This is the second lecture of Chapter 9

Chapter 9

Alternative Architectures (**B**)

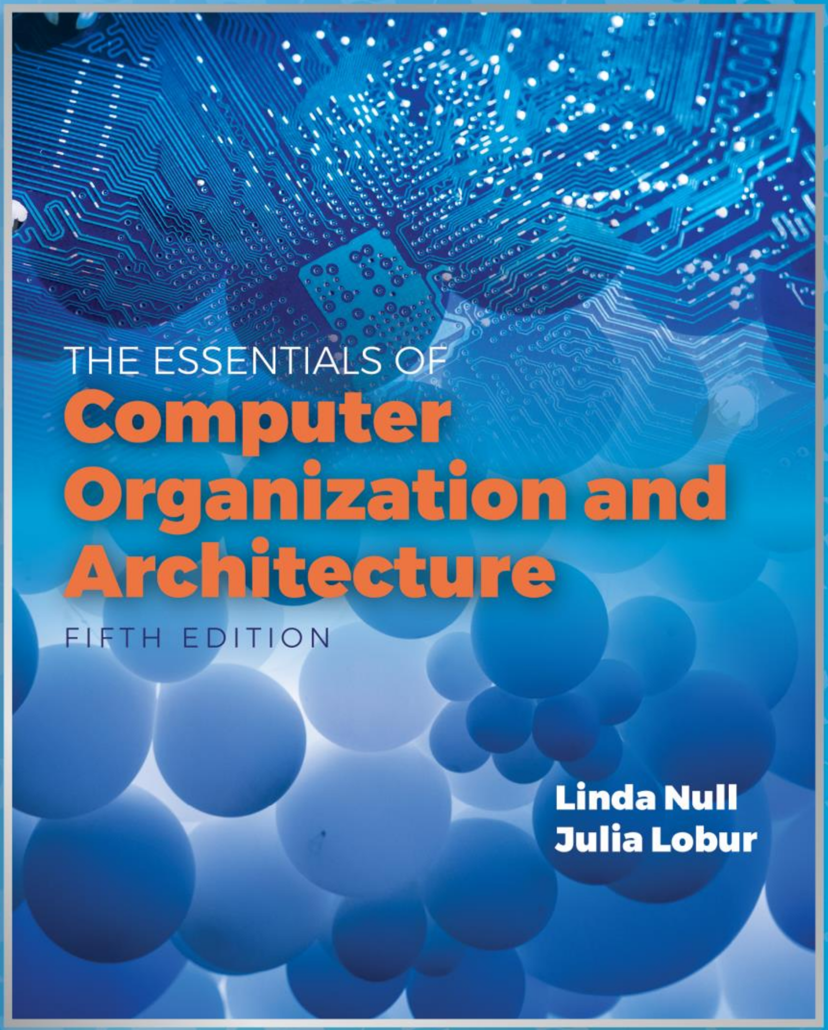

Quick review of last lecture

- RISC Machines
  - RISC vs. CISC,
  - Overlapping Register Windows in RSIC
- Flynn’s Taxonomy
  - Date-driven:
    - Data flow
  - Instruction-driven:
    - SISD, SIMD, MISD, **MIMD**, SPMD
    - Shared Memory, Distributed Memory

9.4 Parallel and Multiprocessor Architectures (1 of 21)

- Parallel processing is capable of economically increasing system throughput while providing better fault tolerance.
- The limiting factor is that no matter how well an algorithm is parallelized, there is always some portion that must be done sequentially.
  - Additional processors sit idle while the sequential work is performed.
- Thus, it is important to keep in mind that an n-fold increase in processing power does not necessarily result in an n-fold increase in throughput.

9.4 Parallel and Multiprocessor Architectures (2 of 21)

9.4.1 Superscalar and VLIW

- Recall that pipelining divides the fetch-decode-execute cycle into stages that each carry out a small part of the process on a set of instructions.
- Ideally, an instruction exits the pipeline during each tick of the clock.
- *Superpipelining* occurs when a pipeline has stages that require less than half a clock cycle to complete.
  - The pipeline is equipped with a separate clock running at a frequency that is at least double that of the main system clock.
- Superpipelining is only one aspect of superscalar design.

9.4 Parallel and Multiprocessor Architectures (3 of 21)

- Superscalar architectures include multiple execution units such as specialized integer and floating-point adders and multipliers.
- A critical component of this architecture is the *instruction fetch unit,* which can simultaneously retrieve several instructions from memory.
- A *decoding unit* determines which of these instructions can be executed in parallel and combines them accordingly.
- This architecture also requires compilers that make optimum use of the hardware.

