



Parallel Computing



Chris 2008/03/14


Outlines

- ▶ **Overview**
 - ▶ What Is Parallel Computing?
 - ▶ Why Use Parallel Computing?
- ▶ **Concepts and terminology**
 - ▶ SISD, SIMD, MISD, MIMD
 - ▶ Amdahl's Law and Gustafson's Law
- ▶ **Parallel Computer Memory Architecture**
 - ▶ Shared Memory
 - ▶ Distributed Memory
 - ▶ Hybrid Distributed-Shared Memory




What is parallel computing ?

- ▶ Traditionally, software has been written for **serial** computation:
 - ▶ To be run on a single computer having a single Central Processing Unit (CPU);
 - ▶ A **problem** is broken into a discrete series of **instructions**.
 - ▶ **Instructions are executed one after another.**
 - ▶ Only **one instruction** may execute at any moment in time



What is parallel computing (cont.)?

- ▶ In the simplest sense, **parallel computing** is the simultaneous use of multiple compute resources to solve a computational problem.
 - ▶ To be run using multiple CPUs
 - ▶ A **problem** is broken into discrete **parts** that can be solved concurrently
 - ▶ Each part is further broken down to a series of **instructions**
 - ▶ **Instructions from each part execute simultaneously on different CPUs**



Why use parallel computing ?

- ▶ The primary reasons for using parallel computing:
 - ▶ Save **time** - wall clock time
 - ▶ Solve **larger** problems
 - ▶ Provide **concurrency** (do multiple things at the same time)
- ▶ Other reasons might include:
 - ▶ Taking advantage of **non-local resources** - using available compute resources on a wide area network, or even the Internet when local compute resources are scarce.
 - ▶ Cost savings - using multiple "**cheap**" computing resources instead of paying for time on a supercomputer.
 - ▶ **Overcoming memory constraints** - single computers have very finite memory resources. For large problems, using the memories of multiple computers may overcome this obstacle.



Outlines

- ▶ **Overview**
 - ▶ What Is Parallel Computing?
 - ▶ Why Use Parallel Computing?
- ▶ **Concepts and terminology**
 - ▶ SISD, SIMD, MISD, MIMD
 - ▶ Amdahl's Law and Gustafson's Law
- ▶ **Parallel Computer Memory Architecture**
 - ▶ Shared Memory
 - ▶ Distributed Memory
 - ▶ Hybrid Distributed-Shared Memory




Concepts and terminology

- ▶ There are different ways to classify parallel computers. One of the more widely used classifications, in use since 1966, is called **Flynn's Taxonomy**.
- ▶ Flynn's taxonomy distinguishes multi-processor computer architectures according to how they can be classified along the two independent dimensions of ***Instruction*** and ***Data***. Each of these dimensions can have only one of two possible states: ***Single*** or ***Multiple***.

Flynn's taxonomy		
	Single Instruction	Multiple Instruction
Single Data	SISD	MISD
Multiple Data	SIMD	MIMD

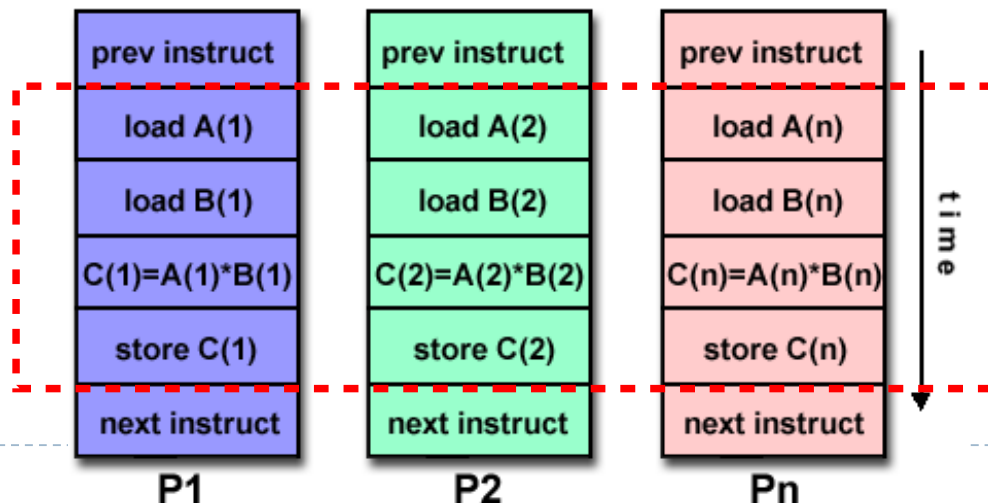
Single Instruction, Single Data (SISD)

