

# Speculative Versioning Cache

T.N. Vijaykumar, *Member, IEEE*, Sridhar Gopal,  
James E. Smith, *Member, IEEE*, and Gurindar Sohi, *Senior Member, IEEE*

**Abstract**—Dependences among loads and stores whose addresses are unknown hinder the extraction of instruction level parallelism during the execution of a sequential program. Such ambiguous memory dependences can be overcome by memory dependence speculation which enables a load or store to be speculatively executed before the addresses of all preceding loads and stores are known. Furthermore, multiple speculative stores to a memory location create multiple speculative versions of the location. Program order among the speculative versions must be tracked to maintain sequential semantics. A previously proposed approach, the Address Resolution Buffer (ARB) uses a centralized buffer to support speculative versions. Our proposal, called the Speculative Versioning Cache (SVC), uses distributed caches to eliminate the latency and bandwidth problems of the ARB. The SVC conceptually unifies cache coherence and speculative versioning by using an organization similar to snooping bus-based coherent caches. Our evaluation for the Multiscalar architecture shows that hit latency is an important factor affecting performance and private cache solutions trade-off hit rate for hit latency.

**Index Terms**—Speculative memory, memory disambiguation, snooping cache coherence protocols, speculative versioning.

## 1 INTRODUCTION

MODERN microprocessors extract instruction level parallelism (ILP) from sequential programs by issuing instructions from an active instruction window. Data dependences among instructions and not the original program order, determine when an instruction may be issued from the window. Dependences involving register data are detected easily because register designators are completely specified within instructions. However, dependences involving memory data (e.g., between a load and a store or two stores) are ambiguous until the memory addresses are computed.

A straightforward solution to the problem of ambiguous memory dependences is to issue loads and stores only after their addresses are determined. Furthermore, a store is not allowed to complete and commit its result to memory until all preceding instructions are known to be free of exceptions. Each such store to a memory location creates a *speculative version* of that location. These speculative versions are held in buffers until they can be committed. Multiple speculative stores to the same location create multiple versions of the location. To improve performance, loads are allowed to bypass buffered stores, as long as they are to different addresses. If a load is to the same address as a buffered store, it can use data bypassed from the store

when the data becomes available. An important constraint of this approach is that a load instruction *cannot* be issued until the addresses of all the preceding stores are determined. This restriction may diminish ILP unnecessarily, especially in the common case where the load is not dependent on preceding stores.

More aggressive uniprocessor implementations issue load instructions as soon as their addresses are known, even if the addresses of all previous stores may not be known. These implementations employ memory dependence speculation [11] and predict that a load does not depend on previous stores. Furthermore, one can also envision issuing and computing store addresses out of order. Such memory dependence speculation enables higher levels of ILP, but more advanced mechanisms are needed to support this speculation and to recover from the misspeculations that inevitably occur. And it means that a significant amount of store data must be buffered and committed in correct program order.

Prior to being committed, each store to a memory location creates a new speculative *version* of that memory location. Multiple speculative stores to a memory location create multiple speculative versions of the location. Program order between the multiple versions of a location has to be tracked to: 1) commit versions in program order and 2) supply (bypass) the correct version as per program order to loads. We call this problem *speculative versioning*. We illustrate the issues in speculative versioning with an example program containing two loads and two stores to the same memory location, *A*.

- load  $R_1, A$
- store 2,  $A$
- load  $R_2, A$
- store 3,  $A$

Sequential program semantics dictate that,

- T.N. Vijaykumar is with the School of Electrical and Computer Engineering, Purdue University, 1285 EE Building, West Lafayette, IN 47907-1285. E-mail: vijay@ecn.purdue.edu.
- S. Gopal is with the Network Operations Software Group of Lucent Technologies, Inc.
- J.E. Smith is with Department of Electrical and Computer Engineering, University of Wisconsin, 4621 Engineering Hall, 1415 Engineering Dr., Madison, WI 53706. E-mail: jes@ece.wisc.edu.
- G. Sohi is with the Department of Computer Sciences University of Wisconsin, 1210 West Dayton St., Madison, WI 53706. E-mail: sohi@cs.wisc.edu.

Manuscript received 13 Feb. 1998; revised 5 June 2001; accepted 6 June 2001. For information on obtaining reprints of this article, please send e-mail to: tpd@computer.org, and reference IEEECS Log Number 112306.

- the load to  $R_1$  should *not* read the value 2 created by the store even if it is executed *after* the store. This requires that the store must not overwrite memory location  $A$  until the load has executed and, hence, must be buffered.
- The load to  $R_2$  should *eventually* read the value 2 created by the store. This requires that the load must be squashed and reexecuted if it executes before the store and incorrectly reads the version stored in memory.
- Memory location  $A$  should *eventually* have the value 3 stored independent of the order of execution of the stores. This requires that multiple stores to the same location have to be buffered before they are committed.
- Program order between the loads and stores has to be tracked to guarantee the above three conditions.

Most modern microprocessors dispatch instructions from a single instruction stream and issue load and store instructions from a common set of hardware buffers (e.g., reservation stations). This allows the hardware to maintain a time-ordering of loads and stores via simple queue mechanisms, coupled with address comparison logic. The presence of store queues provides a simple form of speculative versioning. However, proposed next generation processor designs use replicated processing units that dispatch and/or issue instructions in a distributed manner. These future approaches partition the instruction stream into tasks [4], [15] or traces [14], [18], and use higher level instruction dispatch units to assign the tasks or traces to the processing units for execution. These approaches are aimed at higher levels of ILP and use replicated processing units with hierarchical control for a number of practical engineering reasons [13]. Proposed next generation multiprocessors [12], [9], [17], [16], [2], [19] that provide hardware support for dependence speculation use such hierarchical organization of the multiple processing units.

These highly parallel, hierarchical organizations naturally lead to memory address streams with a similar hierarchical structure. In particular, each individual task or trace generates its own address stream, which can be properly ordered (disambiguated) within the processing unit that generates it and at the higher level, the multiple address streams produced by the processing units must also be properly ordered. It is more challenging to support speculative versioning for this execution model than a superscalar execution model because a processing unit executes loads and stores *without knowing* those executed by the other processing units. Consequently, there is a need for speculative versioning mechanisms that operate within this hierarchical organization.

The Address Resolution Buffer [5] (ARB) provides speculative versioning support for such hierarchical execution models. Each entry in the ARB buffers all versions of the same memory location. However, there are two significant performance limitations of the ARB:

1. The ARB is a single shared buffer connected to the multiple processing units and, hence, every load and store incurs the latency of the interconnection network. Also, the ARB may become a bandwidth

bottleneck since every load and store from the multiple processing units.

2. When a task completes all its instructions, the ARB commits its speculative state into the architectural storage (or copies all the versions created by this task to the data cache). Such write backs generate bursty traffic and may increase the time to commit a task, which delays the issue of new task to that processing unit, limiting the overall performance.

We propose a new solution for speculative versioning called the Speculative Versioning Cache [8] (SVC), for hierarchical execution models. The SVC provides a private cache for each processing unit and the system is organized similar to a snooping bus-based cache coherent Symmetric Multiprocessor (SMP). Memory references that hit in the private cache do not use the bus as in an SMP. Task commits do *not* write back speculative versions *en masse*. Instead, versions are marked, *en masse*, as committed at commit time, without performing any data movement. Each cache line is individually handled when it is accessed the next time.

