text{monitoring interval}/(\text{idle_detect}+1)) \rceil$, 4 bits when idle_detect is 5 (7-bit due to 3-bit left shift necessary explained below). Its value is incremented whenever the power gating counter is less than the break-even point and the unit needs to wake-up from power gating. Both efficiency counters are reset at the beginning of each monitoring interval and updated by incrementers. The total cost of all harmful instances is computed by shifting the harmful efficiency counter (a 3-bit left shift when cost of one harmful instance is 8 units). To calculate whether the scheme is successful, a 6-bit subtractor subtracts the success efficiency counter from the shifted value of the harmful efficiency counter.

Token Counting Guard. A 12-bit register is necessary to hold the value of the token bag.[2] This register is initiated to a fixed value at the beginning of each guarantee interval (100 in our experiments for a 2% bound over 100 monitoring intervals). A 12-bit adder is necessary for updating the value of the token bag at the end of each monitoring interval, and a 6-bit register holds the value to be added/subtracted.

Total Hardware Cost. Individual efficiency counters and token bag registers are required for each power gateable unit. However, the rest of the hardware necessary for the Success Monitor Switch and Token Counting Guard can be shared across closely located power gateable units, because the monitoring of the units can be done at different cycles for different units. We did a careful analysis of the added hardware overhead, by calculating the "latch count equivalent" of all the components within the Success Monitor Switch and the Token Counting Guard. Based on known VHDL-derived latch counts and macro-level area estimates of the processor core, we were able to (very conservatively) bound the power overhead of the added hardware to at most 0.1% of a PowerPC-like core described in Section 4.1.

4. EVALUATION METHODOLOGY

4.1 Simulation Methodology

To evaluate the results of applying our proposed mechanisms for power gating, we use processor functional unit utilization traces from all SPEC2006 benchmarks. All the solutions we compare are run on the same traces (100 million cycles simulation equivalent). To derive the unit-level utilization traces, we instrument a cycle-accurate performance simulator pre-configured to model the details of the published core pipeline depth and superscalar execution semantics of each core within the POWER6 microprocessor [5][3]. Compared to Hu et al., we use a more modern processor model (POWER6-like compared to POWER4-like) and a more modern technology node (65nm compared to an older generation).

We implemented our power gating algorithms and proposed bounding mechanisms in an analysis tool written in C. This tool implements an energy model for power gating that allows us to obtain estimates of leakage energy savings. One limitation of this trace oriented framework is that it allows us only an approximate analysis of the potential performance impact a power gating scheme could have. We assign a fixed cost, in terms of cycles lost, whenever a power gated unit is awoken after being power gated, and we consider two possible wake-up costs. First, if decode-time signals are used to proactively wake up units before they are needed, the penalty is zero cycles. Second, if no such proactive wake-up is performed, we assign a 3-cycle wake-up penalty.

4.2 Energy Model

For estimating the energy saved by our proposed schemes, we refer to the energy model proposed by Hu et al. [2]. In this section we give a brief overview of the energy model.

---

1. 100 intervals*50 cycles/interval*2tokens/100 cycles = 100 tokens

2. The maximum value for the token bag is 3400 (<50*100). It can be calculated from the limit case when the scheme has maximum energy savings until interval i and maximum energy penalty after i: $100+50*i = 100(100-i)$ => $i = 66$ and the maximum token bag value is 3400.

3. The results and analysis presented in this paper are not claimed to be an accurate representation of any real PowerPC product in the marketplace. In fact, as evident from published papers [5], power-gating features are not part of server-class processors like POWER6.

