

# C-AMTE: A Location Mechanism for Flexible Cache Management in Chip Multiprocessors

Mohammad Hammoud, Sangyeun Cho, and Rami Melhem

*Department of Computer Science  
University of Pittsburgh*

---

## Abstract

This paper describes Constrained Associative-Mapping-of-Tracking-Entries (C-AMTE), a scalable mechanism to facilitate flexible and efficient distributed cache management in large-scale chip multiprocessors (CMPs). C-AMTE enables fast locating of cache blocks in CMP cache schemes that employ one-to-one or one-to-many associative mappings. C-AMTE stores in per-core data structures tracking entries to avoid on-chip interconnect traffic outburst or long distance directory lookups. Simulation results using a full system simulator demonstrate that C-AMTE achieves improvement in cache access latency by up to 34.4%, close to that of a perfect location strategy.

*Key words:*

CMP, Shared Scheme, Private Scheme, Associative Mapping, Fixed Mapping, Tracking Entries.

---

## 1. Introduction

Crossing the billion-transistor per chip barrier has had a profound influence on the emergence of chip multiprocessors (CMPs) as a mainstream architecture of choice. As CMPs' realm is continuously expanding, they must provide high and scalable performance. One of the key challenges to obtaining high performance from CMPs is the management of the limited on-chip cache resources (typically the L2 cache) shared by multiple executing threads/processes.

Economic, manufacturing, and physical design considerations suggest tiled CMP architectures (e.g., Tiler's Tile64 and Intel's Teraflops Research

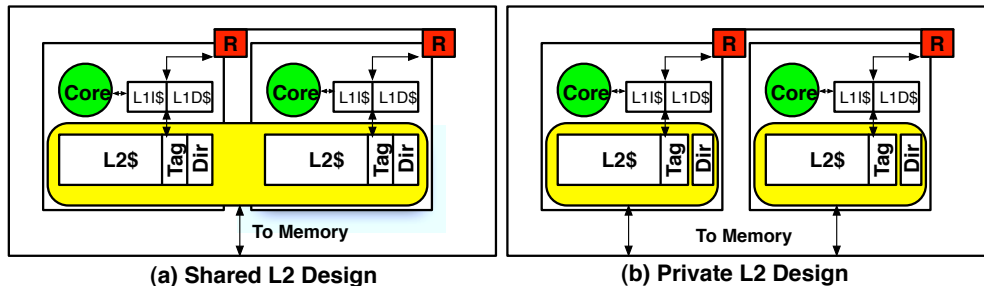


Figure 1: **Two traditional cache organizations.** (a) The shared L2 design backs up all the L1 caches. (b) The private L2 design backs up only the private L1 cache on each tile. (Dir stands for directory and R for router).

Chip) that co-locate distributed cores with distributed cache banks in tiles communicating via a network on-chip (NoC) [13]. A tile typically includes a core, private L1 caches (I/D), and an L2 cache bank. A traditional practice, referred to as the shared scheme, logically shares the physically distributed L2 banks. On-chip access latencies differ depending on the distances between requester cores and target banks creating a Non Uniform Cache Architecture (NUCA) [16]. As an example, the Intel Core<sup>TM</sup> i7 processor introduces NUCA into its platform [25]. Another conventional practice referred to as the private scheme, associates each L2 bank to a single core and provides no capacity sharing among cores. Fig. 1 demonstrates the two designs. For simplicity we show only a dual-core tiled CMP architecture. In addition, we assume a distributed directory protocol.

The private scheme replicates cache blocks at the L2 banks of the requesting cores. Hence, an effective cache associativity which equates the aggregate associativity of the L2 cache banks is provided [6]. That is, a cache block can map to any of the private L2 banks, and if shared amongst cores, can reside in multiple L2 banks. A high bandwidth on-chip directory protocol can be employed to keep the multiple L2 banks coherent. The directory can be held as a duplicate set of L2 tags distributed across tiles by set index [2, 32]. We generally refer to a mapping process that exploits the aggregate associativity of the L2 cache banks as an *associative mapping* strategy. In particular, we designate the mapping strategy of the private scheme as *one-to-many associative mapping* because a single block can be mapped to multiple L2 banks.

In contrast to the private design, the shared scheme maintains the exclu-

siveness of cache blocks at the L2 level. A core maps and locates a cache block, B, to and from a target L2 bank at a tile referred to as the *static home tile* (SHT) of B. The SHT of B is determined by a subset of bits denoted as *home select* bits (or HS bits) from B’s physical address. As such, the shared strategy requires maintaining coherence only at the L1 level. The SHT of B can store B itself and a bit vector indicating which cores had cached copies of B in their L1 private caches. This on-chip coherence practice is referred to as an in-cache coherence protocol [5, 12, 32]. In this work we refer to an entry that tracks copies (either at L1 or L2) of a certain cache block as a *tracking entry*. We, furthermore, identify a mapping process that maps an entry (block or tracking) to a fixed tile as a *fixed mapping* strategy (e.g., the shared design employs fixed mapping).

Recent research work on CMP cache management has recognized the importance of the shared scheme [27, 12, 11, 14, 29, 15]. Besides, many of today’s multi-core processors, the Intel Core<sup>TM</sup>2 Duo processor family [22], Sun Niagara [17], and IBM Power5 [26], have featured shared caches. A shared design, however, suffers from a growing on-chip delay problem. Access latencies to L2 banks are non-uniform and proportional to the distances between requester cores and target banks. This drawback is referred to as the *NUCA problem*.

To mitigate the NUCA problem, many proposals have extended the nominal basic shared design to allow associative mapping (i.e., leveraging the aggregate associativity of the L2 cache banks). For instance, block migration [3, 11, 14, 15, 31] exploits associative mapping by moving frequently accessed blocks closer to requesting cores. We denote such a strategy as *one-to-one associative mapping* due to the fact that the exclusiveness of cache blocks at the L2 level is still preserved (only a single copy of a block is promoted along identical sets over different banks). In contrast to migration, replication duplicates cache blocks at different L2 banks [8, 6, 32]. Accordingly, a replication scheme is said to adopt *one-to-many associative* mapping.

