



# Performance and Power-Consumption Implication of Fine-Grained Synchronization in Multiprocessors

---

A Master of Science Thesis in Computer Systems by Oscar Sierra Merino  
Department of Microelectronics and Information Technology  
Royal Institute of Technology  
Stockholm – May 2.002

Oscar Sierra Merino  
o.sierra@ieee.org  
MSc Thesis author

Vladimir Vlassov  
vlad@it.kth.se  
Main advisor

Mladen Nikitovic  
mladen@it.kth.se  
Secondary advisor

# Abstract

It has been already verified that hardware-supported fine-grain synchronization provides a significant performance improvement over coarse-grained synchronization mechanisms, such as barriers. Support for fine-grain synchronization on individual data items becomes notably important in order to implement thread-level parallelism more efficiently.

One of the major goals of this project is to propose a new efficient way to support fine-grain synchronization mechanisms in multiprocessors. This novel idea is based on the efficient combination of fine-grain synchronization with cache coherence and instruction level parallelism. Both snoopy and directory-based cache coherence protocols have been studied.

The work includes the definition of the complete set of synchronizing memory instructions as well as the architecture of the full/empty tagged shared memory that provides support for these operations.

A detailed model based on a shared memory multiprocessor is presented and systematically described. To achieve this, an existing execution-driven simulator, namely `RSIM`, has been appropriately adapted. The simulation environment is designed for verification and performance evaluation of the proposed solutions.

Some guidelines for implementing a power estimation algorithm as an extended feature of the simulation platform have been presented. The integration of fine-grain synchronization at the cache coherence level is expected to increase the energy consumption of the system.

**Keywords:** fine-grain synchronization, shared memory, instruction-level parallelism, cache coherence, power consumption, execution-driven simulation.

# Table of Contents

1. Overview and motivation .....	8
2. Semantics of synchronizing memory operations .....	9
3. Architectural support for fine-grain synchronization .....	11
3.1. Related work .....	11
3.1.1. The Alewife machine .....	11
3.1.2. The StarT-NG machine .....	12
3.2. Proposed architecture .....	14
3.3. Cache coherence .....	17
4. Integration with snoopy protocols .....	19
4.1. Mapping between processor instructions and bus transactions .....	23
4.2. Management of pending requests .....	25
4.3. Transition rules .....	27
4.3.1. Invalid state .....	27
4.3.2. Modified state .....	28
4.3.3. Exclusive-clean state .....	29
4.3.4. Shared state .....	30
4.4. Summary .....	30
5. Integration with directory-based protocols .....	32
5.1. Mapping between processor instructions and network transactions .....	34
5.2. Management of pending requests .....	36
5.3. Directory transition rules .....	38
5.3.1. Absent state .....	38
5.3.2. Read-only state .....	39
5.3.3. Read-write state .....	41
5.3.4. Read transaction state .....	43
5.3.5. Write transaction state .....	44

5.4. Summary .....	46
6. Simulation framework .....	47
6.1. Features of the simulated platform.....	47
6.2. Simulation methodology .....	48
6.3. Implementation of synchronizing instructions .....	49
6.4. Simulation flowchart.....	52
6.5. Simulation results.....	54
7. Power-consumption estimation .....	57
7.1. Available energy estimation tools.....	57
7.2. Implementing an energy estimation framework in RSIM.....	58
7.3. Planned experiments .....	59
8. Conclusions .....	60
9. Future work.....	61
Appendix A. Preparing binaries for simulation.....	62
Appendix B. Application source (fine-grained version) .....	64
Appendix C. Application source (coarse-grained version) .....	66
Acknowledgements .....	68
References .....	69

# List of Figures

Figure 1: Classification of synchronizing operations (extracted from [72]).....	9
Figure 2: Notation of synchronizing memory operations.....	10
Figure 3: Architecture of a StarT-NG node [27] .....	12
Figure 4: Two sample scenarios of synchronized loads and stores .....	13
Figure 5: Logical structure of shared memory .....	14
Figure 6: Memory map for each processing node .....	15
Figure 7: Organization of a cache supporting fine-grain synchronization.....	15
Figure 8: Cache line containing both ordinary and synchronized data.....	16
Figure 9: Bus-based system architecture .....	19
Figure 10: MESI coherence protocol.....	20
Figure 11: MESI protocol integrated with fine-grain synchronization (explicit full/empty states).....	21
Figure 12: MESI protocol integrated with fine-grain synchronization (implicit full/empty states) .....	22
Figure 13: Sample scenario of mapping between processor instructions and bus transactions .....	24
Figure 14: Resuming of pending requests .....	26
Figure 15: Mesh network-based architecture .....	32
Figure 16: Alewife's coherence protocol state diagram .....	32
Figure 17: Management of pending requests for an absent or read-only memory block.....	36
Figure 18: Management of pending requests for a read-write memory block.....	37
Figure 19: State transitions from the absent state.....	39
Figure 20: State transitions from the read-only state.....	41
Figure 21: State transitions from the read-write state.....	43
Figure 22: State transitions from the read transaction state.....	44
Figure 23: State transitions from the write transaction state.....	45
Figure 24: Simulated system architecture.....	47
Figure 25: Simulation steps .....	48
Figure 26: Simulation steps with a compiler supporting synchronizing instructions .....	49
Figure 27: Alternate load and store instruction format.....	49
Figure 28: Execution flowchart of the simulator.....	52

Figure 29: RSIM_EVENT scheduling .....	53
Figure 30: Instruction lifetime stages .....	53
Figure 31: Normalized execution time for different machine and problem sizes.....	54
Figure 32: Integrated power consumption framework .....	59
Figure 33: Details on how to transform a standard store to a synchronized store .....	63

# List of Tables

Table 1: Notation of synchronized operations.....	10
Table 2: Relevant information stored in ordinary MSHR registers [46] .....	16
Table 3: Additional bus transactions in the MESI protocol.....	22
Table 4: Correspondence between processor instructions and memory requests .....	23
Table 5: Management of coalescing requests .....	25
Table 6: Semantics of the transitions in the directory-based protocol.....	33
Table 7: Network transactions in the directory-based protocol.....	35
Table 8: Correspondence between processor instructions and memory requests .....	35
Table 9: ASI values for synchronizing operations.....	50
Table 10: Specific modifications made to RSIM .....	50
Table 11: Set of full/empty memory instructions.....	51
Table 12: Execution times (in cycles) for 1.000 iterations .....	55
Table 13: Execution times (in cycles) for 16 nodes.....	55
Table 14: Structure of switch capacitance tables.....	58
Table 15: Relevant values of the op3 field .....	63

# 1. Overview and motivation

---

Two types of synchronization operations guarantee correctness in shared-memory multiprocessors: mutual exclusion and conditional synchronization, such as producer-consumer data dependency. Barriers are an example of synchronization directives that ensure the correctness of a producer-consumer behavior. They are coarse-grain in the sense that all processes have to wait in a common point before they can proceed, even though the data they truly depend on was available in a previous execution step.

The main advantage of fine-grain synchronization arises from the fact that synchronization is provided at data-level. As a consequence, false data dependencies and unnecessary process waiting are avoided. Communication overhead due to global barriers is also avoided, because each process communicates only with the processes it depends on. Thus, the serialization of program execution is notably reduced and more parallelism can be exploited. This effect is more noteworthy as the number of processors increases. While the overhead of a fine-grain synchronization operation remains constant, that of a coarse-grain operation typically increases with the number of processors.

As explained in [78], fine-grain synchronization is most commonly provided by three different mechanisms:

- i) language-level support for expressing data-level synchronization operations,
- ii) full/empty bits storing the synchronization state of each memory word,
- iii) processor operations on full/empty bits.

Traditional theory on data-level parallelism has led to the definition of specific structures supporting fine-grain synchronization in data arrays. As an example,  $\mathcal{J}$ -structures provide consumer-producer style of synchronization, while  $\mathcal{L}$ -structures guarantee mutual exclusion access to a data element [4]. Both data types associate a state bit with each element of an array.

Several alternatives exist for handling a synchronization failure. The most immediate are either polling the memory location until the synchronization condition is met or blocking the thread and returning the control at a later stage, which requires more support as it is necessary to save and restore context information. A combination of both is another option, polling first for a given period and then blocking the thread. The waiting algorithm may depend on the type of synchronization being executed [52].

Most research regarding multiprocessors show that fine-grain synchronization is a valuable alternative for improving the performance of many applications. As exposed in [46], evidence is shown on the worthiness of having modest hardware support for fine-grain synchronization. Testing the benefits of aggressive hardware support in fine-grain synchronization is one of the goals of this project.



## 2. Semantics of synchronizing memory operations

---

Synchronization operations require the use of a *tagged* memory, in which each location is associated to a state bit in addition to a 32-bit value. The state bit is known as *full/empty* bit and implements the semantics of synchronizing memory accesses. As a matter of fact, this bit controls the behavior of synchronized loads and stores. A set *full/empty* bit indicates that the corresponding memory reference has been written by a successful synchronized `store`. On the contrary, an unset *full/empty* bit means either that the memory location has never been written since it was initialized or that a synchronized `load` has read it.

A complete categorization of the different synchronizing memory operations is depicted in Figure 1. These instructions are introduced as an extension of the instruction set of *Sparcle* [6], which is in turn based on `SPARC`. The simplest type of operations includes *unconditional load*, *unconditional store*, setting the *full/empty* bit or a combination of these. As they do not depend on the previous value of the *full/empty* bit, *unconditional* operations always succeed.

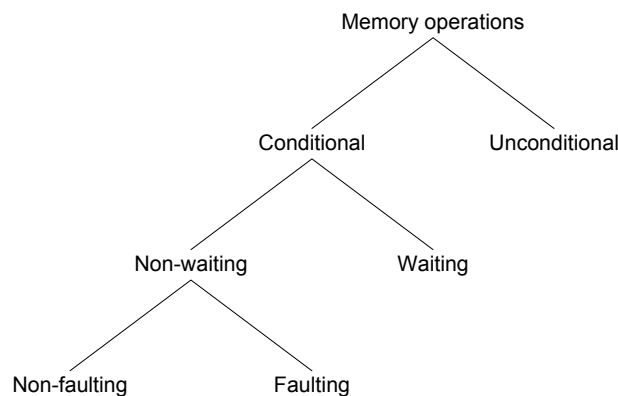


Figure 1: Classification of synchronizing operations (extracted from [72])

*Conditional* operations depend on the value of the *full/empty* state bit to successfully complete. A *conditional* read, for instance, is only performed if the state bit of the location being accessed is set. The complementary applies for a *conditional* write. *Conditional* memory operations can be either *waiting* or *non-waiting*. In the former case, the operation remains pending in the memory until the state miss is resolved. This introduces non-deterministic latencies in the execution of synchronizing memory operations. Lastly, *conditional non-waiting* operations can be either *faulting* or *non-faulting*. While the latter do not treat the miss as an error, *faulting* operations fire a trap on a state miss and either retry the operation immediately or switch to another context.

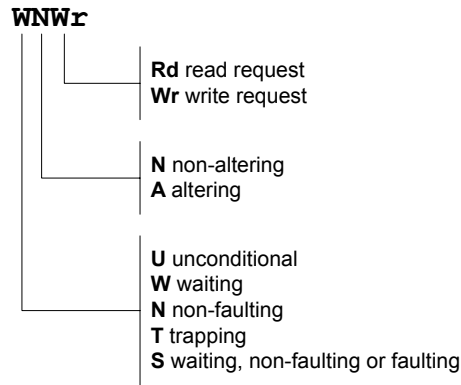


Figure 2: Notation of synchronizing memory operations

All memory operations, regardless of the classification made in Figure 1, can be further catalogued into *altering* and *non-altering* operations. While the former modify the *full/empty* bit after a successful synchronizing event, the latter do not touch this bit in any case. According to this distinction, ordinary memory operations fall into the *unconditional non-altering* category.

The following table shows the notation used for each variant of memory operation and its behavior in the case of a synchronization miss. The notation is further explained in Figure 2.

Table 1: Notation of synchronized operations

Notation	Semantics	Behavior on a synchronization miss
UNRd	Unconditional non-altering read	Never miss
UNWr	Unconditional non-altering write	
UARd	Unconditional altering read	
UAWr	Unconditional altering write	
WNRd	Waiting and non-altering read from <i>full</i>	Placed on the list of pending requests until resolved
WNWr	Waiting and non-altering write to <i>empty</i>	
WARd	Waiting and altering read from <i>full</i>	
WAWr	Waiting and altering write to <i>empty</i>	
NNRd	Non-faulting and non-altering read from <i>full</i>	Silently discarded
NNWr	Non-faulting and non-altering write to <i>empty</i>	
NARd	Non-faulting and altering read from <i>full</i>	
NAWr	Non-faulting and altering write to <i>empty</i>	
TNRd	Faulting and non-altering read from <i>full</i>	Signal trap
TNWr	Faulting and non-altering write to <i>empty</i>	
TARd	Faulting and altering read from <i>full</i>	
TAWr	Faulting and altering write from <i>empty</i>	

## 3. Architectural support for fine-grain synchronization

---

### 3.1. Related work

#### 3.1.1. The Alewife machine

The MIT Alewife machine is a cache-coherent shared memory multiprocessor (see [2] and [4]) with *non-uniform memory access* (NUMA). Although it is internally implemented with an efficient message-passing mechanism, it provides an abstraction of a global shared memory to programmers. The most relevant part of its nodes regarding coherency and synchronization protocols is the *communication and memory management unit* (CMMU), which deals with memory requests from the processor and determines whether a remote access is needed, managing also cache fills and replacements. Cache coherency is achieved through `LimitLESS`, a software extended *directory-based* protocol. The home node of a memory line is responsible for the coordination of all coherence operations for that line.

Support for fine-grain synchronization in Alewife includes full/empty bits for each 32-bit data word and fast user-level messages. Colored `load` and `store` instructions are used to access synchronization bits. The *alternate space indicator* (ASI) distinguishes each of these instructions. Full/empty bits are stored in the bottom four bits of the coherency directory entry (at the memory) and as an extra field in the cache tags (at the cache), so they do not affect DRAM architecture nor network data widths. The Alewife architecture also defines language extensions to support both `J-` and `L-`structures. A specific programming language, namely `Semi-C`<sup>1</sup>, has been defined for this purpose [42].

The aim is that a successful synchronization operation does not incur much overhead with respect to a normal `load` or `store`. In the ideal case, the cost of both types of operations is expected to be the same. This is possible because full/empty bits can be accessed simultaneously with the data they refer to. The cost of a failed synchronization operation depends much on the specific hardware support for synchronization. The overhead of software-supported synchronization operations is expected to be much higher than their hardware counterparts. However, Alewife minimizes this by rapidly switching between threads on a failed synchronization attempt or a cache miss, requiring the use of lockup-free caches.

Handling failed synchronization operations in software has the advantage of being less complex in terms of hardware and more flexible. The basis of Alewife support for fine-grain synchronization is that, as synchronization operations are most probably successful, overhead due to such failures is not expected to notably reduce overall system performance.

---

<sup>1</sup> Semi-C is an extension of the C language that can handle parallel programming constructs.

### 3.1.2. The StarT-NG machine

StarT-NG, an improved version of the StarT machine [9], is a high-performance *message passing* architecture in which each node consists of a commercial symmetric multiprocessor (SMP) that can be configured with up to 3 processors, which are connected to the main memory by a data crossbar. At least one *network interface unit* is present in each node, allowing communicating with a *network router*, which is implemented in a proprietary chip [17].

A low-latency high-bandwidth network interconnects every node in the system. StarT-NG also supports cache-coherent global shared memory. In this case, one processor on each site is used to implement the shared memory model. This functionality can be disabled when shared memory is not needed.

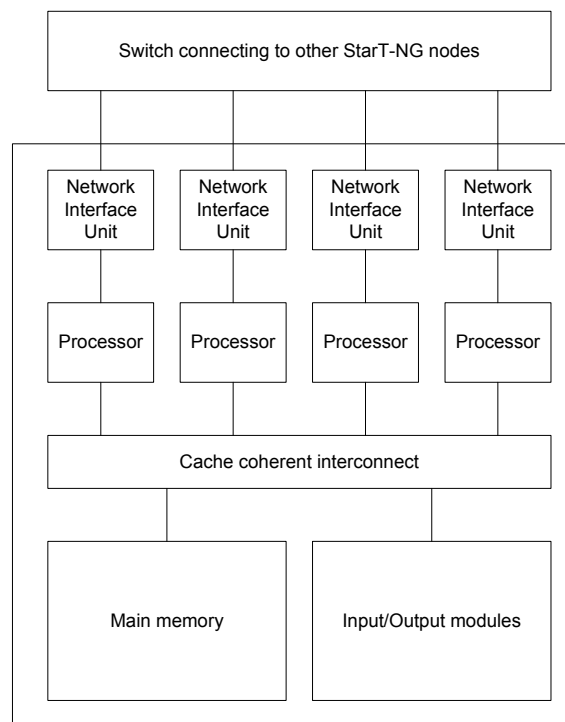
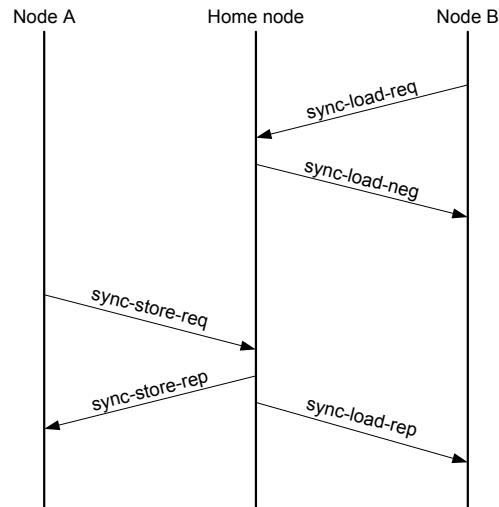


Figure 3: Architecture of a StarT-NG node [27]

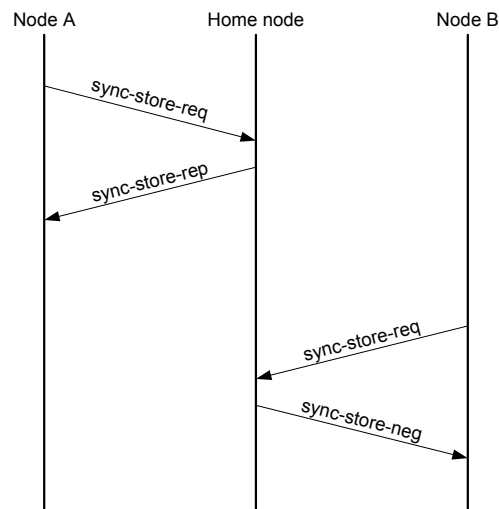
Coherence protocols in StarT-NG are fully implemented in software. As a consequence, the choice of protocols and configuration of the shared memory is notably flexible. The performance of several coherence models has been evaluated. Particularly relevant to this work is the study made in [75], which introduces a cache coherence protocol with support for fine-grained data structures. These data structures are known as I-structures [75].

