

3. Operating-System Structures

Sungyoung Lee

*College of Engineering
KyungHee University*

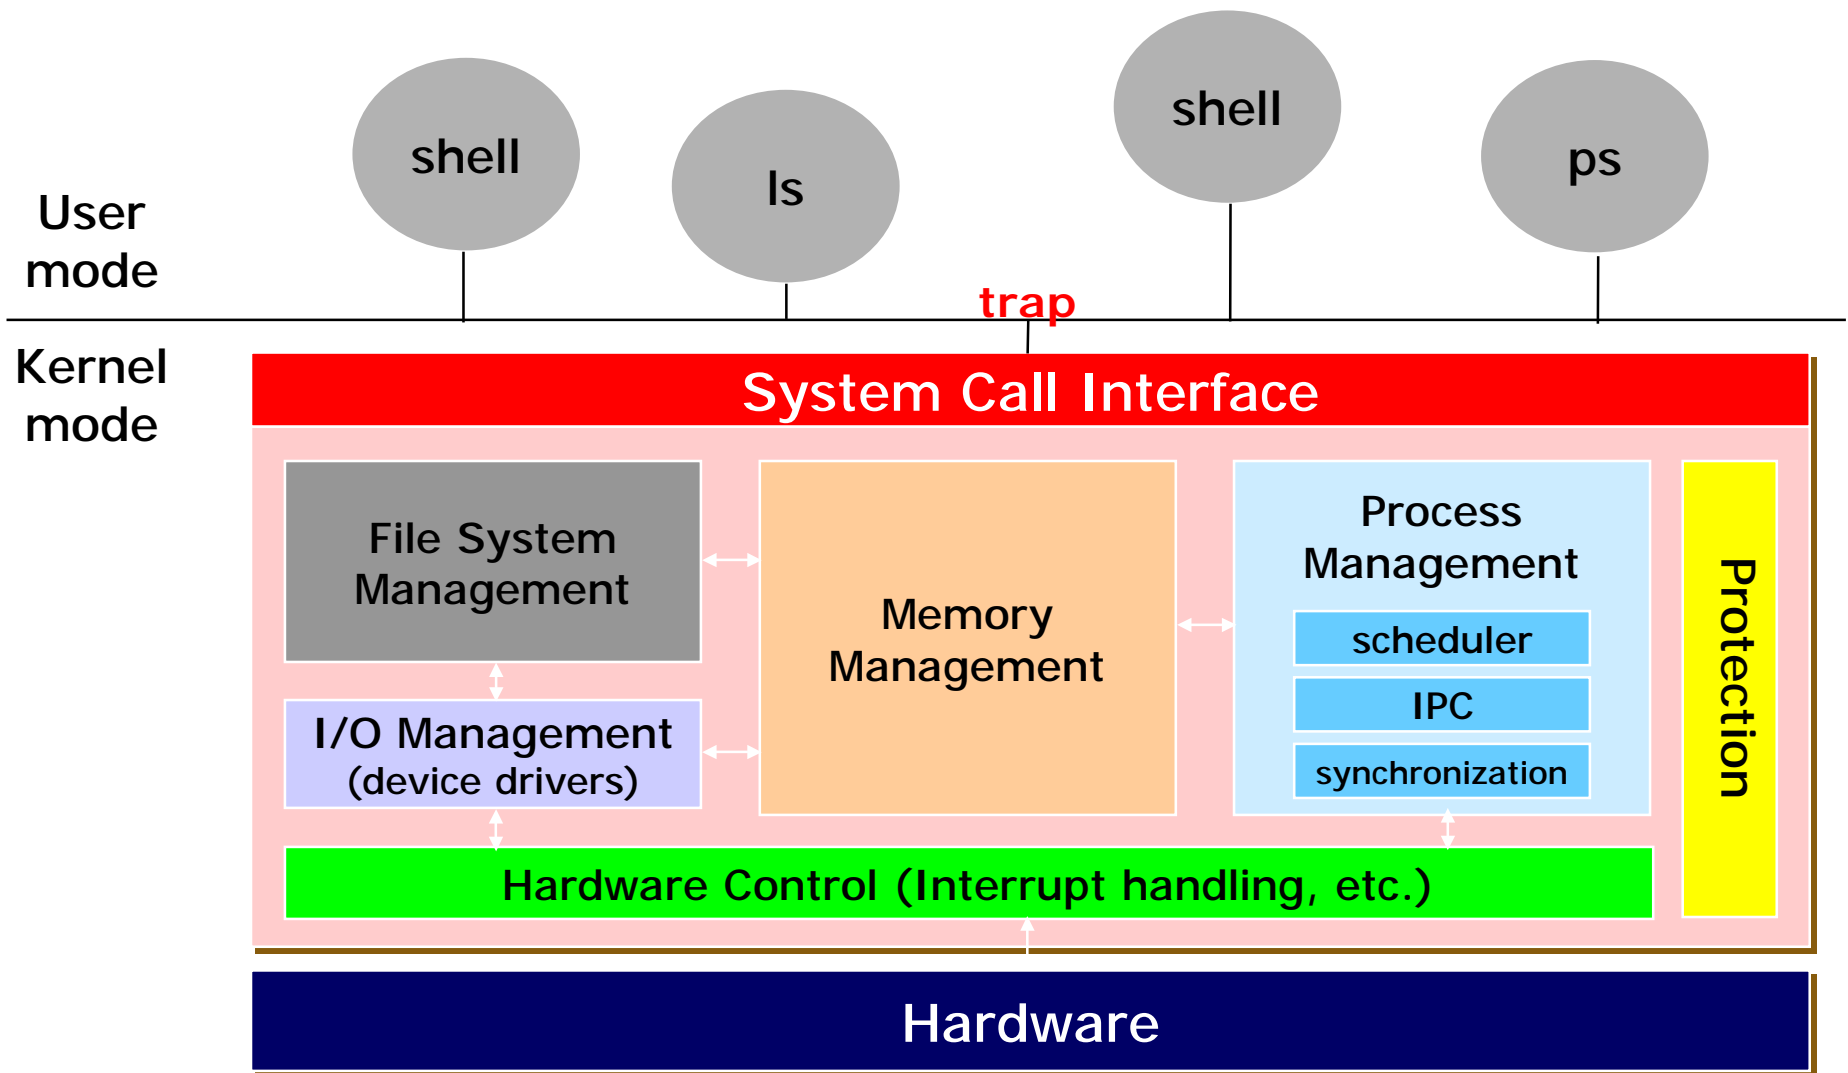
Contents

- n *System Components*
- n *Operating System Services*
- n *System Calls*
- n *System Programs*
- n *System Structure*
- n *Virtual Machines*
- n *System Design and Implementation*
- n *System Generation*

Common System Components

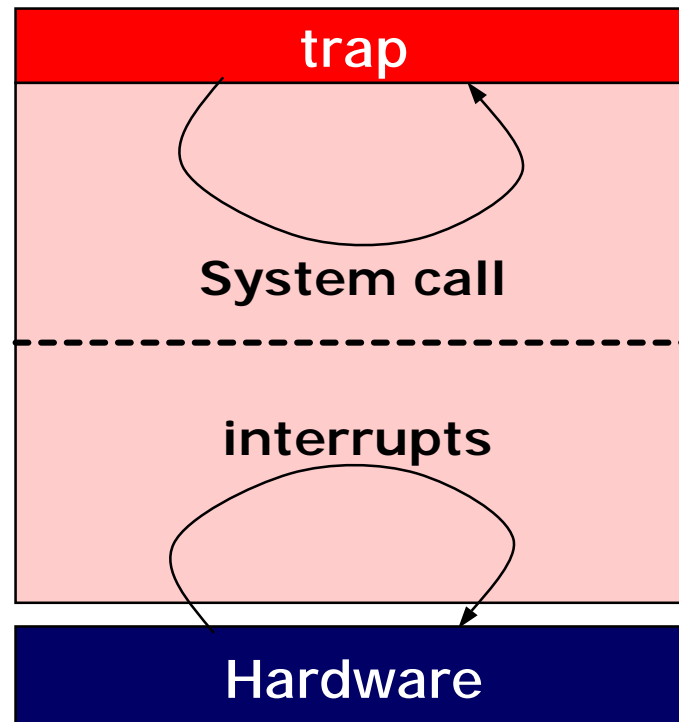
- n Process Management
- n Main Memory Management
