

Using Dead Blocks as a Virtual Victim Cache

Samira Khan
Univ. of Texas at San Antonio
skhan@cs.utsa.edu

Doug Burger
Microsoft Research
dburger@microsoft.com

Daniel A. Jiménez
Univ. of Texas at San Antonio
dj@cs.utsa.edu

Babak Falsafi
EPFL
babak.falsafi@epfl.ch

ABSTRACT

Caches mitigate the long memory latency that limits the performance of modern processors. However, caches can be quite inefficient. On average, a cache block in a 2MB L2 cache is *dead* 59% of the time, i.e., it will not be referenced again before it is evicted. Increasing cache efficiency can improve performance by reducing miss rate, or alternately, improve power and energy by allowing a smaller cache with the same miss rate.

This paper proposes using predicted dead blocks to hold blocks evicted from other sets. When these evicted blocks are referenced again, the access can be satisfied from the other set, avoiding a costly access to main memory. The pool of predicted dead blocks can be thought of as a *virtual victim cache*. For a set of memory-intensive single-threaded workloads, a virtual victim cache in a 16-way set associative 2MB L2 cache reduces misses by 26%, yields a geometric mean speedup of 12.1% and improves cache efficiency by 27% on average, where cache efficiency is defined as the average time during which cache blocks contain live information. This virtual victim cache yields a lower average miss rate than a fully-associative LRU cache of the same capacity. For a set of multi-core workloads, the virtual victim cache improves throughput performance by 4% over LRU while improving cache efficiency by 62%.

Alternately, a 1.7MB virtual victim cache achieves about the same performance as a larger 2MB L2 cache, reducing the number of SRAM cells required by 16%, thus maintaining performance while reducing power and area.

Categories and Subject Descriptors

B.3.2 [Hardware]: Memory Structures—*Cache Memories*

General Terms

Design, Performance

Keywords

microarchitecture, cache management, prediction

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

PACT'10, September 11–15, 2010, Vienna, Austria.

Copyright 2010 ACM 978-1-4503-0178-7/10/09 ...\$10.00.

1. INTRODUCTION

The performance gap between modern processors and memory is a primary concern for computer architecture. Processors have large on-chip caches and can access a block in just a few cycles, but a miss that goes all the way to memory incurs hundreds of cycles of delay. Thus, reducing cache misses can significantly improve performance.

One way to reduce the miss rate is to increase the number of *live* blocks in the cache. A cache block is live if it will be referenced again before its eviction. From the last reference until the block is evicted the block is *dead* [10]. Studies show that cache blocks are dead most of the time. For the benchmarks and 2MB L2 cache used for this study, cache blocks are dead on average 59% of the time.

Dead blocks lead to poor cache efficiency [12, 4] because a block may reside in the cache for a long time between the time it is last accessed and the time it is evicted. In the least-recently-used (LRU) replacement policy, after the last access, a block has to move down from the most-recently-used (MRU) position to the LRU position and then it is evicted. This process will take a long time in a highly associative cache.

Cache efficiency can be improved by replacing dead blocks with a live blocks as soon as possible after a block becomes dead, rather than waiting for it to be evicted. Having more live blocks in the cache improves the system performance by reducing misses since more live blocks means more cache hits. Alternately, a technique that increases the number of live blocks may allow reducing the size of the cache, resulting in a system with the same performance but reduced power and energy needs.

This paper describes a technique to improve cache performance by using predicted dead blocks to hold victims from cache evictions in other sets. The pool of predicted dead blocks can be thought of as a *virtual victim cache* (VVC). Figure 1 graphically depicts the efficiency of a 1MB 16-way set associative L2 cache with LRU replacement for the SPEC CPU 2006 benchmark `456.hammer`. The amount of time each cache block is live is shown as a greyscale intensity. Figure 1(a) shows the unoptimized cache. The darkness shows that many blocks remain dead for large stretches of time. Figure 1(b) shows the same cache optimized with the VVC idea. Now many blocks have more live time so the cache is more efficient.

The VVC idea uses a dead block predictor, i.e., a microarchitectural structure that uses past history to predict whether a given block is likely to be dead at a given time. This study uses a trace based dead block predictor [10].

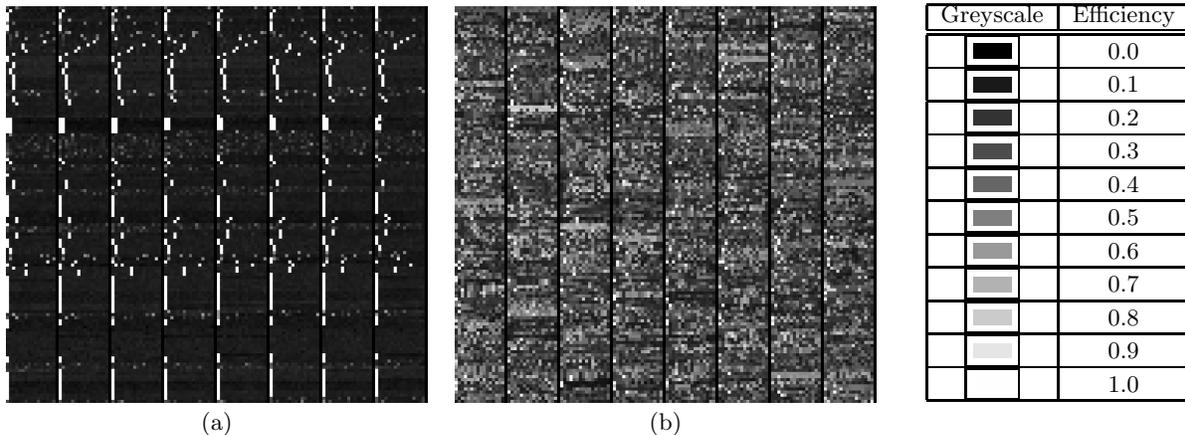


Figure 1: Virtual victim cache increases cache efficiency. Block efficiency (i.e., fraction of time block is live) shown as greyscale intensities for *456.hmmmer* for (a) a baseline 1MB cache and (b) a VVC-enhanced cache; darker blocks are dead longer.

This idea has some similarity to the victim cache [6], but victims are stored in the same cache from which they were evicted, simply moving from one set to another. When a victim block is referenced again, the access can be satisfied from the other set. Another way to view the idea is as an enhanced combination of block insertion (i.e. placement) policy, search strategy, and replacement policy. Blocks are initially placed in one set and migrated to less active set when they become least-recently-used. A more active block is found with one access to the tag array, and a less active block may be found with an additional search.

1.1 Contributions

This paper makes the following contributions:

1. The virtual victim cache, i.e. the concept of caching evicted blocks in predicted dead blocks, is introduced.
2. An improved dead block predictor inspired by branch predictors is introduced. This predictor organization reduces harmful false positive predictions by over 10% on average, significantly improving the performance of the VVC with the potential to improve other optimizations that rely on dead block prediction.
3. For single-threaded workloads, the VVC reduces the number of cache misses per thousand instructions (MPKI) by 26% on average with a 2MB L2 cache, yields a geometric mean speedup of 12.1% over the baseline and improves cache efficiency by 27% on average. On average, the VVC achieves a number of misses halfway between LRU and optimal. The VVC outperforms several other caching strategies.
4. For multi-core workloads, the VVC improves throughput performance by 4% and cache efficiency by 62% while outperforming several other techniques.
5. We show that a policy of simply evicting predicted dead blocks can achieve almost the same performance as the VVC for single-threaded workloads, but results in a 6% reduction in throughput performance for CMP workloads. Thus, the VVC is superior to dead block replacement.

2. RELATED WORK

In this section we discuss related work. Previous work introduced several dead block predictors and applied them to problems such as prefetching and block replacement [10, 12, 8, 5, 1], but did not explore coupling dead block prediction with alternative block placement strategies.

2.1 Dead Block Predictors

The VVC depends accurately identifying dead blocks to replace with victims from other sets. We discuss related work in dead block prediction.