A major shortcoming of using associative mapping for blocks in any CMP cache management scheme is the location process. For example, a migration scheme that promotes a cache block B to a tile different than its home tile, denoted as the *current host* of B, can’t use anymore the HS bits of B’s physical address to locate B. Consequently, different strategies for the location process need to be considered. A tracking entry can always be retained at a centralized directory or at B’s home tile (if the underlying directory protocol is distributed) to enable tracking B after promotion. Hence, if a core

requests B, the repository of the tracking entries is reached first then the query is forwarded to B's host tile to satisfy the request. The disadvantage of this option is the arousal of 3-way cache-to-cache transfers which can degrade the average L2 access latency. An alternative location strategy could be to broadcast queries to all the tiles assuming no tracking entry for B is kept at a specific repository. Such a strategy can, however, burden the NoC and potentially degrade the overall system performance.

This paper proposes Constrained Associative-Mapping-of-Tracking-Entries (C-AMTE), a mechanism that flexibly accelerates cache management in CMPs. In particular, C-AMTE presents *constrained associative mapping* that combines the effectiveness of both, the associative and fixed mapping strategies and applies that to tracking entries to resolve the challenge of locating cache blocks without broadcasting and with minimal 3-way communications.

To summarize, the contributions of C-AMTE are as follows:

- It enables fast location of cache blocks without swamping the NoC.
- It can be applied whenever associative mapping is used for cache blocks, either in case of one-to-one (i.e, migration) or one-to-many (i.e, replication).
- It can be generally applied to cache organizations that extend the conventional private or shared schemes. Furthermore, it opens opportunities for architects to propose more creative cache management designs with no necessity to stick to either private or shared traditional paradigms.

The rest of the paper is organized as follows. Section 2 presents some recent CMP cache management schemes. C-AMTE mechanism is detailed in Section 3. In Section 4 we evaluate C-AMTE, and we conclude in Section 5.

## 2. Related Work

Much work has been done to effectively manage CMP caches. Many proposals advocate CMP cache management at either fine (block) or coarse (page) granularities and base their work on either the nominal shared or private schemes. We briefly discuss below some of the prior work and describe the location process that each proposal employs. *We note that C-AMTE is not an independent CMP scheme that can be run by itself, but yet a location*

*mechanism that can be applied to CMP designs that employ one-to-one or one-to-many associative mapping.*

Beckmann and Wood [3] and Huh et al. [14] studied generational promotion and suggested *Dynamic NUCA* (DNUCA) that migrates blocks towards banks close to requesting processors. To locate migratory blocks, [14] adopts sending concurrent queries to L2 banks. To reduce the number of queries sent over the NoC, [3] staggers the location process by searching L2 banks sequentially in an increasing order of their distances from the requester cores.

Guz et al. [11] presented a new architecture that utilizes migration to divert only *shared data* to cache banks at the center of the chip close to all the cores. To locate migratory blocks, sequential, hybrid (between sequential and broadcast), and sequential with predictor policies have been scrutinized. Kandemir et al. [15] proposed a mechanism that determines a suitable location for a data block, B, within the shared L2 space at any given point during execution and then migrates B to that suitable place. To locate B, a multistep checking scheme was employed.

Zhang and Asanović [31] examined direct promotion (upon first touch) and proposed *Victim Migration* that migrates a cache block, B, from its home tile to the initial requester tile. A victim migration (VM) table per each tile was suggested to keep track of the locations of migratory blocks. Specifically, a migration tag for B is kept in the VM table at B's home tile to point to the current host of B. Later if a sharer core S reaches the home tile of B and fails to find a matching tag in the regular L2 tag array but hits in the associated VM table, the current host of B, pointed out by the matched migration tag, satisfies the request using a 3-way cache-to-cache transfer. Clearly, Victim Migration fails to exploit *distance locality*. That is, the request of a sharer core S might incur significant latency to locate B (due to approaching B's home tile), though B might reside in close proximity to S.

Marty and Hill [19] proposed imposing a two-level virtual coherence hierarchy on a physically flat CMP that harmonizes with virtual machines (VMs) assignments. A key challenge for an intra-VM protocol is to find the home tile of a requested block. For an intra-VM, the home tile is a function of two properties: which tiles belong to a VM and how many tiles belong to a VM. Awkwardly, a dynamic VM reassignment can change both. As such, they suggest co-locating caches with tables within tiles. A table must be looked up before a miss leaves a tile. Each table includes 64 six-bit entries indexed by the six least-significant bits of the block number. Tables would be set by a hypervisor (or OS) at a VM (or process) reassignment.

Hammoud et al. [12] proposed an adaptive controlled migration (ACM) scheme that relies on prediction to collect accessibility information regarding cores that accessed a block B in the past, and then assuming that each of these cores will access B again in the future, dynamically migrates B to a bank that minimizes the overall network hops needed. To locate cache blocks, the cache-the-cache-tag (CTCT) location policy has been suggested. CTCT is a specific version of the C-AMTE mechanism and had been presented in [12] specifically to perform blocks' locations for ACM. This paper generalizes CTCT (now C-AMTE) to enable fast locating of cache blocks in CMP cache schemes that adopt one-to-one (i.e., migration) or one-to-many (i.e., replication) associative mappings.

Cho and Jin [9] proposed an OS-based page allocation algorithm applicable to NUCA architectures. Cache blocks are mapped and located to L2 banks using interleaving on page frame numbers. Chaudhuri [7] suggested PageNUCA which employs data migration at page granularity. Hardvellas et al. [13] presented R-NUCA that relies on OS to classify cache accesses into either private, shared, or instructions and then places and locates each differently at the L2 cache space. Both, PageNUCA and R-NUCA adopt direct location strategies similar to C-AMTE. In Section 3.5 we detail the two schemes and compare and contrast them versus C-AMTE.

Lastly, many researchers explored data replication instead of migration to mitigate the NUCA latency problem. Zhang and Asanović [32] proposed victim replication (VR) scheme based on the nominal shared design. VR keeps replicas of local primary cache victims within *only* the local L2 cache banks. As such, the location process becomes straightforward: local L2 banks are looked up (seeking for replica hits) before potentially checking with blocks' home tiles. However, many other cache schemes don't limit themselves to replicating blocks at only local L2 banks. Chang and Sohi [6] proposed *cooperative caching* based on the private scheme, and created a globally managed shared aggregate on-chip cache. Chisti et al. [8] proposed *CMP-NuRAPID* that controls replication based on usage patterns. Both, [6] and [8] utilize 3-way cache-to-cache transfers to satisfy L2 requests upon misses at local L2 banks.

