Inferential Queueing and Speculative Push for Reducing Critical Communication Latencies

Ravi Rajwar,† Alain Kägi† and James R. Goodman‡

†Microprocessor Research Labs
Intel Corporation
Hillsboro, Oregon 97124, USA
{ravi.rajwar,alain.kägi}@intel.com

‡Department of Computer Sciences
University of Wisconsin-Madison
Madison, Wisconsin 53706, USA
{goodman@cs.wisc.edu}

ABSTRACT

Communication latencies within critical sections constitute a major bottleneck in some classes of emerging parallel workloads. In this paper, we argue for the use of *Inferentially Queued Locks* (IQLs) [31], not just for efficient synchronization but also for reducing communication latencies, and we propose a novel mechanism, *Speculative Push* (SP), aimed at reducing these communication latencies. With IQLs, the processor infers the existence, and limits, of a critical section from the use of synchronization instructions and joins a queue of lock requestors. The SP mechanism extracts information about program structure by observing IQLs. SP allows the cache controller, responding to a request for a cache line that likely includes a lock variable, to predict the data sets the requestor will modify within the associated critical section. The controller then pushes these lines from its own cache to the target cache, as well as writing them to memory. Overlapping the protected data transfer with that of the lock can substantially reduce the communication latencies within critical sections. By pushing data in exclusive state, the mechanism can collapse a read-modify-write sequences within a critical section into a single local cache access. The write-back to memory allows the receiving cache to ignore the push. Neither mechanism requires any programmer or compiler support nor any instruction set changes. Our experiments demonstrate that IQLs and SP can improve performance of applications employing frequent synchronization.

Categories and Subject Descriptors

C.1.2 [**Processor Architectures**]: Multiple Data Stream Architectures (Multiprocessors)—*Multiple-instruction-stream, multiple-data-stream processors (MIMD)*

General Terms

Performance, Design

Keywords

Synchronization, data forwarding, inferential queueing

ICS'03, June 23-26, 2003, San Francisco, California, USA.

1 INTRODUCTION

The shared-memory programming model is now widely established as a leading paradigm for parallel computing. The shared memory abstraction is particularly attractive for irregular applications, where reasoning about program behavior and predicting performance may be difficult. Under the shared-memory model, in addition to holding values, memory also provides the means for synchronization and coordination of activities among processors. When multiple processors attempt to access a set of variables simultaneously and at least one processor updates at least one of the variables, a *data race* may occur wherein the execution outcome depends on the relative speed of the operations and the result of memory accesses becomes unpredictable. The most common method used to resolve data races and to enforce mutually exclusive accesses to regions of code, known as *critical sections*, is through the use of a lock.

Optimizing lock accesses associated with an actively shared critical section is both crucial and subtle: crucial because naïve locking algorithms can lead to disastrous performance [2, 3, 18, 17], and subtle because multiple processors may access the lock even while that lock guarantees exclusive access to the data it protects. The protected data is often modified with the result that the efficient lock handling only reveals subsequent delays in accessing the protected data. A wide range of synchronization mechanisms have emerged over the years [34, 12, 3, 13, 28, 26], and while no common mechanism is available in all architectures, virtually all architectures provide a hardware means for acquiring a lock atomically. Beyond this basic capability, numerous mechanisms have been proposed for enhancing the efficiency of locking in hardware, but to date very few implemented multiprocessor systems have incorporated these ideas.

It has been argued that parallel applications that spend too much time in critical sections could be restructured to minimize synchronization and in fact this is true of many structured scientific applications. These applications are computationally intensive, highly regular, and generally display easily exploitable parallelism. However, among new classes of programs emerging as important applications for parallel systems, online transaction processing workloads display radically different behavior from the traditional scientific applications: they are characterized by high communication miss rates [4, 20, 33]. A study by Ranganathan et al. [33] showed a large fraction of misses are generated within critical sections. For a given configuration of a 4-way system running the Oracle database engine, 20% of the execution time was spent servicing communication misses to dirty data in remote caches. These workloads are characterized by fine-grain updates of control data and frequent synchronization protecting such data. The protected data sets migrate among processors with the passing of the lock and contribute to a large portion of the access latencies to

dirty data in remote caches. They also noted most of these cache misses targeted only a small fraction of the total number of cache lines experiencing misses.

Gharachorloo et al. also note a large portion of execution latencies spent in critical sections [11]. With larger number of processors, faster processor speeds, and relatively increasing remote access latencies, processor stalls induced due to communication misses within critical sections will only increase and processors will be unable to generate misses early enough so as to hide memory access latencies to actively shared data.

Paper contributions. This paper addresses the problems outlined above by targeting the lock operations in conjunction with the data accesses protected by these locks. The paper makes two contributions:

1. **A case for Inferentially Queued Locks.** This paper advocates the use of *Inferentially Queued Locks (IQLs)* [31], that is, devoting hardware to build an orderly queue of lock contenders. The queue is speculative because the processor infers the existence, and limits, of a critical section from the use of synchronization instructions. We show for the first time how to implement IQLs for a general directory-based system and without making assumptions about the coherence interconnect.
2. **Speculative Push**. This paper proposes a new technique, *Speculative Push (SP)*, for reducing the miss latency associated with data accesses within critical sections. SP allows the cache controller of a processor currently holding a lock not only to defer momentarily its response to a request for the cache line holding the lock, but to provide additional modified cache lines at the same time, anticipating misses likely to occur immediately after the requestor has acquired the lock. By overlapping the transfer of the data with the lock, the communication latency experienced within a critical section can be reduced. To our knowledge, this is the first hardware technique to convert data misses in a critical section from a multi-hop transaction to a local access.

Intuition for IQLs. IQLs are motivated by two observations about performance loss due to lock interference. First, when a processor requests a lock for purposes of acquiring the lock, it will likely spin-wait upon discovering the lock is already held. If the response to this request is delayed briefly, any increase in the probability of the lock having been released will increase the success rate for the initial attempt, and thereby reduce total communication. The latency to acquire a held lock is optimal if such a request is serviced immediately after the lock is released. Second, if such a request is serviced immediately, that is, while the lock is still held, the release of the lock will probably be delayed, because releasing it will require the lock to be re-obtained in a writable state. Again, any delay resulting in an increase in the probability of the lock having been released is likely to improve performance, not only by reducing total communication, but also by avoiding a delay in releasing the lock. IQLs extend the notion of buffering external requests by applying it to cache lines inferred to contain a synchronization variable. By delaying the service for a small and bounded period, and servicing the deferred request as soon as the lock is inferred to be released, many critical sections can be fully executed and the lock released without interference from other processors. In addition, the transfer of locks occurs directly between the two nodes involved without the coherence network being in the critical path. In the presence of frequent synchronization to migratory locks there is great benefit in optimizing the lock access: network contention is reduced, thus having a positive effect on memory system performance.

Intuition for SP. The performance of any scheme optimizing data transfer within critical sections depends on the accuracy of correctly predicting which processor will acquire the lock and use the data. An advantage of IQLs is the early, accurate knowledge of the next owner of a lock. IQLs allow inference of the presence and extent of critical sections in programs. Assuming such knowledge of locations and sizes of these critical sections in a program, information associated with these locks can easily be tracked. Once the pairing between critical sections and the data they protect has been established, SP forwards the actively shared data to the requesting processor, along with the lock. On acquiring the lock, the requestor finds in its cache the data it was unable to prefetch. SP forwards the data in an exclusive state thus allowing the processor to modify the cache line without experiencing a further delay it would otherwise suffer if it had initially retrieved the line for reading.

For critical section accesses, SP has inherent advantages over more traditional approaches for latency reduction such as prefetching and compiler-assisted data forwarding. In situations of contended locks, prefetching is not sufficient as the processor would spin waiting for the lock and would generate data requests only once the lock has been acquired. SP transfers data as soon as the data is ready to be forwarded and does not interfere with the execution of the processor performing the push. The speculative mechanism can also adapt to run-time behavior. Speculative Push provides a double performance gain for data that is read before being written to: its initial access is overlapped with the lock transfer, and it does not have to be upgraded for writing.

