

Fairness via Source Throttling: A Configurable and High-Performance Fairness Substrate for Multicore Memory Systems

EIMAN EBRAHIMI, The University of Texas at Austin
CHANG JOO LEE, Intel
ONUR MUTLU, Carnegie Mellon University
YALE N. PATT, The University of Texas at Austin

7

Cores in chip-multiprocessors (CMPs) share multiple memory subsystem resources. If resource sharing is unfair, some applications can be delayed significantly while others are unfairly prioritized. Previous research proposed separate fairness mechanisms for each resource. Such resource-based fairness mechanisms implemented independently in each resource can make contradictory decisions, leading to low fairness and performance loss. Therefore, a coordinated mechanism that provides fairness in the entire shared memory system is desirable.

This article proposes a new approach that provides fairness in the entire shared memory system, thereby eliminating the need for and complexity of developing fairness mechanisms for each resource. Our technique, Fairness via Source Throttling (FST), estimates unfairness in the *entire memory system*. If unfairness is above a system-software-set threshold, FST throttles down cores causing unfairness by limiting the number of requests they create and the frequency at which they do. As such, our *source-based* fairness control ensures fairness decisions are made in tandem in the entire memory system. FST enforces thread priorities/weights, and enables system-software to enforce different fairness objectives in the memory system.

Our evaluations show that FST provides the best system fairness and performance compared to three systems with state-of-the-art fairness mechanisms implemented in both shared caches and memory controllers.

Categories and Subject Descriptors: C.1.0 [Processor Architectures]: General; C.5.3 [Computer System Implementation]: Microcomputers—*Microprocessors*; C.1.2 [Processor Architectures]: Multiple Data Stream Architectures—*Multiprocessors*

General Terms: Design, Performance

Additional Key Words and Phrases: Chip-multiprocessors, memory system, fairness

ACM Reference Format:

Ebrahimi, E., Lee, C. J., Mutlu, O., and Patt, Y. N. 2012. Fairness via source throttling: A configurable and high-performance fairness substrate for multicore memory systems. *ACM Trans. Comput. Syst.* 30, 2, Article 7 (April 2012), 35 pages.

DOI = 10.1145/2166879.2166881 <http://doi.acm.org/10.1145/2166879.2166881>

1. INTRODUCTION

Chip-multiprocessor (CMP) systems commonly share a large portion of the memory subsystem between different cores. Main memory and shared caches are two examples

We gratefully acknowledge the support of the Cockrell Foundation, AMD, Intel, Oracle, GSRC, NSF, and CyLab.

Authors' addresses: E. Ebrahimi, C. J. Lee, and Y. N. Patt, Department of Electrical and Computer Engineering, University of Texas at Austin; email: E.ebrahimi@mail.utexas.edu; O. Mutlu, Department of Electrical and Computer Engineering, Carnegie Mellon University.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2012 ACM 0734-2071/2012/04-ART7 \$10.00

DOI 10.1145/2166879.2166881 <http://doi.acm.org/10.1145/2166879.2166881>



Fig. 1. Disparity in slowdowns due to unfairness in the memory system.

of shared resources. Memory requests from different applications executing on different cores of a CMP can interfere with and delay each other in the shared memory subsystem. Compared to a scenario where each application runs alone on the CMP, this intercore interference causes the execution of simultaneously running applications to slow down. However, sharing memory system resources affects the execution of different applications very differently because the resource management algorithms employed in the shared resources are unfair [Mutlu and Moscibroda 2007]. As a result some applications are unfairly slowed down significantly more than others.

Figure 1 shows examples of vastly differing effects of resource-sharing on simultaneously executing applications on a 2-core CMP system (Section 4 describes our experimental setup). When *bzip2* and *art* run simultaneously with equal priorities, the intercore interference caused by the sharing of memory system resources slows down *bzip2* by 5.2X compared to when it is run alone while *art* slows down by only 1.15X. In order to achieve system level fairness or quality of service (QoS) objectives, the system software (operating system or virtual machine monitor) expects proportional progress of equal-priority applications when running simultaneously. Clearly, disparities in slowdown like those shown in Figure 1 due to sharing of the memory system resources between simultaneously running equal-priority applications is unacceptable since it would make priority-based thread scheduling policies ineffective [Fedorova et al. 2007; Mutlu and Moscibroda 2007].

To mitigate this problem, previous papers [Hsu et al. 2006; Iyer 2004; Iyer et al. 2007; Kim et al. 2004; Mutlu and Moscibroda 2008, 2007; Nesbit et al. 2006, 2007] on fair memory system design for multicore systems mainly focused on partitioning a particular shared resource (cache space, cache bandwidth, or memory bandwidth) to provide fairness in the use of that shared resource. However, none of these prior papers directly target a *fair* memory system design that provides fair sharing of all resources together. We define a memory system design as fair if the slowdowns of equal-priority applications running simultaneously on the cores sharing that memory system are the same (this definition has been used in several prior papers [Cazorla et al. 2004; Gabor et al. 2006; Luo et al. 2001a; Mutlu and Moscibroda 2007; Snively and Tullsen 2000]). As shown in previous research [Bitirgen et al. 2008], employing separate uncoordinated fairness techniques together does not necessarily result in a fair memory system design. This is because fairness mechanisms in different resources can contradict each other. Our goal in this article is to develop a low-cost architectural technique that allows system software fairness policies to be achieved in CMPs by enabling fair sharing of the *entire memory system*, without requiring multiple complicated, specialized, and possibly contradictory fairness techniques for different shared resources.

Basic Idea. To achieve this goal, we propose a fundamentally new mechanism that (1) gathers dynamic feedback information about the unfairness in the system and (2)

uses this information to dynamically adapt the rate at which the different cores inject requests into the shared memory subsystem such that system-level fairness objectives are met. To calculate unfairness at runtime, a slowdown value is estimated for each application in hardware. Slowdown is defined as T_{shared}/T_{alone} , where T_{shared} is the number of cycles it takes to run simultaneously with other applications and T_{alone} is the number of cycles it would have taken the application to run alone. Unfairness is calculated as the ratio of the largest slowdown to the smallest slowdown of the simultaneously running applications. If the unfairness in the system becomes larger than the *unfairness threshold* set by the system software, the core that interferes most with the core experiencing the largest slowdown is throttled down. This means that the rate at which the most interfering core injects memory requests into the system is reduced, in order to reduce the intercore interference it generates. If the system software's *fairness goal* is met, all cores are allowed to throttle up to improve system throughput while system unfairness is continuously monitored. The fairness metric/goal, unfairness threshold, and throttling rates are all configurable by system software. This configurable hardware substrate enables the system software to achieve different QoS/fairness policies: it can determine the balance between fairness and system throughput, dictate different fairness objectives, and enforce thread priorities in the entire memory system.

Summary of Evaluation. We evaluate our technique on both 2-core and 4-core CMP systems in comparison to three previously-proposed state-of-the-art shared hardware resource management mechanisms. Experimental results across ten multiprogrammed workloads on a 4-core CMP show that our proposed technique improves average system performance by 25.6%/14.5% while reducing system unfairness by 44.4%/36.2% compared respectively to a system with no fairness techniques employed and a system with state-of-the-art fairness mechanisms implemented for both shared cache capacity [Nesbit et al. 2007] and the shared memory controller [Mutlu and Moscibroda 2008].

Contributions. We make the following contributions:

1. We introduce a low-cost, hardware-based and system-software-configurable mechanism to achieve fairness goals specified by system software in the entire shared multicore memory system.
2. We introduce a mechanism that collects dynamic feedback on the unfairness of the system and adjusts request rates of the different cores to achieve the desired fairness/performance balance. By performing source-based fairness control, this work eliminates the need for complicated individual resource-based fairness mechanisms that are implemented independently in each resource and that require coordination.
3. We qualitatively and quantitatively compare our proposed technique to multiple prior works in fair shared cache partitioning and fair memory scheduling. We find that our proposal, while simpler, provides significantly higher system performance and better system fairness compared to previous proposals.

2. BACKGROUND AND MOTIVATION

We first present brief background on how we model the shared memory system of CMPs. We then motivate our approach to providing fairness in the entire shared memory system by showing how employing resource-based fairness techniques does not necessarily provide better overall fairness.

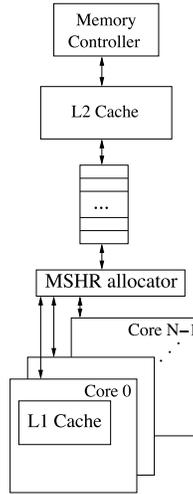


Fig. 2. Shared CMP memory system.

2.1. Shared CMP Memory Systems

In this article, we assume that the last-level (L2) cache and off-chip DRAM bandwidth are shared by multiple cores on a chip as in many commercial CMPs [AMD 2009; Wechsler 2006; Intel 2008; Tendler et al. 2001]. Each core has its own L1 cache. Miss Status Holding/information Registers (MSHRs) [Kroft 1981] keep track of all requests to the shared L2 cache until they are serviced. When an L1 cache miss occurs, an access request to the L2 cache is created by allocating an MSHR entry. Once the request is serviced by the L2 cache or DRAM system as a result of a cache hit or miss respectively, the corresponding MSHR entry is freed and used for a new request. Figure 2 gives a high level view of such a shared memory system. The number of MSHR entries for a core indicates the total number of outstanding requests allowed to the L2 cache and DRAM system. Therefore increasing/decreasing the number of MSHR entries for a core can increase/decrease the rate at which memory requests from the core are injected into the shared memory system.

2.2. Motivation

Most prior papers on providing fairness in shared resources focus on partitioning of a single shared resource. However, by partitioning a single shared resource, the demands on other shared resources may change such that neither system fairness nor system performance is improved. In the following example, we describe how constraining the rate at which an application's memory requests are injected to the shared resources can result in higher fairness and system performance than employing fair partitioning of a single resource.

Figure 3 shows the memory-related stall time¹ of applications A and B either running alone on one core of a 2-core CMP (parts (a)–(d)), or, running concurrently with equal priority on different cores of a 2-core CMP (parts (e)–(j)). For simplicity of explanation, we assume an application stalls when there is an outstanding memory request, focus on requests going to the same cache set and memory bank, and assume

¹Stall-time is the amount of execution time in which the application cannot retire instructions. Memory-related stall time caused by a memory request consists of: 1) time to access the L2 cache, and if the access is a miss 2) time to wait for the required DRAM bank to become available, and finally 3) time to access DRAM.

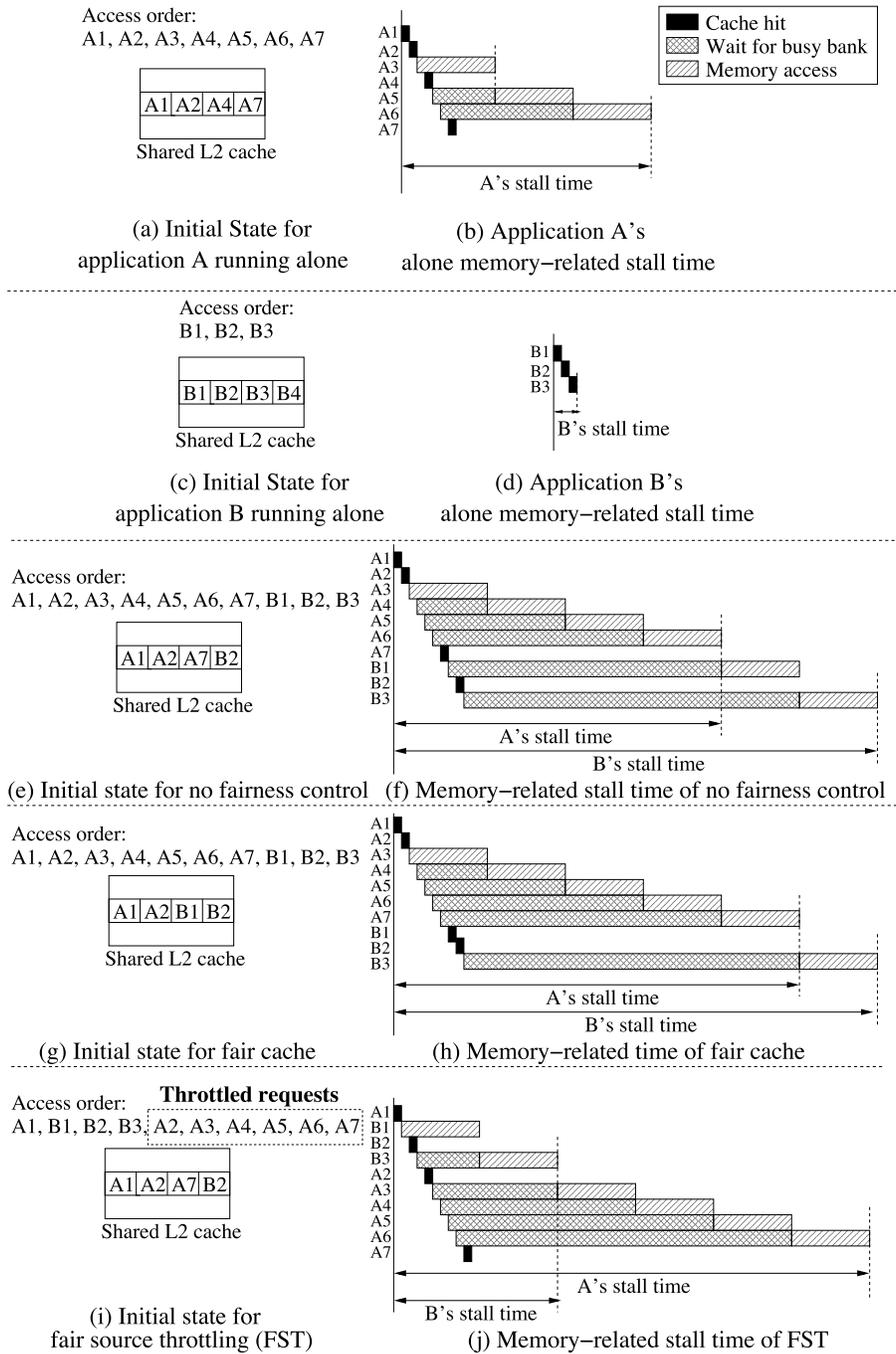


Fig. 3. Access pattern and memory-related stall time of requests when application A running alone (a, b), application B running alone (c, d), A and B running concurrently with no fairness control (e, f), fair cache (g, h), and fair source throttling (i, j).

all shown accesses to the shared cache occur before any replacement happens. Application A is very memory-intensive, while application B is much less memory-intensive as can be seen by the different memory-related stall times they experience when running alone (Figures 3(a)–(d)). As prior work has observed [Mutlu and Moscibroda 2008], when a memory-intensive application with already high memory-related stall time interferes with a less memory-intensive application with much smaller memory-related stall time, delaying the former improves system fairness because the additional delay causes a smaller slowdown for the memory-intensive application than for the nonintensive one. Doing so can also improve throughput by allowing the less memory-intensive application to quickly return to its compute-intensive portion while the memory-intensive application continues waiting on memory.

