COP 6611 Advanced Operating System

Introduction

Chi Zhang
czhang@cs.fiu.edu

1

Outline

- Goals
- Hardware Concepts
- Software Concepts
- The Client-Server Model

2

Definition of a Distributed System (1)

- A distributed system is:

  A collection of independent computers that appears to its users as a single coherent system.
- Goals
  - Utilizing distributed resources
  - Transparency
  - Open / Extensible
  - Scalable

3

Definition of a Distributed System (2)

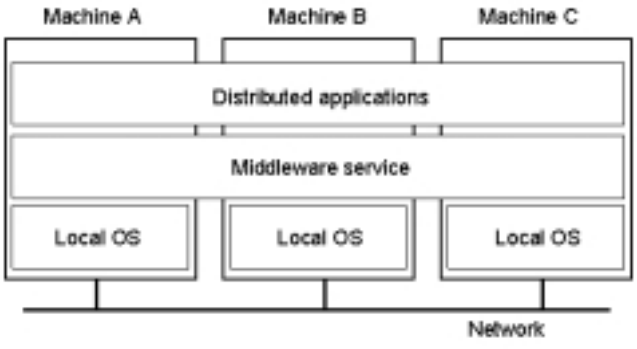

A distributed system organized as middleware.
Note that the middleware layer extends over multiple machines.

4

Transparency in a Distributed System

Transparency	Description
Access	Hide differences in data representation and how a resource is accessed
Location	Hide where a resource is located
Migration	Hide that a resource may move to another location
Relocation	Hide that a resource may be moved to another location while in use
Replication	Hide that a resource may be shared by several competitive users
Concurrency	Hide that a resource may be shared by several competitive users
Failure	Hide the failure and recovery of a resource
Persistence	Hide whether a (software) resource is in memory or on disk

Different forms of transparency in a distributed system.

5

Scalability Problems

- More users / resources
- Geographically scalable
- Administratively scalable

Concept	Example
Centralized services	A single server for all users
Centralized data	A single on-line telephone book
Centralized algorithms	Doing routing based on complete information

Examples of scalability limitations.

6

Scaling Techniques (1)

- Hiding communication latencies
  - Asynchronous communication
  - Moving server tasks to clients
- (Server) Distribution
- Replication
  - Increase availability
  - Load balancing
  - Access nearby copy

7

Scaling Techniques (2)

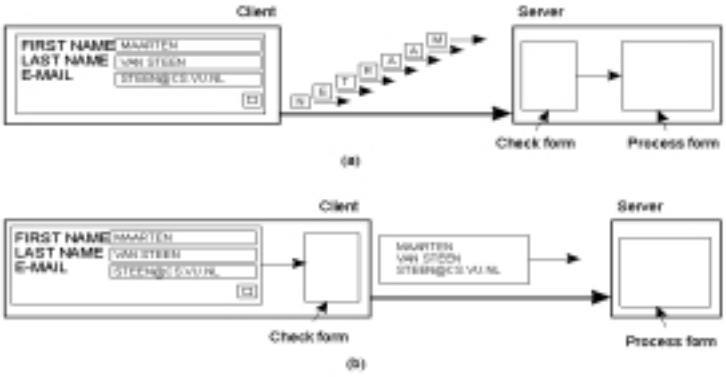

The difference between letting:

a) a server or

b) a client check forms as they are being filled

8

Scaling Techniques (3)

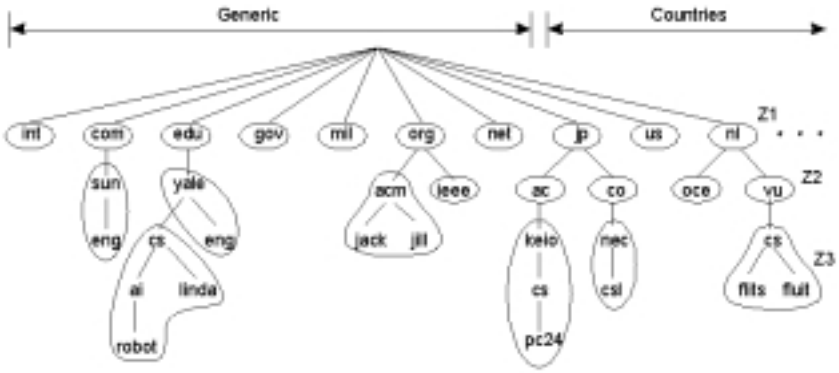

An example of dividing the DNS name space into zones.

