

## **Cache Coherence and Synchronization-ISSUES**

**AUTHOR-DR.R.HEMAMALANI**  
**CO-AUTHOR-M.KAMARUNISHA**  
**R.JOTHI**

### **INTRODUCTION**

Cache coherence refers to the consistency and synchronization of data stored in different caches within a multiprocessor or multicore system. In such systems, each processor or core typically has its own cache memory to improve performance.

#### **Definition**

---

Coherence defines the behavior of reads and writes to a single address location.

One type of data occurring simultaneously in different cache memory is called cache coherence, or in some systems, global memory.

In a multiprocessor system, consider that more than one processor has cached a copy of the memory location X. The following conditions are necessary to achieve cache coherence:

1. In a read made by a processor P to a location X that follows a write by the same processor P to X, with no writes to X by another processor occurring between the write and the read instructions made by P, X must always return the value written by P.
2. In a read made by a processor P1 to location X that follows a write by another processor P2 to X, with no other writes to X made by any processor occurring between the two accesses and with the read and write being sufficiently separated, X must always return the value written by P2. This condition defines the concept of coherent view of memory. Propagating the writes to the shared memory location ensures that all the caches have a coherent view of the memory. If processor P1 reads the old value of X, even after the write by P2, we can say that the memory is incoherent.

The above conditions satisfy the Write Propagation criteria required for cache coherence. However, they are not sufficient as they do not satisfy the Transaction Serialization condition. To illustrate this better, consider the following example:

A multi-processor system consists of four processors - P1, P2, P3 and P4, all containing cached copies of a shared variable S whose initial value is 0. Processor P1 changes the value of S (in its cached copy) to 10 following which processor P2 changes the value of S in its own cached copy to 20. If we ensure only write propagation, then P3 and P4 will certainly see the changes made to S by P1 and P2. However, P3 may see the change made

by P1 after seeing the change made by P2 and hence return 10 on a read to *S*. P4 on the other hand may see changes made by P1 and P2 in the order in which they are made and hence return 20 on a read to *S*. The processors P3 and P4 now have an incoherent view of the memory.

Therefore, in order to satisfy Transaction Serialization, and hence achieve Cache Coherence, the following condition along with the previous two mentioned in this section must be met:

- Writes to the same location must be sequenced. In other words, if location *X* received two different values *A* and *B*, in this order, from any two processors, the processors can never read location *X* as *B* and then read it as *A*. The location *X* must be seen with values *A* and *B* in that order.<sup>[5]</sup>

The alternative definition of a coherent system is via the definition of sequential consistency memory model: "the cache coherent system must appear to execute all threads' loads and stores to a *single* memory location in a total order that respects the program order of each thread".<sup>[3]</sup> Thus, the only difference between the cache coherent system and sequentially consistent system is in the number of address locations the definition talks about (single memory location for a cache coherent system, and all memory locations for a sequentially consistent system).

Another definition is: "a multiprocessor is cache consistent if all writes to the same memory location are performed in some sequential order".<sup>[6]</sup>

Rarely, but especially in algorithms, coherence can instead refer to the locality of reference. Multiple copies of same data can exist in different cache simultaneously and if processors are allowed to update their own copies freely, an inconsistent view of memory can result.

### The Cache Coherence Problem

In a multiprocessor system, data inconsistency may occur among adjacent levels or within the same level of the memory hierarchy. For example, the cache and the main memory may have inconsistent copies of the same object.

As multiple processors operate in parallel, and independently multiple caches may possess different copies of the same memory block, this creates **cache coherence problem**. **Cache coherence schemes** help to avoid this problem by maintaining a uniform state for each cached block of data.

Let *X* be an element of shared data which has been referenced by two processors, P1 and P2. In the beginning, three copies of *X* are consistent. If the processor P1 writes a new data *X1* into the cache, by using **write-through policy**, the same copy will be written immediately into the shared memory. In this case, inconsistency occurs between cache memory and the main memory. When a **write-back policy** is used, the main memory will be updated when the modified data in the cache is replaced or invalidated.

In general, there are three sources of inconsistency problem –

- Sharing of writable data
- Process migration
- I/O activity

### Snoopy Bus Protocols

Snoopy protocols achieve data consistency between the cache memory and the shared memory through a bus-based memory system. **Write-invalidate** and **write-update** policies are used for maintaining cache consistency.