Figures 3(e) and (f) show the initial L2 cache state, access order and memory-related stall time when no fairness mechanism is employed in any of the shared resources. Application A's large number of memory requests arrive at the L2 cache earlier, and as a result, the small number of memory requests from application B are significantly delayed. This causes large unfairness because the compute-intensive application B is slowed down significantly more than the already-slow memory-intensive application A. Figures 3(g) and (h) show that employing a fair cache increases the fairness in utilization of the cache by allocating an equal number of ways from the accessed set to the two equal-priority applications. This increases application A's cache misses compared to the baseline with no fairness control. Even though application B gets more hits as a result of fair sharing of the cache, its memory-related stall time does not reduce due to increased interference in the main memory system from application A's increased misses. Application B's memory requests are still delayed behind the large number of memory requests from application A. Application A's memory-related stall time increases slightly due to its increased cache misses, however, since application A already had a large memory-related stall time, this slight increase does not incur a large slowdown for it. As a result, fairness improves slightly, but system throughput degrades because the system spends more time stalling rather than computing compared to no fair caching.

In Figure 3, if the unfair slowdown of application B due to application A is detected at runtime, system fairness can be improved by limiting A's memory requests and reducing the frequency at which they are issued into the shared memory system. This is shown in the access order and memory-related stall times of Figures 3(i) and (j). If the frequency at which application A's memory requests are injected into the shared memory system is reduced, the memory access pattern can change as shown in Figure 3(i). We use the term *throttled requests* to refer to those requests from application A that are delayed when accessing the shared L2 cache due to A's reduced injection rate. As a result of the late arrival of these throttled requests, application B's memory-related stall time significantly reduces (because A's requests no longer interfere with B's) while application A's stall time increases slightly. Overall, this ultimately improves both system fairness and throughput compared to both no fairness control and just a fair cache. Fairness improves because the memory-intensive application is delayed such that the less intensive application's memory related-stall time does not increase significantly compared to when running alone. Delaying the memory-intensive application does not slow it down too much compared to when running alone, because even when running alone it has high memory-related stall time. System throughput improves because the total amount of time spent computing rather than stalling in the entire system increases, as can be seen by comparing the stall times in Figures 3(f) and (h) to Figure 3(j).

The *key insight* is that both system fairness and throughput can improve by detecting high system unfairness at runtime and dynamically limiting the number of or

delaying the issuing of memory requests from the aggressive applications. In essence, we propose a new approach that performs source-based fairness in the entire memory system rather than individual resource-based fairness that implements complex and possibly contradictory fairness mechanisms in each resource. Sources (i.e., cores) can collectively achieve fairness by throttling themselves based on dynamic unfairness feedback, eliminating the need for implementing possibly contradictory/conflicting fairness mechanisms and complicated coordination techniques between them.

3. FAIRNESS VIA SOURCE THROTTLING

To enable fairness in the entire memory system, we propose *Fairness via Source Throttling* (FST). The proposed mechanism consists of two major components: (1) *runtime unfairness evaluation* and (2) *dynamic request throttling*.

3.1. Runtime Unfairness Evaluation Overview

The goal of this component is to dynamically obtain an estimate of the unfairness in the CMP memory system. We use the following definitions in determining unfairness:

- (1) We define a memory system design as fair if the slowdowns of equal-priority applications running simultaneously on the cores of a CMP are the same, similarly to previous works [Cazorla et al. 2004; Gabor et al. 2006; Luo et al. 2001a; Mutlu and Moscibroda 2007; Snaveley and Tullsen 2000].
- (2) We define slowdown as T_{shared}/T_{alone} where T_{shared} is the number of cycles it takes to run simultaneously with other applications and T_{alone} is the number of cycles it would have taken the application to run alone on the same system.

The main challenge in the design of the runtime unfairness evaluation component is obtaining information about the number of cycles it would have taken an application to run alone, while it is running simultaneously with other applications. To do so, we estimate the number of extra cycles it takes an application to execute due to intercore interference in the shared memory system, called T_{excess} . Using this estimate, T_{alone} is calculated as $T_{shared} - T_{excess}$. The following equations show how Individual Slowdown (IS) of each application and Unfairness of the system are calculated.

$$IS_i = \frac{T_i^{shared}}{T_i^{alone}}, \quad Unfairness = \frac{MAX\{IS_0, IS_1, \dots, IS_{N-1}\}}{MIN\{IS_0, IS_1, \dots, IS_{N-1}\}}.$$

Section 3.3 explains in detail how the runtime unfairness evaluation component is implemented and in particular how T_{excess} is estimated. Assuming for now that this component is in place, we next explain how the information it provides is used to determine how each application is throttled to achieve fairness in the entire shared memory system.

3.2. Dynamic Request Throttling

This component is responsible for dynamically adjusting the rate at which each core/application² makes requests to the shared resources. This is done on an interval basis, as shown in Figure 4.

An interval ends when each core has executed a certain number of instructions from the beginning of that interval. During each interval (for example Interval 1 in Figure 4) the runtime unfairness evaluation component gathers feedback used to estimate the slowdown of each application. At the beginning of the next interval (Interval 2), the

²Since each core runs a separate application, we use the words core and application interchangeably in this article.

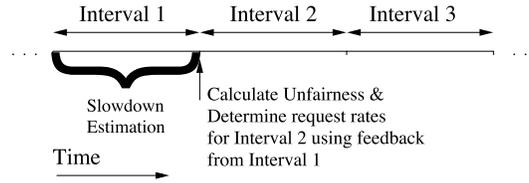


Fig. 4. FST's interval-based estimation and throttling.

feedback information obtained during the prior interval is used to make a decision about the request rates of each application for that interval. More precisely, slowdown values estimated during Interval 1 are used to estimate unfairness for the system. That unfairness value is used to determine the request rates for the different applications for the duration of Interval 2. During the next interval (Interval 2), those request rates are applied, and unfairness evaluation is performed again. The algorithm used to adjust the request rate of each application using the unfairness estimate calculated in the prior interval is shown in Algorithm 1. To ease explanations, Algorithm 1 is simplified for dual-core configurations. Section 3.5 presents the more general algorithm for more than two cores.

We define multiple possible levels of aggressiveness for the request rate of each application. The dynamic request throttling component makes a decision to increase/decrease or keep constant the request rate of each application at interval boundaries. We refer to increasing/decreasing the request rate of an application as throttling the application up/down.

ALGORITHM 1: Dynamic Request Throttling

```

if Estimated Unfairness > Unfairness Threshold then
  Throttle down application with the smallest slowdown
  Throttle up application with the largest slowdown
  Reset Successive Fairness Achieved Intervals
else
  if Successive Fairness Achieved Intervals = threshold then
    Throttle all applications up
    Reset Successive Fairness Achieved Intervals
  else
    Increment Successive Fairness Achieved Intervals
  end if
end if

```

At the end of each interval, the algorithm compares the unfairness estimated in the previous interval to the unfairness threshold that is defined by system software. If the fairness goal has not been met in the previous interval, the algorithm reduces the request rate of the application with the smallest individual slowdown value and increases the request rate of the application with the largest individual slowdown value. This reduces the number and frequency of memory requests generated for and inserted into the memory resources by the application with the smallest estimated slowdown, thereby reducing its interference with other cores. The increase in the request rate of the application with the highest slowdown allows it to be more aggressive in exploiting Memory-Level Parallelism (MLP) [Chou et al. 2004; Glew 1998] and as a result

reduces its slowdown. If the fairness goal is met for a predetermined number of intervals (tracked by a *Successive Fairness Achieved Intervals* counter in Algorithm 1), the dynamic request throttling component attempts to increase system throughput by increasing the request rates of all applications by one level. This is done because our proposed mechanism strives to increase throughput while maintaining the fairness goals set by the system software. Increasing the request rate of all applications might result in unfairness. However, the unfairness evaluation during the interval in which this happens detects this occurrence and dynamically adjusts the requests rates again.

Throttling Mechanisms. Our mechanism increases/decreases the request rate of each application in multiple ways: (1) Adjusting the number of outstanding misses an application can have at any given time. To do so, an *MSHR quota*, which determines the maximum number of MSHR entries an application can use at any given time, is enforced for each application. Reducing MSHR entries for an application reduces the pressure caused by that application's requests on all shared memory system resources by limiting the number of concurrent requests from that application contending for service from the shared resources. This reduces other simultaneously running applications' memory-related stall times and gives them the opportunity to speed up. (2) Adjusting the frequency at which requests in the MSHRs are issued to access L2. Reducing this frequency for an application reduces the number of memory requests per unit time from that application that contend for shared resources. This allows memory requests from other applications to be prioritized in accessing shared resources even if the application that is throttled down does not have high MLP to begin with and is not sensitive to reduction in the number of its MSHRs. We refer to this throttling technique as *frequency throttling*. We use both of these mechanisms to reduce the interference caused by the application that experiences the smallest slowdown on the application that experiences the largest slowdown.

3.3. Unfairness Evaluation Component Design

T_{shared} is simply the number of cycles it takes to execute an application in an interval. Estimating T_{alone} is more difficult, and FST achieves this by estimating T_{excess} for each core, which is the number of cycles the core's execution time is lengthened due to interference from other cores in the shared memory system. To estimate T_{excess} , the unfairness evaluation component keeps track of intercore interference each core incurs.

Tracking Intercore Interference. We consider three sources of intercore interference: (1) cache, (2) DRAM bus and bank conflict, and (3) DRAM row-buffer.³ Our mechanism uses an *InterferencePerCore* bit-vector whose purpose is to indicate whether or not a core is delayed due to intercore interference. In order to track interference from each source separately, a copy of *InterferencePerCore* is maintained for each interference source. A main copy which is updated by taking the union of the different *InterferencePerCore* vectors is eventually used to update T_{excess} as described below. When FST detects intercore interference for core i at any shared resource, it sets bit i of the *InterferencePerCore* bit-vector, indicating that the core was delayed due to interference. At the same time, it also sets an *InterferingCoreId* field in the corresponding *interfered-with* memory request's MSHR entry. This field indicates which

³On-chip interconnect can also experience intercore interference [Das et al. 2009]. Feedback information similar to that obtained for the three sources of intercore interference we account for can be collected for the on-chip interconnect. That information can be incorporated into our technique seamlessly, which we leave as part of future work.

Table I. Hardware Cost of FST on a 4-Core CMP System

	Cost for N cores	Cost for N = 4
<i>ExcessCycles</i> counters	$N \times N \times 16$ bits/counter	256 bits
Interference pollution filter per core	$2048 \text{ entries} \times N \times (1 \text{ pollution bit} + (\log_2 N) \text{ bit processor id})/\text{entry}$	24,576 bits
<i>InterferingCoreId</i> per MSHR entry	$32 \text{ entries/core} \times N \times 2 \text{ interference sources} \times (\log_2 N) \text{ bits/entry}$	512 bits
<i>InterferencePerCore</i> bit-vector	$(3 \text{ interference sources} + 1 \text{ main copy}) \times N \times N \times 1 \text{ bit}$	64 bits
Shadow row-buffer address register	$N \times \# \text{ of DRAM banks (B)} \times 32 \text{ bits/address}$	1024 bits
<i>Successive Fairness Achieved Intervals To Wait To Throttle Up</i> per core <i>Inst Count Each Interval</i> per core	$(2 \times N + 1) \times 16$ bits/counter	144 bits
Core id per tag store entry in K MB L2	$16384 \text{ blocks/Megabyte} \times K \times (\log_2 N) \text{ bit/block}$	65,536 bits
Total hardware cost for N-core system	Sum of the above	11.24 KB
Percentage area overhead (as fraction of the baseline K MB L2 cache)	$\text{Sum (KB)} \times 100 / (K \times 1024)$	$11.24\text{KB}/2048\text{KB} = \mathbf{0.55\%}$

core interfered with this request and is later used to reset the corresponding bit in the *InterferencePerCore* vector when the *interfered-with* request is scheduled/serviced. We explain this process in more detail for each resource below in Sections 3.3.1–3.3.3. If a memory request has not been interfered with, its *InterferingCoreId* will be the same as the core id of the core it was generated by.

