



Computer Science and Artificial Intelligence Laboratory
Technical Report

MIT-CSAIL-TR-2015-011

April 9, 2015

A Cache Model for Modern Processors
Nathan Beckmann and Daniel Sanchez



A Cache Model for Modern Processors

Modeling high-performance cache replacement on modern last-level caches

Nathan Beckmann
Massachusetts Institute of Technology
beckmann@csail.mit.edu

Daniel Sanchez
Massachusetts Institute of Technology
sanchez@csail.mit.edu

ABSTRACT

Modern processors use high-performance cache replacement policies that outperform traditional alternatives like least-recently used (LRU). Unfortunately, current cache models use stack distances to predict LRU or its variants, and do not capture these high-performance policies. Accurate predictions of cache performance enable many optimizations in multicore systems. For example, cache partitioning uses these predictions to divide capacity among applications in order to maximize performance, guarantee quality of service, or achieve other system objectives. Without an accurate model for high-performance replacement policies, these optimizations are unavailable to modern processors.

We present a new probabilistic cache model designed for high-performance replacement policies. This model uses absolute reuse distances instead of stack distances, which makes it applicable to arbitrary age-based replacement policies. We thoroughly validate our model on several high-performance policies on synthetic and real benchmarks, where its median error is less than 1%. Finally, we present two case studies showing how to use the model to improve shared and single-stream cache performance.

1. INTRODUCTION

On-chip caches are fundamental to cope with the long latency, high energy, and limited bandwidth of main memory accesses. Most cache space is consolidated in a large last-level cache (LLC) shared among cores, which consumes significant resources (e.g., over 50% of chip area [22]). Rising core counts place mounting pressure on the memory system, and have led to a resurgence of cache architecture research. As a result, modern LLCs have become quite sophisticated in their array organization and replacement policy.

Understanding the LLC’s behavior is critical to achieve system objectives. Accurate predictions of cache behavior enable a large number of optimizations, including of single-threaded performance [4, 8], shared cache performance [3, 26, 29, 37], fairness [26, 28], quality of service [3, 17], security [27], etc.. These optimizations can be performed in many ways, e.g. by explicitly partitioning the shared cache in hardware [2, 23, 33] or software [7], or through job scheduling to avoid interference in cache accesses [24, 44, 45].

Unfortunately, cache behavior is difficult to predict because it depends on many factors, both of the application (e.g., its access pattern) and the cache (e.g., its size, associativity, and replacement policy). Existing cache models [1, 8, 25, 35, 43] tend to focus on traditional, set-associative caches using

simple replacement policies like least-recently used (LRU), pseudo-LRU, or random replacement. *But modern processors do not use LRU (or pseudo-LRU) for the LLC.*

Modern LLCs instead employ high-performance replacement policies that greatly improve cache performance over traditional policies like LRU (Sec. 2.2). These designs are already available in commercial processors [13, 41], and a new model is needed to understand their behavior.

We present a cache model that accurately predicts the behavior of high-performance replacement policies on modern LLCs. Our model leverages two key observations: First, each core’s private cache hierarchy filters accesses before they reach the LLC, capturing most temporal locality [19]. Thus, LLC accesses are free of short-term temporal correlations. (This is also why LRU is a poor policy for the LLC; LRU relies on temporal locality, and at the LLC, there is little.) Second, modern LLCs use hashing to map lines to sets, reducing hotspots [21, 41]. Existing models often focus on the behavior of individual sets and the effects of low associativity, i.e. on conflict misses, which are hard to capture. But with hashing, modern LLCs have near-uniform behavior across sets and high effective associativity [32], making conflict misses a second-order concern (Sec. 2.1).

These two observations mean that *modern LLCs can be modeled as a pseudo-random process*: by capturing highly correlated accesses, private caches essentially randomize the accesses seen by the LLC; and hashing and high effective associativity mean that the replacement candidates constitute a representative sample of cached lines. We therefore model caching as a probabilistic process operating on individual lines. Specifically, our model computes the probability of various events (hits and evictions) occurring to individual lines as they age. This approach lets us model arbitrary age-based policies, including familiar policies like LRU and several high-performance policies, and abstracts away the details of array organization.

These features set our model apart from existing models. Existing models target set-associative LRU caches, and many model the performance of individual *sets* using *stack distances*, which measure the number of unique addresses between references to the same line. Stack distances are meaningful for set-associative LRU caches, but have little meaning for modern LLCs. In particular, it is unclear how to model high-performance policies through stack distances (Sec. 2.3).

Our model is built from logically distinct components, each of which captures a separate aspect of cache behavior (Sec. 4).

For instance, to model a new replacement policy, our model only requires simple changes to a single parameter, the ranking function (Sec. 5). We present an efficient, practical implementation of our model (Secs. 6 and 7) and thoroughly validate it against synthetic and real benchmarks (Sec. 8). Finally, we present two case studies: (i) cache partitioning with high-performance replacement (Sec. 9); and (ii) improving the replacement policy itself (Sec. 10). In summary, we offer an efficient way to predict the performance of high-performance policies, allowing them to enjoy the many benefits that prior modeling work has demonstrated for LRU.

2. BACKGROUND

We first review the relevant background in modern LLC architecture, replacement policies, and cache modeling.

2.1 Modern cache architecture

Modern multicores feature multilevel cache hierarchies. Each core has one or two levels of small, fast, private caches (L1 and L2). Private caches are backed by a much larger last-level cache (LLC) that contains the bulk of cache capacity.

Cache architecture varies greatly across levels. The purpose of the private levels is to cache the program’s most frequently accessed data at low latency and high bandwidth, so they are kept small and simple—and therefore fast. Specifically, both the L1 and L2 adopt a familiar, set-associative architecture and use simple replacement policies like pseudo-LRU.

The cache architecture is more sophisticated at the LLC. Since off-chip misses are very expensive, various caching techniques are attractive for the LLC that are a poor design for the private levels. The most significant change is the replacement policy, discussed in detail below. In addition, the LLC adopts an array organization that ensures high effective associativity. For instance, by hashing addresses before indexing the array, the cache pseudo-randomly spreads accesses across sets and reduces conflicts [21, 22, 38, 41]. Other designs further increase effective associativity without adding ways. For instance, skew-associative caches [36] use a different hash function for each way, mixing replacement candidates across sets; and zcaches [32] use cuckoo hashing to select many more replacement candidates than physical ways.

For the purposes of cache modeling, there are two important implications. First, private caches capture most “hot” data, so the LLC’s access stream is stripped of short-term temporal correlations: a second access to the same line requires that the line be first evicted from the L1 and L2. Second, since modern LLCs use hashing and achieve high effective associativity, replacement candidates form a representative sample of cached lines [4, 29, 32]. Details of array organization, which have been the focus of prior studies [1, 8, 25, 35, 43], are relatively unimportant in modern LLCs. Our model leverages this insight by modeling replacement as a probabilistic process affecting individual lines, not sets.

2.2 High-performance replacement policies

The largest single difference between the traditional caches targeted by existing models and modern LLCs is the replacement policy. Traditional replacement policies (e.g., LRU, pseudo-LRU, least-frequently used, or random) all perform

poorly at the LLC because private caches filter out most locality [19]. LRU is the most common, but has common performance pathologies. For instance, when an application’s working set does not fit in the cache, LRU causes thrashing and performance cliffs; e.g., iterating sequentially over a 1 MB array gets zero hits with less than 1 MB of cache, but every access suddenly hits at 1 MB.

Such pathologies are avoided in the optimal replacement policy, Belady’s MIN [6, 25], which relies on a perfect oracle to replace the line that will be reused furthest in the future. Prior work has proposed many mechanisms and heuristics that improve on LRU and attempt to emulate optimal replacement. We observe that, details aside, they use three broad techniques:

- **Recency:** Recency prioritizes recently used lines over old ones. LRU uses recency alone, leading to its pathologies (e.g., thrashing). Most high-performance policies combine recency with other techniques.
- **Protection:** When the working set does not fit in the cache, some policies protect a portion of the working set against eviction to avoid thrashing. This is equivalent to *thrash-resistance* [14, 30].
- **Classification:** Some policies divide accesses into separate classes, and treat lines of each class differently. For example, several policies classify lines as either reused or non-reused [14, 20]. Classification works well when classes have markedly different access patterns.

For instance, DIP [30] enhances LRU by dynamically detecting thrashing using set dueling, and protects lines in the cache by inserting most lines at low priority in the LRU chain. DIP inserts a fixed fraction of lines ($\epsilon = 1/32$) at high priority to avoid stale lines. DRRIP [14] classifies between reused and non-reused lines by inserting lines at medium priority, includes recency by promoting lines on reuse, and protects against thrashing with the same mechanism as DIP. SHIP [42] extends DRRIP with more elaborate classification, based on the memory address, program counter (PC), or instruction sequence. Instruction-based reuse distance prediction [18] uses PC-based classification, but does not do protection. PDP [9] decides how long to protect lines based on the reuse distance distribution, improving on DIP, but does not do classification. Finally, IRGD [39] adapts its policy to the access stream, ranking lines according to their harmonic expected reuse distance. The policies we consider are described in greater detail in Sec. 5.1.