- n File Management
- n I/O System Management
- n Secondary Management
- n Networking
- n Protection System
- n Command-Interpreter System

Operating System Structure



Taking Control of the System

- n Bootstrapping
- n System calls
- n Interrupts



Process Management

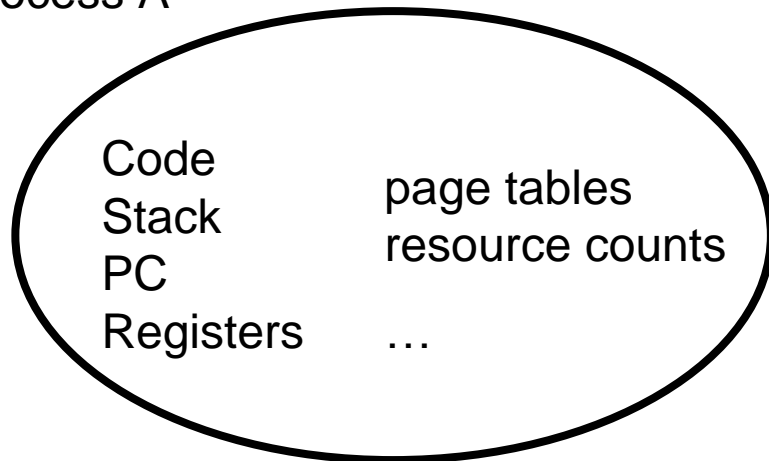
- n *A process is a program in execution*
 - ü A process needs certain resources, including CPU time, memory, files, and I/O devices, to accomplish its task