In this case, we have three processors P1, P2, and P3 having a consistent copy of data element 'X' in their local cache memory and in the shared memory (Figure-a). Processor P1 writes X1 in its cache memory using **write-invalidate protocol**. So, all other copies are invalidated via the bus. It is denoted by 'I' (Figure-b). Invalidated blocks are also known as **dirty**, i.e. they should not be used. The **write-update protocol** updates all the cache copies via the bus. By using **write back cache**, the memory copy is also updated (Figure-c).

### Cache Events and Actions

Following events and actions occur on the execution of memory-access and invalidation commands –

- **Read-miss** – When a processor wants to read a block and it is not in the cache, a read-miss occurs. This initiates a **bus-read** operation. If no dirty copy exists, then the main memory that has a consistent copy, supplies a copy to the requesting cache memory. If a dirty copy exists in a remote cache memory, that cache will restrain the main memory and send a copy to the requesting cache memory. In both the cases, the cache copy will enter the valid state after a read miss.
- **Write-hit** – If the copy is in dirty or **reserved** state, write is done locally and the new state is dirty. If the new state is valid, write-invalidate command is broadcasted to all the caches, invalidating their copies. When the shared memory is written through, the resulting state is reserved after this first write.
- **Write-miss** – If a processor fails to write in the local cache memory, the copy must come either from the main memory or from a remote cache memory with a dirty block. This is done by sending a **read-invalidate** command, which will invalidate all cache copies. Then the local copy is updated with dirty state.
- **Read-hit** – Read-hit is always performed in local cache memory without causing a transition of state or using the snoopy bus for invalidation.
- **Block replacement** – When a copy is dirty, it is to be written back to the main memory by block replacement method. However, when the copy is either in valid or reserved or invalid state, no replacement will take place.

## Directory-Based Protocols

By using a multistage network for building a large multiprocessor with hundreds of processors, the snoopy cache protocols need to be modified to suit the network capabilities. Broadcasting being very expensive to perform in a multistage network, the consistency commands is sent only to those caches that keep a copy of the block. This is the reason for development of directory-based protocols for network-connected multiprocessors.

In a directory-based protocols system, data to be shared are placed in a common directory that maintains the coherence among the caches. Here, the directory acts as a filter where the processors ask permission to load an entry from the primary memory to its cache memory. If an entry is changed the directory either updates it or invalidates the other caches with that entry.

## Hardware Synchronization Mechanisms

Synchronization is a special form of communication where instead of data control, information is exchanged between communicating processes residing in the same or different processors.

Multiprocessor systems use hardware mechanisms to implement low-level synchronization operations. Most multiprocessors have hardware mechanisms to impose atomic operations such as memory read, write or read-modify-write operations to implement some synchronization primitives. Other than atomic memory operations, some inter-processor interrupts are also used for synchronization purposes.

## Cache Coherency in Shared Memory Machines

Maintaining cache coherency is a problem in multiprocessor system when the processors contain local cache memory. Data inconsistency between different caches easily occurs in this system.

The major concern areas are –

- Sharing of writable data
- Process migration
- I/O activity

## Sharing of writable data

When two processors (P1 and P2) have same data element (X) in their local caches and one process (P1) writes to the data element (X), as the caches are write-through local cache of P1, the main memory is also updated. Now when P2 tries to read data element (X), it does not find X because the data element in the cache of P2 has become outdated.

### Process migration

In the first stage, cache of P1 has data element X, whereas P2 does not have anything. A process on P2 first writes on X and then migrates to P1. Now, the process starts reading data element X, but as the processor P1 has outdated data the process cannot read it. So, a process on P1 writes to the data element X and then migrates to P2. After migration, a process on P2 starts reading the data element X but it finds an outdated version of X in the main memory.

### I/O activity

As illustrated in the figure, an I/O device is added to the bus in a two-processor multiprocessor architecture. In the beginning, both the caches contain the data element X. When the I/O device receives a new element X, it stores the new element directly in the main memory. Now, when either P1 or P2 (assume P1) tries to read element X it gets an outdated copy. So, P1 writes to element X. Now, if I/O device tries to transmit X it gets an outdated copy.