We restrict SP to modified lines because, while some shared lines may also result in misses, such shared lines would already be present in the requestor's cache—data in cache lines previously read but not modified by the requestor in a previous execution of the same critical section would probably still be in the cache. Indeed, our experiments suggest that much of the benefit is derived simply from pushing cache lines into caches where they were present in an earlier execution.

Paper outline. In Section 2 we discuss IQLs and show how IQLs can be efficiently supported naturally in modern systems by using existing cache coherence protocols. We use a snoop and a directory-based protocol. Then we extend the base protocols to incorporate mechanisms for SP in Section 3. In Section 4 and Section 5 we present results. We discuss related work in Section 6 and conclude in Section 7.

2 INFERENTIALLY QUEUED LOCKS

Processors with non-blocking caches allow multiple outstanding requests to the memory system. Such processors use special buffers such as miss status holding registers (MSHRs) [21] to track the pending memory requests. Multiprocessor systems use such structures to buffer requests from other processors to caches lines that are in a pending state.

Inferentially Queued Locks (IQLs) extend the notion of buffering external requests by applying it to cache lines inferred to contain a synchronization variable. By delaying the service at most for a small and bounded period, and servicing the deferred request as soon as the lock is inferred to be released, many critical sections can be fully executed and the lock released without interference from other processors.

In particular, the lock transfer occurs directly between the two nodes involved eliminating the coherence network from the critical path. Contrast IQLs with a conventional system that may require

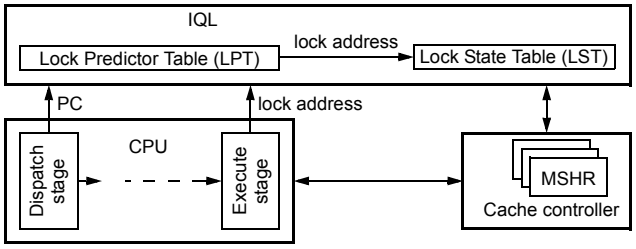

Figure 1. IQL organization

The Lock Predictor Table (LPT) and Lock State Table (LST) are the two new structures.

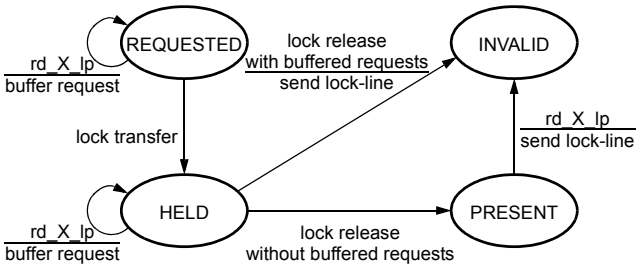

Figure 2. LST state transition diagram

Only a subset of the transitions is shown. Arcs between states represent transitions labeled with event/action pairs. A horizontal line separates event (above) and action (below). If there is no action, the line is omitted.

many additional network transactions to transfer a lock: a directory-based system may not be able to transfer the lock directly from producer to consumer passing instead the lock through an intermediate node (home), a lock release may require upgrade permissions from the directory, or a broadcast-based system may see repeated lock requests. This additional traffic may artificially augment the time that a lock is held, thus increasing the critical path of the overall computation.

The IQL method is speculative because the hardware does not have sufficient information, nor does it require such information, to prove even that a lock is being employed, but infers this information from program behavior, and specifically infers acquire and release points delimiting a critical section.

In the next subsection, we discuss the basic mechanisms involved in IQLs. In Section 2.2 and Section 2.3, we discuss details of implementing IQLs on broadcast and directory systems.

2.1 Basic mechanisms

A processor can *have* a lock in two different ways: (1) a process running on the processor has acquired a logical entity, a lock, and (2) the cache has a shared or exclusive copy of a cache line containing the lock. In the first case we refer to a lock as being *acquired*, *held*, or *released*. In the second case we refer to a lock-line or lock variable as being *present*, *requested*, *sent*, or *received*. A requestor is a processor that has requested a lock-line and for which it may be inferred that this processor is attempting to acquire a lock. A responder is a processor holding a writable copy of a lock-line sought by a requestor. Predicting synchronization events, and distinguishing between simple atomic read-modify-write operations, and the acquisition of a lock have been studied elsewhere [31]. We assume the processor has predicted, using a Lock Predictor Table (LPT) and based on previous executions of the code sequence, the event to be an acquiring of a lock. On a mispredict, performance may be degraded but the program will be correctly executed.

We now outline the basic IQL mechanism as shown in Figure 1. The IQL protocol is invoked on actions, both local and external requests, involving cache lines predicted to contain lock variables. The local processor must track when it acquires a presumed lock and when it has the lock-line in its cache. The processor must also track the release of such a lock and initiate actions triggered by the release. A Lock State Table (LST) is used for this purpose. The LST is indexed by the PC address of the synchronizing instruction identifying the critical section. All predicted locks known to the cache controller through the LST are in one of four states: (1) INVALID: the lock-line is not present in the cache and the local processor (probably) does not hold the lock, (2) PRESENT: the lock-line is present in the local cache but the lock is not held by the local processor, (3) HELD: the local processor holds the lock, and (4) REQUESTED: the local processor has requested the lock but does not have the lock.

To assure that the IQL protocol is only invoked for predicted lock operations, a read-for-exclusive request (rd_X) is specially annotated as lower priority, rd_X_lp. This request can be deferred for a brief but bounded time. Otherwise it is identical to rd_X.

Figure 2 outlines a simplified state transition diagram for the LST. When a processor predicts a lock-acquire, a lock-line already present in the local cache is marked HELD in the LST otherwise space is allocated in the cache, and a rd_X_lp is issued and the corresponding LST entry is marked REQUESTED. When the request is eventually serviced, the LST entry transitions to HELD. A subsequent write to the byte address predicted to be a lock (and thus predicted to correspond to a lock release) results in the LST entry being set to PRESENT. When a lock-line is evicted from the cache for whatever reason, including invalidation, the LST entry is marked INVALID. The LST is consulted on any incoming rd_X_lp request. If the LST entry for the requested cache line is in state PRESENT, the request is handled as any other read-for-exclusive (rd_X) request. If the cache line is in state HELD or REQUESTED, the line is marked for special action upon the release of the inferred lock, and the request is buffered in an MSHR. A subsequent inferred release of the lock triggers the servicing of the buffered rd_X_lp.

Since cache coherence protocols already serialize rd_X requests, a queue of requesting processors is easily constructed with minor support from the coherence protocol. The first rd_X_lp request transfers the apparent owner of the lock line to the requestor thus making it the recipient of any subsequent request. By this means, a queue of requesting processor is formed, with each processor receiving the lock-line sequentially in the order of the original request. Regular priority reads, whether for-exclusive or shared, are handled separately. Interactions between regular- and low-priority requests only occur in rare cases where the lock protocol is being violated, in the presence of false sharing, or the hardware has otherwise misspeculated.

We next discuss how IQLs can be supported on two popular design approaches: broadcast- and directory-based systems. We take two protocols from commercial systems and supplement them to support IQLs.

2.2 IQLs and broadcast-based systems

We use a snoop design similar to the Sun Gigaplane [36, 7]. The protocol uses a split-transaction, pipelined address bus with support for a large number of outstanding transactions and out-of-order responses. The bus implements an invalidation-based three-state (Owned, Shared, Invalid) snooping cache coherence protocol

and the cache implements a MOESI protocol. The cache with the requested block in Owned state (as seen by the bus) will respond to the next request for that block. IQLs can be supported naturally with the above protocol as the protocol already has a notion of queues. Any processor placing a rd_X_lp on the bus will respond to the next processor which places a rd_X_lp for the same block on the bus. With IQLs, the request is not serviced until a lock release. By placing a rd_X_lp on the bus, the processor joins the queue of lock requestors at the tail and is serviced when the earlier requestor releases the lock. IQLs for broadcast systems were discussed earlier elsewhere [31].

