Software-assisted Cache Replacement Mechanisms for Embedded Systems

Prabhat Jain
prabhat@mit.edu

Srinivas Devadas
devadas@mit.edu

Daniel Engels
dwe@mit.edu

Larry Rudolph
rudolph@lcs.mit.edu

Laboratory for Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139

Abstract

We address the problem of improving cache predictability and performance in embedded systems through the use of software-assisted replacement mechanisms. These mechanisms require additional software controlled state information that affects the cache replacement decision. Software instructions allow a program to kill a particular cache element, i.e., effectively make the element the least recently used element, or keep that cache element, i.e., the element will never be evicted.

We prove basic theorems that provide conditions under which kill and keep instructions can be inserted into program code, such that the resulting performance is guaranteed to be as good as or better than the original program run using the standard LRU policy. We developed a compiler algorithm based on the theoretical results that, given an arbitrary program, determines when to perform software-assisted replacement, i.e., when to insert either a kill or keep instruction. Empirical evidence is provided that shows that performance and predictability (worst-case performance) can be improved for many programs.

1 Introduction

On-chip memory, in the form of cache, scratchpad SRAM, (and more recently) embedded DRAM or some combination of the three, is ubiquitous in programmable embedded systems to support software and to provide an interface between hardware and software. Most systems have both cache and scratchpad memory on-chip since each addresses a different need. Caches are transparent to software since they are accessed through the same address space as the larger backing storage. They often improve overall software performance but are unpredictable. Although the cache replacement hardware is known, predicting its performance depends on accurately predicting past and future reference patterns. Of course, these reference patterns vary depending on input data. Scratchpad memory is addressed via an independent address space and thus must be managed explicitly by software, oftentimes a complex and cumbersome problem, but provides absolutely predictable performance. Thus, even though a pure cache system may perform better overall, scratchpad memories are necessary to guarantee that critical performance metrics are always met.

Of course, both caches and scratchpad memories should be available to embedded systems so that the appropriate memory structure can be used in each instance. A static division, however, is guaranteed to be suboptimal as different applications have different requirements. Previous research has shown that even within a single application, dynamically varying the partitioning between cache and scratchpad memory can significantly improve performance [12].

One important aspect to cache design is the choice of the replacement strategy, that controls which cache line to evict from the cache when a new line is brought in. The most commonly used replacement strategy is the Least Recently Used (LRU) replacement strategy, where the cache line that was least recently used is evicted. It is known, however, that LRU does not perform well in many situations, including timeshared systems where multiple processes use the same cache and when there is streaming data in applications. Additionally, the LRU policy often performs poorly for applications in which the cache memory requirements and memory access patterns change during execution. Furthermore, while caches improve average performance, they can cause unpredictable performance. Most cache replacement policies, including LRU, do not provide mechanisms to increase predictability (worst-case performance), making them unsuited for many real-time embedded system applications.

In this paper, we address the problem of improving cache predictability (worst-case performance) and performance through the use of software-assisted replacement mechanisms. The basic mechanism we consider is an augmentation of the least recently used (LRU) replacement method, where additional state in the cache affects the replacement decision. Software can kill a cache element, i.e., effectively make the element the least recently used element, or keep a cache element, i.e., the element will never be evicted from the cache. We consider different variations of these cache kill and keep instructions in this paper.

Our contributions are twofold. First, we provide a theoretical foundation for the development of program analysis and transformation techniques that can automatically add kill and keep instructions to a program. In particular, we prove basic theorems that provide conditions under which kill and keep instructions can be inserted into program code, such that the resulting performance, measured as the hit rate, is guaranteed to be as good as or better than the original program run using the standard LRU pol-

icy. Second, we develop a compiler algorithm based on this theory that, given an arbitrary program, determines when to perform software-assisted replacement, i.e., when to insert either a kill or keep instruction. Empirical evidence is provided that shows that performance and predictability (worst-case performance) can be improved for many programs.

The remainder of the paper is organized as follows. In Section 2, we describe related work in software-controlled caches. In Section 3, we describe our overall strategy for performance improvement. We present our theoretical results in Section 4 that provide the foundation for the compiler algorithm described in Section 5. Preliminary experimental results are presented in Section 6. We provide conclusions and discuss ongoing work in Section 7.

2   Related Work

2.1   Cache Management

Some current microprocessors have cache management instructions that can flush or clean a given cache line, prefetch a line or zero out a given line [9, 11]. Other processors permit cache line locking within the cache, essentially removing those cache lines as candidates to be replaced [3, 4]. Explicit cache management mechanisms have been introduced into certain processor instruction sets, giving those processors the ability to limit pollution. One such example is the Compaq Alpha 21264 [5] where the new load/store instructions minimize pollution by invalidating the cache-line after it is used.