Most of these policies are age-based: lines enter the cache at a given priority, which changes over time according to how long the line has been in the cache. However, the policies vary in details. PDP and IRGD directly count the number of accesses since a line was last referenced, while DRRIP and SHIP use a clever aging mechanism that allows them to rank lines using fewer bits in hardware. Our model targets arbitrary age-based policies, and we age lines by the number of accesses since their last reference (like PDP and IRGD).

No prior work gives a general way to predict these policies. DIP gives a simple performance analysis on a single access pattern (the scanning pattern discussed above), but does not consider performance on realistic access patterns or any




Figure 1: Our model consists of three interdependent probability distributions. Each distribution captures distinct model inputs. Arrows denote dependence, e.g. $A \rightarrow B$ means “A depends on B”.

general method to assess this. PDP employs an analytical model to choose for how long to protect lines in the cache. PDP’s model is simple, but it is inaccurate on common access patterns (Sec. 8.1). Moreover, even if it were accurate, it is limited to modeling protecting distances, and therefore does not solve the general problem of modeling high-performance cache replacement.

2.3 Cache modeling

Most prior work has developed analytical cache models using *stack distance distributions* [1, 8, 25, 35, 43]. The stack distance of an access is the number of *unique* addresses accessed after the last reference to the same address. For example, in the sequence ABCBDDA the stack distance of A is four, as B, C, and D are accessed between both accesses to A. Stack distances are meaningful for LRU: in a fully-associative LRU cache of S lines, accesses with stack distance $\leq S$ will be hits, and all others will be misses.

Stack distances are inherently tied to LRU, but have little meaning for other policies. Our insight is that *absolute reuse distances* [35] can be used to model cache behavior for a much wider range of replacement policies. The absolute reuse distance of an access is the total number of references after the last reference to the same address. For example, in ABCBDDA the absolute reuse distance of access A is six. Absolute reuse distances correspond to lines’ ages: a line with absolute reuse distance d will hit at age d if it is not first evicted.

Absolute reuse distance distributions cannot be trivially translated into miss rates. *Our key innovation is to develop an analytical model that performs this translation for a broad class of policies.* Prior work has used absolute reuse distances only in a limited context: Sen et al. [35] use absolute reuse distances to model random replacement, but use stack distances for LRU. By contrast, we model all policies through a single framework based on absolute reuse distances. Since we are able to model LRU, our model demonstrates that absolute reuse distances dominate stack distances in predictive power.

For brevity, in the rest of the paper *we refer to absolute reuse distances simply as reuse distances.* This terminology is in line with prior cache replacement papers [9, 14, 18], but note that, since prior work in analytical models focuses on stack distance, some use reuse and stack distance as synonyms [8, 35, 43].

Request	A	A	B	C	B	D	B	C
Lines	A	A	A	B	B	D	B	C
Time	1	2	3	4	5	6	7	8

Table 1: Steady state behavior for a 3-line LRU cache on a simple, repeating access pattern. Live lines are colored green, and dead lines red.

3. OVERVIEW

Fig. 1 shows the high-level design of our cache model. The input to the model is the cache architecture (its size, associativity, and replacement policy) and a concise description of the access stream. Specifically, we describe the access stream by its *reuse distance distribution*; i.e., for each distance d , how many accesses have reuse distance d .

From these inputs, our model produces a concise description of the cache’s behavior. Specifically, the model produces the cache’s *hit and eviction distributions*; i.e., for each age a , how many accesses are hit or evicted at age a , respectively. The hit and eviction distributions directly yield the cache’s performance: the cache’s hit rate is the sum over the hit distribution. Additionally, they give a rich picture of the cache’s behavior that can be used to improve cache performance (Sec. 10).

Internally, the model uses three distributions: the hit and eviction distributions (already discussed) and the age distribution of cached lines (i.e., the probability that a randomly selected line has age a). We define age as the number of accesses since a line was last referenced. These distributions are interdependent and related to one another by simple equations. Each incorporates a particular model input and conveys its constraints: (i) the age distribution incorporates the cache size, and constrains modeled capacity; (ii) the hit distribution incorporates the access stream and constrains when lines can hit; and (iii) the eviction distribution incorporates the replacement policy and constrains how long lines stay in the cache. The model is solved by iterating to a fixed point. When the distributions converge, the hit and eviction distributions accurately describe the cache’s behavior.

We build the model in stages. First, we introduce the model on a specific example. Second, we develop the model formulation for LRU. Third, we generalize it to other policies. Finally, we show to solve the model efficiently.

3.1 Example

Table 1 shows the behavior of a 3-line LRU cache on a simple, repeating access pattern. Time (measured in accesses) increases from left to right and wraps around (e.g., time 9 would be identical to 1). Live lines, those that eventually hit, are colored green; dead lines, those that do not, are red. For example, the first line is live (green) at time 1 because A hits at time 2. However, D evicts A at time 6, so A is dead (red) in 2–5. Similarly, A evicts D at time 1, so D is dead in 6–9. B and C always hit, so they are always live. This divides the cache into *lifetimes*, the intervals from last use until hit or eviction. For example, A’s lifetime starting at time 1 ends after a single access when A hits at time 2; this

Request	A	A	B	C	B	D	B	C
Ages	<u>1</u>	<u>1</u>	<u>2</u>	<u>3</u>	<u>4</u>	<u>1</u>	<u>2</u>	<u>3</u>
Time	1	2	3	4	5	6	7	8

Table 2: Ages per line in Table 1 (hits are underlined, evictions in italic).

Symbol	Meaning	
$P[y]$	The probability of event y .	
$E[X]$	The expected value of random variable (rv) X .	
$P_X(x)$	The probability that rv X equals x , $P[X=x]$.	
$P_{X,Y}(x,y)$	The joint probability rv X equals x and rv Y equals y .	
$P_X(x y)$	The conditional probability that rv X equals x given y .	
RV	Meaning	
D	Reuse distance of an access.	} Input
H	Age at which a line hits.	
E	Age at which a line is evicted.	} Output
A	Age of a line.	
L	Lifetime of a line (see text).	} Internal

Table 4: Notation used in this paper. D is a property of the access stream; A , H , and E are properties of cached lines.

starts a new lifetime that ends after four accesses when D evicts A at time 6.

We can redraw Table 1 showing the ages of each line at each access (Table 2). In steady state, 6 out of 8 accesses are hits. These hits (underlined) come at ages 1 (A once), 2 (B twice), and 4 (C twice and B once). The other 2 accesses are evictions (*italic*), at ages 3 (D) and 4 (A).

This information lets us compute the hit and eviction distributions, denoted by random variables H and E , respectively. These distributions yield the probability that a (randomly selected) access will hit or be evicted at a given age. For example, one quarter of accesses hit at age 2 (B at times 4 and 6), so the hit distribution at age 2 is $P_H(2) = 1/4$. Table 4 summarizes our notation and the most important distributions in the model.

Table 3 gives the three model distributions for the example in Table 1. For the hit and eviction distributions, we turn counts into probabilities by dividing by 8 accesses. Since every lifetime ends in a hit or eviction, the hit and eviction distributions together sum to 1, but separately sum to less than 1. We express the distribution inclusive of both hits and evictions as the *lifetime distribution*, $P_L(a) = P_H(a) + P_E(a)$ (not shown in Table 3). L is the age at which a line is either hit or evicted, i.e. the age when its lifetime ends. Finally, we can compute the age distribution by counting the ages of lines and dividing by 3 lines \times 8 accesses = 24.

These distributions tell us how the cache performs on this

Age	a	1	2	3	4	5	6	7	8	Sum
Distribution	$P_D(a)$	$1/8$	$2/8$	–	$3/8$	–	–	$1/8$	$1/8$	1
	$P_A(a)$	$1/3$	$7/24$	$5/24$	$1/6$	–	–	–	–	1
	$P_H(a)$	$1/8$	$1/4$	–	$3/8$	–	–	–	–	$3/4$
	$P_E(a)$	–	–	$1/8$	$1/8$	–	–	–	–	$1/4$

Table 3: Steady-state distributions for the cache in Table 1.

access pattern: its hit rate is the sum over the hit distribution, $3/4$. Moreover, the distributions also say how lines behave: e.g., no lines make it past age 4, despite some having reuse distance of 8. Such information will prove valuable in applying the model (Sec. 10).

In this example, we have computed the distributions by brute force, but this method is too expensive for our purposes. We instead model the relationships among distributions analytically and solve them by iteration. This example gives the intuition behind the model, but note that some model assumptions only hold for large caches.

4. BASIC MODEL (FOR LRU)

This section presents the model for caches with LRU replacement. We present the complete equations for the age and hit distributions, and the eviction distribution equations for LRU replacement. Sec. 5 extends the eviction distribution to model arbitrary age-based replacement policies.

4.1 Model assumptions

First, we make a few simplifying assumptions that make cache modeling tractable, chiefly about the independence of certain events. These assumptions are motivated by the properties of modern LLCs (Sec. 2.1).