2.3 IQLs and directory-based systems

We now discuss directory protocol support. Directory protocols allow cache coherence to scale to many processors. Directory protocols store a directory state for each memory block currently cached in any node. Commonly, directory protocols implement an invalidation-based scheme. With directory protocols, state information for a block is obtained from a directory through network transactions, and communication with various cached copies is performed by explicit messages using an arbitrary network. Typically, information for a block can be found in a fixed directory known at the time of the request.

Two popular approaches distribute directories either with memory or with caches. Memory-based schemes store directory information for a cache block at the home node of the block [23, 25, 11]. In cache-based schemes, the sharing information is distributed among the various copies (rather than at the home node). Each cache block contains a pointer to the node with the next cached copy of the block in a distributed linked-list organization [15, 27, 6]. For cache-based directories, IQLs can be supported relatively easily since there already exists a notion of queues for cache blocks. We focus our discussion on memory-based directories.

We use the SGI Origin-2000 coherence protocol [23] to discuss IQL implementation. Other protocols can be adapted likewise. The protocol supports the MESI (Modified, Exclusive, Shared, Invalid) states and is non-blocking: memory does not buffer requests while waiting for other messages to arrive. The protocol also supports request forwarding for three party transactions and silent evictions of clean-exclusive lines. The protocol does not rely on an ordered network. Two virtual channels are provided and deadlock in the request network (due to request forwarding) is broken by the use of back-off messages.

Memory is the owner for all clean lines in the system, thus memory services any request for clean data immediately. In addition, rd_X requests cause transfer of exclusive ownership to the requestor and sending of invalidates to other holders of cached copies. The holders of cached copies subsequently send invalidate acknowledgments to the requestor. Requests to blocks not owned by memory are forwarded (as an intervention) to the owner (and in the case of a rd_X request, the requestor becomes the owner). The directory enters a transitional Busy state for the particular memory block until it receives a revision message from the previous owner. All requests received by the directory for a memory block in Busy state are NACKed. A sharing bit-vector associated with each cache block identifies the processors holding shared copies of the block.

Under the base protocol, the directory enters a transitional busy state while tracking down an exclusive copy in the system. For IQLs, the directory, instead of NACKing another request while in busy state, forwards the request to the previous requestor. If the previous requestor's request is still pending, the intervention is buffered. Forwarding a request to the last requestor guarantees a processor receives at most one intervention for a given memory block. An additional bit, the synch_bit, is used per directory entry to determine whether the IQL protocol should be invoked. The owner pointer is overloaded to store the last requestor when the synch_bit is set. The sharing bit-vector is also used to track members of the inferential queue. Alternate ways to capture such information include using unused state bits in the encoding.

To understand how queues are created we step through a simple example in Section 2.3.1 and then discuss how the directory detects and handles queue breakdowns (in the case of write-backs) in Section 2.3.2.

2.3.1 Constructing queues

Consider three processors, P1, P2, and P3 attempting to enter a critical section. P1 issues a rd_X_lp to the directory. Since the directory exclusively owns the block, it responds to the request with data and enters an Exclusive state. P1 is now the owner of the block. Subsequently, P2 issues a rd_X_lp request to the directory. The directory forwards P2's request to P1, marks P2 as the last requestor, and enters the Busy state. The bit vector now has 2 bits set: P1 and P2. In addition, the synch_bit is also set. Now, P3 sends a rd_X_lp request to the directory. The directory state is Busy. However, the synch_bit is set. Instead of NACKing the request, the directory forwards it to the last requestor, P2. In addition, P3's bit is set in the bit vector. Thus, we have P3 waiting for P2's response, which is waiting for P1's response. A processor buffers an intervention until the lock is released. At the time when the intervention is serviced, a revision message (part of the base protocol) is sent to the directory. On receiving the revision message, the directory unsets the processor's bit in the bit vector. If only one bit is set in the vector, the last requestor automatically becomes the owner. Under this situation, the directory unsets the synch_bit and leaves the Busy state, entering the Exclusive state. A revision message exists for every rd_X_lp request serviced, thus ensuring the directory will eventually transition into an Exclusive or any other stable state. A simplified state transition diagram for the IQL protocol is shown in Figure 3.

2.3.2 Handling queue breakdown

The distributed nature of the IQL queue (the directory does not track the order in which requests are received: it only marks who has requested the block), may cause the queue to break-down due to write-backs. For the example in the previous section, suppose P1 is writing back the block to memory. Under the base protocol, P1 can ignore P2's intervention since the directory can recognize this race condition and detect P2 will not be serviced by P1. The directory does this as it keeps track of the owner P2. However, under IQL, if there are multiple bits set in the bit-vector (the synch_bit is set and a queue exists), the directory cannot determine the identity of the processor which will not receive a response from P1; the directory does not remember it forwarded P2's request to P1. While this race condition is indeed rare, it must be handled correctly, if not efficiently.

We adopt a simple approach to handle queue breakdowns. On receiving a write-back to a block with a synch_bit set and more than two bits set in the bit-vector (if only two bits are set, the directory can uniquely determine the processor that will not receive a response), the directory unsets the synch_bit—the directory is breaking down the queue. Doing so, the protocol behaves like the base protocol with an additional side effect: NACKs are sent to all processors in the bit-vector. When the processor's retry on receiving the NACK, the directory can detect this by a bit in the retried message. At that point, the directory unsets the bit in the bit-vector corresponding to the requestor. When the bit-vector has no more bits set, the directory entry enters an Unowned state. Such a mech-

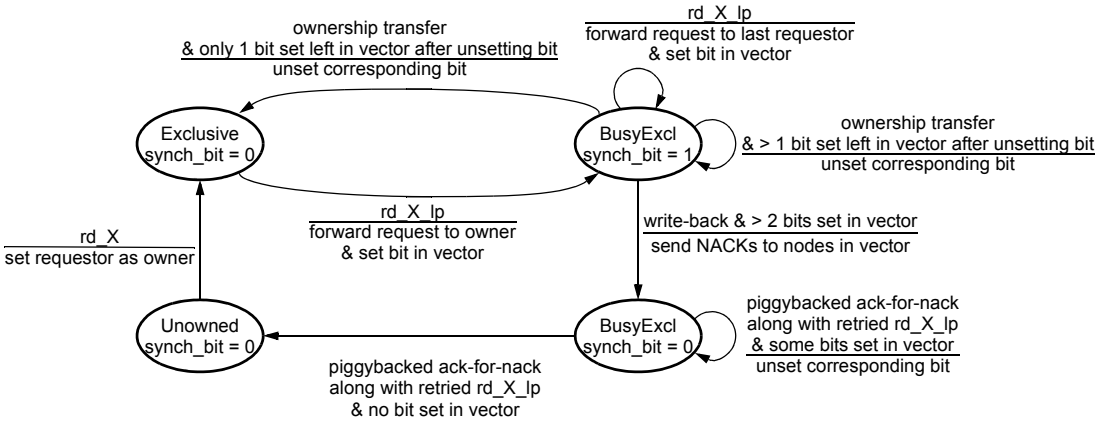

Figure 3. Protocol transitions for inferentially queued locks in a directory-based system

Only a subset of the transitions is shown, this subset focuses on the process of forming a queue. Arcs between states represent transitions labeled with event/action pairs. A horizontal line separates event (above) and action (below). The queue-breakdown sequence is shown when a write-back occurs.

anism guarantees a directory will eventually transition into an Unowned state.

