

Introduction


Chapter 1

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.

Definition of a Distributed System (2)



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

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

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.


Scaling Techniques (1)



The difference between letting:


- a) a server or
- b) a client check forms as they are being filled

Scaling Techniques (2)




An example of dividing the DNS name space into zones.

Hardware Concepts




Different basic organizations and memories in distributed computer systems

Multiprocessors (1)




A bus-based multiprocessor.

Multiprocessors (2)



- a) A crossbar switch
- b) An omega switching network

Homogeneous Multicomputer Systems



- a) Grid
- b) Hypercube


Software Concepts

System	Description	Main Goal
DOS	Tightly-coupled operating system for multiprocessors 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

Uniprocessor Operating Systems



Separating applications from operating system code through a microkernel.

Multiprocessor Operating Systems (1)

```
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.

Multiprocessor Operating Systems (2)

```
monitor Counter {
private:
    int count = 0;
    int blocked_procs = 0;
    condition unblocked;
public:
    int value () { return count;}
    void incr () {
        if (blocked_procs == 0)
            count = count + 1;
        else
            signal (unblocked);
    }
    void decr () {
        if (count == 0) {
            blocked_procs = blocked_procs + 1;
            wait (unblocked);
            blocked_procs = blocked_procs - 1;
        }
        else
            count = count - 1;
    }
}
```


A monitor to protect an integer against concurrent access, but blocking a process.

Multicomputer Operating Systems (1)



General structure of a multicomputer operating system

Multicomputer Operating Systems (2)



Alternatives for blocking and buffering in message passing.


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.

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.

Network Operating System (1)




General structure of a network operating system.

Network Operating System (2)




Two clients and a server in a network operating system.

Network Operating System (3)




Different clients may mount the servers in different places.

Positioning Middleware



General structure of a distributed system as middleware.

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.

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.

Clients and Servers



General interaction between a client and a server.

An Example Client and Server (1)

```
/* Definitions needed by clients and servers. */
#define TRUE 1
#define MAX_PATH 255 /* maximum length of file name */
#define BUF_SIZE 1024 /* how much data to transfer at once */
#define FILE_SERVER 243 /* file 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_PARAM -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; /* sender'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 */
};
```

The *header.h* file used by the client and server.

An Example Client and Server (2)

```
#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 */
    }
}
```

A sample server.

An Example Client and Server (3)

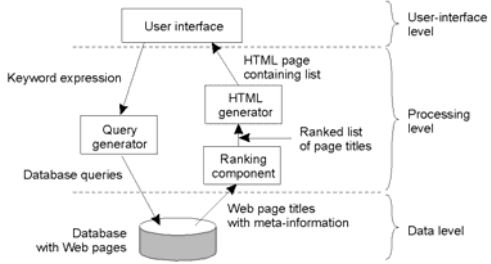
```
#include <header.h>
int copy(char *src, char *dst) { /* procedure to copy file using the server */
    struct message m; /* message buffer */
    long position; /* current file position */
    long client = 110; /* client's address */

    initialize(); /* prepare for execution */
    position = 0;
    do {
        m1.opcode = READ; /* operation is a read */
        m1.offset = position; /* current position in the file */
        m1.count = BUF_SIZE; /* how many bytes to read? */
        strcpy(m1.name, src); /* copy name of file to be read to message */
        send(FILE_SERVER, &m1); /* send the message to the file server */
        receive(client, &m1); /* block waiting for the reply */

        /* Write the data just received to the destination file. */
        m1.opcode = WRITE; /* operation is a write */
        m1.offset = position; /* current position in the file */
        m1.count = m1.result; /* how many bytes to write */
        strcpy(m1.name, dst); /* copy name of file to be written to buf */
        send(FILE_SERVER, &m1); /* send the message to the file server */
        receive(client, &m1); /* block waiting for the reply */
        position += m1.result; /* m1.result is number of bytes written */
    } while(m1.result > 0); /* iterate until done */
    return(m1.result >= 0 ? OK : m1.result); /* return OK or error code */
}
```


A client using the server to copy a file.

Processing Level




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

Multitiered Architectures (1)




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

Multitiered Architectures (2)



An example of a server acting as a client.

Modern Architectures



An example of horizontal distribution of a Web service.