We assume each access has reuse distance d distributed identically and independently according to the distribution $P_D(d)$. This assumption is not exactly correct in practice, but since private caches filter accesses before they reach the LLC, it is a good approximation for large caches ([4], Sec. 8). In return, it greatly simplifies the probability calculations, allowing a simple model to capture diverse access patterns.

Similarly, we model an idealized, *random-candidates* cache, where replacement candidates are drawn at random from cached lines. This is a good model for modern LLCs, where the replacement candidates form a representative sample of cached lines. The random-candidates model is a direct analog of skew-associative caches or zcaches, but is also a good approximation for hashed, set-associative caches with many ways [32]. Although the accuracy of this assumption depends slightly on the cache array architecture, it is a reasonable simplifying assumption for modern LLCs.

4.2 Age distribution

The age distribution is used internally by the model to constrain cache capacity. It is presented first because it is the simplest to compute from the other distributions.

Since ages measure the time since a line was last referenced, a line reaches age a if and only if it is not hit or evicted for at

least a accesses. Hence the probability of a line having age a , $P_A(a)$, is proportional to the fraction of lines that survive at least a accesses in the cache without being referenced. The probability of surviving *exactly* x accesses is given by the lifetime distribution at age x , $P_L(x) = P_H(x) + P_E(x)$, and the probability of surviving *at least* a accesses is the probability over all $x \geq a$. The age distribution $P_A(a)$ is thus proportional to $P[L \geq a]$, which determines the age distribution up to a constant factor.

To find the constant factor, note that every access necessarily produces a line of age 1 (whether a hit or miss), since we define ages as the number of accesses since a line was last referenced. Since ages increase upon each access, there is always exactly one line of age 1 in the cache. Hence if the cache has S lines, the age distribution at age 1 is $P_A(1) = 1/S$.

Combining the two results, we find the age distribution for a cache of S lines is:

$$P_A(a) = \frac{P[L \geq a]}{S} = \sum_{x=a}^{\infty} \frac{P_H(x) + P_E(x)}{S} \quad (1)$$

For example, in Sec. 3.1, by counting cells we find that $P_A(3) = 5/24$. That is, of the twenty-four cells depicted in Table 2, age 3 appears five times. Eq. 1 gives another way to arrive at the same conclusion without counting cells: The probability of a hit or eviction at age 3 or greater is $3/8 + 1/8 + 1/8 = 5/8$ (Table 3). Dividing by $S = 3$ gets $P_A(3) = 5/24$, as desired. This cell-counting argument can be generalized to yield an alternative derivation of Eq. 1 and show $\sum_{a=1}^{\infty} P_A(a) = 1$.

4.3 Hit distribution

We now show how to compute when hits occur for a given access pattern, again assuming the other distributions are known. The hit distribution is perhaps the most important product of the model, since it yields the cache's hit rate.

A line will eventually hit if it is not evicted. Intuitively, a line's hit probability depends on its reuse distance (longer distances have greater opportunity to be evicted) and the cache's eviction distribution (i.e., at what age does the replacement policy evict lines?). Moreover, by definition, a line with reuse distance d must hit at age d , if it hits at all. Thus the hit probability is just the reuse distance probability minus the probability of being evicted, or:

$$\begin{aligned} P_H(d) &= P_D(d) - P[\text{evicted}, D=d] \\ &= P_D(d) - \sum_{a=1}^{d-1} P_{D,E}(d, a) \end{aligned} \quad (2)$$

To proceed, we require the critical insight that candidates of age a look alike to the replacement policy. In other words, the replacement policy does not know candidates' reuse distances, only that lines of age a have reuse distance $d > a$.¹

¹It may be tempting to simply subtract the eviction probability below d in the above equation. That is, $P_H(d) = P_D(d) \times (1 - P[E < d])$. This is incorrect. Eviction age and reuse distance are not independent.

Evicting at age a thus only implies that the line's reuse distance was at least a , i.e. $E = a \Rightarrow D > a$. From this insight:

$$\begin{aligned} P_{D,E}(d, a) &= P_D(d|E=a) P_E(a) \\ (\text{Insight}) \quad &= P_D(d|D > a) P_E(a) \\ (\text{Simplify}) \quad &= \frac{P_D(d) \cdot P_E(a)}{P[D > a]} \end{aligned} \quad (3)$$

Recall that our assumptions model a large, last-level cache, so Eq. 3 is inexact on the toy example shown in Table 1.

Finally, we substitute this result into the above equation: Summing over all a below d counts all the lines of reuse distance d that are evicted before hitting. Since the rest hit, subtracting these from $P_D(d)$ yields the hit probability:

$$P_H(d) = P_D(d) \times \left(1 - \sum_{a=1}^{d-1} \frac{P_E(a)}{P[D > a]} \right) \quad (4)$$

4.4 Eviction distribution in LRU

The eviction distribution accounts for the replacement process and is the most complex part of the model. It models both the selection of replacement candidates (i.e., associativity) and the selection of a victim among candidates (i.e., the replacement policy). For clarity, we begin by presenting the model for LRU only.

To compute the eviction distribution, we assume that candidates are drawn randomly from the cache, as discussed above. Among these, LRU simply evicts the oldest, so the eviction probability at age a is the probability that both an access misses and a is the oldest age among candidates. We also assume that cache misses and the oldest age are independent. This is reasonable because a large cache is negligibly affected by the behavior of a few candidates.

Hence to find the eviction probability at age a , we need to know the probabilities that the access misses and the oldest candidate has age a : $P_E(a) = P[\text{miss}] \cdot P_{\text{oldest}}(a)$. The first factor is trivial, since

$$P[\text{miss}] = 1 - P[\text{hit}] = 1 - \sum_{a=1}^{\infty} P_H(a) \quad (5)$$

The challenge lies in finding the distribution of oldest ages, $P_{\text{oldest}}(a)$. Replacement candidates are drawn randomly from the cache, so each has age identically and independently distributed according to $P_A(a)$. The oldest age is just the maximum of W iid random variables, where W is the associativity. The maximum of iid random variables is a well-known result [11]: the oldest age among replacement candidates is less than a iff all candidates are of age less than a . Thus given W candidates, $P[\text{oldest} < a] = P[A < a]^W$. To get the distribution from the cumulative distribution, differentiate: $P_{\text{oldest}}(a) = P[\text{oldest} < a + 1] - P[\text{oldest} < a]$. Altogether, the eviction distribution at age a for LRU is:

$$\begin{aligned} P_E(a) &= (1 - P[\text{hit}]) P_{\text{oldest}}(a) \\ &= (1 - P[\text{hit}]) (P[\text{oldest} < a + 1] - P[\text{oldest} < a]) \\ \Rightarrow P_E(a) &= (1 - P[\text{hit}]) \left(P[A < a + 1]^W - P[A < a]^W \right) \end{aligned} \quad (6)$$

4.5 Summary

Equations 1, 4, and 6 form the complete model for LRU replacement. The age distribution incorporates the cache size and depends on the hit and eviction distributions. The hit distribution incorporates the access stream and depends on the eviction distribution. The eviction distribution incorporates the cache’s associativity and replacement policy (LRU) and depends on the hit and age distributions.

5. OTHER REPLACEMENT POLICIES

We now extend the eviction distribution to support arbitrary age-based policies, like those discussed in Sec. 2.2.

5.1 Ranking functions

To support other policies, we must abstract the replacement policy in a way that can be incorporated into the model. We do so through a *ranking function*, $\mathcal{R} : \text{age} \rightarrow \mathbb{R}$, which gives an eviction priority to every age. By convention, higher rank means higher eviction priority.

Ranking functions capture many existing policies. For example, LRU’s ranking function is $\mathcal{R}_{\text{LRU}}(a) = a$ (or any other strictly increasing function). This represents LRU because it ensures that older lines will be preferred for eviction. Similarly, a constant ranking function produces random replacement, e.g. $\mathcal{R}_{\text{random}}(a) = 0$. Ranking functions can also capture many high-performance replacement policies.

PDP [9] protects lines up to an age d_p , known as the protecting distance. It prefers to evict lines older than the protecting distance, but if none are available among candidates, it evicts the youngest line. Thus PDP’s ranking function decreases up to the protecting distance (d_p), upon which it jumps to a large value and increases thereafter:

$$\mathcal{R}_{\text{PDP}}(a) = \begin{cases} d_p - a & \text{If } a < d_p \\ a & \text{If } a \geq d_p \end{cases} \quad (7)$$

IRGD [39] ranks lines using an analytical formulation based on the reuse distance distribution. IRGD essentially ranks lines by their expected reuse distance, but since in practice very large reuse distances can’t be measured, IRGD uses a weighted harmonic mean instead of the conventional, arithmetic mean. This lets it ignore immeasurably large reuse distances, since they have small impact on the harmonic mean.² Its ranking function is:

$$\mathcal{R}_{\text{IRGD}}(a) = P[D > a] \times \left(\sum_{x=1}^{\infty} \frac{P_D(a+x)}{a+x} \right)^{-1} \quad (8)$$

Rank functions thus model any age-based policy, but not all high-performance policies are strictly age-based. Our model can support such policies (e.g., RRIP [14, 42]) with specialized extensions. However, this paper presents the general framework, and we leave extensions to future work.

