Bundling: Reducing the Overhead of Multiprocessor Prefetchers

Dan Wallin and Erik Hagersten
Uppsala University, Department of Information Technology
P.O. Box 337, SE-751 05 Uppsala, Sweden
{*dan.wallin, erik.hagersten*}*@it.uu.se*

Abstract

Prefetching has proven to be a useful technique for reducing cache misses in multiprocessors at the cost of increased coherence traffic. This is especially troublesome for snoop-based systems, where the available coherence bandwidth often is the scalability bottleneck.

The bundling technique presented in this paper reduces the overhead caused by prefetching in two ways: piggybacking prefetches with normal requests, and requiring only one device to perform the snoop lookup for each prefetch transaction. This can reduce both the address bandwidth and the number of snoop lookups compared with a non-prefetching system. We describe bundling implementations for two important transaction types: reads and upgrades.

While bundling could reduce the overhead of most existing prefetch schemes, the evaluation of bundling performed in this paper has been limited to two of them: sequential prefetching and Dahlgren's adaptive sequential prefetching. Both schemes have their snoop bandwidth halved for all commercial and scientific benchmarks in the study. The combined effect of bundling applied to these prefetch schemes lowers the cache miss rate, the address bandwidth and the snoop bandwidth, compared with a system with no prefetching, for all applications.

Bundling, will not reduce the data bandwidth introduced by a prefetch scheme. However, we argue that the data bandwidth is more easily scaled than the snoop bandwidth for snoop-based coherence systems.

1 Introduction

Many important applications spend a substantial time waiting for memory transactions [4, 14]. A cache miss rate of 1 percent may add as much as 1.0 to the overall CPI number, assuming an access cost to memory of 400 CPU cycles and a 25 percent ratio of memory instructions. Much research effort has been devoted to reducing the number of cache misses using various latency-hiding and latency-avoiding techniques, such as prefetching.

While most existing prefetch techniques efficiently reduce the amount of cache misses, they also increase the address traffic and snoop lookups, which are scarce resources in a shared-memory multiprocessor. This is especially true for systems based on snooping coherence, where each device has to perform a cache lookup for every global address transaction. The address networks of systems based on directory coherence are more scalable, since the address transactions are sent point-to-point. Still, systems based on snooping are often preferred because of their superior cache-to-cache transfer time. Note, that there is no difference in scalability of the data network between systems based on snooping coherence and systems based on directory coherence, since data packets can be sent point-to-point in both cases. Actually, commercial snoop-based systems have been built where the data network handles 50 percent more data traffic than the available snoop bandwidth supports [5].

The goal of this research is to use prefetching to reduce the cache miss rate *without* increasing the address traffic and snoop lookups. Our proposal, *bundling*, piggybacks prefetch transactions on ordinary cache miss transactions. An alteration of the coherence protocol allows for a very selective snooping, such that only one device needs to perform a snoop lookup on each prefetch transaction. Bundling in combination with existing prefetch proposals may actually decrease the bandwidth of the address bus, and more importantly, also decrease the coherence activity, i.e., the number of address snoops, compared with a system without prefetching.

This paper uses the simple, but effective, adaptive sequential prefetching of read transactions, proposed by Dahlgren et al. [6], as a baseline system. We propose bundling techniques for read and upgrade transactions. Next, the efficiency of bundled prefetch protocols are studied. Bundling reduces the address snoops with 49 percent and the cache misses with 12 percent relative to the original fixed sequential prefetch protocol on average. The bundled

adaptive protocol reduces the address snoops with 41 percent and the cache misses with 7 percent compared with the original adaptive protocol. The data traffic is similar for bundled and non-bundled prefetchers. We also discuss the complexity of the proposal and provide a more detailed description of one possible implementation of bundling on an existing architecture.

We end the paper with a performance discussion of various prefetch protocols and make a comparison between the bundled prefetch protocols and non-prefetching protocols. Combining the adaptive prefetch scheme with bundling allows us to cheaply prefetch read transactions as well as upgrade transactions, resulting in a miss reduction of 30 percent for the scientific benchmarks (SPLASH-2) and a miss reduction of 19 percent for the commercial benchmarks. It also *reduces* the snoop lookups by 20 percent and 12 percent respectively compared to a system without prefetching. It is interesting to note that we lower the miss rate *and* the snoop lookups for all studied applications.

While the evaluation section of this paper studies the effect of bundling on the Dahlgren sequential prefetching and pure sequential prefetching, it does not argue that these necessarily are the strategies of choice. However, it does make the case that various forms of bundling can help reduce the coherence overhead caused by most prefetch strategies. While bundling reads and upgrades may seem obvious, it has not previously been proposed. We are not aware of any prefetch study reporting a *reduction* of both cache misses and address snoops for all studied applications.

2 Background: Multiprocessor Prefetching

Several prefetching techniques have been proposed which efficiently reduce the cache miss rate of multiprocessors. The prefetches can either be handled in software or hardware. However, both software and hardware prefetching lead to certain disadvantages. Software prefetching [17, 18, 19, 23] relies on inserting prefetch instructions in the code and results in an instruction overhead. The hardware prefetching techniques [3, 6, 7, 11, 13, 15, 21] require hardware modifications to the cache controller to speculatively bring additional data into the cache. They often rely on detecting regularly accessed strides. A common approach to avoid unnecessary prefetches in multiprocessors is to adapt the amount of prefetching at run time [11, 13, 21]. These proposals introduce small caches that detect the efficiency of prefetches based on the data structure accessed. Baer and Chen proposed to predict the instruction stream with a look-ahead program counter [3]. A cache-like reference predictor table is used to keep previous predictions of instructions. Correct branch prediction is needed for successful prefetching.

Another hardware prefetch approach is to exploit spatial locality by fetching data close to the originally used cache line. A larger cache line size can achieve this. Unfortunately, enlarging the cache line size is not as efficient in multiprocessors as in uniprocessors since it can lead to a large amount of false sharing and an increase in data traffic. The influence of cache line size on cache miss rate and data traffic has been studied by several authors [9, 10, 12, 24, 25]. To avoid false sharing and at the same time take advantage of spatial locality, sequential prefetching fetches a number of cache lines having consecutive addresses on a read cache miss. The number of additional cache lines to fetch on each miss is called the prefetch degree.