With conceptually simple changes to the way the directory protocol works for certain types of requests, and some additional bits in the directory entry, IQLs can be efficiently supported.

3 SPECULATIVE PUSH

Speculative Push (SP) aims at constructing a link between critical sections and the data they protect. Once the link is established, whenever a request for a lock-line is received, SP also forwards any predicted data to the requesting processor along with the lock-line. This allows the lock requestor to find both the lock and any protected data (which it was unable to prefetch) readily available in writable state in its local cache upon entering the critical section. The requestor now experiences minimal delay in executing its critical section, virtually eliminating latency associated with access to data protected by migratory locks.

SP may initiate data transfer of predicted critical section data even before the requestor has acquired the lock, and the data is forwarded in Exclusive (as opposed to Modified) state. The pushed data is also written back to memory, providing two key benefits: (1) if the target node does not reference the pushed data, the data can be silently evicted thus doing no harm, and (2) the target of the push now has the important option of ignoring the push if no local buffer space is available.

In our simulations, we were very conservative in accepting the push—a pushed cache line never evicted a valid cache line, that is, a push was accepted only if an invalid line was available. This approach is surprisingly effective, in part because migratory lines that are frequently written are frequently invalidated, leaving available invalid lines in the corresponding cache sets. Even if the push is rejected, benefits may accrue. Specifically, when the rejected cache line is subsequently accessed, it can be supplied from the directory, avoiding the three-hop latency that would have occurred without the attempted push.

SP thus provides a double performance gain for data that is read before being written to: its initial access is overlapped with the lock transfer, and it does not have to be upgraded for writing.

As shown earlier in Figure 1, IQL has two structures: the LPT to infer critical sections, and the LST to invoke the IQL protocol partly based on information received from the LPT. Figure 4 shows the SP hardware. SP extends the LST by recording data accesses while the inferred critical section is executed and by deciding if predicted data is to be forwarded along with a released lock.

The basic steps of SP are:

1. Establish and record the association between critical section data and a lock.
2. Enable (or disable) the optimization by assigning a confidence level to pairings determined in Step 1.
3. Perform the speculative push.

3.1 Predicting lock and data pairings

When IQL infers a critical section, the SP hardware starts recording addresses of accesses performed while the processor holds the lock. These addresses become candidates for forwarding and are stored in an extended LST entry shown in Figure 5. Besides the lock address, each extended LST entry stores a valid bit, a saturating counter to establish a confidence level, and data addresses with associated access bits (A).

The LST stores cache-aligned addresses for two reasons: (1) we want to match SP with existing controllers and existing hardware manipulates data at the granularity of a cache line; (2) false sharing with critical data is considered poor programming practice; therefore all data in a cache line are assumed to be either related or unused. SP may record either individual addresses or address ranges. In our experiments at most two cache lines generally caused a write-miss during the execution of a critical section. Therefore, our LST stores individual addresses. Figure 5 shows the LST entry for supporting SP.

When an access results in a write fault, the SP hardware allocates a new entry for the address. If no free entry is available, the

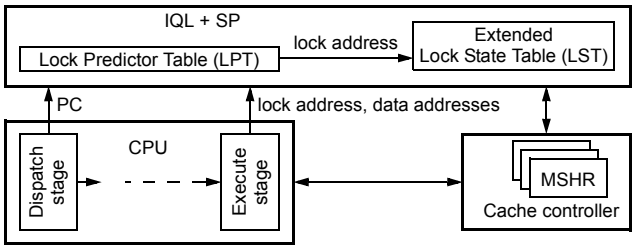

Figure 4. Speculative Push hardware organization

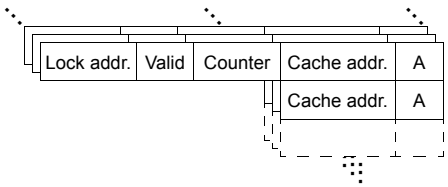

Figure 5. LST entry extended for Speculative Push

entry with the lowest confidence is evicted (see next sub-section). Ties are broken randomly. For the purpose of SP, distinguishing between finding a lock held or simply finding the lock present but not held is not necessary. In both cases, the triggering event causes the candidate cache lines to be pushed.

Two actions identify candidates for future pushes: lines that caused write misses during a critical section and lines that have been pushed into the cache (and therefore do not cause write misses). While not all data written in a critical section is migratory, write-misses capture data written within the critical section and obtained from the memory system. If such misses occur repeatedly, one can speculate the data is migratory. Local private data may also cause a write miss but is more likely to remain in the cache, since it will not be invalidated by another processor. Of course, various critical sections may touch many or few cache lines, but the first lines touched after acquisition of the lock seem especially important.

3.2 Prediction confidence

A saturating counter is added to each LST entry data address to assign confidence in the data for forwarding. Also associated with each candidate address is an access bit (A) that is set each time the critical section accesses this address and cleared each time after the execution of a presumed critical section ends. Before the access bits are cleared, SP inspects them and increases the counter for each set bit or decreases it for each cleared bit. A counter reaching the maximum value enables SP for the data address. This avoids a cache line being repeatedly pushed but never written to.

Repeated evictions of a candidate address also causes the counter to decrement. Repeated evictions is a sign that the addresses accessed inside a critical section vary from one execution to the next preventing effective data forwarding. A counter reaching the minimum value disables the optimization.

More sophisticated predictors, for example using information about remote requests for modified data in the local cache, or combining collective information from multiple nodes about migratory patterns, may improve the effectiveness of SP.

3.3 Speculative Push protocol

Once the data association has been identified, the SP protocol is initiated. The SP sequence involves two separate actions to maintain memory ordering requirements and achieve high performance:

1. Determine the forwarding data path for the push.
2. Order the speculative push in the memory order thereby granting coherence permissions to the target.

We discuss these two actions in Section 3.3.1 and Section 3.3.2. In Section 3.3.3 we discuss an interesting approach to match up speculatively pushed data messages with their appropriate coherence permission messages correctly. The issue arises because no ordering guarantees are assumed from the network or coherence protocol and we demonstrate a simple solution to handle the absence of any ordering guarantees.

3.3.1 Determining the forwarding data path

The predicted data (present in the local cache in modified state) can be pushed by the initiating processor either directly to the target processor (with a write back also sent to memory) or is pushed via memory (i.e., the data is written back to the directory which will forward the data to the target node). If the data is pushed via memory, the initiating processor sends address hints to the target node informing the target to expect unrequested data. The target processor in this case pre-allocates data buffers in anticipation for the pushed data. Doing so prevents the target processor from generating unnecessary requests for data addresses which are going to be pushed along with the lock. If data is pushed directly to the target processor, the target processor receives the actual data rather than hints and allocates cache blocks to sink the data if possible.

When data is directly pushed to the target node, the recipient cannot commit the use of the data until the push has been ordered because the SP is initiated on a network separate from the one used for enforcing serialization (at the snoop bus or directory). Early access may still be beneficial, for example, if data value speculation is being performed in the critical section—data is present in the cache and most probably will be a valid copy.

The choice of the forwarding data path implementation will depend, among others, on the coherence protocol. For example, for snooping systems, pushing data directly to the requestor (instead of via memory) may be beneficial. Modern systems are designed to make snoop bandwidth the performance limiter rather than the data network. Thus, the data network can be utilized to overlap the data transfer with the latency of ordering the SP. On the other hand, for directory systems, address hints may be sufficient since the data has to first go to the directory for ordering.

For the two systems we have simulated, the data is sent directly to the target node for the snoop-based protocol, while it arrives along with the coherence permission in the directory-based scheme (with hints being sent when the lock request is received).

3.3.2 Ordering the speculative push

The precise implementation for correctly ordering (i.e., inserting in the global memory order for correct memory consistency and coherence), the speculative push depend upon the base coherence protocol. We discuss two commercial coherence protocols here and show how a speculative push is ordered.

