

Introduction to Parallel Processing


CS 5334/4390

Shirley Moore, Instructor

January 19, 2017

Definition of Parallel Computing

- Simultaneous use of multiple compute resources to solve a computational problem



Definition (2)

- The compute resources might be
 - A single computer with multiple processors and/or cores;
 - An arbitrary number of computers connected by a network;
 - A combination of both.
- The computational problem should be amenable to
 - being broken apart into discrete pieces of work that can be solved simultaneously;
 - executing multiple program instructions at any moment in time;
 - being solved in less time with multiple compute resources than with a single compute resource.

Why Use Parallel Computing?

- The universe is massively parallel.
- Save time and/or money
- Solve larger problems
 - May not be able to fit data into memory of single computer
- Provide concurrency
 - Necessary to take advantage of all resources on multicore/manycore systems
- Trends indicated by ever faster networks, distributed systems, and multi-processor computer architectures (even at the desktop level) clearly show that ***parallelism is the future of computing.***

Parallel Terminology

- **Node**
 - A standalone "computer in a box". Usually comprised of multiple CPUs/processors/cores. Nodes are networked together to comprise a supercomputer.
- **CPU / Socket / Processor / Core**
 - This varies, depending upon whom you talk to. In the past, a CPU (Central Processing Unit) was a singular execution component for a computer. Then, multiple CPUs were incorporated into a node. Then, individual CPUs were subdivided into multiple "cores", each being a unique execution unit. CPUs with multiple cores are sometimes called "sockets" - vendor dependent. The result is a node with multiple CPUs, each containing multiple cores. The nomenclature is confusing at times.

Terminology (2)

- Task
 - A logically discrete section of computational work. A task is typically a program or program-like set of instructions that is executed by a processor. A parallel program consists of multiple tasks running on multiple processors.
- Pipelining
 - Breaking a task into steps performed by different processor units, with inputs streaming through, much like an assembly line; a type of parallel computing.

Terminology (3)

- Shared memory
 - From a strictly hardware point of view, describes a computer architecture where all processors have direct (usually bus based) access to common physical memory. In a programming sense, it describes a model where parallel tasks all have the same "picture" of memory and can directly address and access the same logical memory locations regardless of where the physical memory actually exists.
- Symmetric Multi-Processor (SMP)
 - Hardware architecture where multiple processors share a single address space and access to all resources.
- Distributed memory
 - In hardware, refers to network based memory access for physical memory that is not common. As a programming model, tasks can only logically "see" local machine memory and must use communication to access memory on other machines where other tasks are executing.

Terminology (4)

- Communication
 - Exchange data between parallel tasks. There are several ways this can be accomplished, such as through a shared memory or over a network.
- Synchronization
 - The coordination of parallel tasks in real time, very often associated with communication. Often implemented by establishing a synchronization point within an application where a task may not proceed further until other tasks reaches the same or logically equivalent point.
- Granularity
 - A qualitative measure of the ratio of computation to communication.
 - **Coarse:** relatively large amounts of computational work are done between communication events
 - **Fine:** relatively small amounts of computational work are done between communication events

Speedup and Overhead

- Speedup = $\frac{\text{Wallclock time of serial execution}}{\text{Wallclock time of parallel execution}}$
- Parallel overhead
 - The amount of time required to coordinate parallel tasks, as opposed to doing useful work. Parallel overhead can include factors such as:
 - Task start-up time
 - Synchronizations
 - Data communications
 - Software overhead imposed by parallel compilers, libraries, tools, operating system, etc.
 - Task termination time

Scalability

- Refers to a parallel system's (hardware and/or software) ability to demonstrate a proportionate increase in parallel speedup with the addition of more processors.
- Factors that contribute to scalability include:
 - Hardware
 - Especially memory-CPU and network bandwidths
 - Application algorithm
 - Parallel overhead
 - Characteristics of your specific implementation

Types of Scalability

- Strong scaling
 - Fixed overall problem size
 - Goal is to reduce time to solve the problem
 - Ideal is perfect linear speedup (e.g., with twice as many processors, the problem is solved in half the time)
- Weak scaling
 - Fixed problem size per processor
 - Scale up overall problem size with number of processors (e.g., with twice as many processors, double the overall problem size)
 - Ideal is to solve the larger problem in the same amount of time
- Iso-efficiency scaling
 - Increase problem size along with increase in number of processors so as to maintain constant efficiency in use of resources
 - Often requires increases problem size per processor

Parallel Programming Models

- There are several parallel programming models in common use:
 - Shared Memory
 - Threads
 - Distributed Memory / Message Passing
 - Data Parallel
 - Hybrid
 - Single Program Multiple Data (SPMD)
 - Multiple Program Multiple Data (MPMD)
- Parallel programming models exist as an abstraction above hardware and memory architectures.