### 3. The Proposed Mechanism

#### 3.1. Description of the Mechanism

Constrained Associative-Mapping-of-Tracking-Entries (C-AMTE) is not an autonomous CMP cache organization that can run by itself but rather a mechanism that can be applied to CMP cache designs that employ one-to-one (i.e., migration) or one-to-many (i.e., replication) associative mappings. A shared NUCA architecture maps and locates a cache block, B, to and from a *home tile* determined by a subset of bits (home select or HS bits) from B's physical address. Accordingly, B might be mapped to a bank far away from the requester core, causing the core significant latency to locate B. Such a problem is referred to as the *NUCA problem*. Migration and replication have been suggested as techniques to alleviate the NUCA problem. To save latency on subsequent requests to B, migration and replication relocate and replicate, respectively B at a tile different than its home tile, denoted as the *host tile* of B, closer to requesting cores. Consequently, B can have, in addition to the home tile, one or more host tiles. To locate B at a host tile, the HS bits of B's physical address can't be used anymore. C-AMTE offers a robust and versatile location strategy to locate B at host tiles.

Assuming a distributed directory protocol, C-AMTE supports storing one tracking entry corresponding to a block B at the home tile of B. We refer to this tracking entry as the *principal* tracking entry. The principal tracking entry points to B and can always be checked by any requester core to locate B at its current host. The principal tracking entry is stored using a fixed mapping strategy because the home tile of B is designated by the HS bits of B's physical address. C-AMTE also supports storing another type of tracking entries for B at requester tiles. We refer to this type of tracking entries as *replicated* tracking entries. A replicated tracking entry at a requester tile also points to the current host of B but can be rapidly checked by a requester core to directly locate B (instead of checking with B's home tile to achieve that). The idea of replicating tracking entries at requester tiles capitalizes on the one-to-many associative mapping strategy traditionally applied for cache blocks. C-AMTE combines associative and fixed mapping strategies and apply that to tracking entries in order to efficiently solve the location problem. Table 1 illustrates the hybrid approach adopted by the C-AMTE mechanism. We refer to such a hybrid mapping process as a *constrained associative mapping* strategy.

	BLOCK MAPPING	TRACKING ENTRIES MAPPING
<i>Private Scheme (P)</i>	Associative (at requesting tiles)	Fixed (at home tiles)
<i>Shared Scheme (S)</i>	Fixed (at home tiles)	Fixed (at home tiles)
<i>Scheme With C-AMTE</i>	Associative (one-to-one or one-to-many depending on the underlying cache scheme)	Constrained=Fixed (at home tiles) + Associative (at requesting tiles)

Table 1: Mapping strategies of private and shared CMP caches and the hybrid mapping approach of C-AMTE.

Based on the above discussion, per each tile,  $T$ , a principal tracking entry is kept for each cache block  $B$  whose home tile is  $T$  but had been mapped/promoted to another tile. Besides, replicated tracking entries are retained at  $T$  to track the locations of other corresponding cache blocks that have been recently accessed by  $T$  but whose home tile is not  $T$ . Though both, principal and tracking entries essentially act as pointers to the current hosts of cache blocks, we differentiate between them for consistency and replacement purposes (more on this shortly). We can add two distinct data structures per each tile to store the two types of the tracking entries. A data structure, referred to as the principal tracking entries (PTR) table, can hold principal tracking entries, and a data structure, referred to as the replicated tracking entries (RTR) table, can hold replicated ones. Alternatively, a single table, could be referred to as the tracking entries (TR) table, can be added to hold both classes of tracking entries pertaining that a hardware extension (i.e., an indicative bit) is engaged to distinguish between the two entries.

Assume a CMP organization with PTR and RTR tables. Whenever a core issues a request to a block  $B$ , its RTR table is checked first for a matching replicated tracking entry. C-AMTE then proceeds as follows:

- On a miss at the RTR table, the home tile of  $B$  is reached and its PTR table is looked up.
  - If a miss occurs at the PTR table,  $B$  is fetched from the main memory and mapped to a tile  $T$  specified by the underlying cache scheme protocol. If  $T$  is not  $B$ 's home tile, principal and replicated tracking entries are stored at the PTR table of  $B$ 's home tile and the RTR table of the requester core, respectively. If, in contrary,  $T$  is  $B$ 's home tile, no tracking entries are kept at either PTR or RTR tables ( $B$  can be located directly using the HS bits of  $B$ 's physical address).
  - If, on the other hand, a hit occurs at the PTR table,  $B$  is located at its current host tile and a replicated tracking entry is stored at



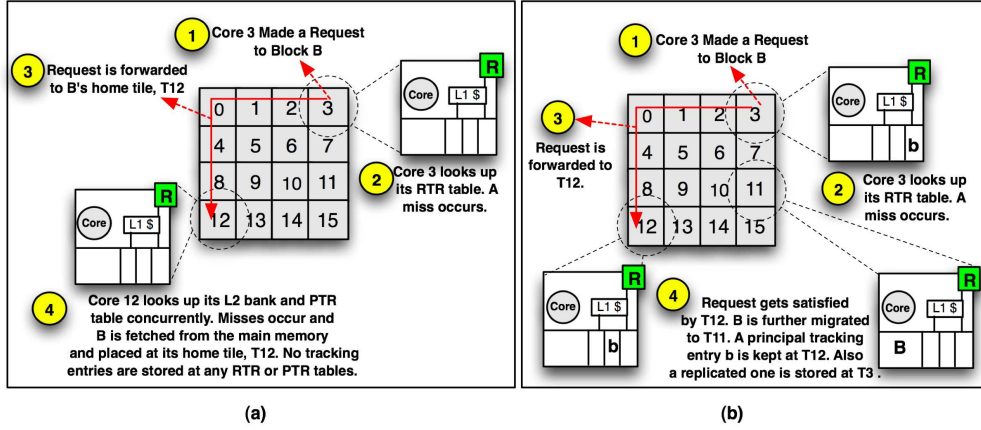


Figure 2: A first example on locating a migratory block B using the C-AMTE mechanism.

the requester’s RTR table.

- On a hit at the RTR table, B is located directly at its current host designated by the matched replicated tracking entry.

Therefore, upon a hit to the requester’s RTR table, a 3-way cache-to-cache transfer, which would have been incurred if we had to approach B’s home tile to locate B, is avoided. Similar logic applies if C-AMTE assumes a single TR table instead of two distinct PTR and RTR ones.