In [8] the use of cache line locking and release instructions is suggested based on the frequency of usage of the elements in the cache lines. In [15] some modified LRU replacement policies have been proposed to improve the second-level cache behavior that look at the temporal locality of the cache lines either in an off-line analysis or with the help of some hardware. In [13], active management of data caches by exploiting the reuse information is discussed along with the active block allocation schemes. In [6], policies in the range of Least Recently Used and Least Frequently Used are discussed.

Our work differs from previous work in that we provide hit rate guarantees when our algorithm is used to insert cache control instructions. Further, the theoretical results that we provide can be used as a basis for developing a varied set of methods for automatic cache control instruction insertion.

2.2   Memory Exploration in Embedded Systems

Cache memory issues have been studied in the context of embedded systems. McFarling presents techniques of code placement in main memory to maximize instruction cache hit ratio [10, 14]. A model for partitioning an instruction cache among multiple processes has been presented [7].

Panda, Dutt and Nicolau present techniques for partitioning on-chip memory into scratchpad memory and cache [12]. The presented algorithm assumes a fixed amount of scratchpad memory and a fixed-size cache, identifies critical variables and assigns them to scratchpad memory. The algorithm can be run repeatedly to find the optimum performance point.

A technique to dynamically partition a cache using column caching was presented in [2]. While column caching can improve predictability for multitasking, it is less effective for single processes. Column caching requires significant cache redesign.

3   Overall Strategy

To use caches more efficiently, application-specific information should be incorporated into the cache line replacement decisions. Program analysis or trace analysis gives indications about future variable accesses that can be used to augment the LRU replacement policy with cache kill or keep instructions. These cache instructions are implemented with some additional cache replacement logic, state and tables. These instructions modify the cache replacement state and the tables. Changing the replacement state influences the replacement policy. A variety of cache control instructions with differing hardware requirements can be used.

3.1   Cache Control Instructions

We consider two forms of cache control instructions: (1) modified load/store instructions that contain the necessary cache control information (2) separate cache control instructions that contain only the cache control information. We discuss the kill, conditional kill, and keep control instructions and their hardware and software requirements.

3.2   Kill Instruction

This form of a kill instruction is a load-store instruction with the kill hint information as part of the instruction. The kill instruction allows a cache line associated with the access to be augmented with a cache line kill state. This additional kill state is used along with the LRU information to choose a cache line other than the LRU cache line for replacement. This instruction provides a mechanism for replacement of data earlier than it would be possible in the LRU policy. It can be used for references that result in the last accesses of the array elements or data structures or for references whose accessed data reuse time is such that early replacement is likely to benefit the overall performance.

3.3   Conditional Kill Instruction

This form of a kill instruction is a load-store instruction with the kill hint information as part of the instruction. The kill hint information contains the condition(s) for the cache line kill state to be updated. We consider the cache line *offset* condition which specifies the cache line offset(s) as a condition. A cache line kill state is updated only if an access generated by the kill instruction satisfies the cache line offset condition. For example, consider a cache line size of 4 words $(0, ..., 3)$, an array A, and a reference $A[i]$. To set the kill state of a cache line when the reference $A[i]$ accesses the fourth word of the cache line, the load-store instruction corresponding to the reference $A[i]$ would have the cache line offset

value 3 as the condition. So, whenever the reference $A[i]$ accesses the fourth word of a cache line, the kill state of that cache line would be set.

3.4 Keep Instruction

This form of a keep instruction is a load-store instruction with the keep hint information as part of the instruction. The keep instruction allows a cache line associated with the access to be augmented with a cache line keep state. This additional keep state provides a mechanism to keep a cache line in the cache longer than it would otherwise be kept in the cache with the LRU replacement policy. The keep state is used along with the LRU information to choose a cache line other than the LRU cache line for replacement. It can be used to keep the time-critical data in the cache for a desired period of time. We consider the use of the keep state as a *flexible keep* state that does not require a release instruction. If the keep state is used as a *fixed keep*, then a corresponding release instruction may be needed.

3.5 Hardware Cost

The use of the above instructions requires modification to the replacement logic to take into account the additional cache line states for replacement decisions. The Kill, Conditional Kill, and Keep instructions described above require only one bit of additional state per cache line in the cache. The Conditional Kill instruction requires a small amount of logic for offset matching.

3.6 Software Cost

The software cost for the above instructions is in the form of the additional flavors of load-store instructions with kill hint, keep hint, and kill with offset condition. The use of the flavors of Kill, Conditional Kill, and Keep instructions does not result in additional accesses in the instruction or data stream during program execution.

4 Theoretical Results

We present theoretical results for the replacement mechanisms that use the additional states to keep or kill the cache lines. We show the conditions under which the replacement mechanisms with the kill and keep states are guaranteed to perform better than the LRU policy.