Shared Memory

- Tasks share a common address space, which they read and write to asynchronously.
- Various mechanisms such as locks / semaphores may be used to control access to the shared memory.
- An advantage of this model from the programmer's point of view is that the notion of data "ownership" is lacking, so there is no need to specify explicitly the communication of data between tasks, simplifying program development.
- A disadvantage in terms of performance is that it becomes more difficult to understand and manage data locality. Keeping data local to the processor that works on it conserves memory accesses, cache refreshes and bus traffic that occurs when multiple processors use the same data.

Threads

- A type of shared memory programming in which each process can have multiple execution paths
- From a programming perspective, threads implementations commonly comprise
 - A library of subroutines that are called from within parallel source code – e.g., Pthreads, Java threads
 - A set of compiler directives embedded in either serial or parallel source code – e.g., OpenMP

POSIX Threads (aka Pthreads)

- Library based; requires parallel coding
- Specified by the IEEE POSIX 1003.1c standard
- C Language only
- Commonly referred to as Pthreads.
- Most hardware vendors now offer Pthreads in addition to their proprietary threads implementations.
- Very explicit parallelism; requires significant programmer attention to detail.

OpenMP

- Compiler directive based
- Jointly defined and endorsed by a group of major computer hardware and software vendors
 - www.openmp.org
- Portable / multi-platform
- Available in C/C++ and Fortran implementations
- Can be very easy and simple to use - provides for "incremental parallelism"
- Until OpenMP 4.0, didn't have support for data locality => can perform poorly

Distributed Memory/Message Passing Model

- A set of tasks that use their own local memory during computation.
 - Multiple tasks can reside on the same physical machine and/or across an arbitrary number of machines.
- Tasks exchange data through communications by sending and receiving messages.
- Data transfer usually requires cooperative operations to be performed by each process.
 - For example, a send operation must have a matching receive operation.

Message Passing Model Implementation

- From a programming perspective, message passing implementations usually comprise a library of subroutines.
 - Calls to these subroutines are embedded in source code.
 - The programmer is responsible for determining all parallelism.
- Historically, a variety of message passing libraries have been available since the 1980s. These implementations differed substantially from each other making it difficult for programmers to develop portable applications.
- In 1992, the MPI Forum was formed with the primary goal of establishing a standard interface for message passing implementations.
 - www.mpi-forum.org
- MPI is now the "de facto" industry standard for message passing, replacing virtually all other message passing implementations used for production work. MPI implementations exist for virtually all popular parallel computing platforms. Not all implementations include everything in the standard.

Hybrid Model

- A hybrid model combines more than one of the previously described programming models.
- Currently, a common example of a hybrid model is the combination of the message passing model (MPI) with the threads model (OpenMP).
 - Threads perform computationally intensive kernels using local, on-node data
 - Communications between processes on different nodes occurs over the network using MPI
- This hybrid model lends itself well to the increasingly common hardware environment of clustered multi/many-core machines.
- Another similar and increasingly popular example of a hybrid model is using MPI with GPU (Graphics Processing Unit) programming.
 - GPUs perform computationally intensive kernels using local, on-node data
 - Communications between processes on different nodes occurs over the network using MPI

Designing Parallel Programs

- Understand the problem
- Decompose the problem
- Design inter-task communication and synchronization (requires understanding dependencies)

Cost of Communication

- Inter-task communication virtually always implies overhead.
- Machine cycles and resources that could be used for computation are instead used to package and transmit data.
- Communications frequently require some type of synchronization between tasks, which can result in tasks spending time "waiting" instead of doing work.
- Competing communication traffic can saturate the available network bandwidth, further aggravating performance problems.