



CS4513 Distributed Computer Systems

Introduction
(Ch 1: 1.1-1.2, 1.4-1.5)




Outline

- Overview
- Goals
- Software
- Client Server

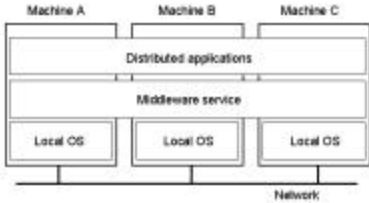


The Rise of Distributed Systems

- Computer hardware prices falling, power increasing
 - If cars the same, Rolls Royce would cost 1 dollar and get 1 billion miles per gallon (with 200 page manual to open the door)
- Network connectivity increasing
 - Everyone is connected with fat pipes
- It is *easy* to connect hardware together
- Definition: a *distributed system* is
 - A collection of independent computers that appears to its users as a single coherent system.




Definition of a Distributed System



Examples:
-The Web
-Processor Pool
-Airline
Reservation


A distributed system organized as middleware.
Note that the middleware layer extends over multiple machines.
Users can interact with the system in a consistent way, regardless of where the interaction takes place



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.




Scalability Problems

- As distributed systems grow, centralized solutions are limited
 - Consider LAN name resolution vs. WAN

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

- Sometimes, hard to avoid (consider a bank)
- Need to collect information in distributed fashion and distributed in a distributed fashion
- Challenges:
 - geography, ownership domains, time synchronization



Scaling Techniques: Hiding Communication Latency

- Especially important for interactive applications
- If possible, do *asynchronous communication*
- Not always possible when client has nothing to do

(a) Synchronous communication: Client sends request, waits for server to finish processing, then client starts its own work.

(b) Asynchronous communication: Client sends request, immediately starts its own work, while server processes the request in the background.

- Instead, can hide latencies

Scaling Techniques: Distribution

Example: DNS name space into zones
 (nl.vu.cs.fluit - z1 gives address of vu gives address of cs)

Example: The Web

Scaling Techniques: Replication

- Copy of information to increase availability and decrease centralized load
- Example: P2P networks (Gnutella +) distribute copies uniformly or in proportion to use
- Example: akamai
- Example: Caching is a replication decision made by client
- Issue: Consistency of replicated information
- Example: Web Browser cache

Outline

- Overview (done)
- Goals (done)
- Software →
- Client Server

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

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

Uniprocessor Operating Systems

- Separating applications from operating system code through a microkernel
- Can extend to multiple computers

Multicomputer Operating Systems

- But no longer have shared memory
 - Can try to provide *distributed shared memory*
 - Tough, coming up
 - Can provide *message passing*

Multicomputer Operating Systems

- Message passing primitives vary widely between systems
 - Example: consider *buffering* and *synchronization*

Multicomputer Operating Systems

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*.
- These issues make synchronization harder. It was easier when we had shared memory.
 - So ... distributed shared memory

Distributed Shared Memory Systems

- Pages of address space distributed among four machines
- Situation after CPU 1 references page 10
- Situation if page 10 is read only and replication is used

Distributed Shared Memory Systems


- Issue: how large should page sizes be? What are the tradeoffs?

- Overall, DSM systems have struggled to provide efficiency and convenience (and been around 15 years)
 - For higher-performance, typically still do message passing
 - Likely will remain that way

Network Operating System

- OSes can be different (Windows or Linux)
- Typical services: rlogin, rcp
 - Fairly primitive way to share files


Network Operating System



- Can have one computer provide files transparently for others (NFS)
 - (try a "df" on the WPI hosts to see. Similar to a "mount network drive" in Windows)



Network Operating System




- Different clients may mount the servers in different places
- Inconsistencies in view make NOSes harder, in general for users than DOSes.
 - But easier to scale by adding computers



Positioning Middleware


- Network OS not transparent. Distributed OS not independent computers.
 - Middleware can help



- Much middleware built in-house to help use networked operating systems (distributed transactions, better comm, RPC)
 - Unfortunately, many different standards



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.
 - If different, compatibility issues
 - If incomplete, then users build their own or use lower-layer services (frowned upon)



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

- DOS most transparent, but closed and only moderately scalable
- NOS not so transparent, but open and scalable
- Middleware provides a bit more transparency than NOS