- n *The operating system is responsible for the following activities in connection with process management*
 - ü Process creation and deletion
 - ü process suspension and resumption
 - ü Provision of mechanisms for:
 - § process synchronization
 - § process communication

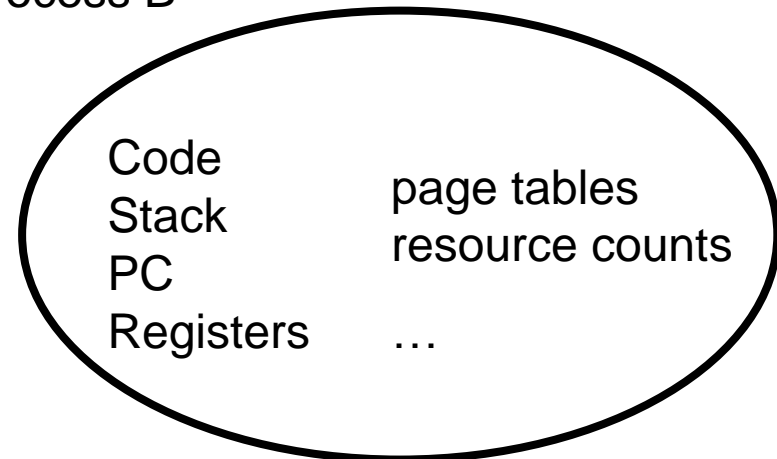
Program vs. Process

- n A program is a passive thing – just a file on the disk with code that is *potentially* runnable
- n A process is one instance of a program *in execution*; at any instance, there may be many processes running copies of a single program (e.g., an editor): each is a *separate, independent* process

Process A



Process B



Main-Memory Management

- n Memory is a large array of words or bytes, each with its own address
 - ü It is a repository of quickly accessible data shared by the CPU and I/O devices

- n Main memory is a volatile storage device
 - ü It loses its contents in the case of system failure

- n The operating system is responsible for the following activities in connections with memory management:
 - ü Keep track of which parts of memory are currently being used and by whom
 - ü Decide which processes to load when memory space becomes available
 - ü Allocate and deallocate memory space as needed

File Management

- n A file is a collection of related information defined by its creator
 - ü Commonly, files represent programs (both source and object forms) and data

- n The operating system is responsible for the following activities in connections with file management:
 - ü File creation and deletion
 - ü Directory creation and deletion
 - ü Support of primitives for manipulating files and directories
 - ü Mapping files onto secondary storage
 - ü File backup on stable (nonvolatile) storage media

n A convenient abstraction for the secondary storage

- ü Defines logical objects (files, directories)
- ü Defines logical operations

n File

- ü Named collection of persistent information
- ü The basic long-term storage unit

n Directory (folder)

- ü Named file that contains names of other files and metadata about those files

I/O System Management

- n The I/O system consists of:
 - ü A buffer-caching system
 - ü A general device-driver interface
 - ü Drivers for specific hardware devices

- n *I/O Abstraction*
 - ü The OS provides a standard interface between programs and devices
 - ü File system, Sockets, I/O devices, etc.

Secondary-Storage Management

- n Since main memory (*primary storage*) is volatile and too small to accommodate all data and programs permanently, the computer system must provide *secondary storage* to back up main memory

- n Most modern computer systems use disks as the principle on-line storage medium, for both programs and data

- n The operating system is responsible for the following activities in connection with disk management:
 - ü Free space management
 - ü Storage allocation
 - ü Disk scheduling

Networking (Distributed Systems)

- n A *distributed* system is a collection processors that do not share memory or a clock
- n Each processor has its own local memory
- n The processors in the system are connected through a communication network
- n Communication takes place using a *protocol*
- n A distributed system provides user access to various system resources
- n Access to a shared resource allows:
 - ü Computation speed-up
 - ü Increased data availability
 - ü Enhanced reliability

Protection System

- n *Protection* refers to a mechanism for controlling access by programs, processes, or users to both system and user resources

- n The protection mechanism must:
 - ü distinguish between authorized and unauthorized usage
 - ü specify the controls to be imposed
 - ü provide a means of enforcement

Protection System

- n Protection is a general mechanism throughout the OS

- n All resources objects need to protection
 - ü Memory
 - ü Processes
 - ü Files
 - ü Devices

- n Protection mechanisms help to detect errors as well as to prevent malicious destruction