- ▶ A serial (non-parallel) computer
- ▶ Single instruction: only one instruction stream is being acted on by the CPU during any one clock cycle
- ▶ Single data: only one data stream is being used as input during any one clock cycle
- ▶ Deterministic execution
- ▶ This is the oldest and until recently, the most prevalent form of computer
- ▶ Examples: most PCs, single CPU workstations and mainframes



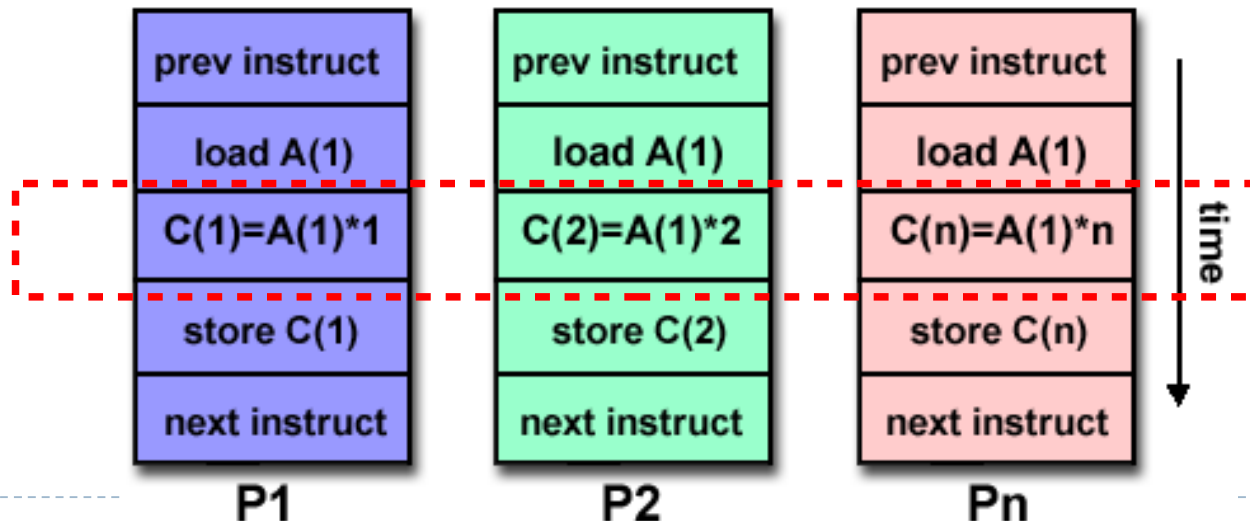
Single Instruction, Multiple Data (SIMD)

- ▶ A type of parallel computer
- ▶ Single instruction: All processing units execute the same instruction at any given clock cycle
- ▶ Multiple data: Each processing unit can operate on a different data element
- ▶ This type of machine typically has an instruction dispatcher, a very high-bandwidth internal network, and a very large array of very small-capacity instruction units.
- ▶ Best suited for specialized problems characterized by a high degree of regularity, such as image processing.
- ▶ Synchronous (lockstep) and deterministic execution
- ▶ Two varieties: Processor **Arrays** and Vector Pipelines



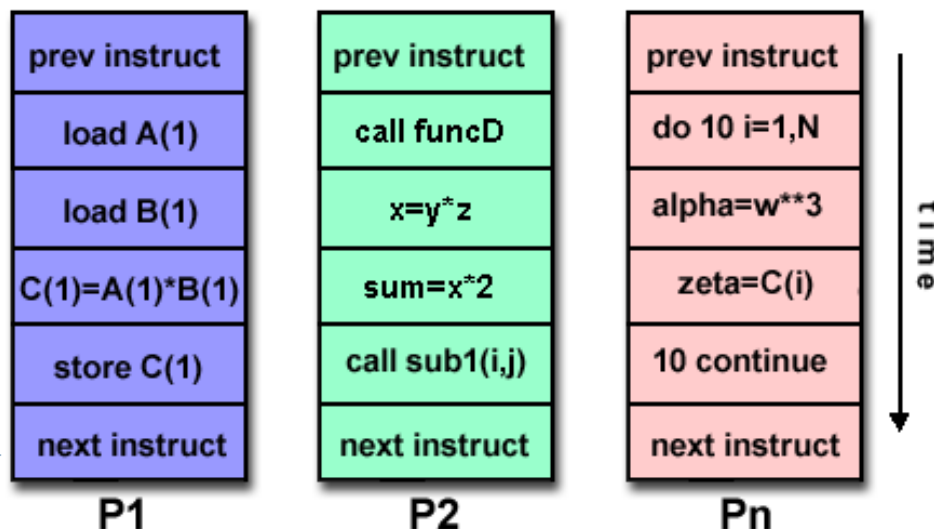
Multiple Instruction, Single Data (MISD)