### 3.2. Illustrative Examples

Fig. 2 demonstrates an example of the C-AMTE mechanism on a tiled CMP platform, assuming an underlying shared scheme and a migration policy that promotes cache blocks towards requesting cores. Fig. 2(a) shows a request made by core 3 to a cache block, B. Core 3 looks up its local RTR table. We assume a miss occurs and the request is subsequently forwarded to B’s home tile, T12. The PTR table and the regular L2 bank at T12 are looked up concurrently. We assume misses occur at both. Consequently, B is fetched from the main memory and mapped to B’s home tile, T12 (following the mapping strategy of the nominal shared scheme). As such, no tracking entries are retained at either PTR or RTR tables. Fig. 2(b) shows a subsequent request made by core 3 to B. B is located at its home tile, T12. Assume after that hit, B is migrated to T11 (closer to T3). Thus, corresponding principal and

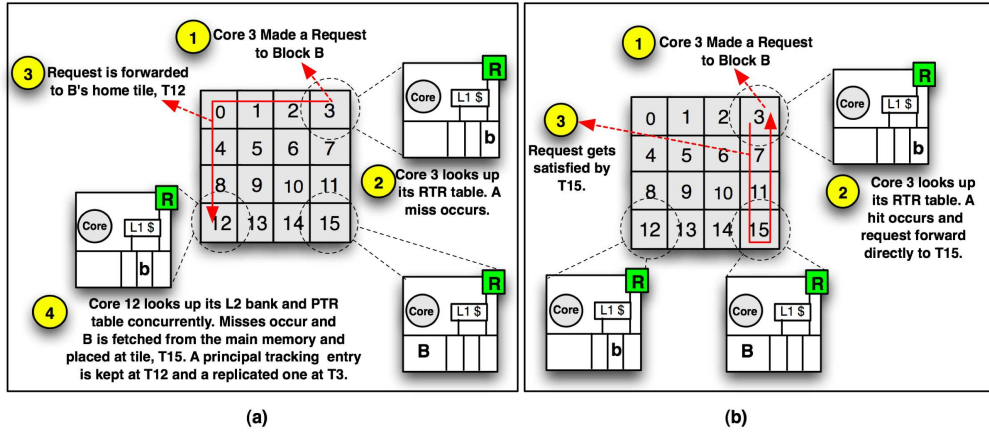


Figure 3: A second example on locating a block B using the C-AMTE mechanism.

replicated tracking entries are stored at T12 and T3, respectively. If at any later time core 3 requests B again, a hit will occur at its RTR table (as long as the entry has not been replaced yet) and B can be located straightforwardly at T11 avoiding thereby 3-way cache-to-cache transfers. Lastly, note that if any other core requests B, T12 can always indirectly satisfy the request and a corresponding tracking entry can be stored at the new requester's RTR table.

Fig. 3 demonstrates C-AMTE in operation assuming a cache scheme that might map cache blocks to tiles different than their home tiles. Fig. 3(a) shows a request made by core 3 to a cache block B. Core 3 looks up its local RTR table. We assume a miss occurs and the request is subsequently forwarded to B's home tile, T12. The PTR table and the regular L2 bank at T12 are looked up concurrently and misses are then incurred. Consequently, B is fetched from the main memory and mapped to T15 (determined by the mapping strategy of the cache scheme). As such, principal and replicated tracking entries are kept at T12 and T3, respectively. Fig. 3(b) shows a request made again by core 3 to B. A hit occurs at T3's RTR table. Consequently, B is directly located at T15. Clearly, the two examples shown in Figures 2 and 3 reveal the efficiency and versatility of C-AMTE as a strategy that exploits distance locality. C-AMTE, in fact, opens opportunities for architects to propose creative block migration, replication, and placement CMP strategies with the required location process being on-hand.

### 3.3. Maintenance and Coherence of the Tracking Entries

The principal and replicated tracking entries need to be kept coherent. We accomplish this by embedding a bit vector with each principal tracking entry at the PTR tables to indicate which cores had cached related replicated tracking entries at their RTR tables (much similar to the in-cache coherence protocol in [5]). For instance, given the example depicted in Fig. 2, each time B is migrated to a different tile, the principal and the replicated tracking entries that correspond to B are updated to point to the new host of B. Besides, C-AMTE can easily preclude potential false misses that can occur when L2 requests fail to hit on cache blocks because they are in transit between L2 banks. When migration is to be performed, a copy, B', of the cache block B is kept at the current bank so as if an L2 request arrives while B is in transit, the request is immediately satisfied without incurring any delay. When B reaches the new host, an acknowledgement message is sent back to the old host to discard B'. The old host keeps track of any tile that accesses B', and when receiving the acknowledgment message, sends an update message to the new host to indicate the new sharers that requested B while it was in transit. The directory state entry of B is consecutively updated. Clearly, enforcing coherence among tracking entries and precluding false misses impose traffic overhead on the network on-chip. Section 4.2 demonstrates the increase in message hops per 1K instructions incurred by the C-AMTE mechanism.

Finally, PTR and RTR tables can employ the LRU replacement policy. However, in case of a single TR table, it is wise to never evict a principal tracking entry in favor of a replicated one (this is the reason of why we suggested distinguishing between the two entries). An eviction of a principal tracking entry causes evictions to the corresponding cache block and all the related replicated tracking entries. Therefore, the TR replacement policy should replace the following three classes of entries in an ascending order: (1) an invalid entry, (2) an LRU replicated tracking entry, (3) and an LRU principal tracking entry. Besides, upon storing a replicated tracking entry, only the first two classes are considered for replacement. If no entry belonging to one of these two classes is detected, a replicated tracking entry is not retained.