Command-Interpreter System

- n Many commands are given to the operating system by control statements which deal with:
 - ü process creation and management
 - ü I/O handling
 - ü secondary-storage management
 - ü main-memory management
 - ü file-system access
 - ü protection
 - ü networking

- n The program that reads and interprets control statements is called variously:
 - ü command-line interpreter
 - ü shell (in UNIX)

- n Its function is to get and execute the next command statement

Command-Interpreter System

n Shell

- ü A particular program that handles the interpretation of users' commands
- ü Helps to manage processes

n Types

- ü A standard part of the OS
 - § MS-DOS, Apple II
- ü A non-privileged process
 - § sh / csh / tcsh / zsh / ksh on UNIX
- ü No command interpreter
 - § MacOS

Operating System Services

n Program execution

- ü system capability to load a program into memory and to run it

n I/O operations

- ü since user programs cannot execute I/O operations directly, the operating system must provide some means to perform I/O

n File-system manipulation

- ü program capability to read, write, create, and delete files

n Communications

- ü exchange of information between processes executing either on the same computer or on different systems tied together by a network (Implemented via *shared memory* or *message passing*)

n Error detection

- ü ensure correct computing by detecting errors in the CPU and memory hardware, in I/O devices, or in user programs

Additional Operating System Functions

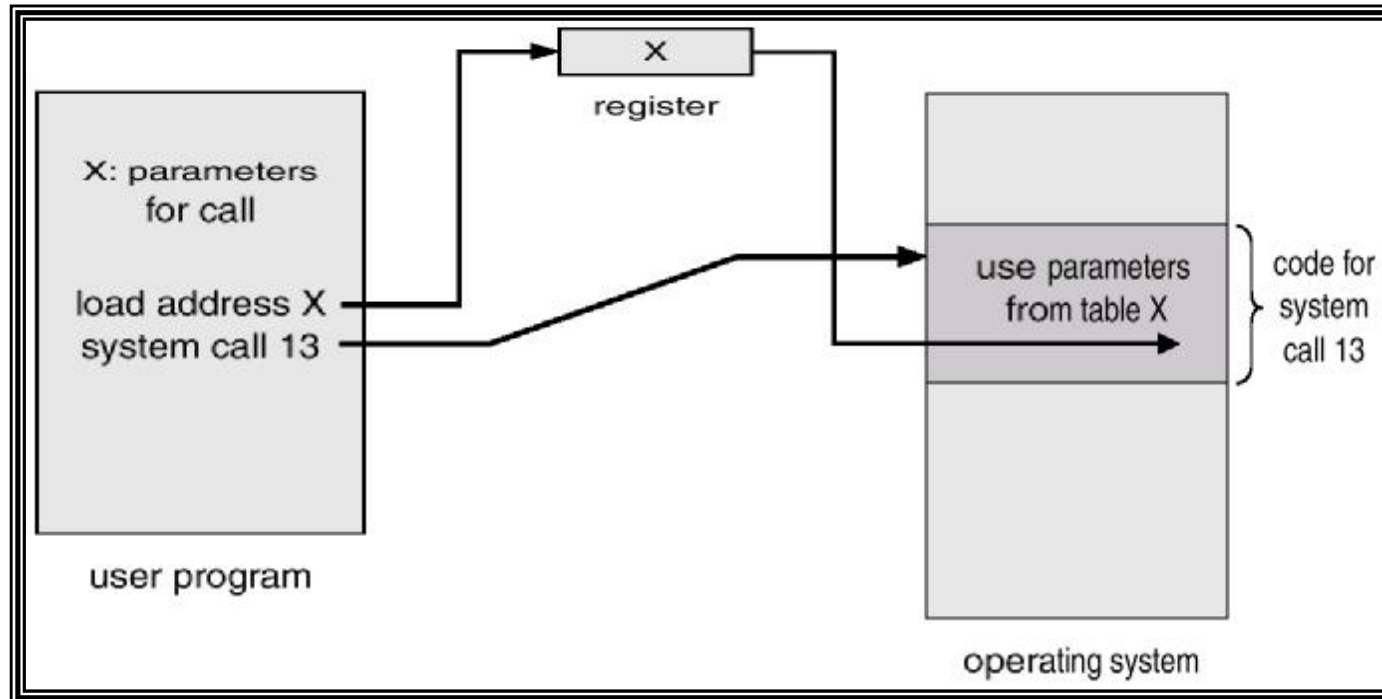
- n Additional functions exist not for helping the user, but rather for ensuring efficient system operations
 - ü Resource allocation
 - § allocating resources to multiple users or multiple jobs running at the same time
 - ü Accounting
 - § keep track of and record which users use how much and what kinds of computer resources for account billing or for accumulating usage statistics
 - ü Protection
 - § ensuring that all access to system resources is controlled