Updating T_{excess} . FST stores the number of *extra cycles* it takes to execute a given interval’s instructions due to intercore interference (T_{excess}) in an *ExcessCycles* counter per core. Every cycle, if the *InterferencePerCore* bit of a core is set, FST increments the corresponding core’s *ExcessCycles* counter. Section 3.3.5 shows how this can be done less frequently.

Algorithm 2 shows how FST calculates *ExcessCycles* for a given core i . The following subsections explain in detail how each source of intercore interference is taken into account to set *InterferencePerCore*. Table I summarizes the required storage needed to implement the mechanisms explained here.

3.3.1. Cache Interference. In order to estimate intercore cache interference, for each core i we need to track the last-level cache misses that are caused by any other core j . To do so, FST uses a pollution filter for each core to approximate such misses. The pollution filter is a bit-vector that is indexed with the lower order bits of the accessed cache line’s address.⁴ In the bit-vector, a set entry indicates that a cache line belonging to the corresponding core was evicted by another core’s request. When a request from core j replaces one of core i ’s cache lines, core i ’s filter is accessed using the evicted line’s address, and the corresponding bit is set. When a memory request from core i misses the cache, its filter is accessed with the missing address. If the corresponding bit is set, the filter predicts that this line was previously evicted due to intercore interference and the bit in the filter is reset. When such a prediction is made, once the interfered-with request is scheduled to DRAM the *InterferencePerCore* bit corresponding to core i is set to indicate that core i is experiencing extra execution cycles due to cache interference. Once the interfered-with memory request is finished receiving service from the memory system and the corresponding cache line is filled, core i ’s filter is accessed and the bit is reset and so is core i ’s *InterferencePerCore* bit.

⁴We empirically determined the pollution filter for each core to have 2K-entries in our evaluations.

ALGORITHM 2: Estimation of T_{excess} for core i

Every cycle

```

if intercore cache or DRAM bus or DRAM bank or
DRAM row-buffer interference then
  set InterferencePerCore bit  $i$ 
  set InterferingCoreId in delayed memory request
end if
if InterferencePerCore bit  $i$  is set then
  Increment ExcessCycles for core  $i$ 
end if

```

Every L2 cache fill for a miss due to interference OR**Every time a memory request which is a row-buffer miss due to interference is serviced**

```

reset InterferencePerCore bit of core  $i$ 
InterferingCoreId of core  $i = i$  (no interference)

```

Every time a memory request is scheduled to DRAM

```

if Core  $i$  has no requests waiting on any bank which is busy servicing another core  $j$  ( $j \neq i$ )
then
  reset InterferencePerCore bit of core  $i$ 
end if

```

3.3.2. DRAM Bus and Bank Conflict Interference. Intercore DRAM bank conflict interference occurs when core i 's memory request cannot access the bank it maps to, because a request from some other core j is being serviced by that memory bank. DRAM bus conflict interference occurs when a core cannot use the DRAM because another core is using the DRAM bus. These situations are easily detectable at the memory controller, as described in Mutlu and Moscibroda [2007]. When such interference is detected, the *InterferencePerCore* bit corresponding to core i is set to indicate that core i is stalling due to a DRAM bus or bank conflict. This bit is reset when no request from core i is being prevented access to DRAM by the other cores' requests.

3.3.3. DRAM Row-Buffer Interference. This type of interference occurs when a potential row-buffer hit of core i when running alone is converted to a row-buffer miss/conflict due to a memory request of some core j when running together with others. This can happen if a request from core j closes a DRAM row opened by a prior request from core i that is also accessed by a subsequent request from core i . To track such interference, a *Shadow Row-buffer Address Register* (SRAR) is maintained for each core for each bank. Whenever core i 's memory request accesses some row R , the SRAR of core i is updated to row R . Accesses to the same bank from some other core j do not affect the SRAR of core i . As such, at any point in time, core i 's SRAR will contain the last row accessed by the last memory request serviced from that core in that bank. When core i 's memory request suffers a row-buffer miss because another core j 's row is open in the row-buffer of the accessed bank, the SRAR of core i is consulted. If the SRAR indicates a row-buffer hit would have happened, then intercore row-buffer interference is detected. As a result, the *InterferencePerCore* bit corresponding to core i is set. Once the memory request is serviced, the corresponding *InterferencePerCore* bit is reset.⁵

⁵To be more precise, the bit is reset "row buffer hit latency" cycles before the memory request is serviced. The memory request would have taken at least "row buffer hit latency" cycles had there been no interference.

3.3.4. Slowdown due to Throttling. When an application is throttled, it experiences some slowdown due to the throttling. This slowdown is different from the intercore interference induced slowdown estimated by the mechanisms of Sections 3.3.1 to 3.3.3. Throttling-induced slowdown is a function of an application’s sensitivity to 1) the number of MSHRs that are available to it, 2) the frequency of injecting requests into the shared resources. Using profiling, we determine for each throttling level l , the corresponding slowdown (due to throttling) f of an application A . At runtime, any estimated slowdown for application A when running at throttling level l is multiplied by f . We find that accounting for this slowdown using this profiling information improves the system performance gained by FST by 4% on 4-core systems, as we evaluate in Section 5.10.

Slowdown due to throttling can also be tracked by maintaining a counter for the number of cycles each application A stalls because it can not obtain an MSHR entry because of its limited MSHR quota. We separately keep track of the number of such cycles and refer to them as those excess cycles which are due to throttling (as opposed to excess cycles due to interference from other applications). We discuss how this information is used later in a more general form of dynamic request throttling presented in Section 3.5, Algorithm 3.

3.3.5. Implementation Details. Section 3.3 describes how separate copies of *InterferencePerCore* are maintained per interference source. The main copy which is used by FST for updating T_{excess} is physically located close by the L2 cache. Note that shared resources may be located far away from each other on the chip. Any possible timing constraints on the sending of updates to the *InterferencePerCore* bit-vector from the shared resources can be eliminated by making these updates periodically. In Section 5.5 we show that making updates as infrequently as even once every 1000 cycles provides negligible loss of fidelity compared to ideally making updates every cycle.

3.4. System Software Support

Different Fairness Objectives. System-level fairness objectives and policies are generally decided by the system software (the operating system or virtual machine monitor). FST is intended as architectural support for enforcing such policies in shared memory system resources. The *fairness goal* to be achieved by FST can be configured by system software. To achieve this, we enable the system software to determine the nature of the condition that triggers Algorithm 1. In the explanations of Section 3.2, the *triggering condition* is

Triggering Condition (1) : Estimated Unfairness > Unfairness Threshold.

System software might want to enforce different triggering conditions depending on the system’s fairness/QoS requirements. To enable this capability, FST implements different triggering conditions from which the system software can choose. For example, the fairness goal the system software wants to achieve could be to keep the maximum slowdown of any application below a threshold value. To enforce such a goal, the system software can configure FST such that the triggering condition in Algorithm 1 is changed to

Triggering Condition (2) : Estimated Slowdown_i > Max. Slowdown Threshold.

Alternatively, per application slowdown thresholds can be specified. In this case, if any application slows down beyond its own specified threshold, Algorithm 1 will be triggered.

Thread Weights. So far, we have assumed all threads are of equal importance. FST can be seamlessly adjusted to distinguish between and provide differentiated services to threads with different priorities. We add the notion of *thread weights* to FST, which are communicated to it by the system software using special instructions. Higher slowdown values are more tolerable for less important or *lower weight* threads. To incorporate thread weights, FST uses *weighted slowdown* values calculated as:

$$\text{Weighted Slowdown}_i = \text{Measured Slowdown}_i \times \text{Weight}_i.$$

By scaling the real slowdown of a thread with its weight, a thread with a higher weight appears as if it slowed down more than it really did, causing it to be favored by FST. Section 5.4 quantitatively evaluates FST with one different fairness goal and threads with different weights.

Thread Migration and Context Switches. FST can be seamlessly extended to work in the presence of thread migration and context switches. When a context switch happens or a thread is migrated, the interference state related to that thread is cleared. When a thread restarts executing after a context switch or migration, it starts at maximum throttle. The interference caused by the thread and the interference it suffers are dynamically re-estimated and FST adapts to the new set of co-executing applications.

3.5. General Dynamic Request Throttling

Scalability to More Cores. When the number of cores is greater than two, a more general form of Algorithm 1 is used. The design of the *unfairness evaluation* component for the more general form of Algorithm 1 is slightly different. This component gathers the following extra information for the more general form of dynamic request throttling presented in Algorithm 3: a) for each core i , FST maintains a set of $N-1$ counters, where N is the number of simultaneously running applications. We refer to these $N-1$ counters that FST uses to keep track of the amount of the intercore interference caused by any other core j in the system for i as *ExcessCycles_{ij}*. This information is used to identify which of the other applications in the system generates the most interference for core i , b) FST maintains the total intercore interference an application on core i experiences due to interference from other cores in a *TotalExcessCyclesInterference_i* counter per core, and c) as described in the last paragraph of Section 3.3.4, those excess cycles that are caused as a result of an application being throttled down are accounted for separately in a *TotalExcessCyclesThrottling_i* counter per core.

Algorithm 3 shows the generalized form of Algorithm 1 that uses the extra information described above to make more accurate throttling decisions in a system with more than two cores. The four most important changes are as follows:

First, when the algorithm is triggered due to unfair slowdown of core i , FST compares the *ExcessCycles_{ij}* counter values for all cores $j \neq i$ to determine which other core is interfering most with core i . The core found to be the most interfering is throttled down. We do this in order to reduce the slowdown of the core with the largest slowdown value, and improve system fairness.

Second, the first line of the algorithm shows how we change the condition that triggers throttling. Throttling is triggered if both the estimated unfairness (*Max. Slowdown/Min. Slowdown*) and the ratio between the slowdowns of core with the largest slowdown (*App_{slow}*) and the core generating the most interference (*App_{interfering}*) are greater than *Unfairness Threshold*. Doing so helps reduce excessive throttling when two applications significantly interfere with each other and alternate between being identified as *App_{slow}* and *App_{interfering}*. Consider the case where application A and B alternate between being *App_{slow}* (which has *Max. Slowdown*) and *App_{interfering}*; and some other (possibly memory nonintensive) application C is the application with

ALGORITHM 3: Dynamic Request Throttling - General Form

```

if Estimated Unfairness > Unfairness Threshold AND  $App_{slow} slowdown / App_{interfering} slowdown > Unfairness Threshold$  then
  if  $App_{slow}$ 's excess cycles due to interference from  $App_{interfering} > App_{slow}$ 's TotalExcessCyclesThrottlingi then
    Throttle down application that causes most interference ( $App_{interfering}$ ) for application with largest slowdown
  end if
  Throttle up application with the largest slowdown ( $App_{slow}$ )
  Reset Successive Fairness Achieved Intervals
  Reset Intervals To Wait To Throttle Up for  $App_{interfering}$ .

  // Preventing bank service denial
  if  $App_{interfering}$  throttled lower than  $Switch_{thr}$  AND causes greater than  $Interference_{thr}$  amount of  $App_{slow}$ 's total interference then
    Temporarily stop prioritizing  $App_{interfering}$  due to row hits in memory controller
  end if
  if  $App_{RowHitNotPrioritized}$  has not been  $App_{interfering}$  for  $SwitchBack_{thr}$  intervals then
    Allow it to be prioritized in memory controller based on row-buffer hit status of its requests
  end if

  for all applications except  $App_{interfering}$  and  $App_{slow}$  do
    if Intervals To Wait To Throttle Up = threshold1 then
      throttle up
      Reset Intervals To Wait To Throttle Up for this app.
    else
      Increment Intervals To Wait To Throttle Up for this app.
    end if
  end for

  else
    if Successive Fairness Achieved Intervals = threshold2 then
      Throttle up application with the smallest slowdown
      Reset Successive Fairness Achieved Intervals
    else
      Increment Successive Fairness Achieved Intervals
    end if
  end if

```

Min. Slowdown. With the throttling condition of Algorithm 1 in place, applications A and B would continuously be throttled up and down in successive intervals without the *Estimated Unfairness* ever dropping below the specified *Unfairness Threshold*. This is because, in the intervals when either is detected to be App_{slow} , *Estimated Unfairness* will be high because of application C's small slowdown. By comparing the slowdowns of applications A and B before throttling is performed overall throughput is improved by avoiding excessive throttling which would not improve the system's *Estimated Unfairness*.

Third, we observe that there are situations where an application suffers slowdown that is incurred as a result of throttling from previous intervals and not due to

intercore interference. To address this, we detect such cases. We restrict throttling down of $App_{interfering}$ to cases where the slowdown that App_{slow} is suffering is mainly caused by intercore interference and is not a result of App_{slow} having been throttled down in previous intervals. If the excess cycles that App_{slow} suffers due to not being able to acquire MSHR entries is greater than the excess cycles caused for it by $App_{interfering}$, we do not throttle down $App_{interfering}$ as this would result in a loss of throughput. In these cases the detected unfairness is resolved by throttling up App_{slow} and reducing its slowdown by allowing it to acquire more MSHR entries.

Fourth, cores that are neither the core with the largest slowdown (App_{slow}) nor the core generating the most interference ($App_{interfering}$) for the core with the largest slowdown are throttled up every $threshold1$ intervals. This is a performance optimization that allows cores to be aggressive if they are not the main contributors to the unfairness in the system.