According to the results of this study, performance improvements in an integrated coherence protocol are two-fold. First, the write-once behavior of I-structures allows writes to be performed without the exclusive ownership of the respective cache line. Once a write has been carried out, stale data in other caches is identified because its *full/empty* bit is unset. In a *directory-based* protocol, a synchronized load in a remote location will find the *full/empty* bit

unset and forward the request to the proper node. This behavior is illustrated in Figure 4, where two nodes (namely, A and B) share a copy of a block on which they perform different operations.



**Scenario 1** Initially, both nodes A and B have a copy of the cache line in the shared state. A synchronized store operation is performed by node A without the *exclusive ownership* of the cache block, which is consequently kept in the *shared* state during the whole process. Pending synchronized loads from node B to the affected slot are resumed after the store is performed.



**Scenario 2** Initially, both nodes A and B have a copy of the cache line in the shared state. A synchronized store operation is successfully performed by node A without the *exclusive ownership* of the cache block. If node B issues a synchronized store, the request will be rejected by the home node after finding the *full-empty* bit set.

Figure 4: Two sample scenarios of synchronized loads and stores

As stated in [74], another advantage of a coherence protocol integrated with fine-grain synchronization is the efficiency in the management of pending requests by reducing the number of transactions needed to perform some particular operations. As an example, a synchronized

load in traditional coherence protocols usually requires the requesting node to obtain the exclusive ownership of the affected block in order to set the *full/empty* bit to the empty state.

### 3.2. Proposed architecture

In a multiprocessor system providing fine-grain synchronization, each shared memory word is tagged with a *full/empty* bit that indicates the synchronization state of the referred memory location. Assuming that a memory word is 32-bit long, this implies an overhead of just 3%. Although many variations exist when implementing this in hardware, the structure of shared memory is conceptually as shown in Figure 5.

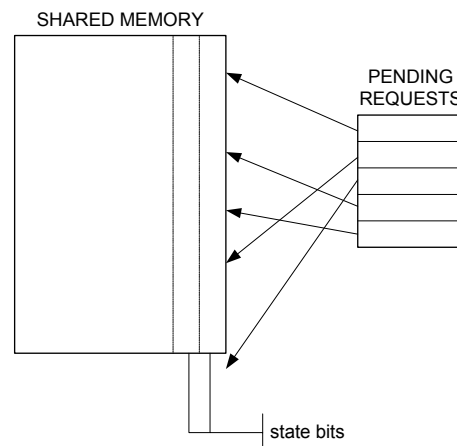


Figure 5: Logical structure of shared memory

Figure 5 shows that each shared memory location has three logical parts, namely:

- i) the shared data itself.
- ii) *state* bits. The *full/empty* bit is placed within the *state* bits. This bit is set to 1 if the corresponding memory location has already been written by a processor and thus contains valid data. If the architecture has cache support other *state* bits such as the *dirty* bit may exist. The *dirty* bit is set if the memory location is not up-to-date, indicating that it has been modified in a remote node.
- iii) the list of pending memory requests. Synchronization misses fired by conditional waiting memory operations are placed in this list. When an appropriate synchronizing operation is performed, the relevant pending requests stored in this list are resumed. If the architecture has cache support, the list of pending memory requests also stores ordinary cache misses. The difference between both types of misses is basically that synchronization misses store additional information, such as the accessed slot's index in the corresponding cache block. These differences are further explained later in this section.

Note that fine-grain synchronization is described here only for shared memory locations. In the presented architecture, the local memory in each processing node does not make use of *full/empty* bits. With this consideration, the memory map of the system seen by each processor is similar to the one sketched in Figure 6.




Figure 6: Memory map for each processing node

Fine-grain synchronization is implemented by atomic *test-and-set* operations. These operations modify the *full/empty* condition bit in the processor's condition bits register<sup>2</sup>. Note that the condition bit is changed regardless of the particular variant of synchronization operation; no matter it is altering and/or trapping.

As stated before, many implementation alternatives are possible. State bits may be stored in the coherence directory entry in the case of a *directory-based* protocol, such as the one implemented in *Alewife*. A proposed structure for a cache supporting *fine-grain* synchronization is depicted in Figure 7.

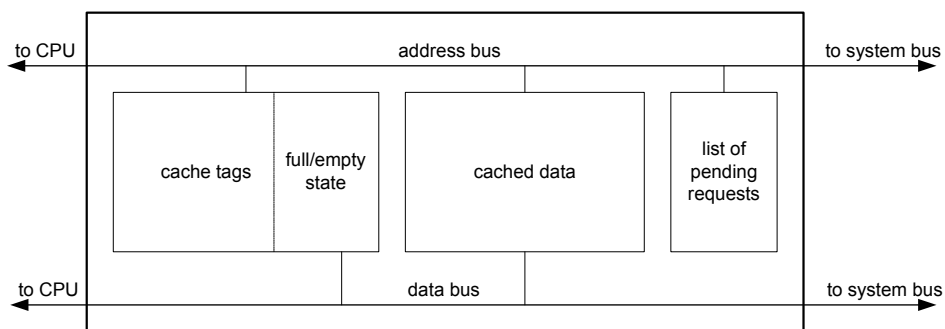


Figure 7: Organization of a cache supporting fine-grain synchronization

When a memory word is cached, its *full/empty* bit must also be stored at the cache side. As a consequence, not only data has to be kept coherent, but also *full/empty* bits. In a system with cache support, an efficient option is to store the *full/empty* bit as an extra field in the cache tag, allowing checking the synchronization state in the same step as the cache lookup. The coherence protocol has then two logical parts, one for the data and another for the synchronization bit.

<sup>2</sup> In *Sparcle*, for instance, the *full/empty* condition bit is stored in the condition bit #0 (see [73]).

Our design assumes that the smallest synchronizing element is a word. As a cache line is usually longer, it may contain multiple elements, including both synchronized and ordinary data (see Figure 8). A tag for a cache line includes the *full/empty* bits for all the synchronized words that are stored in that line. As directory states are maintained at cache line level, this complicates the maintenance of pending memory requests. Effectively, while a *dirty* bit refers to a complete cache line, a *full/empty* bit refers to a single word in a cache line.




Figure 8: Cache line containing both ordinary and synchronized data

In the proposed architecture, lists of pending requests are maintained in hardware at the cache level, more concretely in the *miss status holding registers* (MSHR). With this assumption, waiting memory operations require the architecture to have cache support. However, if cache support is not available, the behavior of waiting operations can be implemented in software by using faulting conditional operations instead. The system kernel is then responsible for maintaining the list of pending requests [39]. In the case of a *directory-based* coherence protocol, an alternative is to store the pending requests as a separate field in the directory entries.

Some modifications have to be made to the cache architecture in case synchronization misses are to be kept in MSHR. More concretely, MSHR in traditional *lockup-free* caches store the information listed in Table 2 (see [46] for a more detailed description). In order to store synchronization misses in these registers, two more fields have to be added containing the slot's index accessed by the operation and the specific variant of synchronized operation that will be performed.

Table 2: Relevant information stored in ordinary MSHR registers [46]

Field	Semantics
Cache buffer address	Location where data retrieved from memory is stored
Input request address	Address of the requested data in main memory
Identification tags	Each request is marked with a unique identification label
Send-to-CPU flags	If set, returning memory data is sent to CPU
In-input stack	Data can be directly read from input stack if indicated
Number of blocks	Number of received words for a block
Valid flag	When all words have been received the register is freed
Obsolete flag	Data is not valid for cache update, so it is disposed

A complete description of a cache coherence mechanism includes the states, the transition rules, the protocol message specification and the description of cache line organization and memory management of pending requests. Other design issues are dealing with conflicting and/or merging synchronization misses, as well as ordering of misses from the same processor.



Our design is based on a multiprocessor system with the following assumptions:

- the CPU implements *out-of-order* execution of instructions,
- each processing node has a *miss-under-miss lockup-free* cache, supporting multiple outstanding memory requests,
- the smallest synchronized data element is a word; this statement does not imply a loss of generality, as the extension of the presented design to other data sizes is straightforward.

### 3.3. Cache coherence

In a multiprocessor system, cache memory local to each processing node can be used to speed up memory operations. It is necessary to keep the caches in a state of coherence by ensuring that modifications to data that is resident in a cache are seen in the rest of the nodes that share a copy of the data. This can be achieved in several ways, which may depend on the particular system architecture. In bus-based systems, for instance, cache coherence is implemented by a snooping mechanism, where each cache is continuously monitoring the system bus and updating its state according to the relevant transactions seen on the bus. On the contrary, mesh network-based multiprocessors use a directory structure to ensure cache coherence. In these systems, each location in the shared memory is associated with a directory entry that keeps track of the caches that have a copy of the referred location. Both *snoopy* and *directory-based* mechanisms can be further classified into *invalidation* and *update* protocols. In the former case, when a cache modifies shared data, all other copies are set as invalid. Update protocols change copies in all caches to the new value instead of marking them as invalid.

The performance of multiprocessor systems is partially limited by cache misses and node interconnection traffic. Consequently, cache coherence mechanisms play an important role in solving the problems associated with shared data. Another performance issue is the overhead imposed by synchronizing data operations. In the case of systems that provide fine-grain synchronization, this overhead is due to the fact that synchronization is implemented as a separate layer over the cache coherence protocol. Indeed, bandwidth demand can be reduced if no data is sent in a synchronization miss. This behavior requires the integration of cache coherence and fine-grain synchronization mechanisms. It is important to remark, however, that both mechanisms are conceptually independent. This means that synchronizing operations can be implemented in machines without cache support and vice-versa.

One of the main objectives of this project is to define a coherence protocol that integrates fine-grain synchronization. This will be done for both *snoopy* and *directory-based* protocols. An event-driven simulator, namely *RSIM*, is used in order to verify and measure the performance of our design. As this simulation platform does not integrate synchronization at the cache coherence level, modifications in its source code are needed.

In the proposed architecture, failing synchronizing events are resolved in hardware. The following architecture requirements must be considered in order to integrate synchronization and cache coherency. Note that most of the hardware needed is usually already available in modern multiprocessor systems.

- i) each memory word has to be associated with a full/empty bit; as in Alewife, this state information can be stored in the coherency directory entry,
- ii) at the cache side, state information is stored as an additional field in the cache tags; a lookup-free cache is needed in order to allow non-blocking loads and stores,
- iii) the cache controller not only has to deal with coherency misses, but also with full/empty state misses; synchronization is thus integrated with cache coherency operations, as opposed to Alewife, in which the synchronization protocol is implemented separately from the cache coherency system.

This approach can be extended to the processor registers by adding a full/empty tag to them. This would allow an efficient execution of synchronization operations from simultaneous threads on the registers. However, additional modifications are needed in the processor architecture to implement this feature.

In order to evaluate the performance improvement of this novel architecture with respect to existing approaches, appropriate workloads must be tested on the devised machine. A challenge task is to find suitable applications that show these results in a meaningful way, so that the effects of the synchronization overhead such as the cost of additional state storage, execution latency or extra network traffic can be studied in detail.

## 4. Integration with snoopy protocols

---




Figure 9: Bus-based system architecture

We consider a bus-based system such that depicted in Figure 9. Note that even though each memory address has conceptually a list of pending operations for that address, at hardware level the lists are distributed between all the processing nodes. The management of deferred lists will be explained later in this section. The description made here is based on the `MESI` protocol, also known as `Illinois` protocol. It is a four-state write-back invalidation protocol with the following state semantics [30]:

- `modified` - this cache has the only valid copy of the block; the location in main memory is invalid.
- `exclusive clean` - this is the only cache that has a copy of the block; the copy in main memory is up-to-date. A signal `s` is available to the controller in order to determine on a `BusRd` if any other cache currently holds the data.
- `shared` - the block is present in an unmodified state in this cache, main memory is up-to-date and zero or more caches may also have a shared copy.
- `invalid` - the block does not have valid data.

The state diagram corresponding to the `MESI` protocol without fine-grain synchronization support is shown in Figure 10.




Figure 10: MESI coherence protocol

In the figure above, we use the notation  $A/B$ , where  $A$  indicates an observed event and  $B$  is an event generated as a consequence of  $A$ . Dashed lines show state transitions due to observed bus transactions, while continuous lines indicate state transitions due to local processor actions. Finally, the notation  $Flush'$  means that data is supplied only by the corresponding cache. Note that this diagram does not consider transient states used for bus acquisition.

The transitions needed to integrate fine-grain synchronization in  $MESI$  are sketched in Figure 11, in which the *full/empty* state of the accessed word is explicitly indicated by splitting the ordinary  $MESI$  states into two groups. The transactions not shown in this figure are not relevant for the corresponding state and do not cause any transition in the receiving node. The notation is the same as in the previous figure, and as it can be appreciated below, no new states are preliminarily required so as to integrate fine-grain synchronization in the coherence protocol.

The description made here considers only waiting non-altering reads and waiting altering writes. Altering reads can be achieved by issuing non-altering reads in combination with an operation that clears the *full/empty* bit without retrieving data. This operation is named *unconditional altering clear*, or  $PrUACl$  according to the nomenclature previously described.  $PrUACl$  operates on a *full/empty* bit without accessing or altering the data corresponding to that state bit.

Clearing of *full/empty* bits is necessary in order to reuse synchronized memory locations (a more detailed description is made in [46]). While a  $PrUARd$  could be used for this end, the  $PrUACl$  instruction completes faster, as it alters the *full/empty* bit without actually reading data from the corresponding location. For this reason,  $PrUACl$  can be seen as an optimized memory instruction.



Waiting operations constitute the most complex sort of synchronizing operations, as they require additional hardware in order to manage deferred list and resume pending synchronization requests. The behaviour of other types of memory operations is a simplified version of waiting operations. Most of the transitions depicted in Figure 11 are identical in the rest of the cases, with the only different being the behaviour when a synchronization miss is detected. Instead of being added to the list of pending requests, other variants of missing operations either fire an exception or are silently discarded.

Two additional bus transactions are needed in order to integrate fine-grain synchronization in the MESI protocol. A detailed description of these bus transactions is presented in Table 3. Coherence of *full/empty* bits is ensured precisely by these two bus transactions (BusSWr and BusSCL).

Table 3: Additional bus transactions in the MESI protocol

Bus transaction	Description
BusSWr	A node has performed an altering waiting write. The effect of this operation in observing nodes is to set the <i>full/empty</i> bit of the referring memory location and resume the relevant pending requests. Resuming of pending requests is further explained in section 4.2.
BusSCL	A node has performed an <i>altering read</i> or an <i>unconditional clear</i> operation. The effect of this operation in observing nodes is to clear the <i>full/empty</i> bit of the referring memory location, thus making it reusable.




Figure 12: MESI protocol integrated with fine-grain synchronization (implicit full/empty states)

A new signal (referred as *c* in Figure 12) is introduced in order to determine whether a synchronized operation misses. This bus signal will be called *shared-word* signal, as it indicates whether any other node is sharing the referring word. The *shared-word* signal can be implemented as a wired-OR

controller line, which is asserted by each cache that holds a copy of the relevant word with the *full/empty* bit set. According to this notation, a waiting read request written in the form  $\text{PrWNRd}(C)$  successfully performs, while an event of the form  $\text{PrWNRd}(\bar{C})$  causes a synchronization miss. Note also that, as each cache line may contain several synchronized data words, it is necessary to specify the specific word to which the synchronized operation is to be performed. Consequently, a negated synchronization signal ( $\bar{C}$ ) causes a requesting read to be appended to the list of pending operations whereas a requesting write to be performed successfully. If the synchronization signal is otherwise asserted ( $C$ ), then a synchronized read is completed successfully whereas a requesting write is suspended.

In addition to the *shared-word* signal already introduced, three more wired-OR signals are required for the protocol to operate correctly, as described in [30]. The first signal (named *s*) is asserted if any processor different than the requesting processor has a copy of the cache line. The second signal is asserted if any cache has the block in a dirty state. This signal modifies the meaning of the former in the sense that an existing copy of a cache line has been modified and then all the copies in other nodes are invalid. A third signal is necessary in order to indicate whether all the caches have completed their snoop, that means, if it is reliable to read the value of the first two signals.

Figure 12 shows a more compact state transition specification in which information about the *full/empty* state of the accessed word is implicit. Instead, the value of the *C* line or the *full/empty* bit is specified as a required condition between parentheses. Figure 11 and Figure 12 do not consider neither transient states needed for bus acquisition nor the effects due to real signal delays.

#### 4.1. Mapping between processor instructions and bus transactions

When a processing node issues a memory operation, the cache located at that node first interprets the request and, in case of a miss, it later translates the operation into one or more bus transactions. The correspondence between the different processor instructions and the memory requests seen on the bus is shown in Table 4. The same notation as in Figure 2 is used.

