

# Exploiting Non-Uniform Memory Access Patterns Through Bitline Segmentation

Ravishankar Rao, Justin Wenck, Diana Franklin<sup>†</sup>, Rajeevan Amirtharajah and Venkatesh Akella  
University of California, Davis  
California Polytechnic State University, San Luis Obispo<sup>†</sup>

## Abstract

*On chip caches in modern processors account for a sizable fraction of the dynamic and leakage power. Much of this power is wasted, required only because the memory cells farthest from the sense amplifiers in the cache must discharge a large capacitance on the bitlines. We reduce this capacitance by segmenting the memory cells along the bitlines, and turning off the segmenters to reduce the overall bitline capacitance.*

*The success of this cache relies on accessing segments near the sense-amps much more often than remote segments. We show that the access pattern to the first level data and instruction cache is extremely skewed. Only a small set of cache lines are accessed frequently. We exploit this non-uniform cache access pattern by mapping the frequently accessed cache lines closer to the sense amp. These lines are isolated by segmenting circuits on the bitlines and hence dissipate lesser power when accessed.*

*Modifications to the address decoder enable a dynamic re-mapping of cache lines to segments. In this paper, we explore the design-space of segmenting the level one data and instruction caches.*

## 1 Introduction

A fundamental shift is occurring in microprocessor design. The unending quest for performance has led to an unquenchable thirst for power. The focus is shifting from pure performance to power-efficient performance. As this shift occurs, architects take a new look at existing structures. We take another look at the design of a cache.

Caches have traditionally been designed such that each element in the cache is accessed in the same amount of time, requiring the same amount of power. In reality, all elements within the cache are not accessed equally, so a cache should be designed to exploit the patterns that emerge in an application.

In essence, this is merely an extension of the original idea of caches. The basic hypothesis is that a subset of the

memory has locality, and that a cache can exploit this locality. Temporal and spatial locality are exploited by storing recently used data items, and data in the successive memory locations, in the cache. Once an item is placed in the cache, however, it becomes equal to all other elements in the cache. This assumes there is no locality within the cache, an assumption that wastes an opportunity for optimization.

Within the cache, some memory addresses will be accessed much more often than others due to temporal locality. Even without the temporal and spatial locality, some sets in the cache may be accessed more often than others because multiple cache sets, from distant locations in memory, happen to map to the same set. If a group of unrelated addresses that map to the same set are accessed more often than others, it also skews the access pattern in the cache.

We propose designing the cache such that the position within the cache determines the power consumption for accessing that set. This offers the opportunity to partition the cache at the circuit level into different power domains and explore micro-architectural techniques to remap the memory accesses into the appropriate power domain to save power. For example, a set that resides physically closer to the output will require less power than a more distant set.

In order to implement a low-overhead segmented bitline cache, we exploit two key characteristics in the design of modern caches. We first observe that caches require more power on the bitlines because every element must drive the signal along the whole length of the bitline, irrespective of its physical location in the cache. If this element had been closest to the output, it would have only needed to drive the distance between itself and the output. We exploit this by placing segmenters on the bitline, allowing an element to drive only the distance from its segment to the output, resulting in a lower power to attain the same performance. We focus on dividing the data array along the bitlines, as they consume a significant fraction of the total cache power.

In this paper, we present a segmented bit line cache implementation that minimizes the power and performance overhead of segmenting. We develop a power model of the segmented bitline cache using CACTI [14] and HSPICE. We then describe and evaluate two methods to dynamically

map sets to bitline segments of the cache to take advantage of different power domains based on the access pattern. Finally, we evaluate these methods in the context of level one instruction cache and data caches on a subset of integer and floating point SPEC2000 benchmarks.

The rest of the paper is organized as follows. We begin by describing related work in power efficient cache design and showing the difference between those approaches and ours in Section 2. We provide empirical evidence for the non-uniform access patterns in Section 3, which provides motivation for the work. Section 4 describes the implementation of our segmented bit line cache, followed by the clustering and mapping algorithms in Section 5. Evaluation methodology and results are presented in Section 6. Finally, we outline our future work and conclusions in Section 7.