2.1.1 Trace Based Predictor

Dead block prediction was introduced by Lai *et al.* [10]. Given sequence of accesses to a given cache block leads to last access of the block, that same sequence of accesses to a different block is likely to lead to the last access of that block. An access is represented by the program counters (PC) of the instructions making the access. The sequence, or *trace* of PCs of the instructions accessing a block is encoded as the fixed-length truncated sum of hashes of these PCs. This trace encoding is called a *signature*. For a given block, the trace for the sequence of PCs begins when the block is filled and ends when the block is evicted. The predictor learns from the trace encoding of the evicted blocks. A table of saturating counters is indexed by block signatures. When a block is replaced, the counter associated with it is incremented. When a block is accessed, the corresponding counter is decremented. A block is predicted dead when the counter corresponding to its signature exceeds a threshold. This predictor is used to prefetch data into predicted dead blocks in the L1 data cache, enabling lookahead prefetching and eliminating the necessity of prefetch buffers. A trace based predictor is also used to optimize a cache coherence protocol [9, 24]. Dynamic self-invalidation involves another kind of block “death” due to coherence events [11]. PC traces are used to detect the last touch and invalidate the shared cache blocks to reduce cache coherence overhead.

2.1.2 Counting and Time Based Predictor

Dead blocks can also be predicted depending on how many times a block has been accessed. Kharbutli and Solihin pro-

pose counting based predictors for the L2 cache [8]. The Live-time Predictor (LvP) keeps track of the number of accesses to each block. This value is stored in the predictor on eviction. A block is predicted dead if it has been accessed more times than the previous generation.

Another approach is to predict a block dead when it is not accessed for a certain amount of time. Hu *et al.* propose a time based dead block predictor [5] that learns the number of cycles a block is live and predicts it dead if it is not accessed for more than twice that number of cycles. This predictor is used to prefetch into the L1 cache and filter a victim cache. Abella *et al.* propose a similar predictor [1] based on number of references rather than cycles for reducing cache leakage power by dynamically turning off L2 cache blocks whose content is not likely to be reused without hurting the performance.

2.1.3 Cache Burst Predictor

Cache bursts [12] can be used with trace based, counting based and time based dead block predictors. A cache burst consists of all the contiguous accesses that a block receives while in the most-recently-used (MRU) position. Instead of each individual references, cache burst based predictor updates the predictor only on each bursts. It also improves prediction accuracy by making prediction only when a block moves out of the MRU position. The dead block predictor needs to store trace or reference count information for each burst only rather than for each reference. But since prediction is made only after a block becomes non MRU, some of the dead time is lost compared to non burst predictors. Cache burst predictors improve prefetching, bypassing and enhancing LRU replacement policy both for L1 Data cache and L2 cache.

2.1.4 Other Dead Block Predictors

Another kind of dead block prediction involves predicting in software [26, 22]. In this approach the compiler collects dead block information and provides hints to the microarchitecture to make cache decisions. If a cache block is likely to be reused again it hints to keep the block in the cache; otherwise, it hints to evict the block.

2.2 Cache Placement and Replacement Policy

Qureshi *et al.* [16] proposed Dynamic Insertion Policy (DIP) to improve cache performance by preventing thrashing in large workloads. The technique proposed in improved block insertion (i.e. placement) policy. The victim selection policy selects the least recently used (LRU) block. The insertion policy decides what position in the LRU stack the incoming block is placed. Two distinguished positions are consider: the LRU and MRU positions. For most workloads, the standard MRU insertion policy is sufficient. However, for workloads whose working sets exceed the capacity of the cache, some fraction of the working set can be retained in the cache by inserting the incoming block in the LRU position. This LRU line moves to the MRU position only if it is referenced again, otherwise it is evicted. This ensures that blocks that are never used between insertion and eviction do not occupy cache space. This insertion policy can hurt performance of LRU friendly workloads, so a technique called *set dueling* is used to dynamically determine the best policy [17]. Set dueling dedicates a small number of sets in the cache to LRU replacement policy and a another small

set of sets to LRU Insertion policy. The policy resulting in fewer misses in the dedicated sets wins and rest of the sets in the cache follows that policy.

Rolán *et al.* propose a technique that places evicted blocks into other less active sets, which is similar to the VVC [21]. However, the performance gains of 5% on 10 SPEC CPU benchmarks are more modest since this technique does not use dead block prediction to precisely identify candidate blocks to hold victims. This technique is not evaluated with multi-core workloads.

Keramidas *et al.* [7] proposed a cache replacement policy that uses reuse distance prediction. This policy tries to evict cache blocks that will be reused furthest in the future. A memory-level parallelism aware cache replacement policy relies on the fact that isolated misses are more costly on performance than parallel misses [18].

The virtual victim cache can be seen as a way of providing extra associativity. The V-way cache also provides variable associativity based on program behavior [20].

3. USING DEAD BLOCKS AS A VIRTUAL VICTIM CACHE

Victim caches [6] work well because they effectively extend the associativity of any hot set in the cache, reducing localized conflict misses. However, victim caches must be small because of their high associativity, and are flushed quickly if multiple hot sets are competing for space. Thus, victim caches do not reduce capacity misses appreciably, nor conflict misses where the reference patterns do not produce a new reference to the victim quickly, but they provide excellent miss reduction for a small additional amount of state and complexity. Larger victim caches have not come into wide use because any additional miss reduction benefits are outweighed by the overheads of the larger structures.

Large caches already contain significant quantities of unused state, however, which in theory can be used for optimizations similar to victim caches if the unused state can be identified and harvested with sufficiently low overhead. Since the majority of the blocks in a cache are dead at any point in time, and since dead-block predictors have been shown to be accurate in many cases, the opportunity exists to replace these dead blocks with victim blocks, moving them back into their set when they are accessed. This *virtual victim cache* approach has the potential to reduce both capacity misses and additional conflict misses: Capacity misses can be reduced because dead blocks are evicted before potentially live blocks that have been accessed less recently (avoiding misses that would occur with full associativity), and conflict misses can be further reduced if hot set overflows can spill into other dead regions of the cache, no matter how many hot sets are active at any one time.

An important question is how the dead blocks outside of a set are found and managed without adding prohibitive overhead. By coupling small numbers of sets, and moving blocks overflowing from one set into the predicted dead blocks (which we call *receiver blocks*) of a “partner set,” a virtual victim cache can be established with little additional overhead. While this approach effectively creates a higher-associativity cache, the power overheads are kept low because only the original set is searched the majority of the time, with the partner sets only searched upon a miss in the original set. The overheads include more tag bits (the log

of the number of partner sets) and more energy and latency incurred on a cache miss, since the partner sets are searched to no avail.

3.1 Description of the VVC

3.1.1 Identifying Potential Receiver Blocks

A trace based dead block predictor keeps a trace encoding for each cache block. The trace is updated on each use of the block. When a block is evicted from the cache, a saturating counter associated with that block’s trace is incremented. When a block is used, the counter is decremented.

Ideally, any victim block could replace any receiver block in the entire cache, resulting in the highest possible usage of the dead blocks as a virtual victim cache. However, this idea would increase the dead block hit latency and energy as every set in the cache would have to be searched for a hit. Thus, there is a trade-off between the number of sets that can store victim blocks from a particular set and the time and energy needed for a hit. We have determined that, for each set, considering only one other partner set to identify a receiver block yields a reasonable balance. Sets are paired into *adjacent sets* that differ in their set indices by one bit.

3.1.2 Placing Victim Blocks into the Adjacent Set

When a victim block is evicted from a set, the adjacent set is searched for invalid or predicted dead receiver blocks. If no such block is found, then the LRU block of the adjacent set is used. Once a receiver block is identified, the victim block replaces it. The victim block is placed into either the most-recently-used or the least-recently-used position in the LRU stack based on set dueling (see Section 3.1.4).

3.1.3 Block Identification in the VVC

If a previously evicted block is referenced again, the tag match will fail in the original set, the adjacent set will be searched, and if the receiver block has not yet been evicted then the block will be found there. The block will then be refilled in the original set from the adjacent set, and the block in the adjacent set will be marked as invalid. A small penalty for the additional tag match and fill will accrue to this access, but this access is considered a hit in the L2 for purposes of counting hits and misses (analogously, an access to a virtually-addressed cache following a TLB miss may still be considered a hit, albeit with an extra delay).