Table 4: Correspondence between processor instructions and memory requests

Request from processor	Bus transaction
PrUNRd	BusRd (ordinary read)
PrUNWr	BusWr (ordinary write)
PrUARd	BusRd + BusSCl
PrUAWr	BusAWr <sup>3</sup>
PrSNRd	BusRd(C) <sup>4</sup>
PrSNWr	BusWr(C)
PrSARd	BusRd(C) + BusSCl
PrSAWr	BusSWr(C)

<sup>3</sup> Neither unconditional altering writes nor conditional non-altering writes are considered in the protocol specification.

<sup>4</sup> The bus transaction BusRd is in this case used in combination with the shared-copy signal.

As seen on Table 4, unconditional non-altering read and write requests generate ordinary read and write transactions on the bus. On the contrary, an unconditional altering read requires a `BusRd` transaction followed by a `BusSCL` transaction. Effectively, apart from retrieving the data from the corresponding memory location, a `PrUARd` request also clears the *full/empty* state bit of the referring location. This is performed by `BusSCL`, which does not access nor modifies the data but only the *full/empty* bit. It is important to observe that an unconditional altering read cannot be performed by just a `BusSCL` transaction, as it just alters the *full/empty* bit without retrieving any data. The last unconditional operation, `PrUAWr`, generates a specific bus transaction, namely `BusAWr`, which unconditionally sets the *full/empty* bit after writing the corresponding data to the accessed memory location.

It is inferred from Table 4 that the behavior of all conditional memory operations depends on the value of the `shared-word` bus signal<sup>5</sup>. A conditional non-altering read, for instance, generates an ordinary read bus transaction after checking whether the `shared-bus` signal is asserted. A conditional altering read generates a `BusSCL` transaction in addition to the ordinary read transaction. Finally, a conditional altering write causes a `BusSWr` transaction to be initiated on the bus. This transaction sets the *full/empty* bit after writing the corresponding data to the referred memory location.




Figure 13: Sample scenario of mapping between processor instructions and bus transactions

Note that all synchronized operations generate the same bus transactions regardless of their particular type (waiting, non-faulting or faulting). The difference resides in the behavior when a synchronization miss is detected and not in the bus transactions issued as a consequence of the request. A sample scenario is shown in the figure below.

<sup>5</sup> This behavior is systematically specified in section 4.3.



## 4.2. Management of pending requests

Each processing node keeps a local deferred list. This list holds both ordinary *presence* misses and synchronization misses. It is possible also for both types of misses to happen for a single access. In this case, not only the accessed line is not present in the cache, but also the synchronization state is not met at the remote location where the copy of the word is held. After a relevant *full/empty* bit change is detected, any operation that matches a required synchronization state is resumed at the appropriate processing node.

Table 5 shows how the management of the deferred list local to a node is done. Concretely, the table specifies the action taken when a given request is received with respect to a pending request already present in the list of deferred operations<sup>6</sup>. A *C* indicates that both requests are conflicting and thus need to be kept separated into two different entries, always ensuring that local order is maintained. On the contrary, an *M* means that both requests can be merged and thus treated as a sole request from the point of view of memory accesses.

Table 5: Management of coalescing requests

		Pending request (already in MSHR)							
		PrUNRd	PrUARd	PrWNRd	PrWARd	PrNNRd	PrNARd	PrTNRd	PrTARd
Incoming request	PrUNRd	M	M	M	M	M	M	M	M
	PrUARd	M	M	M	M	M	M	M	M
	PrWNRd	M	C	M	C	M	C	M	C
	PrWARd	M	C	M	C	M	C	M	C
	PrNNRd	M	C	M	C	M	C	M	C
	PrNARd	M	C	M	C	M	C	M	C
	PrTNRd	M	C	M	C	M	C	M	C
	PrTARd	M	C	M	C	M	C	M	C

As a rule of thumb, a pending write is conflicting with any incoming request, so it can never be merged and requires a separate entry in the list of pending requests<sup>7</sup>. As they are always conflicting, all write requests have been excluded from Table 5. Another important observation is that pending altering reads can only be merged with unconditional operations. Additionally, all non-altering pending read request can be coalesced with any other incoming read request.

Apart from coalescing of requests, it is also crucial to specify how resuming of pending requests is done. As explained at the beginning of this section, coherence of *full/empty* state bits is ensured by

<sup>6</sup> The simulation model only considers PrSNRd, PrSAWr and PrUACl instructions.

<sup>7</sup> It could be possible to make read requests be satisfied by pending writes to the same location. However, this introduces extra complexity in the memory unit in order to meet the consistency model. A write request cannot be satisfied by a pending read request in any case.

proper bus transactions, to be precise,  $\text{BusSWr}$  and  $\text{BusSCl}$ . This means that all caches that have pending requests for a given memory location will know when the synchronization condition is met by snooping into the bus and waiting for a  $\text{BusSWr}$  or a  $\text{BusSCl}$  transaction. When such transaction is noticed, a comparator checks if there is an entry in any  $\text{MSHR}$  matching the received bus transaction. In this case, action is taken so as to resume the pending request.

Due to this feature, it is possible for a cache to have pending requests for a memory location that is not cached or is cached in an *invalid* state. The location will be brought again into the cache when the synchronization miss is solved. The ability of replacing cache lines that have pending requests allows efficient management and resuming of pending requests with minimum risk of saturating the cache hierarchy.

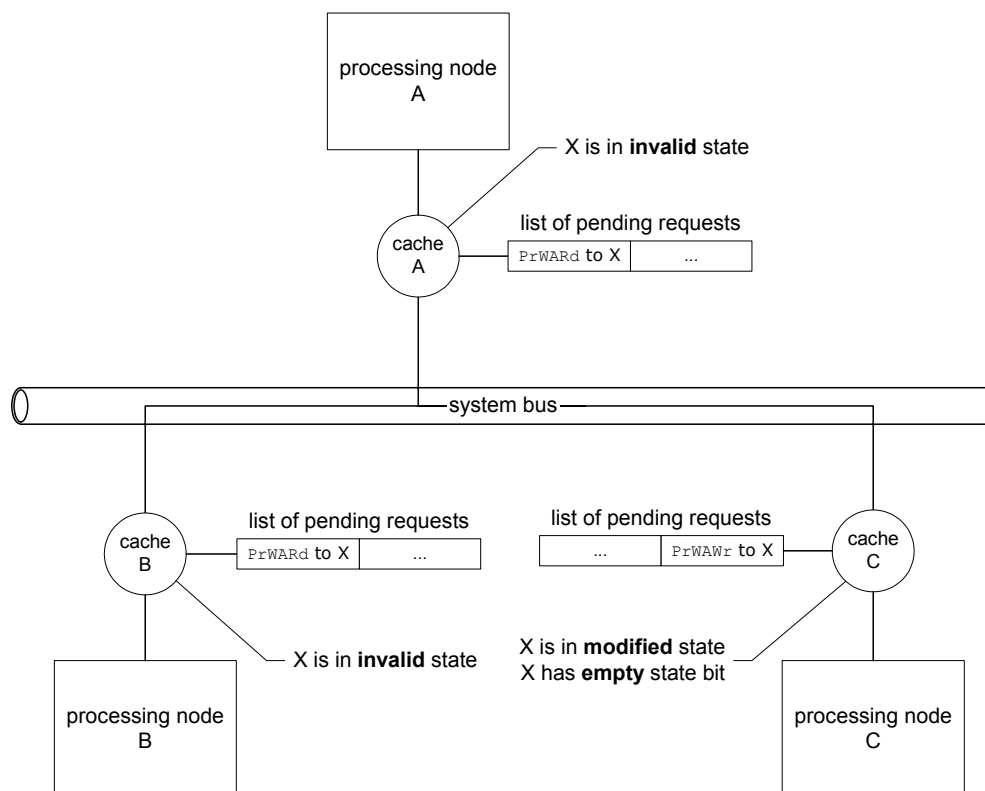


Figure 14: Resuming of pending requests

A representative scenario is shown in Figure 14, in which three nodes have pending requests to a location ( $X$ ) in their  $\text{MSHR}$ . While nodes  $A$  and  $B$  have invalid copies in their caches, node  $C$  has the exclusive ownership of the referred location, whose *full/empty* state bit is unset. After node  $C$  successfully performs a conditional altering write to location  $X$ , this event is notified on the bus by a  $\text{BusSWr}$  transaction. This transaction informs nodes  $A$  and  $B$  that they can resume the pending request to location  $X$ , which happens to be a conditional altering read. As a consequence, only one of these nodes will be able to successfully issue the operation at this point. This is imposed by bus order. For instance, if node  $B$  gets the bus ownership before node  $A$ , the pending request from the former will be resumed and the operation at node  $A$  will stay pending in the  $\text{MSHR}$ .

### 4.3. Transition rules

A detailed explanation of the new transition rules from each coherence state is presented in the following sections. A description in the form of C-styled *pseudo-code* is also presented in each case. Observe that, as with the ordinary coherence misses, the ordering of synchronization misses from different processors is imposed by bus order.

#### 4.3.1. Invalid state

```

SWITCH(incomingRequest) {
  CASE PrUNRd: send(BusRd);
                IF (S) {
                    flushFromOtherCache();
                    nextState = shared;
                } ELSE {
                    readFromMemory();
                    nextState = exclusive;
                }
                BREAK;
  CASE PrUNWr: send(BusRdX);
                nextState = modified;
                BREAK;
  CASE PrWNRd: send(BusRd);
                IF (S && C) {
                    flushFromOtherCache();
                    nextState = shared;
                } ELSE IF (!S && C) {
                    readFromMemory();
                    nextState = exclusive;
                } ELSE {
                    addToDeferredList(); // Wait.
                }
                BREAK;
  CASE PrWAWr: send(BusSWr);
                IF (S && !C) {
                    writeToBus();
                    nextState = shared; // To be evaluated at simulation.
                } ELSE IF (!S && !C) {
                    writeToCache();
                    nextState = modified;
                } ELSE {
                    addToDeferredList(); // Wait.
                }
                BREAK;
  CASE PrUACl: IF (C) {
                    send(BusSCl);
                    nextState = invalid;
                }
                BREAK;
}

```

A successful conditional waiting read request from the local processor (PrWNRd) leads either to the *exclusive-clean* state (if no other cache holds a copy of the block) or to the *shared* state (if more caches have a copy of the accessed block). In any case, a BusRd transaction is generated in order to fetch the data from the corresponding cache or shared memory location. However, if the synchronization condition is not met ( $\bar{C}$ ), then the request is appended to the local deferred list and the state is not changed. This occurs when neither the caches nor the shared memory assert the c line.

*Cache-to-cache* transfers are needed when data is modified in one or more caches and the copy in the shared memory is stale. An alternative is to flush the modified data back to memory and then to the node that requested access, but this approach is obviously slower than the former.

A successful waiting write request from the local processor (`PrWAWr`) leads either to the *modified* state (if no other cache holds a copy of the block) or to the *shared* state (if more caches have a copy of the block). This implies a performance improvement since the next successful synchronized operation to the same cache slot will necessarily be a read and a state transaction will be saved<sup>8</sup>. If the synchronization condition is not met (the line `C` is asserted), then the operation is suspended.

A `PrUACl` request generates a `BusSCL` transaction but does not load the block into cache. This is a design alternative and will be evaluated at the simulation stage of this study.

#### 4.3.2. Modified state

If the *full/empty* bit is set, a conditional waiting read (`PrWNRd`) retrieves the data from the local cache and generates no bus transaction. Otherwise, the request is appended to the local deferred list.

A conditional waiting write (`PrWAWr`) fails if the `C` line is asserted and sets the *full/empty* bit otherwise. In the latter, a `BusSWr` transaction is generated and the relevant pending requests in the local deferred list are resumed. The effect of a `BusSWr` in the other caches is precisely to set the *full/empty* bit and to analyze their deferred list so as to resume the relevant pending requests.

A `PrUACl` request generates a `BusSCL` transaction and unsets the *full/empty* bit. This transaction does not flush the block from cache. This is a design alternative and will be evaluated at the simulation stage of this study.

```
SWITCH(incomingRequest) {
  // Processor requests
  CASE PrUNRd: readFromCache();
               nextState = modified;
               BREAK;
  CASE PrUNWr: writeToCache();
               nextState = modified;
               BREAK;
  CASE PrWNRd: IF (full) {
                 readFromCache();
                 nextState = modified;
               } ELSE {
                 addToDeferredList();
                 nextState = modified;
               }
               BREAK;
  CASE PrWAWr: send(BusSWr);
               IF (empty) {
                 writeToCache();
                 resumePendingReqs();
                 nextState = modified;
               } ELSE {
                 addToDeferredList();
                 nextState = modified;
               }
               BREAK;
  CASE PrUACl: IF (full) {
                 unsetFE();
                 nextState = modified;
               }
               BREAK;
  // Bus signals
  CASE BusRd: flush();
               nextState = shared;
```

<sup>8</sup> If a transition to the *modified* state is performed as in the ordinary MESI protocol, an additional `BusRdX` transaction is then required on the bus.

```

        BREAK;
CASE BusRdX: flush();
             nextState = invalid;
             BREAK;
CASE BusSWr: IF (empty) {
              writeToCache();
              resumePendingReqs();
              nextState = shared;
            }
             BREAK;
CASE BusScl: IF (full) {
              unsetFE();
              nextState = shared;
            }
             BREAK;
}

```

### 4.3.3. Exclusive-clean state

```

SWITCH(incomingRequest) {
  // Processor requests
CASE PrUNRd: readFromCache();
             nextState = exclusive;
             BREAK;
CASE PrUNWr: writeToCache();
             nextState = modified;
             BREAK;
CASE PrWNRd: IF (full) {
              readFromCache();
              nextState = exclusive;
            } ELSE {
              addToDeferredList();
              nextState = exclusive;
            }
             BREAK;
CASE PrWAWr: send(BusSWr);
             IF (empty) {
               writeToCache();
               resumePendingReqs();
               nextState = shared; // To be evaluated at simulation.
             } ELSE {
               addToDeferredList();
               nextState = exclusive;
             }
             BREAK;
CASE PrUACL: IF (full) {
              unsetFE();
              nextState = modified;
            }
             BREAK;
  // Bus signals
CASE BusRd:  flush();
             nextState = shared;
             BREAK;
CASE BusRdX: flush();
             nextState = invalid;
             BREAK;
CASE BusSWr: IF (empty) {
              writeToCache();
              resumePendingReqs();
              nextState = shared;
            }
             BREAK;
CASE BusScl: IF (full) {
              unsetFE();
              nextState = shared;
            }
             BREAK;
}

```

As no other caches hold a copy of this block, a synchronized read (PrWNRd) leads to the same coherence state.

#### 4.3.4. Shared state

```

SWITCH(incomingRequest) {
  // Processor requests
  CASE PrUNRd: readFromCache();
               nextState = shared;
               BREAK;
  CASE PrUNWr: send(BusRdX);
               writeToCache();
               nextState = modified;
               BREAK;
  CASE PrWNRd: IF (full) {
                 readFromCache();
                 nextState = shared;
               } ELSE {
                 addToDeferredList();
                 nextState = shared;
               }
               BREAK;
  CASE PrWAWr: send(BusSWr);
               IF (empty) {
                 writeToCache();
                 resumePendingReqs();
                 nextState = shared; // To be evaluated at simulation.
               } ELSE {
                 addToDeferredList();
                 nextState = shared;
               }
               BREAK;
  CASE PrUAcl: IF (full) {
                 unsetFE();
                 send(BusScl);
                 nextState = shared;
               }
               BREAK;
  // Bus signals
  CASE BusRd:  flush();
               nextState = shared;
               BREAK;
  CASE BusRdX: flush();
               nextState = invalid;
               BREAK;
  CASE BusSWr: IF (empty) {
                 writeToCache();
                 resumePendingReqs();
                 nextState = shared;
               }
               BREAK;
  CASE BusScl: IF (full) {
                 unsetFE();
                 nextState = shared;
               }
               BREAK;
}

```

The same rules apply as for the modified state, with the only exception of the `BusSWr` and `BusScl` bus transactions, which do not cause a state transition in this case.

#### 4.4. Summary

A bus based coherence protocol with fine-grain synchronization support has been introduced. A systematic protocol description is made in the form of state diagrams and *pseudo-code*. Although this implementation considers only waiting non-altering reads and waiting altering writes, the behavior of other memory operations is derived in a straightforward manner, as it is a simplified version of the former.

One of the base ideas of the protocol is that *full/empty* state bit coherence is maintained by bus transactions defined for this purpose, namely `BusSWr` and `BusScl`. An additional bus signal called

`shared-word` is also introduced in order to implement the conditional behavior of synchronizing operations.

A drawback of integrating fine-grain synchronization support at the cache level is the complexity of managing pending synchronization requests. Rules for coalescing and resuming synchronizing requests have been explained in detail. It is expected that this supplementary complexity does not translate in excessive hardware overhead, as most of the required hardware is already present in modern multiprocessors. Consequently, application software making use of synchronizing memory operations will likely experience a noteworthy performance improvement without the need of extensive hardware deployment.

