

LECTURE - XXII
DISTRIBUTED SYSTEMS


Tevfik Koşar

Louisiana State University
April 24th, 2007

Motivation

- **Distributed system** is collection of loosely coupled processors that
 - do not share memory
 - interconnected by a communications network
- **Reasons for distributed systems**
 - Resource sharing
 - sharing and printing files at remote sites
 - processing information in a distributed database
 - using remote specialized hardware devices
 - Computation speedup - **load sharing**
 - Reliability - detect and recover from site failure, function transfer, reintegrate failed site
 - Communication - message passing

A Distributed System




Distributed-Operating Systems

- Users not aware of multiplicity of machines
 - Access to remote resources similar to access to local resources
- Data Migration - transfer data by transferring entire file, or transferring only those portions of the file necessary for the immediate task
- Computation Migration - transfer the computation, rather than the data, across the system

Distributed-Operating Systems (Cont.)

- **Process Migration** - execute an entire process, or parts of it, at different sites
 - Load balancing - distribute processes across network to even the workload
 - Computation speedup - subprocesses can run concurrently on different sites
 - Hardware preference - process execution may require specialized processor
 - Software preference - required software may be available at only a particular site
 - Data access - run process remotely, rather than transfer all data locally

Network Topology



Robustness in Distributed Systems

- Failure detection
- Reconfiguration

Failure Detection

- Detecting hardware failure is difficult
- To detect a link failure, a handshaking protocol can be used
- Assume Site A and Site B have established a link
 - At fixed intervals, each site will exchange an *I-am-up* message indicating that they are up and running
- If Site A does not receive a message within the fixed interval, it assumes either (a) the other site is not up or (b) the message was lost
- Site A can now send an *Are-you-up?* message to Site B
- 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

Distributed File Systems

- Distributed file system (DFS) - a distributed implementation of the classical time-sharing model of a file system, where multiple users share files and storage resources
- A DFS manages set of dispersed storage devices
- Overall storage space managed by a DFS is composed of different, remotely located, smaller storage spaces
- There is usually a correspondence between constituent storage spaces and sets of files

DFS Structure

- **Service** - software entity running on one or more machines and providing a particular type of function to a priori unknown clients
- **Server** - service software running on a single machine
- **Client** - process that can invoke a service using a set of operations that forms its *client interface*
- A client interface for a file service is formed by a set of primitive *file operations* (create, delete, read, write)
- Client interface of a DFS should be transparent, i.e., not distinguish between local and remote files

Naming and Transparency

- **Naming** - mapping between logical and physical objects
- Multilevel mapping - abstraction of a file that hides the details of how and where on the disk the file is actually stored
- A **transparent** DFS hides the location where in the network the file is stored
- For a file being replicated in several sites, the mapping returns a set of the locations of this file's replicas; both the existence of multiple copies and their location are hidden

Naming Structures

- **Location transparency** - file name does not reveal the file's physical storage location
 - File name still denotes a specific, although hidden, set of physical disk blocks
 - Convenient way to share data
 - Can expose correspondence between component units and machines
- **Location independence** - file name does not need to be changed when the file's physical storage location changes
 - Better file abstraction
 - Promotes sharing the storage space itself
 - Separates the naming hierarchy from the storage-devices hierarchy

Naming Schemes — Three Main Approaches

- Files named by combination of their host name and local name; guarantees a unique systemwide name
 - Eg. host:local-name
 - Not location transparent, nor location independent
- Attach remote directories to local directories, giving the appearance of a coherent directory tree; only previously mounted remote directories can be accessed transparently
 - Eg. NFS
- Total integration of the component file systems
 - A single global name structure spans all the files in the system
 - If a server is unavailable, some arbitrary set of directories on different machines also becomes unavailable

Remote File Access

- **Remove-service mechanism** is one transfer approach
- Reduce network traffic by retaining recently accessed disk blocks in a cache, so that repeated accesses to the same information can be handled locally
 - If needed data not already cached, a copy of data is brought from the server to the user
 - Accesses are performed on the cached copy
 - Files identified with one master copy residing at the server machine, but copies of (parts of) the file are scattered in different caches
 - **Cache-consistency problem** - keeping the cached copies consistent with the master file
 - Could be called **network virtual memory**

Cache Location - Disk vs. Main Memory

- Advantages of disk caches
 - More reliable
 - Cached data kept on disk are still there during recovery and don't need to be fetched again
- Advantages of main-memory caches:
 - Permit workstations to be diskless
 - Data can be accessed more quickly
 - Performance speedup in bigger memories
 - Server caches (used to speed up disk I/O) are in main memory regardless of where user caches are located; using main-memory caches on the user machine permits a single caching mechanism for servers and users

Any Questions?



Reading Assignment

- Read chapter 16 and 17 from Silberschatz.

19

Acknowledgements

- “Operating Systems Concepts” book and supplementary material by Silberschatz, Galvin and Gagne.

20