By employing a distributed organization where each processing unit has its own private cache like an SMP, the SVC conceptually unifies hardware cache coherence and speculative versioning under a single set of mechanisms.

Section 2 introduces the hierarchical execution model briefly and identifies the issues in providing support for speculative versioning for such execution models. Section 3 presents the SVC as a progression of designs to ease understanding. Section 4 gives a performance evaluation of the SVC to highlight the importance of a private cache solution for speculative versioning. We derive conclusions in Section 5.

## 2 SPECULATIVE VERSIONING

First, we discuss the issues involved in providing support for speculative versioning for current generation processors. Second, we describe the hierarchical execution model used by the proposed next generation processors. Third, we discuss the issues in providing support for speculative versioning for this execution model and use examples to illustrate them. Finally, we present similarities between multiprocessor cache coherence and speculative versioning for the hierarchical execution model and use this unification to motivate our new design, the speculative versioning cache.

Speculative versioning involves tracking the program order among the multiple buffered versions of a location to guarantee the following sequential program semantics:

- A load must *eventually* read the value created by the most recent store to the same location. This requires that 1) the load must be squashed and reexecuted if it executes before the store and incorrectly reads the previous version and, 2) all stores (to the same location) that follow the load in program order must be buffered until the load is executed.
- A memory location must *eventually* have the correct version independent of the order of the creation of the versions. Consequently, the speculative versions




Fig. 1. Task commits and squashes: example.

of a location must be committed to the architected storage in program order.

### 2.1 Hierarchical Execution Model

In this execution model, the dynamic instruction stream of a program is partitioned into fragments called *tasks*. These tasks form a *sequence* corresponding to their order in the dynamic instruction stream. A higher level control unit *predicts* the next task in the sequence and assigns it for execution to a free processor. Each processor executes the instructions in the task assigned to it and buffers the speculative state created by the task. The Wisconsin Multiscalar [15] is an example architecture that uses the hierarchical execution model.

When a task misprediction is detected, the speculative state of all the tasks in the sequence including and after the incorrectly predicted task are invalidated<sup>1</sup> and the corresponding processors are freed. This is called a *task squash*. The correct tasks in the sequence are then assigned for execution. When a task prediction has been validated, it *commits* by copying the speculative buffered state to the architected storage. Tasks commit one by one in the order of the sequence. Once a task commits, its processor is free to execute a new task. Since the tasks commit in program order, tasks are assigned to the processors in program order.

Fig. 1 illustrates task commits and task squashes. Initially, tasks 0, 1, 99, and 3 are predicted and speculatively executed in parallel by the four processors, as shown in Fig. 1a. When the misprediction of task 99 is detected, tasks 99 and 3 are squashed and their buffered states are invalidated. New tasks 2 and 3 are then executed by the processors, as shown in Fig. 1b. Tasks that are currently executing are said to be *active*. When task 0 completes execution, the corresponding processor is freed and task 4 is assigned for execution, as shown in Fig. 1c. The program order, represented by the sequence among the tasks, enforces an *implicit* total order among the processors; the arrows show this order. When the tasks are speculatively executed in parallel, the multiple speculative load/store streams from the processors are merged in arbitrary order. Providing support for speculative versioning for such execution models requires mechanisms that establish program order among these streams. The following sections outline how the order is established using the sequence among the tasks.

1. An alternative model for recovery invalidates only the dependent chains of instructions by maintaining information at a finer granularity. This paper assumes the simpler model.

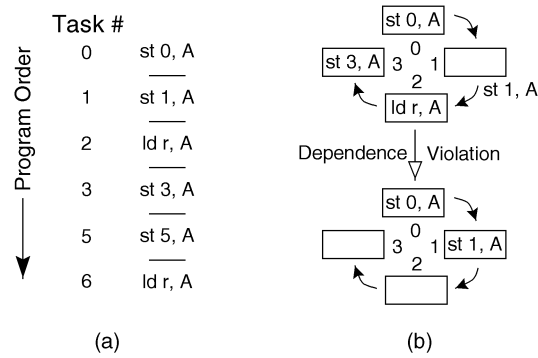


Fig. 2. Speculative versioning example. (a) Example program. (b) Sample execution.

#### 2.1.1 Loads

A task executes a load as soon as its address is available, speculating that stores from previous tasks in the sequence do not write to the same location. The closest previous version of the location is supplied to the load; this version could have been created either by the same task or by a previous task. A load that is supplied a version from a previous task is recorded to indicate a use before a potential definition. If such a definition (a store to the same location from a previous task) occurs, the load was supplied with an incorrect version and memory dependence was violated.

#### 2.1.2 Stores

When a task executes a store to a memory location, it is communicated to all later active tasks in the sequence.<sup>2</sup> When a task receives a new version of a location from a previous task, it squashes if a use before definition is recorded for that location—a memory dependence violation is detected. All tasks after the squashed task are also squashed as on a task misprediction (simple squash model).

#### 2.1.3 Task Commits and Squashes

The oldest active task is nonspeculative and can commit its speculative memory state (versions created by stores from this task) to architected storage. Committing a version involves logically copying the versions from the speculative buffers to the architected storage (data cache). As we assume the simple task squash model, the speculative state associated with a task is invalidated when it is squashed.

### 2.2 Examples for Speculative Versioning

Fig. 2 illustrates the issues involved in speculative versioning using an example program and a sample execution of the program on a four processor hierarchical system. We use the same example in the later sections to explain the SVC design. Fig. 2a shows the loads and stores in the example program and the task partitioning. Other instructions are not of direct relevance here. Fig. 2b shows two snapshots of the memory system during a sample execution of the program. Each snapshot contains four boxes, one for each active task and shows the load or store that has been executed by the corresponding task. The program order among the instructions translates to a sequence among the

2. In reality, the store has to be communicated only until the task that has created the next version, if any, of the location.

TABLE 1  
Versioning: Events and Actions

Event	Actions
Load	Record use before definition by the task; supply the closest previous version.
Store	Communicate store to later tasks; later tasks look for memory dependence violations.
Commit	Write back buffered versions created by the task to main memory.
Squash	Invalidate buffered versions created by the task.

tasks which imposes a total order among the processors executing them; *solid* arrowheads show the program order and *hollow* arrowheads show the execution time order in all the examples.

The first snapshot is taken just before task 1 executes a store to address *A*. Tasks 0 and 3 have already stored 0 and 3 to *A* and task 2 has executed a load to *A*. The load is supplied the version created and buffered by task 0. But, according to the original program, this load *must* be supplied the value 1 created by the store from task 1, i.e., the store to load dependence has been violated. This violation is detected when task 1 stores to address *A* and all the tasks including and after task 2 are squashed and reexecuted. The second snapshot is taken after the tasks have been squashed and restarted.

### 2.3 Coherence and Speculative Versioning

The actions performed on memory accesses and task commits and squashes are summarized in Table 1. The functionality in this table requires the hardware to track the active tasks or processors that executed a load/store to a location and the order among the different copies/versions of this location. Cache coherent Symmetric MultiProcessors (SMP) use similar functionality to track the caches that have a copy of every memory location. SMPs, however, need not track the order among these copies since all the copies are of a single version.