Broadcast-based system. For a typical bus-based protocol, the pushed data could be broadcast once on the bus, with a special annotation allowing the target node to capture the data as it was being written back to memory. In our example bus system, since data is transmitted point-to-point, a write-back operation requires some care to ensure that the data is received, and ordered correctly, at the target node. The data would normally be pushed immediately after the response providing the lock.

To serialize the push, an annotated write-back (the annotation identifies the target of the speculative push) is sent to memory and is serialized in the global memory order. When the annotated write-back appears on the bus the target receives coherence permissions implicitly. The data may or may not have reached the target node at this point. Since the bus provides an implicit ordering, in many cases, the latency can be overlapped to a larger extent than in a directory system.

Directory-based system. Since the data is also being written to memory to allow the receiving node to ignore it if necessary, this three-way communication presents ordering challenges in a directory-based system. In a directory-based protocol, the directory node must be involved as it is the point of serialization for opera-

tions. Responding to the annotated write-back, the directory node communicates with the target node, granting coherence permission (exclusive) or sending a NACK to the target node if necessary.

As noted in the previous paragraph, the directory node must be involved in all data forwarding to guarantee proper memory semantic. However, this perceived disadvantage has the benefit of solving all the race conditions that might occur. In effect, having to include the directory node in all transactions has the property of not adding any new race condition that the directory-based protocol must not already handle. Thus, receiving an annotated write-back is really no different than receiving data evicted from a node's cache. The only difference is that SP requests the directory to forward a copy of the data to the target, an action that the directory can decline to do if necessary.

For directory systems, latency is not completely overlapped since the coherence permissions need to come via the directory. However, the write-back to the directory is overlapped with the lock transfer to the target node. In addition, we are sending address hints to the target node. Doing so allows the target node to pre-allocate local buffer space for sinking the push. The only exposed latency left is the directory lookup and transfer of coherence permissions to the requestor. In our experiments, while a three-hop read took about 360ns, with the speculative push, the latency observed by the target node substantially reduces to about 60ns. In addition, the upgrade traffic that follows after the read of the data is also eliminated since the data is being sent to the target in exclusive state.

3.3.3 Automatically matching multiply pushed data with coherence permissions

The two actions of speculatively pushing data, and ordering the speculative push in the global memory order can occur in any order and different networks may be used for them. In addition, multiple pushes of a cache line to a given processor may happen since no ordering guarantees are assumed from the network. To easily handle such situations, we treat the two actions symmetrically. If a push is rejected, the corresponding coherence permission must also be rejected and vice-versa. Bookkeeping is necessary to track multiple cache lines to ensure consistent responses to both actions. For generality and without assumptions about the network, we require a requestor to include in a rd_X_lp an indication of the number of lines it can track (but not necessarily sink) at any given time. This number could be quite small.

The push/coherence permission information is stored in a small table at the cache controller. Both messages in the pair will occur exactly once, so every push (irrespective of address) received is tracked until its corresponding coherence permission is received, and vice versa. An entry is removed when the pair is matched up. Any push reject can be matched with any coherence reject and so on. As the mechanism is speculative, the push and coherence permissions may arrive as a result of two different attempted pushes. Nevertheless, it is not difficult to guarantee that the processor will have the latest data if the two actions are matched correctly.

A third processor may attempt to read data in a cache line while the line is being pushed, either due to data races or due to misspeculated pushes. The situation is handled efficiently by providing the third processor with the data and conservatively cancelling the push to the target node (which, for directory systems, may require an additional message to be sent to the target node).

4 EXPERIMENTAL METHODOLOGY

We use an execution-driven simulator to perform cycle-by-cycle simulation of an out-of-order processor and a detailed event driven simulation of the memory hierarchy. The simulator models all data movements accurately (in the pipeline as well as in the memory hierarchy) and models port contention at all levels. The processor implements a release consistency memory model [10] similar to the Compaq Alpha 21264 [8]. The processor retires stores to a coalescing write buffer.

Table 1: Benchmarks

Applications	Application type	Input	Procs
Cholesky [35]	Sparse matrix factorization	tk14.O	16
MP3D [35]	Rarefied fluid flow simulation	24,000 mols, 25 iter.	4
Raytrace [40]	3-D rendering	teapot	8
Water-Nsq [40]	N-body molecular dynamics	512 mols, 3 iter.	16

Not all benchmarks display the behavior of high communication miss rates we are targeting. We specifically select four benchmarks (Table 1) that frequently experience communication misses within critical sections. A locking version of *mp3d* is used to demonstrate the effectiveness for applications with frequent synchronization. The test&test&set synchronization primitive is used and is implemented with the load-locked and store-conditional instructions [16]. We selected system sizes such that contention at the locks is at most medium.

4.1 Target systems

We simulate two systems: symmetric multiprocessor (SMP) and distributed shared-memory (DSM) systems. The SMP system is modeled after the Sun Gigaplane [36]: coherence and data traffic is split onto two separate networks. Address requests and associated coherence operations take place on a high bandwidth snoop bus; while a high-speed point-to-point crossbar transfers data among the nodes. The SMP employs a coherence protocol similar to the one used in the Sun Enterprise 10000 system [7]. The DSM implements a MESI cache coherence protocol similar to the SGI Origin 2000 system [23]. We assume a fully connected, point-to-point network in which the messages take a constant latency to traverse one hop. However, port contention is accurately modeled. Table 2

Table 2: Integrated processor and cache subsystem

Processor	
Processor speed	1 GHz (1ns clock)
Reorder buffer	64 entry with a 32 entry load/store queue
Issue mechanism	out-of-order issue/commit of 4 ops/cycle, 64 entry return address stack, aggressively issue loads (~ MIPS R10000)
Branch predictor	8-K entry combining predictor, 8K entry, 4-way BTB
Cache	
L1 instruction cache	64-KByte, 2-way associative, 1-cycle access, 8 outstanding misses
L1 data cache	128-KByte, 2-way associative, write-back, 2 ports, 1-cycle access, 8 outstanding misses
L2 unified cache	1-MByte, 4-way associative, write-back, 10-cycle access, 16 outstanding misses
L1/L2 bus	Runs at processor clock
Line size	64 bytes

Table 3: External network and memory configuration

DRAM memory module	8-byte wide, ~70ns access time for 64-byte line
Snoop-based configuration	OSI protocol on address bus modeled after the Sun Enterprise 10000, MOESI at snoop cache
Address bus	split-transaction, out-of-order responses, 120 outstanding requests, 22ns snoop cycle (including 2ns arbitration)
Data network	pipelined, point-to-point crossbar, 64-bit wide, 80ns transfer latency
Some uncontended latencies	(pin to pin) read miss to memory: ~172ns, read miss to another cache: ~125ns
Directory-based configuration	SGI Origin-2000 based MESI protocol
Directory access	70ns (overlapped with memory access)
Processor and local directory	30ns (directory is integrated with memory and network controllers, point to point)
Directory and remote router	50ns (point to point)
Some uncontended latencies	(pin to pin) read miss to local memory: ~130ns, read miss to remote memory: ~230ns, read miss to remote dirty cache: ~360ns
Network configuration (DSM only)	pipelined point-to-point network
Network width	64 bits
Read latency from network to cache	1ns per word
Setup latency for header packet	5ns
Setup latency for data packet	5ns + 1ns per word

list the parameters of the integrated processor and cache subsystem used in both SMP and DSM systems and Table 3 lists the parameters for the SMP and DSM memory systems.