- ▶ A single data stream is fed into multiple processing units.
- ▶ Each processing unit operates on the data independently via independent instruction streams.
- ▶ Few actual examples of this class of parallel computer have ever existed. One is the experimental Carnegie-Mellon C.mmp computer (1971).
- ▶ Some conceivable uses might be:
 - ▶ multiple frequency filters operating on a single signal stream
 - ▶ multiple cryptography algorithms attempting to crack a single coded message



Multiple Instruction, Multiple Data (MIMD)

- ▶ Currently, **the most common type of parallel computer**. Most modern computers fall into this category.
- ▶ Multiple Instruction: every processor may be executing a **different instruction stream**
- ▶ Multiple Data: every processor may be working with a **different data stream**
- ▶ Execution can be synchronous or asynchronous, deterministic or non-deterministic
- ▶ Examples: most current supercomputers, networked parallel computer "grids" and **multi-processor SMP computers** - including some types of PCs.



Amdahl's Law and Gustafson's Law

- ▶ http://www.zdnet.com.tw/white_board/intel/video-2.htm
- ▶ Sun Ni's Law



Outlines

- ▶ **Overview**
 - ▶ What Is Parallel Computing?
 - ▶ Why Use Parallel Computing?
- ▶ **Concepts and terminology**
 - ▶ SISD, SIMD, MISD, MIMD
 - ▶ Amdahl's Law and Gustafson's Law
- ▶ **Parallel Computer Memory Architecture**
 - ▶ Shared Memory
 - ▶ Distributed Memory
 - ▶ Hybrid Distributed-Shared Memory




Parallel Computer Memory Architectures

- ▶ Shared Memory
- ▶ Distributed Memory
- ▶ Hybrid Distributed-Shared Memory



Shared Memory



▶ Advantages:


- ▶ Global address space provides a **user-friendly programming** perspective to memory
- ▶ **Data sharing** between tasks is both fast and uniform due to the proximity of memory to CPUs

▶ Disadvantages:

- ▶ Primary disadvantage is the lack of scalability between memory and CPUs. Adding more CPUs can geometrically increase **traffic on the shared memory-CPU path**, and for cache coherent systems, geometrically increase traffic associated with cache/memory management.
 - ▶ **Programmer responsibility for synchronization constructs** that insure "correct" access of global memory.
 - ▶ Expense: it becomes increasingly difficult and **expensive** to design and produce shared memory machines with ever increasing numbers of processors.
-



Distributed Memory



▶ Advantages:


- ▶ Memory is **scalable** with number of processors. **Increase the number of processors and the size of memory increases proportionately.**
- ▶ Each processor can rapidly access its own memory without interference and without the overhead incurred with trying to maintain cache coherency.
- ▶ Cost effectiveness: can use commodity, off-the-shelf processors and networking.

▶ Disadvantages:

- ▶ **The programmer is responsible for many of the details associated with data communication between processors.**
 - ▶ It may be difficult to map existing data structures, based on global memory, to this memory organization.
 - ▶ Non-uniform memory access (NUMA) times
-



Hybrid Distributed-Shared Memory



- ▶ The shared memory component is usually a cache coherent SMP machine. Processors on a given SMP can address that machine's memory as global.
 - ▶ The distributed memory component is the networking of multiple SMPs. SMPs know only about their own memory - not the memory on another SMP. Therefore, **network communications are required to move data from one SMP to another.**
 - ▶ Current trends seem to indicate that this type of memory architecture will continue to prevail and increase at the high end of computing for the foreseeable future.
 - ▶ Advantages and Disadvantages: whatever is common to both shared and distributed memory architectures.
-