### 3.4. Hardware Cost and Scalability

The storage overhead incurred by C-AMTE pertains to the usage of principal and replicated tracking entries. As described earlier, C-AMTE incurs at least

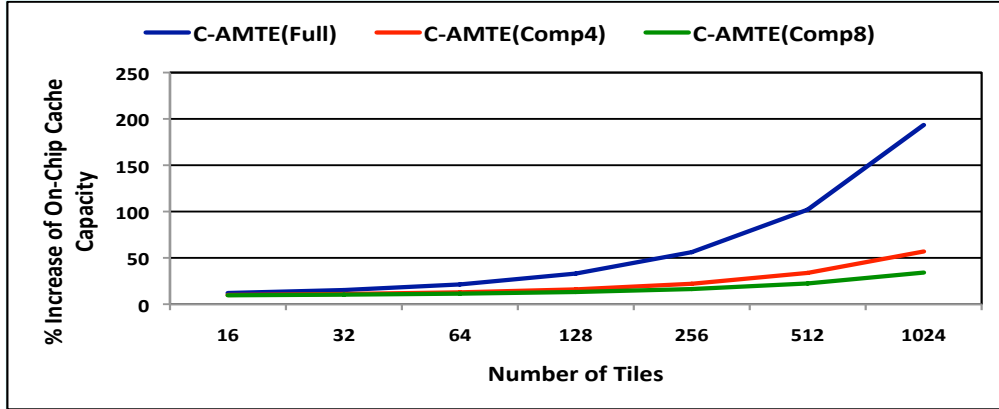


Figure 4: Storage requirements of C-AMTE with a full-map bit vector (C-AMTE(Full)), a compact vector with 1 bit for every 4 cores (C-AMTE(Comp4)), and a compact vector with 1 bit for every 8 cores (C-AMTE(Comp8)).

one principal and one replicated tracking entries per each cache block,  $B$ , when placed at a tile different than its static home tile. On the other hand, C-AMTE incurs at most  $N - 1$  tracking entries (one principal and the rest are replicated) per  $B$  with  $N$  being the number of tiles on the CMP platform. The worst case scenario occurs only when  $B$  exhibits a sharing degree of  $N$ . Assuming split tracking entries tables, each principal tracking entry would include: (1) the tag of  $B$  (typically 22 bits), (2) a bit vector that acts as a directory to keep the principal and the replicated tracking entries coherent (e.g., 16 bits for a 16-tile CMP model), and (3) an ID that points to the tile that is currently hosting  $B$  (e.g., 4 bits for a 16-tile CMP model). On the other hand, a replicated tracking entry includes only  $B$ 's tag and the ID to  $B$ 's current host tile. In contrast, in case of a single TR table, both the principal and the replicated tracking entries would each encompass a tag, a bit vector, an ID, and an indicative bit to distinguish between the two types of entries (required for replacement purposes). Clearly, the bit vector added to each replicated entry becomes in this case redundant. Thus, splitting TR table into RTR and PTR might be preferable for reducing storage overhead.

Assuming a 16-tile CMP where each tile encompasses 32KB I/D L1 caches and a 512KB L2 cache bank and assume PTR and RTR tables each with 8K entries per tile, C-AMTE demonstrates a 12% increase of on-chip cache capacity. To illustrate how the area overhead of C-AMTE scales, Fig. 4 shows

the storage requirements of C-AMTE under 16-tile, 32-tile, 64-tile, 128-tile, 256-tile, 512-tile, and 1024-tile platforms. The figure shows that C-AMTE with full-map bit vector (one bit for every core) per each principal tracking entry (C-AMTE(Full)) scales poorly especially after involving more than 64 cores on a single chip. Clearly, what makes C-AMTE non-scalable to a large number of tiles is the bit vector associated with each principal tracking entry. C-AMTE, however, needs not incorporate full-map vectors. Similar to sparse directories [10] and SGI Origin style design [18], C-AMTE can involve more compact (coarse) vectors to improve upon the poor scalability at a moderate bandwidth increase. For instance, a bit vector can contain one bit for every four cores (C-AMTE(Comp4)), or one bit for every eight cores (PDA(Comp8)) and rely on a broadcast or multicast protocol to track replicated tracking entries.

### *3.5. Qualitative Comparison with Closely Related Designs*

Two of the closely related location strategies are those proposed and utilized by PageNUCA [7] and R-NUCA [13]. Chaudhuri [7] suggested PageNUCA which employs data migration at page granularity. Access patterns of cores are dynamically monitored and pages are migrated to banks that minimize the access time for the sharing cores. To locate the migratory pages at the L2 space, each core maintains at the L1 level two tables (organized exactly as TLBs) that map the original physical frame number of an instruction or data page to the migrated frame number. These tables are referred to as iL1Map and dL1Map, respectively. Upon each L2 request, the appropriate table is looked up before routing the request to the correct L2 bank. An entry in the appropriate L1Map is loaded from another unified map table (L2Map) maintained at the L2 level when the corresponding page table entry is loaded in the TLB at the time of a TLB miss. On a migration, the new physical frame number of a page is sent to the sharing cores (with the help of a sharing vector maintained at a table referred to as PACT) so that they can update their L1Map tables appropriately.

Hardavellas et al. [13] proposed R-NUCA that also relies on OS. R-NUCA classifies cache accesses to either private, shared, or instructions. Private pages are placed at the local L2 banks of the requesting cores, shared at fixed address-interleaved on-chip locations, and instructions at non-overlapping fixed-center clusters of L2 banks. R-NUCA extends page table and TLB entries to distinguish between private and shared pages. A request to L2 after a miss at the L1-I cache is immediately classified as an instruction and

a direct location is simply performed assuming a fixed-center cluster centered at the requesting core. On the other hand, a request to L2 after a miss at an L1-D cache is resolved during the virtual-to-physical translation. The corresponding TLB (or page table in case of a TLB miss) entry is examined to decide upon whether the requested page is private or shared. Subsequently, the request is routed to either the local (if the page is private) or home (if the page is shared) L2 bank.

Clearly, both R-NUCA and PageNUCA employ *direct* location strategies similar to C-AMTE. However, two main things differentiate C-AMTE from them. First, R-NUCA and PageNUCA are page-granular schemes that involve OS while C-AMTE is a block-granular strategy that is fully transparent and doesn't involve OS. Second, R-NUCA and PageNUCA are not pure location strategies but rather placement and migration designs, respectively. R-NUCA changes the original mappings of blocks and PageNUCA migrates blocks from their original home banks. As such, they both require ways to locate the L2 cache blocks that don't reside anymore at their home banks. Their suggested location strategies are *specific* to their proposed mechanisms. In contrast, C-AMTE is a pure location strategy that is general (i.e., not specific to any cache scheme). C-AMTE enables any block-granular placement and migration schemes.

#### 4. Quantitative Evaluation

In this work we assume a baseline block-interleaved shared CMP cache organization. We study C-AMTE with an implementation of the DNUCA scheme [14, 3]. We employ DNUCA on our tiled CMP architecture via allowing migration in vertical and horizontal directions seeking to reduce hit latencies. Each cache block is augmented by four 2-bit saturation counters in correspondence to the four plausible ways: north, south, west, and east. Once a counter saturates, its value is cleared and the block is migrated towards the indicated direction (i.e., promoted up, down, left, or right one tile upon the saturation of the north, south, west, or east counter, respectively). To locate cache blocks after migration, C-AMTE is utilized. We refer to this DNUCA implementation with C-AMTE being incorporated as DNUCA(C-AMTE).