Savings before full discharge. From equation (11) in Hu et al. [2], the energy savings in the i^{th} cycle after power gating and before the component is fully discharged, E^i, can be described by the formula:
$E^i = E_L * i * (\mathrm{DIBL}/mV_t) * \Delta V^0$
where E_L is the leakage power of the unit during one cycle, i is the number of cycles after power gating, DIBL is the drain-induced barrier lowering factor, V_t is the thermal voltage, and ΔV^0 is the voltage droop during the first cycle that the circuit is power gated. For current technology parameters, we obtain $E^i = 0.045 * i * E_L$.
Savings after full discharge. After full discharge, E^i equals E_L.[4]
Parameter values. The break-even point and energy overhead depend on the technology and the particular design macros. Hu et al. [2] used circuit-level simulations to validate viable ranges of these parameters, and they evaluated break-even points between 9-24. We use the same model input values as Hu et al., with two exceptions. We set the ratio between header device to the size of the power-gated circuit (W_H in their model) to 0.125, and we set the ratio between average leakage and switching energy dissipated per cycle (L in their model) to 0.3. Using equations (12) and (14) as in [2], we get a break-even point of 19 and energy overhead of $8 * E_L$.

5. EXPERIMENTAL RESULTS

We evaluate the ability of our mechanisms to bound the *worst-case* energy penalty of applying power gating, and we analyze their impact on the *average* energy savings of power gating (Section 5.1) and the *average* performance loss (Section 5.2). We compare the following power gating solutions:

- The baseline IdleCount [2] that we evaluate for 2 values of the IdleCount idle_detect (5 and 15, data presented for the 5 value due to space constraints).
- IdleCount augmented with the Token Counting Guard (Token_IdleCount). The token bag is updated every 50 cycles and the bound is set for a maximum of 2% leakage energy loss over a guarantee interval of 100 monitoring intervals (a total of 5000 cycles). The length of the monitoring interval (50 cycles) was motivated by wanting to quickly adapt to workload changes and reduce extra logic.
- IdleCount augmented with the Success Monitor Switch (Succ_IdleCount in the figures). The Success Monitor Switch is invoked every 50 cycles.

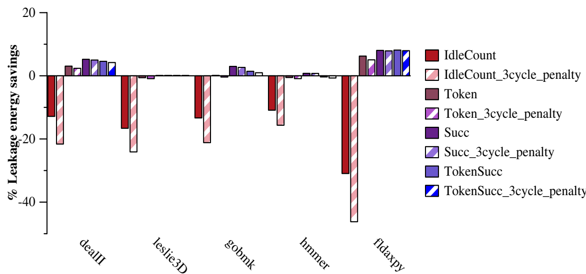

Figure 6. Comparative energy savings for benchmarks wasting energy under IdleCount scheme

4. Although the point of full discharge is not identical with the break-even point for the precise constants used, we approximate that the two are the same. This small approximation allows a single energy savings value for every cycle after the break-even point, and we believe the difference to be in the noise of Hu et al.'s theoretical model [2].

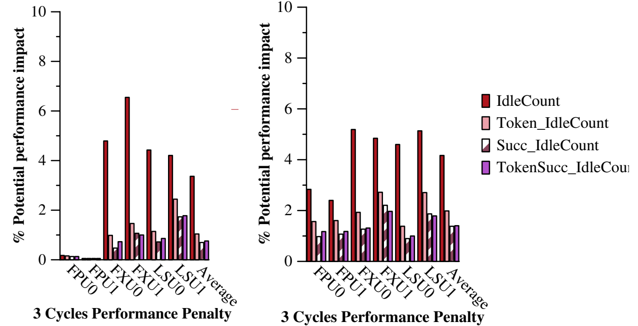

Figure 9. Average Perf. Impact (SpecINT, SpecFP)

- IdleCount with both the Success Monitor Switch and the Token Counting Guard (TokenSucc_IdleCount).

In our experiments, we consider the break-even point to be 19 and we vary the number of performance penalty cycles from zero to three. We present the potential performance impact as the percentage increase in cycles compared to the penalty incurred by the baseline. In reporting energy savings, we explicitly consider the leakage energy consumed during cycles added due to power gating. It is infeasible to evaluate our mechanisms for all possible parameter values or to evaluate them in the context of all previously developed power gating algorithms. However, we believe the chosen parameters and baseline are representative and that our results would be qualitatively similar for other parameters and baselines.