## 5. Integration with directory-based protocols

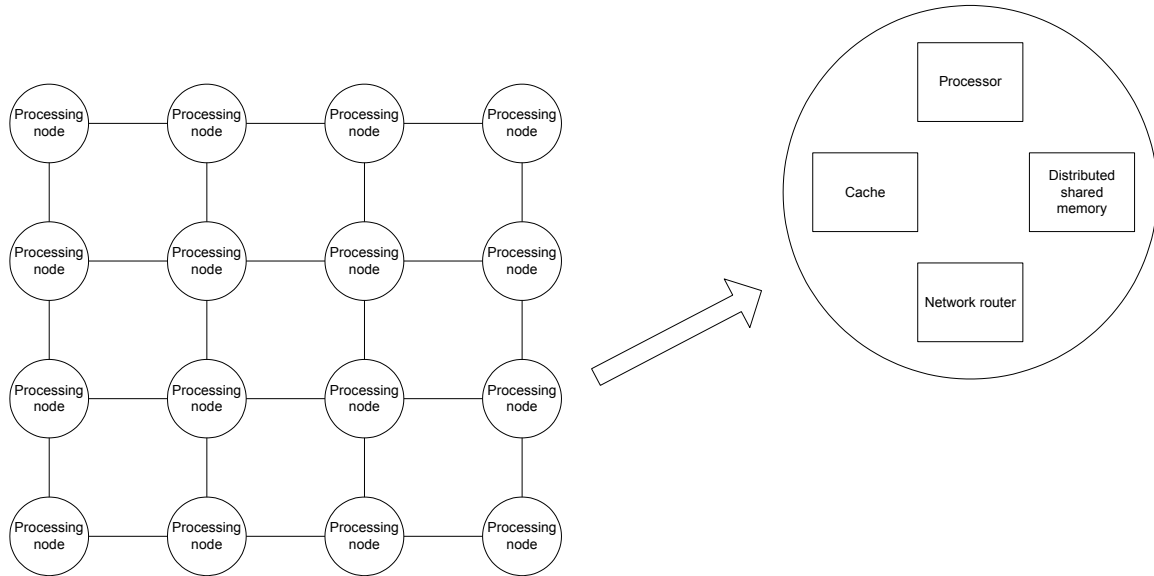


Figure 15: Mesh network-based architecture

In a network-based system, such the one shown in Figure 15, each shared memory block has a directory entry that lists the nodes that have a cached copy of the data. *Full/empty* bits are stored as an extra field in the coherence directory entry. *Point-to-point* messages are used to keep the directory up-to-date and to request permission for a load or a store to a particular location.




Figure 16: Alewife's coherence protocol state diagram



The description made here is based on `Alewife`'s coherence protocol [24]. Our model considers a limited directory protocol, thus restricting the amount of simultaneous copies of a memory block. The following states are defined in `Alewife`'s coherence protocol:

- `Read-Only`: One or more caches have a read-only copy of the block.
- `Read-Write`: Only one cache has a read-write copy of the block.
- `Read Transaction`: Cache is holding a read request (update in progress).
- `Write Transaction`: Cache is holding a write request (invalidation in progress).

The state diagram corresponding to this protocol is shown in Figure 16. The semantics of the transitions depicted in this figure are resumed in Table 6 [50].

Table 6: Semantics of the transitions in the directory-based protocol

Label	Input message	Output message
1	$i \rightarrow \text{RREQ}$	$\text{RDATA} \rightarrow i$
	$i \rightarrow \text{FETCH}$	$\text{RDATA} \rightarrow i$
2	$i \rightarrow \text{WREQ}$	$\text{WDATA} \rightarrow i$
	$i \rightarrow \text{MREQ}$	$\text{MODG} \rightarrow i$
3	$i \rightarrow \text{WREQ}$	$\text{INVR} \rightarrow k_j$
	$i \rightarrow \text{MREQ}$	$\text{INVR} \rightarrow k_j$
4	$i \rightarrow \text{WREQ}$	$\text{INWV} \rightarrow i$
	$i \rightarrow \text{MREQ}$	$\text{INWV} \rightarrow i$
5	$j \rightarrow \text{RREQ}$	$\text{INWV} \rightarrow i$
	$j \rightarrow \text{FETCH}$	$\text{INWV} \rightarrow i$
6	$i \rightarrow \text{REPM}$	-
7	$j \rightarrow \text{RREQ}$	$\text{BUSY} \rightarrow j$
	$j \rightarrow \text{WREQ}$	$\text{BUSY} \rightarrow j$
	$j \rightarrow \text{MREQ}$	$\text{BUSY} \rightarrow j$
	$j \rightarrow \text{FETCH}$	$\text{BUSY} \rightarrow j$
	$j \rightarrow \text{ACKC}$	-
8	$j \rightarrow \text{ACKC}$	$\text{WDATA} \rightarrow i$
	$j \rightarrow \text{REPM}$	$\text{WDATA} \rightarrow i$
	$j \rightarrow \text{UPDATE}$	$\text{WDATA} \rightarrow i$
9	$j \rightarrow \text{RREQ}$	$\text{BUSY} \rightarrow j$
	$j \rightarrow \text{WREQ}$	$\text{BUSY} \rightarrow j$
	$j \rightarrow \text{MREQ}$	$\text{BUSY} \rightarrow j$
	$j \rightarrow \text{FETCH}$	$\text{BUSY} \rightarrow j$
10	$j \rightarrow \text{ACKC}$	$\text{RDATA} \rightarrow i$
	$j \rightarrow \text{REPM}$	$\text{RDATA} \rightarrow i$
	$j \rightarrow \text{UPDATE}$	$\text{RDATA} \rightarrow i$

Although `Alewife` provides support for fine-grain synchronization, these mechanisms are implemented over the cache coherence protocol, which works as if *full/empty* bits do not exist. The cache controller in `Alewife` has limited hardware support for *full/empty* bits storage. Concretely, these bits are saved as an extra field in the cache tags. This has two advantages. First, the memory used to store cache

data does not need to have odd word-length. Second, access to the cache data is slower than access to the cache tags.

When the processor requests a memory access, the *Communications and Memory Management Unit* (CMMU) determines whether the access is local or remote. The CMMU also checks if the access implies a synchronizing operation by analyzing the `AS1` value in the memory operation. The address corresponding to the access is checked against the cache tags file, and both the appropriate tag and the *full/empty* bit are retrieved. At this point one of the following actions is executed<sup>9</sup>:

- a context switch is executed if the access produces a cache miss,
- a *full/empty* trap is fired in the case of a synchronization fault,
- otherwise, the operation is completed successfully.

According to the performance measures made in *Alewife*, the overhead of successful synchronizing operations is not significant [46]. When a synchronization miss is detected, a trap is fired and the corresponding thread either polls the location until the synchronization condition is met or blocks according to a given waiting algorithm. While no additional hardware is required for polling, blocking needs to save and restore context registers. The latter case is notably expensive, as loads take two cycles and stores take three cycles.

By integrating synchronization mechanisms with the coherence protocol, the complexity introduced by thread scheduling is avoided. Instead, synchronization misses are handled similarly to ordinary cache misses. As the hardware needed to deal with the latter has already the capability to store part of the information associated with a synchronization miss, it is expected that the hardware overhead introduced by integrating synchronization mechanisms with cache coherence is not excessive.

## 5.1. Mapping between processor instructions and network transactions

The network transactions used in the proposed protocol are explained in Table 7, which shows both messages sent from a cache to memory and requests sent back from memory to a cache.

Six new messages are introduced in order to implement fine-grain synchronization at the cache level. More concretely, these messages are `SRREQ`, `SWREQ`, `SCREQ` from cache to memory and `SRDENY`, `SWDENY` and `ACKSC` from memory to cache.

---

<sup>9</sup> In all cases, the retrieved *full/empty* bit is placed into the external condition codes so that the processor has access to it.

Table 7: Network transactions in the directory-based protocol

Type of message	Symbol	Semantics
Cache to Memory	RREQ	request to read a word that is not in the cache
	WREQ	request to write a word
	SRREQ	waiting and non-altering read request
	SWREQ	waiting and altering write request
	SCREQ	request to clear the full/empty bit
	UPDATE	returns modified data to memory
Memory to Cache	ACKC	acknowledges that a word has been invalidated
	RDATA	contains a copy of data in memory (response to RREQ)
	WDATA	contains a copy of data in memory (response to WREQ)
	SRDENY	sent if a SRREQ misses; the requesting cache will retry at a later stage
	SWDENY	sent if a SWREQ misses; the requesting cache will retry at a later stage
	INV	invalidates cached words
	ACKSC	acknowledges that the full/empty bit has been unset in all the copies of the block
BUSY	response to any RREQ or WREQ while invalidations are in progress	

As proposed in [74], some fields are needed in the coherence protocol messages in order to integrate fine-grain synchronization. We will make use of some of these proposed additional fields. Specifically, the following fields are required:

- slot's index in the cache line which is being accessed,
- slots in the home directory copy whose list of pending requests is empty; this allows saving protocol messages in some cases where a block is in the `read-write` state (see section 5.3.3 for more details),
- deferred lists in remote caches are sent to the home node when they release the exclusive ownership; this scenario is further explained in section 5.2.

When a processing node issues a memory operation, the cache located at that node interprets the request and translates it into one or more network transactions. The correspondence between the different processor instructions and memory requests sent over the network is shown below.

Table 8: Correspondence between processor instructions and memory requests

Instruction from processor	Initiated network transactions
PrUNRd	RREQ
PrUNWr	WREQ
PrUARd	RREQ + SCREQ
PrUAWr	-
PrSNRd	SRREQ
PrSNWr	CWREQ
PrSARd	SRREQ + SCREQ
PrSAWr	SWREQ

## 5.2. Management of pending requests

Extensive discussion about different alternatives for managing deferred lists is presented in [75]. We propose a hybrid procedure for managing deferred lists in which lists of pending operations are kept either at the home directory or in a distributed manner, depending on the state of the line to which pending operations refer. The rules for coalescing requests are the same as in Table 5.

Lists of pending requests for memory locations that are in an `absent`<sup>10</sup> or `read-only` state are maintained as an additional field in the corresponding home directory. Effectively, in these states it is not possible to adopt a distributed approach, since after a transition to the `read-write` state the home directory will need to have knowledge of the type of pending requests and the nodes that issued this requests.

A sample case of this scenario is shown in Figure 17, in which two nodes, namely A and B, share a copy of a given memory block. If another node takes the exclusive ownership of this block, information about pending requests issued by nodes A and B will be lost unless the home directory has knowledge of those requests. A naive approach is to make the directory keep track of only the nodes with pending requests, because this would require informing all of these nodes each time a *full/empty* state change is detected, thus generating extra traffic. Figure 17 also shows a different memory block for which there is no copy at any other node in the system, thus being in the `absent` state. The same rules apply for this location.

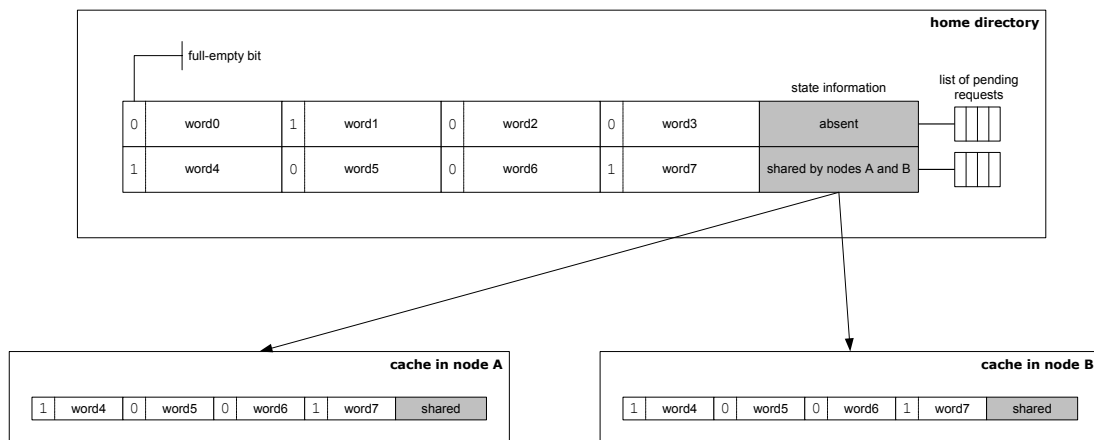


Figure 17: Management of pending requests for an absent or read-only memory block

For locations in a `read-write` state, we adopt a distributed solution in which both the home directory and the remote cache keep track of pending operations. When a remote cache releases its copy of the block, the deferred list kept locally to that cache is sent to the home node and merged with the deferred list at the home directory. The rules for coalescing requests are those in Table 5. An example in which a location is first owned by node A and then flushed from its cache is shown on Figure 18.

<sup>10</sup> The absent state indicates that no cache is holding a copy of the referred memory location. Consequently this location does not fall into any of the four states described on page 33.

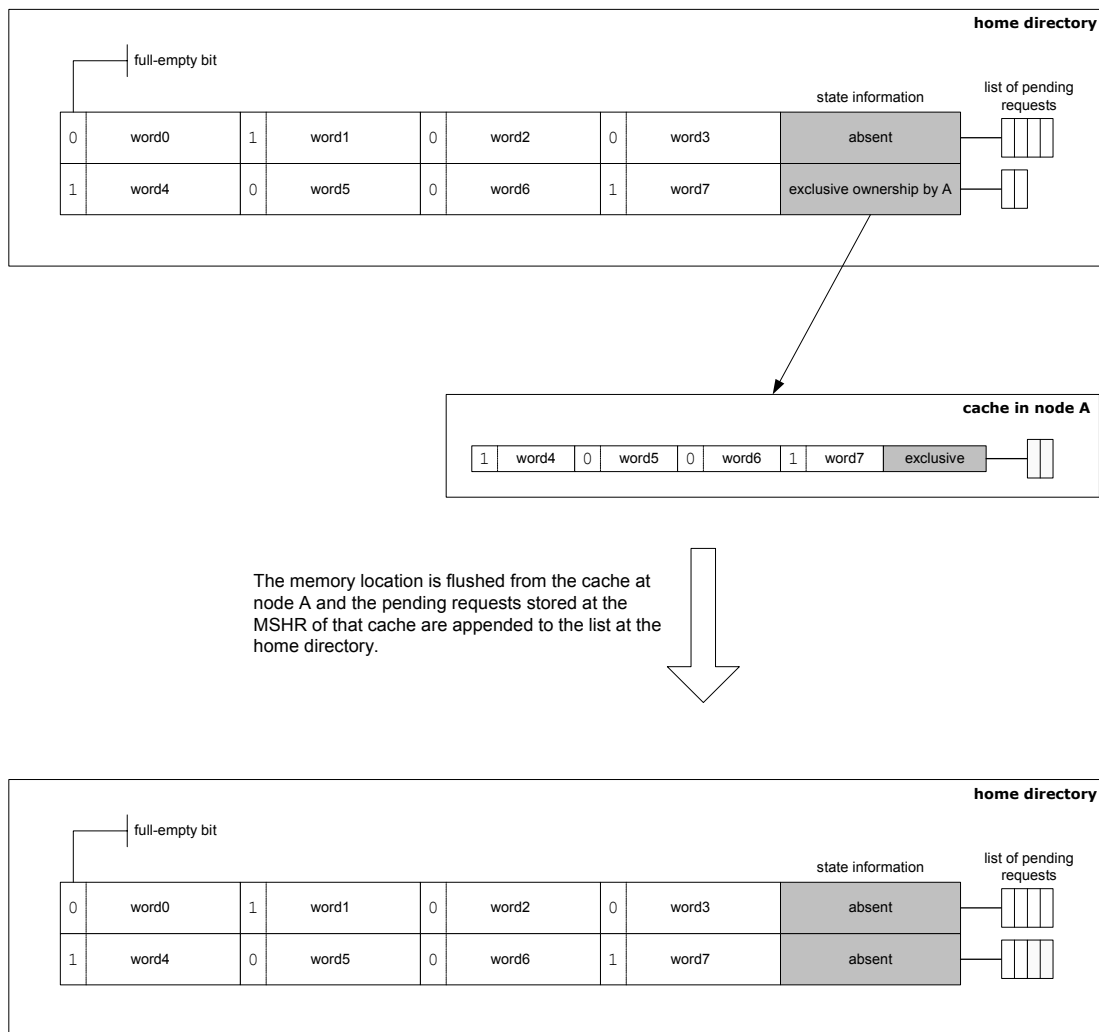


Figure 18: Management of pending requests for a read-write memory block

As in the bus-based scheme, it is also necessary to specify how resuming of pending requests is done. Contrary to the former, coherence of *full/empty* state bits is not always ensured at the home directory. In fact, the home directory does not have a valid copy of the *full/empty* bit of a memory location that is in the *read-write* state. In such case, the directory forwards requests from other nodes to the exclusive owner of the block, where they will be serviced. According to these features, resuming of pending requests is based on the following rules:

- if a block is in the *absent* or *read-only* state, the home directory is responsible for resuming requests, by checking if there is any entry in the deferred list that matches an incoming transaction,
- if a block is in the *read-write* state, the cache having the exclusive ownership knows whether there are pending requests for that block at the home directory. In that case, relevant operations performed at that node are forwarded to the home node in order to check if any pending request can be resumed. Otherwise, the deferred list can be locally managed at the exclusive owner.

Consequently, it is not possible for a cache to have pending requests for a memory location that is not cached. These pending requests are kept and managed at the home directory. This solution is a hybrid approach between a *fully distributed* and a *centralized* deferred list management.

### 5.3. Directory transition rules

A detailed explanation of the transition rules from each coherence state is presented in the following sections. A description in the form of C-styled *pseudo-code* is also presented in each case. In the state diagrams, the notation indicated below is used. A tilde symbol (~) is indicated when no side effect or output message is necessary.

