

L04a. Shared Memory Machines

Shared Memory Machine Model:

- Dance Hall Architecture:
 - CPUs on one side of the interconnection network, and memory on the other side.
 - The entire memory space is accessible from any CPU.
- Symmetric Multiprocessor (SMP) Architecture:
 - The CPUs are connected to the memory via a system bus.
 - The access time from any CPU to the memory is the same.
 - Each CPU is equipped with a private cache.
- Distributed Shared Memory (DSM) Architecture:
 - Each CPU has a piece of memory connected to it to facilitate faster access.
 - Each CPU can access other memory spaces through the interconnection network.
 - Each CPU is equipped with a private cache.



Memory Consistency Model:

- A Memory Consistency Model formulates what can be expected from the system.
- Sequential Consistency (SC):
 - Program order: Memory accesses from each process are ordered.
 - Arbitrary interleaving: The interaction between processes in terms of memory accesses is arbitrary.

Cache Coherence:

- How the system implements the Memory Consistency Model in the presence of private caches.
- Non-cache-coherent (NCC) multiprocessor: This model grants a shared memory space to all the CPUs, while leaving the private cache coherence issue to be dealt with by software.
- Cache-coherent (CC) multiprocessor:
 - Write-invalidate: Whenever a CPU writes to a memory location in its cache that happens to be fetched by other CPUs, the HW broadcasts a message on the system bus to invalidate other CPUs' caches.
 - Write-update: Whenever a CPU writes to a memory location in its cache that happens to be fetched by other CPUs, the HW broadcasts an update on the system bus to replicate this change on other CPUs' caches.