Preventing Bank Service Denial due to FR-FCFS Memory Scheduling. First ready-first come first serve (FR-FCFS) [Rixner et al. 2000] is a commonly used memory scheduling policy which we use in our baseline system. This algorithm prioritizes requests that hit in the DRAM bank row buffers over all other requests. The FR-FCFS policy has the potential to starve an application with low row-buffer locality in the presence of an application with high row-buffer locality (as discussed in prior work [Nesbit et al. 2006; Moscibroda and Mutlu 2007; Mutlu and Moscibroda 2007, 2008]). Even when the interfering application is throttled down, the potential for continued DRAM bank interference exists when FR-FCFS memory scheduling is used, due to the greedy row-hit-first nature of the scheduling algorithm: a throttled-down application with high row-buffer locality can deny service to another application continuously. To overcome this, we supplement FST with a heuristic that prevents this denial of service. Once an application has already been throttled down lower than $Switch_{thr}\%$, if FST detects that this throttled application is generating greater than $Interference_{thr}\%$ of App_{slow} 's total interference, it will temporarily stop prioritizing the interfering application based on row-buffer hit status in the memory controller. We refer to this application as $App_{RowHitNotPrioritized}$. If $App_{RowHitNotPrioritized}$ has not been the most interfering application for $SwitchBack_{thr}$ number of intervals, its prioritization over other applications based on row-buffer hit status will be re-allowed in the memory controller. This is done because if an application with high row-buffer locality is not allowed to take advantage of row buffer hits for a long time, its performance will suffer.⁶

3.6. Hardware Cost and Implementation Details

Table I shows the breakdown of FST's required storage. The total storage cost required by our implementation of FST is 11.24KB which is only 0.55% the size of the L2 cache being used. FST does not require any structure or logic that is on the critical path since all updates to interference-tracking structures can be made periodically at relatively large intervals to eliminate any timing constraints (see Section 5.5).

Figure 5 shows the shared CMP memory system we model for evaluation of FST including additional structures for tracking interference added to the baseline memory system shown in Figure 2. The two boxes on the right of the figure contain interference tracking structures and counters, and the shaded bit positions in the L2 cache lines and MSHR entries on the left are additions to these structures required by FST.

⁶We do this so that we can have minimal changes to the most commonly used scheduling algorithm. FST can be combined with other forms of memory scheduling, which is part of future research and out of the scope of this article.

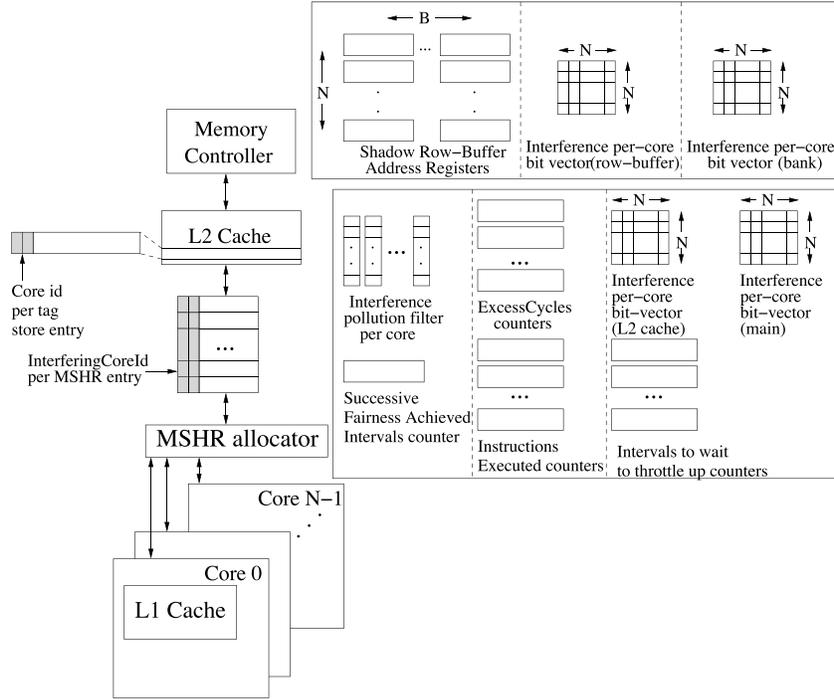


Fig. 5. Changes made to the memory system.

3.7. Lightweight FST

In this section, we describe an alternative FST implementation that requires less hardware cost and is more scalable in terms of hardware requirements to a larger number of cores. In this alternative implementation, we do not keep track of how much interference is caused by each application for any other application which requires N^2 *ExcessCycles* counters (where N is the number of applications), as described in the previous subsection. Instead, we propose maintaining two counters for each core i . One counter tracks the total number of *ExcessCycles* that the application executing on core i generated for *any other* concurrently-executing application. We refer to this counter as *ExcessCyclesGenerated_i*. The other counter tracks the total number of *ExcessCycles* that *any other* concurrently-executing application creates for the application on core i . We refer to this counter as *ExcessCyclesSuffered_i*. This requires a total of $2N$ 16-bit counters to be maintained and allows for a more scalable solution with larger numbers of cores: the number of required counters is linear instead of quadratic in the number of cores.

For the lightweight FST implementation to work with the counters described above, we modify Algorithm 3 as follows. With lightweight FST, the core executing the application that has the largest slowdown App_{slow} is still throttled up when throttling is triggered. However, as opposed to throttling down the core executing the application which causes the most interference for App_{slow} ($App_{interfering}$) in Algorithm 3), we throttle down the core that is executing the application which is generating the most interference for any other concurrently-executing application since $App_{interfering}$ is not known due to the reduced number of counters. This is the core with the highest *ExcessCyclesGenerated_i* counter in a given interval. We evaluate the performance of our lightweight FST in Section 5.7.

Table II. Baseline System Configuration

Execution Core	Out of order processor, 15 stages, Decode/retire up to 4 instructions Issue/execute up to 8 micro instructions 256-entry reorder buffer
Front End	Fetch up to 2 branches; 4K-entry BTB 64K-entry Hybrid branch predictor
On-chip Caches	L1 I-cache: 32KB, 4-way, 2-cycle, 64B line L1 D-cache: 32KB, 4-way, 2-cycle, 64B line Shared unified L2: 1MB (2MB for 4-core), 8-way (16-way for 4-core), 16-bank, 15-cycle (20-cycle for 4-core), 1 port, 64B line size
DRAM Controller	On-chip, FR-FCFS scheduling policy [Rixner et al. 2000] 128-entry MSHR and memory request buffer
DRAM and Bus	667MHz bus cycle, DDR3 1333MHz [MICRON] 8B-wide data bus Latency: 15-15-15ns ($t_{RP-t}RCD-CL$) 8 DRAM banks, 16KB row buffer per bank Round-trip L2 miss latency: Row-buffer hit: 36ns, conflict: 66ns

4. METHODOLOGY

Processor Model. We use an in-house cycle-accurate x86 CMP simulator for our evaluation. We faithfully model all port contention, queuing effects, bank conflicts, and other major DDR3 DRAM system constraints in the memory subsystem. Table II shows the baseline configuration of each core and the shared resource configuration for the 2 and 4-core CMP systems we use.

Workloads. We use the SPEC CPU 2000/2006 benchmarks for our evaluation. Each benchmark was compiled using ICC (Intel C Compiler) or IFORT (Intel Fortran Compiler) with the -O3 option. We ran each benchmark with the reference input set for 200 million x86 instructions selected by Pinpoints [Patil et al. 2004] as a representative portion for the 2-core experiments. Due to long simulation times, 4-core experiments were done with 50 million instructions per benchmark.

We classify benchmarks as highly memory-intensive/with medium memory intensity/nonintensive for our analyses and workload selection. We refer to a benchmark as highly memory-intensive if its L2 Cache Misses per 1K Instructions (MPKI) is greater than ten. If the MPKI value is greater than one but less than ten, we say the benchmark has medium memory-intensity. If the MPKI value is less than one, we refer to it as nonintensive. This classification is based on measurements made when each benchmark was run alone on the 2-core system. Table III shows the characteristics of the benchmarks that appear in the evaluated workloads when run on the 2-core system.

Workload Selection. We used 18 two-application and 10 four-application multiprogrammed workloads for our 2-core and 4-core evaluations respectively. The 2-core workloads were chosen such that at least one of the benchmarks is highly memory intensive. For this purpose we used either *art* from SPEC2000 or *lbm* from SPEC2006. For the second benchmark of each 2-core workload, applications of different memory intensity were used in order to cover a wide range of different combinations. Of the 18 benchmarks combined with either *art* or *lbm*, seven benchmarks have high memory intensity, six have medium intensity, and five have low memory intensity. The ten 4-core workloads were randomly selected with the condition that the evaluated workloads each include at least one benchmark with high memory intensity and at least one benchmark with medium or high memory intensity.

Table III.

Characteristics of 29 SPEC 2000/2006 benchmarks: IPC and MPKI (L2 cache Misses Per 1K Instructions).

Benchmark	Type	IPC	MPKI	Benchmark	Type	IPC	MPKI
art	FP00	0.10	90.89	milc	FP06	0.30	29.33
soplex	FP06	0.28	21.24	leslie3d	FP06	0.41	20.88
lbm	FP06	0.45	20.16	bwaves	FP06	0.46	18.71
GemsFDTD	FP06	0.46	15.63	lucas	FP00	0.61	10.61
astar	INT06	0.37	10.19	omnetpp	INT06	0.36	10.11
mgrid	FP00	0.52	6.5	gcc	INT06	0.45	6.26
zeusmp	FP06	0.82	4.69	cactusADM	FP06	0.60	4.51
bzip2	INT06	1.14	2.61	xalancbmk	INT06	0.71	1.68
h264ref	INT06	1.46	1.28	vortex	INT00	1.01	1.24
parser	INT00	1.24	0.91	apsi	FP00	1.81	0.85
ammp	FP00	1.8	0.75	perlbench	INT06	1.49	0.68
mesa	FP00	1.82	0.61	gromacs	FP06	1.06	0.29
namd	FP06	2.25	0.18	crafty	INT00	1.82	0.1
calculix	FP06	2.28	0.05	gamess	FP06	2.32	0.04
povray	FP06	1.88	0.02				

Table IV. FST Parameters

<i>Fairness Threshold</i>	<i>Successive Fairness Achieved Intervals Threshold</i>	<i>Intervals Wait To Throttle Up</i>	<i>Interval Length</i>	<i>Switch_{thr}</i>	<i>Interference_{thr}</i>	<i>SwitchBack_{thr}</i>
1.4	4	2	25Kinsts	5%	70%	3 intervals

FST parameters used in evaluation. Table IV shows the values we use in our evaluation unless stated otherwise. There are eight aggressiveness levels used for the request rate of each application: 2%, 3%, 4%, 5%, 10%, 25%, 50%, and 100%. These levels denote the scaling of the MSHR quota and the request rate in terms of percentage. For example, when FST throttles an application to 5% of its total request rate on a system with 128 MSHRs, two parameters are adjusted. First, the application is given a 5% quota of the total number of available MSHRs (in this case, 6 MSHRs). Second, the application's memory requests in the MSHRs are issued to access the L2 cache at 5% of the maximum possible frequency (i.e., once every 20 cycles).

Metrics. To measure CMP system performance, we use *Harmonic mean of Speedups (Hspeedup)* [Luo et al. 2001a], and *Weighted Speedup (Wspeedup)* [Snively and Tullsen 2000]. These metrics are commonly used in measuring multiprogram performance in computer architecture [Eyerhan and Eeckhout 2008]. In order to demonstrate fairness improvements, we report *Unfairness* (see Section 3.1), as defined in Gabor et al. [2006] and Mutlu and Moscibroda [2007], and *Maximum Slowdown*, which is the maximum individual slowdown that any application in a workload experiences. The *Maximum Slowdown* metric provides understanding about at most how much any of the applications in a given workload is slowed down due to sharing of memory system resources [Kim et al. 2010a]. Since *Hspeedup* provides a balanced measure between fairness and system throughput as shown in previous work [Luo et al. 2001a], we use it as our primary evaluation metric. In the metric definitions below: N is the number of cores in the CMP system, IPC^{alone} is the IPC measured when an application runs alone on one core in the CMP system (other cores are idle), and IPC^{shared} is the IPC

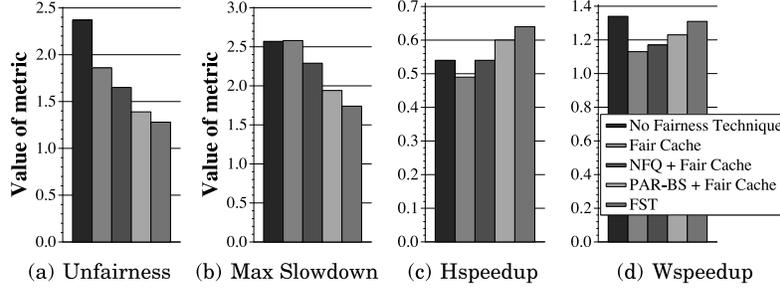


Fig. 6. Average performance of FST on the 2-core system.

measured when an application runs on one core while other applications are running on the other cores.

$$Hspeedup = \frac{N}{\sum_{i=0}^{N-1} \frac{IPC_i^{alone}}{IPC_i^{shared}}}, \quad Wspeedup = \sum_{i=0}^{N-1} \frac{IPC_i^{shared}}{IPC_i^{alone}}.$$

5. EXPERIMENTAL EVALUATION

We evaluate our proposed techniques on both 2-core (Section 5.1) and 4-core systems (all other sections). We compare FST to four other systems in our evaluations: 1) a baseline system with no fairness techniques employed in the shared memory system, using LRU cache replacement and FR-FCFS memory scheduling [Rixner et al. 2000], both of which have been shown to be unfair [Kim et al. 2004; Moscibroda and Mutlu 2007; Nesbit et al. 2006]. We refer to this baseline as *NoFairness*, 2) a system with only fair cache capacity management using the virtual private caches technique [Nesbit et al. 2007], called *FairCache*, 3) a system with a network fair queuing (NFQ) fair memory scheduler [Nesbit et al. 2006] combined with fair cache capacity management [Nesbit et al. 2007], called *NFQ+FairCache*, 4) a system with a parallelism-aware batch scheduling (PAR-BS) fair memory scheduler [Mutlu and Moscibroda 2008] combined with fair cache capacity management [Nesbit et al. 2007], called *PAR-BS+FairCache*.