SMPs typically use snooping [6] to implement a Multiple Reader/Single Writer protocol, which uses a coherence directory that is a collection of sets, each of which tracks the sharers of a line. In a snooping bus-based SMP, the directory is typically implemented in a distributed fashion comprising state bits associated with each cache line. On the other hand, the Speculative Versioning Cache (SVC) implements a Multiple Reader/Multiple Writer protocol that tracks copies of multiple speculative versions of each memory location. This protocol uses a version directory that maintains *ordered* sets for each line, each of which tracks the program order among the multiple speculative versions of a line. This ordered set or list, called the Version Ordering List (VOL), can be implemented in several different ways—the SVC, proposed in this paper, uses explicit pointers in each line to implement it as a linked list (like in SCI [1]). The following sections elaborate on a design that uses pointers in each cache line to maintain the VOL.

The *private* cache organization makes the SVC a feasible memory system for proposed next generation single chip

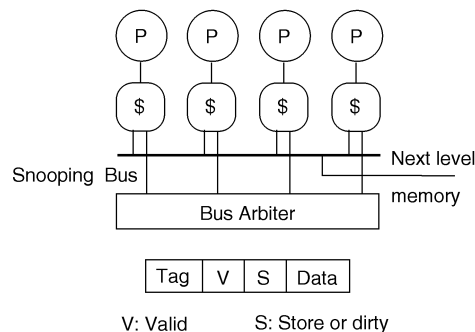


Fig. 3. SMP coherent cache.

multiprocessors that execute sequential programs on tightly coupled processors using automatic parallelization [12], [17]. Previously, ambiguous memory dependences limited the range of programs chosen for automatic parallelization. The SVC provides hardware support to overcome ambiguous memory dependences and enables more aggressive automatic parallelization of sequential programs.

## 3 SVC DESIGN

In this section, we present the Speculative Versioning Cache (SVC) as a progression of designs to ease understanding. Each design improves performance over the previous one by tracking more information. We begin with a brief review of snooping bus-based cache coherence and then present a base SVC design which provides support for speculative versioning with minimal modifications to the cache coherence scheme. We then highlight the performance bottlenecks in the base design and introduce optimizations one by one in the rest of the designs.

### 3.1 Snooping Bus-Based Cache Coherence

Fig. 3 shows a 4-processor SMP with private L1 caches that uses a snooping bus to keep the caches consistent. Each cache line comprises an address tag that identifies the data that is cached, the data that is cached, and two bits (valid and store) representing the state of the line. The valid (*V*) bit is set if the line is valid. The store (*S*) or dirty bit is set when a processor stores to the line.

A cache line is in one of three states: Invalid, Clean, and Dirty. A request (load or store) from a processor to its L1 cache *hits* if a valid line with the requested tag is in an appropriate state; otherwise, it *misses*. Cache misses issue bus requests while cache hits do not. More specifically, a load from a clean or dirty line and a store to a dirty line result in cache hits. Otherwise, the load(store) misses and the cache issues a *BusRead(BusWrite)* request. The L1 caches and the next level memory snoop the bus on every request. If a cache has a valid line with the requested tag, it issues an appropriate response according to a *coherence protocol*. A store to a clean line misses and the cache issues a *BusWrite* request. An invalidation-based coherence protocol invalidates copies of this line in all other caches, if any.

#### 3.1.1 SMP Finite State Machine

The finite state machine shown in Fig. 4 implements an invalidation-based MRSW protocol that invalidates all but




Fig. 4. SMP invalidation-based protocol FSM. (a) PU request. (b) Bus request.

the requestor's copy of a line on a *BusWrite* request. This protocol allows a dirty line to be present in only one of the caches. However, a clean line can be present in multiple caches simultaneously. The cache with the dirty line supplies the data on a *BusRead* request. A cache issues a *BusWback* request to cast out a dirty line on a replacement. This simple protocol can be extended by adding an *exclusive* bit to the state of each line to cut down traffic on the shared bus. If a cache line has the exclusive bit set, then it has the only valid copy of the line and can perform a store to that line locally. The SVC designs we discuss in the following sections also use an invalidation-based protocol.

Fig. 5 shows snapshots of the cache lines with tag or address *A* in an SMP with four processors, *W*, *X*, *Y*, and *Z*. The state of the cache line is shown in a box corresponding to that cache. An empty box corresponding to a cache

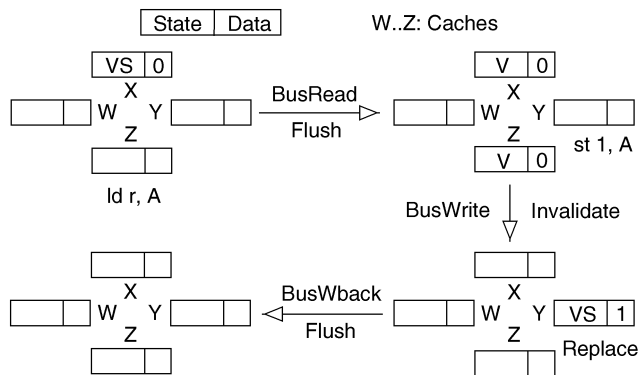


Fig. 5. Cache coherence example.

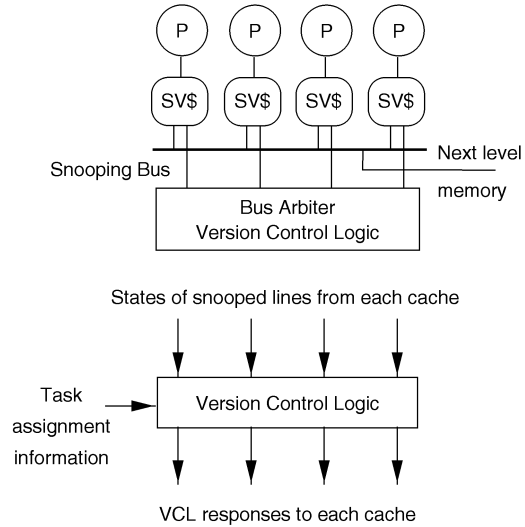


Fig. 6. Speculative versioning cache.

represents that the line is not present in that cache. The first snapshot is taken before processor *Z* issues a load from *A* and misses in its private cache. The cache issues a *BusRead* request and cache *X* supplies the data on the bus. The second snapshot shows the final state of the lines; they are clean. Later, processor *Y* issues a *BusWrite* request to perform a store to *A*. The clean copies in caches *X* and *Z* are invalidated and the third snapshot shows the final state. Now, if cache *Y* chooses to replace this line, it casts out the line to memory by issuing a *BusWback* request; the final state is shown in the fourth snapshot; only the next level memory contains a valid copy of the line.

### 3.2 Base SVC Design

The organization of the private L1 caches in the SVC design is shown in Fig. 6; all the SVC designs use the same organization. The base design minimally modifies the memory system of the snooping bus-based cache coherent SMP to support speculative versioning for processors based on the hierarchical execution model. We assume that memory dependences among loads and stores executed by an individual processor are ensured by a conventional load-store queue; our design guarantees program order among loads and stores from different processors. The base design also assumes that the cache line size is one word; a later design relaxes this assumption. First, we introduce the modifications to the SMP coherent cache and then discuss how the individual operations listed in Table 1 are performed.

1. Each cache line maintains an extra state bit called the load (*L*) bit, as shown in Fig. 7. The *L* bit is set when a task loads from a line before storing to the line—a



Fig. 7. Base SVC design: structure of a line.