To distinguish receiver blocks from other blocks, we keep an extra bit with each block that is true if the block is a receiver block, false otherwise. When a set’s tags are searched for a normal cache access, receiver blocks from the adjacent set are prevented from matching to maintain correctness. Note that keeping this extra bit is equivalent to keeping an extra tag bit in a higher associativity cache.

Figure 2(a) shows what happens when a LRU block is evicted from a set s . If the adjacent set v has any predicted dead or invalid block in it, the victim block replaces that block, otherwise the LRU block of set v is used. Similarly, Figure 2(b) depicts a VVC hit. If the access results in a miss in the original set s , that block can be found in a receiver block of the adjacent set v . Algorithm 1 shows the complete algorithm for the VVC.

3.1.4 Set Dueling for Placement

When a block is placed into an adjacent set, we must

decide where to place it. Some workloads benefit from placing the receiver of a victim block into the most-recently-used (MRU) position, while others benefit from placement in the LRU position depending on the behavior of the program. For example, benchmarks such as `187.facerec` and `401.bzip2` perform well when evicted blocks are placed in MRU position, but `179.art`, `188.amp` do not. For some workloads, the MRU insertion policy might be better because a victim block is not be accessed immediately, but it will be accessed soon so it will still be in the adjacent set making its way down the LRU stack. For other workloads, the LRU insertion policy might be better because 1) when a block will be accessed again, it will be accessed soon, or 2) many victim blocks placed in adjacent sets will not used for a long time and might as well be evicted instead of more useful data; after all, this is the point of the LRU replacement policy.

We used set dueling [17] to determine whether to place victims into the MRU or LRU position. A small number of sets in the cache are dedicated to LRU placement policy and a another small set of sets to MRU placement. An 11-bit counter is incremented whenever there is a miss to one of the dedicated LRU sets, or decremented when there is a miss to one of the dedicated MRU sets. The LRU placement policy is only used if the counter falls below 0; thus, the policy resulting in fewer misses in the dedicated sets wins and rest of the sets in the cache follow that policy. In this work we use 32 dedicated sets for each policy.

Algorithm 1 Virtual Victim Cache with Trace based Predictor

```

On an access to set  $s$  with address  $a$ , PC  $pc$ 
if the access is a hit in block  $blk$  then
     $blk.trace \leftarrow updateTrace(blk.trace, pc)$ 
     $isDead \leftarrow lookupPredictor(blk.trace)$ 
    if  $isDead$  then
        mark  $blk$  as dead
    return
end if
else /* search adjacent set for a dead block hit */
     $v = adjacentSet(s)$ 
    access set  $v$  with address  $a$ 
    if the access is a hit in a dead block  $dbl$  then
        bring  $dbl$  back into set  $s$ 
    return
end if
else /* this access is a miss */
     $replbk \leftarrow$  block chosen by LRU policy
     $updatePredictor(replbk.trace)$ 
     $v = adjacentSet(s)$ 
    place  $replbk$  in an invalid/dead/LRU block in set  $v$ 
    into position (MRU vs. LRU) determined by set dueling
    place block for address  $a$  into  $replbk$ 
     $replbk.trace \leftarrow updateTrace(pc)$ 
return
end if

```

3.1.5 Why Not Just Evict Predicted Dead Blocks?

A natural question arises when considering dead block prediction: why not simply evict predicted dead blocks out of the cache instead of evicting LRU blocks into predicted dead blocks? Indeed, this technique has been proposed as an opti-

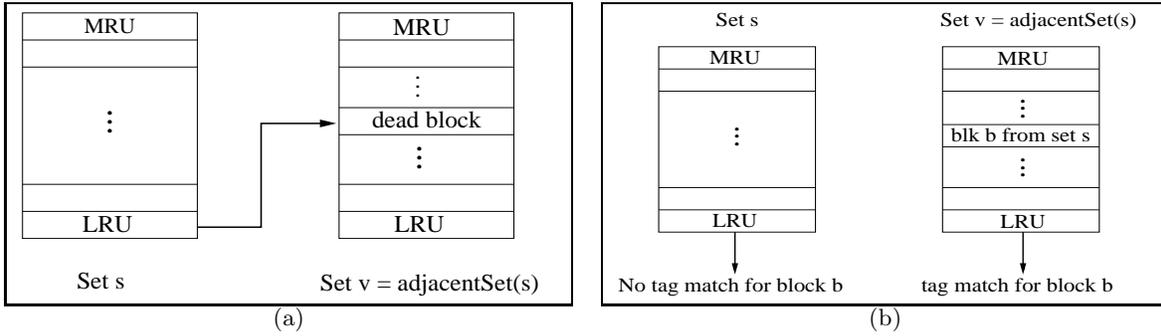


Figure 2: (a) Placing evicted block into an adjacent partner set, and (b) hitting in the virtual victim cache

mization for single-threaded workloads [8]. Note that blocks that are predicted dead and evicted from a set may be cached in the VVC; it seems counterintuitive to replace one dead block with another dead block.

Nevertheless, the VVC does give an advantage over evicting predicted dead blocks for two reasons. The first reason is that the VVC reduces conflict misses by effectively increasing associativity for hot sets. The second reason is because the VVC is robust in the presence of mispredictions. Consider a set S and its adjacent set S' . One of these sets is likely to be more active than the other due to the nonuniform distribution of accesses to sets. The fact that a block b in S is accessed makes it somewhat more likely that S is the more active set, since it was just accessed and S' was not. Suppose b is incorrectly predicted dead. Caching it in S' would be the right move since there are likely to be true dead blocks in S' , while simply discarding b from S would be the wrong move and lead to a miss.

When prediction accuracy is high, replacing dead blocks is a reasonable policy. However, as we will see in Section 6.3, when prediction accuracy suffers because of shared cache contention, replacing dead blocks is disastrous for performance while the VVC provides a performance improvement.

3.2 Implementation Issues

Adjacent sets differ in one bit, bit k . The set adjacent to set index s is s exclusive-ORed with 2^k . A value of $k = 3$ provides good performance, although performance is largely insensitive to the choice of k . Victims replace receiver blocks in the MRU position of the adjacent set and are allowed to be evicted just as any other block in the set. Evicted receiver blocks are not allowed to return to their original sets, i.e., evicted blocks may not “ping-pong” back and forth between adjacent sets.

Each cache block keeps the following additional information: whether or not it is a receiver block (1 bit), whether or not the block is predicted dead (1 bit), and the truncated sum representing the trace for this block (14 bits). The dead block predictor additionally keeps two tables of two-bit saturating counters indexed by traces. The predictor tables consume an additional 2^{14} entries $\times 2$ bit counters $\times 2$ tables = 64 kilobits, or 8 kilobytes.

In the presence of multiple processors with coherent caches, the virtual victim cache can slightly delay snoops since tags must be checked both in the original set and the adjacent set. The VVC can avoid most of the sequential tag checks by associating a bit with each set that is true if a victim from

that set could be cached in another set. Thus, an extra tag check would be unnecessary for most snoops.

4. SKEWED DEAD BLOCK PREDICTOR

In this section we discuss a new dead block predictor based on the reference trace predictor of Lai *et al.* [10] as well as skewed table organizations [23, 14].

4.1 Reference Trace Dead Block Predictor

The reference trace predictor collects a trace of the instructions used to access a particular block. The theory is that, if a sequence of memory instructions to a block leads to the last access of that block, then the same sequence of instructions should lead to the last access of other blocks. The reference trace predictor encodes the path of memory access instructions leading to a memory reference as the truncated sum of the instructions’ addresses. This truncated sum is called a *signature*. Each cache block is associated with a signature that is cleared when that cache block is filled and updated when that block is accessed. The signature is used to access a table of two-bit saturating counters. When a block is accessed, the corresponding counter is decremented and then the signature is updated. When a block is evicted, the counter is incremented. Thus, a counter is only incremented by a signature resulting from the last access to a block.