5.1. 2-core System Results

Figure 6 shows system performance and unfairness averaged (using geometric mean) across 18 workloads evaluated on the 2-core system. Figure 7 shows the Hspeedup performance of FST and other fairness techniques normalized to that of a system without any fairness technique for each of the 18 evaluated 2-core workloads. FST provides the highest system performance (in terms of Hspeedup) and the best unfairness among all evaluated techniques. We make several key observations:

1. Fair caching's unfairness reduction comes at the cost of a large degradation in system performance. Also average maximum slowdown which indicates the highest slowdown any application in a workload experiences due to sharing of memory system resources is increased slightly. These two phenomena occur because fair caching changes the memory access patterns of applications. Since the memory access scheduler is unfair, the fairness benefits of the fair cache itself are reverted by the memory scheduler.
2. NFQ+FairCache together reduces system unfairness by 30.2% compared to *NoFairness* and reduces maximum slowdown by 10.9%. However, this degrades Wspeedup

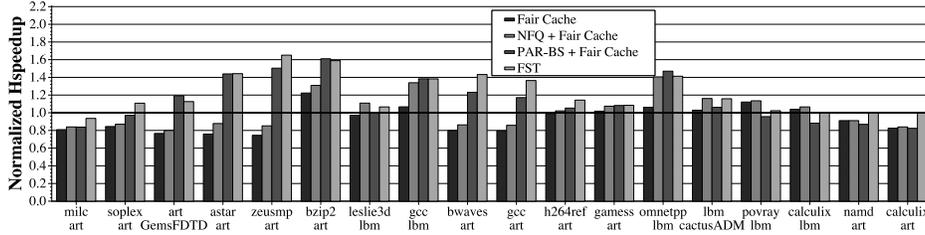


Fig. 7. Hspeedup of 18 2-core workloads normalized to no fairness control.

(by 12.3%). The combination of PAR-BS and fair caching improves both system performance and fairness compared to the combination of NFQ and a fair cache. The main reason is that PAR-BS preserves both DRAM bank parallelism and row-buffer locality of each thread better than NFQ, as shown in previous work [Mutlu and Moscibroda 2008]. Compared to the baseline with no fairness control, employing PAR-BS and a fair cache reduces unfairness and maximum slowdown by 41.3%/24.5% and improves Hspeedup by 11.5%. However, this improvement comes at the expense of a large (7.8%) Wspeedup degradation.

- NFQ+FairCache and PAR-BS+FairCache both significantly degrade system throughput (Wspeedup) compared to employing no fairness mechanisms. This is due to two reasons both of which lead to the delaying of memory nonintensive applications (Recall that prioritizing memory nonintensive applications is better for system throughput [Mutlu and Moscibroda 2008; Nesbit et al. 2006]). First, the fairness mechanisms that are employed separately in each resource interact negatively with each other, leading to one mechanism (e.g. fair caching) increasing the pressure on the other (fair memory scheduling). As a result, even though fair caching might benefit system throughput by giving more resources to a memory nonintensive application, increased misses of the memory-intensive application due to fair caching causes more congestion in the memory system, leading to both the memory-intensive and nonintensive applications to be delayed. Second, even though the combination of a fair cache and a fair memory controller can prioritize a memory nonintensive application's requests, this prioritization can be temporary. The deprioritized memory-intensive application can still fill the shared MSHRs with its requests, thereby denying the nonintensive application entry into the memory system. Hence, the nonintensive application stalls because it cannot inject enough requests into the memory system. As a result, the memory nonintensive application's performance does not improve while the memory-intensive application's performance degrades (due to fair caching), resulting in system throughput degradation.
3. FST reduces system unfairness and maximum slowdown by 46.1%/32.3% while also improving Hspeedup by 20% and degrades Wspeedup by 1.8% compared to *NoFairness*. Unlike other fairness mechanisms, FST improves both system performance and fairness, without large degradation to Wspeedup. This is due to two major reasons. First, FST provides a coordinated approach in which both the cache and the memory controller receive less frequent requests from the applications causing unfairness. This reduces the starvation of the applications that are unfairly slowed down as well as interference of requests in the memory system, leading to better system performance for almost all applications. Second, because FST uses MSHR quotas to limit requests injected by memory-intensive applications that cause unfairness, these memory-intensive applications do not deny other applications' entry into the memory system. As such, unlike other fairness techniques that do not consider fairness in memory system buffers (e.g., MSHRs), FST ensures that

Table V. Summary of Results on the 2-Core System

	Unfairness	Maximum Slowdown	Hspeedup	Wspeedup
FST Δ over No Fairness Mechanism	-46.1%	-32.3%	20%	-1.8%
FST Δ over Fair Cache	-31.3%	-32.6%	30.2%	16.1%
FST Δ over NFQ + Fair Cache	-22.8%	-24.1%	19.7%	11.9%
FST Δ over PAR-BS + Fair Cache	-8.2%	-10.4%	7.5%	6.4%

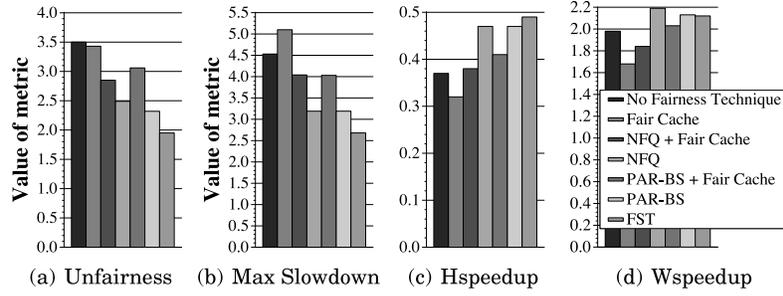


Fig. 8. Average performance of FST on the 4-core system.

unfairly slowed-down applications are prioritized in the entire memory system, including all the buffers, caches, and schedulers.

Table V summarizes our results for the 2-core evaluations. Compared to the previous technique that provides the highest system throughput (i.e. *NoFairness*), FST provides a significantly better balance between system fairness and performance. Compared to the previous technique that provides the best fairness (*PAR-BS+FairCache*), FST improves both system performance and fairness. We conclude that FST provides the best system fairness as well as the best balance between system fairness and performance.

5.2. 4-core System Results

5.2.1. Overall Performance. Figure 8 shows unfairness and system performance averaged across the ten evaluated 4-core workloads (results in all sections that follow are evaluated on the same system and same set of benchmarks). FST provides the best fairness (in terms of both smallest unfairness and smallest maximum slowdown) and Hspeedup among all evaluated techniques,⁷ while providing Wspeedup that is within 3.5% that of the best previous technique. Overall, FST reduces unfairness and maximum slowdown by 44.4%/41%⁸ and increases system performance by 30.4% (Hspeedup) and 6.9% (Wspeedup) compared to *NoFairness*. Compared to NFQ, the previous technique with the highest system throughput (Wspeedup), FST reduces unfairness and max slowdown by 22%/16.1% and increases Hspeedup by 4.2%. FST's large performance improvement is mainly due to the large reduction in unfairness.⁹

⁷In this subsection we also include data points for NFQ alone and PAR-BS alone with no FairCache to show how the uncoordinated combination of fairness techniques at different shared resources can result in degradation of both performance and fairness compared to when only one is employed.

⁸Similarly, FST also reduces the coefficient of variation, an alternative unfairness metric [Zhang et al. 2009], by 45%.

⁹Since relative slowdowns of different applications are most important to improving unfairness and performance using FST, highly accurate T_{excess} estimations are not necessary for such improvements. However, we find that with the mechanisms proposed in this article the application which causes the most interference for the most-slowed-down application is on average identified correctly in 70% of the intervals.

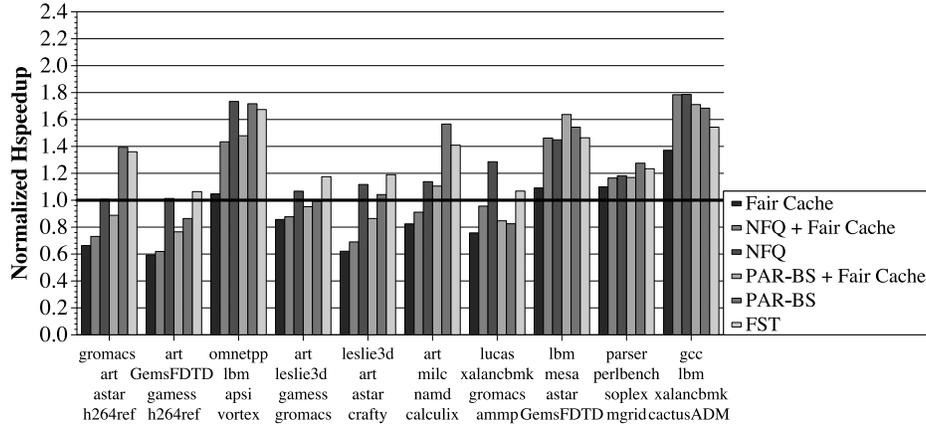


Fig. 9. Normalized speedup of 10 4-core workloads.

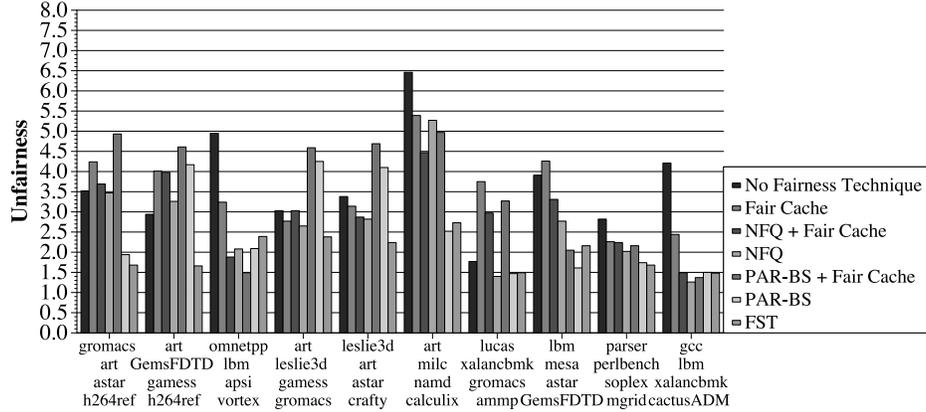


Fig. 10. Unfairness of 10 4-core workloads.

Note that the overall trends in the 4-core system are similar to those in the 2-core system except that previous fairness mechanisms do not significantly improve fairness in the 4-core system. As explained in detail in Section 5.2.2, this happens because previous fairness mechanisms prioritize nonintensive applications in individual resources regardless of whether or not those applications are actually being slowed down.

Figure 9 shows the harmonic speedup performance of FST and other fairness techniques normalized to that of a system without any fairness technique for each of the ten workloads. Figure 10 shows the system unfairness of all the techniques for each of the ten workloads. We make two major conclusions. First, FST improves system performance (both Hspeedup and Wspeedup) and fairness compared to no fairness control for all workloads. Second, FST provides the highest Hspeedup compared to the previous technique with the highest average system performance (NFQ) on seven of the ten workloads, and provides the best fairness compared to the previous technique with the best system fairness (PAR-BS) on seven of the ten workloads.

5.2.2. Case Study. To provide more insight into the performance and fairness improvements of FST, we analyze one 4-core workload in detail. This workload is a mix of applications of different levels of memory intensity. *Art* and *leslie* are both

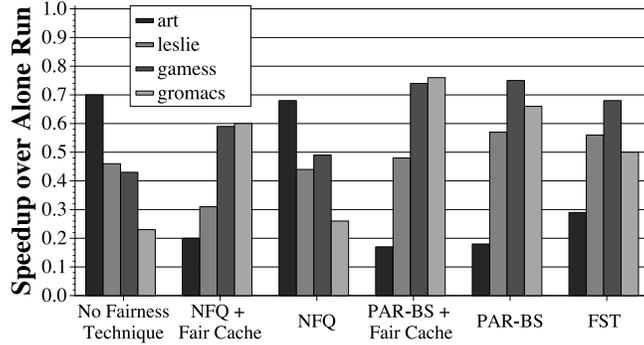


Fig. 11. Case Study: individual application behavior.

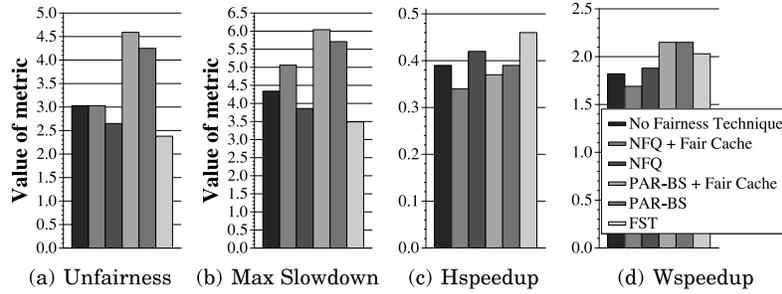


Fig. 12. Case study: system behavior.

highly memory-intensive, while *gamess* and *gromacs* are nonintensive (as shown in Table III). When these applications are run simultaneously on a 4-core system with no fairness control, the two memory-intensive applications (especially *art*) generate a large amount of memory traffic. *Art*'s large number of memory requests to the shared resources unfairly slows down all other three applications, while *art* does not slow down significantly. Figures 11 and 12 show individual benchmark performance and system performance/fairness, respectively (note that Figure 11 shows speedup over alone run which is the inverse of individual slowdown, defined in Section 3.1). Several observations are in order.