**processor id: input message, preconditions / side effects / output message, local actions**

#### 5.3.1. Absent state

```

SWITCH (incomingRequest) {
  CASE RREQ(i): addNodeToDirectory(i); // "i" is the sending node id.
                send(RDATA, i);      // send requested data to node.
                nextState = readOnly;
                BREAK;
  CASE WREQ(i): IF (ackCounter == 0) {
                  addNodeToDirectory(i);
                  send(WDATA, i);
                  nextState = readWrite;
                } ELSE {
                  addNodeToDirectory(i);
                  nextState = writeTransaction;
                }
                BREAK;
  CASE SRREQ(i): IF (full) {
                  addNodeToDirectory(i);
                  send(RDATA, i);
                  nextState = readOnly;
                } ELSE {
                  send(RDENY, i);
                  addToDeferredList();
                  nextState = absent;
                }
                BREAK;
  CASE SWREQ(i): IF (empty && deferredListEmpty()) {
                  addNodeToDirectory(i);
                  send(WDATA, i);
                  nextState = readOnly;
                } ELSE IF (empty && !deferredListEmpty()) {
                  addNodeToDirectory(i);
                  send(WDATA, i);
                  resumePendingReqs();
                  nextState = readOnly;
                } ELSE {
                  send(WDENY, i);
                  addToDeferredList();
                  nextState = absent;
                }
                BREAK;
  CASE SCREQ(i): unsetFE();
                  send(ACKSC, i);
                  nextState = absent;
                  BREAK;
  CASE ACKC(i):  ackCounter--;
                  nextState = absent;
                  BREAK;
}

```

If a SRREQ is received and the synchronization state is met (the FE bit is set), the requesting cache is added to the directory and the requested data is sent in a RDATA message. The state is then changed to read-only. If the SRREQ fails (the FE bit is unset), a RDENY message is sent

back to the requesting cache and the operation is appended to the deferred list in the home node. The state is not changed and the requesting cache waits until the home node solves the synchronization miss and sends back the requested data.




Figure 19: State transitions from the absent state

If a `SWREQ` is received and the synchronization condition is met (`FE` is not set), the requesting cache is added to the directory and a `WDATA` message is sent to it. Any relevant pending request in the local deferred list is resumed and appropriate data is sent to its corresponding cache, which is also added to the directory. The state is then changed to `read-only`, and not to `read-write` as it could be expected. This optimization allows other processing nodes to read this data without any state transition. If the *full/empty* bit is set on a `SWREQ`, then a `WDENY` is replied and the operation is suspended. The state in the home directory is not changed.

If a `SCREQ` is received, the *full/empty* bit is reset and an `ACKSC` sent back to the requesting cache. The state in the home directory is not changed.

### 5.3.2. Read-only state

```

SWITCH (incomingRequest) {
  CASE RREQ(i): IF (hasPointerInDirectory(i)) {
    send(RDATA, i); // "i" is the sending node id.
  } ELSE IF (!directoryFull()) {
    addNodeToDirectory();
    send(RDATA, i);
  } ELSE {
    ++ackCounter;
    j = evictRandomDirectoryEntry(); // j is the evicted line.
    send(INV, j);
    addNodeToDirectory();
  }
}

```

```

        send(RDATA, i);
    }
    nextState = readOnly;
    BREAK;
CASE WREQ(i): IF (hasPointerInDirectory(i) && (numberOfEntries() > 1)) {
    ackCounter += numberOfEntries() - 1;
    FOR (j = 0; j < numberOfEntries(); j++) {
        If (i != j)
            send(INV, j);
    }
    clearDirectory();
    addNodeToDirectory(i);
    nextState = writeTransaction;
} ELSE IF (hasPointerInDirectory(i)
    && (numberOfEntries() == 1)
    && (ackCounter != 0)) {
    nextState = writeTransaction;
} ELSE IF (hasPointerInDirectory(i)
    && (ackCounter == 0)) {
    send(WDATA, i);
    nextState = readWrite;
} ELSE { // if the line is not in the directory
    ackCounter += n;
    FOR (j = 0; j < numberOfEntries(); j++) {
        send(INV, j);
    }
    clearDirectory();
    addNodeToDirectory(i);
}
}
BREAK;
CASE SRREQ(i): IF (full && hasPointerInDirectory(i)) {
    send(RDATA, i);
} ELSE IF (full && !directoryFull()) {
    addNodeToDirectory(i);
    send(RDATA, i);
} ELSE IF (full && directoryFull()) {
    ++ackCounter;
    j = evictRandomDirectoryEntry(); // j is the evicted line.
    send(INV, j);
    addNodeToDirectory(i);
    send(RDATA, i);
} ELSE IF (empty) {
    send(RDENY, i);
    addToDeferredList();
}
}
nextState = readOnly;
BREAK;
CASE SWREQ(i): IF (empty & deferredListEmpty()) {
    addNodeToDirectory(i);
    send(WDATA, i);
} ELSE IF (empty & !deferredListEmpty()) {
    addNodeToDirectory(i);
    send(WDATA, i);
    resumePendingReqs();
} ELSE {
    send(WDENY, i);
    addToDeferredList();
}
}
nextState = readOnly;
BREAK;
CASE SCREQ(i): IF (numberOfEntries() > 1) {
    ackCounter += numberOfEntries() - 1;
    FOR (j = 0; j < numberOfEntries(); j++) {
        If (i != j)
            send(SCREQ, j);
    }
    clearDirectory();
    addNodeToDirectory(i);
    nextState = writeTransaction;
} ELSE IF (hasPointerInDirectory(i)) {
    unsetFE();
    send(ACKSC, i);
    nextState = readOnly;
}
}
BREAK;
CASE ACKC(i): ackCounter--;
nextState = readOnly;
BREAK;
}

```

If a SRREQ is received and the synchronization condition is met, an RDENY message is replied and the request is appended to the local deferred list. If the synchronization condition is met and



the requesting cache is already in the directory, an RDATA message is sent back with the requested memory location. If the requesting cache is not in the directory and there are still free directory entries, the cache is added to the directory. Otherwise, a random cache is replaced with the requesting cache and an INV message is sent to the removed cache. The home directory state is not changed in any case.




Figure 20: State transitions from the read-only state

If a SWREQ is received and the *full/empty* is set, then a WDENY message is replied. Otherwise, the requesting cache is added to the directory and a WDATA message is sent back. In any case, the home directory state is not changed.

If a SCREQ is received and no more caches share this block, then the *full/empty* bit is cleared and the request is acknowledged with an ACKSC message. The state is not changed in this case. However, if more caches have a copy of this block, their *full/empty* bits must be reset before acknowledging the operation. Consequently, the state is changed to write-transaction and an SCREQ message is sent to each cache with a copy of the block. Note that the SCREQ operation is particularly time-expensive, as it works as a barrier for all the involved caches.

### 5.3.3. Read-write state

```

SWITCH (incomingRequest) {
  CASE RREQ(j): IF (!hasPointerInDirectory(j)) { // there is only one node
    ++ackCounter; // in the directory (the
    send(INV, i); // owner, namely "i")
    clearDirectory();
    addNodeToDirectory(j);
    nextState = readTransaction;
  }
}
  
```

```

        BREAK;
CASE WREQ(j): IF (!hasPointerInDirectory(j)) {
    ++ackCounter;
    send(INV, i);
    clearDirectory();
    addNodeToDirectory(j);
    nextState = writeTransaction;
}
        BREAK;
CASE SRREQ(j): IF (!hasPointerInDirectory(j) && full) {
    ++ackCounter;
    send(INV, i);
    clearDirectory();
    addNodeToDirectory(j);
    nextState = readTransaction;
} ELSE IF (empty) {
    send(RDENY, j);
    addToDeferredList();
    nextState = readWrite;
}
        BREAK;
CASE SWREQ(j): IF (!hasPointerInDirectory(j) && empty) {
    ++ackCounter;
    send(INV, i);
    clearDirectory();
    addNodeToDirectory(j);
    nextState = writeTransaction;
} ELSE IF (full) {
    send(WDENY, j);
    addToDeferredList();
    nextState = readWrite;
}
        BREAK;
CASE SCREQ(j): send(SCFWD, i);
nextState = readWrite;
        BREAK;
CASE ACKC(j): ackCounter--;
nextState = readOnly;
        BREAK;
CASE UPDATE(i, Dpack): addToDeferredList(Dpack);
resumePendingReqs();
nextState = readOnly;
        BREAK;
}

```

If a SRREQ is received, either a RDATA or RDENY message is replied depending on whether the synchronization condition is met. The state is not changed in any case.

If a SWREQ is received from a cache other than the current owner of the block and the synchronization condition is met, the request is forwarded to the owner. Additional functionality is required in the cache protocol, as the owner is expected to answer this type of forwarded requests. Another design alternative is to centralize all the synchronized writes in the home node. This avoids the need of forwarded requests but introduces an overhead associated with the excess traffic generated by the caches that request a synchronized write to the home node even though they are exclusive owners for that block.

Our design assumes that caches can service forwarded requests. In this case, when a cache with exclusive ownership performs a synchronized write to a block, it only communicates this action to the home node in the case where the deferred list at the home node is not empty. This knowledge is implicit in an extra bit at the cache side, which is set when there are pending requests for the referred location at the home directory, as proposed in [74].




Figure 21: State transitions from the read-write state

As with the SWREQ operation, a SCREQ coming from a cache different than the owner is forwarded to the cache that has the read-write privilege.

#### 5.3.4. Read transaction state

```

SWITCH (incomingRequest) {
  CASE RREQ(i):  send(BUSY, i);
                 nextState = readTransaction;
                 BREAK;
  CASE WREQ(i):  send(BUSY, i);
                 nextState = readTransaction;
                 BREAK;
  CASE SRREQ(i): send(BUSY, i);
                 nextState = readTransaction;
                 BREAK;
  CASE SWREQ(i): send(BUSY, i);
                 nextState = readTransaction;
                 BREAK;
  CASE SCREQ(i): send(BUSY, i);
                 nextState = readTransaction;
                 BREAK;
  CASE ACKC(i):  ackCounter--;
                 nextState = readOnly;
                 BREAK;
  CASE UPDATE(i): --ackCounter;
                  send(RDATA, i);
                  nextState = readOnly;
                  BREAK;
}

```

No new transitions are added from this state. All synchronized operations are ignored and a BUSY message is sent back to the requesting cache.




Figure 22: State transitions from the read transaction state

### 5.3.5. Write transaction state

```

SWITCH (incomingRequest) {
  CASE RREQ(i):  send(BUSY, i);
                 nextState = writeTransaction;
                 BREAK;
  CASE WREQ(i):  send(BUSY, i);
                 nextState = writeTransaction;
                 BREAK;
  CASE SRREQ(i): send(BUSY, i);
                 nextState = writeTransaction;
                 BREAK;
  CASE SWREQ(i): send(BUSY, i);
                 nextState = writeTransaction;
                 BREAK;
  CASE SCREQ(i): send(BUSY, i);
                 nextState = writeTransaction;
                 BREAK;
  CASE ACKC(i):  IF (ackCounter == 1) {
                   ackCounter = 0;
                   send(WDATA, cacheInDirectory());
                   nextState = readWrite;
                 } ELSE {
                   --ackCounter;
                   nextState = writeTransaction;
                 }
                 BREAK;
  CASE ACKSC(i): IF (ackCounter == 1) {
                   ackCounter = 0;
                   send(ACKSC, cacheInDirectory());
                   nextState = readWrite;
                 } ELSE {
                   --ackCounter;
                   nextState = writeTransaction;
                 }
                 BREAK;
  CASE UPDATE(i): IF (ackCounter == 1) {
                   ackCounter = 0;
                   send(WDATA, cacheInDirectory());
  
```

```

        nextState = readWrite;
    } ELSE {
        --ackCounter;
        nextState = writeTransaction;
    }
    BREAK;
}

```

As in the previous case, all SRREQ, SWREQ and SCREQ messages are replied with a BUSY message.




Figure 23: State transitions from the write transaction state

A new transition is specified in the protocol in the case when a cache has requested to clean the *full/empty* bit of all caches with a copy of a block. In this case the acknowledgement counter already present in the implementation is used to keep track of the caches that have not yet cleared their copy of the *full/empty* bit. After a cache clears its copy of the *full/empty* bit, it sends an ACKSC message to the directory and invalidates the corresponding cache line. When all the ACKSC messages have been received from the home node, then the operation is acknowledged to the requesting cache.

## 5.4. Summary

A directory based protocol with support for fine-grain synchronization has been systematically specified in the form of state diagrams and *pseudo-code*. As in the bus based directory protocol (section 4), only waiting non-altering reads and waiting altering writes are considered in this implementation. The operation of other variants of synchronized accesses can be easily inferred because they are a simplified version of the former.

Six new network messages are introduced in order to implement fine-grain synchronization at the cache coherence level. Some optimizations reducing the number of messages are proposed, requiring additional functionality in the protocol so that caches can service forwarded requests that are sent to the directory from other caches.

We propose a deferred list management scheme in which lists of pending requests can be either kept at the home directory or distributed between the directory and the caches. This solution is a compromise between a *distributed* approach and a *centralized* design and minimizes the number of protocol messages sent over the network. The same rules as in the bus-based approach are applied for coalescing of pending requests.

## 6. Simulation framework

---

As a practical working model of the proposed coherence protocols, a *directory-based* protocol with fine-grain synchronization support has been partially implemented and simulated. This experimental model is based on the Rice Simulator for ILP<sup>11</sup> Multiprocessors (RSIM) simulator<sup>12</sup> and runs on Solaris 2.5 or above.

### 6.1. Features of the simulated platform

RSIM is a discrete event-driven simulator based on the YACSIM library [43]. This means that most of the resources in the simulated architecture are activated as events only when they have some tasks to execute. As an exception, both processor and caches are simulated as an event that is executed on every cycle. This decision is based on the facts that those units are likely to have nearly continuous activity.

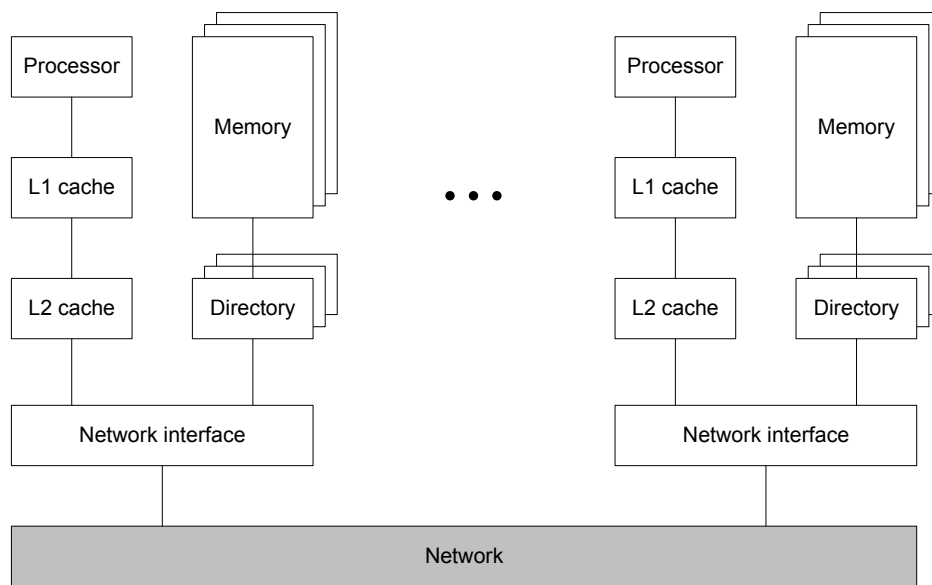


Figure 24: Simulated system architecture

Figure 24 shows the network and memory system hierarchy in the simulation platform. The key features of simulated systems are listed below [59].

- Multiple instruction issue
- Out-of-order scheduling
- Branch prediction support
- Non-blocking loads and stores

---

<sup>11</sup> Instruction-level parallelism.

<sup>12</sup> Available at <http://rsim.cs.uiuc.edu/rsim/> (accessed November 2.001).

- Optimized memory consistency implementations
- Two-level cache hierarchy
- Multiple outstanding cache requests
- CC-NUMA shared-memory system
- *Directory-based* cache coherence protocol with fine-grain synchronization support
- Routed two-dimensional mesh network

Additionally, contention effects are modeled at all resources in the processor, caches, memory banks, processor-memory bus and network.

## 6.2. Simulation methodology

The steps required to perform a general simulation with the developed platform are depicted in the figure below.




Figure 25: Simulation steps

The starting point is the source code of the program to be run under the simulator. As this code is supposed to successfully compile on an ordinary compiler, no language-level support for expressing data-level synchronization operations is available at this step. However, it is necessary to somehow distinguish these operations. Using unique assembler instructions through the source code achieve this<sup>13</sup>.

An 8-bit *Alternate Space Indicator* (ASI) is defined in the SPARC architecture in order to tag loads and stores with 256 different values. As some of these values are user-defined, they can be used for synchronizing instructions. As a consequence, synchronizing operations are distinguished by particular ASI values. The ASI parameter determines the specific variant of synchronizing instruction that will be executed. A sample C program is presented below.

```

int main(int argc, char **argv) {
    int sVar; /* synchronized variable */
    /* The values of ASI for synchronizing operations are:
       ASI_UE 0x91
       ASI_FF 0x96
       ASI_EFT 0x9d
    */
    asm("wr    %g0, 0x9d, %asi"); /* WRASI instruction */
    sVar = 5; /* synchronized store (STWA_EFT) */

    /* The complete assembler sequence looks like this:
       wr    %g0, 0x9d, %asi
       mov   xxx, %o0           !xxx is the data to be stored
       stwa %o0, [%fp - yyy], %asi !yyy is an appropriate offset
    */
}
  
```

<sup>13</sup> For example, with calls to the function *asm* in the C language.



In order to get a binary SUN's `cc` compiler has to be used<sup>14</sup> with the `-xarch=v8plusa` option. Otherwise `WRASI` will not be recognized as a valid instruction. A binary is obtained with<sup>15</sup>:

```
cc -xarch=v8plusa synch.c -o synch
```

The simulation process currently requires the use of a hexadecimal editor in order to manually modify the op-code of the desired memory instructions so that they are synchronized (refer to Appendix A). As stated above, these memory instructions are easily recognized because they are preceded by a write to the `ASI` register. A disassembler must be used in conjunction with the hexadecimal editor in order to determine the appropriate offset of the memory operations in the binary file. As no `SPARC32+` disassembler is openly available, a standard `SPARC` disassembler was modified in order to recognize the new instructions. A future improvement would be to extend the compiler in order to support the complete set of synchronizing memory instructions. With this extension, the simulation steps would be simplified as shown in the next figure.




Figure 26: Simulation steps with a compiler supporting synchronizing instructions

Once a binary with synchronizing instructions is obtained, a `predecoder` is executed on it. A new binary in a loosely encoded format, which can be interpreted by the simulator, is thus obtained.