Outline

- Overview (done)
- Goals (done)
- Software (done)
- Client Server ↪



Clients and Servers

- Thus far, have not talked about organization of *processes*
 - Again, many choices but most agree upon client-server



- If can do so without connection, quite simple
 - If underlying connection is unreliable, not trivial
 - Resend? What if receive twice
- Use TCP for reliable connection (apps on Internet)
 - Not always appropriate for high-speed LAN connection (4513)

Example Client and Server: Header

```

/* Definitions needed by clients and servers. */
#define TRUE 1
#define MAX_PATH 255 /* maximum length of file buffer */
#define BUF_SIZE 1024 /* how much data to transfer at once */
#define FILE_SERVER 240 /* the server's network address */

/* Definitions of the allowed operations */
#define CREATE 1 /* create a new file */
#define READ 2 /* read data from a file and return it */
#define WRITE 3 /* write data to a file */
#define DELETE 4 /* delete an existing file */

/* Error codes */
#define OK 0 /* operation performed correctly */
#define E_BAD_OPCODE -1 /* unknown operation requested */
#define E_BAD_PATH -2 /* error in a parameter */
#define E_IO -3 /* disk error or other I/O error */

/* Definition of the message format */
struct message {
    long source; /* server's identity */
    long dest; /* receiver's identity */
    long opcode; /* requested operation */
    long count; /* number of bytes to transfer */
    long offset; /* position in file to start I/O */
    long result; /* result of the operation */
    char name[MAX_PATH]; /* name of file being operated on */
    char data[BUF_SIZE]; /* data to be read or written */
};
    
```

- Used by both the client and server.

Example Client and Server: Server

```

#include <header.h>
void main(void) {
    struct message m1, m2; /* incoming and outgoing messages */
    int r; /* result code */

    while(TRUE) { /* server runs forever */
        receive(FILE_SERVER, &m1); /* block waiting for a message */
        switch(m1.opcode) { /* dispatch on type of request */
            case CREATE: r = do_create(&m1, &m2); break;
            case READ: r = do_read(&m1, &m2); break;
            case WRITE: r = do_write(&m1, &m2); break;
            case DELETE: r = do_delete(&m1, &m2); break;
            default: r = E_BAD_OPCODE;
        }

        m2.result = r; /* return result to client */
        send(m1.source, &m2); /* send reply */
    }
}
    
```

Example Client and Server: Client

```

#include <header.h>
int copy(char *src, char *dst) { /* procedure to copy file using the server */
    struct message m; /* message buffer */
    long position; /* current file position */
    long client = 111; /* client's address */

    while(1) /* prepare for execution */
    {
        position = 0;

        do {
            m.opcode = READ; /* operation is a read */
            m.offset = position; /* current position in the file */
            m.count = BUF_SIZE; /* how many bytes to read? */
            strcpy(m.name, src); /* copy name of file to be read to message */
            send(FILE_SERVER, &m); /* send the message to the file server */
            r = wait(&client, &m); /* block waiting for the reply */

            /* While the data just received to the destination file. */
            m.opcode = WRITE; /* operation is a write */
            m.offset = position; /* current position in the file */
            m.count = m.result; /* how many bytes to write */
            strcpy(m.name, dst); /* copy name of file to be written to buf */
            send(FILE_SERVER, &m); /* send the message to the file server */
            r = wait(&client, &m); /* block waiting for the reply */
            position += m.result; /* m.result is number of bytes written */
        } while(m.result > 0); /* repeat until done */

        return(m.result > 0 ? OK : m.result); /* return OK or error code */
    }
}
    
```


- One issue, is how to clearly differentiate

Client-Server Implementation Levels




- Example of an Internet search engine
 - UI on client
 - Processing can be on client or server
 - Data level is server, keeps consistency

Multitiered Architectures



- Thin client (a) to Fat client (e)
 - (d) and (e) popular for NOS environments


Multitiered Architectures: 3 tiers



- Server may act as a client
 - Example would be transaction monitor across multiple databases



Modern Architectures: Horizontal



- Rather than vertical, distribute servers across nodes
 - Example of Web server "farm" for load balancing
 - Clients, too (peer-to-peer systems)