System Calls

- n System calls provide the interface between a running program and the operating system
 - ü Generally available as assembly-language instructions
 - ü Languages defined to replace assembly language for systems programming allow system calls to be made directly (e.g., C, C++)

- n Three general methods are used to pass parameters between a running program and the operating system
 - ü Pass parameters in *registers*
 - ü Store the parameters in a table in memory, and the table address is passed as a parameter in a register
 - ü *Push* (store) the parameters onto the *stack* by the program, and *pop* off the stack by operating system

Passing of Parameters As A Table



Types of System Calls

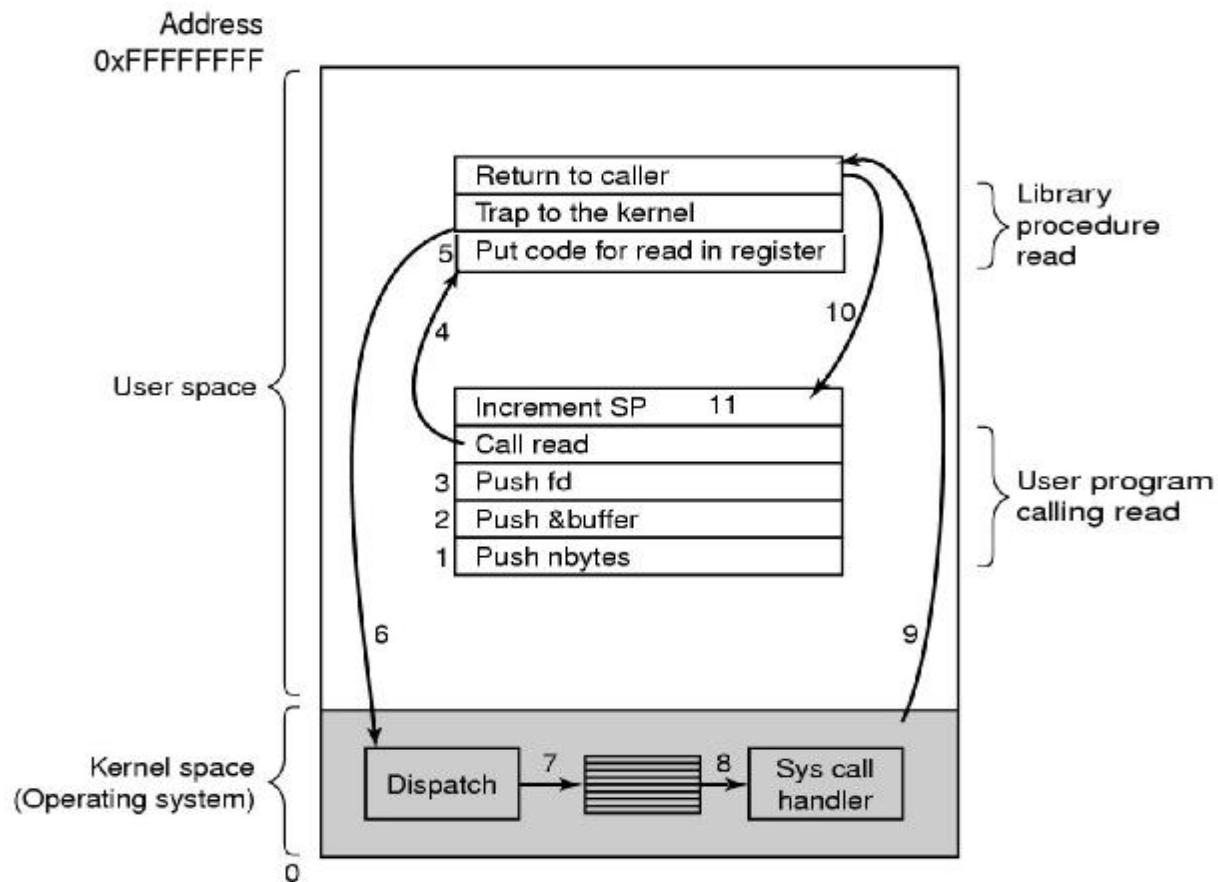
- n Process control
- n File management
- n Device management
- n Information maintenance
- n Communications

