Lecture 3

Hardware Architecture

Institute of Computer Science & Information Technology,
Faculty of Management & Computer Sciences,
The University of Agriculture, Peshawar, Pakistan.

Workstation

- It is a high performance computer system than mainstream personal computers, that is basically designed for a single user and has advanced graphics capabilities, large storage capacity and a powerful central processing unit.
- It is more capable than a personal computer but is less advanced than a server.
- A server can manage a large network of peripheral PCs or workstations and handle immense data-processing and reporting tasks.

Network of Workstations (NOW)

- High speed networks and rapidly improving performance make networks of workstations an increasingly appealing vehicle for parallel computing.

- By relying solely on commodity hardware and software, network of workstations offer parallel processing at a relatively low cost.

Network of Workstations (1)

- Commodity hardware, sometimes known as off-the-shelf hardware is a computer device or IT component that is relatively inexpensive, widely available and basically interchangeable with other hardware of its type.

Network of Workstations (2)

- A network of workstations and multiprocessor may be realized as a processor bank, a number of processors dedicated for the purpose of providing computing cycles.

- Alternatively, it may consists of a dynamically varying set of machines on which idle cycles are used to perform long-running computations.

Network of Workstations (3)

- In the latter case, the (hardware) cost is essentially zero, since many organization already have extensive workstation networks in place.

- In terms of performance, improvements in processor speed and network bandwidth and latency allow networked workstations to deliver performance approaching or exceeding supercomputer performance for an increasing class of applications.

Distributed Memory

- Distributed memory refers to a multiprocessor computer system in which each processor has its own private memory.

- Computational tasks can only operate on local data, and if remote data are required, the computational task must communicate with one or more remote processors.

Distributed Memory

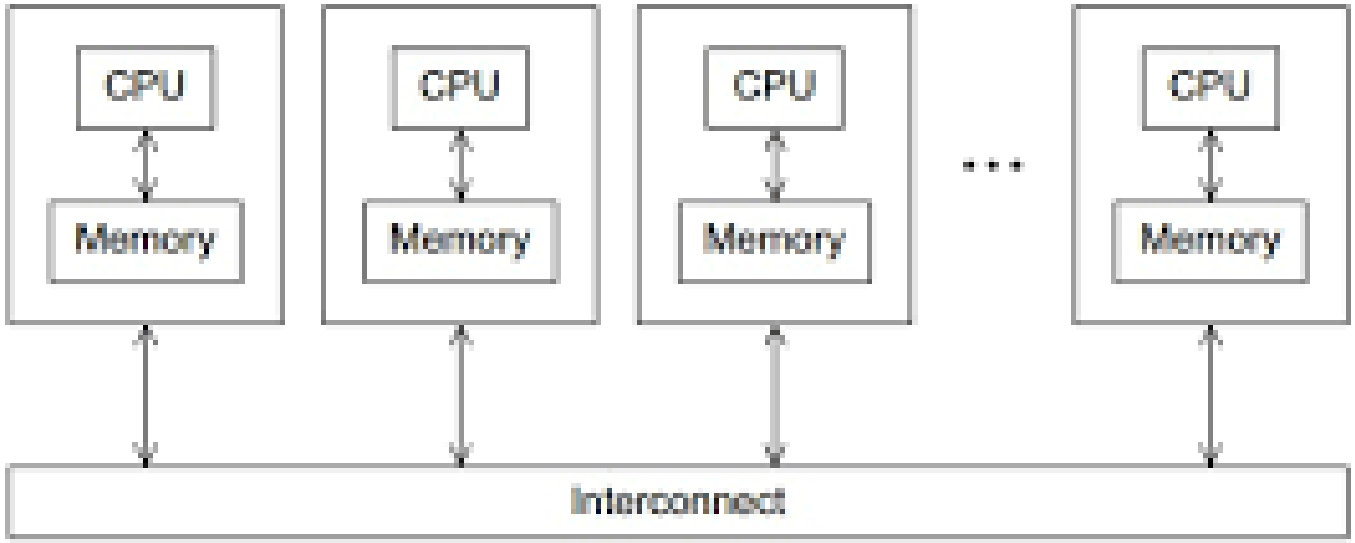

FIGURE 3.1

A distributed-memory system

In a distributed memory system there is a typically a processor, a memory and some form of interconnection that allows programs on each processor to interact with each other.