4.1 Kill+LRU Replacement Policy

In this replacement policy, each element in the cache has an additional one-bit state (K_l) called the kill state associated with it. The K_l bit can be set under software or hardware control. On a hit the elements in the cache are reordered along with their K_l bits the same way as in an LRU policy. On a miss, instead of replacing the LRU element in the cache, an element with its K_l bit set is chosen to be replaced and the new element is placed at the most recently used position and the other elements are reordered as necessary.

We consider two variations of this replacement policy to choose an element with the K_l bit set for replacement: (1) the least recent element that has its K_l bit set is chosen to be replaced; (2) the most recent element that has its K_l bit set is chosen to be replaced. The K_l bit is reset when there is a hit on an element with its K_l bit set unless the current access sets the K_l bit. We assume that the K_l bit is changed – set or reset – for an element upon an access to that element. The access to the element has an associated hint that determines the K_l bit after the access and the access does not affect the K_l bit of other elements in the cache.

Definitions: For a fully-associative cache C with associativity m, the cache state is an ordered set of elements. Let the elements in the cache have a position number in the range $[1, ..., m]$ that indicates the position of the element in the cache. Let $pos(e)$, $1 \leq pos(e) \leq m$ indicate the position of the element e in the ordered set. If $pos(e) = 1$, then e is the most recently used element in the cache. If $pos(e) = m$, then the element e is the least recently used element in the cache. Let $C(LRU, t)$ indicate the cache state C at time t when using the LRU replacement policy. Let $C(KIL, t)$ indicate the cache state C at time t when using the Kill+LRU policy. Let X and Y be sets of elements and let X_0 and Y_0 indicate the subsets of X and Y respectively with K_l bit reset. Let the relation $X \preceq_0 Y$ indicate that the $X_0 \subseteq Y_0$ and the order of common elements ($X_0 \cap Y_0$) in X_0 and Y_0 is the same. Let X_1 and Y_1 indicate the subsets of X and Y respectively with K_l bit set. Let the relation $X \preceq_1 Y$ indicate that $X_1 \subseteq Y_1$ and the order of common elements ($X_1 \cap Y_1$) in X_1 and Y_1 is the same. Let d indicate the number of distinct elements between the access of an element e at time t_1 and the next access of the element e at time t_2.

Lemma 1 *If the condition $d \geq m$ is satisfied, then the access of e at t_2 would result in a miss in the LRU policy.*

Proof: On every access to a distinct element, the element e moves by one position towards the LRU position m. So, after $m - 1$ distinct element accesses, the element e reaches the LRU position m. At this time, the next distinct element access replaces e. Since $d \geq m$, the element e is replaced before its next access, therefore the access of e at time t_2 would result in a miss. ∎

Lemma 2 *The set of elements with K_l bit set in $C(KIL, t) \preceq_1 C(LRU, t)$ at any time t.*

Proof: The proof is based on induction on the cache states $C(KIL, t)$ and $C(LRU, t)$. We can show that after every access to the cache (hit or miss), the new cache states $C(KIL, t+1)$ and $C(LRU, t+1)$ maintain the relation $\preceq_1$ for the elements with the K_l bit set. Please refer to the Appendix A.1 for a detailed proof. ∎

We show the proof of the theorem below for variation (1); the proof for variation (2) is similar.

Theorem 1 *For a fully associative cache with associativity m if the K_l bit for any element e is set upon an access at time t_1 only if the number of distinct elements d between the access at time t_1*

and the next access of the element e at time t_2 is such that $d \geq m$, then the Kill+LRU policy variation (1) is as good as or better than LRU.

Proof: The proof is based on induction on the cache states $C(LRU, t)$ and $C(KIL, t)$. We can show that after every access to the cache, the new cache states $C(LRU, t+1)$ and $C(KIL, t+1)$ maintain the $\preceq_0$ relation for the elements with the K_l bit reset. In addition, the new cache states $C(KIL, t+1)$ and $C(LRU, t+1)$ maintain the relation $\preceq_1$ based on Lemma 2. Therefore, every access hit in $C(LRU, t)$ implies a hit in $C(KIL, t)$ for any time t. Please refer to Appendix A.2 for a detailed proof. ∎

4.2 Kill+Keep+LRU Replacement Policy

In this replacement policy, each element in the cache has two additional states associated with it. One is called a kill state represented by a K_l bit and the other is called a keep state represented by a K_p bit. The K_l and K_p bits cannot both be 1 for any element in the cache at any time. The K_l and K_p bits can be set under software or hardware control. On a hit the elements in the cache are reordered along with their K_l and K_p bits the same way as in an LRU policy. On a miss, if there is an element with the K_p bit set at the LRU position, then instead of replacing this LRU element in the cache, the most recent element with the K_l bit set is chosen to be replaced by the element at the LRU position (to give the element with the K_p bit set the most number of accesses before it reaches the LRU position again) and all the elements are moved to bring the new element at the most recently used position. On a miss, if the K_p bit is 0 for the element at the LRU position, then the elements in the cache are reordered along with their K_l and K_p bits in the same way as in an LRU policy. There are two variations of this policy: (a) *Flexible Keep:* On a miss, if there is an element at the LRU position with the K_p bit set and if there is no element with the K_l bit set, then replace the LRU element (b) *Fixed Keep:* On a miss, if there is an element at the LRU position with the K_p bit set and if there is no element with the K_l bit set, then replace the least recent element with its K_p bit equal to 0.