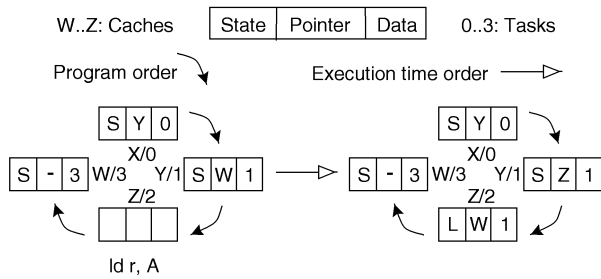


Fig. 8. Base SVC design: example load.

potential violation of memory dependence in case a previous task stores to the same line.

- Each cache line maintains a pointer that identifies the processor (or L1 cache) that has the next copy/version, if any, in the Version Ordering List (VOL) for that line. Thus, the VOL for a line is stored in a distributed fashion among the private L1 cache lines. It is important to note that the pointer identifies a processor rather than a task. Storing the VOL explicitly in the cache lines using pointers may not be necessary for the base design. However, it is necessary to explicitly store the VOL for the advanced designs and we introduce it in the base design to ease the transition to the advanced designs.
- The SVC uses combinational logic called the Version Control Logic (VCL) that provides support for speculative versioning using the VOL. A processor request that hits in the private L1 cache does not need to consult the VOL and, hence, does not issue a bus request; the VCL is also not used. Cache misses issue a bus request that is snooped by the L1 caches and the next level memory. The states of the requested line in each L1 cache and the VOL are supplied to the VCL. The VCL uses the bus request, the program order among the tasks and the VOL to compute appropriate responses for each cache. Each cache line is updated based on its initial state, the bus request and the VCL response. A block diagram of the Version Control Logic is shown in Fig. 6. For the base design, the VCL responses are similar to that of the disambiguation logic in the ARB [5]. The disambiguation logic searches for previous or succeeding stages in a line to execute a load or store, respectively.

### 3.2.1 Loads

Loads are handled in the same way as in an SMP except that the *L* bit is set if the line was initially invalid. On a *BusRead* request, the VCL locates the closest previous version by searching the VOL in the reverse order beginning from the requestor; this version, if any, is supplied to the requestor. If a previous version is not buffered in any of the L1 caches, the next level memory supplies the data. Task assignment information is used to determine the position of the requestor in the VOL. The VCL can search the VOL in reverse order because it has the entire list available and the list is short.

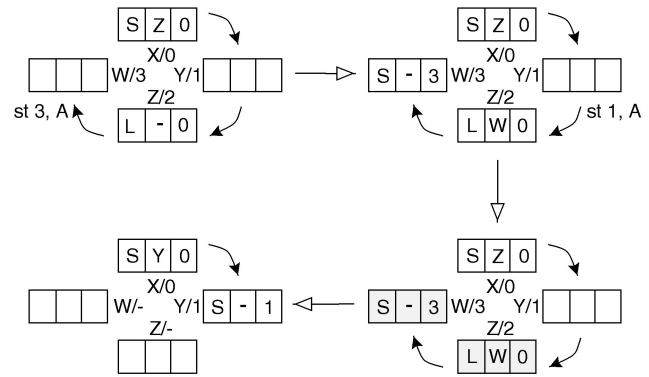


Fig. 9. Base SVC design: example stores.

We illustrate the load executed by task 2 to address *A* in the example program. Fig. 8 shows two snapshots: one before the load executes and one after the load completes. Each box shows the line with tag or address *A* in an L1 cache (the valid bit is not explicitly shown). The number adjacent to a box gives a processor/cache identifier and a task identifier. The processor identifiers are used by the explicit pointers in each line to represent the VOL, whereas, the task identifiers serve only to ease the explanation of the examples. Task 2 executes a load that misses in cache *Z* and results in a bus request. The VCL locates cache *Z* in the VOL for address *A* using program order and then searches the VOL in the reverse order to find the correct version to supply, which is the version in cache *Y* (the version created by task 1).

### 3.2.2 Stores

The SVC performs more operations on a store miss as compared to a cache coherent SMP. When a *BusWrite* request is issued on a store miss, the VCL sends invalidation responses to the caches beginning from the requestor's immediate successor (in task assignment order) to the cache that has the next version (including it, if it has the *L* bit set). This invalidation response allows for multiple versions of the same line to exist and also serves to detect memory dependence violations. A cache sends a task squash signal to its processor when it receives an invalidation response from the VCL and the *L* bit is set in the line.

We illustrate the stores executed by tasks 1 and 3 in the example program. Fig. 9 shows four snapshots of the cache lines with address *A*. The first snapshot is taken before task 3 executes a store that results in a *BusWrite* request. Since task 3 is the most recent in program order, the store by task 3 does not result in any invalidations. Note that a store to a line does not invalidate *all* other cache lines (unlike an SMP) to allow for multiple versions of the same line. The second snapshot is taken after the store from task 3 completes and before task 1 executes its store. Based on task assignment information, the VCL sends an invalidation response to each cache from the one after cache *Y* until the one before cache *W*, which has the next version of the line (cache *W* is not included since it does not have the *L* bit set)—VCL sends an invalidation response to cache *Z*. But, the load executed by task 2, which follows the store by task 1 in program order, has already executed. Cache *Z* detects a




Fig. 10. Base SVC protocol FSM. (a) PU request. (b) Bus request.

memory dependence violation since the  $L$  bit is set when it receives an invalidation response from the VCL. Tasks 2 and 3 are squashed as shown in the third snapshot by shaded boxes. The final snapshot is taken after the store by task 1 has completed.

### 3.2.3 Task Commits and Squashes

The base SVC design handles task commits and squashes in a naive manner. When a processor commits a task, all dirty lines in its L1 cache are *immediately* written back to the next level memory and all other lines are invalidated. To write back all the dirty lines immediately, a list of the stores executed by the task is maintained by the processor. When a task is squashed, *all* lines in the corresponding cache are invalidated.

### 3.2.4 Base SVC Finite State Machine

Support for the SVC is similar to an SMP cache; a cache line is in one of three states: Invalid, Clean and Dirty. The state encodings and the finite state machine are shown in Fig. 10. The  $L$  bit is not shown in the figure. The finite state machine used by the L1 cache controllers in the base SVC design includes the VCL responses and a few modifications to those used by an SMP. A *BusWrite* request is generated on a store to a clean or invalid line. Caches that receive the *Invalidate* response from the VCL invalidate their cache line and send a squash signal to the PU if the  $L$  bit is set in their line. The VCL allows multiple speculative versions to exist by not invalidating all lines as an MRSW protocol would. Another difference between the state machines in Figs. 4 and 10 is that a dirty line that supplies data on a *BusRead* request remains dirty since the next level memory cannot store the speculative data yet.

A valid line can be replaced by the head task’s cache. Other caches cannot replace a valid line because it contains information necessary to guarantee correct execution. Hence, either a different replacement victim is chosen or the PU request stalls until the task becomes nonspeculative (i.e., the head task). On task commits, all clean lines are invalidated and dirty lines are written back to the next level memory by issuing *BusWback* requests. All lines are invalidated on task squashes.

### 3.3 Base Design Performance Drawbacks

The base design just described has two significant performance limitations that make it less desirable: 1) write backs performed when a processor commits a task lead to bursty bus traffic that may increase the time to commit the task and delay issuing a new task to that processor 2) clean lines are also invalidated when a task commits or squashes because the buffered versions could be stale for the new task allocated on the same processor; the correct version may be present in other caches. Consequently, *every* task begins execution with a *cold* L1 cache, increasing the bandwidth demand. The following advanced designs eliminate these problems by tracking additional information.