Hardware Concepts

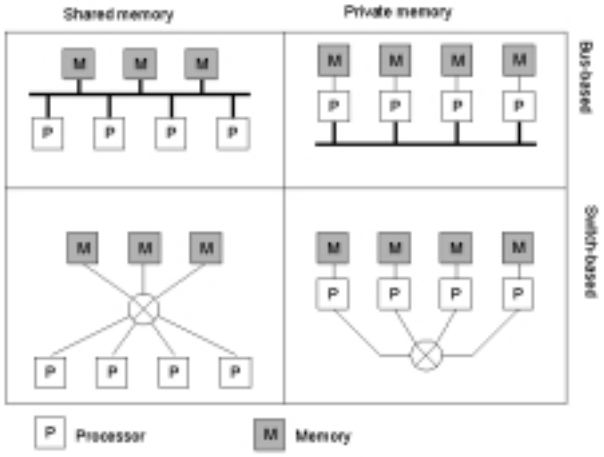

Homogeneous vs. Heterogeneous

Multiprocessors (1)

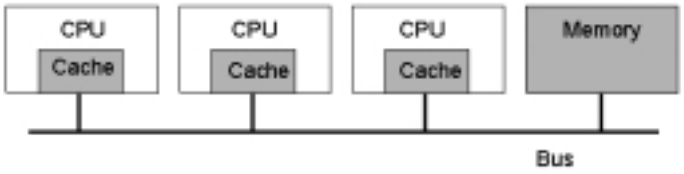

A bus-based multiprocessor.

11

Multiprocessors (2)

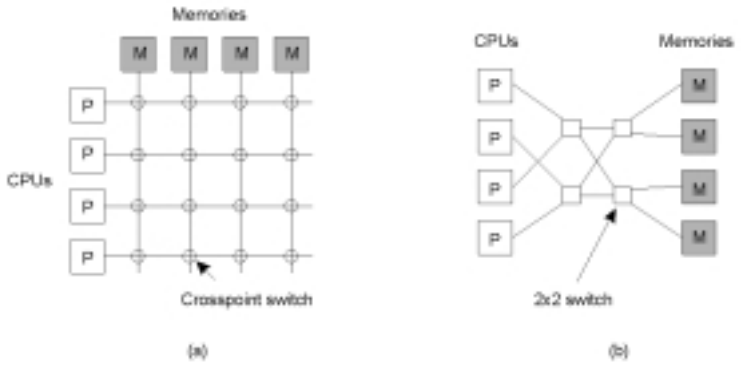

a) A crossbar switch
b) An omega switching network

12

Homogeneous Multicomputer Systems

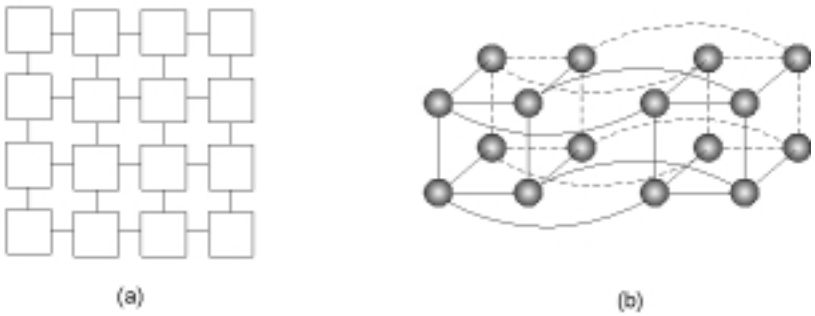

a) Grid
b) Hypercube

System Area Networks (SAN)

13

Heterogeneous Multicomputer Systems

- Most distributed systems
- Lack a global system view
  - Operating system does not know how many computers in the system
  - Applications are aware of distributed ness (without middleware)
  - Applications cannot assume the same performance or services everywhere

14

Software Concepts

System	Description	Main Goal
DOS	Tightly-coupled operating system for multi-processors and homogeneous multicomputers	Hide and manage hardware resources
NOS	Loosely-coupled operating system for heterogeneous multicomputers (LAN and WAN)	Offer local services to remote clients
Middleware	Additional layer atop of NOS implementing general-purpose services	Provide distribution transparency