## 2 Related Work

In this section we discuss several related research projects. In particular, two important areas of research are related to this work: circuit level optimizations of the bitlines for power savings and architectural techniques for power efficient caches.

At the circuit level, bitline isolation has been extensively studied. Kamble and Ghose [8] propose a local bitline for segments which are then connected to a common line across isolating switches. Higher metal layers are used to implement the common bitline. Since these layers have a lower capacitance, it reduces the overall bitline capacitance and hence saves dynamic power.

Yang and Kim propose a low power SRAM using a hierarchical bitline and local sense amplifiers (HBLSA-SRAM) [15]. The conventional SRAM bitline is divided into a several sub-bitlines each with its local sense amplifier. The sub-bitlines are connect to the read/write/pre-charge circuits through a common bit line. The HBLSA-SRAM reduces the write power in bitlines without noise degradation by applying a low swing signal to the bit line and a full swing signal to the sub-bitline.

Although bitline isolation and HBLSA-SRAM are similar to a segmentation of the bitline, they offer equal critical paths for the different segments. While this makes it possible to reduce the overall dynamic power dissipation, the techniques cannot be used to exploit the application characteristics. Our approach, which uses a segmented bitline has decreasing critical paths for segments closer to the sense amps, and hence dissipate lower power.

Several techniques have been suggested to exploit spatial locality to reduce leakage power. The spatial pattern predictor [5], predicts the usage within a cache line of data caches using a table-based predictor. The authors propose group-prefetching which predicts neighboring data items to be fetched with the data requested on a cache miss.

Yang and Falsafi [16] note that significant energy is wasted in statically pulling up the bitlines in all cache sub-arrays. They show the potential leakage power savings in not pre-charging the unaccessed sub-arrays. The authors study the energy and performance trade-offs of bitline isolation, and propose prediction techniques to exploit its full potential.

Other groups have suggested power-savings optimizations to reduce the power consumption when areas of the cache are not in use. Drowsy caches [7] sense when sets are not being used in the cache and put lines that have not been accessed lately into a low power *drowsy* mode. While these lines retain the data, they have to be reinstated to the high power mode before reading. This results in a slight performance penalty. Decay caches [9] suggest turning-off cache lines when they hold data not likely to be used. We focus on saving power even while in-use. Our design could be used in conjunction with either the drowsy cache or the decay cache to further reduce power.

NUCA [10] explored larger wire delay dominated secondary caches with several hundred banks of memory. These caches have latencies of 41 to 255 cycles in a 50nm process technology based on SIA Roadmap [2] predictions. Mapping, searching and routing of cache lines provide interesting design trade-offs at these latencies. NuRAPID [6] extended this work by exploring the use of sequential tag access for placement. While NUCA and NuRAPID exploit wire delay to develop a non-uniform access time cache for performance improvements, we suggest a non-uniform power dissipating cache for power savings.

## 3 Motivation

In order for us to gain power savings from our segmented bitline cache, the frequently accessed sets must be mapped to the lower-powered segments. Without segmentation overhead, a uniform access pattern would net no savings. Non-uniform cache access patterns in caches are critical to the success of this design. This section provides preliminary data which demonstrate these patterns.

We simulated an out-of-order processor using SimpleScalar [4]. The cache parameters were taken from the Itanium L1 cache [3]. The instruction and data caches are 4-way set associative, 16KB caches with a 64B set size. A subset (floating-point and integer) of the SPEC2000 benchmark suite was run to completion to generate detailed access distribution data.

Figures 1 and 2 give histograms showing the number of sets and the frequency at which they were accessed for L1 instruction and data caches. There were a total of 64 sets. A uniform access pattern would have resulted in about 1.6% of access to each set. We can see that for the L1 Instruction cache shown in Figure 1, *mcf*, *gcc*, and *quake* have 1-




Figure 1. Histogram: L1 I-Cache




Figure 2. Histogram: L1 D-Cache

3 cache sets that account for at least 10% of the accesses each. For *mcf*, about 78% of the sets are accessed less than 1% of the time. These applications are good candidates for our approach. For some applications, like *bzip*, most cache sets are accessed at the mean access rate. They will be poor candidates.