5.1 Worst-Case and Average-Case Energy

Figure 6 compares the energy savings for the benchmarks that exhibited an energy penalty under IdleCount. We do not present the data for the other benchmarks, because our schemes have negligible impact on their behavior. There are two bars presented for each of the compared schemes. The solid bar represents the energy savings for a system using decode signals to mask the performance penalty while the striped bar represents the energy savings of a system with a 3-cycle penalty paid on unit wake-up. Success_IdleCount eliminates the energy penalties incurred by the baseline system and transforms them into energy savings for all benchmarks. TokenSucc_IdleCount has a marginal energy penalty of below 1% only for *hmmer*. Token_IdleCount exhibits a marginal energy penalty of below 1% for *leslie3D*, *gobmk*, and *hmmer*. Note that the bound that was set for both Token_IdleCount and TokenSucc_IdleCount schemes was a maximum loss of 2%, which means they respect the bound and function correctly.
Figure 7 and Figure 8 compare the average energy savings for an idle_detect threshold of 5, for SPECint and SPECfp benchmarks, respectively. We assume that the wake-up penalty is three cycles; results (not shown) for a zero-cycle penalty are similar. We see that our schemes save slightly more energy than the baseline.

5.2 Performance

Figure 9 shows the worst-case potential performance impact of applying power gating, averaged across the SPECint and SPECfp benchmarks. We call it potential impact because much of it is likely to be masked due to other causes of stalls. All our proposed schemes decrease the baseline performance impact to a large extent, because disabling power gating when it is not saving energy also leads to a decrease in performance penalty.

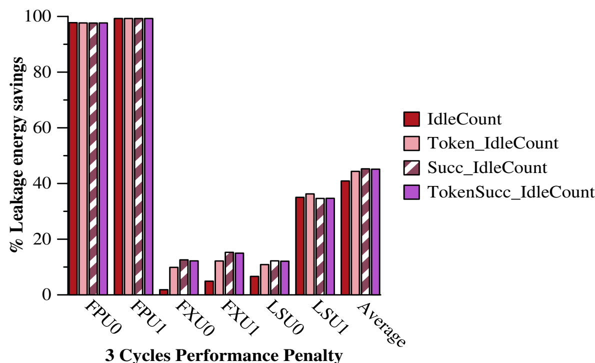

Figure 7. Average energy savings (SPECint)

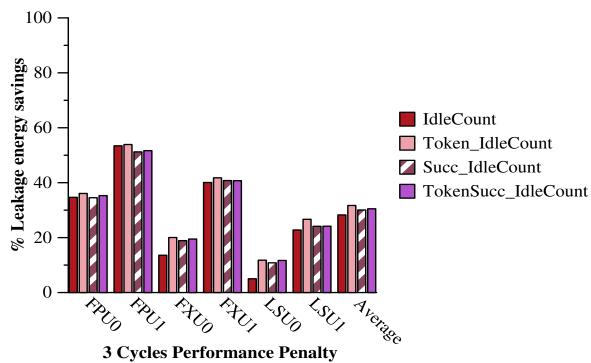

Figure 8. Average energy savings (SPECfp)

6. RELATED WORK

Power gating mechanisms and algorithms. Narendra and Chandrakasan provide an overview of power gating and other techniques to decrease leakage energy [6]. Intel's Nehalem processor [4] incorporates power gating techniques at the core granularity (instead of the finer functional unit granularity we consider). Because Nehalem's power gating control algorithms have not been published, we cannot estimate whether adding our guard mechanisms would be beneficial. At the granularity of functional units, there have been several proposed power gating algorithms [1, 7, 10], but all are capable of incurring energy penalties in the absence of a guard mechanism. At the circuit level, Jiang et al. [3] evaluate the benefits and costs of implementing power gating, while our focus here is on microarchitectural level mechanisms.