4.2 Explanation of metrics

Speedup is measured as the ratio of the parallel execution time of the base case to the execution time of the optimized case. Therefore, a speedup greater than one implies a performance gain. Attributing stall cycles to specific components is a complex task for multiprocessor systems with out-of-order processors where many events occur concurrently. We use an approximate approach. For every cycle, we compute the ratio of instructions committed that cycle to the maximum commit rate. This fraction of cycle time is attributed to the busy time for the processor. The remaining fraction is attributed as stall time to the first instruction that could not be committed that cycle. The fractions of stall cycles are normalized to the running time of IQL-only case. The stall categories are:

- *WMB*: stall at write-memory barrier
- *MB*: stall at memory barrier
- *C SECTION*: stall associated with data accesses within a critical section
- *MEM ACC*: stall associated with shared-memory accesses outside critical sections
- *CPU*: the remainder of useful cycles while the processor is busy computing

The shared memory accesses do not include lock variable accesses. This is done specifically to avoid any bias in latency accounting due to differences in synchronization primitives.

5 RESULTS

5.1 Performance

Table 4 presents the main results. The numbers in parentheses show the running time in millions of processor cycles. All the other numbers in the table represent speedups. SP does not hurt performance for most of our benchmarks. The exception is for `cholesky` running on an DSM system where its performance drop by only 1%. SP helps `mp3d` and `raytrace` substantially compared to the IQL-only case, speeding up their performance by a factor of up to 1.51. We also have the results for the SMP system (not shown). For the most part, the trends are identical but the magnitudes are different. In particular, SP speeds up `mp3d` by 21%, but slows down `cholesky` by 2%. The exception is `raytrace`, which displays an improvement of only 3%. The latencies in the SMP system are such that the locking behavior of `raytrace` does not affect its performance much. `Water-nsq` communicates comparatively far more infrequently than the other benchmarks do (see Figure 6 and Figure 7), but nevertheless still suffers some migratory data-related stalls and benefits, if little, from SP.

Table 4 also compares the performance of one system implementing IQL only and the other one without it (i.e., programs use test&test&set built with load-linked and store-conditional instructions but otherwise run without hardware support). To our knowledge this paper is the first one to describe and to simulate a DSM system with IQL. The results are impressive; IQL improves the performance of all our benchmarks and, in particular, reduces the running time of `raytrace` by more than half. These results follow a trend similar to numbers published elsewhere [18, 17]. Besides differences in the systems being simulated, our methodologies differ also in that they used QOLB as their mechanism, which requires special instructions and the recompilation of the programs.

5.2 Speculative Push stall cycles breakdown

Figure 6 and Figure 7 shows how the stall cycles are attributed to the different components of the system. These figures show that most of the performance gains stem from the reduction of the following two components: C SECTION and WMB. SP is able to

Table 4: IQL and SP performance for a DSM system

	Cholesky	MP3D	Raytrace	Water-Nsq
Without IQL (base)	(11)	(373)	(121)	(18)
IQL	1.13	1.21	2.75	1.04
IQL+SP	1.12	1.60	4.15	1.11

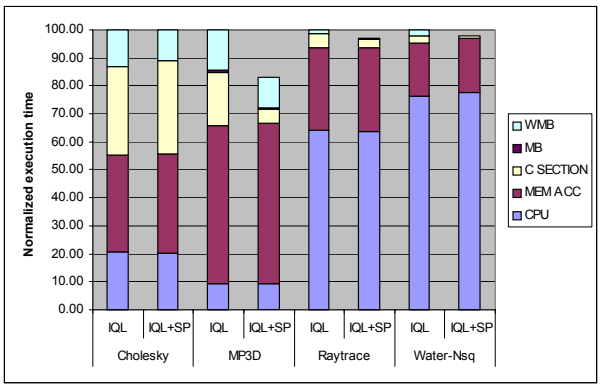

Figure 6. SMP stall contributions

Normalized stall contributions are expressed as percentage of the IQL running time for the specific contribution.

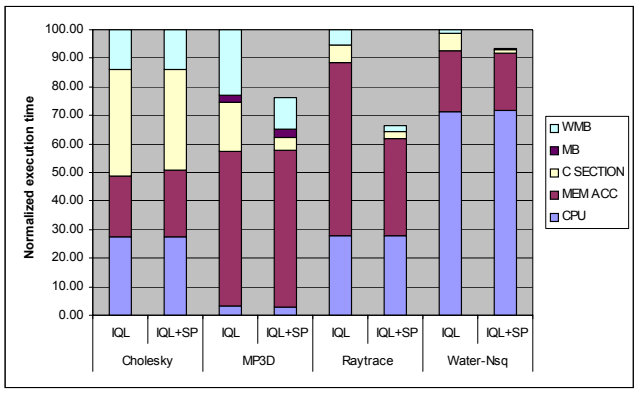

Figure 7. DSM stall contributions

Normalized stall contributions are expressed as percentage of the IQL running time for the specific contribution.

eliminate nearly all latencies associated with loading shared-memory locations. This observed behavior reduces the stall cycles associated with critical section execution (C SECTION). SP is also able to reduce considerably the wait time at write memory barriers (WMB) located at the end of each critical section. In our experiments, write memory barriers are almost always strictly confined to the end of each critical section and are there to ensure all the memory operations effected inside the critical section have performed globally before releasing the lock.

5.3 Speculative Push characteristics

Table 5 breaks down speculative pushes into four categories: (1) Used: the pushed data was accessed, (2) Evicted: the pushed data was evicted before being used, (3) Rejected: the pushed data was rejected, and (4) Invalidated: the pushed data was invalidated by another processor before being used. For the most part SP performs relatively well: more than 70% of all pushes are useful for the execution of `mp3d`, `raytrace`, and `water-nsq`. The exception to this trend being `cholesky`. `Cholesky` is the control benchmark and its critical section behavior does not lend itself to the type of optimization we are studying in this paper. Indeed, `cholesky` rarely suffers write-faults on the same address on successive executions of the same critical section. For this benchmark, little correlation exists between a lock address and data addresses accessed while the lock is held. Our predictor detects this patterns and turns off the SP optimization. This behavior leads to little performance improvement for this benchmark (Table 4) and many evicted pushes (Table 5).

We also varied the number of cache lines the SP mechanism could forward with a lock. For our benchmarks, we found most of the benefits could be had with a single cache line. Adding a second cache line improved the performance of SP by no more than 2% and more than two cache lines lead to insignificant improvements.

5.4 Speculative Push versus flush

We compare the performance of SP against a technique that consists of flushing data back to memory at the end of a critical section. The basic idea is to avoid the penalty of accessing remote dirty data and instead to attempt finding the desired data at memory directly. To implement the flushing mechanism, we rely again on our predictors to identify critical section boundaries and data set associated with a lock. Upon receiving a lock request, we write the data back to memory only unless the lock is still held. If the lock is still held, we wait until the lock is released to write the data back to memory.

We ran experiments both with the SMP and DSM configurations. We found (numbers are not shown here) the flush technique to be ineffective in our simulated SMP system because our assumptions are such that the transfer latency to and from memory is much larger than the transfer latency between two caches. Thus flush techniques and self-invalidation techniques will degrade performance of SMP systems.

Table 6 contrasts the performance of our two DSM variants: flush to memory and SP. We observe that, for our set of benchmarks, flush achieves a smaller speedup compared to SP. These results also show that flush is able to achieve performance gains over the IQL-only case. However, flush either requires special instructions and recompilation or requires a predictor, a unit that it shares with the SP mechanism.