To demonstrate the potential performance gain from C-AMTE we compare DNUCA(C-AMTE) against the baseline shared (S) scheme and three other DNUCA implementations that are only different in their location processes. First, we consider DNUCA with a broadcast location strategy. That

COMPONENT	PARAMETER
<i>Cache Line Size</i>	64 B
<i>L1 I/D-Cache Size/Associativity</i>	16KB each/2way
<i>L1 Read Penalty (on hit per tile)</i>	1 cycle
<i>L1 Replacement Policy</i>	LRU
<i>L2 Cache Size/Associativity</i>	512KB per L2 bank or 8MB aggregate/16way
<i>L2 Bank Access Penalty</i>	12 cycles
<i>L2 Replacement Policy</i>	LRU
<i>Latency Per NoC Hop</i>	3 cycles
<i>Memory Latency</i>	300 cycles

Table 2: **System parameters**

NAME	INPUT
<i>SPECJbb</i>	Java HotSpot (TM) server VM v 1.5, 4 warehouses
<i>Ocean</i>	514×514 grid (16 threads)
<i>Barnes</i>	32K particles (16 threads)
<i>Lu</i>	2048×2048 matrix (16 threads)
<i>Radix</i>	3M integers (16 threads)
<i>Bodytrack</i>	4 frames and 1K particles (16 threads)
<i>Fluidanimate</i>	5 frames and 300K particles (16 threads)
<i>Swaptions</i>	64 swaptions and 20K simulations (16 threads)
<i>MIX1</i>	Hmmer (reference) (16 copies)
<i>MIX2</i>	Sphinx (reference) (16 copies)
<i>MIX3</i>	Barnes, Ocean, Radix, Lu, Milc (ref), Mcf (ref), Bzip2 (ref), and Hmmer (2 threads/copies each)
<i>MIX4</i>	Barnes, FFT (4M complex numbers), Lu, and Radix (4 threads each)

Table 3: **Benchmark programs**

is, queries to all tiles are sent upon every L2 request to locate the required block. We denote this implementation as DNUCA(B). Second, a 3-way cache-to-cache transfer strategy is employed similar to the one in [31]. This implementation is designated as DNUCA(3W). Lastly, we consider DNUCA with an ideal location strategy to set an upper bound for C-AMTE and see how close it draws to a perfect approach. The ideal strategy assumes that cores have oracle knowledge about the on-chip residences of blocks. Hence, every L2 request is routed directly to the correct L2 bank. We refer to such an implementation as DNUCA(Ideal).

#### 4.1. Methodology

We present our results based on a detailed full system simulation using Virtutech’s Simics 3.0.29 [1]. We use our own CMP cache modules fully developed in-house. We implement the XY-routing algorithm and accurately model congestion for all types of messages. A tiled CMP architecture comprised of 16 UltraSPARC-III Cu processors is simulated running with Solaris 10 OS. Each processor uses an in-order core model. The tiles are organized as a  $4 \times 4$  grid connected by a 2D mesh network on-chip (NoC). A 3-cycle latency (in addition to the NoC congestion delay) per hop is incurred when a datum traverses through the mesh network [32, 31]. Each tile encompasses a switch, an aggregate 32KB I/D L1 cache, a 512KB L2 cache bank, and a tracking table (TR) with 16K entries. The latency to lookup a TR table is hidden under the delay to enqueue the request in the port scheduler of the local switch [7]. Lastly, for coherence enforcement at the L1 cache level, a distributed in-cache MESI-based directory protocol is employed (fully verified and tested). Table 2 details our configuration’s experimental parameters.

We use a mixture of multithreaded and multiprogramming workloads to study the five designs, S, DNUCA(B), DNUCA(3W), DNUCA(C-AMTE), and DNUCA(Ideal). For multithreaded workloads we use the commercial benchmark SpecJBB [28], five shared memory programs from the SPLASH-2 suite [30] (Ocean, Barnes, Lu, Radix, and FFT), and three applications from the PARSEC suite [4] (Bodytrack, Fluidanimate, and Swaptions). We composed multiprogramming workloads using the above considered SPLASH-2 benchmarks and five other applications from SPEC2006 [28] (Hmmer, Sphinx, Milc, Mcf, and Bzip2). Table 3 shows the data sets and other important features of the simulated workloads. Lastly, the programs are fast forwarded to get past of their initialization phases. After various warm-up periods, each SPLASH-2 and PARSEC benchmark is run until the completion of its main loop, and each of SpecJBB, MIX1, MIX2, MIX3, and MIX4 is run for 20 billion user instructions.

#### 4.2. Results

Fig. 5 demonstrates the average L2 access latency (AAL) of S, DNUCA(B), DNUCA(3W), DNUCA(C-AMTE), and DNUCA(Ideal) schemes normalized to S. The incurred latency per L2 access is defined depending on three scenarios. First, it can involve only the L2 access time. This happens when a hit occurs to a local L2 bank from a requesting core. Second, it can incorporate distance latency (computed in terms of the number of hops traversed between



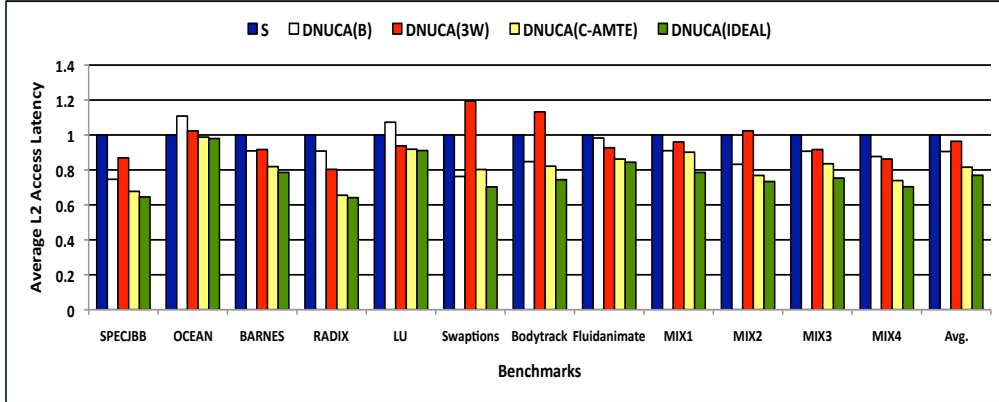


Figure 5: Average L2 access latency of the baseline shared scheme (S), DNUCA(B), DNUCA(3W), DNUCA(C-AMTE), and DNUCA(Ideal) normalized to S (B= Broadcast, 3W = 3 Way).