When a block is accessed and then the signature is updated, the table of counters is consulted. If the counter exceeds a threshold (e.g. 2), then the block is predicted dead. Each cache block stores a single bit prediction. For comparison, we use a 15-bit signature indexing a 32K-entry table of counters. The predictor exclusive-ORs the first 15 bits of each PC with the next 15 bits and adds this quantity to a running 15-bit trace.

The original reference trace predictor of Lai *et al.* uses data addresses as well as instruction addresses, requiring a large table because of the high number of signatures. Subsequent work found that using only instruction addresses was sufficient and allows smaller tables [12]; thus, we use only instruction addresses for all of the predictors in this paper.

4.2 A Skewed Organization

In the original trace-based dead block predictor, a single table is indexed with the signature. For this study, we explore an organization that uses the idea of a skewed organization [23, 14] to reduce the impact of conflicts in the table. The predictor keeps two 16K-entry tables of 2-bit counters,

Table 1: Microarchitectural simulator parameters

Parameter	Configuration
Issue width	4
Reorder Buffer	128 entry
Load/Store Queue	32 entry
Private L1 I-Cache	64KB, 2 way LRU, 64B blocks, 1 cycle hit
Private L1 D-Cache	64KB, 2 way LRU, 64B blocks, 3 cycle hit
Shared L2 Cache	2 MB, 16 way LRU, 64B blocks, 12 cycle hit
Virtual addresses	64 bits
Main Memory	270 cycle
Number of Cores	4

Table 2: Dead block predictor parameters

Parameter	Configuration
Trace encoding	15 bit
Predictor table entries	32,768
Predictor entry	2 bit
Predictor overhead	8KB
Cache overhead	64KB
Total overhead	76 KB

each indexed by a different 14-bit hash of the 15-bit block signature. Each access to the predictor yields two counter values. The sum of these values is used as a confidence that is compared with a threshold; if the threshold is met, then the corresponding block is predicted dead. This organization offers an improvement over the original reference trace predictor because two distinct traces might conflict in one table, but are less likely to conflict in both tables, so the effect of destructive conflicts is reduced.

Figure 3 shows the difference in the design of the original and skewed reference trace predictors.

5. EXPERIMENTAL METHODOLOGY

This section outlines the experimental methodology used in this study.

5.1 Simulation Environment

We use SPEC CPU 2000 and SPEC CPU 2006 benchmarks. We use a cycle-accurate microarchitectural multi-core simulator. This simulator is a heavily modified version of SimpleScalar [3]. This infrastructure enables collecting instructions-per-cycle figures as well as misses per kilo-instruction, dead block predictor accuracy, and cache efficiency.

Table 1 shows the configuration of the simulated machine. Each benchmark is compiled for the Alpha EV6 instruction set. For SPEC CPU 2000, we use the Alpha executables that were at one time available from `simplescalar.com` compiled with DEC C V5.9008, Compaq C++ V6.2-024, and Compaq FORTRAN V5.3-915. For SPEC CPU 2006, we use binaries compiled with the GCC 4.11 compilers for C, C++, and FORTRAN. For most experiments, we model a 16-way set-associative cache to remain consistent with other previous work [9, 12, 16, 18, 17], but Section 6.7 shows that our

technique maintains significant improvement at lower associativities.

We use SimPoint [15] to identify a single two billion instruction characteristic region (i.e. *simpoint*) of each benchmark. For single-threaded workloads, the infrastructure simulates two billion instructions, using the first 500 million to warm microarchitectural structures and reporting the results on the next 1.5 billion instructions.

5.2 Simulating CMP Workloads

We use a multi-core simulator to measure the performance of the virtual victim cache in the presence of multiple threads. The simulator simulates four cores accessing a shared L2 cache. We choose ten combinations of SPEC CPU benchmarks to run simultaneously to represent chip-multiprocessor (CMP) workloads. These combinations are identified in the x -axis of the relevant graphs in Section 6. For the multi-core simulations, we warm the caches for 500 million instructions, then simulate each of the four threads until every thread has committed 250 million instructions. Although some threads continue to execute beyond 250 million instructions to model cache contention for other threads that have not yet reached the limit, we only report statistics based on the first 250 million instructions. We report IPC throughput, i.e., $\sum_{i=1}^n IPC_i$, giving the sum of the individual IPCs for n benchmarks ($n = 4$ in our case), normalized to the baseline LRU cache. Our approach is consistent with methodology from recent papers on caches in multi-core microarchitectures [19, 13].

We choose a memory-intensive subset of the benchmarks based on the following criteria: a benchmark is used if it (1) does not cause an abnormal termination in the baseline sim-outorder simulator for the chosen *simpoint*, and (2) if increasing the size of the L2 cache from 1MB to 2MB results in at least a 5% speedup. Benchmarks that experience negligible improvement from a higher capacity cache are unlikely to be affected positively or negatively by our optimization.

5.3 Dead Block Predictor Details

5.3.1 Accounting for State in the Predictor

The dead block predictor keeps two 16K-entry tables of 2-bit counters. Each cache block includes additional VVC metadata: 1 bit that is true if the block is predicted dead, 1 bit that is true if the block is a receiver block, and 15 bits for the current trace encoding for that block. The overhead of the predictor and VVC metadata is 76KB which is 3.4% of the total 2MB cache space (including both the data and tag arrays). An accounting of the predictor overhead is given in Table 2

5.4 Estimating Dead Block Hit latency

An L2 cache access takes 12 cycles in the simulated environment. CACTI 5.1 [25] estimates that an additional tag match in the adjacent set, made once the initial tag match in the original set has failed, consumes an extra 2 cycles.

An additional sequential tag match latency is simulated for VVC hits. Experiments show that IPC is insensitive to additional L2 hit latency as long as that latency is a small fraction of the miss latency. For instance, negligible change results when pessimistically assuming that VVC hits take double the normal hit latency because a) most accesses hit

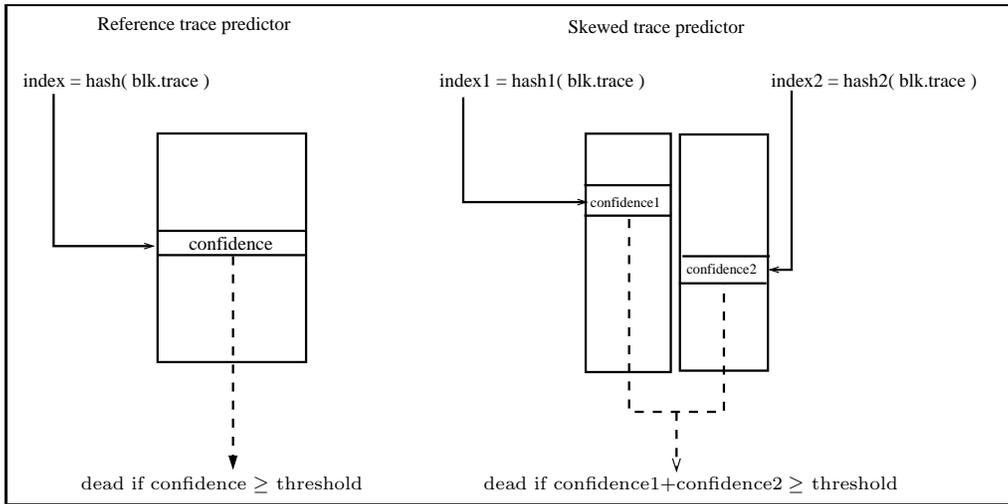


Figure 3: Block diagrams for dead block predictors

in the normal L2 cache, and b) a VVC hit at twice the L2 hit latency avoids a much larger L2 miss latency.

5.5 Measuring Cache Efficiency

Cache efficiency is a statistic defined by Burger *et al.* [4] to quantify the average amount of time blocks in the cache contain live data. Cache efficiency is computed as:

$$E = \frac{\sum_{i=0}^{A \times S - 1} U_i}{N \times A \times S}$$

where N is the total number of cycles executed, A is the number of blocks per set, S is the number of sets in the cache, and U_i is the total number of cycles for which cache block i contains live data, i.e., data that will be referenced again before it is evicted. Thus, cache efficiency is the average of the live cycles for each cache block. The performance simulation infrastructure collects block live times and produces cache efficiency as an output.