Sequential prefetching in multiprocessors was first studied by Dahlgren [6]. The increased prefetch traffic tends to hurt multiprocessors more than uniprocessors. This is especially a problem in bus-based multiprocessors where the available snoop bandwidth is limited [22]. Dahlgren proposed two types of sequential prefetching schemes, a fixed version and an adaptive version. Prefetches are only generated on read misses in these schemes and the focus on the study was on prefetching to the second level cache.

The fixed sequential prefetch scheme issues prefetches to the K consecutive cache lines on each cache read miss. If the consecutive cache lines are not already present in a readable state in the cache, a prefetch message for each missing cache line is generated on the interconnect. The prefetch degree K is fixed to a positive integer in this scheme. The fixed sequential prefetch scheme requires only small changes to the cache controller of the prefetching cache. In addition to this, a special prefetch request has to be handled by the interconnect and the memory system.

The adaptive sequential prefetch scheme is identical to the fixed sequential prefetch scheme, except that the prefetch degree K can be varied during run time. The prefetch degree is varied based on the success of previous prefetches. Dahlgren's approach to finding an optimal value of K is to count the number of useful prefetches. The protocol uses two counters that keep track of the total number of prefetches and the number of useful accesses to prefetched cache lines. Prefetched cache lines are tagged for later detection. Every sixteenth prefetch the useful prefetch counter is checked. If the number of useful prefetches is larger than twelve, K is incremented. K is decremented if the number of useful prefetches is lower than eight or divided by two if less than three prefetches are useful. The scheme also has a method of turning prefetching on, since no detection can be carried out if the prefetch degree is lowered so that no prefetches are performed.

Both schemes proposed by Dahlgren reduce the cache misses significally. However, both the address and the data traffic are increased [6].

3 Bundling

Snooping protocols result in a much lower latency for cache-to-cache transfers than directory-based protocols. However, the available snoop bandwidth of snoop-based systems limits their scalability. Data packets do not need the broadcast capabilities and can be returned on a general network, such as a crossbar switch or a point-to-point network, and do not suffer from such limitations. An example is the architecture of Sunfire 6800, which has a data interconnect capable of transferring 14.4 GB/s while its snooping address network only supports 9.6 GB/s worth of address snoops [5]. The main goal of our proposal is to limit the snoop bandwidth consumed by the address network, while the amount of bandwidth used in the data network is considered to be less critical.

A simple form of bundling applied only to read transactions has previously been studied together with the capacity prefetching technique [24]. However, no evaluation of the possible performance gains of read bundling has previously been performed. This paper also extends the previous publication with a thorough discussion on how to implement bundled reads on an existing architecture, the SunFire 6800, as well as to also introduce bundling for upgrades. Bundling is also studied together with an adaptive prefetch scheme.

3.1 Packet Format

Traditionally, hardware prefetchers, e.g., sequential prefetchers, send address transactions for the original cache miss as well as for all prefetch transactions. The number of address transactions sent on the network can be significantly reduced if the original transaction, and its associated prefetch transactions are bundled into a single transaction. Each original transaction has to be extended with a prefetch bit mask that indicates which extra cache lines to prefetch beyond the original one according to Figure 1. While this would reduce the number of address transactions on the bus, it would not reduce the number of snoop lookups each cache has to perform, since existing coherence protocols require a snoop lookup also for each prefetch address.

Figure 1. Bundled transaction.

3.2 Read Bundling

The increase in snoop lookups has led us to alter the semantics of the coherence for two types of coherence activities, reads and upgrades, to limit the number of caches that snoops each prefetch transaction.

A single bundled read transaction consists of the address *A* of the cache miss and information about the address offsets to the *K* prefetches. The address offsets relative to the original address are encoded in a prefetch bit mask. All devices on the bus need to perform a snoop lookup for address *A*, but only the device owning cache line *A* performs snoop lookups for the prefetched cache lines. This device will reply with data for each prefetched address for which it is the owner. Otherwise, an empty NACK data packet will be supplied for the prefetched cache line. Since the states of the other caches are not affected by the prefetch transaction, they do not need to snoop the prefetch addresses.

This extension will not add any states to the ordinary cache protocol, but will require the main memory to add an one-bit state: Owner. The Owner bit is cleared on a *ReadExclusive* or an *Invalidate* request and is set again on a *Write-back* request. This will allow the memory controller to perform the bundling optimization as well.

3.3 Upgrade Bundling

The semantics of upgrade requests, i.e., a write request to cache lines in state Shared, has also been altered to limit the number of snoop lookups for bundled upgrade requests. First, we introduce one additional flavor of the Owner state in the cache coherence protocol. Second, a similar extension is made to the memory states.

The two flavors of the Owner state, Owner_2 (Owner 2) and Owner_m (Owner many) are used to keep track of how many shared copies that are available in the system. The first time a cache line becomes downgraded from the Modified state, the cache line will enter the Owner_2 state. In the Owner_2 state, we know that there is at most one other cache sharing the data. If additional read requests are issued to the same address by another processor, cache lines in the Owner_2 state will change their states to the Owner_m state. Cache lines in this state can be shared by an unknown number of caches. Figure 2 shows a transition state diagram.

On a prefetched upgrade request, the invalidate requests for each of the *K* consecutive cache lines being in the Shared state in the requesting device are bundled with the *Invalidate* request of address *A* on the bus. Address *A* is snooped by all devices, possibly causing a cache invalidation. If a device has address *A* in the Owner_2 state, it will also invalidate each of the prefetch cache lines it currently has in the Owner_2 state. Since cache lines in the Owner_2 state are shared by at most one other device, i.e., the requesting device, we know that the copy in the requesting device now is the only copy left. The device owning address *A* will send a reply to the requesting node indicating which of the bundled upgrade cache lines it now safely can put into the Modified

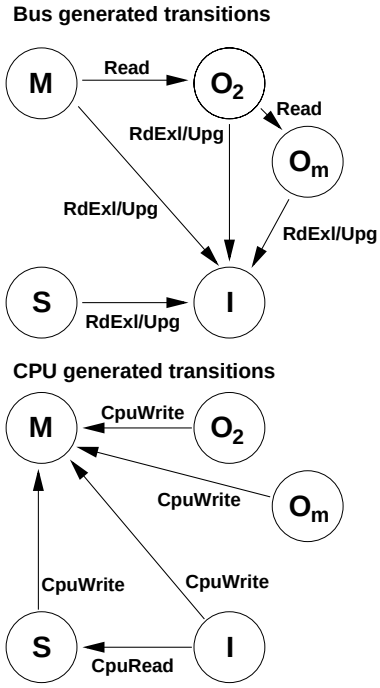

Figure 2. Transition state diagrams for bus and CPU generated transitions.

state. Cache lines being invalidated in the Owner$_m$ state can not be handled the same way since we do not know the number of sharers. In this case, only the original address will be invalidated.