	S	DNUCA(B)	DNUCA(3W)	DNUCA(C-AMTE)	DNUCA(IDEAL)
<i>SPECjbb</i>	5.3	87.5	5.7	4.8	2.4
<i>Ocean</i>	2.5	35.8	3	3.1	2.4
<i>Barnes</i>	3.6	55.1	4.4	4	2.9
<i>Radix</i>	6.9	136.4	9.8	13.5	9.4
<i>Lu</i>	70	905.4	78.3	76	70.5
<i>Swaptions</i>	4.8	64.3	6.6	7.2	3.2
<i>Bodytrack</i>	5.2	95	8.5	11.3	4.9
<i>Fluidanimate</i>	11.3	174.9	11.88	11.82	10.3
<i>MIX1</i>	35.5	573.8	37.3	37.6	27.4
<i>MIX2</i>	22.1	370.2	32.6	47	19
<i>MIX3</i>	11.6	168	16.4	14.9	10.3
<i>MIX4</i>	50.8	691.7	54.8	80	26

Table 4: Message-Hops per 1K instructions

the requester and the target tiles and the observed NoC congestion delay) and the L2 access time. This occurs upon a hit to a remote L2 bank. Third, it can involve memory latency because of a miss on L2. DNUCA(C-AMTE) achieves AAL improvement over S by an average of 18.4%, and by as much as 34.4% for Radix. This makes DNUCA(C-AMTE) significantly close to DNUCA(Ideal) which accomplishes, in contrast, an average AAL improvement of 23%. DNUCA(C-AMTE) doesn't draw nearer to DNUCA(Ideal) because of two main reasons: (1) misses to TR tables by requester cores and (2) overhead to keep the principal and the replicated tracking entries coherent after blocks' migrations. Consequently, DNUCA(C-AMTE) generates a higher NoC traffic which causes more NoC delay and, subsequently, inferior

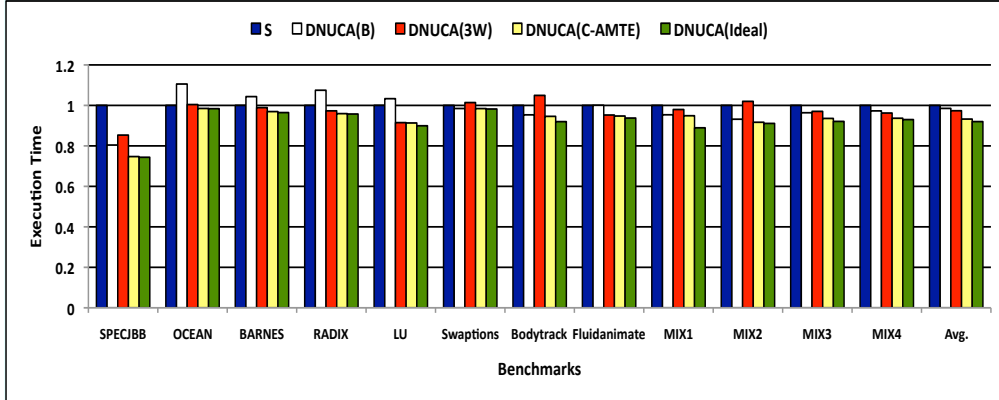


Figure 6: Execution times of the baseline shared scheme (S), DNUCA(B), DNUCA(3W), DNUCA(C-AMTE), and DNUCA(Ideal) normalized to S (B=Broadcast, 3W = 3 Way).

AAL accomplishment. Table 4 shows the number of message-hops per 1K instructions experienced by all the studied schemes for the examined benchmark programs. A message-hop is defined as one message travelling one hop on a router in the 2D mesh NoC.

As shown in Fig. 5, DNUCA(B) and DNUCA(3W) provide AAL improvements over S by averages of 9.4% and 3.6%, respectively. DNUCA(B) is similar to DNUCA(Ideal) in that it offers a direct locations for cache blocks. However, DNUCA(B) profoundly outbursts the NoC with superfluous queries. This causes more NoC delay which translates to a lower AAL improvement. Two factors determine the eligibility of an application to accomplish a higher or a lower AAL under DUNCA(B): (1) the gain,  $G$ , out of direct locations to cache blocks and (2) the loss,  $L$ , from congestion delay. When  $L$  is offset by  $G$ , DNUCA(B) improves AAL (e.g., SpecJBB), otherwise, a degradation over S is observed (e.g., Ocean). DNUCA(3W), on the other hand, fails to exploit distance locality and is expected, accordingly, not to surpass S. Nonetheless, most of the applications experience AAL improvement under DNUCA(3W) (SpecJBB, Barnes, Radix, Lu, Fluidanimate, MIX1, MIX3, MIX4). This improvement comes, in fact, from the fewer off-chip accesses attained by DNUCA. Computer programs exhibit large asymmetry in cache sets' usages [23, 20]. DNUCA inadvertently equalizes the non-uniformity across cache sets via the employment of the one-to-one associative mapping.

To that end, Fig. 6 presents the execution times of S, DNUCA(B), DNUCA(3W), DNUCA(C-AMTE), and DNUCA(Ideal) normalized to S. Across all benchmarks, DNUCA(B), DNUCA(3W), DNUCA(C-AMTE), and DNUCA(Ideal) outperform S by averages of 1.4%, 2.6%, 6.7%, and 8%, respectively. Though DNUCA(B) accomplished 9% and 9.2% AAL reductions over S for Barnes and Radix respectively, this didn't effectively translate to an improvement in the overall system performance.

## 5. Concluding Remarks and Remaining Work

Cache management in CMP is crucial to fuel its performance growth. This paper proposes C-AMTE, a mechanism that effectively simplifies the process of locating cache blocks in CMP caching schemes that employ either one-to-one or one-to-many associative mappings. C-AMTE stores tracking entries that correspond to cache blocks at per-core data structures for direct locations at subsequent accesses. We demonstrated the effectiveness of C-AMTE by applying it to the DNUCA [3, 14] migration scheme (i.e., a scheme that adopts one-to-one associative mapping). A performance improvement of up to 25.2% has been achieved, close to that of a perfect location strategy.