Theorem 2 *(a) The* Flexible Keep *variation of the Kill + Keep + LRU policy is as good as or better than the LRU policy. (b) Whenever there is an element at the LRU position with the K_p bit set if there is also a different element with the K_l bit set, then the* Fixed Keep *variation of the Kill + Keep + LRU policy is as good as or better than LRU.*

Proof: We just give a sketch of the proof here, since the cases are similar to the ones in the Kill+LRU Theorem. We assume a fully-associative cache C with associativity m. Let $C(KKL, t)$ indicate the cache state C at time t when using the Kill+Keep+LRU policy. A different case from the Kill+LRU theorem is where the current access of an element e results in a miss in $C(KKL, t)$ and the element x at the LRU position has its K_p bit set.

Consider the *Flexible Keep* variation of the Kill+Keep+LRU policy. If there is no element with the K_l bit set, then the element x is replaced and the case is similar to the Kill+LRU policy. If

there is at least one element with the K_l bit set in $C(KKL, t)$, let the most recent element with its K_l bit set is y. Let the cache state $C(KKL, t) = \{L_1, y, L_2, x\}$. The new state is $C(KKL, t+1) = \{e, L_1, x, L_2\}$. There is no change in the order of the elements in L_1 and L_2, so the relationship $C(LRU, t+1) \preceq_0 C(KKL, t+1)$ holds for the induction step. This implies the statement of Theorem 2(a).

Consider the *Fixed Keep* variation of the Kill+Keep+LRU policy. If there is at least one element with the K_l bit set in $C(KKL, t)$, let the most recent element with its K_l bit set be y. Let the cache state $C(KKL, t) = \{L_1, y, L_2, x\}$. The new state is $C(KKL, t+1) = \{e, L_1, x, L_2\}$. There is no change in the order of the elements in L_1 and L_2, so the relationship $C(LRU, t+1) \preceq_0 C(KKL, t+1)$ holds for the induction step. This implies the statement of Theorem 2(b). ∎

4.3 Set-Associative Caches

Theorem 1 and Theorem 2 can be generalized to set-associative caches.

Theorem 3 *For a set-associative cache with associativity m if the K_l bit for any element e mapping to a cache-set i is set upon an access at time t_1 only if the number of distinct elements d, mapping to the same cache-set i as e between the access at time t_1 and the next access of the element e at time t_2, is such that $d \geq m$, then the Kill+LRU policy variation (1) is as good as or better than LRU.*

Proof: Let the number of sets in a set-associative cache C be s. Every element maps to a particular cache set. After an access to the element e that maps to cache set i, the cache state for the cache sets 0 to $i-1$ and $i+1$ to $s-1$ remains unchanged. The cache set i is a fully-associative cache with associativity m. So, using Theorem 1, the Kill+LRU policy variation (1) is as good as or better than LRU for the cache set i. This implies the statement of Theorem 3. ∎

Theorem 4 *(a) The* Flexible Keep *variation of the Kill + Keep + LRU policy is as good as or better than the LRU policy. (b) Whenever there is an element at the LRU position with the K_p bit set in a cache-set i, if there is a different element with the K_l bit set in the cache-set i, then the* Fixed Keep *variation of the Kill + Keep + LRU policy is as good as or better than LRU.*

Proof: We omit the proof of this theorem because it is similar to Theorem 3. ∎

5 Algorithm

The theoretical results presented in the previous section can be used directly if we can determine or estimate the number of distinct memory references d between two given accesses to a variable or data structure. For any pair of accesses, we do not, necessarily, have to determine d precisely, but rather we need to determine if d is less than m, where m is the associativity of the cache. We will define d to be ∞ for the last access to a variable.

We use a compiler-based static analysis strategy. In the case where d is hard to estimate due to conditionals in the program or due to incomplete data layout information, a lower bound on d can be used to guarantee performance better than LRU.

The strategy has to produce the kill and keep hints that can be incorporated into the program source code or used during the compilation phase to insert appropriate kill and keep instructions into the generated code. In the sequel, a reference corresponds to the source code expression or a load/store instruction that may result in different accesses. For example, given an array A, the expression $A[i]$ in the program is a reference that would generate different accesses to the array if $A[i]$ is part of a loop that is iterated multiple times.

Compiler Algorithm

1. Perform life-time analysis on variables in the program, and identify the last use of all variables in the program.

2. Choose variables as kill candidates that are local to a procedure, or shared in a small number of procedures. Also choose global or local variables that have a relatively short life-time and variables that are accessed infrequently.