Energy management. Zeng et al. [11] introduce *currentcy* as a unifying abstraction used by the operating system for the management of devices that consume energy. Applications are allocated a certain currentcy that they can use during an epoch to satisfy their energy requirements. At a high level, the currentcy abstraction is similar to our token mechanism; however, we use tokens to bound the worst-case behavior of a power gating algorithm.

Other mechanisms for reducing leakage. Input vector control [9] applies an input pattern to a combinatorial circuit so as to decrease the leakage power while maintaining the same functional behavior. Adaptive body bias [8] is a post-silicon technique that reduces the effects of process variation on leakage power. These circuit-level schemes are complementary to power gating.

7. CONCLUSIONS

Power gating has potential drawbacks that, if not overcome, could preclude its use. We have presented two mechanisms for overcoming these problems. The key result is that we can bound the worst-case energy penalty incurred by a power gating scheme, with minimal impact on power gating's average energy savings. We believe that our mechanisms reduce the barrier to adopting power gating and make it an attractive solution for reducing leakage energy. An interesting avenue of future work is applying our mechanisms to speculative techniques other than power gating. Our mechanisms apply generally, and they can be used to reduce the risk of other speculative techniques that have high upsides but large potential losses in the case of excessive mispredictions.

ACKNOWLEDGMENTS

This material is based upon work supported (in part) by the Defense Advanced Research Projects Agency under its Agreement No. HR0011-07-9-0002 and the National Science Foundation under grant CCF-0811290.

REFERENCES

[1] S. Dropsho, V. Kursun, D. H. Albonesi, S. Dwarkadas, and E. G. Friedman. Managing Static Leakage Energy in Microprocessor Functional Units. In *Proc. of the 35th Annual International Symposium on Microarchitecture*, Nov. 2002.

[2] Z. Hu, A. Buyuktosunoglu, V. Srinivasan, V. Zyuban, P. Bose, and H. Jacobson. Microarchitectural Techniques for Power Gating of Execution Units. In *Proc. of the International Symposium on Low Power Electronics and Design*, Aug. 2004.

[3] H. Jiang, M. Marek-Sadowska, and S. R. Nassif. Benefits and Costs of Power-Gating Technique. In *Proc. of the International Conference on Computer Design*, 2005.

[4] R. Kumar and G. Hinton. A Family of 45nm IA Processors. In *Proc. of the International Solid-State Circuits Conference*, Feb. 2009.

[5] H. Q. Le et al. IBM POWER6 Microarchitecture. *IBM Journal of Research and Development*, 51(6), 2007.

[6] S. G. Narendra and A. Chandrakasan, editors. *Leakage in Nanometer CMOS Technologies*. Springer-Verlag, 2006.

[7] S. Rele, S. Pande, S. Onder, and R. Gupta. Optimizing Static Power Dissipation by Functional Units in Superscalar Processors. In *Proc. of the International Conference on Compiler Construction*, Apr. 2002.

[8] J. W. Tschanz et al. Adaptive Body Bias for Reducing Impacts of Die-to-Die and Within-Die Parameter Variations on Microprocessor Frequency and Leakage. *IEEE Journal of Solid-State Circuits*, 37(11), Nov. 2002.

[9] Y. Ye, S. Borkar, and V. De. A New Technique for Standby Leakage Reduction in High-Performance Circuits. In *Proc. of the Symposium on VLSI Circuits*, 1998.

[10] A. Youssef, M. Anis, and M. Elmasry. Dynamic Standby Prediction for Leakage Tolerant Microprocessor Functional Units. In *Proc. of the 39th Annual International Symposium on Microarchitecture*, Dec. 2006.

[11] H. Zeng et al. ECOSystem: Managing Energy as a First Class Operating System Resource. In *Proc. of the Tenth International Conference on Architectural Support for Programming Languages and Operating Systems*, Oct. 2002.