6. EXPERIMENTAL RESULTS

In this section we discuss results of our experiments. We investigate the virtual victim cache with a skewed dead block predictor in the context of a baseline 2MB 16-way set associative cache as well as a 2MB fully associative cache and a 2MB cache enhanced with a 64KB victim cache. Note that both the fully associative cache and the 64KB victim cache are infeasible in hardware, requiring 32K entry and 1K entry associative memories, respectively. We choose a 64KB victim cache because it requires approximately the same amount of SRAM, including the tag array, as the extra structures of the VVC.

We also compare the virtual victim cache with two other techniques:

1. Dynamic insertion policy (DIP) using set dueling [17]. This technique, described in detail in Section 2.2, samples from 32 dedicated sets performing block placement into the MRU position as well as 32 dedicated sets placing blocks into the LRU position. The policy incurring the fewest misses is used for the rest of the cache.

2. A replacement policy based on dead block prediction. We use our skewed dead block predictor to drive a replacement policy in which a predicted dead block is chosen as a victim, or the LRU block if there is no predicted dead block. This policy is similar to the technique described by Kharbutli and Solihin [8].

6.1 Reduction in L2 Misses

Figure 4 shows the impact of the VVC on L2 misses per thousand instructions (MPKI). Figure 4(a) shows the raw MPKI values for each benchmark and cache configuration. The average MPKIs for the baseline and fully associative LRU caches are approximately 9.7. The real victim cache yields 8.9 MPKI. The VVC provides an average MPKI of 7.2, an improvement over the baseline of 26% and over the real victim cache of 19%. The VVC provides its best reduction in raw MPKI for `181.mcf`, reducing misses by approximately 30 MPKI. DIP achieves 7.5 MPKI while dead block replacement achieves 7.1 MPKI. As a lower limit on MPKI, we include results for Belady’s MIN optimal replacement policy [2] which results in an average MPKI of 4.9. The VVC achieves an average MPKI only 46% higher than optimal, while LRU is 96% higher than optimal.

Note that it is not enough to compare MPKI for VVC with the other techniques because the VVC incurs a slight overhead when there is a hit to the adjacent set. However, we also report improved performance in a cycle-accurate simulator taking into account this overhead.

6.2 Single-thread IPC Improvement

Reducing cache misses translates into improved performance. Figure 5 shows the instructions-per-cycle rates given by the VVC as well as other techniques. In 11 of the 16 single-threaded workloads, the virtual victim cache outperforms DIP. In a different 11 of the 16 benchmarks, the virtual victim cache outperforms dead block replacement.

Figure 6 shows the speedup computed by dividing the improved IPC by the baseline IPC for each benchmark, as well as the geometric mean speedup. The geometric mean speedup for the VVC is 12.1%, compared with 5.2% for the real victim cache and 7.1% for the fully associative cache, 8.7% for DIP, and 10.3 for dead block replacement.

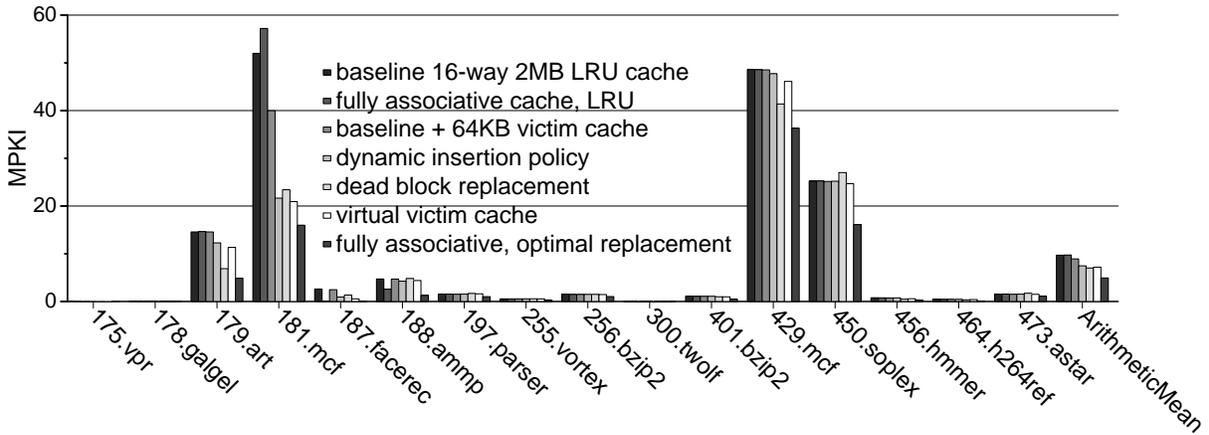


Figure 4: L2 cache misses per thousand instructions

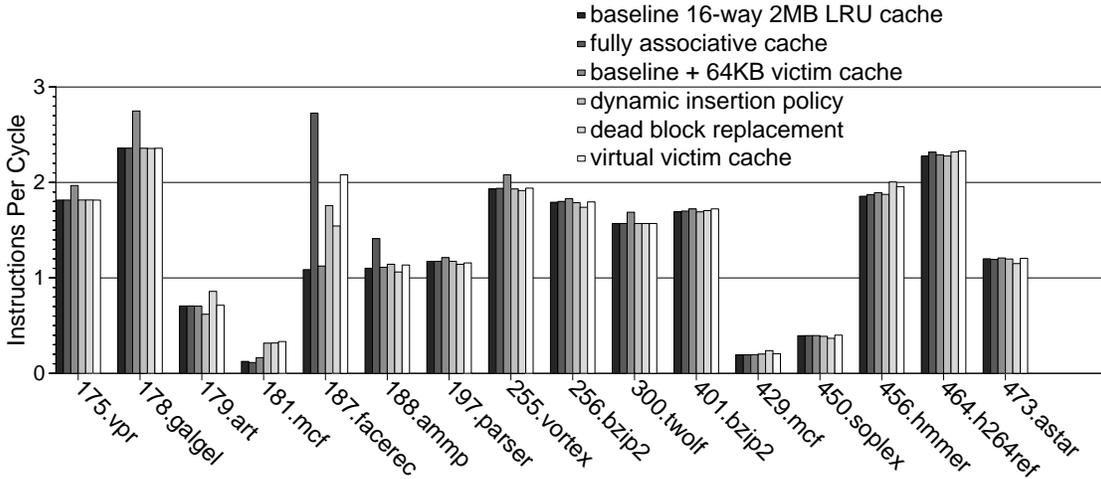


Figure 5: IPC improvement

Two benchmarks in particular, `181.mcf` and `187.facerec`, yield remarkable speedups of 167% and 91%, respectively, while other benchmarks show more modest improvements. No benchmark is significantly slowed down by the VVC; `197.parser` is the worst case in terms of slowdown, with a speedup of -1.8%. Note that dead block replacement slows down several benchmarks, the worst of which is `450.soplex` with a speedup of -7.0%.

Although the difference between the performance of the virtual victim cache and dead block replacement is small, there is a very important reason to prefer the virtual victim cache: for a shared multi-core cache, dead block replacement experiences a significant performance degradation while the virtual victim cache continues to provide a performance improvement as we will see in Section 6.3.

6.3 Prediction Accuracy and Coverage

In this section we quantify the improvement given by our new dead block predictor using a skewed organization derived from a branch predictor design. Mispredictions come in two varieties: false positives and false negatives. False positives are more harmful for most applications of dead block prediction because they falsely allow an optimization to use a live block for some other purpose. False negatives,

while not harmful to performance, limit the potential applicability of the optimization. Thus, we would like to reduce both kinds of errors.

The coverage of a dead block predictor is the number of positive predictions divided by the total number of predictions. If a dead block predictor is consulted on every cache access, then the coverage represents the fraction of cache accesses when the optimization may be applied. Higher coverage means more opportunity for the optimization.