Upgrade bundling should work well in all programs where at most two processors share a cache line. This behavior occurs in programs experiencing migratory sharing, which has been identified as one of the major sources of global invalidations in multiprocessors [12].

3.4 Details of the Bundled Snooping Protocol Implementation

This section presents an implementation of read bundling on a specific architecture. Although this section assumes a cache coherence implementation similar to that of Sunfire 6800 [5], read bundling implementations should be fairly similar in other modern snooping architectures.

In Sunfire 6800, snooping cache coherence is implemented using logically duplicated cache state tags: the snoop state (OSI) and the access state (MSI). A similar scheme is also used in the Sun E6000 servers [20]. The action taken for each snooped transaction depends on the snoop state. A service request may be put in the service queue for the cache as the result of the snoop state lookup, e.g., an *Invalidate* request or a *Copy-Back* request. The snoop state is changed before the next address transaction lookup is started. Entries may be added to the queue also when a cache snoops its own transaction, e.g., a *My-Read* request or a *My-ReadExclusive* request. Eventually, each service request will access the cache and change its access state accordingly. The cache's own requests are not removed from the head of the queue until the corresponding data reply is received and can thus temporarily block the completion of later service requests [20].

The UltraSPARC III processor, used in Sunfire 6800, implements the two logically duplicated states as a combination of the snoop result from the lookup in the main cache state and in the much smaller transition state table, which contains the new state caused by the snooped transactions still in the request queue. The hit in the transitional state table has precedence over a hit in the main cache state.

Bundled read prefetches in Sunfire 6800 will only retrieve the data if the owner of the original transaction also is the owner of the prefetched data. All caches snoop the address of the original read transaction address in order to determine if they are the owner. Only the owner will add an entry in its service queue: a *Read-Prefetch* request. Thus, the snoop bandwidth will not increase for the other caches. When the transaction reaches the head of the request queue, it is expanded to behave like one *Copy-Back* request for each prefetch cache line. If the cache is not the owner of a prefetch line, it will reply with a null-data NACK packet to the requesting CPU, which will simply remove the *My-Read* request from its service queue. The requesting cache must assume a shared snoop state for each prefetched cache line when the original request is sent out. This may create false invalidate requests if the null data is returned from the owner cache.

The owner state bit must be added to each cache line in main memory. The SunFire 6800 already has some state (gI, gS and gO) associated with each cache line for other reasons [5]. The extra bits used for these states are gained by calculating the ECC code over a larger data unit and are available at no additional cost. There is one such unused bit in memory that comes handy for the owner state bit. That bit should be cleared on the first *ReadExclusive* request to the cache line and set again on its *Write-Back* request.

3.5 Implementation Complexity

Our experience from designing commercial shared-memory systems has taught us that much of the implementation complexity lies in the details of an implementation – often at a level of details far below the description found in research papers. Here, we will nevertheless try to carry out a complexity discussion at a higher level.

Read bundling will introduce most complexity of the two bundling schemes. That is also why we covered it in some more detail in Section 3.4. While this kind of bundling

will neither alter the core of the coherence protocol nor add new states, it will introduce more corner cases in its implementation, e.g., the invalidation of prefetched data that are NACKed. However, we feel that our detailed description would solve such corner cases at the cost of a reasonable amount of logic adjacent to the service queue. Another major cost for read bundling could be adding the one Owner state to the memory if the existing system not already has a memory state associated with each cache line in memory.

The upgrade bundling does add more states to the caches and the memory. However, it does not alter the core coherence scheme, since both Owner states behave the same way from a global coherence point of view. Adding one state may force you to add one extra state bit, if there should not be any unused pattern using the current bits. The mechanisms to handle the corner cases for read bundling should be sufficient for implementing also upgrade bundling.

4 Simulation Environment

The Simics [16] full-system simulator is used in all the experiments. The simulation is execution-driven and models the in-order SPARC v9 ISA. We have implemented an invalidation-based MOSI (Modified, Owner, Shared, Invalid) protocol extension to Simics as a baseline cache coherence protocol. In all experiments, we model a bus-based 16-processor system with one level of 4-way associative unified data and instruction caches per CPU. Since our goal is to reduce the second level cache misses and we assume an inclusive cache hierarchy with a write-through first level cache, we have chosen to only model one cache level. A single level simulation will yield the same number of cache misses as for a two level hierarchy, but the miss-ratio will be lower since more read accesses will reach the cache compared with a multi-level cache hierarchy.

The cache miss characterization in this paper is similar to the one proposed by Eggers and Jeremiassen [8]. We include the conflict misses in the capacity miss category.

The studies are performed on the SPLASH-2 benchmarks [25] and two commercial workloads, SPECjbb2000 [2] and ECperf [1].

The cache size is chosen to match the data footprints. The SPLASH-2 programs are rather old benchmarks with small data footprints. Therefore, the cache size for the SPLASH-2 simulations is chosen accordingly to only 64 KB. At this cache size, the number of communication misses (false, true and upgrade misses) and non-communication misses (cold and capacity misses) are roughly equal and the different cache miss categorizes could be observed and evaluated for all applications. If a larger cache size is used, the diagrams are entirely dominated by cold misses and with a smaller cache size, very few communication misses occur. The workloads are chosen according to the default values specified in the SPLASH-2 release [25] with some minor changes: the Cholesky benchmark is optimized for the cache size, the FFT benchmark is run with 65536 data points, the Raytrace benchmark allocates a total of 64 MB global memory, and the Radiosity benchmark uses the small test scene provided in the distribution instead of the default room scene to limit the simulation time. The benchmarks are run using 16 parallel threads, and the measurements are started right after the child processes are created in all applications except Barnes and Ocean, where the measurements are started after two time steps.