An overview of

- DOS (Distributed Operating Systems)
- NOS (Network Operating Systems)
- Middleware

15

Uniprocessor Operating Systems

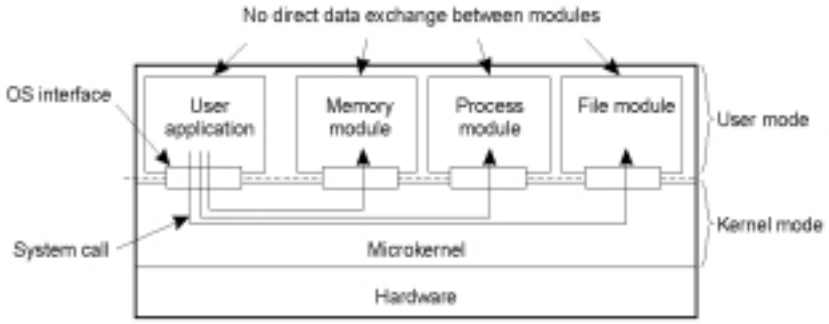

Separating applications from operating system code through a microkernel.

16

Multiprocessor Operating Systems (1)

- High Performance with Multiple CPUs
  - One thread per CPU
  - Communication through shared memory location
- Protect data against simultaneous accesses
  - Semaphore
  - Monitor

17

Multiprocessor Operating Systems (2)

```
monitor Counter {
private:
 int count = 0;
public:
  int value() { return count;}
  void incr () { count = count + 1;}
  void decr() { count = count – 1;}
}
```

A monitor to protect an integer against concurrent access.

18

Multicomputer Operating Systems (1)

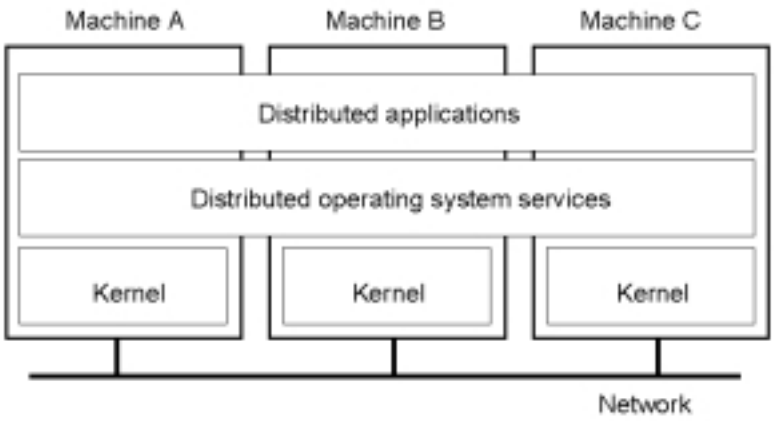

General structure of a multicomputer operating system

19

Multicomputer Operating Systems (2)

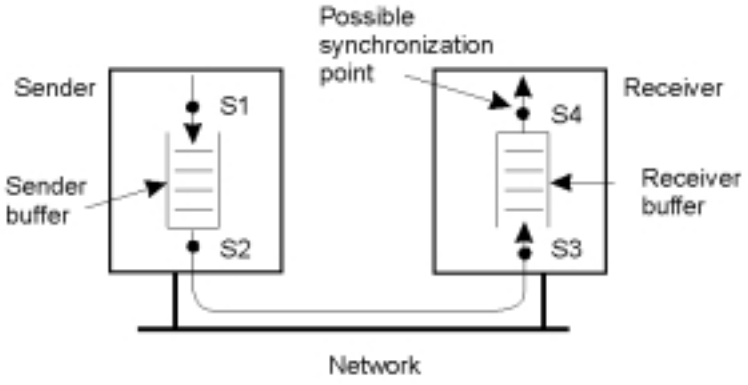

Alternatives for blocking and buffering in message passing.

20

Multicomputer Operating Systems (3)

Synchronization point	Send buffer	Reliable comm. guaranteed?
Block sender until buffer not full	Yes	Not necessary
Block sender until message sent	No	Not necessary
Block sender until message received	No	Necessary
Block sender until message delivered	No	Necessary

Relation between blocking, buffering, and reliable communications.