System Calls

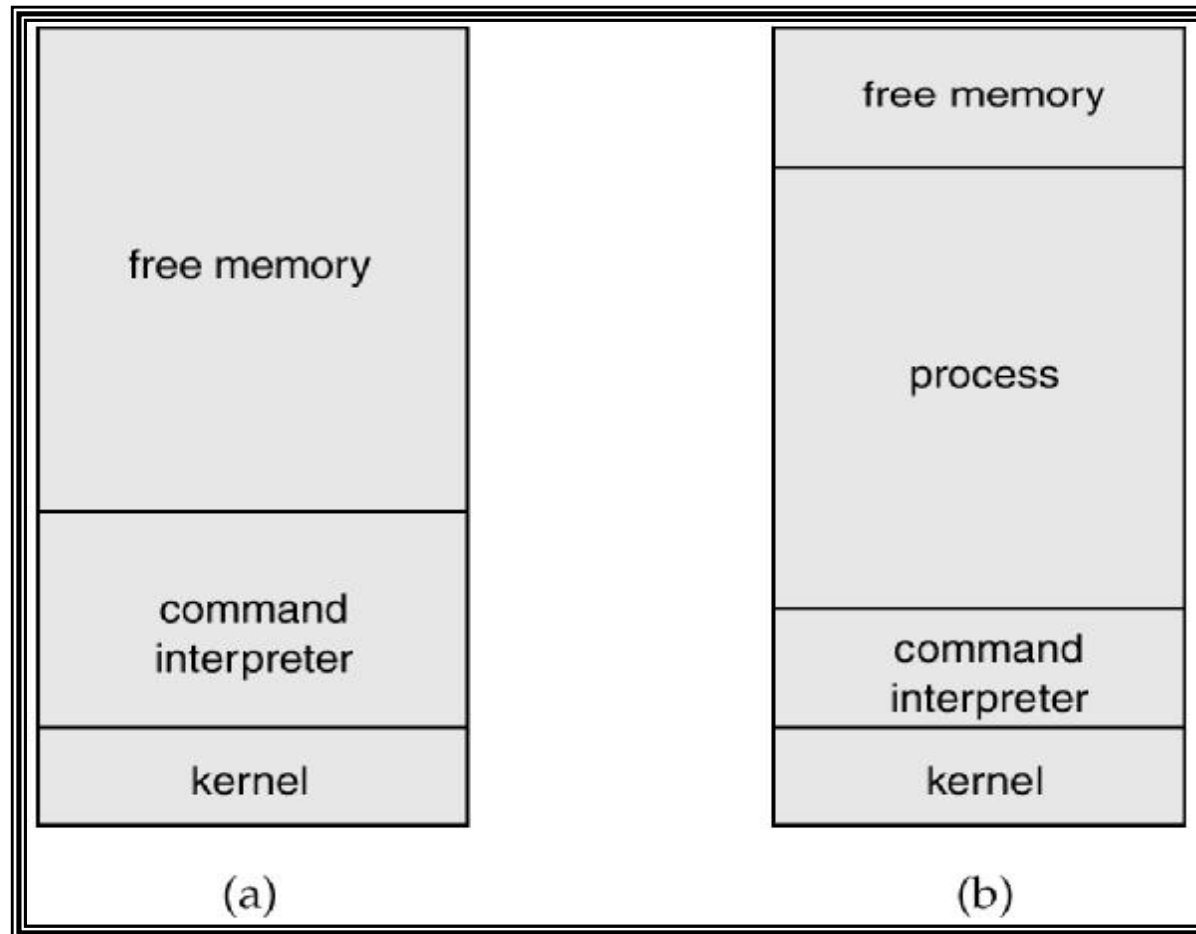
Process Management	fork	CreateProcess	Create a new process
	waitpid	WaitForSingleObject	Wait for a process to exit
	execve	(none)	CreateProcess = fork + execve
	exit	ExitProcess	Terminate execution
	kill	(none)	Send a signal
File Management	open	CreateFile	Create a file or open an existing file
	close	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
	chmod	(none)	Change the file access permission
File System Management	mkdir	CreateDirectory	Create a new directory
	rmdir	RemoveDirectory	Remove an empty directory
	link	(none)	Make a link to a file
	unlink	DeleteFile	Destroy an existing file
	mount	(none)	Mount a file system
	umount	(none)	Unmount a file system
	chdir	SetCurrentDirectory	Change the current working directory