### Uniform Memory Access (UMA)

Uniform Memory Access (UMA) architecture means the shared memory is the same for all processors in the system. Popular classes of UMA machines, which are commonly used for (file-) servers, are the so-called Symmetric Multiprocessors (SMPs). In an SMP, all system resources like memory, disks, other I/O devices, etc. are accessible by the processors in a uniform manner.

### Non-Uniform Memory Access (NUMA)

In NUMA architecture, there are multiple SMP clusters having an internal indirect/shared network, which are connected in scalable message-passing network. So, NUMA architecture is logically shared physically distributed memory architecture.

In a NUMA machine, the cache-controller of a processor determines whether a memory reference is local to the SMP's memory or it is remote. To reduce the number of remote memory accesses, NUMA architectures usually apply caching processors that can cache the remote data. But when caches are involved, cache coherency needs to be maintained. So these systems are also known as CC-NUMA (Cache Coherent NUMA).

### Cache Only Memory Architecture (COMA)

COMA machines are similar to NUMA machines, with the only difference that the main memories of COMA machines act as direct-mapped or set-associative caches. The data blocks are hashed to a location in the DRAM cache according to their addresses. Data that is fetched remotely is actually stored in the local main memory. Moreover, data blocks do not have a fixed home location, they can freely move throughout the system.

COMA architectures mostly have a hierarchical message-passing network. A switch in such a tree contains a directory with data elements as its sub-tree. Since data has no home location, it must be explicitly searched for. This means that a remote access requires a traversal along the switches in the tree to search their directories for the required data. So, if a switch in the network receives multiple requests from its subtree for the same data, it combines them into a single request which is sent to the parent of the switch. When the requested data returns, the switch sends multiple copies of it down its subtree.

### COMA versus CC-NUMA

Following are the differences between COMA and CC-NUMA.

- COMA tends to be more flexible than CC-NUMA because COMA transparently supports the migration and replication of data without the need of the OS.
- COMA machines are expensive and complex to build because they need non-standard memory management hardware and the coherency protocol is harder to implement.
- Remote accesses in COMA are often slower than those in CC-NUMA since the tree network needs to be traversed to find the data.

### REFERENCES-

1. *E. Thomadakis, Michael (2011). The Architecture of the Nehalem Processor and Nehalem-EP SMP Platforms (PDF). Texas A&M University. p. 30. Archived from the original (PDF) on 2014-08-11.*
2. ^ Jump up to:<sup>a b</sup> Yan, Solihin. *Fundamentals of parallel multicore architecture*. OCLC 884540034.
3. ^ Jump up to:<sup>a b</sup> Sorin, Daniel J.; Hill, Mark D.; Wood, David Allen (2011-01-01). *A primer on memory consistency and cache coherence*. Morgan & Claypool Publishers. OCLC 726930429.
4. ^ Jump up to:<sup>a b c</sup> Patterson and Hennessy. *Computer Organization and Design - 4th Edition*. ISBN 978-0-12-374493-7.
5. ^ Neupane, Mahesh (April 16, 2004). "Cache Coherence" (PDF). Archived from the original (PDF) on 20 June 2010.
6. ^ Steinke, Robert C.; Nutt, Gary J. (2004-09-01). "A Unified Theory of Shared Memory Consistency". *J. ACM*. **51** (5): 800–

849. *arXiv:cs/0208027*. doi:10.1145/1017460.1017464. ISSN 0004-5411. S2CID 3206071.
7. ^ "Ravishankar, Chinya; Goodman, James (February 28, 1983). "Cache Implementation for Multiple Microprocessors"" (PDF). *Proceedings of IEEE COMPCON*: 346–350.
  8. ^ RasmusUlfsnes (June 2013). "Design of a Snoop Filter for Snoop-Based Cache Coherency Protocols" Archived 2014-02-01 at the Wayback Machine (PDF). *diva-portal.org*. Norwegian University of Science and Technology. Retrieved 2014-01-20.
  9. ^ "Lecture 18: Snooping vs. Directory Based Coherency" (PDF). *Berkeley.edu*. Retrieved 14 May 2023.
  10. ^ Kriouile (16 September 2013). *Formal Analysis of the ACE Specification for Cache Coherent Systems-on-Chip*. In *Formal Methods for Industrial Critical Systems*. Springer Berlin Heidelberg. ISBN 978-3-642-41010-9.
  11. ^ Ltd, Arm. "AMBA / AMBA 5". *Arm Developer*. Retrieved 2021-04-27.