Figure 7 shows the coverage and false positive rates of the old and new predictors. On average, our skewed predictor covers 17.2% of accesses, compared with 18.0% for the Lai *et al.* predictor. However, *all* of the extra coverage attributable to the Lai *et al.* predictor is due to false positive mispredictions. On average, our predictor has a false positive rate of 3.5%, compare with 4.3% for the Lai *et al.* predictor; thus, harmful false positives are significantly reduced.

Figure 8 shows the false positive misprediction rate for each predictor as the hardware budget for the prediction tables varies from 128 bytes to 128KB. We choose the 8KB hardware budget as our predictor because it represents a good trade-off between area and performance, but clearly there is potential to improve accuracy with a larger hardware budget. As mentioned above, at the 8KB hardware budget,

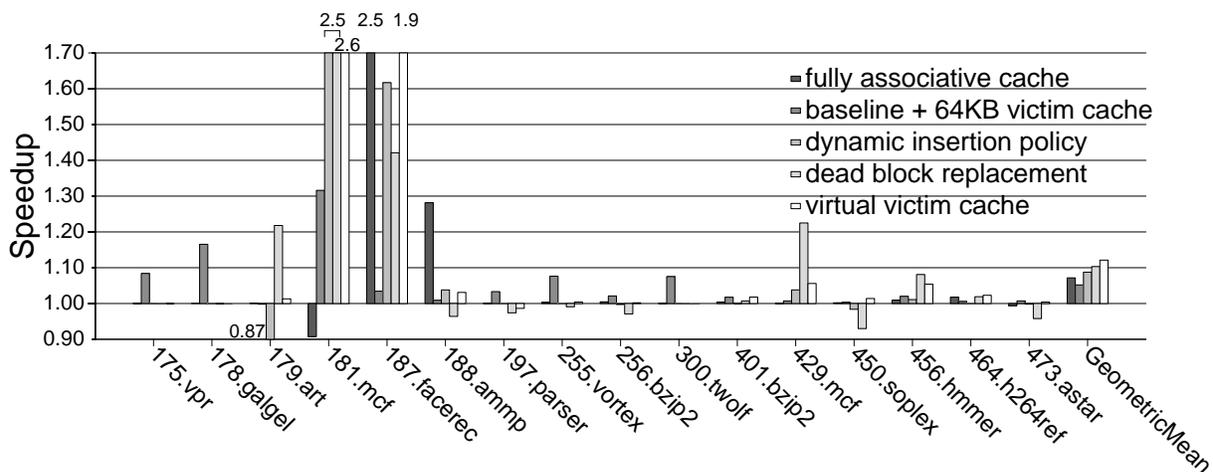


Figure 6: Speedup

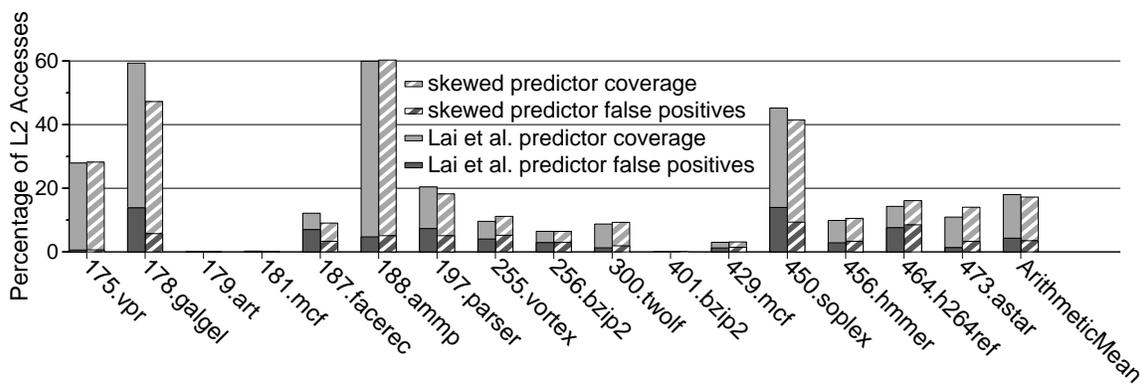


Figure 7: Predictor coverage and false positive rate

the skewed predictor has a false positive misprediction rate of 3.5% compared with 4.3% for the Lai *et al.* predictor. At a 128KB hardware budget, the new predictor has a false positive misprediction rate of 3.4% compared with 4.2% for the Lai *et al.* predictor.

With an 8KB budget, the original Lai *et al.* predictor allowed the VVC to achieve a 5.4% geometric mean speedup over all the benchmarks. The skewed predictor improves this geometric mean speedup to 12.1% (not graphed for space reasons).

We did investigate other dead block predictors such as the reference counting predictors [8] and cache bursts, however neither of these predictors provided any accuracy or performance improvements in our study. The skewed organization resulted in significant improvements in accuracy and performance; we believe that future improvements to dead block predictors will result in improved performance for the VVC as well as other optimizations.

6.4 Methodological Issues

Following reviewers suggestions, we ran new simulations to investigate the potential impact of our methodology. In the first experiment, we use multiple simpoints instead of one single simpoint to compare the baseline cache with the VVC. The VVC yields an 18% improvement in harmonic mean IPC over the baseline and a 6% geometric mean speedup for the single-threaded benchmarks. In the second experi-

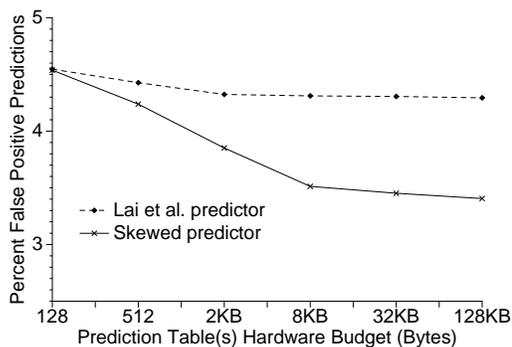


Figure 8: Arithmetic mean predictor false positive mispredictions for a range of hardware budgets

ment, we used a 32-way cache as the baseline, keeping the VVC at 16-way associativity and using the multiple simpoint methodology. We found no difference in geometric mean speedup to the fifth decimal place, i.e., the extra associativity gave no additional benefit for the single-threaded benchmarks.

6.5 CMP Throughput Improvement

We simulate the various cache configurations in a multi-core simulator. Figure 9 shows the normalized IPC through-

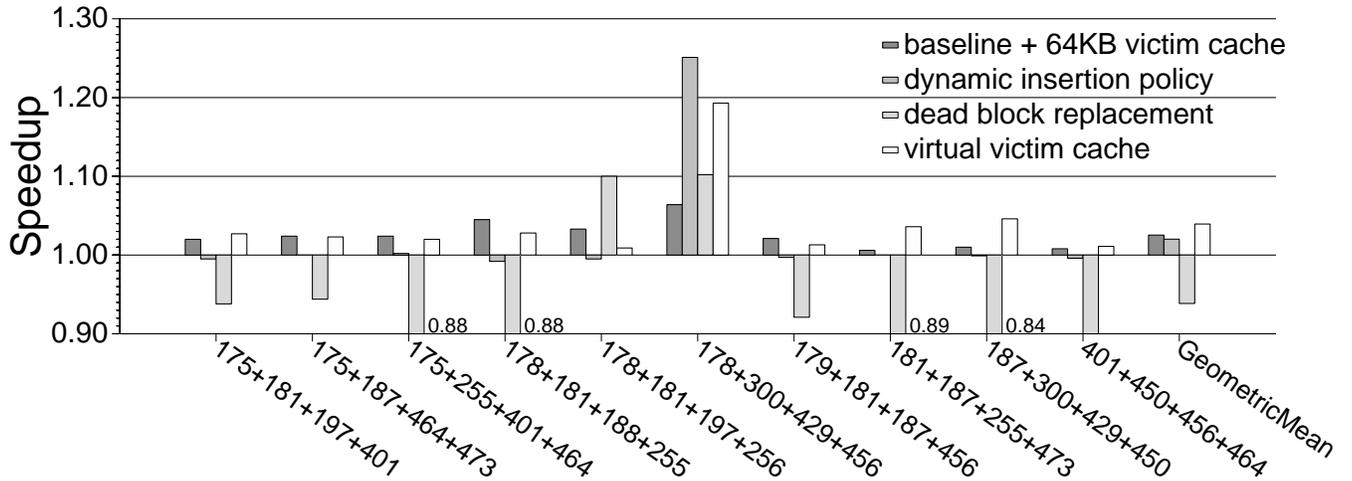


Figure 9: Normalized throughput IPC for multiple threads and a 2MB cache

put for ten combinations of benchmarks representing CMP workloads.