21

Distributed Shared Memory Systems (1)

a) **Pages of address space distributed among four machines**

b) **Situation after CPU 1 references page 10**

c) **Situation if page 10 is read only and replication is used**

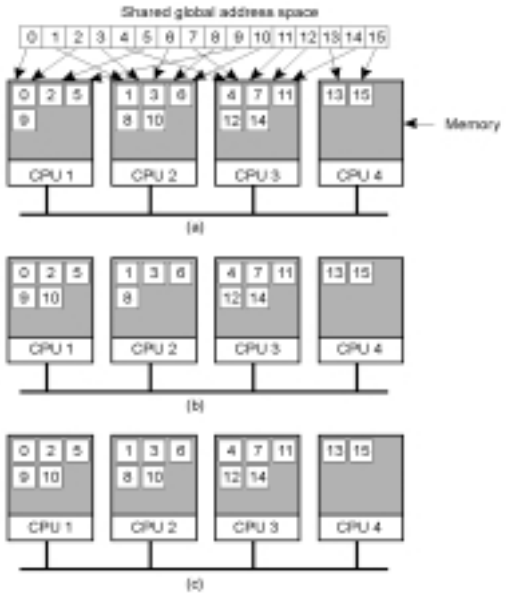

Distributed Shared Memory Systems (2)

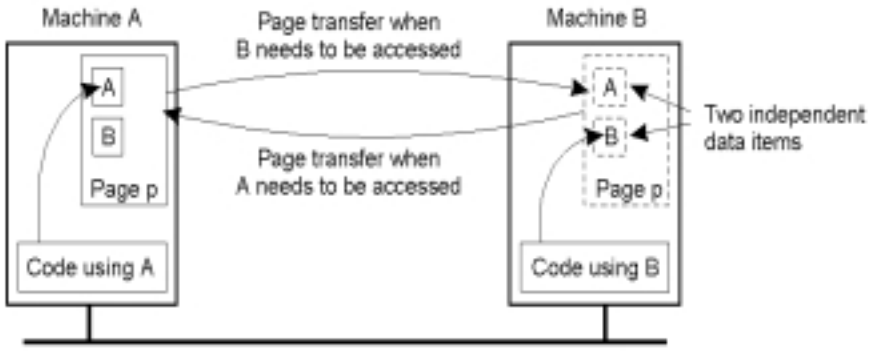

False sharing of a page between two independent processes.

23

Network Operating System (1)

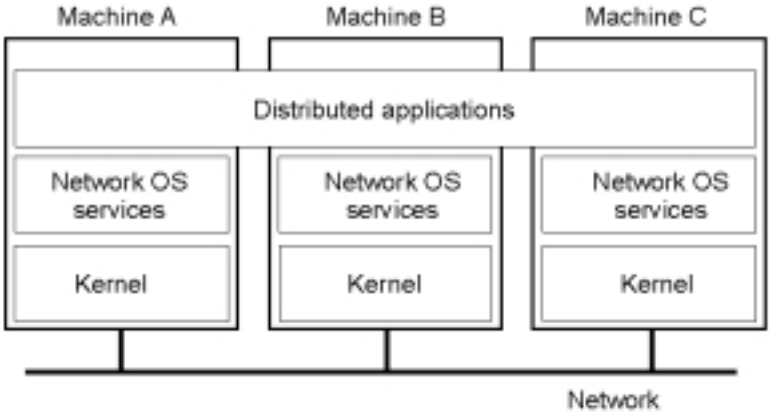

General structure of a network operating system.

24

Network Operating System (2)

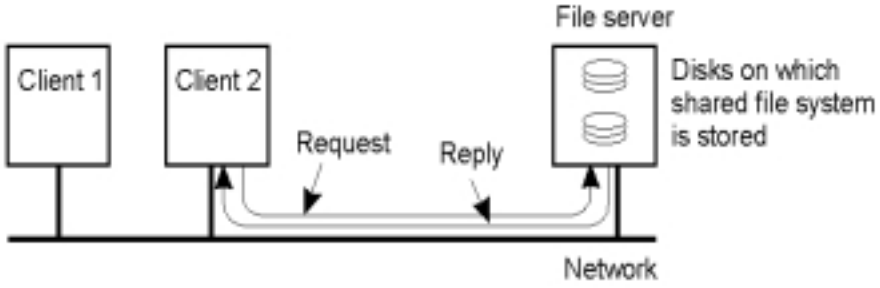

Two clients and a server in a network operating system.