ECperf and SPECjbb2000 are both commercial Java-based middleware benchmarks. ECperf is a benchmark modeling Java Enterprise Application Servers that use a number of Java 2 Enterprise Edition (J2EE) APIs in a web application. ECperf is a complicated, multi-tier benchmark that runs on top of a database server and an application server. SPECjbb2000 evaluates the performance of server-side Java. It can be run on any Java Virtual Machine. Both are commercial benchmarks, which set heavy demands on the memory and cache system. The SPECjbb2000 and ECperf workloads are configured according to Karlsson et al. [14]. The commercial benchmarks have larger data footprints and therefore the cache size is chosen to 1 MB for these applications for a more realistic mixture of cache misses. ECperf models the number of successfully completed "benchmark business operations" during a time period. Such operations include business transactions such as a customer making an order, updating an order or checking the status of an order. The ECperf transactions take long time and a total of 10 transactions are run with a 3-transaction warm-up period. SPECjbb2000 transactions take less time and we simulate 50,000 transactions including 10,000 transactions of warm-up time.

5 Efficiency in Prefetch Protocols

Figure 3 shows the efficiency of bundling for the fixed sequential prefetch scheme and the Dahlgren adaptive prefetch scheme. The prefetch strategies are evaluated in terms of cache misses, snoop lockups and data traffic for all benchmarks. The adaptive protocol *Dr* is the one described by Dahlgren [6]. The abbreviation *Dr* indicates that this is the adaptive protocol proposed by *D*ahlgren, which only prefetches on *r*eads. The *F3r* protocol is a *f*ixed scheme that issues prefetches to the next *3* consecutive cache lines for *r*ead misses, i.e., one read request followed by three prefetches.

The *F3ru* configuration prefetches the three consecutive addresses on each *r*ead and on each write generating *u*pgrades. The *F3Bru* configuration is the bundled protocol prefetching three cache lines on each read and upgrade miss. The *B* indicates that the protocol uses bundling. The