3. Determine a lower bound on d between each adjacent pair of references of kill candidate variables.

   In the case of a set-associative cache, we require information about what cache set and block each reference is mapped to. For statically-allocated variables, the layout of variables assumed by the compiler along with the sizes of the variables accessed along each (control-flow) path is used to determine a lower bound on d. For large dynamically-allocated variables such as arrays, allocated in a contiguous region of memory, parts of the array will map to all the sets of the cache. When these arrays are accessed between the accesses to a kill candidate variable, we use the size of the arrays and the number and type of accesses to determine how many times, if any, the set corresponding to the kill candidate variable is touched by an array access. We ignore any intermediate accesses to small dynamically-allocated variables, though we might generate kills for these variables. Thus, given a pair of accesses to a kill candidate variable, we determine how many accesses to other variables fall into the same set on each control flow path. The minimum number of distinct accesses over all control flow paths is the value used for d.

   For a fully associative cache, the lower bound on d is the minimum number of distinct block references along any (control-flow) path from the first reference to the second. We do not have to deal with the mapping to sets.

4. For kill candidate variables, if any adjacent pair of references has $d \geq m$, associate a kill instruction with the first reference.

5. Kill instructions are generated after the last access of all variables (since $d = \infty$). We use conditional Kill instructions in this case.

6. Identify keep references and associate keep instructions with such references. The keep variable references can be critical variables that are accessed frequently and/or have a long lifetime, e.g., a group of index variables that are used in a loop. The keep variable references can also be a set of variables that are accessed by the program at regular intervals, but the interval is such that these variables would be evicted from the cache. The keep variables can also be specified by the user.

In order to guarantee performance as good or better than LRU, check to see that at each point in the program, the number of killed references in the cache is equal to or greater than the number of keep references for the fixed keep method. A limit can be specified for the maximum number of keep variables per cache set, and the hardware can enforce the limit. We do not need to check the number of killed references for the flexible keep method. Amongst the keep variables, the variables that are referenced more often are given higher priority for selection.

6 Experimental Results and Analysis

For our experiments, we compiled the benchmarks for a MIPS-like PISA processor instruction set used by the SimpleScalar 3.0 [1] tool set. We generated traces for the benchmarks using SimpleScalar 3.0 [1] and chose sub-traces (instruction + data) from the middle of the generated trace. We used a hierarchical cache simulator, *hiercache*, to simulate the trace assuming an L1 cache and memory. In our experiments, we measured the L1 hit rate and the performance of some of the Spec95 benchmarks for various replacement policies.

We describe our experiments using the Spec95 Swim benchmark as an example. We chose a set of arrays as candidates for the kill and keep related experiments. The arrays we considered were *u, v, p, unew, vnew, pnew, uold, vold, pold, cu, cv, z, h, psi*. These variables constitute 29.15% of the total accesses. We did the experiment with different associativities of $2, 4, 8$ and cache line sizes of $2, 4, 8$ words and a cache size 16K bytes. The overall L1 hit rate results for the Swim benchmark are shown in Figure 1.

In Figure 1, the x-axis *a,b* indicates the associativity and the cache line size in words. The column labeled *LRU* shows the hit rate over all accesses (not just the array accesses) with the LRU policy. The column labeled *Kill* shows the hit rate for the Kill+LRU replacement policy. The columns labeled *KK1, KK2, KK3* show the hit rate for the Kill+Keep+LRU replacement policy with the *Flexible Keep* variation. In *KK1*, the array variables *unew, vnew, pnew* are chosen as the keep candidates. In *KK2*, the array variables *uold, vold, pold* are chosen as the keep candidates. In *KK3*, only the array variable *unew* is chosen as the keep candidate. The hit rates of the variables of interest for an associativity of 4 and a cache line size of 8 words are shown in Figure 2 for the same columns as described above. The modified program with cache control instructions does not have any more instruction or data accesses than the original program. In Figure 2, the columns *%Imprv* show the percentage improvement in hit rate for the variables over the LRU policy.

Figure 1: **Overall Hit Rates for the Spec95 Swim Benchmark (L1 Cache Size 16 KBytes)**

Vars	#Ref	#Kill	#Cond Kill
u	28	6	10
v	28	5	12
p	24	2	11
unew	13	6	6
vnew	13	4	8
pnew	13	4	8
uold	13	5	7
vold	13	5	7
pold	13	5	7
cu	15	4	7
cv	15	2	10
z	13	5	5
h	13	4	5
psi	5	0	2

Figure 3: **Number of Kill Instructions for the array variables in the Spec95 Swim benchmark for L1 cache size 16 KB, associativity 4, and cache line size 8 words**

Figure 3 shows the number of Kill (labeled as #Kill) and Conditional Kill (labeled as #Cond Kill) instructions generated corresponding the number of references (labeled as #Ref) for the array variables of the Spec95 Swim benchmark.