25

Introduction to Middleware

- NOS + MiddlewareNOS: Scalable / Open
  - DOS: Transparent
- Middleware models
  - Distributed file systems / distributed database
  - Remote procedures calls (RPC)
  - Distributed Objects
- Middleware services
  - Naming, persistence, transactions …

26

Positioning Middleware

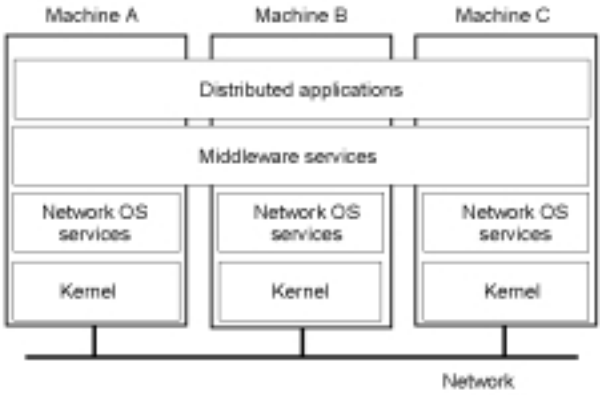

General structure of a distributed system as middleware.

27

Middleware and Openness

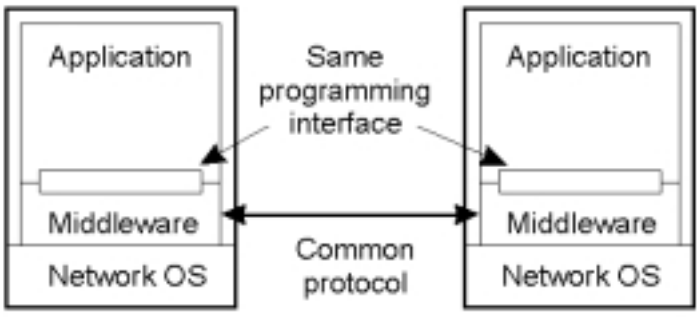

In an open middleware-based distributed system, the protocols used by each middleware layer should be the same, as well as the interfaces they offer to applications.

28

Comparison between Systems

Item	Distributed OS		Network OS	Middleware-based OS
	Multiproc.	Multicomp.		
Degree of transparency	Very High	High	Low	High
Same OS on all nodes	Yes	Yes	No	No
Number of copies of OS	1	N	N	N
Basis for communication	Shared memory	Messages	Files	Model specific
Resource management	Global, central	Global, distributed	Per node	Per node
Scalability	No	Moderately	Yes	Varies
Openness	Closed	Closed	Open	Open

A comparison between multiprocessor operating systems, multicomputer operating systems, network operating systems, and middleware based distributed systems.

29

Clients and Servers

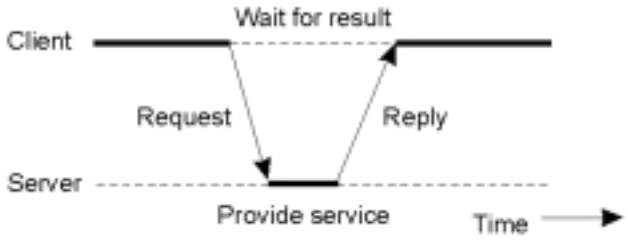

General interaction between a client and a server.

Underlying protocols: Reliable or Unreliable?

30

Processing Level

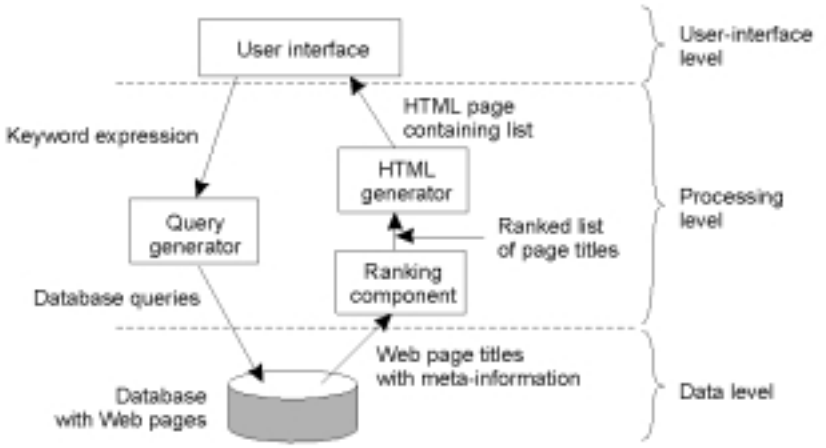

The general organization of an Internet search engine into three different layers