- (1) NFQ+FairCache significantly degrades system performance by 12.3% (Hspeedup) and 7.1% (Wspeedup) compared to no fairness control. This combination slows down the memory-intensive applications too much, resulting in a 16.7% increase in maximum slowdown compared to employing no fairness technique. The largest slowdowns are experienced by the memory-intensive *art* and *leslie* because they both get less cache space due to FairCache, and are deprioritized in DRAM due to NFQ prioritizing infrequent requests (with earlier virtual finish times) from *gamess* and *gromacs*. On the other hand, when NFQ alone is employed, the memory nonintensive applications' performance is slightly improved by prioritizing them in DRAM at small reductions to the performance of the memory-intensive applications. NFQ alone improves system performance by 6.7%/3.1% (HS/WS) and reduces unfairness/maximum slowdown by 12.7%/10.9%. However, NFQ does not address interference caused in the shared cache so its gains are limited.

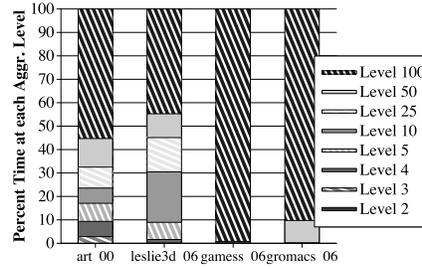


Fig. 13. Case Study: application throttling levels.

- (2) With PAR-BS+FairCache and PAR-BS, *art* is heavily deprioritized with unfair improved performance for the less memory intensive applications resulting in improved overall system throughput (Wspeedup). These two techniques are an example of where unfair treatment of applications in a workload may increase system throughput at the cost of large increases to unfairness and maximum slowdown (51.5%/39% and 40.4%/31.6% for PAR-BS+FairCache and PAR-BS respectively) and degradation of average system turnaround time (Hspeedup) compared to not using any fairness technique. These techniques overly deprioritize memory intensive applications (specifically *art*) because they do not explicitly detect when such applications cause slowdowns for others. They simply prioritize nonintensive applications all the time regardless of whether or not they are actually slowed down in the memory system. In contrast, our approach explicitly detects when memory-intensive applications are causing unfairness in the system. If they are not causing unfairness, FST does not deprioritize them. As a result, their performance is not unnecessarily reduced. This effect is observed by examining the most memory-intensive application's (*art*'s) performance with FST. With FST, *art* has higher performance than with any of the other fairness techniques.
- (3) FST increases system performance by 17.5%/11.6% (HS/WS) while reducing unfairness/maximum slowdown by 21.4%/19.5% compared to no fairness control. In this workload, the memory-intensive *art* and *leslie* cause significant interference to each other in all shared resources and to *gromacs* in the shared cache. Unlike other fairness techniques, FST dynamically tracks the interference and the unfairness in the system in a fine-grained manner. When the memory-intensive applications are causing interference and increasing unfairness, FST throttles the offending *hog* application(s). In contrast, when the applications are not interfering significantly with each other, FST allows them to freely share resources in order to maximize each application's performance. The fine-grained dynamic detection of unfairness and enforcement of fairness mechanisms only when they are needed allow FST to achieve higher system performance (Hspeedup) and a better balance between fairness and performance than other techniques.

To provide insight into the dynamic behavior of FST, Figure 13 shows the percentage of time each core spends at each throttling level. FST significantly throttles down *art* and *leslie* much of the time (but not always) to reduce the intercore interference they generate for each other and the less memory intensive applications. As a result, *art* and *leslie* spend almost 25%/30% of their execution time at 10% or less of their full aggressiveness. Also, a lot of the time *art* can prevent bank service to the continuous accesses of *leslie* to the same bank. FST detects this and disallows *art*'s requests to be prioritized based on row-buffer hits for 74% of all intervals, preventing *art* from causing bank service denial, as described in Section 3.5. Note that *art* spends

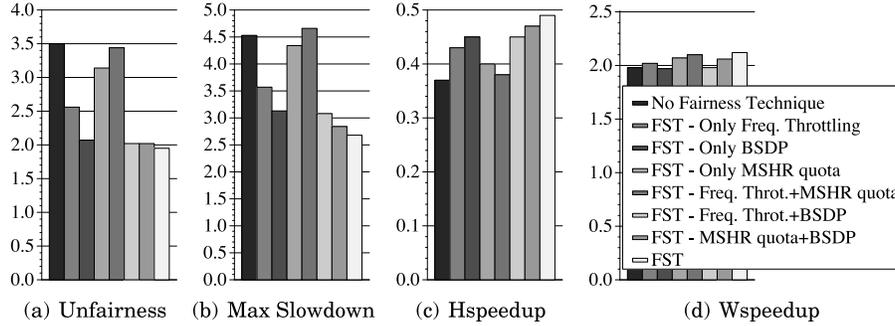


Fig. 14. Effects of different throttling mechanisms for FST.

approximately 55% of its time at throttling level 100, which shows that FST detects times when *art* is not causing large interference and does not penalize it. Figure 13 also shows that FST detects interference caused by not only *art* but also other applications. *leslie*, *gromacs*, and even *gamess* are detected to generate intercore interference for other applications in certain execution intervals. As such, FST dynamically adapts its fairness control decisions to the interference patterns of applications rather than simply prioritizing memory nonintensive applications. Therefore, unlike other fairness techniques, FST does not overly deprioritize *art* in the memory system.

We conclude that FST provides a higher-performance approach to attaining fairness than coarsely tracking the memory-intensity of applications and deprioritizing memory-intensive applications without dynamic knowledge of interference and unfairness. FST achieves this by tracking unfairness in the system and making fairness/throttling decisions based on that tracking in a finer-grained manner.

5.3. Effect of Throttling Mechanisms

As described in Section 3.2, FST uses the combination of two mechanisms to throttle an application up/down and increase/decrease its request rate from the shared resources. (1) Applying an MSHR quota to each application, (2) Adjusting the frequency at which requests in the MSHRs are issued to access L2. Section 3.5 explains how to prevent bank service denial from FR-FCFS memory scheduling within FST. Figure 14 shows the effect of each of the different throttling mechanisms, the effect of bank service denial prevention (BSDP), and FST on the 4-core system. Several observations are in order:

- (1) Employing BSDP always improves performance regardless of the throttling mechanism that is used. BSDP’s improvements are due to the resolution of a problem we refer to as the *over-throttling* problem. As explained in Section 3.5 memory intensive applications that also have high row-buffer locality can cause significant interference even if they are throttled when the memory controller uses an FR-FCFS scheduling algorithm. When this occurs (using the terminology of Section 3.5), FST detects an already throttled down application to be $App_{interfering}$ and continuously throttles it down further because the estimated unfairness remains high and App_{slow} stays the same. We call this over-throttling of $App_{interfering}$. BSDP resolves this issue by eliminating the cause of bank service denial due to FR-FCFS scheduling.

In Figure 14, the fourth and fifth bars from the left in each subgraph show the importance of BSDP. Without BSDP, enabling MSHR quotas destroys fairness (subfigures (a) and (b)) and degrades system performance in terms of harmonic

Table VI. Sensitivity of Alone Performance (IPC) to # of MSHRs

# of MSHRs	1	2	3	5	6	12	32	64	128
sphinx3 (IPC)	0.13	0.23	0.28	0.29	0.29	0.30	0.30	0.30	0.30
milc (IPC)	0.10	0.22	0.36	0.38	0.39	0.40	0.40	0.40	0.40
lbm (IPC)	0.04	0.10	0.22	0.26	0.30	0.39	0.45	0.46	0.48

mean of speedups (sub-figure (c)) as a result of unfair treatment of memory-intensive applications in some workloads. The large increase in unfairness is mainly due to workloads that contain the application *art*. *Art* is a highly memory-intensive workload with high row-buffer locality. As such, as we described in Section 3.5 it can cause bank service denial for concurrently executing applications even when it is throttled down. Additionally, *art*'s performance is very sensitive to the number of MSHR entries at its disposal. As a result, it can get overthrottled as described above when MSHR quotas are employed for throttling. The fourth and fifth bars from the left in Figure 14 show that while the overthrottling problem that exists for the workloads including *art* does not result in an average loss of system throughput (Wspeedup, subfigure (d)) across all the workloads, it does have a large impact on system fairness and average system turnaround time (as shown by subfigures (a)–(c)). We conclude that BSDP is necessary for significant improvements to system fairness when MSHR quotas are employed.

- (2) Without BSDP, the combination of MSHR quota and frequency throttling perform worse than using MSHR quota alone. The reason for this is the overthrottling of memory-intensive benchmarks in the absence of BSDP. When both throttling mechanisms are employed, the negative effect of overthrottling dominates average performance in our evaluated workloads. This leads to the combination of the two throttling mechanisms performing worse than MSHR alone in the absence of BSDP.
- (3) Using MSHR quotas is more effective than using frequency throttling alone when BSDP is employed. Using *MSHR quotas* together with BSDP achieves 97% of the performance improvement and 95% of the fairness improvement provided by FST. However, as Table VI shows that different applications are affected differently by small adjustments to their MSHR quota values. Applications with high memory-level-parallelism such as *lbm* are sensitive to the number of MSHRs they have available to them: small changes to their MSHR quota results in large slowdowns. On the other hand, applications such as *sphinx3* and *milc* do not make use of many MSHRs even when running alone as they do not have high degrees of memory-level parallelism. For such memory-intensive applications with low MLP, applying MSHR quotas as the throttling mechanism reduces the request rates only at the smallest throttling levels (MSHR quotas of 1 or 2). Therefore, using the second throttling mechanism (frequency throttling) that reduces the frequency at which requests are sent to L2 provides better, fine-grained control of request injection rate.

We conclude that using all mechanisms of FST is better than each throttling mechanism alone in terms of both fairness and performance.

5.4. Evaluation of System Software Support

Enforcing Thread Priorities. As explained in Section 3.4, FST can be configured by system software to assign different weights to different threads. As an example of how FST enforces thread weights, we ran four identical copies of the *GemsFDTD* benchmark on a 4-core system and assigned them thread weights of 1, 1, 4 and 8 (recall

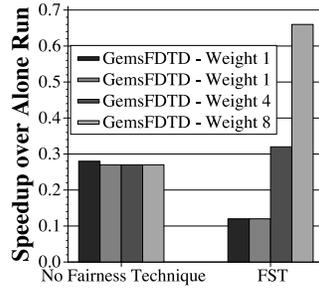


Fig. 15. Enforcing thread weights with FST.

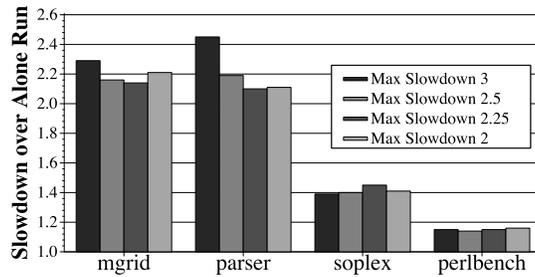


Fig. 16. Enforcing maximum slowdown with FST.

that a higher-weight thread is one the system software wants to prioritize). Figure 15 shows that with no fairness technique each copy of *GemsFDTD* has an almost identical slowdown as the baseline does not support thread weights and treats the applications identically in the shared memory system. However, FST prioritizes the applications proportionally to their weights, favoring applications with higher weight in the shared memory system. FST also slows down the two copies with the same weight by the same amount. We conclude that FST approximately enforces thread weights, thereby easing the development of system software which naturally expects a CMP to respect thread weights in the shared memory system.

Enforcing an Alternative Fairness Objective (Maximum Tolerable Slowdown). Section 3.4 explained how FST can be configured to achieve a maximum slowdown threshold as determined by system software, that dictates the maximum tolerable slowdown of any individual application executing concurrently on the CMP. Figure 16 shows an example of how FST enforces this fairness objective when four applications are run together on a 4-core system. The figure shows each application’s individual slowdown in four different experiments where each experiment uses a different maximum slowdown threshold (ranging from 2 to 3) as set by the system software. As tighter goals are set by the system software, FST throttles the applications accordingly to achieve (close to) the desired maximum slowdown. The fairness objective is met until the maximum slowdown threshold becomes too tight and is violated (for *mgrid* and *parser*), which happens at threshold value 2. We conclude that FST can enforce different system-software-determined fairness objectives.

In Algorithm 3, throttling is triggered when estimated system unfairness is greater than a system-software-specified threshold. Figure 17 shows average system performance and fairness when using a system-software-specified maximum slowdown target (Triggering Condition 2 from Section 3.4) compared to FST with an unfairness

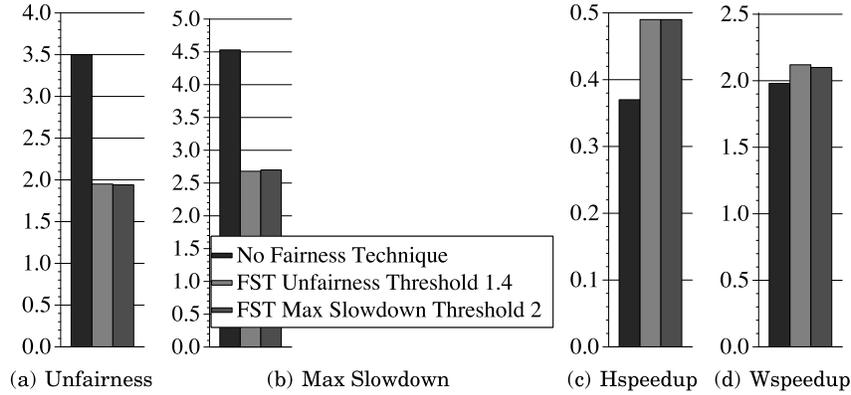


Fig. 17. Comparing overall results with different system level QoS targets.

target (Triggering Condition 1 from Section 3.4, which is the system-software target we use in all other experiments in this article). We conclude that similar system performance and fairness benefits can be gained using either system software goal: maximum tolerable slowdown or maximum tolerable unfairness.