Lastly, having established the effectiveness of C-AMTE, optimizations to reduce hardware cost, a sensitivity study to different RT table sizes (or alternatively RTR and PTR tables), alternatives on evicting principal tracking entries, protocols on replacing blocks, and applying C-AMTE to more CMP caching schemes, specifically to schemes that incorporate one-to-many associative mapping (e.g., replication schemes), are among the obvious future directions.

## References

- [1] Virtutech AB. Simics Full System Simulator “<http://www.simics.com/>”
- [2] L. Barroso et al. “Piranha: A Scalable Architecture Based on Single-Chip Multiprocessing,” *ISCA*, May 2000.
- [3] B. M. Beckmann and D. A. Wood. “Managing Wire Delay in Large Chip-Multiprocessor Caches,” *MICRO*, Dec. 2004.
- [4] C. M. Bienia, S. Kumar, J. P. Singh, and K. Li. “The PARSEC Benchmark Suite: Characterization and Architectural Implications,” *PACT*, Oct. 2008.
- [5] L. Censier and P. Feautrier. “A New Solution to Coherence Problems in Multicache Systems,” *IEEE Trans. Comput. C-27 (12): 1112- 1118*, Dec. 1978.
- [6] J. Chang and G. S. Sohi. “Cooperative Caching for Chip Multiprocessors,” *ISCA*, June 2006.
- [7] M. Chaudhuri. “PageNUCA: Selected Policies for Page-grain Locality Management in Large Shared Chip-multiprocessor Caches,” *HPCA*, pp. 227-238, Feb. 2009.

- [8] Z. Chishti, M. D. Powell, and T. N. Vijaykumar. “Optimizing Replication, Communication, and Capacity Allocation in CMPs,” *ISCA*, June 2005.
- [9] S. Cho and L. Jin. “Managing Distributed Shared L2 Caches through OS-Level Page Allocation,” *MICRO*, Dec 2006.
- [10] A. Gupta, W. D. Weber, and T. Mowry. “Reducing Memory and Traffic Requirements for Scalable Directory-Based Cache Coherence Schemes,” *Int’l Conference on Parallel Processing*, August 1990.
- [11] Z. Guz, I. Keidar, A. Kolodny, U. C. Weiser. “Utilizing Shared Data in Chip Multiprocessors. with the Nahalal Architecture,” *SPAA*, June 2008.
- [12] M. Hammoud, S. Cho, and R. Melhem. “ACM: An Efficient Approach for Managing Shared Caches in Chip Multiprocessors,” *HiPEAC*, pp. 319–330, Jan. 2009.
- [13] N. Hardavellas, M. Ferdman, B. Falsafi, and A. Ailamaki. “Reactive NUCA: Near-Optimal Block Placement and Replication in Distributed Caches,” *ISCA*, June 2009.
- [14] J. Huh, C. Kim, H. Shafi, L. Zhang, D. Burger, and S. W. Keckler. “A NUCA Substrate for Flexible CMP Cache Sharing,” *ICS*, June 2005.
- [15] M. Kandemir, F. Li, M. J. Irwin, and S. W. Son. “A Novel Migration-Based NUCA Design for Chip Multiprocessors,” *SC*, Nov. 2008.
- [16] C. Kim, D. Burger, and S. W. Keckler. “An Adaptive, Non-Uniform Cache Structure for Wire-Delay Dominated On-Chip Caches,” *ASPLOS*, pp. 211–222, Oct. 2002.
- [17] P. Kongetira, K. Aingaran, and K. Olukotun. “Niagara: A 32-Way Multithreaded Sparc Processor,” *IEEE Micro*, 25(2): 21–29, March-April 2005.
- [18] J. Laudon and D. Lenoski. “The SGI Origin: A ccNUMA Highly Scalable Server,” *ISCA*, June 1997.
- [19] M. R. Marty and M. D. Hill. “Virtual Hierarchies to Support Server Consolidation,” *ISCA*, June 2007.
- [20] M. K. Qureshi, D. Thompson, and Y. N. Patt. “The V-WAY Cache: Demand-Based Associativity via Global Replacement,” *ISCA*, pp. 544–555, June 2005.
- [21] N. Rafique, W. Lim, M. Thottethodi. “Architectural Support for Operating System-Driven CMP Cache Management,” *PACT*, Sep. 2006.
- [22] Research at Intel. “Introducing the 45nm Next-Generation Intel Core<sup>TM</sup> Microarchitecture,” *White Paper*.
- [23] D. Rolán, B. B. Fraguera, and R. Doallo. “Adaptive line placement with the set balancing cache,” *MICRO*, pp. 529–540, Dec. 2009.
- [24] A. Ros, M. E. Acacio, and J. M. García. “Scalable Directory Organization for Tiled CMP Architectures,” *CDES*, July 2008.
- [25] K. Strandberg. “Which OS? Considerations for Performance-asymmetric, Multi-core Platforms,” *Research at Intel, White Paper*.
- [26] B. Sinharoy, R. N. Kalla, J. M. Tandler, R. J. Eickemeyer, and J. B. Joyner. “POWER5 System Microarchitecture,” *IBM J. Res. & Dev.*, 49(1):–25, July. 2005.
- [27] S. Srikantaiah, M. Kandemir, and M. J. Irwin. “Adaptive Set Pinning: Managing Shared Caches in Chip Multiprocessors,” *ASPLOS*, pp. 135–144, March 2008.
- [28] Standard Performance Evaluation Corporation. <http://www.specbench.org>.
- [29] D. Tam, R. Azimi, L. Soares, and M. Stumm. “Managing Shared L2 Caches on Multicore Systems in Software,” *WIOSCA*, 2007.
- [30] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta. “The SPLASH-2 Programs: Characterization and Methodological Considerations,” *ISCA*, pp. 24–36, July 1995.
- [31] M. Zhang and K. Asanović. “Victim Migration: Dynamically Adapting Between Private and Shared CMP Caches,” *TR-2005-064*, MIT, Oct. 2005.
- [32] M. Zhang and K. Asanović. “Victim Replication: Maximizing Capacity while Hiding Wire Delay in Tiled Chip Multiprocessors,” *ISCA*, June 2005.