31

Multitiered Architectures (1)

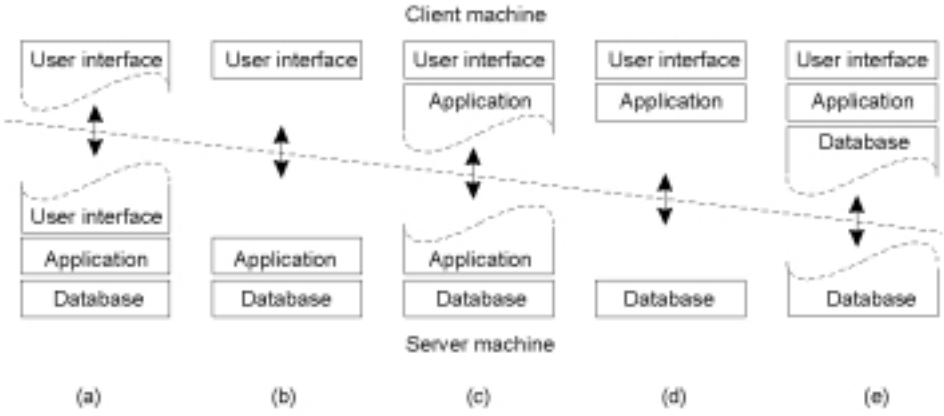

Alternative client-server organizations (a) – (e).

32

Multitiered Architectures (2)

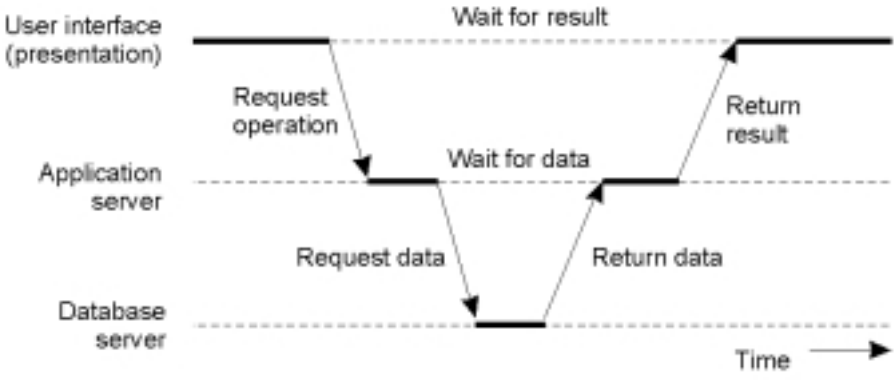

An example of a server acting as a client.

33

Modern Architectures

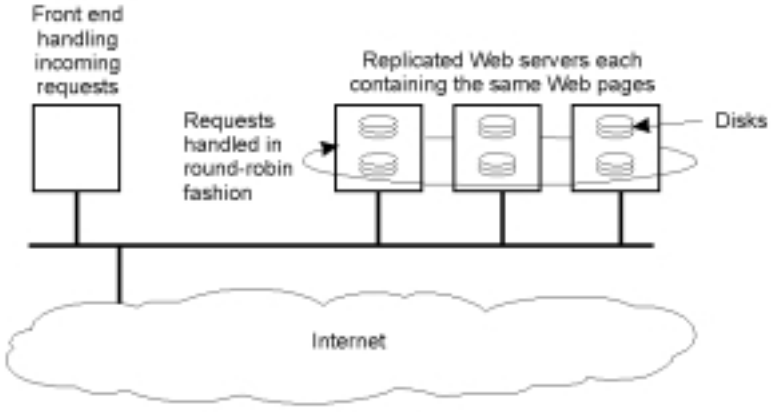

An example of horizontal distribution of a Web service.

34