We evaluate sensitivity to the unfairness threshold which is part of the system software support in Section 5.8 separately.

5.5. Effects of Implementation Constraints

Shared resources may be located far away from each other on the chip. In order to eliminate timing constraints on the sending of updates to the *InterferencePerCore* bit-vector from the shared resources, such updates can be made periodically. Every *UpdateThreshold* cycles, all shared resources send their local copies of *InterferencePerCore* to update the main copy at the L2. Once the updates are applied to the main copy by taking the union of all bit-vectors, FST checks the main copy of *InterferencePerCore*. If the *InterferencePerCore* bit of a core is set, FST increments the *ExcessCycles* counter corresponding to the core by the *UpdateThreshold* value.

Figure 18 shows the effect of periodic updates and sensitivity to chosen period lengths on the performance and fairness improvements of FST. The figure shows that even with updates occurring once every 1000 cycles, system performance is almost identical and fairness improvements are within 2.5% of FST with updates made every cycle. This is because memory system interference generally results in excess cycles in the order of hundreds of cycles. As such, our mechanism can tolerate updates happening periodically without incurring big losses in fidelity. We conclude that using periodic updates (even when made at relatively long periods) eliminates any timing constraints on the sending of updates to the *InterferencePerCore* bit-vector and does not significantly effect the performance and fairness improvements of FST.

5.6. Effects of Different Sources of Interference

Figure 19 shows the effect of taking into account interference from each of the interference sources we discuss in Section 3.3. The figure shows that from the different interference sources discussed in Section 3.3, FST's performance is mostly sensitive to whether or not DRAM bank interference is included in the estimations. Without taking into account DRAM bank interference, FST only improves performance by 5.1% (Hspeedup) and reduces unfairness by 13.8% respectively. On the other hand, if we

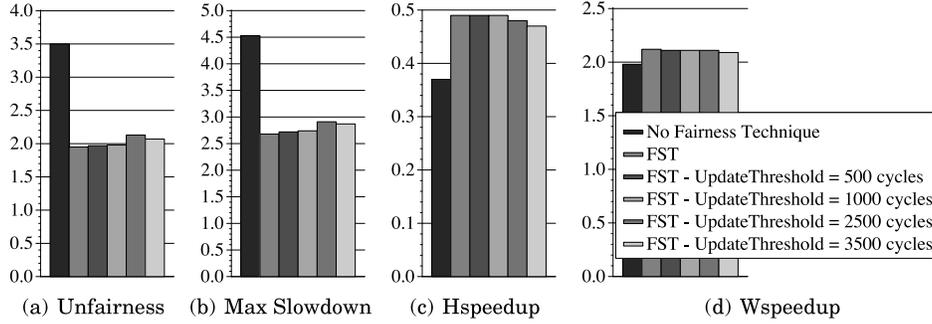


Fig. 18. Effect of periodic updates on FST's performance and unfairness.

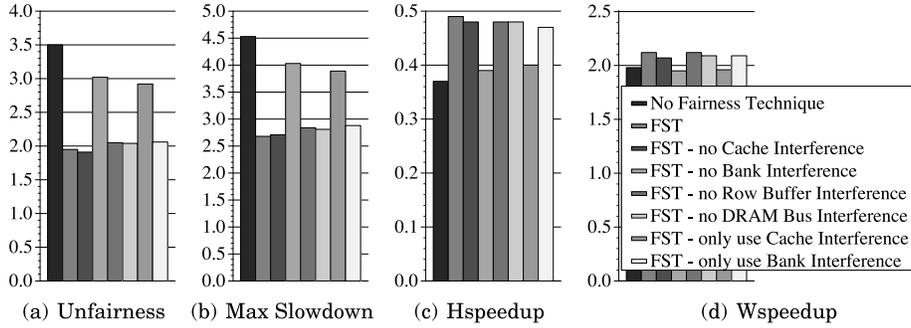


Fig. 19. Sensitivity of FST to taking into account different interference sources.

have an FST implementation that does not take into account cache or DRAM row-buffer interference (i.e., one that takes into account only DRAM bank interference), we can achieve 97% of the total performance improvements of FST and 94% of its total unfairness reduction. As we have observed in Section 3.6, a significant portion of the hardware required to implement FST is required for accounting for cache interference and DRAM row-buffer interference. As a result, this gives opportunity for a much less expensive implementation of FST based only on DRAM bank interference.

5.7. Evaluation of Lightweight FST

Figure 20 compares the performance of the lightweight FST implementation described in Section 3.7 to that of the baseline and the original full-blown FST we have been evaluating so far. The figure shows that the lightweight implementation that requires $2N$ cycles for tracking *ExcessCycles* information provides 98% of the system performance and 95% of the system fairness benefits of the original FST which requires N^2 counters. We conclude that this lightweight version of FST can be a more scalable yet high performance option to consider for systems with a larger number of cores.

5.8. Sensitivity to Unfairness Threshold

Figure 21 shows how FST's average fairness and performance changes with different unfairness thresholds on our evaluated 4-core workloads. Lowering the unfairness threshold set by the system-software continuously improves fairness and performance until the unfairness threshold becomes too small. With a very small unfairness threshold (1.05), FST becomes 1) very aggressive at throttling down cores to reach the very tight unfairness goal, 2) too sensitive to inaccuracies in slowdown estimation and

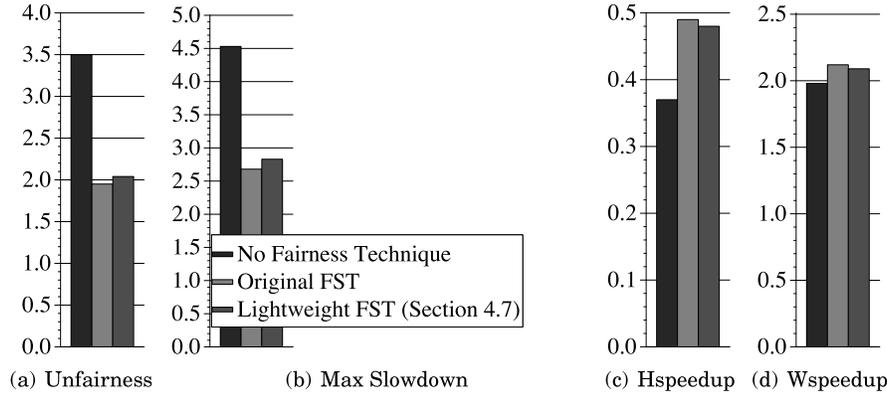


Fig. 20. Comparing overall results of original and lightweight FST.

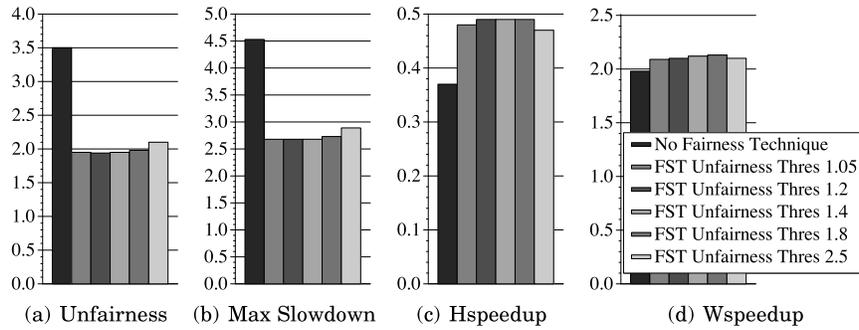


Fig. 21. Sensitivity of FST to unfairness threshold.

therefore triggers throttling of sources unnecessarily. As a result, both system performance and fairness slightly degrade. On the other hand, as the threshold increases, unfairness in the system also increases (because throttling is employed less often) and performance decreases beyond some point (because memory hog applications start causing starvation to others, leading to lower system utilization). Overall, the unfairness threshold provides a knob to the system software, using which the system software can determine the fairness-performance balance in the system. We find an unfairness threshold of 1.4 provides the best fairness and performance for our 4-core workloads.

5.9. Effect of Multiple Memory Controllers

Figure 22 shows the effect of using FST on a system with two memory controllers. Such a system has higher available off-chip bandwidth and therefore less intercore interference and less unfairness than a system with one controller. Yet, even in such a system, FST provides significant improvements in system fairness and performance compared to the baseline and combination of state-of-the-art fairness mechanisms at the different resources.

5.10. Evaluation of Using Profile Information

Figure 23 shows the effect of using profile information to account for slowdown due to throttling as described in Section 3.3.4. The figure shows system performance

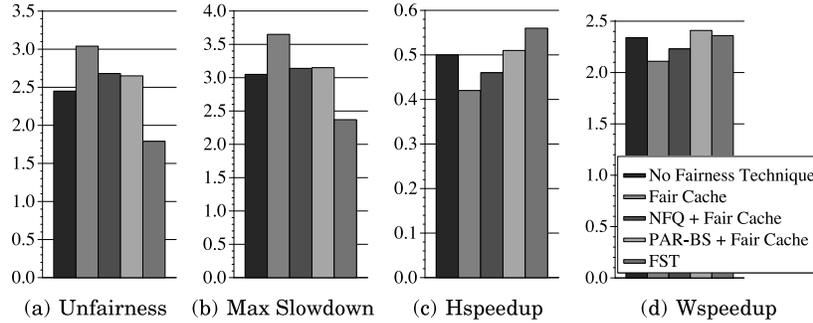


Fig. 22. Effect of FST on a system with two memory controllers.

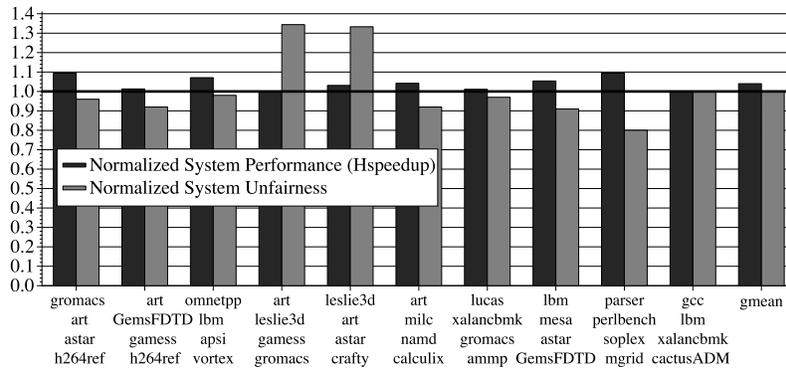


Fig. 23. Effect of using profile information for throttling related slowdown.

(Hspeedup shown on the first bar) and system unfairness (shown on the second bar) of a system using FST with profile information normalized to that of a system using FST without profile information. On average, using such profile information improves system performance by 4% and leaves system unfairness unchanged across the 4-core workloads. However, such profile information is not completely accurate in accounting for slowdowns due to throttling in all intervals since the factors described in Section 3.3.4 are obtained by comparing performance of complete runs of each application at different throttling levels. Due to the inaccuracies that exist, the use of this information results in increased system unfairness in two of the workloads.

6. RELATED WORK

To our knowledge, this article provides the first comprehensive and flexible hardware-based solution that enables system-software-specified fairness goals to be achieved in the entire shared memory system of a multicore processor, without requiring fairness mechanisms to be implemented individually in each shared resource.

Prior work in providing fairness in different shared resources of CMP systems focused on fair caching [Iyer 2004; Iyer et al. 2007; Kim et al. 2004; Nesbit et al. 2007], fair memory scheduling [Kim et al. 2010b; Mutlu and Moscibroda 2007, 2008; Nesbit et al. 2006], and fair on-chip interconnects [Das et al. 2009; Grot et al. 2009; Lee et al. 2008]. We have already provided extensive qualitative and quantitative comparisons showing that our mechanism significantly improves system fairness and performance compared to systems employing the combination of state-of-the-art fair cache capacity

management [Nesbit et al. 2007] and fair memory scheduling [Mutlu and Moscibroda 2008; Nesbit et al. 2006].

Bitirgen et al. [2008] propose implementing an artificial neural network that learns each application's performance response to different resource allocations. Their technique searches the space of different resource allocations among coexecuting applications to find a partitioning in the shared cache and memory controller that improves performance. In contrast to FST, this mechanism requires that resource-based fairness/partitioning techniques are already implemented in each individual resource. In addition, it requires relatively more complex, black-box implementation of artificial neural networks in hardware.

Herdrich et al. [2009] observe that the interference caused by a lower-priority application on a higher-priority application can be reduced using existing clock modulation techniques in CMP systems. However, their proposal does not consider or provide fairness to equal-priority applications. Zhang et al. [2009] propose a software-based technique that uses clock modulation and prefetcher on/off control provided by existing hardware platforms to improve fairness in current multicore systems compared to other software techniques. Neither of these prior works propose an online algorithm that dynamically controls clock modulation to achieve fairness. In contrast, FST provides: 1) hardware-based architectural mechanisms that continuously monitor shared memory system unfairness at runtime and 2) an online algorithm that, upon detection of unfairness, throttles interfering applications using two new hardware-based throttling mechanisms (instead of coarse-grained clock modulation) to reduce the interfering applications' request rates.