The results show that the performance improves in some cases with the use of our software-assisted replacement mechanisms that use kill and keep instructions. The results in Figure 2 show that the hit rates associated with particular variables can be improved very significantly using our method. The bold numbers in the KK1, KK2, and KK3 columns in Figure 2 indicate the hit rate of the variables that were the only keep variables for these columns. Choosing a particular variable and applying our method can result in an substantial improvement in hit rate and therefore performance for the code fragments where the variable is accessed. For example, for a variable vnew, the hit rate for LRU was 37.37, but we could improve it to 75.86 using the Keep method. This is particularly relevant when we need to meet real-time deadlines in embedded processor systems across code fragments, rather than optimizing performance across the entire program.

Figure 4(a) shows the overall hit rate and performance for some Spec95 benchmarks. We show hit rate and performance for a single-issue pipelined processor that stalls on a cache miss. The number of cycles are calculated by assuming 1 cycle for instruction accesses and 1 cycle for on-chip memory and a 10 cycle latency for off-chip memory. The columns show hit rate and number of cycles assuming 10 cycles for off-chip memory access and 1 cycle for on-chip memory access. The last column shows the performance improvement of Kill+Keep over LRU. Figure 4(b) shows the overall worst-case hit rate and performance for the same Spec95 benchmarks. The worst-case hit rate is measured over 10 sets of input data. The columns show hit rate and number of cycles assuming 10 cycles for off-chip memory access and 1 cycle for on-chip memory access. The last column shows the performance improvement of Kill+Keep over LRU.

The programs that do not have much temporal reuse of its data (e.g., some integer benchmarks) do not benefit from kill+LRU replacement in terms of the hit rate improvement, but if the same programs have some data that can benefit by keeping some variables in the cache, then the Kill+Keep strategy can help in improving the hit rates for the keep variables without degrading performance.

7   Conclusions and Ongoing Work

The main contributions of our work are in laying theoretical groundwork for the development of techniques for inserting cache control instructions into programs, and the development of an algorithmic compiler analysis method to automatically insert cache control instructions for improved performance. The compiler analysis method uses the information derived from an analysis of the program during program compilation. This method guarantees that the performance of the modified program is at least as good as the performance of the original program under the LRU replacement policy, when performance is measured in terms of hit rate. The use of cache control instructions improves the performance of a program when executed using a cache, and it improves the predictability of the program by improving its worst-case performance over a range of input data. The increased predictability afforded by cache control instructions makes caches more amenable for use in real-time embedded systems. Our preliminary experiments show that significant improvements are possible using our technique. Many different variants of this technique are possible, and we are currently exploring these variants.

Acknowledgments

Vars	LRU	Kill	%Imprv	KK1	%Imprv	KK2	%Imprv	KK3	%Imprv
u	85.94	88.57	3.06	86.47	0.62	86.32	0.44	86.67	0.84
v	83.13	88.13	6.01	84.37	1.49	84.20	1.28	86.86	4.49
p	84.24	87.93	4.37	85.42	1.39	84.27	0.03	86.75	2.97
unew	28.67	39.72	38.56	**74.16**	158.68	28.80	0.47	**84.77**	195.71
vnew	37.37	47.17	26.21	**75.86**	102.97	42.83	14.58	43.97	17.64
pnew	31.15	55.51	78.22	**75.28**	141.68	47.31	51.88	48.23	54.84
uold	42.62	50.78	19.15	47.59	11.67	**67.74**	58.95	47.59	11.66
vold	54.92	62.60	13.99	62.44	13.70	**75.05**	36.66	62.48	13.77
pold	47.35	59.25	25.13	55.55	6.41	**71.41**	33.56	55.71	6.65
cu	75.81	79.73	5.54	79.35	5.11	77.96	3.19	79.73	5.54
cv	75.75	82.15	10.28	82.15	10.21	81.45	7.58	82.15	10.28
z	73.81	84.57	14.85	84.50	14.85	78.26	6.12	84.57	14.85
h	74.18	85.53	18.38	85.53	18.38	83.16	12.28	85.53	18.38
psi	92.38	92.84	0.49	92.84	0.49	92.84	0.49	92.84	0.49

Figure 2: **Hit Rates for the array variables in the Spec95 Swim Benchmark for L1 cache size 16 KB, associativity 4, and cache line size 8 words. The bold numbers in the** `KK1`, `KK2`, **and** `KK3` **columns indicate the hit rate of the keep variables for these columns.**

Bench-mark	LRU Hit %	LRU Cycles	KK Hit %	KK Cycles	% Imprv Cycles
tomcatv	94.05	1105082720	95.35	1014793630	8.90
applu	97.88	113204089	97.89	113171529	0.03
swim	91.16	12795379	93.52	11188059	14.37
mswim	95.01	10627767	96.30	9715267	9.39