Table 5: Breakdown of push characteristics (shown as percentage of pushes)

	used		evicted		rejected		invalidated	
	SMP	DSM	SMP	DSM	SMP	DSM	SMP	DSM
Cholesky	8.00%	12.02%	92.00%	88.98%	0.00%	0.00%	0.00%	0.00%
MP3D	99.46%	99.56%	0.43%	0.40%	0.09%	0.04%	0.02%	0.00%
Raytrace	80.80%	71.64%	16.55%	28.33%	2.64%	0.00%	0.01%	0.03%
Water	93.80%	99.61%	6.20%	0.08%	0.00%	0.31%	0.00%	0.00%

Table 6: SP and flush performance for DSM

	Cholesky	MP3D	Raytrace	Water-Nsq
IQL+Flush	1.00	1.13	1.28	1.04
IQL+SP	0.99	1.32	1.51	1.07

6 RELATED WORK

A software technique for eliminating latencies associated with critical data accesses involves collocating lock and data in the same cache line. However, since locks may be read even within critical sections, it is not generally recommended that locks be allocated in the same line as data they are protecting. If a mechanism is provided to eliminate or defer accesses to the lock until the end of a critical section, then data can be collocated with a lock profitably, allowing the data to be implicitly transferred along with the lock when it is acquired. Bitar and Despain first proposed collocation [5]. Goodman, Vernon, and Woest [12] made collocation more attractive by establishing the ability to defer access to the lock by an acquiring processor until the lock had been released. This method, Queue-On-Lock-Bit (QOLB) was a synchronizing prefetch operation in the sense that it provided for lock and data to be forwarded "as soon as possible, but no sooner." Using instruction set and programmer support, QOLB maintained a queue of lock requestors in hardware. Kägi, Burger, and Goodman [18] demonstrated that collocation captured consistent and substantial gains in performance for a set of benchmarks on a distributed shared memory system. Collocation however requires substantial programmer involvement and at times, major restructuring of the application data structures. In addition, coupling the lock and data in the same cache line limits the size of the collocated data. QOLB led to other proposals for queued locks, notably MCS [28] and for the DASH multiprocessor [25].

DASH provided a concept of queued locks in hardware for memory based directories. However, the directory was always in the critical path—on a lock release, the lock was sent to the directory which in turn picked a random waiter and serviced it. With IQLs, once a request has been forwarded, the directory is no longer in the critical path.

We have previously proposed Implicit-QOLB [31], an early version of the IQL mechanism, which works by speculating about a program's access patterns—specifically of synchronization operations—and uses the notion of delayed responses to improve the throughput of synchronization. This work focused only on bus-based systems and did not address communication latencies within critical sections. In the present work, which first appeared as an unpublished manuscript in April 2000 [32] we demonstrate that the method is even more effective on a directory-based protocol, and that the SP mechanism can leverage the notion of IQL to achieve still larger performance gains.

Stenström, Brorsson, and Sandberg [38] and Cox and Fowler [9] independently proposed cache coherence protocol optimizations for migratory sharing patterns. Such behavior is exhibited primarily by data protected by locks or monitors. Both approaches succeed by merging an invalidation request for the migratory cache block with the preceding read-miss request. These mechanisms do not reduce the critical miss latency experienced on the first read miss, though reduced contention may have the indirect effect of reducing read miss latencies.

Mowry and Gupta [29] proposed a compiler prefetch heuristic for tolerating latency in shared-memory multiprocessors. The compiler interpreted explicit synchronization operations as a hint that data communication may be taking place. The approach was quite successful for programs with regular access patterns and structures. They mention that it is potentially easy for a programmer to use semantic information about an application and identify critical data structures in small applications but state, "[s]uch focusing in on critical data structures will be much harder for compilers." An additional issue with software prefetching for critical section data is the lack of knowledge regarding the migratory patterns of data: determining which processor should be prefetching data is nearly impossible statically because it depends on the (dynamic) selection of a winner among competitors to acquire the lock. By the time this decision has been made, it is already too late to avoid delay by prefetching needed data. Trancoso and Torrellas [39] attempted to reduce latencies within critical sections through the use of prefetching and data forwarding. They inserted prefetch and forwarding instructions by hand. Their techniques suffer from many of the same limitations of software approaches, specifically, the need for hardware and compiler support for new instructions and the inability to evaluate and exploit run-time behavior. Their results were pessimistic, concluding that complex, forwarding-based optimizations could not be justified.

Abdel-Shafi et al. [1] evaluated producer-initiated communication and proposed remote writes for data accesses associated with synchronization operations. The combination of software prefetching and remote writes provided good performance gains for a set of benchmarks. The mechanisms however, required software and programmer support to identify candidates for remote writes.

Similar data forwarding mechanisms have been proposed in the literature: the *forwarding write* [30], and the DASH *deliver* [25]. DASH also had a producer-prefetch mechanism for pushing data to a set of consumers in shared state. Kaxiras and Goodman [19] proposed speculative pre-send as an approach for data forwarding.

Ranganathan et al. [33] proposed the use of flush primitives to write back dirty data modified in critical sections to memory. They also added prefetches at the beginning of critical sections. Their mechanisms relied on compiler and programmer support to identify critical data to be flushed. However they state that late prefetches and contention effects limited additional performance benefits. Similar flush primitives have also been proposed by Hill et al. [14] and Skeppstedt and Stenström [37]. Mechanisms to reduce invalidation latencies by employing prediction to flush cache lines have also been proposed [24, 22]. These techniques reduce a three-hop transaction to a two-hop transaction while Speculative Push converts a three-hop transaction into a local cache access.

7 CONCLUDING REMARKS

In this paper, we have studied two mechanisms for reducing communication latencies inside critical sections. First, we proposed the first IQL implementation for a directory-based system that makes no assumptions about message ordering. Our results corroborate earlier published results for snoop-based systems running on a system with high network utilization and showing large gains for benchmarks that have either high lock contention or frequent, fine-grained locking. Second, we described SP, a mechanism to overlap lock transfer with data believed to be associated with that lock, thus attempting to convert all global data accesses performed in a critical section into local cache accesses. We showed that SP offers additional benefits on top of those provided by IQLs.

`Mp3d` was chosen as a benchmark specifically because it exhibits the kind of behavior we were targeting, and both mechanisms succeeded in reducing communication delays. The net result was

that the application saw a speedup of 21% for the bus and 32% for the directory system over an aggressive base case of IQLs. Indeed, all the benchmarks saw reductions in shared memory stalls within critical sections, though for some this delay was so small that the reduction did little to improve overall performance. Benchmarks with highly contended locks (such as `raytrace`), show large speedups in some cases. The SP mechanism can provide further speedups in overall performance by substantially reducing the network traffic and three-hop transactions.

We conclude that the two mechanisms can combine to reduce the communication delays within critical sections by more than 50%. In addition, speculative push can quite often collapse the read-modify-write sequences within a critical section into a local cache access. While the total reduction in stalls varies depending on the percentage of time the processor is stalled for communication latencies, the reduction was consistent across all benchmarks.

ACKNOWLEDGEMENTS

This work was supported by NSF Grant CCR-9810114.

REFERENCES

[1] H. Abdel-Shafi, J. Hall, S. V. Adve, and V. S. Adve. An evaluation of fine-grain producer-initiated communication in cache-coherent multiprocessors. In *Proceedings of the Third International Symposium on High-Performance Computer Architecture*, pages 204–215, Feb. 1997.

[2] T. E. Anderson. The performance implications of spin-waiting alternatives for shared-memory multiprocessors. In *Proceedings of the 1989 International Conference on Parallel Processing*, volume II (software), pages 170–174, Aug. 1989.

[3] T. E. Anderson. The performance of spin lock alternatives for shared-memory multiprocessors. *IEEE Transactions on Parallel and Distributed Systems*, 1(1):6–16, Jan. 1990.

[4] L. A. Barroso, K. Gharachorloo, and E. Bugnion. Memory system characterization of commercial workloads. In *Proceedings of the 25th Annual International Symposium on Computer Architecture*, pages 3–14, June 1998.

[5] P. Bitar and A. M. Despain. Multiprocessor cache synchronization: Issues, innovations, evolution. In *Proceedings of the 13th Annual International Symposium on Computer Architecture*, pages 424–433, June 1986.

[6] T. Brewer and G. Astfalk. The evolution of the HP/Convex Exemplar. In *Proceedings of the 42nd IEEE Computer Society International Conference (COMPCON)*, pages 81–86, Feb. 1997.