From the ranking function and age distribution, we can produce a *rank distribution* that gives the probability a line will

²Takagi et al. [39] express IRGD somewhat differently, but the two formulations are equivalent.

have a given rank. It is then possible to generalize Eq. 6. While LRU evicts the oldest replacement candidates, in general the cache evicts the maximum rank among candidates.

5.2 Generalized eviction distribution

Generalizing the eviction distribution is a straightforward substitution from “oldest age” in LRU to “maximum rank” in general. If a line of age a is evicted, then the maximum rank among candidates must be $\mathcal{R}(a)$. Additionally, \mathcal{R} may rank several ages identically (i.e., $\mathcal{R}(a) = \mathcal{R}(b)$), so we must ensure that the candidate had age a (not age b).

This consideration is important because, in practice, continuous ranks are quantized in units of Δr , increasing the possibility that several ages map to indistinguishable ranks. For example, if ranks can take values in $[0, 256]$ (e.g., LRU with 8-bit ages), then an efficient model implementation might quantize ranks into regions as $[0, 8), [8, 16) \dots [248, 256]$. Each region has size $\Delta r = 8$, and many ages may have the “same rank” as far as the model is concerned.

We account for indistinguishable ranks by using the joint distribution of rank and age to avoid double counting:

$$P_E(a) = (1 - P[\text{hit}]) \cdot P_{\text{max rank}, A}(\mathcal{R}(a), a) \quad (9)$$

The joint distribution is in turn:

$$P_{\text{max rank}, A}(\mathcal{R}(a), a) = P_{\text{max rank}}(\mathcal{R}(a)) \cdot \frac{P_A(a)}{P_{\text{rank}}(\mathcal{R}(a))} \quad (10)$$

$P_A(a)/P_{\text{rank}}(\mathcal{R}(a))$ is the fraction of ranks belonging to age a in the rank quantum containing $\mathcal{R}(a)$ (roughly its Δr -neighborhood). Multiplying by this fraction eliminates double counting. This equation should simply be thought as the analog to $P_{\text{oldest}}(a)$ in LRU.

As in LRU, the challenge lies in finding $P_{\text{max rank}}(r)$. To start, we compute the *rank distribution* in the cache from the age distribution. Since ranks depend on age, the probability that a line’s rank equals r is just the total probability of ages with rank r :

$$P_{\text{rank}}(r) = \sum_{a: \mathcal{R}(a)=r} P_A(a) \quad (11)$$

Next, the cumulative distribution of maximum rank is computed just as $P_{\text{oldest}}(a)$ in LRU as the maximum of iid rvs:

$$P[\text{max rank} < r] = P[\text{rank} < r]^W \quad (12)$$

Finally, the distribution of maximum rank is obtained by discrete differentiation [31]:

$$P_{\text{max rank}}(r) = \frac{P[\text{max rank} < r + \Delta r] - P[\text{max rank} < r]}{\Delta r} \quad (13)$$

(In LRU, the oldest age distribution uses $\Delta r = 1$.)

These formulae fill in all the terms to compute the general-

Algorithm 1. The cache model simultaneously solves the cache’s age, hit, and eviction distributions by iteration.

Inputs: S - Cache size; W - associativity; \mathcal{R} - rank function; rdd - reuse distance distribution; prd - with classification, total probability of class (Appendix B), or 1 otherwise.
Returns: Hit and eviction distributions, hit and evict.

```

1: function Model
2:   age, hit, evict, h' ← Seed           ▷ Initialize distributions.
3:   while not Converged:
4:     h ← h'                             ▷ Hit rate from last iteration.
5:     h' ← 0                             ▷  $x'$  is solution of  $x$  for this iteration.
6:     crd ← rdd[1]                       ▷ Cumulative  $D$  probability,  $P[D \leq a]$ .
7:     evBelow ← 0                         ▷ Prob. line evicted at  $D=a$  in Eq. 4.
8:     age'[1] ← prd/S
9:     for  $a \leftarrow 1$  to  $N$ :
10:      hit'[a] ← rdd[a] (1 - evBelow)
11:      evict'[a] ← (1-h) maxRankDist[ $\mathcal{R}(a)$ ] age'[a]/rankDist[ $\mathcal{R}(a)$ ]
12:      age'[a+1] ← age'[a] - (hit'[a] + evict'[a])/S
13:      h' ← h' + hit'[a]
14:      evBelow ← evBelow + evict'[a]/(prd - crd)
15:      crd ← crd + rdd[a+1]
16:   age, hit, evict ← Average(age', hit', evict')
17:   return hit, evict

```

ized eviction distribution:

$$P_E(a) = (1 - P[\text{hit}]) \times \left(\frac{P_A(a)}{P_{\text{rank}}(\mathcal{R}(a))} \right) \times \left(\frac{P[\text{rank} < \mathcal{R}(a) + \Delta r]^W - P[\text{rank} < \mathcal{R}(a)]^W}{\Delta r} \right) \quad (14)$$

5.3 Discussion

The complete cache model is given by the age (Eq. 1), hit (Eq. 4), and eviction (Eq. 14) distributions. These equations describe a cache using an arbitrary, age-based replacement policy. Our model forms a system of equations that describe a valid solution, but does not yield this solution directly.

The implicit nature of our model has benefits. The equations organize the model into logical components. Each distribution is responsible for a specific model input: the age distribution for the cache size, the hit distribution for the access pattern, and the eviction distribution for replacement (both associativity and replacement policy). This makes it easy to adapt the model to new cache architectures. For example, a new replacement policy only requires a new ranking function, and all appropriate changes naturally propagate through the eviction, hit, and age distributions. Likewise, new applications change only the reuse distance distribution.

However the drawback is that, since these equations are not explicit, their solution is not entirely obvious. We solve the system through iteration to a fixed point, discussed next.

6. MODEL SOLUTION

All components of the model are interdependent, and a general, closed-form solution is unlikely to exist. We solve the model by iterating to a fixed point, *simultaneously* solving the three distributions age by age (Algorithm 1). This simultaneous solution tightly couples the solution of each distribution to the others, maintaining their relationships. That is, each distribution is computed from the others as the solution evolves, rather than from the distributions at the last iteration. Only the hit rate and rank distribution are fixed across iterations. We find this tight coupling improves convergence time.

At each iteration, Algorithm 1 solves Eq. 1, Eq. 4 and Eq. 14 for age a in constant time, building on the solution from age $a - 1$. Algorithm 1 uses the following recurrence relation derived from Eq. 1:

$$P_A(a + 1) = P_A(a) - \frac{P_H(a) + P_E(a)}{S} \quad (15)$$

This allows the age distribution to be updated to reflect the hit and eviction distributions as the solution evolves, which in turn influences the solution of the hit and eviction distributions. The hit and eviction distributions are thus constrained, and negative feedback loops are imposed on over-estimation. Sums in other equations are similarly broken across iterations so that each age is solved in constant time. For example, the variable `evBelow` is the inner sum in Eq. 4.

We seed the first iteration with sensible but arbitrary parameters (e.g., hit rate of 50%). To avoid oscillating around a stable solution, in each iteration we average the old and new distributions using an exponentially weighted moving average. We have empirically determined that a coefficient of $1/3$ yields good performance. We detect convergence when the hit rate stays within a 10^{-3} range for ten iterations. Finally, the model sets a floor of 0 for all probabilities during solution. In practice, Algorithm 1 reliably converges to the steady-state distributions after a few iterations (typically 20-40) on hundreds of thousands of distributions from real workloads (Sec. 8).

While involved, iteration is computationally cheap: in practice, we use and monitor coarse ages (see below) for which N -point distributions with $N \approx 64$ -256 suffice, and each iteration runs in linear time on the size of the distributions.

6.1 Convergence

Our model is designed with generality in mind, but this comes at the cost of complicating some theoretical properties. Eq. 1, Eq. 4, and Eq. 14 form a non-linear system (particularly Eq. 14) operating in many dimensions (N points per distribution and multiple distributions). Moreover, the model accepts arbitrary N -point vectors as input (the reuse distance distribution and ranking function). Demonstrating the convergence of fixed point iteration for non-linear systems is difficult. Generally, it involves reducing the system to a contraction mapping of some relevant model parameter [15]. Although several intuitive parameters are attractive (e.g., hit rate or modeled cache size), we cannot yet prove a contraction mapping on these parameters in general—indeed, it seems that for some degenerate ranking functions (not those in Sec. 5.1), the model does not converge.

We instead take an empirical approach. We evaluate our model on a diverse suite of real applications and demonstrate its accuracy and utility in that context. Since the model is solved online at regular intervals, our evaluation represents hundreds of thousands of model solutions. Thus we conclude that the model converges on distributions seen in practice. We leave rigorous convergence conditions to future work.

6.2 Increased step size

Reuse distances in practice can be quite large, and naive iteration over all ages would be quite expensive. Moreover,




Figure 2: Model solution with increased step size on a synthetic benchmark. Solid lines show a full-detail solution; dashed lines show a solution with $N = 32$.

age-by-age iteration is wasteful, since there are large age ranges where few events occur or the event probabilities are roughly constant. Modeling such ranges in detail is unnecessary, since they can be approximated by assuming constant event probabilities throughout the range. This observation allows the solution to take a large step over many ages at once and greatly reduces N , the number of points that need to be solved in Algorithm 1. Reducing N is important when applying the model online, e.g. as part of a runtime system (Sec. 9).