The virtual victim cache achieves a normalized throughput IPC of 1.04, compared with 1.025 for the real victim cache, 1.020 for DIP, and 0.94 for dead block replacement. Although dead block replacement performs well for single-threaded workloads, it performs quite poorly here while the VVC continues to improve performance. We believe that this is because dead blocks become less predictable as several different program access patterns target the same set. The dead block replacement policy can only consider blocks in one set for eviction. If it often makes the wrong decision for hot sets, performance suffers. However, with the VVC, the adjacent set is somewhat less likely to be hot than the original due to the reasoning in Section 3.1.5. Thus, it is more likely that there will be true dead blocks in the adjacent set than in the original set.

Figure 10 illustrates the difference in terms of predictor accuracy for CMP versus single-threaded workloads. The figure shows the percentage of accesses that were incorrectly predicted as the last access to a block, potentially resulting in a live block being replaced. For single-threaded workloads, the average false positive rate is 3.5%, so dead block replacement and the VVC can have similar performance. (A notable exception to this trend is `429.mcf`, which achieves better performance than the VVC despite relatively poor accuracy.) However, for CMP workloads, the false positive rate is 13.0%, resulting in poor performance for the dead block replacement policy but sustaining a benefit for the VVC.

6.6 Improvement in Cache Efficiency

As stated in the introduction, the VVC improves cache efficiency by making sure more cache blocks are live. Figure 11 quantifies this improvement for each CMP benchmark combination and each single-threaded benchmark. The baseline average cache efficiency for the single-threaded workloads is 0.412, compared with 0.523 for the VVC, an improvement of 27%. For the CMP workloads, the baseline average cache efficiency is 0.243, while the VVC achieves an average efficiency of 0.394, an improvement of 62%.

6.7 Reduction in Cache Area

We have shown that the VVC can deliver improved performance by reducing the number of L2 cache misses. Alternately, the VVC can reduce the area consumed by the L2 cache, leading to equivalent performance with reduced power and area requirements. Figure 12 illustrates the ability of the VVC to deliver equivalent performance with a reduced capacity cache. Figure 12(a) shows the average MPKI for the baseline cache and the VVC for a variety of cache sizes obtained by increasing the associativity of the cache from 8 through 16. At a capacity of 1.7MB representing an associativity of 13, the VVC achieves an average MPKI of 9.9, just above the MPKI of the 2MB baseline cache at 9.7. At a capacity of 1.8MB representing an associativity of 14, the VVC outperforms the baseline with an MPKI of 9.1. Figure 12(b) shows how these MPKIs translate into performance. The baseline harmonic mean IPC for a 2MB cache is 0.64, compared with 0.63 for a 1.7MB VVC and 0.68 for a 1.8MB VVC.

At the 1.7MB capacity, the VVC reduces the number of SRAM cells required by 16% including the SRAM cells for data, tags, predictor structures and metadata. At the 1.8MB capacity, the VVC reduces the SRAM cells needed by 9.5%.

6.8 Increase in Tag Array Access

The improved performance of the VVC comes at the expense of extra reads to the tag array. Most accesses hit in the cache; however, every time a tag match fails in a set, the adjacent set must also be searched. In two billion executed instructions, the L2 tag array is read on average 77 million times in the baseline cache and 97 million times in the VVC, an increase of 26%. To put it another way, the number of tag array reads in the baseline cache is 3.9% of the total number of instructions executed, versus 4.9% for the VVC.

7. CONCLUSION AND FUTURE WORK

This paper has explored a cache management strategy that is a combination of block placement, search, and block replacement driven migrating LRU blocks into predicted dead blocks in other sets. The virtual victim cache significantly reduces misses and improves performance for several

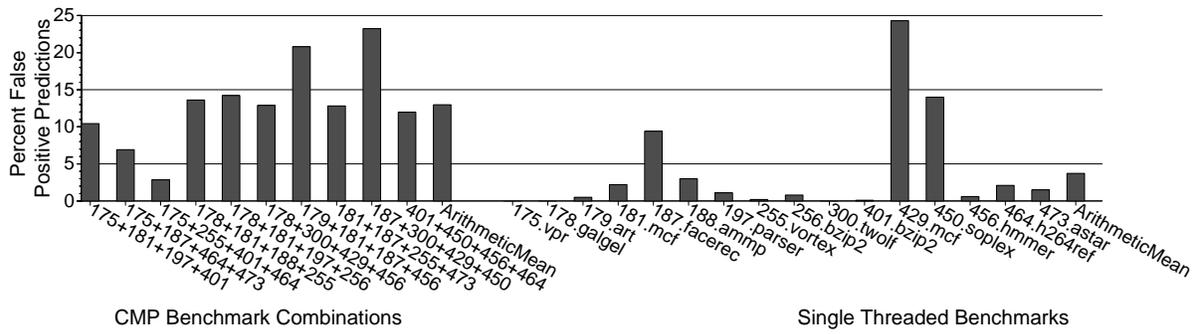


Figure 10: Shared cache contention results in more false positive predictions.