(a)

LRU Hit %	LRU Cycles	KK Hit %	KK Cycles	% Imprv Cycles
93.96	122920088	95.17	113604798	8.20
97.80	635339200	97.80	635200200	0.02
90.82	100394210	93.19	88016560	14.06
94.92	81969394	96.15	75290924	8.87

(b)

Figure 4: **Overall Hit Rates and Performance for benchmarks: (a) For a given input (b) Worst case. L1 cache size 16 KB, associativity 4, and cache line size 8 words. We assume off-chip memory access requires 10 processor clock cycles, as compared to a single cycle to access the on-chip cache.**

References

[1] Doug Burger and Todd M. Austin. The simplescalar tool set, version 2.0. Technical report, University of Wisconsin-Madison Computer Science Department, 1997.

[2] D. Chiou, S. Devadas, P. Jain, and L. Rudolph. Application-Specific Memory Management for Embedded Systems Using Software-Controlled Caches. In *Proceedings of the 37^{th} Design Automation Conference*, June 2000.

[3] Cyrix. Cyrix 6X86MX Processor. May 1998.

[4] Cyrix. Cyrix MII Databook. Feb 1999.

[5] R. Kessler. The Alpha 21264 Microprocessor: Out-Of-Order Execution at 600 Mhz. In *Hot Chips 10*, August 1998.

[6] Donghee Lee, Jongmoo Choi, Jong-Hun Kim, Sam H. Noh, Sang Lyul Min, Yookun Cho, and Chong Sang Kim. On the existence of a spectrum of policies that subsumes the least recently used (LRU) and least frequently used (LFU) policies.

In *Proceedings of the international conference on Measurement and modeling of computer systems*, 1999.

[7] Y. Li and W. Wolf. A Task-Level Hierarchical Memory Model for System Synthesis of Multiprocessors. In *Proceedings of the 34^{th} Design Automation Conference*, pages 153–156, June 1997.

[8] N. Maki, K. Hoson, and A Ishida. A Data-Replace-Controlled Cache Memory System and its Performance Evaluations. In *TENCON 99. Proceedings of the IEEE Region 10 Conference*, 1999.

[9] C. May, E. Silha, R. Simpson, H. Warren, and editors. *The PowerPC Architecture: A Specification for a New Family of RISC Processors*. Morgan Kaufmann Publishers, Inc., 1994.

[10] S. McFarling. Program Optimization for Instruction Caches. In *Proceedings of the 3^{rd} Int'l Conference on Architectural Support for Programming Languages and Operating Systems*, pages 183–191, April 1989.

[11] Sun Microsystems. UltraSparc User's Manual. July 1997.

[12] P. R. Panda, N. Dutt, and A. Nicolau. *Memory Issues in Embedded Systems-on-Chip: Optimizations and Exploration.* Kluwer Academic Publishers, 1999.

[13] Edward S. Tam, Jude A. Rivers, Vijayalakshmi Srinivasan, Gary S. Tyson, and Edward S. Davidson. UltraSparc User's Manual. *IEEE Transactions on Computers*, 48(11):1244–1259, November 1999.

[14] H. Tomiyama and H. Yasuura. Code Placement Techniques for Cache Miss Rate Reduction. *ACM Transactions on Design Automation of Electronic Systems*, 2(4):410–429, October 1997.

[15] W. A. Wong and J.-L Baer. Modified LRU policies for improving second-level cache behavior. In *Proceedings of the Sixth International Symposium on High-Performance Computer Architecture*, 1999.

A Appendix

A.1 Lemma 2 Proof

Proof: At $t = 0, C(KIL, 0) \preceq_1 C(LRU, 0)$.

Assume that at time t, $C(KIL, t) \preceq_1 C(LRU, t)$.

At time $t + 1$, let the element that is accessed be e.

Case H: The element e results in a hit in $C(KIL, t)$. If the K_l bit for e is set, then e is also an element of $C(LRU, t)$ from the assumption at time t. Now the K_l bit of e would be reset unless it is set by this access. Thus, we have $C(KIL, t+1) \preceq_1 C(LRU, t+1)$. If the K_l bit of e is 0, then there is no change in the order of elements with the K_l bit set. So, we have $C(KIL, t + 1) \preceq_1 C(LRU, t + 1)$.