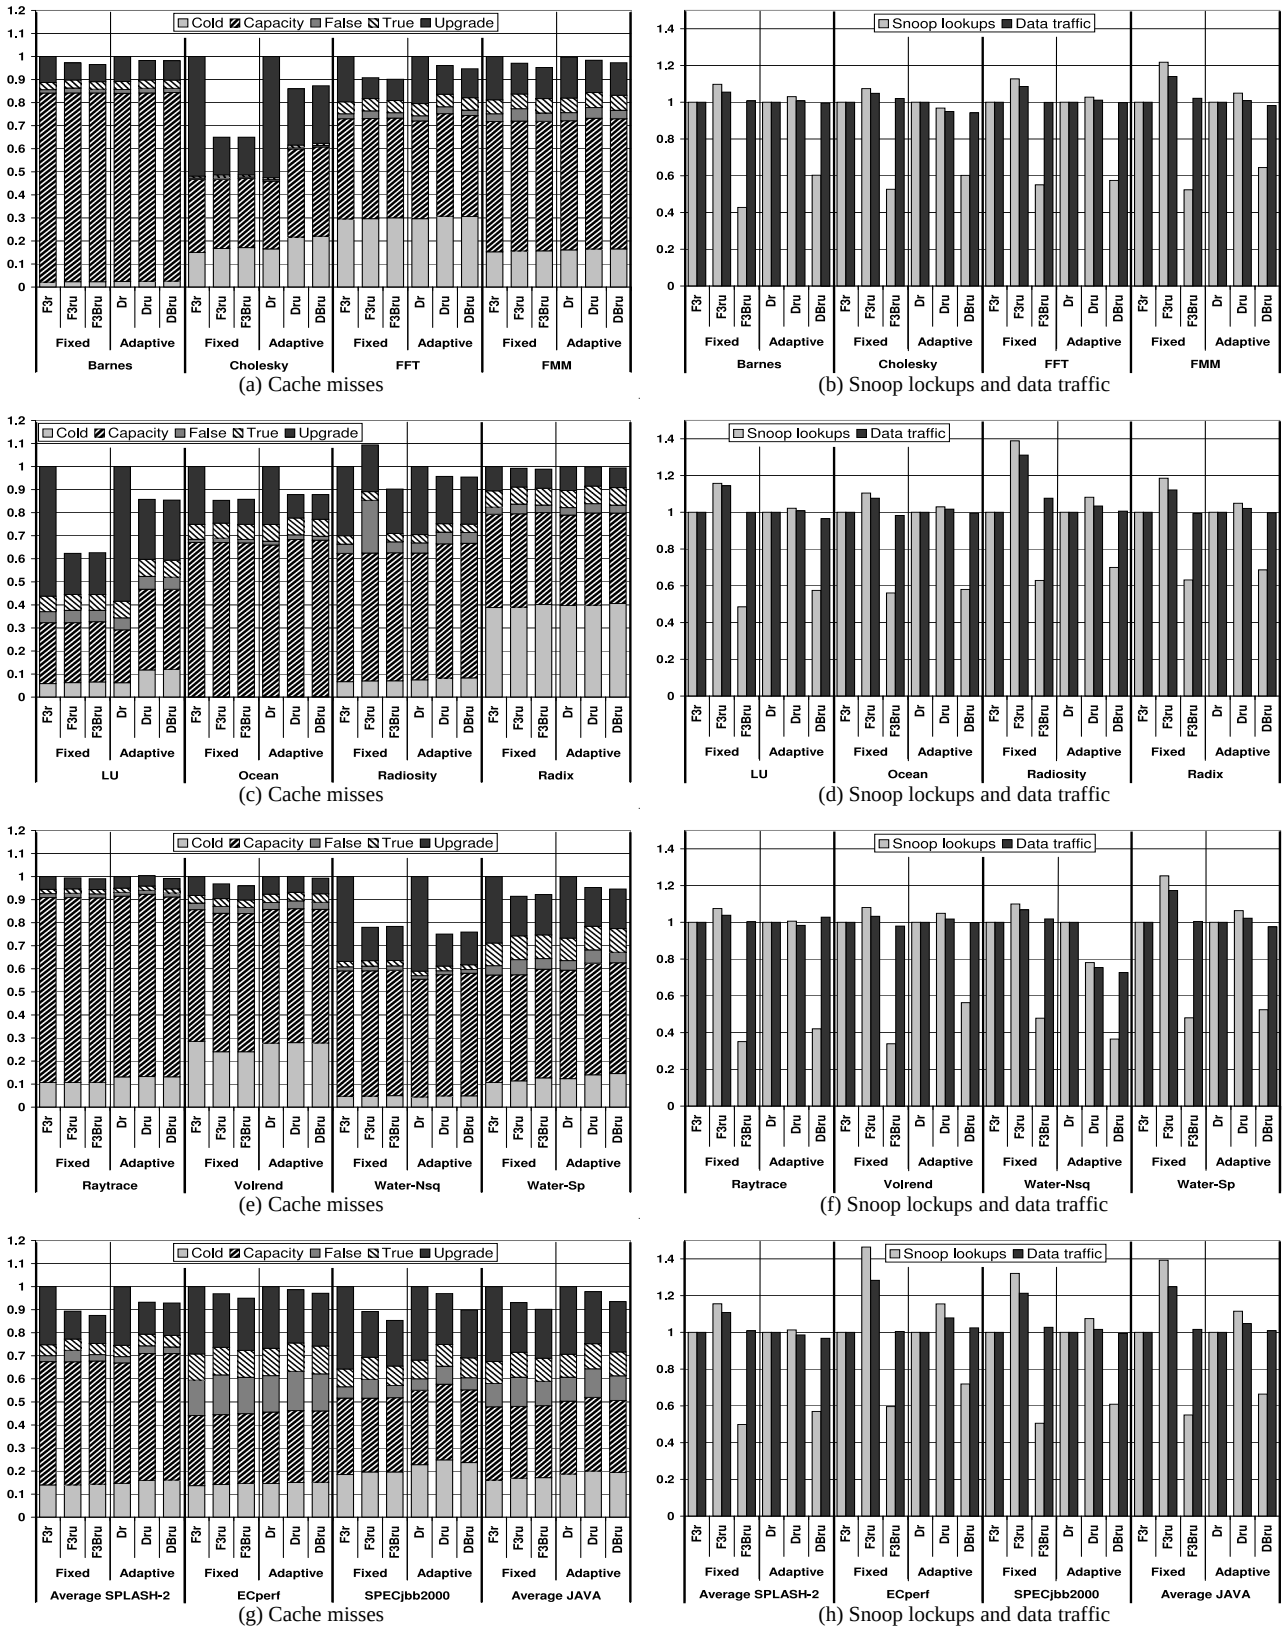

(a) Cache misses (b) Snoop lockups and data traffic

(c) Cache misses (d) Snoop lockups and data traffic

(e) Cache misses (f) Snoop lockups and data traffic

(g) Cache misses (h) Snoop lockups and data traffic

Figure 3. Influence of bundling on a fixed and an adaptive sequential prefetch scheme. Cache misses, snoop lockups and data traffic for three fixed, *F3r*, *F3ru*, *F3Bru* and three adaptive, *Dr*, *Dru*, *DBru* prefetch schemes are presented. The fixed schemes are normalized relative to the baseline fixed prefetch scheme *F3r* and the adaptive schemes relative to the baseline adaptive prefetch scheme *Dr*.

Dru and *DBru* are similar to the *F3ru* and *F3Bru* protocols except that they use adaptive prefetch degrees. The cache misses, data traffic and snoop lookups are normalized relative to the baseline fixed and adaptive schemes *F3r* and *Dr* in the figure. Hence, we can easily study the efficiency of the bundling proposals on the prefetch schemes.

Some of the bundled prefetch requests will get NACKed. Will not the effect on the cache miss rate be negative compared with non-bundled prefetching, since less prefetches are completed? There seems to be a fairly small difference. Actually, sometimes bundling seems to have a positive effect. If the owner of the original transaction is not also the owner of the prefetch data, this may indicate that they do not have a common history and do not belong to the same software object. Not prefetching could therefore be the action of choice. A small positive effect can be seen by comparing the miss rate for *F3ru* and *F3Bru* or *Dru* and *DBru* in ECperf and SPECjbb2000.

There are more cold misses in the adaptive protocol *Dru* than in *Dr* for some applications, e.g., Cholesky, LU, Water-Sp and SPECjbb2000. The reason for this is that useful prefetches are detected also for upgrades in the *Dru* scheme. Upgrades generally take advantage of a smaller prefetch degree than reads. This makes the *Dru* more restrictive at prefetching and causes the cold misses to increase in this protocol compared to the *Dr* protocol.

Prefetching on reads and upgrades leads to increased data traffic compared with only prefetching on reads. However, when the reads and upgrades are bundled, the data traffic becomes smaller. The cache misses are generally lower for protocols prefetching on both reads and upgrades than the baseline read prefetch protocols. On average, the bundled protocols has fewer cache misses than the non-bundled protocols. This is an effect of bundling since unnecessary prefetches for data belonging to separate software objects are avoided. Using bundling, prefetches can be issued on upgrades, thus decreasing the cache misses, without the negative effect on data traffic. There is also a positive effect of bundling on false sharing since the more restrictive prefetching prevents unnecessary prefetches in, e.g., Radiosity, Water-Sp and SPECjbb2000.

The largest difference between the bundled and non-bundled protocols is in address snoops. The bundled adaptive protocol *DBru* requires much less address snoops than the *Dr* and *Dru* protocols for all applications. The decrease in address snoops is 42 percent between *Dr* and *DBru* on average for all applications. The decrease is even larger compared with the non-bundled protocol prefetching also on upgrades. The fixed protocol always issues three additional prefetches on each miss. This makes the total number of prefetches much larger in this protocol. Bundling is therefore more efficient at reducing address snoops in this protocol. The average reduction is 49 percent between the *F3r* and *F3Bru* protocols. The bundling technique is more efficient at reducing address snoops in protocols which issue a large amount of prefetches per cache miss.

6 Performance Discussion

So what are the performance effects of bundling on a real system? Generally, the net effect on execution time of prefetch schemes can be shown to be either positive or negative depending on the parameters chosen for the simulation study. A prefetching computer system with plentiful of coherence bandwidth will be dominated by the positive effects of the lower cache miss rate, while the negative effects from increased queuing delay in the interconnect will dominate if the coherence traffic is close to the available bandwidth of the system.

