

Operating Systems


Fall 2008

Tiina Niklander

Material

- Course book: A.S. Tanenbaum, *Modern Operating Systems, 3rd. ed.*, Prentice-Hall, 2007
 - Any other large OS book, f.e. Stallings, Silbershatz, or Deitel, should be feasible alternative
- Lectures & exercises on the course web page <http://www.cs.helsinki.fi/u/niklande/opetus/kj/>

Distributed Systems and Networking specialization area




Course structure

- | | |
|--------------------------------|--|
| ■ 1: Overview | ■ 8: I/O Management |
| ■ 2: Processes and Threads | ■ 9: Multiple Processor Systems |
| ■ 3: Virtual Memory and Paging | ■ 10: Example: Linux, Windows, Symbian |
| ■ 4: Page Replacement | ■ 11: Design Issues |
| ■ 5: Segmentation | ■ 12: Recapitulation, hints for exam |
| ■ 6: File Systems, part 1 | |
| ■ 7: File Systems, part 2 | |

Tanenbaum: Sections 1,2,3,4 ,5,8,10,11,12,13,14

Distributed Systems and Networking



Master level: Operating systems, distributed systems, networking.

Bachelor level: basic communication & protocols, concurrency concepts, security, computer organisation

Schedule

- Lectures (in Finnish) twice a week
 - Read the book carefully,
 - Focus on mastering the points given in slides
- Exercises once a week
 - Problems available one week before
 - Must be solved in advance, that is BEFORE the meeting (not during it!)
- One weekly problem to be submitted on paper
- EXAM: Wed 15.10. 16.00

Operating System Overview

Concepts (Abstractions)

- Process
 - Is a container that holds all information needed to run a program
 - Process table entry contains all process related info
- Address Spaces
 - Memory locations 0..MAX available for one process
 - Protected from each other
- File
 - Directories to contain files: working directory, path
 - File descriptor
 - File system

Introduction


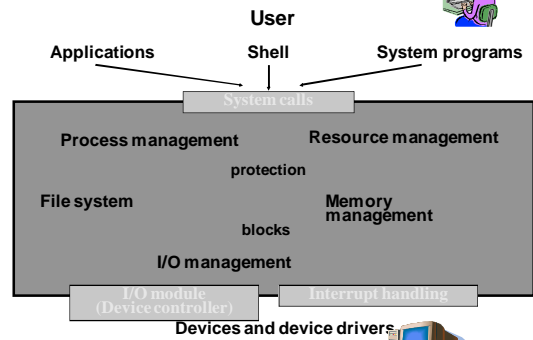


Figure 1-1. Where the operating system fits in.


- A computer system consists of
 - hardware
 - system programs
 - application programs

OS structure



Main Tasks of Operating System

- Provide abstractions
 - Abstractions (like file, address space) to user programs
 - Implement and manage the abstract objects
 - Hide the ugly and messy details in these
- Manage resources
 - Multiplexing in time and in space!
 - Each program gets time with the resource
 - Each program gets space on the resource
 - Optimize the hardware usage to improve performance



Hardware

Computer Hardware Review

- Set of instructions
- Registers
 - Program counter
 - Stack pointer
 - Process Status Word (PSW)
- Kernel mode vs User mode
- System call (service call, trap)

13

Memory hierarchy

Typical access time	Typical capacity
1 nsec	<1 KB
2 nsec	4 MB
10 nsec	512-2048 MB
10 msec	200-1000 GB
100 sec	400-800 GB

Figure 1-9. A typical memory hierarchy. The numbers are very rough approximations.

Pentium 4 cache:
8 KB data, 12 KB code/text, external 256 KB

nano = 10⁻⁹, micro = 10⁻⁶, milli = 10⁻³

Interrupts

- I/O device controller provides a clean interface to the device
- Device driver is the software that communicates with controller
- Interrupt is the way for controller to inform driver about status change

14

Locality of reference

Spatial and temporal locality:

- For example in a loop a small set of instructions is executed several times
- Certain part of code only uses a small set of variables (data)

→ When program makes a memory reference (data or instructions), it is likely to refer again to the same location or a location near by