In the data cache shown in Figure 2, the accesses are not quite as skewed. Only one application has sets taking more than 10% of the total accesses. In addition, two applications, *mcf* and *equake*, have the bulk of their accesses near the mean rate.

Individual set accesses are only a part of the problem. Segmenting a cache would incur overheads in terms of additional hardware and delay, degrading the gains due to lower power on highly accessed elements. We address two issues in the sections ahead. First, we develop an implementation that minimizes performance overhead while maximizing the power savings. Second, we cluster the highly accessed sets together in the lowest-power partition while still providing a low-overhead means to map addresses to bitline segments.

## 4 Segmented Bitlines

Data and tag arrays in caches are typically partitioned into banks and further divided into blocks and sub-blocks. A detailed evaluation of various configurations can be found in [1]. The different partitions give different power-performance trade-offs. We used eCACTI [12] to partition a 16KB, 4 way set associative cache with a 64B block size. The tool gives various designs, each optimal for a specific parameter. While the optimal time configuration divides the wordline of the data array into eight, the lowest power design has four partitions for the wordline. The most area efficient configuration has two partitions of the bitline and does not divide the wordline. In this paper, since we primarily explore level 1 caches, we use the lowest access time configuration.

We then add segmenters to the bitline and study two, four

and eight segment bitline. Although this approach adds delay to the critical path of the bitline segment furthest from the sense amplifiers, it reduces the length of the bitline for the nearest cache rows thereby reducing the capacitance which has to be discharged, and hence the bitline power. The overall power consumption depends on:

- The power consumption overhead of the additional hardware
- The increased power of accessing the segments further from the sense-amplifiers.
- The reduced power consumption to access the nearer segments.
- The access distribution of the segments for different applications.

In this section, we first describe our proposed design. We then analyze the penalties and savings for cache rows as a function of location along the bitline segments.

### 4.1 Cache Design and Operation

In a standard SRAM cache a physical bitline is one continuous wire segment that connects many SRAM cells to the sense amplifier (SA), precharge and write circuitry. Every time the bitline voltage levels need to be changed, either for a read or a write, the entire bitline must be charged or discharged. As CMOS processes continue to scale down, the parasitic interconnect capacitance begins to dominate overall cache energy. Bitlines could consume up to 50% of the total cache power [8]. When accessing an SRAM cell that is physically close to the read/write/precharge circuitry, energy is wasted by charging/discharging the other end of the bitline. To prevent this waste of energy we use a segmented bitline shown in Figure 3. BL is the bitline and nBL is its complement. WL refers to a wordline. Each box has 16 SRAM cells connected between the bitline and its complement. The block diagram shows the bitline broken

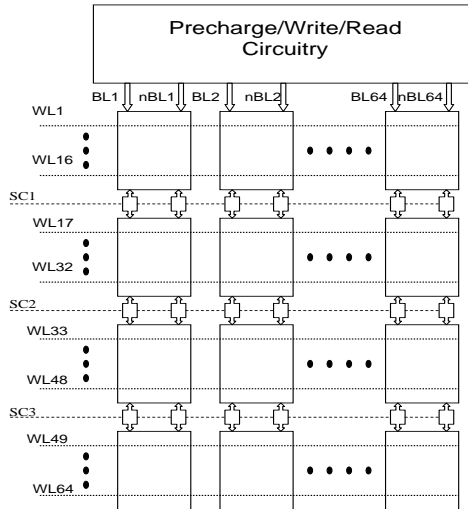


Figure 3. Segmented Bitlines

into 4 segments. We define a segment to be a portion of the bitline that can be isolated from any other bitline segment. The bitline can easily be broken into 2, 8, or any other combination of segments. Each segment is separated by a segmenter, which consists of a full CMOS transmission gate. Each transmission gate has a control signal (SC), which is operated each cycle according to the address of the data being accessed. When the clock is low all segmenters are turned on to enable the precharge of the entire bitline. Although it would take less power to pre-charge only part of the bitline up to the segment being accessed, in high performance caches, pre-charging overlaps address decoding and hence the segment address may not be decoded during pre-charge.