The concept is quite similar to adaptive step size in numerical solutions of differential equations [31]. For example, Fig. 2 shows a solution of all three model distributions for a synthetic benchmark (solid lines). These distributions are *coarsened* by increasing the step size (dashed lines), producing a good approximation with much less computation. Indeed, Fig. 2 shows that $N = 32$ is sufficient to model this access pattern, even though ages go up to 500. This is possible because there are large regions (e.g., ages 0-100) where few events occur. There is no reason to model these regions in great detail. Instead, we adaptively divide ages into regions, modeling regions of high activity at fine granularity and others at coarse granularity.

We then model the total probability within each coarsened region. For example, rather than computing the hit probability at a single age (e.g., $P[H = a]$), we compute the hit probability over several ages (e.g., $P[a \leq H < b]$). Remarkably, the model equations are basically unchanged by coarsening. For example, if regions are split at ages a_1, a_2, a_3, \dots then the coarsened hit equation is:

$$P[a_i \leq H < a_{i+1}] \approx \left(1 - \sum_{j=1}^{i-1} \frac{P[a_j \leq E < a_{j+1}]}{P[D > a_j]} \right) \quad (16)$$

This equation is identical in form to the fine-grain hit equation (Eq. 4), except now operating on regions rather than individual ages. Other model equations are similar (Appendix A).

Another important question is how to choose the regions. The choice must balance two competing goals: modeling regions with high activity at fine granularity, while modeling other ages in sufficient detail for the model to converge. We




Figure 3: An example implementation of our model. A lightweight hardware monitor on each tile samples a small fraction of LLC accesses. Our model runs periodically in software.

address this in two steps. First, we divide all ages evenly into $N/2$ regions. For example, with 8-bit ages and $N = 64$, we first create the 32 regions divided at ages: 0, 8, 16 . . . 256. Second, we further divide these regions $N/2$ times to try to equalize the probability of hits and evictions in each region. We sort regions by their probability of hits and evictions, and recursively divide the largest in equal-probability halves $N/2$ times. We find this procedure chooses regions that yield efficient and accurate solutions.

7. IMPLEMENTATION

We now describe how to integrate our cache model into a full system which we evaluate in simulation. In our validation (Sec. 8) and case studies (Sec. 9 and Sec. 10), the model is used to dynamically model or reconfigure the cache, as shown in Fig. 3. A lightweight hardware monitor samples a small fraction of accesses and produces the application’s reuse distance distribution. Periodically (e.g., every 50 ms), a software runtime models the cache’s behavior from the sampled reuse distance distribution, which is then used to predict the cache’s behavior over the next interval.

This configuration represents just one use case for the model; it can also be used to model cache behavior offline. However, Fig. 3 is the most demanding use case, since it imposes the most stringent run-time requirements and the model must contend with sampling error.

7.1 Application profiling

Our model works with several low-overhead, hardware monitoring schemes. For instance, PDP [9] proposes a FIFO that samples a small fraction of accesses and, when an access hits in the FIFO, records its depth in the FIFO as the reuse distance. Other monitors record stack distance and could be adapted to work with our model. For example, utility monitors [29] are a small, tag-only LRU cache that record stack distance. Geometric monitors [5] extend utility monitors to efficiently monitor very large stack distances. Stack distances can then be approximately translated into reuse distances in software [35]. Alternatively, the monitor itself can be modified to include timestamps and record reuse distance directly. In all cases, these monitors impose small overheads, typically around 1% of LLC area.

Finally, other schemes can sample the reuse distance distribution without adding hardware. Software-only schemes can sample access patterns, e.g. through injecting page faults [46]. Offline profiling can record the access trace, e.g. through compiler hooks [8] or dynamic binary translation [40]. These




Figure 4: The cache model on three synthetically generated traces (A, B, C) driving small caches using LRU, PDP, and IRGD replacement. Simulation results are shown as solid lines; model predictions as dashed lines; PDP’s model as dotted lines.

Cores	Westmere-like OOO [34], 2.4 GHz
L1 caches	32 KB, 8-way set-associative, split D/I, 1-cycle latency
L2 caches	128 KB priv. per-core, 8-way set-assoc, inclusive, 6-cycle
L3 cache	Shared, non-inclusive, 20-cycle; 16-way, hashed set-assoc
Coherence	MESI, 64 B lines, no silent drops; sequential consistency
Memory	200 cycles, 12.8 GBps/channel, 1 channel

Table 5: Configuration of the simulated systems.

schemes enable our model to work when hardware support is unavailable.

7.2 Overheads

Our model requires modest run-time overheads and small monitoring overheads, similar to prior schemes. The model takes only a few arithmetic operations per age region per iteration (~ 25). With $N = 128$ and 30 iterations on average, the model completes in under 100 K arithmetic operations. Since the model runs infrequently (e.g., every 50 ms), this overhead is small ($< 1\%$ of a single core’s cycles). If this overhead is too large, N can be reduced or the re-configuration interval can be increased, typically at little performance loss ([3], Fig. 6). Alternatively, the model can be solved in the background with low-priority threads that interfere minimally with active applications [16]. Finally, computation can be reduced by specializing the model to particular ranking functions.

8. VALIDATION

We now validate our model on synthetic and real benchmarks, showing that it is accurate over diverse replacement policies, access patterns, and cache sizes.

8.1 Synthetic

Fig. 4 compares the model against simulation of synthetic traces. These experiments demonstrate the model’s accuracy in an environment that satisfies its assumptions (Sec. 4.1). Each trace represents a different access mix: B and A expose performance cliffs in LRU and PDP, respectively, while C is a cache-friendly access pattern. Their reuse distance distributions are shown on the left. On the right there is one graph per replacement policy. On each graph, simulation results are shown (solid lines) along with model predictions

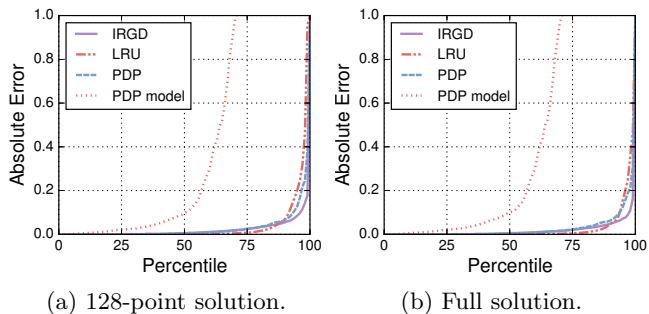


Figure 5: Model error distribution over 250 K-access intervals. Our model is accurate, while PDP’s model (dotted line) has large error.

(dashed lines). The PDP graph also includes the predictions of PDP’s analytical model (dotted lines).

Our model is accurate on every configuration. The dashed lines are often not visible because model predictions are indistinguishable from simulation results. By contrast, PDP’s model exhibits significant error from simulation, failing to distinguish between B and C and mispredicting A badly.

8.2 Execution-driven

These results carry over to real benchmarks on a full system, where model assumptions only approximately hold.

Methodology: We use zsim [34] to evaluate our model. We perform execution-driven simulation of SPEC CPU2006 benchmarks on OOO cores using 16-way hashed, set-associative last-level caches and parameters given in Table 5, chosen to mimic Ivy Bridge-IP. We run each benchmark for 10 B instructions after fast-forwarding 10 B, and we perform enough runs to achieve 95% confidence intervals $\leq 1\%$. All results hold for skew-associative caches [36], zcaches [32], and for systems with prefetching. This methodology also applies to later case studies.

We evaluate a large range of cache sizes, from 128 KB to 128 MB, and solve the model every 250 K accesses from the sampled reuse distance distribution in that interval. This yields many samples—over 400 K model solutions in all.

Results: Fig. 5 shows the distribution of modeling error ($|\text{predicted hit rate} - \text{actual hit rate}|$) in each interval for LRU, PDP, and IRGD. We show results (a) for coarsened solutions with $N = 128$ (Sec. 6.2) and (b) full solutions.




Figure 6: Sensitivity of model to step size for LRU, PDP, and IRGD, measured by the number of steps, N (Sec. 6.2). $N = \infty$ means full, age-by-age solution.

For 128-point solutions, median error is 0.1%/0.4%/0.6% for LRU/PDP/IRGD, respectively; mean error is 3.7%/2.9%/2.1%; and 90th percentile error is 6.9%/6.5%/5.5%. For full solutions, median error is 0%/0.3%/0.5%; mean error is 2.6%/2.6%/1.8%; and 90th percentile error is 4.7%/6.3%/4.8%. Overall, the model is accurate, and there is modest error from coarsening solutions. Fig. 6 shows the mean and 90th percentile error for different values of N . (Median error is negligible in all cases and not shown.) Our model is fairly insensitive to coarsening, although reducing N below 64 noticeably degrades accuracy on LRU.