Jahre and Natvig [2009] observe that adjusting the number of available MSHRs can control the total miss bandwidth available to each thread running on a CMP. However, this prior work does not show how this observation can be used by an online algorithm to dynamically achieve a well-defined fairness or performance goal. In contrast to this prior work, our work (1) provides architectural support for achieving different well-defined system-software fairness objectives while also improving system performance, (2) shows that using complementary throttling mechanisms and preventing bank service denial due to FR-FCFS, as done by FST, provides better fairness/performance than simply adjusting the number of available MSHRs (see Section 5.3), (3) shows that FST's approach of throttling sources based on unfairness feedback provides better system fairness/performance than implementing different fairness mechanisms in each individual shared resource.

Zhuravlev et al. [2010] take a pure software-based scheduling approach to the resource contention problem for multicore memory systems. This article proposes detecting which pairs of applications are likely to interfere less with each other and scheduling them to execute on cores that share as small a number of resources as possible. Tang et al. [2011] show the negative impacts of memory subsystem resource sharing on real datacenter applications. They also show that pure software-based intelligent thread-to-core mappings can reduce the amount of memory subsystem interference different applications suffer and improve their performance. The mechanisms we propose in this work are orthogonal to those proposed by Zhuravlev et al. and Tang et al. as we address the problem of intercore memory system interference in a finer-grained fashion using a hardware/software cooperative approach: First, the mix of applications to be scheduled may be such that whatever software schedule is chosen high intercore interference will exist among the applications sharing multiple memory system resources. In such cases, pure software-based scheduling approaches can not be as effective. However, FST can provide performance and fairness improvements since it throttles applications fine-grained manner. Second, even if intercore interference can be somewhat reduced using better scheduling, after a number of applications have

been scheduled to share some memory system resources, an FST like approach can further improve system fairness and performance by dynamically controlling memory system interference at a finer-grained level.

Prior work on SMT processors [Cazorla et al. 2004; Luo et al. 2001a, 2001b; Tullsen and Eggers 1996] propose fetch policies to improve performance and/or fairness in such processors. These techniques are not applicable to the problem we address, since they mainly address sharing of execution pipeline resources and not the shared memory system. Eyerman and Eeckhout [2009] propose a technique to estimate the execution times of simultaneously running threads had they been run alone. This work estimates interference in the execution resources and does not deal with memory system interference in a detailed manner. As such, our proposed memory interference/slowdown estimation and source throttling techniques are orthogonal to this prior work.

Several prior papers investigated how to handle prefetch requests in shared resources [Ebrahimi et al. 2009, 2011a; Srinath et al. 2007]. Even though we do not consider prefetching in this article, our recent work [Ebrahimi et al. 2011a] describes how our FST proposal can be adapted to systems that employ prefetching. Finally, FST could be combined with very recent proposals for handling memory interference for multithreaded applications [Ebrahimi et al. 2011b] and using channel partitioning [Muralidhara et al. 2011], and we leave this for future work

7. CONCLUSION

We proposed a low-cost architectural technique, Fairness via Source Throttling (FST), that allows system-software fairness policies to be achieved in CMPs by enabling fair sharing of the entire memory system. FST eliminates the need for and complexity of multiple complicated, specialized, and possibly contradictory fairness techniques for different shared resources. The key idea of our solution is to gather dynamic feedback information about the slowdowns experienced by different applications in hardware at runtime and, based on this feedback, collectively adjust the memory request rates of sources (i.e., cores) to balance applications' slowdowns. Our solution ensures that fairness decisions in the entire memory system are made in tandem, thereby significantly improving both system performance and fairness compared to the state-of-the-art resource-based fairness techniques implemented independently for different shared resources. We have also shown FST is configurable by system software, allowing it to enforce thread priorities and achieve different fairness objectives. We conclude that FST provides a promising low-cost substrate that can not only improve the performance and fairness of future multicore systems but also ease the task of future multicore system software in managing shared on-chip hardware resources.

ACKNOWLEDGMENTS

Many thanks to Khubaib, Boris Grot, Carl Waldspurger, the HPS research group, and the anonymous reviewers for their comments and suggestions. We also thank Benjamin Lee for his feedback in earlier stages of this work. We acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources.

REFERENCES

- ADVANCED MICRO DEVICES. 2009. AMD's six-core Opteron processors. <http://techreport.com/articles.x/17005>.
- BITIRGEN, R., IPEK, E., AND MARTINEZ, J. F. 2008. Coordinated management of multiple interacting resources in chip multiprocessors: A machine learning approach. In *Proceedings of the Annual ACM/IEEE International Symposium on Microarchitecture*.
- CAZORLA, F. J., RAMIREZ, A., VALERO, M., KNJNENBURG, P. M. W., SAKELLARIOU, R., AND FERNANDEZ, E. 2004. QoS for high-performance SMT processors in embedded systems. *IEEE Micro*.

- CHOU, Y., FAHS, B., AND ABRAHAM, S. 2004. Microarchitecture optimizations for exploiting memory-level parallelism. In *Proceedings of the 31st Annual International Symposium on Computer Architecture*.
- DAS, R., MUTLU, O., MOSCIBRODA, T., AND DAS, C. R. 2009. Application-aware prioritization mechanisms for on-chip networks. In *Proceedings of the Annual ACM/IEEE International Symposium on Microarchitecture*.
- EBRAHIMI, E., LEE, C. J., MUTLU, O., AND PATT, Y. 2009. Coordinated control of multiple prefetchers in multi-core systems. In *Proceedings of the Annual ACM/IEEE International Symposium on Microarchitecture*.
- EBRAHIMI, E., LEE, C. J., MUTLU, O., AND PATT, Y. 2010. Fairness via source throttling: A configurable and high-performance fairness substrate for multi-core memory systems. In *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*.
- EBRAHIMI, E., LEE, C. J., MUTLU, O., AND PATT, Y. 2011. Prefetch-aware shared-resource management for multi-core systems. In *Proceedings of the Annual International Symposium on Computer Architecture*.
- EBRAHIMI, E., MIFTAKHUTDINOV, R., FALLIN, C., LEE, C. J., JOAO, J. A., MUTLU, O., AND PATT, Y. 2011. Parallel application memory scheduling. In *Proceedings of the Annual ACM/IEEE International Symposium on Microarchitecture*.
- EYERMAN, S. AND L. EECKHOUT, L. 2008. System-level performance metrics for multiprogram workloads. *IEEE Micro* 28, 3, 42–53.
- EYERMAN, S. AND L. EECKHOUT, L. 2009. Per-thread cycle accounting in SMT processors. In *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*.
- FEDOROVA, A., SELTZER, M., AND SMITH, M. D. 2007. Improving performance isolation on chip multi-processors via an operating system scheduler. In *Proceedings of the 16th International Conference on Parallel Architectures and Compilation Techniques*.
- GABOR, R., WEISS, S., AND MENDELSON, A. 2006. Fairness and throughput in switch on event multi-threading. In *Proceedings of the Annual ACM/IEEE International Symposium on Microarchitecture*.
- GLEW, A. 1998. MLP yes! ILP no! In *ASPLOS Wild and Crazy Idea Session'98*.
- GROT, B., KECKLER, S. W., AND MUTLU, O. 2009. Preemptive virtual clock: A flexible, efficient, and cost-effective QoS scheme for networks-on-a-chip. In *Proceedings of the Annual ACM/IEEE International Symposium on Microarchitecture*.
- HERDRICH, A., ILLIKKAL, R., IYER, R., NEWELL, D., CHADHA, V., AND MOSES, J. 2009. Rate-based QoS techniques for cache/memory in CMP platforms. In *Proceedings of the International Conference on Supercomputing*.
- HSU, L. R., REINHARDT, S. K., IYER, R., AND MAKINENI, S. 2006. Communist, utilitarian and capitalist cache policies on cmps: Caches as a shared resource. In *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques*.
- INTEL. 2008. First the tick, now the tock: Next generation Intel microarchitecture (Nehalem). Intel technical white paper.
- IYER, R. 2004. CQoS: A framework for enabling QoS in shared caches of CMP platforms. In *Proceedings of the International Conference on Supercomputing*.
- IYER, R., ZHAO, L., GUO, F., ILLIKKAL, R., MAKINENI, S., NEWELL, D., SOLIHIN, Y., HSU, L., AND REINHARDT, S. 2007. QoS policies and architecture for cache/memory in CMP platforms. In *Proceedings of the ACM SIGMETRICS Joint International Conference on Measurement and Modeling of Computer Systems*.
- JAHRE, M. AND NATVIG, L. 2009. A light-weight fairness mechanism for chip multiprocessor memory systems. In *Proceedings of the ACM International Conference on Computing Frontiers*.
- KIM, S., CHANDRA, D., AND SOLIHIN, Y. 2004. Fair cache sharing and partitioning in a chip multiprocessor architecture. In *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques*.
- KIM, Y., HAN, D., MUTLU, O., AND HARCHOL-BALTER, M. 2010a. ATLAS: A scalable and high-performance scheduling algorithm for multiple memory controllers. In *Proceedings of the International Symposium on High-Performance Computer Architecture*.
- KIM, Y., PAPAMICHAEL, M., MUTLU, O., AND HARCHOL-BALTER, M. 2010b. Thread cluster memory scheduling: Exploiting differences in memory access behavior. In *Proceedings of the Annual ACM/IEEE International Symposium on Microarchitecture*.
- KROFT, D. 1981. Lockup-free instruction fetch/prefetch cache organization. In *Proceedings of the Annual International Symposium on Computer Architecture*.

- LEE, J. W., NG, M. C., AND ASANOVIC, K. 2008. Globally-synchronized frames for guaranteed quality-of-service in on-chip networks. In *Proceedings of the Annual International Symposium on Computer Architecture*.
- LUO, K., FRANKLIN, M., MUKHERJEE, S., AND SEZNE, A. 2001a. Boosting SMT performance by speculation control. In *Proceedings of the International Parallel and Distributed Processing Symposium*.
- LUO, K., GUMMARAJU, J., AND FRANKLIN, M. 2001b. Balancing throughput and fairness in SMT processors. In *Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software*.
- MICRON. Datasheet: 2Gb DDR3 SDRAM, MT41J512M4 - 64 Meg x 4 x 8 banks, <http://download.micron.com/pdf/datasheets/dram/ddr3>.
- MOSCIBRODA, T. AND MUTLU, O. 2007. Memory performance attacks: Denial of memory service in multi-core systems. In *Proceedings of the 16th USENIX Security Symposium*.
- MURALIDHARA, S. P., SUBRAMANIAN, L., MUTLU, O., KANDEMIR, M., AND MOSCIBRODA, T. 2011. Reducing memory interference in multicore systems via application-aware memory channel partitioning. In *Proceedings of the Annual ACM/IEEE International Symposium on Microarchitecture*.
- MUTLU, O. AND MOSCIBRODA, T. 2007. Stall-time fair memory access scheduling for chip multiprocessors. In *Proceedings of the Annual ACM/IEEE International Symposium on Microarchitecture*.
- MUTLU, O. AND MOSCIBRODA, T. 2008. Parallelism-aware batch scheduling: Enhancing both performance and fairness of shared DRAM systems. In *Proceedings of the Annual International Symposium on Computer Architecture*.
- NESBIT, K. J., AGGARWAL, N., LAUDON, J., AND SMITH, J. E. 2006. Fair queuing memory systems. In *Proceedings of the Annual ACM/IEEE International Symposium on Microarchitecture*.
- NESBIT, K. J., LAUDON, J., AND SMITH, J. E. 2007. Virtual private caches. 2007. In *Proceedings of the Annual International Symposium on Computer Architecture*.
- PATIL, H., COHN, R., CHARNEY, M., KAPOOR, R., SUN, A., AND KARUNANIDHI, A. 2004. Pinpointing representative portions of large intel itanium programs with dynamic instrumentation. In *Proceedings of the Annual ACM/IEEE International Symposium on Microarchitecture*.
- RIXNER, S., DALLY, W. J., KAPASI, U. J., MATTSON, P., AND OWENS, J. D. 2000. Memory access scheduling. In *Proceedings of the Annual International Symposium on Computer Architecture*.
- SNAVELY, A. AND TULLSEN, D. M. 2000. Symbiotic job scheduling for a simultaneous multithreading processor. In *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*.
- SRINATH, S., MUTLU, O., KIM, H., AND PATT, Y. 2007. Feedback directed prefetching: Improving the performance and bandwidth-efficiency of hardware prefetchers. In *Proceedings of the International Symposium on High-Performance Computer Architecture*.
- TANG, L., MARS, J., VACHHARAJANI, N., HUNDT, R., AND SOFFA, M. L. 2011. The impact of memory subsystem resource sharing on datacenter applications. In *Proceedings of the Annual International Symposium on Computer Architecture*.
- TENDLER, J., DODSON, S., FIELD, S., LE, H., AND SINHARROY, B. 2001. POWER4 system microarchitecture. IBM technical white paper.
- TULLSEN, D. M. AND EGGERS, S. J. 1996. Exploiting choice: Instruction fetch and issue on an implementable simultaneous multithreading processor. In *Proceedings of the Annual International Symposium on Computer Architecture*.
- WECHSLER, O. 2006. Inside Intel core microarchitecture. Intel technical white paper.
- ZHANG, X., DWARKADAS, S., AND SHEN, K. 2009. Hardware execution throttling for multi-core resource management. In *Proceedings of USENIX*.
- ZHURAVLEV, S., BLAGODUROV, S., AND FEDOROVA, A. 2010. Addressing shared resource contention in multicore processors via scheduling. In *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*.

Received March 2011; revised December 2011; accepted January 2012