

More on Parallel Architectures

COMP375 Computer Architecture
and Organization

Goals

The student should be able to:

- Explain the advantages and disadvantages of different parallel architectures.
- Determine the software environment most appropriate for each architecture.
- Determine the communication cost for various message passing architectures.

Flynn's Parallel Classification

- **SISD** – Single Instruction Single Data
 - standard uniprocessors
- **SIMD** – Single Instruction Multiple Data
 - vector and array processors
- **MISD** - Multiple Instruction Single Data
 - systolic processors
- **MIMD** – Multiple Instruction Multiple Data


MIMD Computers

- Shared Memory
 - Attached processors
 - Symmetric Multi-Processor (SMP)
 - Multiple processor systems
 - Multiple instruction issue processors
 - Multi-Core processors
 - NUMA
- Separate Memory
 - Clusters
 - Message passing multiprocessors
 - Hypercube
 - Mesh

Symmetric Multi-Processors (SMP)

- The system has multiple CPU chips.
- All CPUs are identical and can access all of the common memory.
- Any thread can execute on any CPU.
- Only one copy of the OS is required.
- Many SMP computers are available for servers or workstations.
- Depending on the design, external interrupts can be handled by a particular CPU or by any CPU.

Shared Memory Multiprocessor




Cache Coherence

- While main memory is shared, each processor has its own cache.
- The same data values can be in both caches.
- If a processor changes a data value, the value in the other processor will not reflect the correct logical value.
- Not caching shared values will correct this problem, but will reduce performance.

Snoopy Cache


- The cache watches the bus to see all reads and writes from the other processor.
- If the other CPU writes to an address in the local cache, the local cache removes the value from cache.
- If the other CPU reads an address in the local CPU's cache and the value has been changed, the local CPU intercepts the read and sends its data to both the memory and other CPU.

Cache Coherence



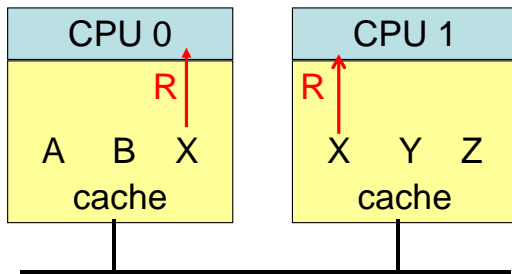
Dual processors with a shared value in each cache. RAM is not shown.

Cache Coherence



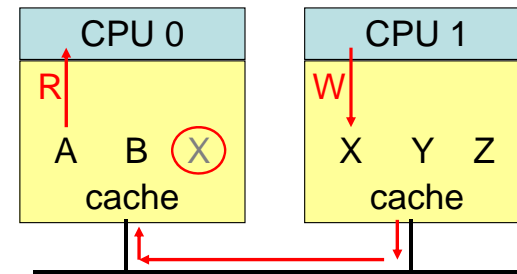
Non-interfering reads and writes

Cache Coherence




Non-interfering simultaneous reads

Cache Coherence




Write invalidates the same address in the other cache

Cache Coherence



Non-interfering reads and writes

Cache Coherence




Read of a value in the other cache causes the value to be transferred from that cache.