9.4 Parallel and Multiprocessor Architectures (4 of 21)

- Very long instruction word (VLIW) architectures differ from superscalar architectures because the VLIW compiler, instead of a hardware decoding unit, packs independent instructions into one long instruction that is sent down the pipeline to the execution units.
- One could argue that this is the best approach because the compiler can better identify instruction dependencies.
- However, compilers tend to be conservative and cannot have a view of the run time code.

9.4 Parallel and Multiprocessor Architectures (5 of 21)

9.4.2 Vector Processors

- Vector computers are processors that operate on entire vectors or matrices at once.
  - These systems are often called supercomputers.
- Vector computers are highly pipelined so that arithmetic instructions can be overlapped.
- Vector processors can be categorized according to how operands are accessed.
  - Register-register vector processors require all operands to be in registers.
  - Memory-memory vector processors allow operands to be sent from memory directly to the arithmetic units.

9.4 Parallel and Multiprocessor Architectures (6 of 21)

- A disadvantage of register-register vector computers is that large vectors must be broken into fixed-length segments so they will fit into the register sets.
- Memory-memory vector computers have a longer startup time until the pipeline becomes full.
- In general, vector machines are efficient because there are fewer instructions to fetch, and corresponding pairs of values can be prefetched because the processor knows it will have a continuous stream of data.

9.4 Parallel and Multiprocessor Architectures (7 of 21)

9.4.3 Interconnection Networks

- MIMD systems can communicate through shared memory or through an interconnection network.
- Interconnection networks are often classified according to their topology, routing strategy, and switching technique.
- Of these, the topology is a major determining factor in the overhead cost of message passing.
- Message passing takes time owing to network latency and incurs overhead in the processors.

9.4 Parallel and Multiprocessor Architectures (8 of 21)

- Interconnection networks can be either static or dynamic.
- Processor-to-memory connections usually employ dynamic interconnections. These can be blocking or nonblocking.
  - Nonblocking interconnections allow connections to occur simultaneously.
- Processor-to-processor message-passing interconnections are usually static, and can employ any of several different topologies, as shown on the following slide.