1. The first advanced design, the *ECS design* (Section 3.5), makes task commits and squashes more efficient. To ease the understanding of this design, we first present an intermediate design, the *EC design* (Section 3.4), that makes task commits efficient by distributing the write backs of dirty lines over time. Also, it retains read-only data in the L1 caches across task commits by careful bookkeeping. However, it assumes that mis-predictions do not occur. Then, we present the *ECS design* that extends the *EC design* to allow task squashes. Task squashes are as simple as in the base design, but are more efficient as they retain non-speculative data in the caches across task squashes.
2. The second advanced design (Section 3.7) boosts the hit-rate of the *ECS design* by allowing requests to *snarf* [7], [3] the bus to account for reference spreading. Snarfing involves copying the data supplied on a bus request issued by another processor in an attempt to combine bus requests indirectly.
3. The final design (Section 3.8) is realistic and allows the size of a cache line to be more than one word.

### 3.4 Implementing Efficient Task Commits (EC)

The *EC design* avoids expensive cache flushes on task commits by maintaining an extra state bit, called the commit bit, in each cache list. Task commits do not stall until all lines with speculative versions are written back. The *EC design* eliminates write back bursts on the bus during task commits. Also, no extra hardware is necessary to maintain a list of stores performed by each task. Further, the *EC design* improves cache utilization by keeping the L1 caches warm across tasks.

The structure of a cache line in the *EC design* is shown in Fig. 11. When a processor commits a task, the  $C$  bit is set in *all* its cache lines. This operation is entirely local to the L1 cache and does not issue a bus request. A dirty




Fig. 11. EC design: structure of a line.

committed line is written back, if necessary, when it is accessed the next time either on a processor request or on a bus request. Therefore, committed versions could remain in the caches until much later in time since the task that created the version committed. The order among committed and uncommitted versions is still maintained by the explicit pointers in the line. This order among the versions is necessary to write back the correct committed version and to supply the correct version on a bus request. The EC design uses an additional state bit, the sTale (*T*) bit, to retain read-only data across tasks. First, we discuss how loads and stores are handled when caches have both committed and uncommitted versions and then discuss the stale bit.

### 3.4.1 Loads and Stores

Loads to committed lines are handled like cache misses and issue a bus request. The VCL searches the VOL in the reverse order beginning from the requestor for the closest previous *uncommitted* version; this version, if any, is supplied to the requestor. If no such version is found, the VCL supplies the most recent *committed* version, if any. This version is the first committed version that is encountered on the reverse search. All other committed versions need not be written back and are invalidated. On a store miss, committed versions are purged in a similar fashion.

We illustrate the load executed by task 2 in the example program. Fig. 12 shows two snapshots: one before the load executes and one after the load completes. Versions 0 and 1 have been committed (the *C* bit is set in the lines in caches *X* and *Y*). Task 2 executes a load that misses in cache *Z* and results in a bus request. The VCL knows that task 2 is the head task and determines that cache *X* has the most recent committed version. Cache *X* supplies the data which is also written back to the next level memory. Other committed versions (version 0) are invalidated and are never written back to memory. The VCL also inserts the new copy of version 1 into the VOL by modifying the pointers in the lines accordingly—the second snapshot shows the modified VOL.

Fig. 13 illustrates the actions performed on a store miss. The first snapshot is taken before a store is executed by task 5. Versions 0 and 1 have been committed. Task 5 executes a store that misses in cache *Y* and results in a *BusWrite* request even though the line has a committed version. The




Fig. 12. EC Design: Example load.




Fig. 13. EC Design: Example store.

VCL purges all committed versions of this line—it determines that version 1 has to be written back to the next level memory and the other versions (version 0) can be invalidated. Purging the committed versions also makes space for the new version (version 5). The modified VOL shown in the second snapshot contains only the two uncommitted versions.

### 3.4.2 Stale Copies

The EC design makes task commits efficient by delaying to commit each cache line until a later time. Therefore, a cache line could have a *stale copy* because versions more recent than the version buffered by the committed task could be present in other caches. The base SVC design does not introduce stale copies because it invalidates all nondirty lines whenever a task commits. The EC design uses the stale (*T*) bit to distinguish stale copies from correct copies and avoids issuing a bus request on accesses to correct copies. This additional information allows the EC design to retain read-only data (correct copies) across task commits. First, we illustrate that stale and correct copies are indistinguishable without the *T* bit and then show how the *T* bit is used.

Fig. 14 shows two execution time lines—one that leaves a correct copy of address *A* (shown using solid lines) in cache *Z* and another that leaves a stale copy of address *A* in the same cache (shown using dashed lines). The first time line shows a sample execution of a modified version of our example program—task 3 in Fig. 2 does not execute the store. The second time line shows an execution of our original program. The first snapshot is the same for both time lines. The second snapshot in the second time line is taken after tasks 0 and 1 have committed. The *C* bit is set in their caches and new tasks 4 and 5 have been allocated. The final snapshot in both time lines are taken when tasks 4 to 7 are active and before task 6 executes a load. In the first time line, the data in cache *Z* is a correct copy, since no versions were created after version 1; the load can be supplied data

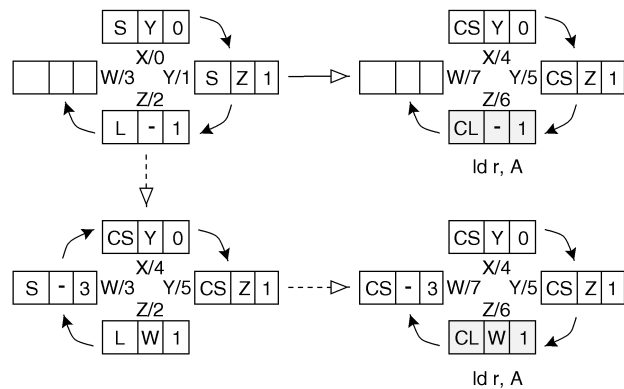


Fig. 14. EC design: correct and stale copies.



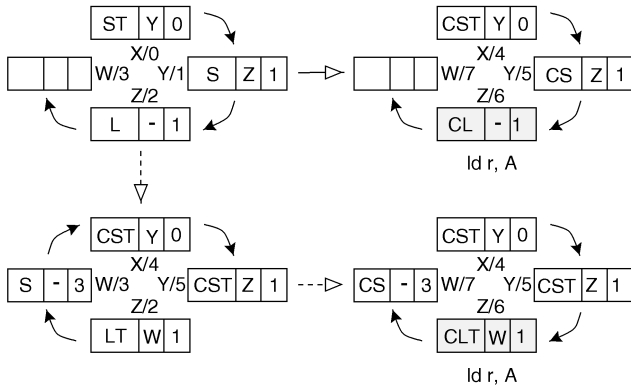


Fig. 15. EC design: Using the stale bit.

by just resetting the *C* bit and without issuing a bus request. In the second time line, the copy in cache *Z* is stale since the creation of version 3 and, hence, the load misses resulting a bus request. However, cache *Z* cannot distinguish between these two scenarios and has to issue a request in both cases to consult the VOL and obtain a copy of the correct version.