To get an idea of the timing and operation of the segmented cache line we will describe two simple cases. Consider a 4 segment bitline. In the first case, we read from a cell in the first segment, and in the second, we read from the fourth segment. In the first case, when the clock is low, all of the segmenters are on, and the entire bitline is precharged. When the clock goes high, precharging stops, and segment one is isolated from the other segments. The selected wordline then goes high, which causes one of the bitlines to discharge. After a fixed time the sense amps are turned on and the result of the read is latched. For the second case, we begin again with the clock low and all segments connected and precharged. When the clock goes high, all the segmenters stay on to allow the fourth segment to have a path to the sense amp. When the wordline goes high, one of the bitlines begins to discharge, the sense amp is enabled and the result is again latched.




Figure 4. Cache Power for accessing different segments

## 4.2 Power Model

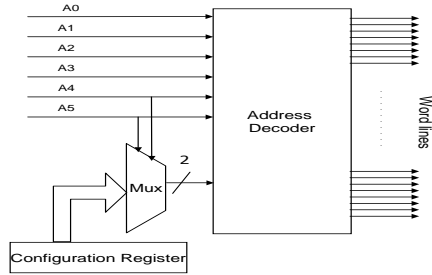
The power model was developed using of CACTI and HSPICE. The only difference between a conventional cache and a segmented bitline cache are the additional segmenters along the bitline. Hence, the power consumption of the rest of the cache is assumed to be as given by CACTI. A column with 64 SRAM cells was simulated in HSPICE using TSMC  $0.18\mu\text{m}$  technology to obtain the bitline power dissipation. The bitlines with two, four and eight segments were also simulated to account for the delay and power consumption of the segmenters.

Figure 4 plots the total cache power dissipation for accessing the different segments for a two, four and eight segmented bitline cache at 100MHz and 1.7V voltage. As seen from the figure, for the two segment case, the power dissipated to access the second segment is slightly higher than the unsegmented cache power. Similarly, in the four segment case, the third and fourth segments require higher power. Thus, our gains will be constrained by the percentage of accesses to the higher-power segments.

## 5 Clustering and Mapping

In a segmented bitline cache, the cache is divided into two, four or eight bitline segments. We define clustering as determining which cache sets can be partitioned together into a cluster. These clusters are then mapped to the bitline segments. Clustering and mapping introduce several challenges.

First, we do not have perfect knowledge of which sets will be accessed. We need a mechanism for predicting which sets should be placed into the low-power segments.



**Figure 5. Address Decoder**

Second, even if we had perfect knowledge, it would not necessarily be feasible to allow arbitrary permutation of sets to clusters. This would require a more complex logic to decode a set in the cache, increasing access times eroding some of our gains due to the shorter bitlines. We need a balance between flexibility in which cache lines may be clustered together and the speed of finding elements in the cache.

Finally, access patterns change with time requiring re-mapping at regular intervals. Performing a re-mapping can be an extremely high-overhead operation. Either the sets that need to be re-mapped must be copied to their new locations, or those sets must be flushed from the cache. Swapping cache lines can be very expensive in terms of power consumption, and flushing the sets from the cache would result in an increased miss-rate. Once again, a trade-off is required in balancing the cost of re-mapping with the additional power savings that could be obtained.

## 5.1 Clustering

Clustering can be done either statically or dynamically. With static clustering, a single clustering is used across all applications and throughout the run-time of the application.

Alternately, the clustering could be changed at run time, which is called dynamic clustering.

The simplest implementation of static clustering is to cluster the cache lines based on address bits. With 64 cache lines, for instance, 6 bits are required to determine the row address. Now, a subset of these bits determine the segments and the others the line in the segment. For example, with 4 segments 2 bits determine the segment and 4 bits determine which line within the segment is being accessed. The modified decoder for an address-based clustering is shown in Figure 5. The figure shows the case when MSBs determine the segment address. A configuration register stores the mapping of the clusters to the segments. The address bits control the mux (in this case, two 4:1 muxes are required.), which selects the new segment address. The only increase in the critical path of the decoder is the multiplexer delay.