9.4 Parallel and Multiprocessor Architectures (9 of 21)

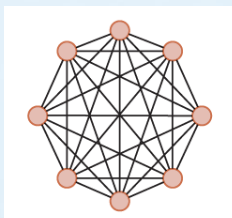

Completely Connected

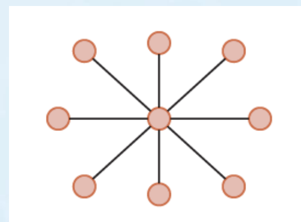

Star

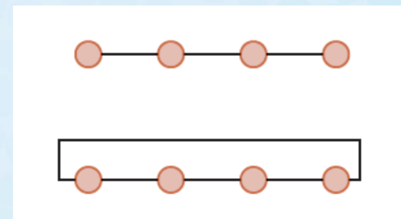

Linear and Ring

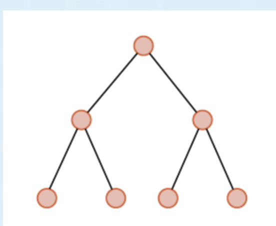

Tree

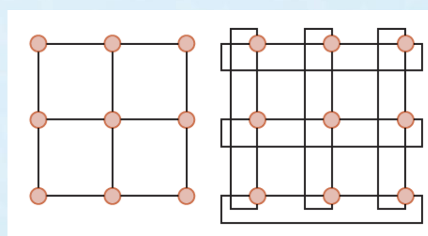

Mesh and Mesh Ring

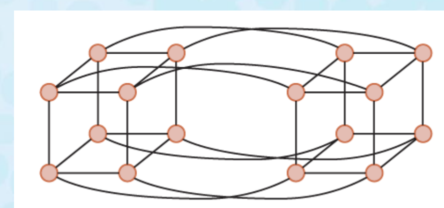

Four-Dimensional Hypercube

9.4 Parallel and Multiprocessor Architectures (9 of 21)

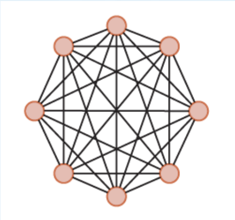

Completely Connected

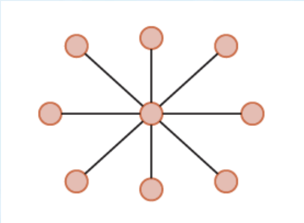

Star

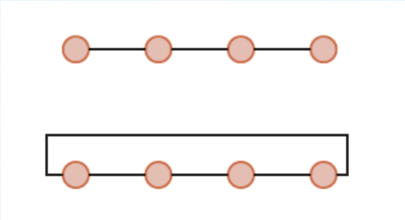

Linear and Ring

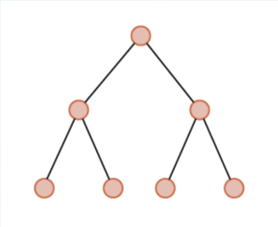

Tree

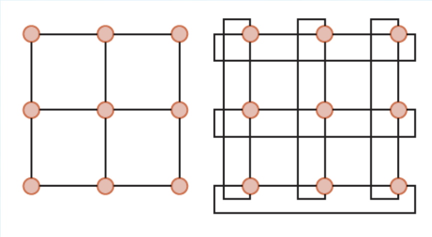

Mesh and Mesh Ring

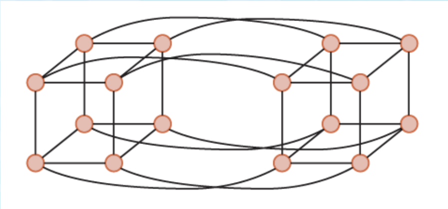

Four-Dimensional Hypercube

9.4 Parallel and Multiprocessor Architectures (10 of 21)

- Dynamic routing is achieved through buses or switching networks that consist of crossbar switches or 2 × 2 switches.

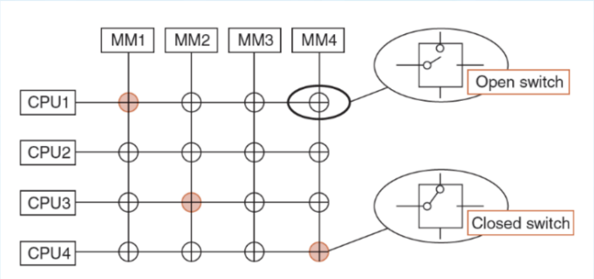

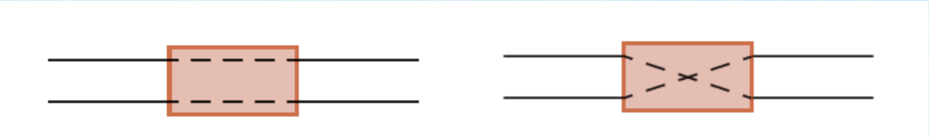

Through

Cross

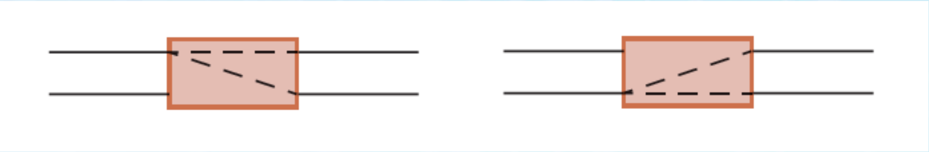

Upper Broadcast

Lower Broadcast

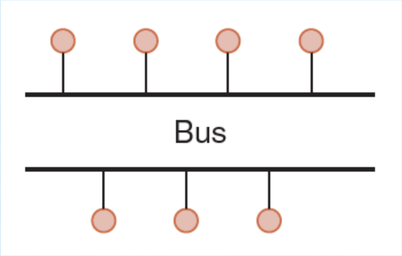

A Bus-Based Network

9.4 Parallel and Multiprocessor Architectures (11 of 21)

- Multistage interconnection (or shuffle) networks are the most advanced class of switching networks.
- They can be used in loosely-coupled distributed systems, or in tightly-coupled processor-to-memory configurations.