### 6.3. Implementation of synchronizing instructions




Figure 27: Alternate load and store instruction format

As specified in [73], the only `SPARC` instructions that access memory `load`, `store`, `prefetch`, `swap`, and `compare-and-swap`. An implicit `ASI` value is provided by normal `load` and `store`. On the contrary, an explicit `ASI` is provided by alternate `load` and `store`. This explicit value is given either

<sup>14</sup> `gcc` version 3.1 and above also supports `SPARC v8+` and `v9` architectures, but it's currently under development.

<sup>15</sup> Note that the application has to be compiled using the `RSIM` application library.

in the ASI register or in the `imm_asi` instruction field (see Figure 27). The 6-bit field `op3` determines the specific `load` or `store` instruction.

Synchronization instructions are implemented on SPARC by defining them as colored load and stores, as specified by the ASI field. ASI values corresponding to synchronizing instructions are presented in Table 9. Each instruction category is assigned four consecutive ASI values. Two of these values specify altering instructions, while the other two represent non-altering accesses. A total of 16 synchronizing memory instructions are introduced<sup>16</sup>.

Table 9: ASI values for synchronizing operations

ASI value range	Instruction category
0x90 to 0x93	Unconditional
0x94 to 0x97	Conditional waiting
0x98 to 0x9B	Conditional non-faulting
0x9C to 0x9F	Conditional faulting

Table 10 summarizes the most relevant modifications already made to RSIM.

Table 10: Specific modifications made to RSIM

Source file	Changes
<code>MemSys/cache2.c</code>	Added extra fields in cache lines storing full/empty bits <sup>17</sup>
<code>MemSys/directory.c</code>	Added extra fields in directory storing full/empty bits
<code>MemSys/l1cache.c</code>	Cache behavior integrated with full/empty bits and synchronizing operations
<code>MemSys/mshr.c</code>	Specification of new types of memory operations and extension of MSHR registers <sup>18</sup>
<code>MemSys/setup_cohe.c</code>	Implementation of a coherence protocol integrated with fine-grain synchronization at both L1 and L2 caches
<code>Processor/except.cc</code>	New type of <i>soft</i> exception fired by trapping conditional instructions
<code>Processor/funcs.cc</code>	Functional implementation of conditional instructions
<code>Processor/memunit.cc</code>	Behavior of the memory unit when dealing with synchronizing instructions
<code>Processor/units.cc</code>	Specification of functional units used by synchronizing instructions, access types and latencies
<code>predecode/predecode_instr.cc</code>	New values of ASI for synchronizing instructions
<code>predecode/predecode_table.cc</code>	Association between new instructions and the functions corresponding to its implementation

<sup>16</sup> Although loads and stores are defined on both integers and floating-point data supporting byte, half-word (16-bit), word (32-bit), double-word (64-bit) and quad-word (128-bit) accesses, only integer word memory operations are supported by the simulation platform.

<sup>17</sup> Currently, the simulation platform only supports cache lines of 64-bit length.

<sup>18</sup> The management of synchronizing pending requests has not been implemented yet.

Table 11: Set of full/empty memory instructions<sup>19</sup>

Alternate superspace		Alternate space	ASI	Instruction	Operation (atomic)	RSIM (extension)
ASI_UNCOND (0x90..0x93)		ASI_UU	0x90	unconditional load or store	set F/E to F/E condition bit; read or write;	LD*_UU ST*_UU
		ASI_UF	0x93	unconditional store and set full	set F/E to F/E condition bit; write and set full;	ST*_UF
		ASI_UE	0x91	unconditional load and set empty	set F/E to F/E condition bit; read and set empty;	LD*_UE
		ASI_UR	0x92	reserved	-	
ASI_COND (0x94..0x9F)	ASI_COND_WAIT (0x94..0x97)	ASI_EE	0x94	waiting conditional store from empty	write when empty;	ST*_EE
		ASI_EF	0x95	waiting conditional altering store from empty	(write and set full) when empty;	ST*_EF
		ASI_FE	0x96	waiting conditional altering load from full	(read and set empty) when full;	LD*_FE
		ASI_FF	0x97	waiting conditional load from full	read when full;	LD*_FF
	ASI_COND_NOWAIT (0x98..0x9B)	ASI_EE_N	0x98	non-faulting conditional store from empty	set F/E to F/E condition bit; write if empty else skip;	ST*_EEN
		ASI_EF_N	0x99	non-faulting conditional altering store from empty	set F/E to F/E condition bit; (write and set full) if empty else skip;	ST*_EFN
		ASI_FE_N	0x9A	non-faulting conditional altering load from full	set F/E to F/E condition bit; (read and set empty) if full else skip;	LD*_FEN
		ASI_FF_N	0x9B	non-faulting conditional load from full	set F/E to F/E condition bit; read if full else skip;	LD*_FFN
	ASI_COND_FAULT (0x9C..0x9F)	ASI_EE_T	0x9C	faulting conditional store from empty	set F/E to F/E condition bit; write if empty else trap;	ST*_EET
		ASI_EF_T	0x9D	faulting conditional altering store from empty	set F/E to F/E condition bit; (write and set full) if empty else trap;	ST*_EFT
		ASI_FE_T	0x9E	faulting conditional altering load from full	set F/E to F/E condition bit; (read and set empty) if full else trap;	LD*_FET
		ASI_FF_T	0x9F	faulting conditional load from full	set F/E to F/E condition bit; read if full else trap;	LD*_FFT

<sup>19</sup> Extracted from [72]. An asterisk indicates a data type, such as floating point (F) or unsigned word (UW). Note that all instructions correspond to the set of LD\*A operations in SPARC.

## 6.4. Simulation flowchart

The figure below shows the different stages of a single simulation. The names of the functions and the source code filenames corresponding to the listed operations are specified at the top of the boxes. The core of the simulation platform consists of a loop in which scheduled events are executed. The iterations through this loop continue until the event list is empty, meaning that the simulation has finished.

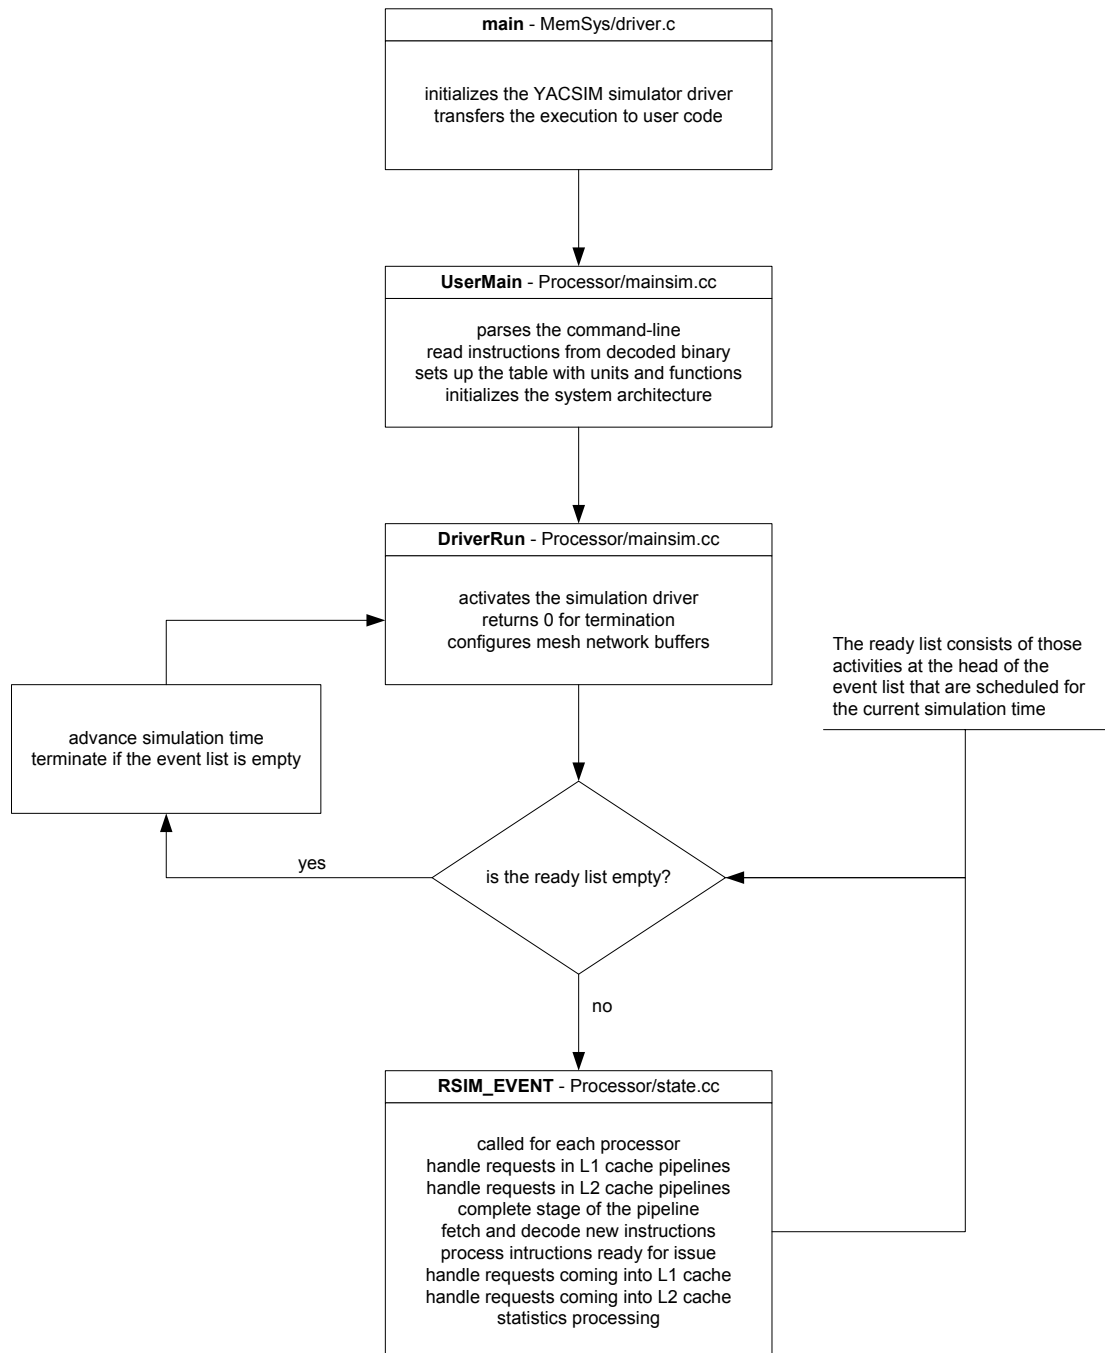


Figure 28: Execution flowchart of the simulator

As stated in Figure 29, the `RSIM_EVENT` function simulates processor and cache operation. As it is likely that both processor and caches have nearly continuous activity, `RSIM_EVENT` is scheduled every

cycle. However, in order to avoid non-deterministic behavior this function is scheduled to occur with an offset of 0.5 with respect to the processor cycles. As a consequence, `RSIM_EVENT` is executed at the midpoint between subsequent processor cycles in the simulation timeline.

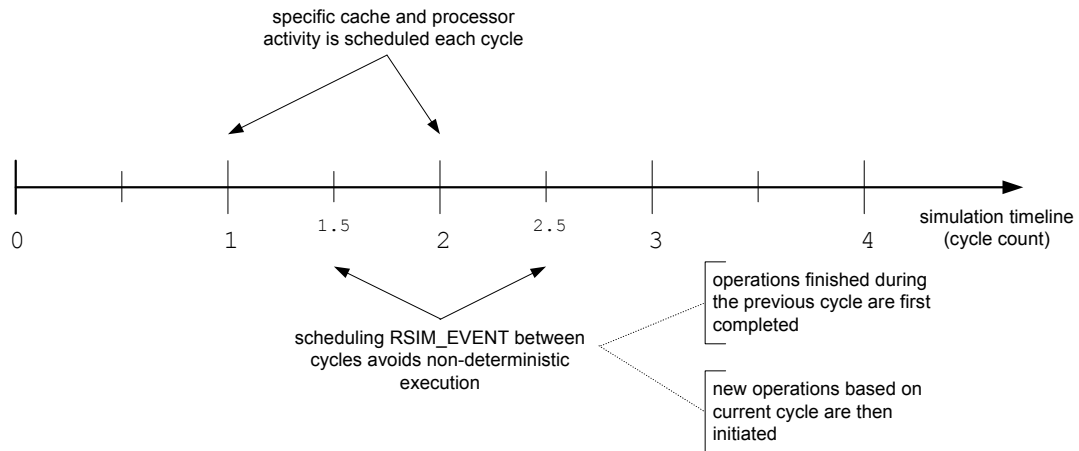


Figure 29: `RSIM_EVENT` scheduling

Figure 30 depicts instruction lifetime stages and the operations performed in each of these stages.

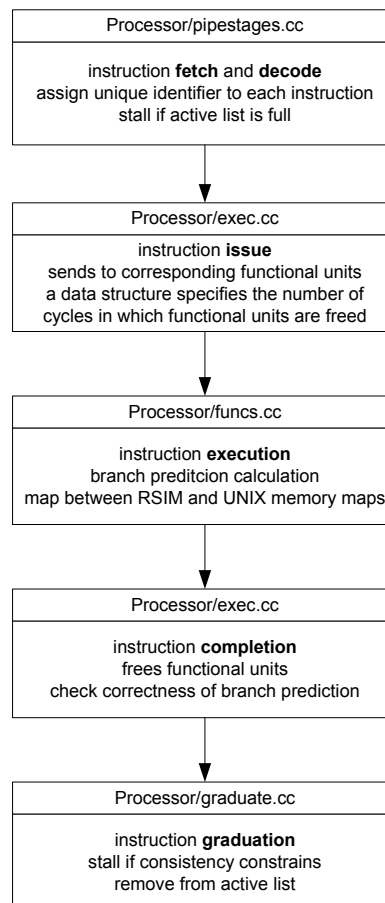




Figure 30: Instruction lifetime stages

## 6.5. Simulation results

A small application that makes use of the set of fine-grain operations implemented by the simulation platform has been developed. The application core is a loop in which a node issues a trapping altering store and the rest of processors perform a non-altering load in parallel. Both the number of nodes and number of iterations are customizable by command-line parameters.



a) Fine-grained application



b) Coarse-grained application

Figure 31: Normalized execution time for different machine and problem sizes

The operation of the fine-grained version of the application is defined below in the form of *pseudo-code* (complete source code is presented in Appendix B).

```

parse_command_line_parameters(); // Configure number of nodes and iterations.
allocate_shared_memory();       // To be used for the shared data array.
turn_on_memory_system_simulation(); // For collection of statistics.
create_processes();             // Each process is run on a different node.

```

```

LOOP {
  if (process_id != 0) { // For the specified number of iterations.
    read_synchronized_var(); // Processes other than the main process
  } else { // perform a waiting non-altering read.
    write_synchronized_var(); // The main process performs a trapping
  } // altering write.
}
}

```

The results of a series of experiments are graphically depicted in Figure 31<sup>20</sup>. While the left plot shows execution times for different machine sizes, the right plot compares execution times for various problem sizes (i.e. number of iterations). Diverse components of execution time are distinguished by a different shade in each bar of these plots. Table 12 shows a tabulated version of the simulation results for a constant number of iterations, while Table 13 shows the same results for a fixed number of processing nodes.

Table 12: Execution times (in cycles) for 1.000 iterations

Number of processors	Fine-grained version	Coarse-grained version
16	32176	2455349
8	26665	1845747
4	25983	1150828
2	24822	548853

As seen in Figure 31 (a) and Table 12, the execution time of the fine-grained version increases linearly with the problem size. Additionally, as the number of nodes decreases, the execution time slightly degrades. Effectively, the more nodes take part in a synchronization operation, the higher is the completion time. Both the cost of storage required for synchronization data and the traffic caused by these operations in the mesh network increase with the number of nodes. It is also important to observe that the most significant part of the execution time is due to cache and remote memory accesses.

Table 13: Execution times (in cycles) for 16 nodes

Number of iterations	Fine-grained version	Coarse-grained version
100000	3002126	-
10000	302068	-
1000	32176	2455349
100	4918	246280

A coarse-grained version of the same application has been implemented using barriers (see the complete source code in Appendix C). Its operation in the form of *pseudo-code* is detailed here for reference.

```

parse_command_line_parameters(); // Configure number of nodes and iterations.

```

<sup>20</sup> The results have been derived from the statistics collection utilities distributed with RSIM.

```
allocate_shared_memory(); // To be used for the shared data array.
initialize_barrier(barrier); // Rendez-vous point for all the processes.
turn_on_memory_system_simulation(); // For collection of statistics.
create_processes(); // Each process is run on a different node.

LOOP { // For the specified number of iterations.
  if (process_id != 0) { // Processes other than the main process
    wait_at_barrier(barrier); // wait for the write to complete and
    read_shared_var(); // perform an ordinary read.
  } else { // The main process performs a trapping
    write_shared_var(); // altering write.
    wait_at_barrier(barrier);
  }
}
```

The results, presented in Figure 31 (b) and Table 13, show that the dependence between the execution time and the number of nodes is remarkably higher in the coarse-grained version, which do not make use of a cache coherence protocol integrated with synchronization operations. This is because the overhead imposed by barriers is notably higher and so is its relevance in comparison with the total execution time. As a consequence, as the number of nodes increases, the execution time rises faster in the coarse-grained version. Note also that the accumulated times at the *FPU stall* and *busy* states are insignificant in comparison to the times spent for *remote accesses* and *barrier synchronization*.



## 7. Power-consumption estimation

---

Battery life is a key feature in modern mobile computing devices. As a consequence, power consumption has to be minimized when possible. This must be done without compromising the overall performance of the system.