[7] A. Charlesworth, A. Phelps, R. Williams, and G. Gilbert. Gigaplane-XB: Extending the Ultra Enterprise family. In *Proceedings of the Symposium on High Performance Interconnects V*, pages 97–112, Aug. 1997.

[8] Compaq Computer Corporation. *Alpha 21264 Hardware Reference Manual*, July 1999.

[9] A. L. Cox and R. J. Fowler. Adaptive cache coherency for detecting migratory shared data. In *Proceedings of the 20th Annual International Symposium on Computer Architecture*, pages 98–108, May 1993.

[10] K. Gharachorloo, D. Lenoski, J. Laudon, P. B. Gibbons, A. Gupta, and J. L. Hennessy. Memory consistency and event ordering in scalable shared-memory multiprocessors. In *Proceedings of the 17th Annual International Symposium on Computer Architecture*, pages 15–26, May 1990.

[11] K. Gharachorloo, M. Sharma, S. Steely, and S. V. Doren. Architecture and design of AlphaServer GS320. In *Proceedings of the Ninth Symposium on Architectural Support for Programming Languages and Operating Systems*, pages 13–24, Nov. 2000.

[12] J. R. Goodman, M. K. Vernon, and P. J. Woest. Efficient synchronization primitives for large-scale cache-coherent shared-memory multiprocessors. In *Proceedings of the Third Symposium on Architectural Support for Programming Languages and Operating Systems*, pages 64–75, Apr. 1989.

[13] G. Graunke and S. Thakkar. Synchronization algorithms for shared-memory multiprocessors. *IEEE Computer*, 23(6):60–69, June 1990.

[14] M. Hill, J. Larus, S. Reinhardt, and D. Wood. Cooperative shared memory: Software and hardware for scalable multiprocessors. *ACM Transactions on Computer Systems*, 11(4):300–318, Nov. 1993.

[15] Institute of Electrical and Electronics Engineers, New York, NY. *IEEE Standard for the Scalable Coherent Interface (SCI)*, Aug. 1993. ANSI/IEEE Std. 1596-1992.

[16] E. H. Jensen, G. W. Hagensen, and J. M. Broughton. A new approach to exclusive data access in shared memory multiprocessors. Technical Report UCRL-97663, Lawrence Livermore National Laboratory, Livermore, CA, Nov. 1987.

[17] A. Kägi. *Mechanisms for Efficient Shared-Memory, Lock-Based Synchronization*. PhD thesis, University of Wisconsin, Madison, WI, May 1999.

[18] A. Kägi, D. Burger, and J. R. Goodman. Efficient synchronization: Let them eat QOLB. In *Proceedings of the 24th Annual International Symposium on Computer Architecture*, pages 170–180, June 1997.

[19] S. Kaxiras and J. R. Goodman. Improving CC-NUMA performance using instruction-based prediction. In *Proceedings of the Fifth International Symposium on High-Performance Computer Architecture*, pages 161–170, Jan. 1999.

[20] K. Keeton, D. Patterson, Y. He, R. Raphael, and W. Baker. Performance characterization of a Quad Pentium Pro SMP using OLTP workloads. In *Proceedings of the 25th Annual International Symposium on Computer Architecture*, pages 15–26, June 1998.

[21] D. Kroft. Lockup-free instruction fetch/prefetch cache organization. In *Proceedings of the Eighth Annual International Symposium on Computer Architecture*, pages 81–87, May 1981.

[22] A.-C. Lai and B. Falsafi. Selective, accurate, and timely self-invalidation using last-touch prediction. In *Proceedings of the 27th Annual International Symposium on Computer Architecture*, June 2000.

[23] J. Laudon and D. E. Lenoski. The SGI Origin: A ccNUMA highly scalable server. In *Proceedings of the 24th Annual International Symposium on Computer Architecture*, pages 241–251, June 1997.

[24] A. R. Lebeck and D. A. Wood. Dynamic self-invalidation: Reducing coherence overhead in shared-memory multiprocessors. In *Proceedings of the 22nd Annual International Symposium on Computer Architecture*, pages 48–59, June 1995.

[25] D. Lenoski, J. Laudon, K. Gharachorloo, W.-D. Weber, A. Gupta, J. L. Hennessy, M. Horowitz, and M. Lam. The Stanford DASH multiprocessor. *IEEE Computer*, 25(3):63–79, Mar. 1992.

[26] B.-H. Lim and A. Agarwal. Reactive synchronization algorithms for multiprocessors. In *Proceedings of the Sixth Symposium on Architectural Support for Programming Languages and Operating Systems*, pages 25–35, Oct. 1994.

[27] T. Lovett and R. Clapp. STiNG: A CC-NUMA computer system for the commercial marketplace. In *Proceedings of the 23rd Annual International Symposium on Computer Architecture*, pages 308–317, May 1996.

[28] J. M. Mellor-Crummey and M. L. Scott. Algorithms for scalable synchronization on shared-memory multiprocessors. *ACM Transactions on Computer Systems*, 9(1):21–65, Feb. 1991.

[29] T. Mowry and A. Gupta. Tolerating latency through software-controlled prefetching in shared-memory multiprocessors. *Journal of Parallel and Distributed Computing*, 12(2):87–106, June 1992.

[30] D. K. Poulsen and P.-C. Yew. Data prefetching and data forwarding in shared memory multiprocessors. In *Proceedings of the 1994 International Conference on Parallel Processing*, volume II (software), pages 276–280, Aug. 1994.

[31] R. Rajwar, A. Kägi, and J. R. Goodman. Improving the throughput of synchronization by insertion of delays. In *Proceedings of the Sixth International Symposium on High-Performance Computer Architecture*, pages 168–179, Jan. 2000.

[32] R. Rajwar, A. Kägi, and J. R. Goodman. Using speculative push to reduce communication latencies in critical sections. Unpublished manuscript, Apr. 2000 (Now available as Technical Report CS-TR-1472, Computer Sciences Department, University of Wisconsin, Madison, WI, Mar. 2003).

[33] P. Ranganathan, K. Gharachorloo, S. Adve, and L. A. Barroso. Performance of database workloads on shared-memory systems with out-of-order processors. In *Proceedings of the Eighth Symposium on Architectural Support for Programming Languages and Operating Systems*, pages 307–318, Oct. 1998.

[34] L. Rudolph and Z. Segall. Dynamic decentralized cache schemes for MIMD parallel processors. In *Proceedings of the 11th Annual International Symposium on Computer Architecture*, pages 340–347, June 1984.

[35] J. P. Singh, W.-D. Weber, and A. Gupta. SPLASH: Stanford parallel applications for shared memory. *Computer Architecture News*, 20(1):5–44, Mar. 1992.

[36] A. Singhal, D. Broniarczyk, F. M. Cerauskis, J. Price, L. Yuan, G. Cheng, D. Doblar, S. Fosth, N. Agarwal, K. Harvey, E. Hagersten, and B. Liencres. Gigaplane: A high performance bus for large SMPs. In *Proceedings of the Symposium on High Performance Interconnects IV*, pages 41–52, Aug. 1996.

[37] J. Skeppstedt and P. Stenström. A compiler algorithm that reduces read latency in ownership-based cache coherence protocols. In *Proceedings of the 1995 International Conference on Parallel Architectures and Compilation Techniques*, 1995.

[38] P. Stenström, M. Brorsson, and L. Sandberg. An adaptive cache coherence protocol optimized for migratory sharing. In *Proceedings of the 20th Annual International Symposium on Computer Architecture*, pages 109–118, May 1993.

[39] P. Trancoso and J. Torrellas. The impact of speeding up critical sections with data prefetching and forwarding. In *Proceedings of the 1996 International Conference on Parallel Processing*, volume III (software), pages 79–86, Aug. 1996.

[40] 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 22nd Annual International Symposium on Computer Architecture*, pages 24–36, June 1995.