The EC design uses the stale (*T*) bit to distinguish between these two scenarios and avoids the bus request whenever a copy is not stale. The design maintains the invariant: the most recent version of an address and its copies have the *T* bit reset and the other copies and versions have the *T* bit set. This invariant is easily guaranteed by resetting the *T* bit in the most recent version, or a copy thereof, when it is created and setting the *T* bit in the previous versions, if any. The *T* bits are updated on the *BusWrite* request issued to create a version or a *BusRead* request issued to copy a version and, hence, do not generate additional bus traffic. Since stores in different tasks can be executed out of program order, an active task could execute a store to a copy that has the *T* bit set (the copy is not stale for this task, but is stale for the next task allocated to the same processor). Fig. 15 shows the two time lines in our example with the status of the *T* bit. Cache *Z* can distinguish between the correct copy (*T* bit is not set) and the stale copy (*T* bit is set). The load hits if a correct copy is present and no bus request is issued.

The EC design eliminates the serial bottleneck in flushing the L1 cache on task commits by using the commit (*C*) bit. Also, this design retains nondirty lines after task commits as long as they are not stale. More generally, read-only data used by a program is fetched *only once* into the L1 caches and never invalidated unless chosen to be replaced on a cache miss. Further a task commits by just setting the *C* bit in all lines in its L1 cache.

### 3.5 Implementing Efficient Task Squashes (ECS)

The ECS design extends the EC design to allow task squashes for the EC design. Also, the ECS design makes the task squashes more efficient than in the base design by retaining non-speculative data in the caches across squashes using another state bit, the Architectural (*A*) bit. The structure of a line in the ECS design is shown in Fig. 16.

When a task squashes, all uncommitted lines (lines with the *C* bit reset) are invalidated by resetting the valid (*V*) bit.




Fig. 16. ECS design: structure of a line.

The invalidation makes the pointers in these lines and their VOLs inexact. The VOL has a (dangling) pointer in the last valid (or unsquashed) copy or version of the line and the status of the *T* bit in the lines are incorrect. The ECS design repairs the VOL of such a line when the line is accessed later either on a processor request or on a bus request. Updating the *T* bits is not necessary because it is only a hint to avoid a bus request and a squash would not incorrectly reset a stale version to be correct. However, the ECS design updates the *T* bit on this bus request by consulting the repaired VOL.

Fig. 17 illustrates VOL repair with an example time line with three snapshots. The first snapshot is taken just before the task squash occurs. Tasks 3 and 4 are squashed; only version 3 is invalidated. The VOL with incorrect *T* bits and the dangling pointer are shown in the second snapshot. Task 2 executes a load that misses in cache *W* and results in a bus request. The VCL resets the dangling pointer and the *T* bit in cache *Y*. The VCL then determines the version to supply the load. Also, the most recent committed version (version 0) is written back to the next level memory. The third snapshot is taken after the load has completed.

#### 3.5.1 Squash Invalidations

The base design invalidates nondirty lines in the L1 cache on task squashes. This includes both speculative data from previous tasks and architectural data from the next level memory (or the committed tasks). The base design invalidates these lines because it does not track the creator of a speculative versions for each line and, hence, cannot determine whether the version in a line has been committed or squashed. Squashing non-speculative data leads to higher miss rates for tasks that are squashed and restarted multiple times.

To distinguish between copies of speculative and architectural versions, we add the architectural (*A*) bit to each cache line, as shown in Fig. 16. The *A* bit is set in a copy if either the next level memory or a committed version supplies data when a bus request issued to obtain the copy;

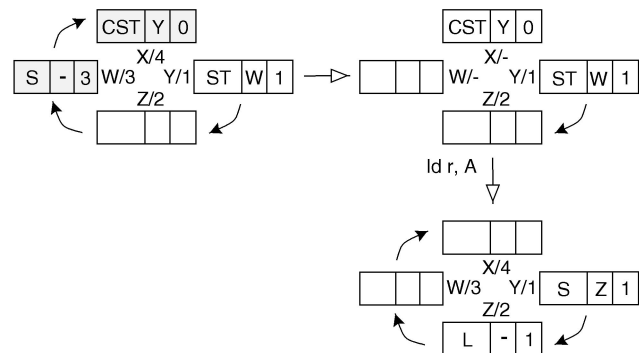


Fig. 17. ECS design: VOL repair.

else the  $A$  bit is reset. One of the VCL responses on a bus request specifies whether the  $A$  bit should be set or reset. Copies of architectural versions are not invalidated on task squashes, i.e., the ECS design only invalidates lines that have both the  $A$  and  $C$  bits reset. Further, a copy of a speculative version used by a task becomes an architectural copy when the task commits. However, the  $A$  bit is not set until the line is accessed by a later task, when the  $C$  bit is reset and the  $A$  bit is set.

### 3.6 Final SVC Design

The final Speculative Versioning Cache (SVC) design is shown in Fig. 6 and the structure of each SVC line is shown in Fig. 16. In addition, the final design uses a hybrid update-invalidate coherence protocol that dynamically selects either write-update or write-invalidation for each bus request. The write-update protocol helps to reduce the inter-task communication latency through memory, i.e., the latency between a store from a task and a dependent load from one of the later tasks, that executes concurrently. On the other hand, the write-invalidate protocol could reduce the bus bandwidth requirement.

Our solution for speculative versioning for hierarchical execution models resolves complexity by decoupling the directory from the disambiguation logic. The directory is distributed as in an SMP and the disambiguation logic is the Version Control Logic. We summarize the final SVC design by first describing the finite state machine used by the cache controllers and then the VCL. The finite state machine has been intentionally made simple to project the important features of the design; specifically, it builds on an invalidation-based coherence protocol.

#### 3.6.1 Final SVC Finite State Machine

The finite state machine used by the L1 cache controllers of the final SVC is shown in Fig. 18. Each line in this SVC is in one of five states: Invalid, Active Clean, Active Dirty, Passive Clean, and Passive Dirty. The two active states correspond to the two valid states in Fig. 10—lines in these two states have been accessed by the task that is executing on the corresponding PU. Adding the  $C$  bit to the base SVC adds the two passive states; committed versions and their copies are in the passive states. The bits set/reset in each state are shown in the figure. To simplify the exposition, the  $L$ ,  $T$ ,  $A$ , and  $X$  bits are not discussed unless of significance. Also, replacements and cache to cache transfers of read-only data are not shown. Lines in the clean states can be replaced without a bus request. Dirty lines generate a  $BusWback$  request on a replacement. Similar to the base design, active lines can be replaced only by the head task.

On a task commit, an active dirty line becomes passive and, hence, its writeback is postponed until later. However, an active clean line is either invalidated or retained on a task commit depending on whether it is stale or not; again this is postponed until later by just setting the  $C$  bit. On a task squash, an active dirty line is invalidated resulting in a temporary inconsistency in the line's VOL which is fixed by the VCL later. An active clean line is either invalidated or retained on a task squash depending on whether it is a copy of a speculative version or the architectural version.




Fig. 18. Final SVC Protocol FSM. (a) PU request. (b) Bus request.

Load/store requests to invalid and active lines are handled similar to the base design. A load to a line in the passive clean state does not result in a bus request unlike the base design which would generate a  $BusRead$  request for the load (as the line would be invalid). On a  $BusRead$  request, the data could be supplied either by a line that is in the active dirty state or the passive dirty state. On a bus request, a line in passive dirty state is invalidated whether it is flushed or not. Similar to the base design, the VCL determines the copies/versions to be invalidated on a  $BusWrite$  request. However, an invalidate response received by a line in one of the passive states will not generate a squash unlike a line in one of the active states. The most recent committed version is written back to the next level memory.