It is important to emphasize that these results are for 250 K-access *intervals*—it is *not* the case that 10% of *benchmarks* have error above 6%. Rather, 10% of intervals have this error. The distinction is critical. Many benchmarks have an unstable access pattern, and their hit rate changes rapidly between intervals. Fig. 7 shows an example of *gcc* on a 1 MB LLC, plotting its hit rate over its first 25 M LLC accesses (Fig. 7a). *gcc*’s hit rate fluctuates wildly across intervals, and our model does a good job of tracking these fluctuations. (Recall that the model only uses the application’s reuse distance distribution; it does not observe the cache’s current hit rate.) However, since *gcc*’s behavior is unstable, its reuse distance distribution is not representative of its equilibrium behavior, and model error is large in some intervals. But this model error tends to average out in the long-run, and indeed our model is quite accurate at predicting *gcc*’s miss curve over its entire execution (Fig. 7b).


As a result, model error is slightly reduced in the long run (i.e., over 10 B instructions). For 128-point solutions, the long-run mean model error is 2.3%/1.9%/1.1% for LRU/PDP/IRGD, respectively, while the 90th percentile error is 5.5%/4.3%/3.4%. Results for full solutions are similarly reduced, with long-run mean error of 1.6%/1.7%/0.9% and 90th percentile error of 3.7%/4.2%/2.9%.

Hence, the model error presented in Fig. 5 is a conservative assessment of our model’s accuracy. We find that in practice the model is accurate and useful at predicting cache behavior, as we now illustrate in two case studies.

9. CASE STUDY: CACHE PARTITIONING

In this section and the next, we show how to apply our model to improve cache performance with high-performance replacement policies. This section focuses on shared caches, while the next focuses on single-threaded performance.

Prior work has demonstrated that explicit partitioning of



(a) *gcc* over 25 M accesses on a 1 MB LLC. (b) Long-run miss ratio miss ratio vs. LLC size.

Figure 7: Analysis of *gcc*. Benchmarks are unstable over time, introducing model error on some intervals. Model is accurate on long-run performance.

cache capacity can greatly improve shared cache performance [29]. Indeed, partitioned LRU outperforms thread-aware variants of high-performance replacement policies [4, 33]. Cache partitioning works by isolating each application’s data in a partition of well-defined size. Partitioning eliminates interference between applications, and since applications vary widely in how well they use cache capacity, it can greatly improve performance if partitions are sized correctly.

Partitioning algorithms use each application’s *miss curve* to size partitions. Miss curves allow the partitioning algorithm to choose sizes that minimize misses. (The optimal solution is NP-complete, but good approximations exist [3, 26, 29].) Moreover, since partitioning is done in software, it can be used to achieve system objectives other than performance; e.g., quality of service [17], fairness [28], and security [27].

In principle, cache partitioning is complementary to high-performance cache replacement: cache partitioning allocates capacity among applications, and the replacement policy then maximizes each application’s hit rate on its partition. However, cache partitioning is only effective when miss curves are available—otherwise, software cannot predict the effect of different partition sizes, and therefore cannot size them to achieve system objectives.

Cache partitioning and high-performance policies are thus in unintended conflict. Our model resolves this conflict. It can model high-performance replacement policies at arbitrary cache sizes, and thus predicts the replacement policy’s miss curve. These miss curves are given to the partitioning algorithm, which can then choose partition sizes that maximize performance (or achieve other objectives).

We evaluate the performance of a shared 4 MB LLC on a 4-core system running 100 random mixes of SPEC CPU2006 applications. We compare four schemes: (i) unpartitioned LRU, (ii) a representative thread-aware high-performance policy (TA-DRRIP [14]), (iii) LRU with utility-based cache partitioning (UCP) [29], and (iv) a high-performance policy (IRGD) with utility-based cache partitioning (UCP+IRGD). All schemes reconfigure every 50 ms and we employ a fixed work methodology [12].

Results: Fig. 8 shows the distribution of weighted and harmonic speedup over the 100 mixes [10], normalized to unpar-

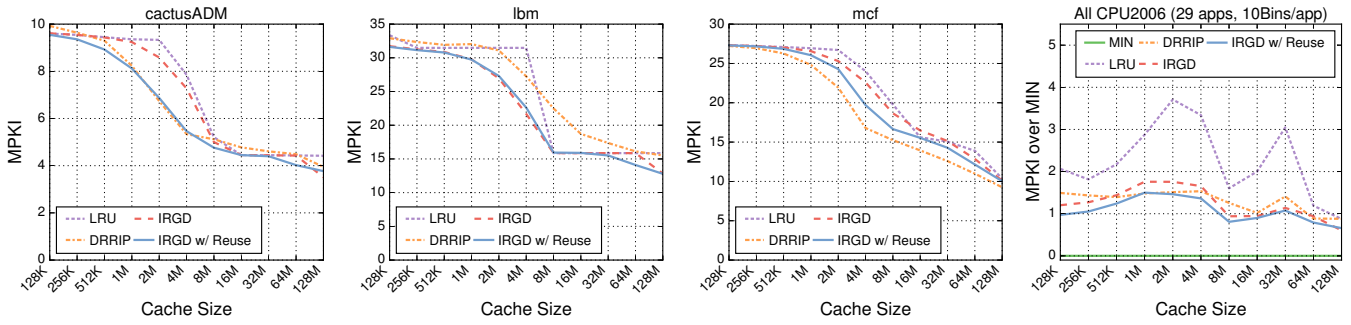


Figure 9: Misses per thousand instructions (MPKI) vs. cache size for (left) three applications and (right) all of SPEC CPU2006. Lower is better. Our model extends IRGD [39] to support reused/non-reused classification, making it outperform DRRIP on most applications and on gmean MPKI over all of SPEC CPU2006.

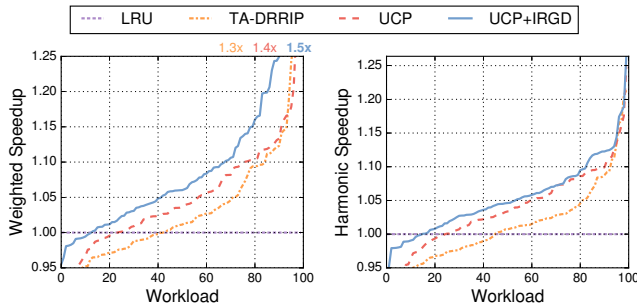


Figure 8: Weighted and harmonic speedup for 100 random mixes of SPEC CPU2006 apps. Our model enables partitioning of IRGD, improving performance over the state-of-the-art.

tioned LRU. TA-DRRIP and UCP outperform the LRU baseline respectively by 3.1%/5.1% on gmean weighted speedup, and 1.1%/3.6% on gmean harmonic speedup. Meanwhile, IRGD+UCP improves gmean weighted speedup by 9.5% and gmean harmonic speedup by 5.3%. Our model combines the single-stream benefits of high-performance cache replacement and the shared-cache benefits of partitioning, outperforming the state-of-the-art.

10. CASE STUDY: IMPROVING CACHE REPLACEMENT

As noted in Sec. 2, classifying lines into classes (e.g., reused vs. non-reused) and treating each class differently can significantly improve replacement policy performance [14, 18, 42]. However, many otherwise attractive policies do not exploit classification, or do so in limited ways, e.g. using fixed heuristics that prioritize certain classes. For instance, PDP [9] suggests classification as a promising area of future work, but lacks a modeling framework to enable it; and IRGD [39] requires several additional monitors to perform classification. Our model provides a simple analytical framework for classification. Specifically, we extend IRGD (Eq. 8) to support reused vs. non-reused classification.

We support classification by modeling the reuse distance distribution of each class, in addition to the hit and eviction distributions. Doing so requires modest changes to the model. Since policies like PDP and IRGD are reuse distance-based, this is sufficient for them to support classification. The de-

tailed changes to the model are presented in Appendix B. Crucially, our model captures the highly dynamic interactions *between* classes, which is essential for dynamic classifications like reused vs. non-reused.

Model-based classification has a further advantage. Unlike some prior high-performance policies (e.g., DRRIP [14]), we do not assume that reused lines are preferable to non-reused lines. We instead model the behavior of each class, and rely on the underlying policy (e.g., IRGD) to make good decisions within each class. This approach allows for direct comparison of candidates *across* classes according to their rank, removing the need for heuristics or tunables to penalize or prefer different classes. We thus avoid pathological performance on applications where, contra expectations, reused lines are less valuable than non-reused lines.

Results: Fig. 9 shows the LLC misses per thousand instructions (MPKI) for three representative benchmarks and average performance over all 29 SPEC CPU2006 on caches of size 128 KB to 128 MB (note log scale on x -axis). Lower MPKIs are better. Our model improves on IRGD on most SPEC CPU2006 benchmarks, and never degrades performance. On *cactusADM*, our model enables IRGD to match DRRIP’s performance, since this benchmark exhibits clear differences between reused and non-reused lines. On *lbm*, our model provides no additional benefit, but unlike DRRIP, we do not degrade performance compared to LRU. Finally, on *mcf*, our model improves performance, but does not match the performance of DRRIP.