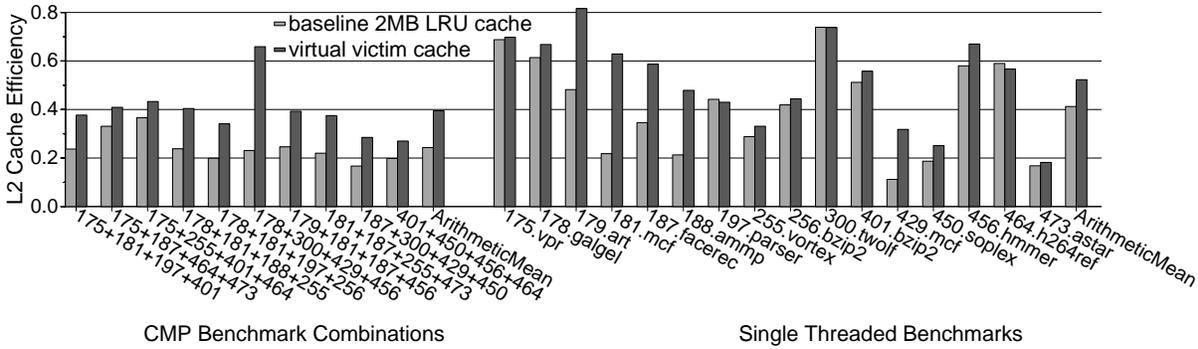


Figure 11: Cache efficiency

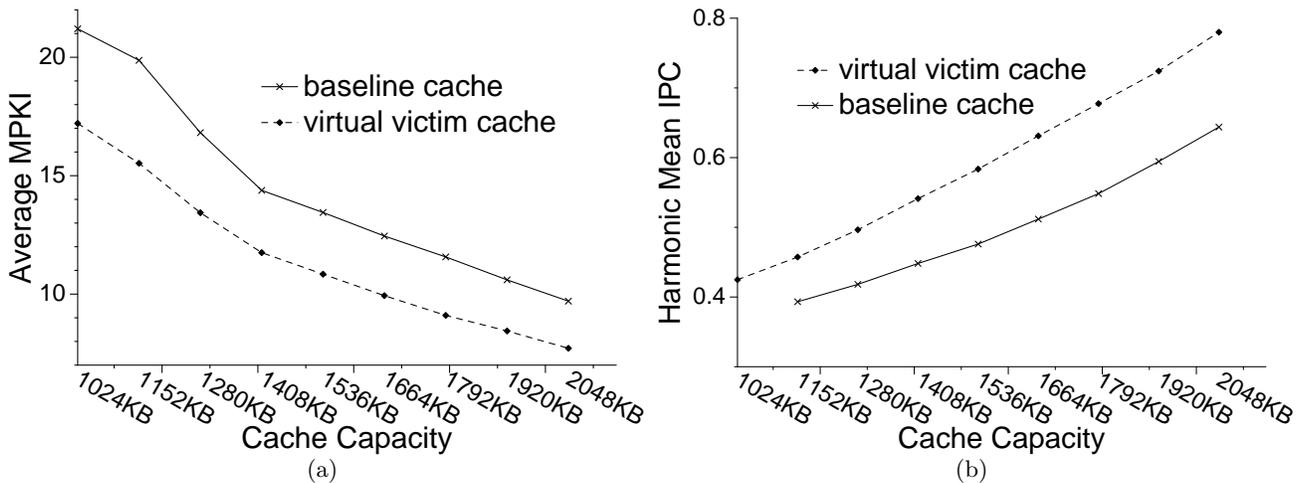


Figure 12: Reduction in cache area

benchmarks, provides modest improvement to most benchmarks, and does not slow down any benchmarks. It significantly improves performance in multi-threaded workloads where a simpler dead block replacement policy yields performance worse than LRU. It does this with a small amount of extra state and a modest increase in the number of accesses to the tag array. Alternately, the virtual victim cache allows reducing the size of the L2 cache while maintaining performance.

We see several future directions for this work. Adapting a skewed organization inspired by branch prediction research has improved dead block predictor accuracy. It might be

possible to adapt other predictor organizations to improve dead block predictor accuracy and coverage. Reducing the number of accesses to the tag array through a more intelligent search strategy could improve the power behavior of the cache. The VVC allows reducing the associativity and size of the cache while maintaining performance, but the potential for reducing the number of sets has not been explored. An adaptive placement policy improves the performance of the VVC, but perhaps other policies would provide additional improvement. So far, the VVC does not distinguish between victims that are likely to be used again and those that are

not. A more discriminating technique could further improve performance by filtering out cold data.

8. ACKNOWLEDGEMENTS

Daniel A. Jiménez and Samira M. Khan are supported by grants from NSF: CCF-0931874 and CRI-0751138.

9. REFERENCES

- [1] J. Abella, A. González, X. Vera, and M. F. P. O’Boyle. Iatac: a smart predictor to turn-off l2 cache lines. *ACM Trans. Archit. Code Optim.*, 2(1):55–77, 2005.
- [2] L. A. Belady. A study of replacement algorithms for a virtual-storage computer. *IBM Systems Journal*, 5(2):78–101, 1966.
- [3] D. Burger and T. M. Austin. The SimpleScalar tool set version 2.0. Technical Report 1342, Computer Sciences Department, University of Wisconsin, June 1997.
- [4] D. Burger, J. R. Goodman, and A. Kagi. The declining effectiveness of dynamic caching for general-purpose microprocessors. *Technical Report 1261*, 1995.
- [5] Z. Hu, S. Kaxiras, and M. Martonosi. Timekeeping in the memory system: predicting and optimizing memory behavior. *SIGARCH Comput. Archit. News*, 30(2):209–220, 2002.
- [6] N. P. Jouppi. Improving direct-mapped cache performance by the addition of a small fully-associative cache and prefetch buffers. *SIGARCH Comput. Archit. News*, 18(3a), 1990.
- [7] G. Keramidas, P. Petoumenos, and S. Kaxiras. Cache replacement based on reuse-distance prediction. In *ICCD*, pages 245–250, 2007.
- [8] M. Kharbutli and Y. Solihin. Counter-based cache replacement and bypassing algorithms. *IEEE Transactions on Computers*, 57(4):433–447, 2008.
- [9] A.-C. Lai and B. Falsafi. Selective, accurate, and timely self-invalidation using last-touch prediction. In *International Symposium on Computer Architecture*, pages 139–148, 2000.
- [10] A.-C. Lai, C. Fide, and B. Falsafi. Dead-block prediction & dead-block correlating prefetchers. *SIGARCH Comput. Archit. News*, 29(2):144–154, 2001.
- [11] A. R. Lebeck and D. A. Wood. Dynamic self-invalidation: reducing coherence overhead in shared-memory multiprocessors. *SIGARCH Comput. Archit. News*, 23(2):48–59, 1995.
- [12] H. Liu, M. Ferdman, J. Huh, and D. Burger. Cache bursts: A new approach for eliminating dead blocks and increasing cache efficiency. In *Proceedings of the IEEE/ACM International Symposium on Microarchitecture*, pages 222–233, Los Alamitos, CA, USA, 2008. IEEE Computer Society.
- [13] G. H. Loh. Extending the effectiveness of 3d-stacked dram caches with an adaptive multi-queue policy. In *MICRO 42: Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 201–212, New York, NY, USA, 2009. ACM.
- [14] P. Michaud, A. Seznec, and R. Uhlig. Trading conflict and capacity aliasing in conditional branch predictors. In *Proceedings of the 24th International Symposium on Computer Architecture*, pages 292–303, June 1997.
- [15] E. Perelman, G. Hamerly, M. Van Biesbrouck, T. Sherwood, and B. Calder. Using simpoint for accurate and efficient simulation. *SIGMETRICS Perform. Eval. Rev.*, 31(1):318–319, 2003.
- [16] M. K. Qureshi, A. Jaleel, Y. N. Patt, S. C. S. Jr., and J. S. Emer. Adaptive insertion policies for high performance caching. In *34th International Symposium on Computer Architecture (ISCA 2007), June 9-13, 2007, San Diego, California, USA*. ACM, 2007.
- [17] M. K. Qureshi, D. N. Lynch, O. Mutlu, and Y. N. Patt. A case for mlp-aware cache replacement. In *ISCA ’06: Proceedings of the 33rd annual international symposium on Computer Architecture*, pages 167–178, Washington, DC, USA, 2006. IEEE Computer Society.
- [18] M. K. Qureshi, D. N. Lynch, O. Mutlu, and Y. N. Patt. A case for mlp-aware cache replacement. In *ISCA ’06: Proceedings of the 33rd annual international symposium on Computer Architecture*, pages 167–178, Washington, DC, USA, 2006. IEEE Computer Society.
- [19] M. K. Qureshi and Y. N. Patt. Utility-based cache partitioning: A low-overhead, high-performance, runtime mechanism to partition shared caches. In *MICRO 39: Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 423–432, Washington, DC, USA, 2006. IEEE Computer Society.
- [20] M. K. Qureshi, D. Thompson, and Y. N. Patt. The v-way cache: Demand based associativity via global replacement. *SIGARCH Comput. Archit. News*, 33(2):544–555, 2005.
- [21] D. Rolán, B. B. Fraguera, and R. Doallo. Adaptive line placement with the set balancing cache. In *MICRO 42: Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 529–540, New York, NY, USA, 2009. ACM.
- [22] J. B. Sartor, S. Venkiteswaran, K. S. McKinley, and Z. Wang. Cooperative caching with keep-me and evict-me. *Annual Workshop on Interaction between Compilers and Computer Architecture*, 0:46–57, 2005.
- [23] A. Seznec. A case for two-way skewed-associative caches. In *ISCA ’93: Proceedings of the 20th annual international symposium on Computer architecture*, pages 169–178, New York, NY, USA, 1993.
- [24] S. Somogyi, T. F. Wenisch, N. Hardavellas, J. Kim, A. Ailamaki, and B. Falsafi. Memory coherence activity prediction in commercial workloads. In *WMPPI ’04: Proceedings of the 3rd workshop on Memory performance issues*, pages 37–45, New York, NY, USA, 2004. ACM.
- [25] S. Thoziyoor, N. Muralimanohar, J. H. Ahn, and N. P. Jouppi. Cacti 5.1. *Technical report HPL-2008-20*, 2008.
- [26] Z. Wang, K. S. McKinley, A. L. Rosenberg, and C. C. Weems. Using the compiler to improve cache replacement decisions. In *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques*, page 199, Los Alamitos, CA, USA, 2002. IEEE Computer Society.