As the architectures described in this work are suitable for integration as a CMP (chip multiprocessor) in a mobile device, it becomes important to evaluate the power consumption overhead introduced by systems with fine-grain synchronization support.

### 7.1. Available energy estimation tools

Traditionally, consumption estimation applications worked at gate level, thus needing a complete design of the multiprocessor in the form of a netlist. Examples of such tools are `PowerMill` [68] or `QuickPower` [63]. While providing extremely accurate estimations, these tools are only appropriate for the final stages of the design process, as they require a complete *hardware-level* description of the system. Moreover, the execution time of *hardware-level* simulation tools make them unsuitable for the study of tradeoffs between different system parameters.

Another approach has to be considered if power consumption estimations are needed in early design stages, where HDL descriptions and circuit designs are not yet available. *Macro modeling* can significantly reduce execution times of power estimation applications. Many tools developed for this purpose divide the system into several functional units and calculate their power consumption behavior separately [51]. More concretely, the power consumption of a functional unit is specified as a function of various parameters, such as the number of ports or the number of transitions in the input vector. These power consumption functions can then be integrated into an existing *high-level cycle-accurate* simulator, for instance, `RSIM`.

*Instruction-level* power analysis is useful in order to relate power consumption with application source code. Tools following this approach [66] are based on an estimation of the energy consumption associated with each particular instruction. A representative example of such application is `Myrmigki`, which also models hardware activity by keeping track of the transitions produced at the ports of the functional units.

`Wattch` [18] and `SimplePower` [76] implement more precise capacitive models and achieve accuracies as low as 5% with respect to measured values. They implement energy models based on transition counts and empirical data. Each functional unit is associated with a table that specifies the power consumption for a given state transition. If the number of possible transitions is too large, similar transitions can be grouped thus reducing the size of this table.

## 7.2. Implementing an energy estimation framework in RSIM

As shown in Figure 28, `RSIM` simulates the behavior of processor and caches at every cycle. This is implemented as part of the `RSIM_EVENT` function, which also includes statistics processing routines. A power consumption model can be added to the simulator platform by extending the statistics collection functions with an additional set of statistics<sup>21</sup>.

More concretely, apart from simulating the operation of the multiprocessor system, `RSIM` keeps track of detailed usage statistics for each functional unit. These statistics can be used to estimate consumed power with help of a set of tables indicating the energy associated for each possible transition, as it is done in tools like `Wattch` [18] or `SimplePower` [76]. The structure of these tables should be similar to Table 14, and they could actually be imported to `RSIM` with minor changes. Tables for novel functional units, such as caches with fine-grain synchronization support, should be obtained either experimentally or analytically by evaluating the power consumption overhead introduced by the additional hardware they require<sup>22</sup>.

Table 14: Structure of switch capacitance tables

Previous input vector	Current input vector	Switch capacitance (pF)
00...00	00...01	$C_0$
00...00	00...10	$C_1$
.	.	.
.	.	.
.	.	.
11...11	11...11	$C_n$

Additionally to these tables, it is necessary to keep track of the particular transitions and the number of accesses performed in each functional unit. The `YACSIM` [43] library provides useful functions for creating statistics records and calculating meaningful values such as the mean or the sum of stored samples.

Figure 32 sketches the module hierarchy of the simulator after implementing the changes mentioned above. Each functional unit is extended with an adequate power consumption model, which is invoked in case the corresponding unit has been used at the given cycle. The energy estimations for each unit are then collected by the statistics functions and summed according to the following expression:

$$E_{\text{CYCLE}} = E_{\text{CPU}} + E_{\text{L1 CACHE}} + E_{\text{L2 CACHE}} + E_{\text{MEM}} + E_{\text{NET}}$$

Energy estimation algorithms are expected to increase the simulation execution time by a factor up to 30% [18]. Gathering power consumption statistics at predefined intervals instead of every cycle notably reduces this overhead, but causes a loss of accuracy. Consequently, a compromise between

<sup>21</sup> The actual implementation of power consumption statistics in `RSIM` is not included in this work.

<sup>22</sup> Note that these tables are technology dependent, so the hardware level layout of the system must also be considered.

efficiency and accuracy has to be reached depending on the particular goals of the simulation being performed.

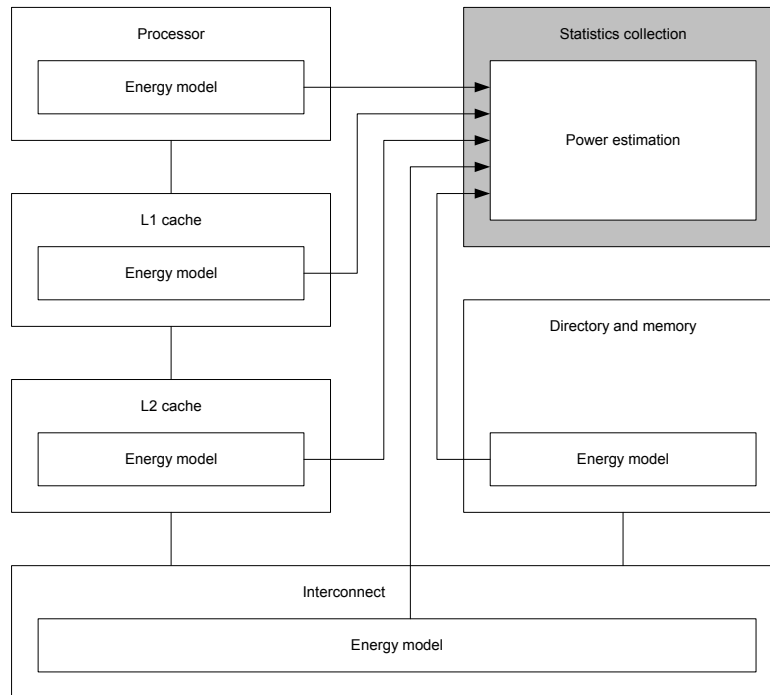


Figure 32: Integrated power consumption framework

### 7.3. Planned experiments

It is essential to validate the power model once it has been implemented in the simulator platform. This is achieved by comparing the results of a set of simulations with those given by already existing tools such as `Wattch` [18] or `SimplePower` [76]. If a hardware schematic of the system is available, comparisons can also be done with more accurate *hardware-level* tools (see [63] and [68]).

A power estimation framework is useful for studying the impact of fine-grain synchronization on crucial parameters of mobile devices. For instance, dissipation limits the maximum power consumption allowed per cycle. Additionally, battery life is governed by the total energy required by an application. It would be extremely valuable to evaluate the tradeoffs of fine-grain synchronization on both performance and power. Adequate benchmark applications have to be written for this purpose.

## 8. Conclusions

---

Fine-grain synchronization is a valuable mechanism for speeding up the execution of parallel algorithms by avoiding false data dependencies and unnecessary process waiting. However, the implementation of fine-grain synchronization introduces additional complexity at both hardware and software system components.

A novel architecture with support for fine-grain synchronization at the cache coherence level is introduced. We propose a model that can be efficiently implemented in modern multiprocessors. The hardware overhead required by this architecture is not expected to be excessive.

Coherence protocols with support for fine-grain synchronization have been systematically described for both bus-based and directory-based multiprocessors. This work includes as well description of the rules for management and resuming pending requests, which is a key issue for the correct operation of the presented architecture.

Although it has not been completely developed yet, the simulation platform has been tested with a sample application making use of a small set of conditional operations. A coarse-grained version of the same application has been written and its simulation results compared to those of the fine-grained version, showing the performance improvements provided by the latter. These preliminary results verify the worthiness of implementing fine-grain synchronization at the cache coherence level.

Some guidelines for implementing a power estimation algorithm as an extended feature of the simulation platform have been presented. The integration of fine-grain synchronization at the cache coherence level is expected to increase the energy consumption of the system, but this has to be quantitatively studied.

## 9. Future work

---

Some features such as sophisticated management of pending requests have been specified but not yet source coded. Additionally, further debugging of the simulation platform is required<sup>23</sup>. This will not only verify the correct functioning of the protocols, but also evaluate design options that were taken during the specification process. Protocol verification with automatic verifier tools is also desirable.

Further simulation is required in order to obtain more precise quantitative data related with the performance of the proposed set of synchronization memory operations. In particular, the statistics collection functions implemented in the simulator platform should be modified so that the cost of storage required for synchronization data and the latency of fine-grain synchronization operations can be measured and easily compared with traditional synchronization mechanisms. Other important parameters to be measured are extra traffic caused by these operations and saturation that may be present at different levels of the memory hierarchy.

Extending a standard C compiler in order to support the complete set of synchronizing memory instructions would greatly simplify the steps required to perform a single simulation. An alternative is to extend the `predecoder` so that it recognizes `Alewife` binaries, which would eliminate the necessity of writing applications from scratch for making use of synchronizing operations. Note however that `Alewife` doesn't make use of the full set of proposed instructions.

Another pending task is to implement the full set of synchronizing instructions under `RSIM`. The evaluation of different coherence protocols other than `MESI` would also be very valuable, as well as developing extensive statistics collection in order to understand the tradeoffs involved in the proposed architectures. The implementation of a power estimation algorithm is also an appealing task, as considerable source code can be reused from already existing energy estimation tools.

---

<sup>23</sup> Regarding this subject, the source code has also been compiled with `gcc` instead of `SUN's cc`, allowing thus to debug under `gdb`, which provides many more debugging features than `dbx`.

## Appendix A. Preparing binaries for simulation

---

This section describes in detail how to prepare a binary that makes use of fine-grain synchronization to be used with the extended `RSIM` simulator, supposing that we start from the following C source code. Note that the synchronized store is marked by a previous write to the `ASI` register, which is likewise performed by a call to the `asm` function.

```
int main(int argc, char **argv) {
    int sVar; /* synchronized variable */

    /* The values of ASI for synchronizing operations are:
       ASI_UE  0x91
       ASI_FF  0x96
       ASI_EF_T 0x9d
    */
    asm("wr   %g0, 0x9d, %asi");          /* WRASI instruction */

    sVar = 5;                             /* synchronized store (STWA_EFT) */

    /* The complete assembler sequence looks like this:
       wr   %g0, 0x9d, %asi
       mov  xxx, %o0           !xxx is the data to be stored
       stwa %o0, [%fp - yyy], %asi !yyy is an appropriate offset
    */
}
```

In order to get an ordinary `SPARC` binary, `SUN`'s `cc` compiler has to be invoked with the `xarch=v8plusa` option. Otherwise the store to the `ASI` register will not be recognized as a valid instruction. A binary is obtained with:

```
cc -xarch=v8plusa synch.c -o synch
```

A disassembler is now used to calculate the file offset in which the store is located. The op-code of this store will be changed so that it is marked as synchronized. It is essential for the disassembler to support the `SPARC32+` instruction set. The relevant disassembled output is listed below. Some instruction fields may vary depending on the particular system.

ADDRESS	INSTRUCTION	DECODED
0x000107f0	0x87826000	wr %g0, 157, %asi
0x000107f4	0x90102005	mov 5, %o0
0x000107f8	0xd027bfff	st %o0, [%fp - 8]
0x000107fc	0x81c7e008	ret

As it is preceded by a `WR` instruction to the `ASI` register, it is straightforward to find the store whose opcode needs to be modified. The detailed instruction format of this store is shown in the following figure, in which it is also depicted the field to be changed so that the store is labelled as synchronized. This change can be easily made with a standard hexadecimal editor<sup>24</sup>.

---

<sup>24</sup> `khexedit` has been used in this study.




Figure 33: Details on how to transform a standard store to a synchronized store

As deduced from Figure 33, in this example the byte at offset 0x000107f9 has to be changed from value 0x27 to value 0xa7. The resulting binary can be then predecoded and used as the input of the simulation platform.

Table 15 shows the relevant values of the field op3 and the corresponding operation associated to those values [73].

Table 15: Relevant values of the op3 field

Op-code	Operation	op3 field (binary)
LDUW	Load Unsigned Word	00 0000
LDUWA	Load Unsigned Word from Alternate space	01 0000
STW	Store Word	00 0100
STWA	Store Word from Alternate space	01 0100

## Appendix B. Application source (fine-grained version)

---

Below is the source code listing of the sample application used in order to test the set of fine-grained synchronized memory operations. The `MEMSYS_OFF` and `MEMSYS_ON` calls make the simulator ignore non-relevant initialisation steps. Note also that after allocating shared memory space for a given variable with `shmalloc`, it is necessary to define the home node that owns this space by using the `AssociateAddrNode` function.

```
#include <rsim apps.h>
#include <stdio.h>
#include <stdlib.h>

/* #define __sparc_v9__ */

int NUM_PROCS = 1; /* number of processors */
int ITERATIONS = 1; /* number of iterations */
int DEBUG = 0; /* print debugging info */

int *sVar_; /* shared array of size ITERATIONS */
int proc_id; /* private variable */
int phase; /* private variable */
extern char *optarg;

main(int argc, char **argv) {
    int c, i, j, dummy;

    MEMSYS_OFF; /* turn off detailed simulation for initialization */

    while ((c = getopt(argc, argv, "p:i:d")) != -1)
        switch (c) {
            case 'p':
                NUM_PROCS = atoi(optarg);
                break;
            case 'i':
                ITERATIONS = atoi(optarg);
                break;
            case 'd':
                DEBUG = 1;
                break;
            case 'h':
            default:
                fprintf(stdout, "SYNCH - OPTIONS\n");
                fprintf(stdout, "\tp - Number of processors\n");
                fprintf(stdout, "\ti - Number of iterations\n");
                fprintf(stdout, "\td - Print debugging info\n");
                fprintf(stdout, "\th - Help\n");
                return;
        }

    sVar_ = (int*) shmalloc(ITERATIONS * sizeof(int));
    AssociateAddrNode(sVar_, sVar_ + ITERATIONS, 0, "sVar");

    if (sVar_ == NULL) {
        fprintf(stdout, "Unable to malloc shared region\n");
        exit(-1);
    }

    if (DEBUG)
        fprintf(stdout, "Running with %d processors and %d iterations...\n",
                NUM_PROCS, ITERATIONS);
    MEMSYS_ON;

    proc_id = 0;
    for (i=0; i<NUM_PROCS-1; i++) {
        if (fork() == 0) {
            proc_id = getpid();
            break;
        }
    }

    if (proc_id == 0) {
        StatReportAll();
        StatClearAll();
    }
    endphase();
}
```



```

newphase(++phase);          /* beginning of new phase */

for (j=0; j<ITERATIONS; j++) {
  if (proc_id == 0) {       /* the main thread stores (STWA EFT) the value */
    /* Values of ASI for synchronizing operations:
    ASI_UE  0x91
    ASI_FF  0x96
    ASI_EF_T 0x9d
    */
    asm("wr    %g0, 0x9d, %asi");
    sVar [j] = 9;          /* synchronized store (STWA EFT) */

    if (DEBUG)
      fprintf(stdout, "%d: Stored value %d from sVar_[%d]\n",
              proc_id, sVar_[j], j);

    /* The complete assembler sequence looks like this:
    mov  0x9d, %o1
    wr   %o1, 0x0, %asi
    mov  xxx, %o0                !xxx is the data to be stored
    stwa %o0, [%fp - yyy], %asi !yyy is an appropriate offset
    */
  } else { /* the rest of the threads try to LDWA_FF the value */
    /* Values of ASI for synchronizing operations:
    ASI_UE  0x91
    ASI_FF  0x96
    ASI_EF_T 0x9d
    */
    asm("wr    %g0, 0x96, %asi");
    dummy = sVar_[j];        /* synchronized load (LDWA_FF) and store to a
                              standard dummy variable */

    if (DEBUG)
      fprintf(stdout, "%d: Read value %d from sVar_[%d]\n",
              proc_id, sVar_[j], j);
  }
}

if (DEBUG)
  fprintf(stdout, "\nProcessor %d about to finish!\n\n", proc_id);

exit(0);                    /* completed successfully */
}

```

## Appendix C. Application source (coarse-grained version)

Below is the source code listing of the sample application implementing the same functionality of the fine-grained version by using barriers. Calls to the directives `MEMSYS_OFF` and `MEMSYS_ON` are used as in the fine-grained version (Appendix B). Barriers are initialised and activated by invoking `TreeBarInit` and `TREEBAR`, respectively.

```
#include <rsim apps.h>
#include <stdio.h>
#include <stdlib.h>

/* #define __sparc_v9__ */

int NUM_PROCS = 1; /* number of processors */
int ITERATIONS = 1; /* number of iterations */
int DEBUG = 0; /* print debugging info */

int *sVar_; /* shared array of size ITERATIONS */
TreeBar barrier; /* tree barrier */
int proc_id; /* private variable */
int phase; /* private variable */
extern char *optarg;

main(int argc, char **argv) {
    int c, i, j, dummy;

    MEMSYS_OFF; /* turn off detailed simulation for initialization */

    while ((c = getopt(argc, argv, "p:i:d")) != -1)
        switch (c) {
            case 'p':
                NUM_PROCS = atoi(optarg);
                break;
            case 'i':
                ITERATIONS = atoi(optarg);
                break;
            case 'd':
                DEBUG = 1;
                break;
            case 'h':
            default:
                fprintf(stdout, "SYNCH - OPTIONS\n");
                fprintf(stdout, "\tp - Number of processors\n");
                fprintf(stdout, "\ti - Number of iterations\n");
                fprintf(stdout, "\td - Print debugging info\n");
                fprintf(stdout, "\th - Help\n");
                return;
        }

    sVar_ = (int*) shmalloc(ITERATIONS * sizeof(int));
    AssociateAddrNode(sVar_, sVar_ + ITERATIONS, 0, "sVar");

    if (sVar_ == NULL) {
        fprintf(stdout, "Unable to malloc shared region\n");
        exit(-1);
    }

    TreeBarInit(&barrier, NUM_PROCS); /* initialize tree barrier */

    if (DEBUG)
        fprintf(stdout, "Running with %d processors and %d iterations...\n\n",
                NUM_PROCS, ITERATIONS);
    MEMSYS_ON;

    proc_id = 0;
    for (i=0; i<NUM_PROCS-1; i++) {
        if (fork() == 0) {
            proc_id = getpid();
            break;
        }
    }

    if (proc_id == 0) {
        StatReportAll();
        StatClearAll();
    }
    endphase();
}
```

```
newphase(++phase);          /* beginning of new phase */

for (j=0; j<ITERATIONS; j++) {
  if (proc id == 0) {       /* the main thread stores the value */
    sVar [j] = 9;          /* ordinary store */
    if (DEBUG)
      fprintf(stdout, "%d: Stored value %d from sVar_[%d]\n",
              proc_id, sVar [j], j);
    TREEBAR(&barrier, proc_id);
  } else { /* the rest of the threads try to load the value */
    TREEBAR(&barrier, proc id);
    dummy = sVar [j];      /* ordinary load and store to dummy variable */
    if (DEBUG)
      fprintf(stdout, "%d: Read value %d from sVar_[%d]\n",
              proc_id, sVar [j], j);
  }
}

if (DEBUG)
  fprintf(stdout, "\nProcessor %d about to finish!\n\n", proc_id);

exit(0);                    /* completed successfully */
}
```

# Acknowledgements

I would like to acknowledge the contributions of many people who have helped bring this work to fruition. My main advisor, Vladimir Vlassov, urged me on by way of his untiring support and seemingly belief in me. I am also indebted to Mats Brorsson, Mladen Nikitovic, Sven Karlsson, Guillaume Delannoy, Erland Nilsson and Andrés Martinelli, as well as the rest of the staff at the Laboratory of Electronics and Computer Systems (IMIT/KTH), including the system administrators, who bore numerous queries from me.