- This is the principle behind the usage of caches

Interrupt in details

Sta Fig 1.11

Computer Hardware Review

**Structure of a large Pentium system:
multiple buses**

18

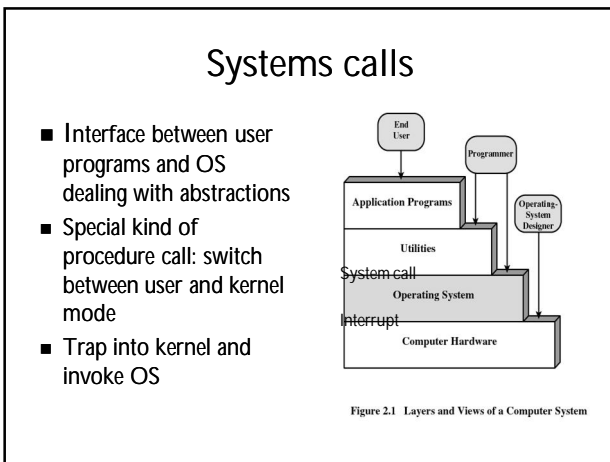
System calls

POSIX: Some System Calls For Process and File Management

Process management	
Call	Description
pid = fork()	Create a child process identical to the parent
pid = waitpid(pid, &statloc, options)	Wait for a child to terminate
s = execve(name, argv, environp)	Replace a process' core image
exit(status)	Terminate process execution and return status

File management	
Call	Description
fd = open(file, how, ...)	Open a file for reading, writing or both
s = close(fd)	Close an open file
n = read(fd, buffer, nbytes)	Read data from a file into a buffer
n = write(fd, buffer, nbytes)	Write data from a buffer into a file
position = lseek(fd, offset, whence)	Move the file pointer
s = stat(name, &buf)	Get a file's status information

22

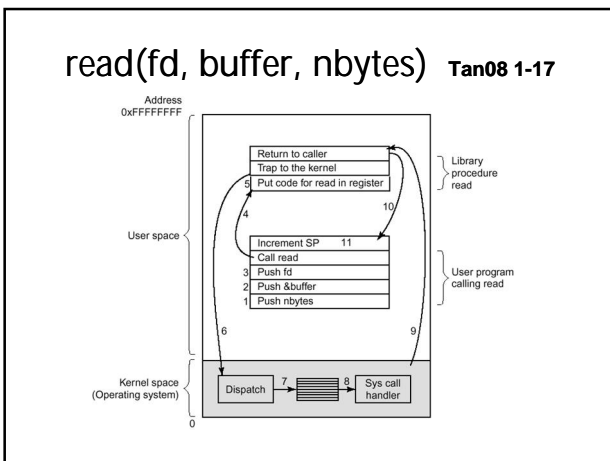


POSIX: Some System Calls For Directory Management and Misc

Directory and file system management	
Call	Description
s = mkdir(name, mode)	Create a new directory
s = rmdir(name)	Remove an empty directory
s = link(name1, name2)	Create a new entry, name2, pointing to name1
s = unlink(name)	Remove a directory entry
s = mount(special, name, flag)	Mount a file system
s = umount(special)	Unmount a file system

Miscellaneous	
Call	Description
s = chdir(dirname)	Change the working directory
s = chmod(name, mode)	Change a file's protection bits
s = kill(pid, signal)	Send a signal to a process
seconds = time(&seconds)	Get the elapsed time since Jan. 1, 1970

23



A stripped down shell Tan08 F.1-19

```

#define TRUE 1

while (TRUE) {
    type_prompt( );           /* repeat forever */
    read_command(command, parameters); /* display prompt on screen */
                                /* read input from terminal */

    if (fork() != 0) {        /* fork off child process */
        /* Parent code. */
        waitpid(-1, &status, 0); /* wait for child to exit */
    } else {
        /* Child code. */
        execve(command, parameters, 0); /* execute command */
    }
}
    
```


"Riisuttu" komentotulkki

Some Win32 API Calls

UNIX	Win32	Description
fork	CreateProcess	Create a new process
waitpid	WaitForSingleObject	Can wait for a process to exit
execve	(none)	CreateProcess = fork + execve
exit	ExitProcess	Terminate execution
open	CreateFile	Create a file or open an existing file
close	CloseHandle	Close a file
read	ReadFile	Read data from a file
write	WriteFile	Write data to a file
lseek	SetFilePointer	Move the file pointer
stat	GetFileAttributesEx	Get various file attributes
mkdir	CreateDirectory	Create a new directory
rmdir	RemoveDirectory	Remove an empty directory
link	(none)	Win32 does not support links
unlink	DeleteFile	Destroy an existing file
mount	(none)	Win32 does not support mount
umount	(none)	Win32 does not support mount
chdir	SetCurrentDirectory	Change the current working directory
chmod	(none)	Win32 does not support security (although NT does)
kill	(none)	Win32 does not support signals
time	GetLocalTime	Get the current time