### 3.7 Hit Rate Optimizations

The base and ECS designs incur severe performance penalties due to *reference spreading*. When a uniprocessor program is executed on multiple processors with private L1 caches, successive accesses to the same line that hit after

missing once in a shared L1 cache could result in a series of misses. This phenomenon is also observed for parallel programs where the miss rate for read-only shared data with private caches is higher than that with a shared cache. We use *snarfing* [7], [3] to mitigate this problem. Our SVC implementation snarfs data on the bus if the corresponding cache set has a free line available. However, an active task's cache can only snarf the version that the task can use unlike an SMP coherent cache. The VCL determines whether a task can copy a particular version or not and informs caches of an opportunity to snarf data on a bus request.

### 3.8 Realistic Line Size

The base and ECS designs assume that the line size of the L1 caches is one word. The final SVC design however allows lines to be longer than a word. Similar to an SMP coherent cache, we observe effects due to false sharing. In addition to causing higher bus traffic, false sharing leads to more squashes when a store to a cache line from a task is executed out-of-order with a load from a different byte or word in the same line from a later task. We mitigate the effects of false sharing by using a technique similar to the sector cache [10]. Each line is divided into subblocks and the *L* and *S* bits are maintained for each subblock. The size of a subblock or versioning block is less than that of the address block (storage unit for which an address tag is maintained). Also, when a store miss results in a *BusWrite* request, mask bits that indicate the versioning blocks modified by the store are also made available on the bus.

## 4 PERFORMANCE EVALUATION

We report performance results for the SVC using some of the SPEC95 benchmarks. The goal of our implementation and evaluation is to prove the SVC design not just to analyze its performance. We underline the importance of a private cache solution by first showing how performance degrades rapidly as the hit latency for a shared cache solution is increased; the Address Resolution Buffer (ARB) is the shared cache solution we use for this evaluation. We mitigate the commit time bottlenecks in the ARB (by using an extra stage that contains architectural data) to isolate the effects of pure hit latency from other performance bottlenecks.

### 4.1 Methodology and Configuration

All the results in this paper were collected on a simulator that faithfully models a Multiscalar processor. The simulator dynamically switches between a functional and a detailed cycle-by-cycle model to provide accurate and fast simulation of a program. The memory system model includes a fair amount of detail including an L2 cache, DRAM banks and interconnects between the different levels of memory hierarchy. The Multiscalar processor used in the experiments has four processors each of which can issue two instructions out-of-order. Each processor has two simple integer ALUs, 1 complex integer unit, one floating point unit, 1 branch unit and 1 address calculation unit, all of which are assumed to be completely pipelined. Inter-

processor register communication latency is one cycle and each processor can send as many as two registers to its neighbor in every cycle.

Loads and stores from each processor are executed in program order by using a load/store queue of 16 entries each. The L2 cache is 4MB in size with a 10-cycle access time and a total peak bandwidth of 16-bytes per processor clock to the L1 data, instruction and task caches. Main memory uses a RAMBUS-like interface that operates at half the speed of the processors to provide a peak bandwidth of eight bytes every bus clock.

The ARB is a fully-associative set of 32-byte lines with a total of 8KB storage per stage and five stages; the shared data cache that backs up the ARB is 2-way set associative and 64KB in size. All the caches and memory are 4-way interleaved. Both the ARB and the L1 data cache have 16 MSHRs/writebuffers each; each buffer can combine up to eight accesses to the same line. Disambiguation is performed at the byte-level. The base ARB hit time is varied from one to three cycles in the experiments. Both the tags and data RAMs are single ported in all the caches.

The private caches that comprise the SVC are connected together and with the L2 cache by a 256-bit split-transaction snooping bus where a transaction requires four bus clocks to determine whether another cache or next level memory will supply the data. Unlike conventional SMPs whose snooping buses are off chip, the SVC's snooping bus is on chip and therefore can be fast. Bus arbitration occurs only once for cache to cache data transfers. Extra bus cycles are used to flush a committed version to the next level memory. Each processor has its own private L1 cache with 16KB of 4-way set-associative storage in 32-byte lines. Both loads and stores are nonblocking with eight MSHRs/writebuffers per cache. Each buffer can combine up to four accesses to the same line. Disambiguation is performed at the byte-level. L1 cache hit time is fixed at one cycle. The tag RAM is dual ported to support snooping while the data RAM is single ported.

### 4.2 Benchmarks

We used the following programs from the SPEC95 benchmark suite with train inputs except in the cases listed: *compress*, *gcc (ref/jump.i)*, *vortex*, *perl*, *jpeg (test/specmun.ppm)*, *mgrid (test/mgrid.in)*, *apsi*, *fpppp*, and *turb3d*. All programs were stopped after executing one billion instructions. From past experience, we know that for these programs performance change is not significant beyond one billion instructions.

### 4.3 Experiments

Fig. 19 presents the instructions per cycle (IPC) for a Multiscalar processor with either the ARB or the SVC. The configurations keep total data storage of the SVC and ARB/cache storage roughly the same. The percentage miss rates for the ARB and the SVC are shown on top of the IPC bar clusters (in that order). For the SVC, an access is counted as a miss if data is supplied by the next level memory; data transfers between the L1 caches are not counted as misses.




Fig. 19. SPEC95 IPCs for ARB and SVC.

From these experiments, we make three observations: 1) the hit latency of data memory significantly affects ARB performance, 2) the SVC trades-off hit rate for hit latency and the ARB trades-off hit latency for hit rate to achieve performance, and 3) for the same total data storage, the SVC performs better than the ARB having a hit latency of two or more cycles, as shown in Fig. 19. The graphs in these figures show that performance improves in the range of 5 percent to 20 percent when decreasing the hit latency of the ARB from three cycles to one cycle. This improvement indicates that techniques that use private caches to improve hit latency are an important factor in increasing overall performance, even for latency tolerant processors like a Multiscalar processor.

The distribution of storage for the SVC produces higher miss rates than for the ARB. We attribute the increase in miss rates for the SVC to two factors. First, distributing the available storage results in reference spreading (Section 3.7) and replication of data reduces available storage. We measured that about 25 percent of SVC's hits shown in Fig. 19 comes from snarfed data, indicating that there is significant reference spreading. The significant percentage of snarfed-data hits also emphasizes the importance of snarfing for the SVC. Second, a latest version of a line that caches fine-grain shared data between Multiscalar tasks constantly moves from one L1 cache to another (migratory data). Such fine-grain communication increase the number of total misses as well. Our measurements show that about 7 to 12 percent of all accesses (both loads and stores) are satisfied by dirty blocks in other caches, indicating that there is moderate amounts of fine-grain sharing.

## 5 CONCLUSION

Speculative versioning is important to overcome limits on Instruction Level Parallelism (ILP) due to ambiguous memory dependences in a sequential program. Our proposal, called the Speculative Versioning Cache(SVC),

uses distributed caches to eliminate the latency and bandwidth problems of a previous solution, the Address Resolution Buffer, which uses a centralized buffer. The SVC conceptually unifies cache coherence and speculative versioning by using an organization similar to snooping bus-based coherent caches. Our evaluation for the Multiscalar architecture shows that hit latency is an important factor affecting performance and private cache solutions trade-off hit rate for hit latency. The SVC provides hardware support to break ambiguous memory dependences allowing proposed next generation multiprocessors to use aggressive parallelizing software for sequential programs.