Case M: The element e results in a miss in $C(KIL, t)$. Let y be the least recent element with the K_l bit set in $C(KIL, t)$. If e results in a miss in $C(LRU, t)$, let $C(KIL, t) = \{M_1, y, M_2\}$ and $C(LRU, t) = \{L, x\}$. M_2 has no element with K_l bit set. If the K_l bit of x is 0, $\{M_1, y\} \preceq_1 \{L\}$ implies $C(KIL, t+1) \preceq_1 C(LRU, t+1)$. If the K_l bit of x is set and $x = y$ then $\{M_1\} \preceq_1 \{L\}$ and that implies $C(KIL, t + 1) \preceq_1 C(LRU, t+1)$. If the K_l bit of x is set and $x \neq y$, then $x \notin M_1$ because that violates the assumption at time t. Further, $y \in L$ from the assumption at time t and this implies $C(KIL, t + 1) \preceq_1 C(LRU, t + 1)$. ∎

A.2 Theorem 1 Proof

Proof: We consider the Kill + LRU policy variation (1) for a fully-associative cache C with associativity m. We show that $C(LRU, t) \preceq_0 C(KIL, t)$ at any time t.

At $t = 0$, $C(LRU, 0) \preceq_0 C(KIL, 0)$.

Assume that at time t, $C(LRU, t) \preceq_0 C(KIL, t)$.

At time $t + 1$, let the element accessed is e.

Case 0: The element e results in a hit in $C(LRU, t)$. From the assumption at time t, e results in a hit in $C(KIL, t)$ too. Let $C(LRU, t) = \{L_1, e, L_2\}$ and $C(KIL, t) = \{M_1, e, M_2\}$. From the assumption at time t, $L_1 \preceq_0 M_1$ and $L_2 \preceq_0 M_2$. From the definition of LRU and Kill + LRU replacement, $C(LRU, t + 1) = \{e, L_1, L_2\}$ and $C(KIL, t + 1) = \{e, M_1, M_2\}$. Since $\{L_1, L_2\} \preceq_0 \{M_1, M_2\}$, $C(LRU, t + 1) \preceq_0 C(KIL, t + 1)$.

Case 1: The element e results in a miss in $C(LRU, t)$, but a hit in $C(KIL, t)$. Let $C(LRU, t) = \{L, x\}$ and $C(KIL, t) = \{M_1, e, M_2\}$. From the assumption at time t, $\{L, x\} \preceq_0 \{M_1, e, M_2\}$ and it implies that $\{L\} \preceq_0 \{M_1, e, M_2\}$. Since $e \notin L$, we have $\{L\} \preceq_0 \{M_1, M_2\}$. From the definition of LRU and Kill + LRU replacement, $C(LRU, t + 1) = \{e, L\}$ and $C(KIL, t + 1) = \{e, M_1, M_2\}$. Since $L \preceq_0 \{M_1, M_2\}$, $C(LRU, t + 1) \preceq_0 C(KIL, t + 1)$.

Case 2: The element e results in a miss in $C(LRU, t)$ and a miss in $C(KIL, t)$ and there is no element with K_l bit set in $C(KIL, t)$. Let $C(LRU, t) = \{L, x\}$ and $C(KIL, t) = \{M, y\}$. From the assumption at time t, there are two possibilities: (a) $\{L, x\} \preceq_0 M$, or (b) $L \preceq_0 M$ and $x = y$. From the definition of LRU and Kill + LRU replacement, $C(LRU, t+1) = \{e, L\}$ and $C(KIL, t+1) = \{e, M\}$. Since $L \preceq_0 M$ for both sub-cases (a) and (b), we have $C(LRU, t + 1) \preceq_0 C(KIL, t + 1)$.

Case 3: The element e results in a miss in $C(LRU, t)$ and a miss in $C(KIL, t)$ and there is at least one element with the K_l bit set in $C(KIL, t)$. There are two sub-cases (a) there is an element with the K_l bit set in the LRU position, (b) there is no element with the K_l bit set in the LRU position. For sub-case (a), the argument is the same as in Case 2. For sub-case (b), let the LRU element with the K_l bit set be in position i, $1 \leq i < m$. Let $C(LRU, t) = \{L, x\}$ and $C(KIL, t) = \{M_1, y, M_2\}, M_2 \neq \phi$. From the assumption at time t, $\{L, x\} \preceq_0 \{M_1, y, M_2\}$, which implies $\{L\} \preceq_0 \{M_1, y, M_2\}$. Since y has the K_l bit set, $y \in L$ using Lemma 2. Let $\{L\} = \{L_1, y, L_2\}$. So, $\{L_1\} \preceq_0 \{M_1\}$ and $\{L_2\} \preceq_0 \{M_2\}$. Using Lemma 1, for the LRU policy y would be evicted from the cache before the next access of y. The next access of y would result in a miss using the LRU policy. So, $\{L_1, y, L_2\} \preceq_0 \{M_1, M_2\}$ when considering the elements that do not have have the K_l bit set. From the definition of LRU and Kill + LRU replacement, $C(LRU, t + 1) = \{e, L_1, y, L_2\}$ and $C(KIL, t + 1) = \{e, M_1, M_2\}$. Using the result at time t, we have $C(LRU, t + 1) \preceq_0 C(KIL, t + 1)$. ∎