The contention bottleneck makes it very difficult to simulate the potential performance gain of prefetchers. In applications that spend a lot of their time waiting for memory transactions, the simulated execution time will vary very much depending on the bandwidth assumptions. For these applications, the wall clock time will follow the amount of cache misses if the available bandwidth is chosen large. If the simulated bandwidth is small, the execution time will instead depend on a combination of the required bandwidth and the amount of cache misses.

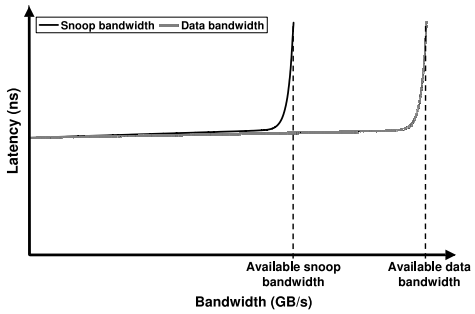

Figure 4. Effects of contention on latency.

Singhal et al. performed measurements of the impact of contention on memory latency in the snoop-based Sun E6000 family of servers [20]. A similar diagram has been drawn in Figure 4. The figure indicates that contention only has a modest influence on memory access time as long as the systems available bandwidth has not been reached. However, when the interconnect is contended, the access time increases largely. This makes it difficult to come up with one unbiased speedup number associated with a good-news/bad-news proposal, such as prefetching [1].

[1] Or rather, it is quite easy to come up with any desired number depending on what you would like to prove.