On net, our model makes a slight but significant improvement to IRGD’s overall performance. The rightmost figure plots the MPKI over optimal cache replacement (Belady’s MIN [6]) averaged over all of SPEC CPU2006. Without classification, DRRIP outperforms IRGD on many cache sizes. With classification, IRGD outperforms DRRIP at every cache size. On average across sizes, IRGD (without classification) and DRRIP perform similarly, closing 44% and 42% of the MPKI gap between LRU and MIN, respectively. With classification, IRGD closes 52% of this gap.

11. CONCLUSION

We have presented a cache model for modern LLCs with high-performance replacement policies. Our model is motivated by observations of modern cache architecture that

allow us to abstract away details of array organization and focus on modeling the replacement policy. As a result, we capture a broad class of policies at relatively low complexity. We have presented an efficient implementation of the model and thoroughly evaluated its accuracy and implementation tradeoffs. Finally, we showed how to use the model to improve cache performance over state-of-the-art techniques.

12. REFERENCES

- [1] A. Agarwal, J. Hennessy, and M. Horowitz. An analytical cache model. *ACM Transactions on Computer Systems (TOCS)*, 7(2), 1989.
- [2] D. H. Albonesi. Selective cache ways: On-demand cache resource allocation. In *MICRO-32*, 1999.
- [3] N. Beckmann and D. Sanchez. Jigsaw: Scalable Software-Defined Caches. In *PACT-22*, 2013.
- [4] N. Beckmann and D. Sanchez. Talus: A Simple Way to Remove Cliffs in Cache Performance. In *HPCA-21*, 2015.
- [5] N. Beckmann, P.-A. Tsai, and D. Sanchez. Scaling Distributed Cache Hierarchies through Computation and Data Co-Scheduling. In *HPCA-21*, 2015.
- [6] L. A. Belady. A study of replacement algorithms for a virtual-storage computer. *IBM Sys. J.*, 5(2), 1966.
- [7] D. Chiou, P. Jain, L. Rudolph, and S. Devadas. Application-specific memory management for embedded systems using software-controlled caches. In *DAC-37*, 2000.
- [8] C. Ding and Y. Zhong. Predicting whole-program locality through reuse distance analysis. In *PLDI*, 2003.
- [9] N. Duong, D. Zhao, T. Kim, R. Cammarota, M. Valero, and A. V. Veidenbaum. Improving Cache Management Policies Using Dynamic Reuse Distances. In *MICRO-45*, 2012.
- [10] S. Eyerhan and L. Eeckhout. System-level performance metrics for multiprogram workloads. *IEEE micro*, 28(3), 2008.
- [11] C. M. Grinstead and J. L. Snell. *Introduction to probability*. American Mathematical Soc., 1998.
- [12] A. Hilton, N. Eswaran, and A. Roth. In *Proc. MoBS*.
- [13] S. Jahagirdar, V. George, I. Sodhi, and R. Wells. Power management of the third generation Intel Core micro architecture formerly codenamed Ivy Bridge. In *Hot Chips*, 2012.
- [14] A. Jaleel, K. Theobald, S. C. S. Jr, and J. Emer. High Performance Cache Replacement Using Re-Reference Interval Prediction (RRIP). In *ISCA-37*, 2010.
- [15] M. Javidi. Iterative methods to nonlinear equations. *Applied Mathematics and Computation*, 193(2), 2007.
- [16] M. Kamruzzaman, S. Swanson, and D. M. Tullsen. Inter-core prefetching for multicore processors using migrating helper threads. In *ASPLOS-XVI*, 2011.
- [17] H. Kasture and D. Sanchez. Ubik: Efficient Cache Sharing with Strict QoS for Latency-Critical Workloads. In *ASPLOS-XIX*, 2014.
- [18] G. Keramidas, P. Petoumenos, and S. Kaxiras. Cache replacement based on reuse-distance prediction. In *ICCD*, 2007.
- [19] G. Keramidas, P. Petoumenos, and S. Kaxiras. Where replacement algorithms fail: a thorough analysis. In *Proc. CF-7*, 2010.
- [20] S. M. Khan, Z. Wang, and D. A. Jiménez. Decoupled dynamic cache segmentation. In *HPCA-18*, 2012.
- [21] M. Kharbutli, K. Irwin, Y. Solihin, and J. Lee. Using Prime Numbers for Cache Indexing to Eliminate Conflict Misses. In *HPCA-10*, 2004.
- [22] N. Kurd, S. Bhamidipati, C. Mozak, J. Miller, T. Wilson, M. Nemani, and M. Chowdhury. Westmere: A family of 32nm IA processors. In *ISSCC*, 2010.
- [23] R. Manikantan, K. Rajan, and R. Govindarajan. Probabilistic shared cache management (PriSM). In *ISCA-39*, 2012.
- [24] J. Mars, L. Tang, R. Hundt, K. Skadron, and M. Soffa. Bubble-up: Increasing utilization in modern warehouse scale computers via sensible co-locations. In *MICRO-44*, 2011.
- [25] R. L. Mattson, J. Gecsei, D. R. Slutz, and I. L. Traiger. Evaluation techniques for storage hierarchies. *IBM Systems journal*, 9(2), 1970.
- [26] M. Moreto, F. J. Cazorla, A. Ramirez, R. Sakellariou, and M. Valero. FlexDCP: A QoS framework for CMP architectures. *ACM SIGOPS Operating Systems Review*, 43(2), 2009.
- [27] D. Page. Partitioned Cache Architecture as a Side-Channel Defence Mechanism. *IACR Cryptology ePrint archive*, (2005/280), 2005.
- [28] A. Pan and V. S. Pai. Imbalanced cache partitioning for balanced data-parallel programs. In *MICRO-46*, 2013.
- [29] M. Qureshi and Y. Patt. Utility-based cache partitioning: A low-overhead, high-performance, runtime mechanism to partition shared caches. In *MICRO-39*, 2006.
- [30] M. K. Qureshi, A. Jaleel, Y. N. Patt, S. C. Steely, and J. Emer. Adaptive insertion policies for high performance caching. In *ISCA-34*, 2007.
- [31] L. Richard and J. Burden. Douglas Faires, Numerical analysis, 1988.
- [32] D. Sanchez and C. Kozyrakis. The ZCache: Decoupling Ways and Associativity. In *MICRO-43*, 2010.
- [33] D. Sanchez and C. Kozyrakis. Vantage: Scalable and Efficient Fine-Grain Cache Partitioning. In *ISCA-38*, 2011.
- [34] D. Sanchez and C. Kozyrakis. ZSim: Fast and Accurate Microarchitectural Simulation of Thousand-Core Systems. In *ISCA-40*, 2013.
- [35] R. Sen and D. A. Wood. Reuse-based online models for caches. In *Proc. SIGMETRICS*, 2013.
- [36] A. Sezne. A case for two-way skewed-associative caches. In *ISCA-20*, 1993.
- [37] S. Srikantaiah, R. Das, A. K. Mishra, C. R. Das, and M. Kandemir. A case for integrated processor-cache partitioning in chip multiprocessors. In *SC09*, 2009.
- [38] Sun Microsystems. UA 2007 Family Programmer's Reference Manual. Technical report, 2009.
- [39] M. Takagi and K. Hiraki. Inter-reference gap distribution replacement: an improved replacement algorithm for set-associative caches. In *ICS'04*, 2004.
- [40] D. K. Tam, R. Azimi, L. B. Soares, and M. Stumm. RapidMRC: approximating L2 miss rate curves on

- commodity systems for online optimizations. In *ASPLOS-XIV*, 2009.
- [41] H. Wong. Intel Ivy Bridge Cache Replacement Policy, <http://blog.stuffedcow.net/2013/01/ivb-cache-replacement/>, 2013.
- [42] C.-J. Wu, A. Jaleel, W. Hasenplaugh, M. Martonosi, S. C. Steely Jr, and J. Emer. SHiP: Signature-based hit predictor for high performance caching. In *MICRO-44*, 2011.
- [43] M.-J. Wu, M. Zhao, and D. Yeung. Studying Multicore Processor Scaling via Reuse Distance Analysis. In *ISCA-40*, 2013.
- [44] H. Yang, A. Breslow, J. Mars, and L. Tang. Bubble-flux: precise online qos management for increased utilization in warehouse scale computers. In *ISCA-40*, 2013.
- [45] X. Zhang, E. Tune, R. Hagmann, R. Jnagal, V. Gokhale, and J. Wilkes. CPI 2: CPU performance isolation for shared compute clusters. In *EuroSys*, 2013.
- [46] P. Zhou, V. Pandey, J. Sundaresan, A. Raghuraman, Y. Zhou, and S. Kumar. Dynamic tracking of page miss ratio curve for memory management. In *ASPLOS-XI*, 2004.