## ACKNOWLEDGMENTS

The authors would like to thank Scott Breach, Andreas Moshovos, Subbarao Palacharla and the anonymous referees for their comments and valuable suggestions on earlier drafts of the paper. This work was supported in part by US National Science Foundation Grants CCR-9303030 and MIP-9505853, ONR Grant N00014-93-1-0465, and by US Army Intelligence Center and Fort Huachuca under Contract DABT63-95-C-0127 and ARPA order no. D346 and a donation from Intel Corp. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the US Army Intelligence Center and Fort Huachuca, or the US Government.

## REFERENCES

- [1] IEEE Standard for Scalable Coherent Interface (SCI) pp. 1596-1992, 1993.
- [2] M. Cintra, J.F. Martinez, and J. Torrellas, "Architectural Support for Scalable Speculative Parallelization in Shared-Memory Multiprocessors," *Proc. 27th Ann. Int'l Symp. Computer Architecture*, pp. 13-24, June 2000.
- [3] F. Dahlgren, "Boosting the Performance of Hybrid Snooping Cache Protocols," *Proc. 22nd Ann. Int'l Symp. Computer Architecture*, pp. 60-69, 1995.
- [4] M. Franklin and G.S. Sohi, "The Expandable Split Window Paradigm for Exploiting Fine-Grain Parallelism," *Proc. 19th Ann. Int'l Symp. Computer Architecture, Assoc. Computing Machinery*, pp. 58-67, May 1992.
- [5] M. Franklin and G.S. Sohi, "(ARB): A Hardware Mechanism for Dynamic Reordering of Memory References," *IEEE Trans. Computers*, vol. 45, no. 5, pp. 552-571, May 1996.
- [6] J.R. Goodman, "Using Cache Memory to Reduce Processor-Memory Traffic," *Proc. 10th Ann. Int'l Symp. Computer Architecture*, pp. 124-131, 1983.
- [7] J.R. Goodman and P.J. Woest, "The Wisconsin Multicube: A New Large-Scale Cache-Coherent Multiprocessor," *Proc. 15th Ann. Int'l Symp. Computer Architecture*, pp. 422-431, 1988.
- [8] S. Gopal, T. Vijaykumar, J.E. Smith, and G.S. Sohi, "Speculative Versioning Cache," *Proc. Fourth IEEE Symp. High-Performance Computer Architecture*, pp. 195-205, Jan. 1998.
- [9] L. Hammond, M. Willey, and K. Olukotun, "Data Speculation Support for a Chip Multiprocessor," *Proc. Eighth Int'l Conf. Architectural Support for Programming Languages and Operating Systems*, pp. 58-69, Nov. 1998.
- [10] J.S. Liptay, "Structural Aspects of the System/360 Model 85 Part II: The Cache," *IBM Systems J.*, vol. 7, no. 1, pp. 15-21, 1968.
- [11] A. Moshovos, S.E. Breach, T.N. Vijaykumar, and G.S. Sohi, "Dynamic Speculation and Synchronization of Data Dependences," *Proc. 24th Ann. Int'l Symp. Computer Architecture*, pp. 181-193, June 1997.

- [12] K. Olukotun, B.A. Nayfeh, L. Hammond, K. Wilson, and K.-Y. Chang, "The Case for a Single-Chip Multiprocessor," *Proc. Seventh Int'l Symp. Architectural Support for Programming Languages and Operating Systems*, pp. 2-11, Oct. 1996.
- [13] S. Palacharla, N.P. Jouppi, and J.E. Smith, "Complexity-Effective Superscalar Processors," *Proc. 24th Ann. Int'l Symp. Computer Architecture*, pp. 206-218, June 1997.
- [14] E. Rotenberg, Q. Jacobson, Y. Sazeides, and J. Smith, "Trace Processors," *Proc. 30th Ann. Int'l Symp. Microarchitecture, Assoc. Computing Machinery*, pp. 138-147, Dec. 1997.
- [15] G.S. Sohi, S.E. Breach, and T.N. Vijaykumar, "Multiscalar Processors," *Proc. 22nd Ann. Int'l Symp. Computer Architecture*, pp. 414-425, June 1995.
- [16] J.G. Steffan, C.B. Colohan, A. Zhai, and T.C. Mowry, "A Scalable Approach to Thread-Level Speculation," *Proc. 27th Int'l Symp. Computer Architecture*, pp. 1-12, June 2000.
- [17] J.G. Steffan and T.C. Mowry, "The Potential for Thread-Level Data Speculation in Tightly-Coupled Multiprocessors," *Proc. Fourth Int'l Symp. High-Performance Computer Architecture*, pp. 2-13, Feb. 1998.
- [18] S. Vajapeyam and T. Mitra, "Improving Superscalar Instruction Dispatch and Issue by Exploiting Dynamic Code Sequences," *Proc. 24th Ann. Int'l Symp. Computer Architecture*, pp. 1-12, June 1997.
- [19] Y. Zhang, L. Rauchwerger, and J. Torrellas, "Hardware for Speculative Parallelization of Partially-Parallel Loops in DSM Multiprocessors," *Fifth Int'l Symp. High-Performance Computer Architecture (HPCA-5)*, Jan. 1999,



**T.N. Vijaykumar** received the PhD degree from the University of Wisconsin, Madison, in 1998. He joined the faculty of the School of Electrical and Computer Engineering after graduation. His research interests are in computer architecture, VLSI microarchitectures, instruction level parallelism, processor and memory hierarchy hardware, and compiler optimizations. Currently, at Purdue, Dr. Vijaykumar leads projects investigating program-behavior based intelligent caching schemes and coleads projects on low power architectures and multi-threaded architectures. He is a member of the IEEE and the IEEE Computer Society.



**Sridhar Gopal** received the Masters degree from the University of Wisconsin in 1998. Currently, he defines software engineering processes and tools for the Network Operations Software Group of Lucent Technologies, Inc. He has operated in several software engineering roles including project management, product design, vendor management, and customer support. He has worked in different areas including network management, compiler optimization, and linux integration. His research interests include processor microarchitecture, memory system architecture, and quantitative performance analysis. He is also interested in developing solutions based on open source software for problems in markets dominated by monopolies.

**James E. Smith** is a professor in the Department of Electrical and Computer Engineering at the University of Wisconsin, Madison. During the past two decades he has been involved with a number of computer research and development projects in both industry and academia. He has investigated several aspects of high-performance processor implementations, including branch prediction, instruction issuing methods, and precise interrupt techniques. He also proposed and developed decoupled superscalar processor architectures and implementations. Dr. Smith and his research group at the University of Wisconsin are currently studying new microarchitecture paradigms based on distributed processing elements. His most recent research project is the development of the virtual machine abstraction as a technique for providing high-performance through codesign and tight coupling of virtual machine hardware and software. Dr. Smith was the recipient of the 1999 Eckert-Mauchly Award for contributions to the field of computer architecture. He is a member of the IEEE and the IEEE Computer Society.



**Gurindar Sohi** received the PhD degree from the University of Illinois in 1985 after graduation he joined the University of Wisconsin faculty and is currently a professor in both the Computer Sciences and Electrical and Computer Engineering departments. His research has been in the area of architectural and microarchitectural techniques for high-performance microprocessors, including instruction-level parallelism, out-of-order execution with precise exceptions, non-blocking caches, decentralized microarchitectures, speculative multithreading, and memory dependence speculation. He is a senior member of the IEEE and the IEEE Computer Society.

► For more information on this or any computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.