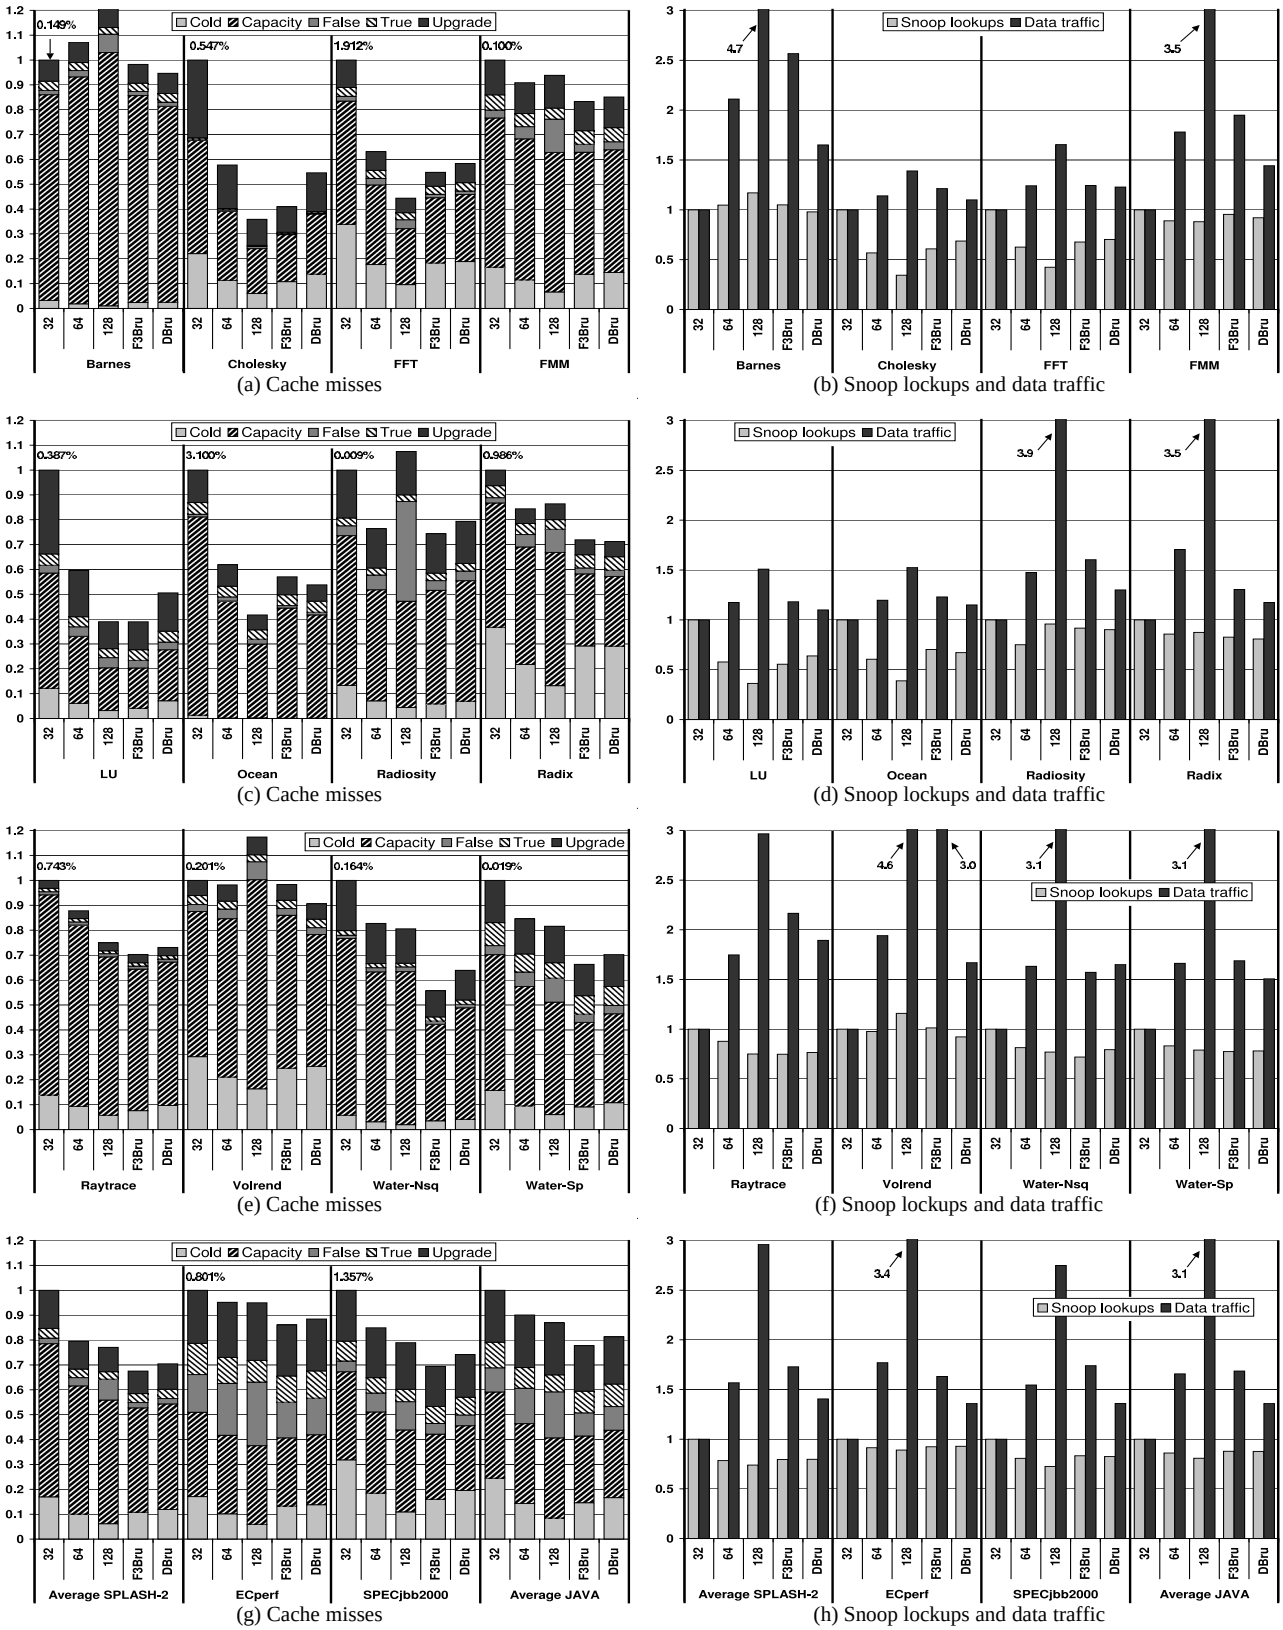

(a) Cache misses (b) Snoop lookups and data traffic

(c) Cache misses (d) Snoop lookups and data traffic

(e) Cache misses (f) Snoop lookups and data traffic

(g) Cache misses (h) Snoop lookups and data traffic

Figure 5. Cache misses, snoop lookups and data traffic for three non-prefetching protocols with different cache line sizes and the bundled fixed and adaptive sequential prefetchers. The results are normalized relative to the 32 B non-prefetching configuration. The cache miss ratios are indicated for the 32 B configuration for each application.

However, this paper is not about evaluating prefetching per se, but to evaluate the effects bundling has on prefetch algorithms. We have shown that bundling can cut the snoop bandwidth roughly in half and the cache misses by ten percent for the two baseline prefetch algorithms, while the data traffic is largely unaffected (Figure 3). We could easily have chosen simulation parameters to demonstrate a 100 percent speedup by limiting the snoop bandwidth. We could also have shown no speedup if the data bandwidth had been made the major bottleneck, or just a modest speedup if there are plentiful of both.

One could argue that we have chosen very primitive prefetch algorithms as our baseline systems, and this is why we can demonstrate the large cut in bandwidth. This may be a valid argument, which has lead us to compare the performance effects of bundling in combination with these primitive prefetch algorithms, with the non-prefetching 32 B protocol. Looking for example at Ocean in Figure 5, its drop in miss ratio from 3.1 to 1.7 percent for *DBru* would result in a CPI reduction from 4.1 to 2.7 (35 percent) using the rough memory access latency numbers given in the introduction. If the available snoop bandwidth had been the dominating bottleneck, a speedup of 33 percent could have been achieved, and a data bandwidth bottleneck would have yielded a slowdown of 15 percent. In Table 1, similar values are presented for all the applications. The table shows that if the available snoop and data bandwidths are large enough, the prefetching scheme will reduce the execution time for all applications. Even larger speedups can be expected if the scarce resource is snoop bandwidth. If instead the data bandwidth is the limiting factor, the performance will decrease with the bundled adaptive scheme.

	Cache miss bottleneck	Snoop BW bottleneck	Data BW bottleneck
Barnes	0.7	2.1	-65.1
Cholesky	16.0	31.3	-10.0
FFT	27.3	29.7	-22.8
FMM	1.3	8.0	-44.3
LU	13.8	36.2	-10.0
Ocean	34.9	32.8	-15.0
Radiosity	0.2	9.8	-30.0
Radix	14.3	19.2	-17.4
Raytrace	11.5	23.5	-89.4
Volrend	1.6	7.8	-66.9
Water-Nsq	5.1	20.6	-65.0
Water-Sp	0.6	21.9	-50.7
ECperf	5.1	7.2	-35.8
SPECjbb2000	14.9	17.5	-36.0

Table 1. The performance difference in percent between the bundled adaptive protocol *DBru* and the original non-prefetching 32 B protocol depending on whether the performance bottleneck is the number of cache misses, the snoop bandwidth or the data bandwidth. The cache miss bottleneck assumes 25 percent memory references and a 400 CPU cycle miss penalty. Non-memory refererences are assumed to have a CPI of 1.0.

In Figure 5, three non-prefetching protocols with cache line sizes of 32, 64 and 128 B are compared to the bundled fixed and adaptive prefetchers, *F3Bru* and *DBru*. The figures are normalized relative the 32 B configuration since this is the cache line size used in the prefetching schemes.

The bundled adaptive prefetch configuration, *DBru*, has 28 percent less cache misses, 19 percent fewer snoop lookups and 40 percent more data traffic than the 32 B configuration on average for all studied applications. The bundled fixed prefetcher, *F3Bru*, reduces the cache misses with 31 percent, the snoop lookups with 19 percent and increases the data traffic with 72 percent compared with the 32 B configuration on average for all applications.

However, based on the simulation results for protocols of different cache line sizes, we can conclude that the most probable choice of an "optimum" cache line for the design of a multiprocessor using our benchmarks is 64 B. At the 64 B cache line size, the cache misses and address snoops on average are significantly decreased compared with a 32 B cache line, while the data traffic is much smaller than in the 128 B configuration. The 64 B cache line size is also the design choice of the SunFire 6800 servers. The bundled adaptive prefetch configuration, *DBru*, has 11 percent less cache misses, the same amount of snoop lookups and 11 percent less data traffic than the 64 B configuration on average for all applications.