We still have two choices for static clustering - hard-code which two address bits determine the segment or allow any subset of bits to determine the segment address. The latter requires a full cross-bar in the critical path to choose the address bits that control the mux. This would significantly increase the delay of the decoder. In this paper, we propose using the MSBs for the segment address.

## 5.2 Mapping

Once we have a clustering scheme, we need to map the cluster into bitline segments. The basic idea is to map the most frequently accessed clusters into the lowest power segments. We associate a counter with every cluster and increment it each time the cluster is accessed.

We study two versions of mapping: static mapping and dynamic re-mapping. In static mapping, a single mapping of clusters to segments, obtained by profiling, is used throughout the application.

The dynamic re-mapping uses the access counters to re-map at regular intervals. Two versions of dynamic re-mapping are explored in this paper. In the first version, the segment access information of only the previous interval is used for re-mapping. We refer to this case as 'dynamic counter flush' or **dcf** because the counters used for re-mapping are flushed each time a re-mapping occurs. In this second case called 'dynamic no counter flush' or **dncf**, the cumulative access counts to the cluster from the beginning of the application to the current interval is used to decide re-mapping.

Both static and dynamic techniques have their advantages. A static profile-based mapping has the overall knowledge of the accesses to the clusters, while the dynamic techniques can better capture the changing program behavior. The two cases of dynamic re-mapping further explore the use of history to decide the mapping. Dynamic cache re-mapping could be an expensive task. The remapped clusters need to be either copied or invalidated. For instance, let us assume that a cluster A is mapped to segment 1 and cluster B to segment 2. Assume that a dynamic re-mapping results in swapping the two clusters. To achieve this, we could either invalidate the contents of the two segments or swap them. Swapping entire cache segments every time a re-mapping occurs would be a power-intensive task. Invalidating the contents would result in a miss to the cache location. In this paper, we propose to limit the re-mapping to context switches. We assume that during a context switch, much of the state in the cache is lost, so invalidating the cache will incur a negligible performance penalty.




Figure 6. Static vs. Dynamic :I-cache

## 6 Results

Simulations were run with SimpleScalar. Counters were added to generate cache line access statistics and re-mapping was implemented on every 1 million cycles. The experiments were run on a subset of the SPEC2000 benchmarks, drawn from both the floating-point and integer suites, using a 16KB 4-way set associative cache for integer and data caches. We used SimPoint3.0 [13] to reduce the simulation time. SimPoint provides representative intervals and weights for each application. Multiple standard simulation points each 100 million instructions long were simulated. For each benchmark, all the intervals were simulated, and the results were weighted using the interval weights given by SimPoint. The weighted percentage of access to each segment was then used as an activity factor to scale the power of the corresponding segment. We perform a series of experiments to evaluate the benefits of these techniques. The graphs show the total cache power savings relative to an unsegmented cache. First, we compare static mapping based on a profile and dynamic re-mapping, with that of a dynamic clustering. We then vary the number of segments from 1 to 8 segments. The operation of a segmented bitline cache at higher frequency is discussed.

### 6.1 Static vs Dynamic Re-mapping

Figures 6 and 7 compare three choices of mapping: static, dynamic with cumulative counters, and dynamic with counters that reset each interval. We present the results for eight segments, and compare with that of a best case dynamic clustering labeled as oracle.

For the instruction cache, the average power savings from a static mapping is 6.3% while the two dynamic re-mapping techniques provide approximately 7.6% average power savings. The dynamic clustering indicates that there




Figure 7. Static vs. Dynamic :D-cache




Figure 8. I-Cache: Number of Segments

is a potential to save up to 10% power on average. For some applications like *applu*, the dynamic techniques achieve close to maximum power savings, while others like *gzip* indicate that the potential is much more than what is obtained by the current dynamic re-mapping techniques. Also, note that maximum power savings are obtained for applications like *mcf* and *gcc* for which access patterns are most-skewed.

