

Distributed Operating Systems

COMP755

“A distributed system is one where you can't get your work done because a computer you didn't even know existed, has crashed.”

- Leslie Lamport

Spectrum of Multiprocessor OS

- Some distributed OS only share files. Accessing data on another computer is very visible to the user. The textbook calls this a **Network Operating System**.
- Other distributed OS try to make a collection of computers look like one large single computer. The textbook calls this a **Distributed Operating System**.
- There are many systems in between.

Types of Multiprocessor Systems

- SMP or dual core computer
- Hypercube, separate memory multiprocessor
- Beowulf Cluster
- Cluster of Workstations
- Ad Hoc collection of computers
- Client/Server systems

Beowulf Cluster

- A Beowulf cluster is a group of usually identical PC computers running an Open Source Unix-like operating system. They are networked into a small LAN.
- Usually consists of one server node and several client nodes connected together via Ethernet
- Server node controls the whole cluster and serves files to the client nodes

Beowulf Programming

- Beowulf appears to the user like a multi-processor computer instead of a bunch of PCs.
- Beowulf clusters often use Parallel Virtual Machine (PVM) or Message Passing Interface (MPI) for parallel programming



Beowulf in Literature

- Epic poem written about 800 – 1000 CE
- Earliest surviving English literature
- The hero, Beowulf, battles monsters and dragons.
- More about “*wurd*” than “*worm*”



Why distributed systems?

- People are distributed, data are distributed
- Performance / Cost
- Scalability
- Modularity
- Availability & Reliability

Characteristics of a Distributed System


- Resource sharing
 - Hardware like printers, disks, scanners
 - Data
 - Web pages
 - software libraries
 - corporate data
 - cooperative work
- Openness
 - interfaces are published
 - uniform communication mechanism
 - possibly heterogeneous

Characteristics

- Concurrency
- Scalability
- Message Passing
- Lack of global information
- Fault Tolerance
 - availability

A distributed system can be built from

1. Microsoft Windows PCs
2. Linux computers
3. Sun Sparc workstations
4. All of the above
5. None of the above



Transparency

- **access** - local and remote objects accessed by identical operations
- **location** - do not have to know location of objects or processes
- **concurrency** - multiple access without interference
- **replication** - multiple copies without user or application knowledge

Transparency (cont.)

- **failure** - continue despite hardware or software failure
- **migration** - objects can move without user knowledge
- **reconfiguration** - system can be changed without user knowledge
- **scaling** - increasing the system size does not effect structure or applications

Goals

- Efficiency
 - Minimize communications
 - load balancing
- Flexibility
 - friendliness
 - ability to evolve.
- Consistency
 - transparency
- Robustness
 - equipped to handle exceptional situations and errors.

Robustness

- Failure detection
- Reconfiguration

Failure Detection

- Detecting hardware failure is difficult
- To detect a link failure, a handshaking protocol can be used
- If Site A does not receive a reply, it can repeat the message or try an alternate route to Site B

Failure Detection (cont)

- If Site A does not ultimately receive a reply from Site B, it concludes some type of failure has occurred
- Types of failures:
 - Site B is down
 - The direct link between A and B is down
 - The alternate link from A to B is down
 - The message has been lost
- However, Site A cannot determine exactly **why** the failure has occurred

Reconfiguration

- When Site A determines a failure has occurred, it must reconfigure the system:
 1. If the link from A to B has failed, this must be broadcast to every site in the system
 2. If a site has failed, every other site must also be notified indicating that the services offered by the failed site are no longer available
- When the link or the site becomes available again, this information must again be broadcast to all other sites