7 Conclusion

Prefetching is useful for reducing cache misses in multiprocessors. Also rather small cache miss rates of less than 1.0 percent can harm the overall performance severely in multiprocessors as the gap between processor speed and memory access time grows. However, many prefetching schemes largely increase the address snoops and data traffic.

Snoop-based systems are generally limited by how fast snoop lookups can be performed. By using the two bundling techniques proposed in this paper for read and upgrade transactions, the address snoops in prefetch protocols can be reduced. Bundling, lumps several snoop transactions together in a way that requires most of the caches to snoop only one of the transactions. We have investigated the efficiency of bundling in two different prefetchers, one fixed sequential prefetch scheme and one adaptive prefetch scheme proposed by Dahlgren [6]. Compared with the original adaptive Dahlgren proposal, the bundled adaptive Dahlgren protocol requires 42 percent less address snoops on average for all studied applications

We show that combining bundling with the Dahlgren

adaptive scheme gives a protocol that reduces the cache misses with 28 percent, the snoop lookups with 19 percent and increases the data traffic with 40 percent compared with the original 32 B cache line non-prefetching protocol on average for all studied applications.

This is the first prefetch paper that reports a *reduction* in the coherence activity, e.g., snoop lookups, as well as cache misses for all studied applications. Bundling requires only small changes to the coherence protocol. The technique is not limited to sequential prefetch schemes but could be used together with more sophisticated prefetch proposals.

Acknowledgement

We would like to express gratitude to Jim Nilsson at Chalmers University of Technology for providing us with the original version of the cache coherence protocol model for Simics.

This work is supported in part by Sun Microsystems, Inc., and the Parallel and Scientific Computing Institute (PSCI), Sweden.

References

[1] http://ecperf.theserverside.com/ecperf/.

[2] http://www.spec.org/osg/jbb2000/.

[3] J.-L. Baer and T.-F. Chen. An Effective On-Chip Preloading Scheme to Reduce Data Access Penalty. In *Proceedings of Supercomputing*, 1991.

[4] L. A. Barroso, K. Gharachorloo, and E. Bugnion. Memory System Characterization of Commercial Workloads. In *Proceedings of the International Symposium on Computer Architecture*, 1998.

[5] A. Charlesworth. The Sun Fireplane System Interconnect. In *Proceedings of Supercomputing*, 2001.

[6] F. Dahlgren, M. Dubois, and P. Stenström. Sequential Hardware Prefetching in Shared-Memory Multiprocessors. *IEEE Transactions on Parallel and Distributed Systems*, 6(7):733–746, 1995.

[7] F. Dahlgren and P. Stenström. Evaluation of Hardware-Based Stride and Sequential Prefetching in Shared-Memory Multiprocessors. *IEEE Transactions on Parallel and Distributed Systems*, 7(4):385–398, 1996.

[8] S. J. Eggers and T. E. Jeremiassen. Eliminating False Sharing. In *Proceedings of the International Conference on Parallel Processing*, 1991.

[9] S. J. Eggers and R. H. Katz. The Effect of Sharing on the Cache and Bus Performance of Parallel Programs. In *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*, 1989.

[10] J. R. Goodman. Using Cache Memory to Reduce Processor-Memory Traffic. In *Proceedings of the International Symposium on Computer Architecture*, 1983.

[11] E. H. Gornish. *Adaptive and Integrated Data Cache Prefetching for Shared-Memory Multiprocessors*. PhD thesis, University of Illinois at Urbana-Champaign, 1995.

[12] A. Gupta and W.-D. Weber. Cache Invalidation Patterns in Shared-Memory Multiprocessors. *IEEE Transactions on Computers*, 41(7):794–810, 1992.

[13] E. Hagersten. *Toward Scalable Cache-Only Memory Architectures*. PhD thesis, Royal Institute of Technology, Stockholm, 1992.

[14] M. Karlsson, K. Moore, E. Hagersten, and D. A. Wood. Memory System Behavior of Java-Based Middleware. In *Proceedings of the International Symposium on High Performance Computer Architecture*, 2003.

[15] D. M. Koppelman. Neighborhood Prefetching on Multiprocessors Using Instruction History. In *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques*, 2000.

[16] P. S. Magnusson, M. Christensson, J. Eskilson, D. Forsgren, G. Hållberg, J. Högberg, F. Larsson, A. Moestedt, and B. Werner. Simics: A Full System Simulation Platform. *IEEE Computer*, 35(2):50–58, 2002.

[17] T. C. Mowry. Tolerating Latency in Multiprocessors Through Compiler-Inserted Prefetching. *ACM Transactions on Computer Systems*, 16(1):55–92, 1998.

[18] T. C. Mowry and A. Gupta. Tolerating Latency Through Software-Controlled Prefetching in Shared-Memory Multiprocessors. *Journal of Parallel and Distributed Computing*, 12(2):87–106, 1991.

[19] T. C. Mowry, M. S. Lam, and A. Gupta. Design and Evaluation of a Compiler Algorithm for Prefetching. In *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*, 1992.

[20] A. Singhal, D. Broniarchyk, F. Cerauskis, J. Price, L. Yuan, C. Cheng, D. Doblar, S. Fosth, N. Agarwal, K. Harvey, and E. Hagersten. Gigaplane: A High Performance Bus for Large SMPs. In *Proceedings of IEEE Hot Interconnects*, 1996.

[21] M. K. Tcheun, H. Yoon, and S. R. Maeng. An Effective On-Chip Preloading Scheme to Reduce Data Access Penalty. In *Proceedings of the International Conference on Parallel Processing*, 1997.

[22] D. M. Tullsen and S. Eggers. Limitations of Cache Prefetching on a Bus-Based Multiprocessor. In *Proceedings of the International Symposium on Computer Architecture*, 1993.

[23] D. M. Tullsen and S. J. Eggers. Effective Cache Prefetching on Bus-Based Multiprocessors. *ACM Transactions on Computer Systems*, 13(1):57–88, 1995.

[24] D. Wallin and E. Hagersten. Miss Penalty Reduction Using Bundled Capacity Prefetching in Multiprocessors. In *Proceedings of the International Parallel and Distributed Processing Symposium*, 2003.

[25] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta. The SPLASH-2 Programs: Characterization and Methodological Considerations. In *Proceedings of the International Symposium on Computer Architecture*, 1995.