A Two-Stage Omega Network

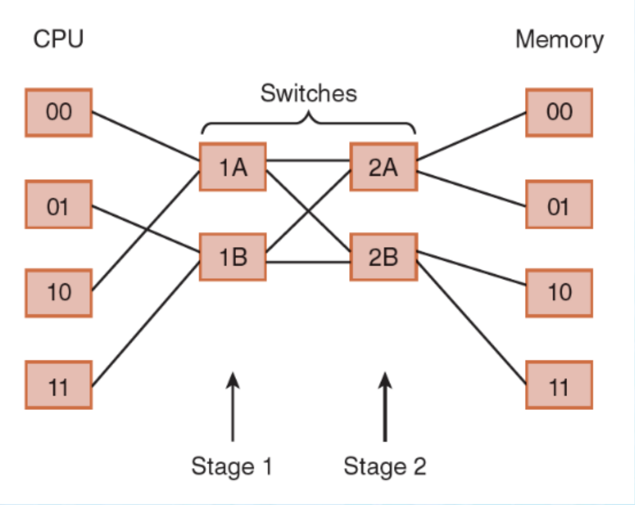

9.4 Parallel and Multiprocessor Architectures (12 of 21)

- There are advantages and disadvantages to each switching approach.
  - Bus-based networks, while economical, can be bottlenecks. Parallel buses can alleviate bottlenecks, but are costly.
  - Crossbar networks are nonblocking, but require n^2 switches to connect n entities.
  - Omega networks are blocking networks, but exhibit less contention than bus-based networks. They are somewhat more economical than crossbar networks, n nodes needing $\log_2 n$ stages with n / 2 switches per stage.