The data caches have less power savings. Static mappings provide about 4% savings, whereas dynamic re-mapping performs slightly better at 4.5% power savings. The dynamic clustering results indicate a potential power savings of about 5.7% on average. This is consistent with the more uniform access distribution we noted earlier for data caches. Again, as noted earlier, we see a direct relation between power savings and non-uniformity of the access distribution. Applications like *mcf* and *equake* give least power savings because most of their cache sets are accessed at about the mean rate.




Figure 9. D-Cache: Number of Segments




Figure 10. Power Dissipated at 500 MHz

## 6.2 Number of Segments

The number of segments is an interesting design trade-off. With more segments, you can save more from the low-power segment, but you lose more from the high-power segment. We vary the number of segments from 2 through 8. Since the dynamic remapping with counter resetting performs better than static mapping and does not require prior knowledge, we present only the **dcf** results. Figures 8 and 9 show the results for the instruction cache and the data cache respectively. For instruction caches, increasing the number of segments increases the power savings. While eight segments gives about 7.6% power savings on average, the savings reduce slightly to 7.4% and 4.1% for 4 and 2 segments respectively. This is because instruction caches showed greater non-uniformity in their access patterns, allowing them to greatly benefit from the extra segments. The data caches show a good inflection point giving best power savings at four segments. A more uniform access pattern of the data cache leads to more accesses to the high-power segments, decreasing power savings for eight segments.

## 6.3 Increasing Frequency

The results presented in the previous sections are at 100 MHz where the segmenter delay was small compared to the clock cycle. As the frequency is increased, the voltage to which the bitline (or its complement) discharges is reduced. Adding segmenters to this bitline, further increases the bitline delay, thereby reducing the discharge voltage of segments even more. This reduces the power savings obtained by segmentation because the power dissipated on the bitlines is directly proportional to the voltage to which the bitline discharges. Its effect on the applications is shown in Figures 11 and 12. The figures give the power savings for different applications for the **dcf** mapping for an eight segment bitline.

The decrease in bitline discharge also potentially reduces the reliability of segments further away from the sense-amp circuitry since reduced swing increases the likelihood of sense-amp failure. Without quantifying reliability, we assume here that it is inversely proportional to the discharge voltage of the bitline. The lower the final voltage to which a bitline discharges, the greater is its reliability. This in turn is proportional to the number of segmenters between the segment and the sense-amp circuitry.

Segments that are closer to the sense-amp have higher reliability than the unsegmented bitline. For instance, in an eight segment bitline, the first three segments have a higher discharge than an unsegmented bitline. So using our remapping techniques, if most accesses are limited to these segments, an overall higher reliability is possible. Even so, in the eight segment case, half the cache has a lower reliability than an unsegmented bitline. It would be ideal if we could obtain the potential power savings of an eight segment bitline, while retaining the reliability of an unsegmented bitline. To achieve this goal we simulated unequal segments in our bitline. Three cases of unequal segments are considered in this paper:

- A two segment bitline, with eight rows in the first segment and 56 in the second segment.
- A three segment bitline, with eight rows in the first segment, 16 in the second and 40 rows in the third segment
- A five segment bitline, with eight rows each in the first four segments, and a fifth segment with the remaining 32 rows.

The power dissipated by the various configurations are shown in Figure 10. The dark lines represent unequal segments, while the lighter lines represent equal segment cases. Note that unequal segments have fewer segmenters than an




Figure 11. I-Cache: Increased frequency




Figure 12. D-Cache: Increased frequency




Figure 13. I-Cache: Unequal Segments




Figure 14. D-Cache: Unequal Segments

eight segment bitline, while retaining regions with much lower power than an unsegmented bitline. Since they have fewer segmenters, the bitlines discharge to lower voltages, and thus, they do not have regions with reliability significantly lower than an unsegmented bitline. The power savings due to these unequal segment configurations is compared to that of eight and four segment bitlines, for the **dcf** mapping in Figures 13 and 14. As seen from the figures in both the data and instruction caches, for most applications, power savings are between those of an eight and a four segment bitline.