One portion of this work is the result of the collaboration with the Electrical and Computer Engineering Department at the University of Massachusetts (UMASS), and specifically with Csaba Andreas Moritz. I am grateful to him and his students Diganta Roychowdhury and Raksit Ashok for their perceptive comments and valuable material from the beginning of this work.

A special thanks to the former members of the MIT Alewife Project Donald Yeung and Matthew Frank for providing excellent, incisive advise and fast response despite the huge amount of daily mail they have to answer.

There are many that from behind the scenes have encouraged my work and made my stay in Stockholm more than pleasant. Much gratitude to my friends all around the globe, especially those who are constantly close to me even if they live many thousands kilometres away.

And last, but not least, I really appreciate the infinite patience and constant support in countless significant ways of my family, particularly my parents Araceli and Laudi, my sister Isa and my brothers Javi and José. This work is dedicated to them.

# References

- [1] Abdel-Shafi, H.; Hall, J.; Adve, S.V. and Adve, V.S.: “An Evaluation of Fine-Grain Producer-Initiated Communication in Cache-Coherent Multiprocessors”, Proceedings of HPCA-3, February 1997
- [2] Agarwal, A.: “The MIT Alewife Machine: Architecture and Performance”, 25 years of the International Symposia on Computer Architecture (selected papers), Association for Computing Machinery, August 1998, pages 103-110
- [3] Agarwal, A.; Beng-Hong Lim; Kranz, D. and Kubiawicz, J.: “APRIL: A Processor Architecture for Multiprocessing”, Laboratory for Computer Science, Massachusetts Institute of Technology, 1990
- [4] Agarwal, A.; Bianchini, R.; Chaiken, D.; Chong, F.T.; Johnson, K.L.; Kranz, D.; Kubiawicz, J.D.; Beng-Hong Lim; Mackenzie, K. and Yeung, D.: “The MIT Alewife Machine: Architecture and Performance”, Laboratory for Computer Science, Massachusetts Institute of Technology, 1999
- [5] Agarwal, A.; Bianchini, R.; Chaiken, D.; Johnson, K.; Kranz, D.; Kubiawicz, J.; Lim, B.H.; Mackenzie, K. and Yeung, D.: “The MIT Alewife Machine: Architecture and Performance”, Proceedings of the 22nd Annual International Symposium on Computer Architecture (ISCA'95), June 1995, pages 2-13
- [6] Agarwal, A.; Kubiawicz, J.D.; Kranz, D.; Lim, B.H.; Yeung, D.; D'Souza, G. and Parkin, M.: “Sparcle: An Evolutionary Processor Design for Large-Scale Multiprocessors”, Laboratory for Computer Science, Massachusetts Institute of Technology, 1993
- [7] Alverson, G.; Alverson, R.; Callahan, D.; Koblenz, B.; Porterfield, A. and Smith, B.: “Exploiting Heterogeneous Parallelism on a Multithreaded Multiprocessor”, Proceedings of the 6th ACM International Conference on Supercomputing, 1992
- [8] Alverson, R.; Callahan, D.; Cummings, D.; Koblenz, B.; Porterfield, A. and Smith, B.: “The Tera Computer System”, Proceedings of the International Conference on Supercomputing, June 1990, pages 1-6
- [9] Ang, B.S. et. al.: “StarT the Next Generation: Integrating Global Caches and Dataflow Architecture”, in Advanced Topics in Dataflow Computing and Multithreading, IEEE Press, 1995
- [10] Arvind and Iannucci, R.A.: “Two Fundamental Issues in Multiprocessing”, Technical Report TM 330, MIT, Laboratory for Computer Science, Oct. 1987
- [11] Ashok, R.; Roychowdhury, D. and Andras, C.: “Synchronization Coherence Protocols: Unifying Synchronization and Caching in Multiprocessors”, Department of Electrical and Computer Engineering, University of Massachusetts, February 2002
- [12] Beng-Hong Lim: “A Synchronization Library for ASIM”, Alewife Systems Memo #12, MIT Laboratory for Computer Science, January 1992

- 
- [13] Beng-Hong, L.: "Instruction Set Summary of the APRIL/SPARC Processor", Alewife Systems Memo #2, MIT Laboratory for Computer Science, Aug. 1990
- [14] Bitar, P. and Despain, A.M.: "Multiprocessor Cache Synchronization: Issues, Innovations and Evolution", Computer Science Division, University of California, 1986
- [15] Bokhari, S.H. and Mavriplis, D.J.: "The Tera Multithreaded Architecture and Unstructured Meshes", Technical Report ICASE Interim Report No. 33, NASA/CR-1998-208953, Institute for Computer Applications in Science and Engineering, Mail Stop 403, NASA Langley Research Center Hampton, VA, December 1998
- [16] Boon Seong Ang, Arvind and Chiou, D.: "StarT the Next Generation: Integrating Global Caches and Dataflow Architecture", Laboratory for Computer Science, Massachusetts Institute of Technology, CSG Memo 354, 1994
- [17] Boughton, R.D.: "Artic Routing Chip", in Parallel Computer Routing and Communications Proceedings of the First International Workshop, PCRW '94, volume 853 of Lecture Notes in Computer Science, pages 310-317, May 1994
- [18] Brooks, D.; Tiwari, V. and Martonosi, M.: "Wattch: A Framework for Architectural-Level Power Analysis and Optimizations", Proceedings of the 27th International Symposium on Computer Architecture, May 2000
- [19] Cache tutorial, Intel Corporation, 1st edition, 1991
- [20] Carpinelli, J. D.: "Computer Systems: Organization and architecture", Addison Wesley Longman, 1st edition, 2001
- [21] Chaiken, D.: "Cache Coherence Protocol Specification", Alewife Systems Memo #5, MIT Laboratory for Computer Science, Apr. 1990
- [22] Chaiken, D.; Fields, C.; Kurihara, K. and Agarwal, A.: "Directory-Based Cache Coherence in Large-Scale Multiprocessors", Computer Magazine, June 1990, pp. 49-59
- [23] Chaiken, D.; Johnson, K.: "NWO User's Manual", Alewife Systems Memo #36, MIT Laboratory for Computer Science, September 1994
- [24] Chaiken, D.; Kubiawicz, J.; Agarwal, A.: "LimitLESS Directories: A Scalable Cache Coherence Scheme", Laboratory for Computer Science, Massachusetts Institute of Technology, 1991
- [25] Chaiken, D.; Lim, B.H. and Nussbaum, D.: "ASIM User's Manual", Alewife Systems Memo #13, MIT Laboratory for Computer Science, November 1991
- [26] Chaiken, D.L.: "Cache Coherence Protocols for Large-Scale Multiprocessors", Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1990
- [27] Chiou, D.; Ang, B.S.; Arvind et. al.: "StarT-NG: Delivering Seamless Parallel Computing", EuroPar '95, August 1995
- [28] CSG Memos, <http://www.csg.lcs.mit.edu/pubs/memos/> (accessed November 2001)

- 
- [29] Culler, D. E. and Pal Singh, J.: "Parallel Computer Architecture: A hardware/software approach", Morgan Kaufmann Publishers, 1st edition, 1999
- [30] David E. Culler, Jaswinder Pal Singh: "Parallel Computer Architecture: A hardware/software approach", Morgan Kaufmann Publishers, 1999
- [31] Dennis, J.B.: "A Preliminary Architecture for a Basic Data Flow Processor", 25 years of the International Symposia on Computer Architecture (selected papers), Association for Computing Machinery, August 1998, pages 2-4
- [32] Dennis, J.B.: "Machines and Models for Parallel Computing", International Journal of Parallel Programming, June 1993
- [33] Dubois, M. and Thakkar, S.: "Cache Architectures in Tightly Coupled Multiprocessors", Computer Magazine, June 1990, pp. 9-11
- [34] Gaudiot, J.L. and Bic, L.: "Advanced Topics in Data-Flow Computing", Prentice Hall, 1st edition, 1991
- [35] Hammond, L.; Hubbert, B.A.; Siu, M.; Prabhu, M.K.; Chen, M. and Olukotun, K.: "The Stanford Hydra CMP", IEEE Computer, March-April 2000, pages 71-84
- [36] Handy, J.: "The Cache Memory Book: The authoritative reference on cache design", Academic Press, 2nd edition, 1998
- [37] Handy, J.: "The Cache Memory Book: The Authoritative Reference on Cache Design", Academic Press, 2nd edition, 1998
- [38] Hennessy, J. L. and Patterson, D. A.: "Computer Architecture: A quantitative approach", Morgan Kaufmann Publishers, 2nd edition, 1996
- [39] Hoare, C.A.R.: "Monitors: An Operating System Structuring Concept", Communications of the ACM, October 1974, pages 549-557
- [40] Hughes, C.J.; Pai, V.S.; Ranganathan, P. and Adve, S.V.: "Rsim: Simulating Shared-Memory Multiprocessors with ILP Processors", IEEE Computer, February 2002, pp 44-49
- [41] Jagannathan, R.: "Dataflow Models", Computer Science Laboratory, SRI International, draft chapter to appear in Parallel and Distributed Computing Handbook, McGraw-Hill, 1st edition, 1995
- [42] Johnson, K.: "Semi-C Reference Manual", Alewife Systems Memo #20, MIT Laboratory for Computer Science, version 0.6, Feb. 1992
- [43] Jump, J. R.: "YACSIM Reference Manual", Rice University Electrical and Computer Engineering Department, March 1993. Available at <http://www-ece.rice.edu/~rsim/rppt.html> (accessed November 2001)
- [44] Kai Li and Hudak, P.: "Memory Coherence in Shared Virtual Memory Systems", ACM Transactions on Computer Systems, November 1989, vol 7, num 4, pp. 321-359

- 
- [45] Kernighan, B. W. and Ritchie, D. M.: "The C Programming Language", Prentice Hall, 2nd edition, 1988
- [46] Kranz, D.; Lim, B.H.; Agarwal, A. and Yeung, D.: "Low-cost Support for Fine-Grain Synchronization in Multiprocessors", in *Multithreaded Computer Architecture: A Summary of the State of the Art*, Kluwer Academic Publishers, 1994, pages 139–166
- [47] Kroft, D.: "Lockup-Free Instruction Fetch/Prefetch Cache Organization", 25 years of the International Symposia on Computer Architecture (selected papers), Association for Computing Machinery, August 1998, pages 20-21
- [48] Kroft, D.: "Lockup-Free Instruction Fetch/Prefetch Cache Organization", 25 years of the International Symposia on Computer Architecture (selected papers), Association for Computing Machinery, August 1998, pages 195-201
- [49] Kubiawicz, J.: "Modifications to SPARC for the Alewife Multiprocessor", Alewife Systems Memo #3, MIT Laboratory for Computer Science, Mar. 1990
- [50] Kubiawicz, J.: "Users Manual for the Alewife 1000 Controller", Alewife Systems Memo #19, MIT Laboratory for Computer Science, version 0.69, Dec. 1991
- [51] Lajolo, M.; Raghunathan, A. and Dey, S.: "Efficient power co-estimation techniques for system-on-chip design", *Proceedings of the Conference on Design, Automation and Test*, January 2000
- [52] Lim, B.H. and Agarwal, A.: "Reactive synchronization Algorithms for Multiprocessors", *Proceedings of the Sixth International Conference on Architectural Support for Programming Languages and Operating Systems*, 1994
- [53] Martínez, J.F. and Torrellas, J.: "Speculative Locks for Concurrent Execution of Critical Sections in Shared-Memory Multiprocessors", *International Symposium on Computer Architecture*, June 2001
- [54] Mehta, H.; Owens, R.M. and Irwin, M.J.: "Energy characterization based on clustering", *Proceedings of the 33rd Design Automation Conference*, June 1998
- [55] MIT Alewife Project, <http://www.cag.lcs.mit.edu/alewife/> (accessed November 2001)
- [56] Murdocca, M. J. and Huring, V. P.: "Principles of Computer Architecture", Prentice Hall, 1st edition, 2000
- [57] Nandy, S.K. and Narayan, R.: "An Incessantly Coherent Cache Scheme for Shared Memory Multithreaded Systems", Laboratory for Computer Science, Massachusetts Institute of Technology, Computation Structure Group Memo 356, 1994
- [58] Nussbaum, D.: "ASIM Reference Manual", Alewife Systems Memo #28, MIT Laboratory for Computer Science, November 1991
- [59] Pai, V.S.; Ranganathan, P. and Adve, S.V.: "RSIM Reference Manual", Department of Electrical and Computer Engineering, Rice University, version 1.0, August 1997



- 
- [60] Pai, V.S.; Ranganathan, P. and Adve, S.V.: "RSIM: An Execution-Driven Simulator for ILP-Based Shared-Memory Multiprocessors and Uniprocessors", Proceedings of the Third Workshop on Computer Architecture Education, February 1997
- [61] Pai, V.S.; Ranganathan, P. and Adve, S.V.: "RSIM: An Execution-Driven Simulator for ILP-Based Shared-Memory Multiprocessors and Uniprocessors", Department of Electrical and Computer Engineering, Rice University, Oct. 1997
- [62] Papadopoulos, G.M. and Culler, D.E.: "Monsoon: An Explicit Token-Store Architecture", 25 years of the International Symposia on Computer Architecture (selected papers), Association for Computing Machinery, August 1998, pages 74-76
- [63] QuickPower, <http://www.mentorgraphics.ca> (accessed May 2002)
- [64] RSIM Project, <http://rsim.cs.uiuc.edu/rsim/> (accessed October 2001)
- [65] Simunic, T.; Benini, L., De Micheli, G. and Hans, M.: "Source code optimization and profiling of energy consumption in embedded systems", Proceedings of the 13th Conference on International Symposium on System Synthesis, September 2000
- [66] Stanley-Marbell, P. and Hsiao, M.: "Fast, flexible, cycle-accurate energy estimation", Proceedings of the 2001 International Symposium on Low Power Electronics and Design, August 2001
- [67] Stenström, P.: "A Survey of Cache Coherence Schemes for Multiprocessors", Computer Magazine, June 1990, pp. 12-25
- [68] Synopsys PowerMill, [http://www.synopsys.com/products/etg/powermill\\_ds.html](http://www.synopsys.com/products/etg/powermill_ds.html) (accessed May 2002)
- [69] Tremblay, M. and O'Connor, J.M.: "UltraSparc I: A Four-Issue Processor Supporting Multimedia", IEEE Micro, 1996, pp. 42-50
- [70] Tullsen, D.M.; Eggers, S.J. and Levy, H.M.: "Simultaneous Multithreading: Maximizing On-Chip Parallelism", 25 years of the International Symposia on Computer Architecture (selected papers), Association for Computing Machinery, August 1998, pages 115-116
- [71] Tullsen, D.M.; Lo, J.L., Eggers, S.J. and Levy, H.M.: "Supporting Fine-Grained Synchronization on a Simultaneous Multithreading Processor", Proceedings of the 5th International Symposium on High Performance Computer Architecture, January 1999, pages 54-58
- [72] Vlassov, V. and Moritz, C.A.: "Efficient Fine Grained Synchronization Support Using Full/Empty Tagged Shared Memory and Cache Coherency", Technical Report TRITA-IT-R 00:04, Dept. of Teleinformatics, Royal Inst. of Technology, Dec. 2000
- [73] Weaver, D.L. and Germond, T.: "The SPARC Architecture Manual", PTR Prentice Hall, version 9, 1994
- [74] Xiaowei, Shen and Boon, S. Ang: "Implementing I-structures at Cache Level Coherence Level", MIT Laboratory for Computer Science, 1995

- [75] Xiaowei, Shen: "Implementing Global Cache Coherence In \*T-NG", MSc. Thesis at the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, May 1995
- [76] Ye, W.; Vijaykrishnan, N.; Kandemir, M. and Irwin, M.J.: "The design and use of SimplePower", Proceedings of the 37th Conference on Design Automation, June 2000
- [77] Yeager, K.C.: "The MIPS R10000 Superscalar Microprocessor", IEEE Micro, April 1996, pages 28-40
- [78] Yeung, D. and Agarwal, A.: "Experience with Fine-Grain Synchronization in MIMD Machines for Preconditioned Conjugate Gradient", Principles and Practice of Parallel Programming, 1993, pages 187-197