Distributed Memory(1)

- The key issue in programming distributed memory system is how to distribute data over the memories.
- Depending on the problem solved, the data can be distributed statically, or it can be moved through the nodes.
- Data can be moved on demand, or data can pushed to the new nodes in advance.

Distributed Memory(2)

- Data can be kept statically in nodes if most computations happen locally, and only changes on edges have to be reported to other nodes.
- An example of this is simulation where data is modeled using a grid (non-interactive workloads), and each node simulates a small part of the larger grid.
- On every iteration, nodes inform all neighboring nodes of the new edge data

Shared vs. Distributed Memory

- There are two kinds of multiple-processor systems exist:

1. Multi-Processors (Share/Distributed)
2. Multi-Computers(Distributed)

- A multiprocessor is a system with a processes eating from a shared memory.
- A multicomputer, on the other hand is design in which process utilizes their own distributed memory.

Computer Clusters

- A computer cluster is a set of computers that work together so that they can be viewed as a single system.

- Unlike grid computers (where each node set to perform a different task/application), computer clusters have each node set to perform the same task, controlled and scheduled by software.

Computer Clusters (1)

- The components of a cluster are usually connected to each other through fast local area networks, with each node (computer used as a server) running its own instance of an operating system.

- In most circumstances, all of the nodes use the same hardware and the same operating system, although in some setups(e.g., using Open Source Cluster Application Resources(OSCAR)), different operating systems can be used on each computer or different hardware.

Computer Clusters (1)

- OSCAR is a Linux based software installation for high performance cluster computing.

- Clusters are usually deployed to improve performance and availability over that of a single computer, while typically being much more cost-effective than single computers of **comparable speed or availability**.

Computer Clusters (2)

- Computer Clusters emerged as a result of convergence of a number of computing trends including the availability of low cost microprocessor, high speed networks and software for high performance distributed computing.

Computer Clusters (3)

- Prior to the advent of clusters, single unit fault tolerant mainframes with modular redundancy were employed; but the lower upfront cost of clusters, and increased speed of network fabric has favored the adoption of clusters.

- In contrast to high-reliability mainframes clusters are cheaper to scale out but also have increased complexity in error handling.

Challenges in Clusters

- One of the challenges in the use of a computer cluster is the cost of administrating it which can at times be as high as the cost of administrating N independent machines, if the cluster has N nodes.

Challenges in Clusters (1)

- Some other challenges are mentioned below:
- When a large multi-user cluster needs to access very large amounts of data, task scheduling becomes a challenge.
- In a heterogeneous CPU-GPU cluster with a complex application environment, the performance of each job depends on the characteristics of the underlying cluster.
- Therefore, mapping tasks onto CPU cores and GPU devices provides significant challenges.

Challenges in Clusters (2)

- When a node in a cluster fails, strategies such as fencing may be employed to keep the rest of the system operational.
- Fencing is the process of isolating a node or protecting shared resources when a node appears to be malfunctioning.
- There are two classes of fencing methods; one disables a node itself, and the other disallows access to resources such as shared disks.

Cluster Implementation

- Linux supports various **cluster software** for application clustering there is distcc and MPICH.
- Linux Virtual Server, Linux HA- director-based clusters that allows incoming requests for services to be distributed across multiple cluster nodes.
- MOSIX, LinuxPMI, Kerrighed, OpenSSI are **full-blown cluster** integrated into the kernel that provide for automatic process migration among homogeneous nodes.
- OpenSSI, openMosix and Kerrighed are single-system image implementation.

Cluster Implementation (1)

- Microsoft Windows computer cluster Server 2003 based on the Windows Server platform provides pieces for High Performance Computing like the Job Scheduler, MSMPI library and management tools.