## 7 Conclusions

In this paper, we present architectural techniques that are combined with simple circuit modifications of a cache. We exploit the non-uniform access patterns of different applications to obtain dynamic power savings. We explore partitioning the cache lines into clusters and mapping them to segments dissipating different power.

In this paper, a re-mapping occurs on a context switch to limit the performance and power overhead. Instead, the re-

mapping could be guided by the changing program phases. SimPoint has tools to detect phase changes in applications at runtime [11]. By integrating these tools clusters could be re-mapped only during the phase changes.

Embedded processors running media applications typically execute a small set of instructions repeatedly. The instruction cache would thus have a much greater physical locality. Hence a segmented bitline cache could be more appropriate in this context. Also, since the operating frequencies in embedded processors are lower, for the same size cache, more delay can be tolerated. Hence, the bitline segments can be operated at much lower frequencies resulting in greater power savings. We intend to evaluate the dynamic re-mapping and possible implementations of dynamic clustering in the case of embedded processors. Finally, the effect of bitline segmenting on leakage power will be evaluated.

## References

- [1] B. S. Amrutur and M. A. Horowitz. Speed and power scaling of srams. *IEEE Journal of Solid-State Circuits*, 35(2):175–185, February 2000.

- [2] S. I. Association. The national technology roadmap for semiconductors. 1999.
- [3] D. Bradley, P. Mahoney, and B. Stackhouse. The 16kb single-cycle read access cache on a next-generation 64b itanium microprocessor. In *International Solid State Circuits Conference*, 2002.
- [4] D. C. Burger and T. M. Austin. The simplescalar tool set, version 2.0. Technical Report CS-TR-1997-1342, University of Wisconsin, Madison, June 1997.
- [5] C. F. Chen, S.-H. Yang, B. Falsafi, and A. Moshovos. Accurate and complexity-effective spatial pattern prediction. In *10th International Symposium on High-Performance Computer Architecture*, Feb 2004.
- [6] Z. Chisti and M. P. adn T.N.Vijaykumar. Distance associativity for high-performance energy efficient non-uniform cache architectures. In *Proceedings of the 36th International Symposium on Microarchitecture*, 2003.
- [7] K. Flautner, N. S. Kim, S. Martin, D. Blaauw, and T. Mudge. Drowsy caches: Simple techniques for reducing leakage power. In *29th Annual International Symposium on Computer Architecture*, may 2002.
- [8] K. Ghose and M. B. Kamble. Reducing power in superscalar processor caches using subbanking, multiple line buffers and bit-line segmentation. In *International Symposium on Low Power Electronics and Design*, pages 70–75, 1999.
- [9] S. Kaxiras, Z. Hu, and M. Martonosi. Cache decay: Exploiting generational behavior to reduce cache leakage power. In *28th Annual International Symposium on Computer Architecture*, June-July 2001.
- [10] C. Kim, D. Burger, and S. Keckler. An adaptive, non-uniform cache structure for wire-delay dominated on-chip caches. In *10th International Conference on Architectural Support for Programming Languages and Operating Systems*, 2002.
- [11] J. Lau, S. Schoenmackers, and B. Calder. Transition phase classification and prediction. In *11th International Symposium on High Performance Computer Architecture*, Feb 2005.
- [12] M. Mamidipaka and N. Dutt. eacti: An enhanced power estimation model for on-chip caches. In *Center for Embedded Computer Systems Technical Report TR-04-28*, Sept 2004.
- [13] T. Sherwood, E. Perelman, G. Hamerly, and B. Calder. Automatically characterizing large scale program behavior. In *10th International Conference on Architectural Support for Programming Languages and Operating Systems*, Oct 2002.
- [14] S. J. Wilton and N. P. Jouppi. Cacti: An enhanced cache access and cycle time model. In *IEEE Journal of Solid-State Circuits*, May 1996.
- [15] B.-D. Yang and L.-S. Kim. A low-power sram using hierarchical bit line and local sense amplifiers. In *IEEE Journal of Solid-State Circuits*, June 2005.
- [16] S.-H. Yang and B. Falsafi. Near-optimal precharging in high-performance nanoscale cmos caches. In *36th International Symposium on Microarchitecture*, 2003.