An 8×8 Omega Network

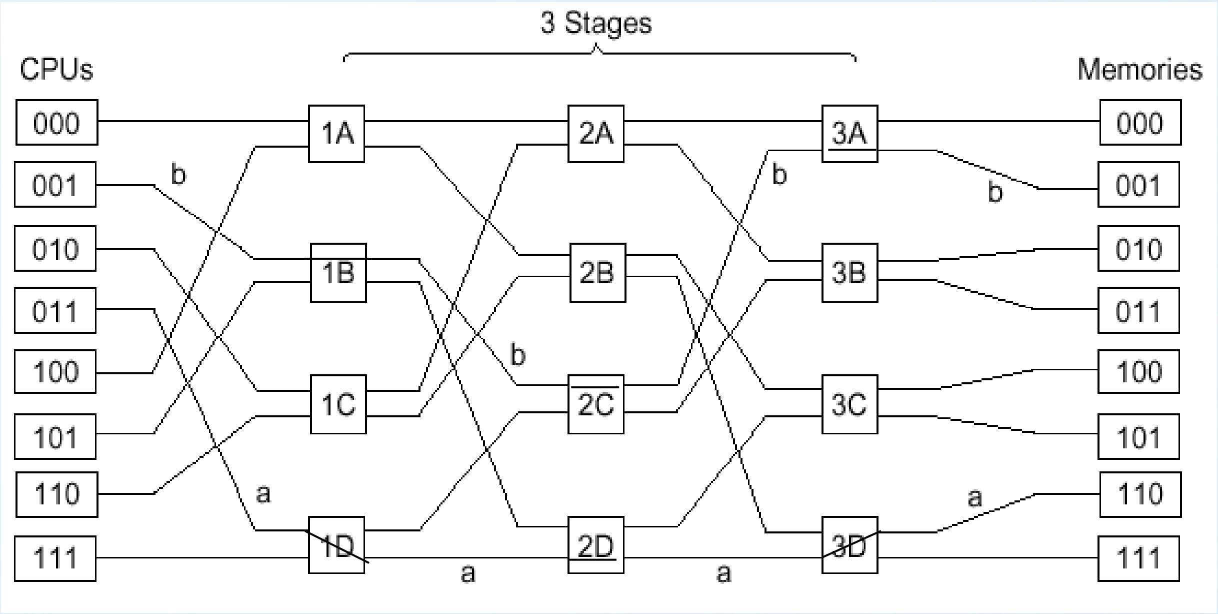

In n×n Omega Network

- $\log_2 n$ stages
- n/2 2×2 switches per stage
- Perfect shuffle ISC

Routing

- a: 011→110
- b: 001→001

The use of 2×2 switches

Straight　Crossover　Upper broadcast　Lower broadcast

Routing in Omega Network

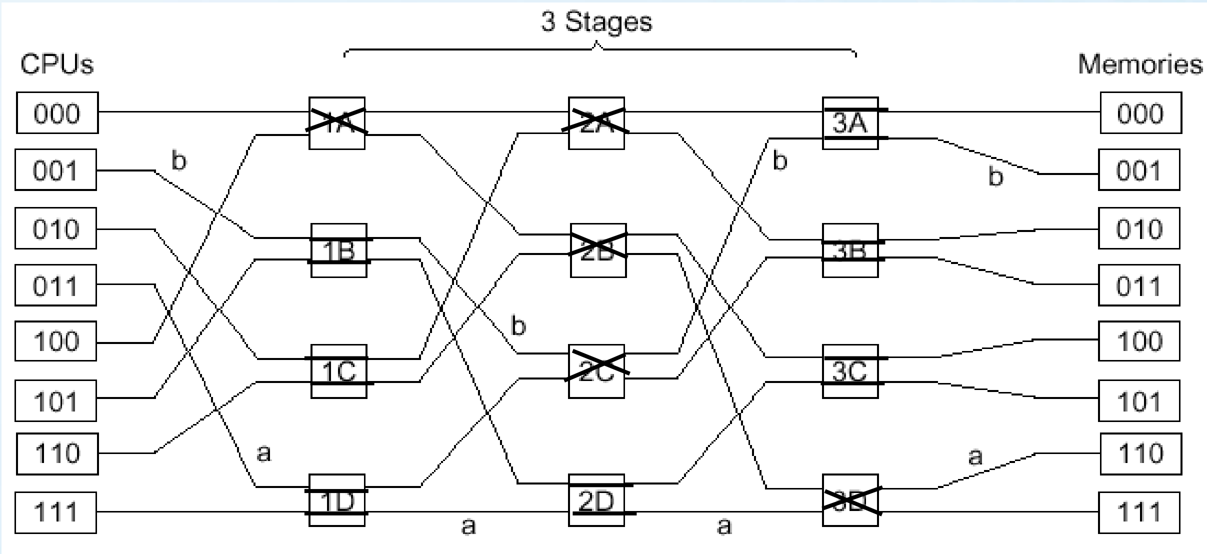

000 → 111
001 → 011
010 → 000
011 → 001
100 → 010
101 → 101
110 → 100
111 → 110

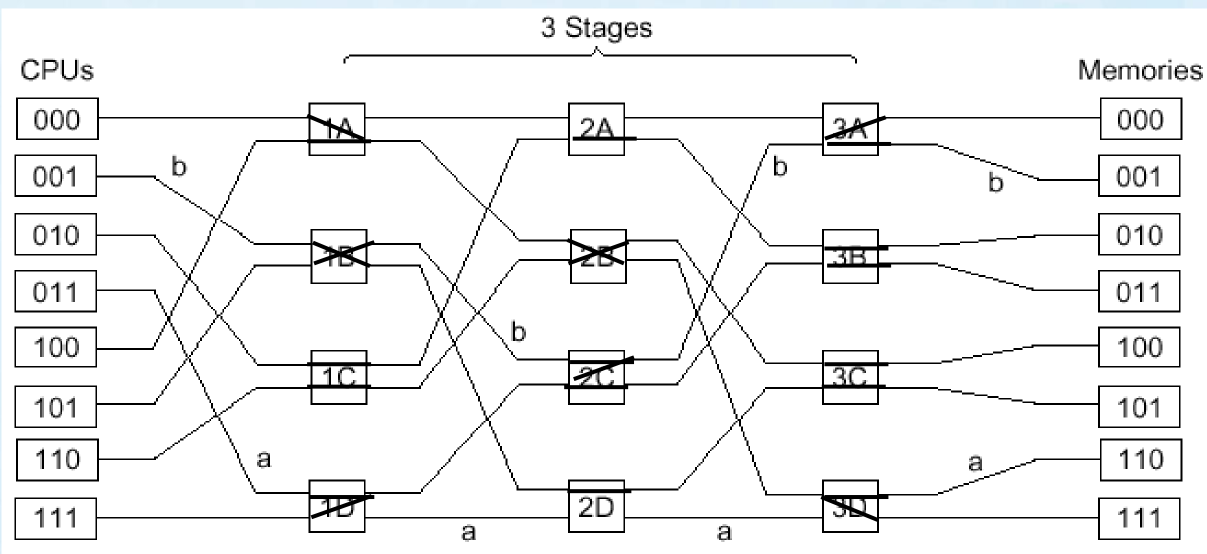

000 → 110
001 → 101
010 → 010
011 → 000
100 → 111
101 → 001
110 → 100
111 → 011

9.4 Parallel and Multiprocessor Architectures (13 of 21)

9.4.4 Shared Memory Multiprocessors

- Tightly-coupled multiprocessor systems use the same memory. They are also referred to as shared memory multiprocessors.
- The processors do not necessarily have to share the same block of physical memory.
- Each processor can have its own memory, but it must share it with the other processors.
- Configurations such as these are called *distributed shared memory multiprocessors*.