25

Monolithic Systems



- Whole OS as one program in kernel mode
 - Collection of procedures, linked together
- Basic structure of a monolithic system
 - Main program invokes requested service procedure
 - Service procedures carry out the system calls
 - Utility procedures help service procedures
- There might be loadable extensions

Operating System Structures

Layered systems

Layer	Function
5	The operator
4	User programs
3	Input/output management
2	Operator-process communication
1	Memory and drum management
0	Processor allocation and multiprogramming

Structure of the THE operating system

- Generalization of the procedure approach
- Each layer has specific tasks
- Example: THE system
 - The layers mainly a design aid

Some examples

- Monolithic System
- Layered System
- Microkernel
- Client-Server Model
- Virtual Machines
- Exokernels

Micro-kernel





Figure 1-26. Structure of the MINIX 3 system.

- Split the OS to small, well-defined modules
- Only the microkernel module runs in kernel mode
- Other modules run in user mode
- Goal: increase availability
 - A buggy driver cannot crash the whole computer
 - MINIX has reincarnation server to automatically replace failed module
- Disadvantage:
 - A lot of mode switches within OS itself

Client-Server Model




- Client-server model is an abstraction that can be used on single computer or networked computers
- Clients send message saying what it wants over network or e.g. using microkernel

Exokernels

- Partitioning of the machine
- Each virtual machine gets only a subset of the existing resources (not the image of the whole machine)

Virtual machines



- Virtual machine abstraction was designed 40 years ago
- Basic idea:
 - Virtual machine monitor does multiplexing of the virtual machines
 - Each virtual machine is 'exact' copy of the bare hardware and has its own OS and applications

C and Metric Units

- Must be able to use:
 - Reading C-code examples
 - Unit conversions

Exp.	Explicit	Prefix	Exp.	Explicit	Prefix
10 ⁻³	0.001	milli	10 ³	1,000	Kilo
10 ⁻⁶	0.000001	micro	10 ⁶	1,000,000	Mega
10 ⁻⁹	0.000000001	nano	10 ⁹	1,000,000,000	Giga
10 ⁻¹²	0.0000000000001	pico	10 ¹²	1,000,000,000,000	Tera
10 ⁻¹⁵	0.0000000000000001	femto	10 ¹⁵	1,000,000,000,000,000	Peta
10 ⁻¹⁸	0.0000000000000000001	atto	10 ¹⁸	1,000,000,000,000,000,000	Exa
10 ⁻²¹	0.0000000000000000000001	zepto	10 ²¹	1,000,000,000,000,000,000,000	Zetta
10 ⁻²⁴	0.000000000000000000000001	yocto	10 ²⁴	1,000,000,000,000,000,000,000,000	Yotta

Virtualization today




Figure 1-29. (a) A type 1 hypervisor. (b) A type 2 hypervisor.

- Type 1 hypervisor on top of hardware multiplex all OS in parallel
- Type 2 hypervisor on top of host OS multiplex only guest OSes

- Virtual machine monitor renamed as hypervisor
- Java virtual machine uses the virtualization abstraction for code portability

History of Operating Systems

History of Operating Systems (1)

The diagram illustrates the workflow of an early batch system. It consists of six stages labeled (a) through (f):

- (a) A person brings a stack of cards to a computer labeled 1401.
- (b) The 1401 computer has a card reader and a tape drive.
- (c) A person puts a tape on the 1401.
- (d) The 1401 outputs a tape to a computer labeled 7094. The 7094 has an input tape and an output tape.
- (e) The 7094 outputs a tape to another computer labeled 1401.
- (f) The 1401 computer has a printer.

Early batch system

- bring cards to 1401
- read cards to tape
- put tape on 7094 which does computing
- put tape on 1401 which prints output

37

History of Operating Systems (2)

- First generation 1945 - 1955
 - vacuum tubes, plug boards
- Second generation 1955 - 1965
 - transistors, batch systems
- Third generation 1965 - 1980
 - ICs and multiprogramming
- Fourth generation 1980 - present
 - personal computers

38