APPENDIX

A. COMPLETE COARSENEDED MODEL

The coarsened model equations are obtained by assuming constant event probabilities within a region. Ages are broken into regions at the ages $1 = a_1 < a_2 < \dots < a_{N+1}$, giving N coarsened age regions. We start from the model equations (Eq. 1, Eq. 4, and Eq. 14) and sum over each region to derive the coarsened equations. For example, the coarsened hit equation for region i is derived as:

$$\begin{aligned} P[a_i \leq H < a_{i+1}] &= \sum_{a=a_i}^{a_{i+1}-1} P[H = a] \\ &= \sum_{a=a_i}^{a_{i+1}-1} \left(P[D = a] \times \left(1 - \sum_{x=1}^{a-1} \frac{P[E = x]}{P[D > x]} \right) \right) \end{aligned} \quad (17)$$

We now model the eviction probability (i.e., the second factor in the above equation) as approximately constant in this region, and pull it out of the sum.

$$\approx \left(\sum_{a=a_i}^{a_{i+1}-1} P[D = a] \right) \times \left(1 - \sum_{x=1}^{a_i} \frac{P[E = x]}{P[D > x]} \right) \quad (18)$$

Applying this approximation to the second factor, we break the sum over x across prior regions, and we simplify the first factor.

$$= P[a_i \leq D < a_{i+1}] \times \left(1 - \sum_{j=1}^{i-1} \frac{P[a_j \leq E < a_{j+1}]}{P[D > a_j]} \right) \quad (19)$$

The equation for the coarsened eviction distribution is similar, although it is complicated by ranking functions. A ranking function may map distinct age regions to overlapping rank regions, so there is no strict rank-ordering of ages between two regions. The correct solution to this problem is to compute the rank distribution at full fidelity, i.e. age-by-age. But doing so sacrifices much of the benefit of coarsening, since it means each iteration of the solution must consider every age. We instead coarsen ranks by representing every age within a region by the mean rank over that region, denoted $\bar{\mathcal{R}}_i$ for the i^{th} region. If regions are chosen poorly, this approach can be inaccurate, but when regions are chosen intelligently it is a good approach. The eviction distribution is thus computed:

$$\begin{aligned} P[a_i \leq E < a_{i+1}] &= \sum_{a=a_i}^{a_{i+1}-1} P[E = a] \\ &= P[\text{miss}] \times \sum_{a=a_i}^{a_{i+1}-1} \frac{P[A = a]}{P[\text{rank} = \mathcal{R}(a)]} \times P_{\max \text{rank}}(\mathcal{R}(a)) \\ &\approx P[\text{miss}] \times \sum_{a=a_i}^{a_{i+1}-1} \frac{P[A = a]}{P[\text{rank} = \bar{\mathcal{R}}_i]} \times P_{\max \text{rank}}(\bar{\mathcal{R}}_i) \end{aligned} \quad (20)$$

We model the age distribution as approximately constant in the region, so using Eq. 11 we obtain:

$$\approx P[\text{miss}] \times \frac{P[A = a_i]}{P[\text{rank} = \bar{\mathcal{R}}_i]} \times P_{\max \text{rank}}(\bar{\mathcal{R}}_i) \quad (21)$$

The age distribution is simple to derive assuming the hit and eviction distributions are constant in the region. The simplest way to derive our coarsened approximation of the age distribution is as an extension of the recurrence relation Eq. 15. In a single step, the age distribution $P_A(a)$ decreases by $P_H(a) + P_E(a)$. Hence in the coarsened region from a_i to a_{i+1} , it decreases by the total hit and eviction probability in the region. Thus Eq. 15 becomes:

$$P_A(a_i) \approx P_A(a_{i-1}) - \frac{P[a_{i-1} \leq H < a_i] + P[a_{i-1} \leq E < a_i]}{S} \quad (22)$$

B. MODELING CLASSIFICATION

For the most part, each class c can be modeled independently using a per-class reuse distance distribution $P_{D,c}(d)$ and rank function $\mathcal{R}_c(a)$ representing how the policy treats class c . Because classes share the cache, we need two additional refinements. First, the age distribution must change to reflect the relative frequency of accesses to each class. Recall from Sec. 4.2 that $P_A(1) = 1/S$ because each access creates a line of age 1. With classification, $P_{A,c}(1)$ is proportional to the total probability of accesses to class c (Algorithm 1, line 8):

$$P_{A,c}(1) = \frac{\sum_{d=1}^{\infty} P_{D,c}(d)}{S} \quad (23)$$

Thus, summing across classes, the probability at age 1 is unchanged at $1/S$.

Second, since any access can force an eviction and victims are chosen from all classes, the eviction distribution uses the cache's total hit rate and combined rank distribution. These are obtained by summing over all classes:

$$P[\text{hit}] = \sum_c \sum_{a=1}^{\infty} P_{H,c}(a) \quad (24)$$

$$P_{\text{rank}}(r) = \sum_c \sum_{a: \mathcal{R}_c(a)=r} P_{A,c}(a) \quad (25)$$

The complete model equations for access class c are:

$$P_{A,c}(a) = \frac{P[L \geq a, c]}{S} = \sum_{x=a}^{\infty} \frac{P_{H,c}(x) + P_{E,c}(x)}{S} \quad (26)$$

$$P_{H,c}(d) = P_{D,c}(d) \cdot \left(1 - \sum_{x=1}^{d-1} \frac{P_{E,c}(x)}{P[D > x, c]} \right) \quad (27)$$

$$P_{E,c}(a) = (1 - P[\text{hit}]) \cdot \frac{P_{\max \text{rank}}(\mathcal{R}_c(a)) \cdot P_{A,c}(a)}{P_{\text{rank}}(\mathcal{R}_c(a))} \quad (28)$$

along with Eq. 24 and Eq. 25. These are simply the original model equations (Eq. 1, Eq. 4, and Eq. 14) with joint probability on class c where appropriate. Note that although we focus on reused/non-reused classification, the above is a general framework that admits other classifications.

B.1 Predictive reused/non-reused classification

Performing reused/non-reused classification is challenging because being reused is a dynamic property of the cache, not an independent property of the access stream. This means we must somehow derive the reuse distance distribution of reused lines $P_{D,\text{reused}}(d)$ to perform classification, e.g. by using Eq. 8 or computing the protecting distance a la PDP. Modeling this distribution is possible because reused lines are just those that hit, and the model gives the distribution of hits in the cache (Eq. 4). So with some information about how lines behave after hitting, we can compute the reuse distance distribution of reused lines. We first model this behavior exactly, and then show a practical implementation.

Exact: To model the behavior of lines after they have hit, we need to introduce a new distribution D' , which gives a line's *next* reuse distance after the current one. D' is distributed identically to D unless the some other condition is specified. Specifically, D' is dependent on the prior reuse distance D , and D is a property of the cache *independent of the replacement policy*. Moreover, $H = d$ implies a prior reuse distance of d , so knowing the distribution of D' conditioned on D lets us reason about how lines will behave after they are reused.

In particular, the *reuse distance transition probabilities* $P[D' = d|D = x]$ let us compute the reuse distance distribution of reused lines. These are the probabilities that lines with *previous* reuse distance $D=x$ will have a *next* reuse distance of $D'=d$. Since reused lines are those that have hit in the cache, the reuse distance distribution of reused lines is:

$$P_{D,\text{reused}}(d) = \sum_{x=1}^{\infty} P_H(x) \cdot P_{D'}(d|D=x) \quad (29)$$

The distribution of non-reused lines is just the difference of this and the full reuse distance distribution.

The all-to-all transition probabilities are difficult to gather in practice, however, since they require a two-dimensional histogram indexed by previous reuse distance \times next reuse distance. The resulting distributions are expensive to sample and can have significant noise.

Simplified: We account for this by instead using a simpler distribution: the reuse distance distribution of lines that were previously reused at a distance d less than some limit ℓ . We choose the limit to serve as an effective proxy for being reused. Specifically, ℓ covers 90% of hits as computed by the model, $P[H < \ell] \approx 0.9$. We chose this parameter empirically; performance is relatively insensitive to its value.

We then compute the reuse distance distribution of reused lines, $P_{D,\text{reused}}(d)$, from whether the reuse came before or after the limit:

$$P_{D,\text{reused}}(d) \approx P[D' = d, H < \ell] + P[D' = d, H \geq \ell] \quad (30)$$

We compute each term assuming the hit probability of the *current* access is independent of the reuse distance of the

next access.³ For example, the former term:

$$\begin{aligned} P[D' = d, H < \ell] &= P[\text{hit}|D' = d, D < \ell] \cdot P[D' = d, D < \ell] \\ \text{(Independence)} \quad &= P[\text{hit}|D < \ell] \cdot P[D' = d, D < \ell] \\ \text{(Simplify)} \quad &= \frac{P[H < \ell] \cdot P[D' = d, D < \ell]}{P[D < \ell]} \\ &= P[H < \ell] \cdot P[D' = d|D < \ell] \quad (31) \end{aligned}$$

This allows us to gather only two distributions, $P_D(d)$ and $P[D' = d, D < \ell]$, in order to model reused/non-reused classification.⁴ For this model to be informative, ℓ must be chosen intelligently. We choose ℓ to cover most hits (90%).

Although this limit scheme is a rough approximation of the full transition probabilities, our evaluation (Sec. 10) shows that it captures the benefits of reused/non-reused classification with much simpler modeling and monitoring hardware.

³The same assumption is implicit in the exact derivation above, and follows nearly an identical form to the derivation that follows.

⁴The third distribution, $P[D' = d, D \geq \ell]$, is their difference.