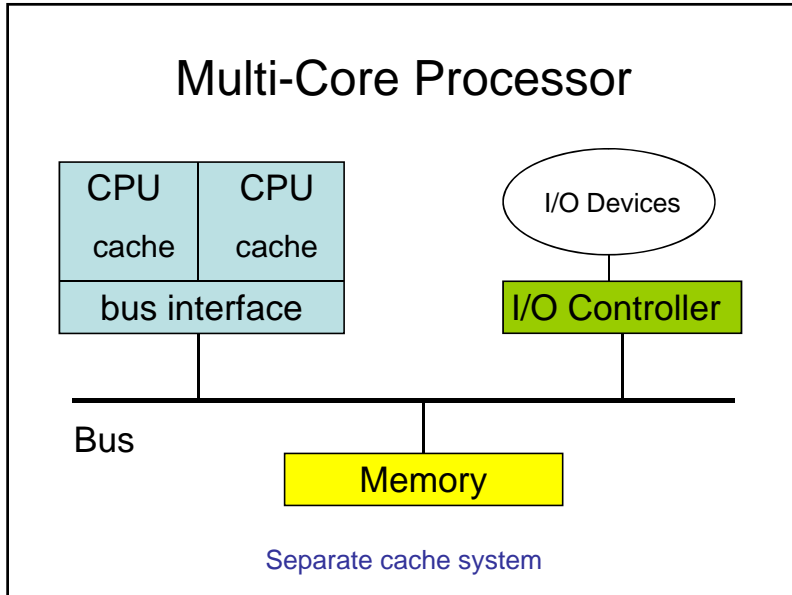
Multi-Core Processors

- A multi-core system has two CPUs on the same chip.
- Some multi-core systems share the cache while others have separate caches.
- Cache coherence is easier to resolve with both caches on the same chip. Communication on the bus is not required.

Multi-Core Processor



Shared cache system



Intel® Pentium® processor Extreme Edition

A dual-core and hyper-threaded processor

images from www.intel.com

Moore's Law

- Moore's Law states that the number of transistors on a chip will double every 18 months.
- Intel currently makes quad-core processors
- The number of "cores" on a chip is directly related to the number of transistors.


$$\#cores = \#now * 2^{years/1.5}$$

If the number of possible cores is described by $4 * 2^{years/1.5}$ how many cores will be possible in 6 years.

1. 16
2. 24
3. 32
4. 64
5. 128

Incentive for Dual Core


- Intel reports that underclocking a single core by 20 percent saves half the power while sacrificing just 13 percent of the performance.



Source: IEEE Spectrum April, 2008


If a single core can run 100 MIPS, what is the maximum throughput of a dual core if each core executes 13% slower?

- 87 MIPS
- 100 MIPS
- 113 MIPS
- 174 MIPS
- 200 MIPS



Amdahl's law

- The speedup of a program using multiple processors in parallel computing is limited by the time needed for the sequential fraction of the program.
- named after computer architect Gene Amdahl



- If 95% of a program can be parallelized, the theoretical maximum speedup using parallel computing would be 20x no matter how many processors are used

Maximum Speedup

- If you can parallelize a fraction **P** of a program, then **1-P** must be sequential.
- If you have **N** processors, the maximum speedup is

$$\text{speedup} = \frac{1}{1 - P + \frac{P}{N}}$$

Comparing Multiprocessors

- If you only have one thread to execute, no multiprocessor is advantageous.
- Systems with separate cache run best when the threads are unrelated.
- Systems with shared cache run best when each thread accesses the same addresses.
- Multiple Instruction Issue shares the resources of one processor with multiple threads.

Separate Memory Systems

- Each processor has its own memory that none of the other processors can access.
- Each processor has unhindered access to its memory.
- No cache coherence problems.
- Scalable to very large numbers of processors
- Communication is done by passing messages over special I/O channels

Virtual Shared Memory

- In virtual shared memory several processors appear to share the same large memory address space.
- Memory is divided into pages with some pages stored locally and some remote.
- When a remote page is accessed, a page fault is created. The requested page is sent to the processor.
- Requires no additional HW.

BlueGene/L Supercomputer


- World's fastest computer (<http://www.top500.org>)
- 596 TeraFLOPS measured
- 73.7 Terabytes of RAM
- Located at the Lawrence Livermore National Laboratory
- Uses about \$1M a year in electricity
- Occupies 2,500 square feet
- Faster than Japan's Earth Simulator

BlueGene/L



If your PC can execute 100MegaFLOPS,
how many PCs are required to generate
596 TeraFLOPS?

1. 596
2. 596,000
3. 5,960,000
4. 59,600,000
5. 596,000,000



BlueGene/L Architecture

- MIMD non-shared memory design
- 64K IBM PowerPC processors with communications built on the chip
- 64 x 32 x 32 three-dimensional torus of compute nodes.
- 1.4 Gb/sec inter-node communication
- 256 MB of memory per node.

BlueGene/L Composition