Invoking a System Call

```
count = read (fd, buffer, nbytes);
```



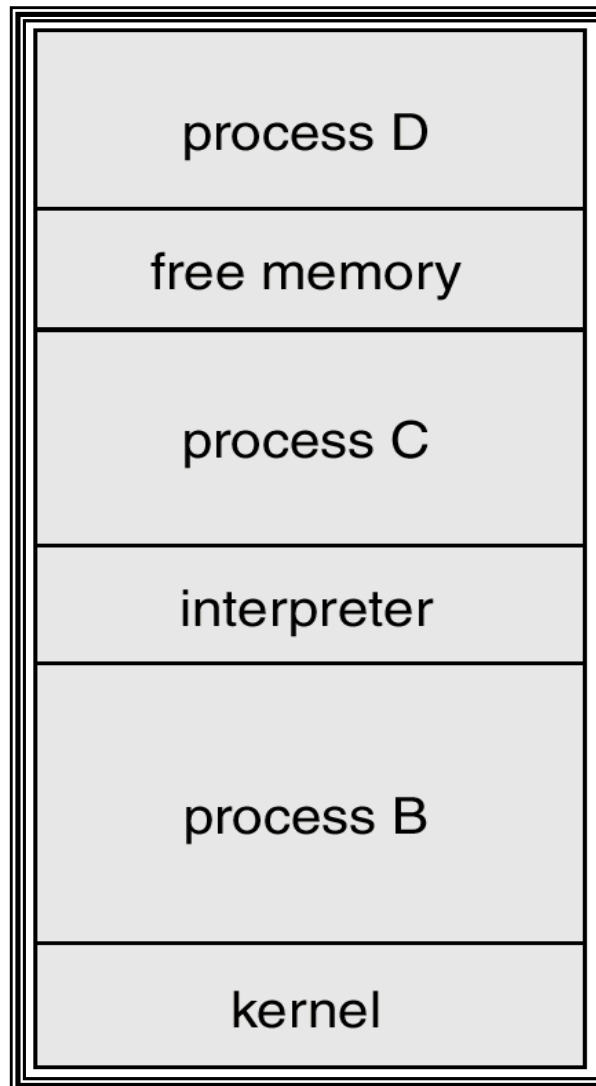
MS-DOS Execution



At System Start-up

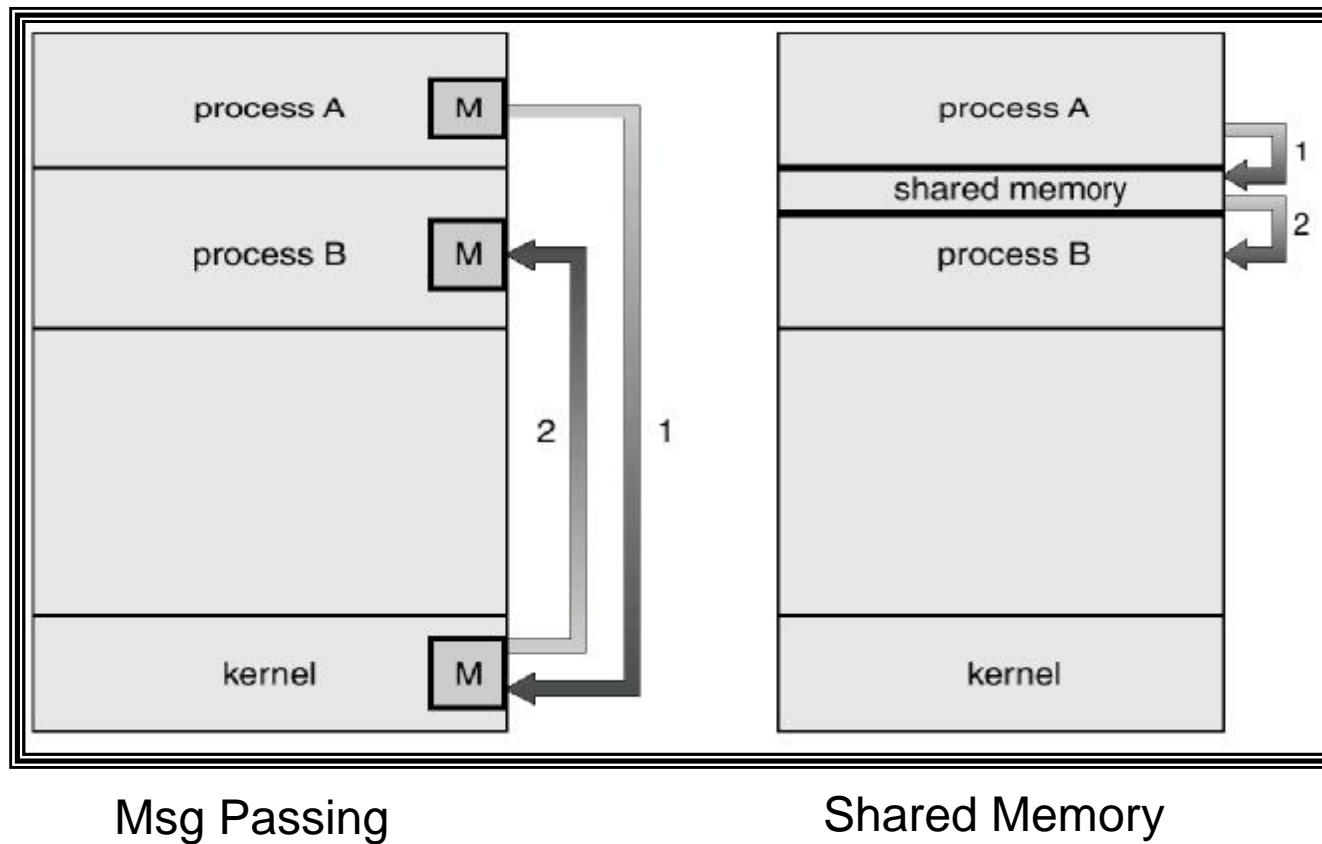
Running a Program

UNIX Running Multiple Programs



Communication Models

- Communication may take place using either message passing or shared memory



System Programs

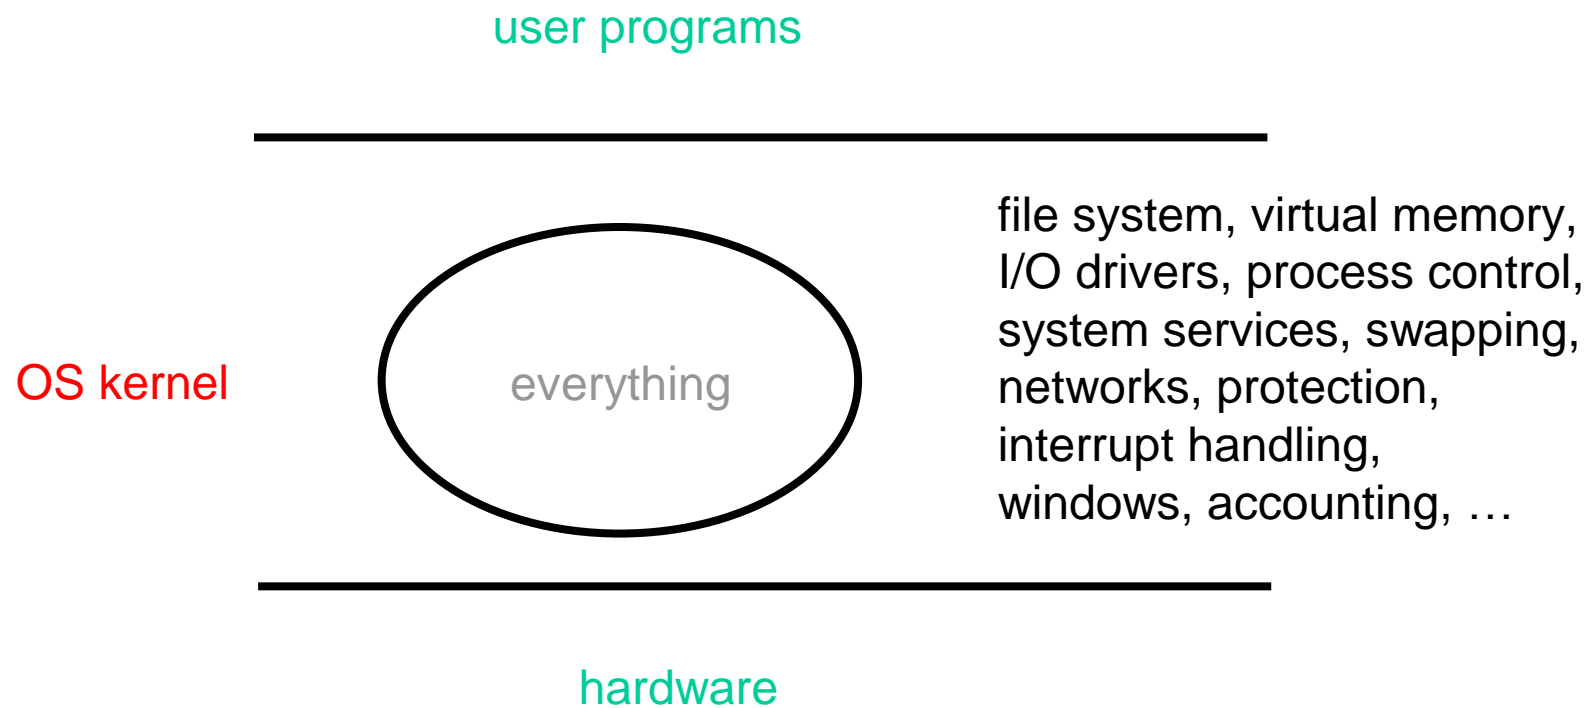
- n System programs provide a convenient environment for program development and execution. They can be divided into:
 - ü File manipulation
 - ü Status information
 - ü File modification
 - ü Programming language support
 - ü Program loading and execution
 - ü Communications
 - ü Application programs

- n Most users' view of the operation system is defined by system programs, not the actual system calls

- n An OS consists of all of these components, plus lots of others, plus system service routines, plus system programs(privileged and non-privileged), plus ...
- n The big issue:
 - ü How do we organize all of this?
 - ü What are the entities and where do they exist?
 - ü How does these entities cooperate?
- n Basically, how do we build a complex system that's:
 - ü Performance
 - ü Reliable
 - ü Extensible

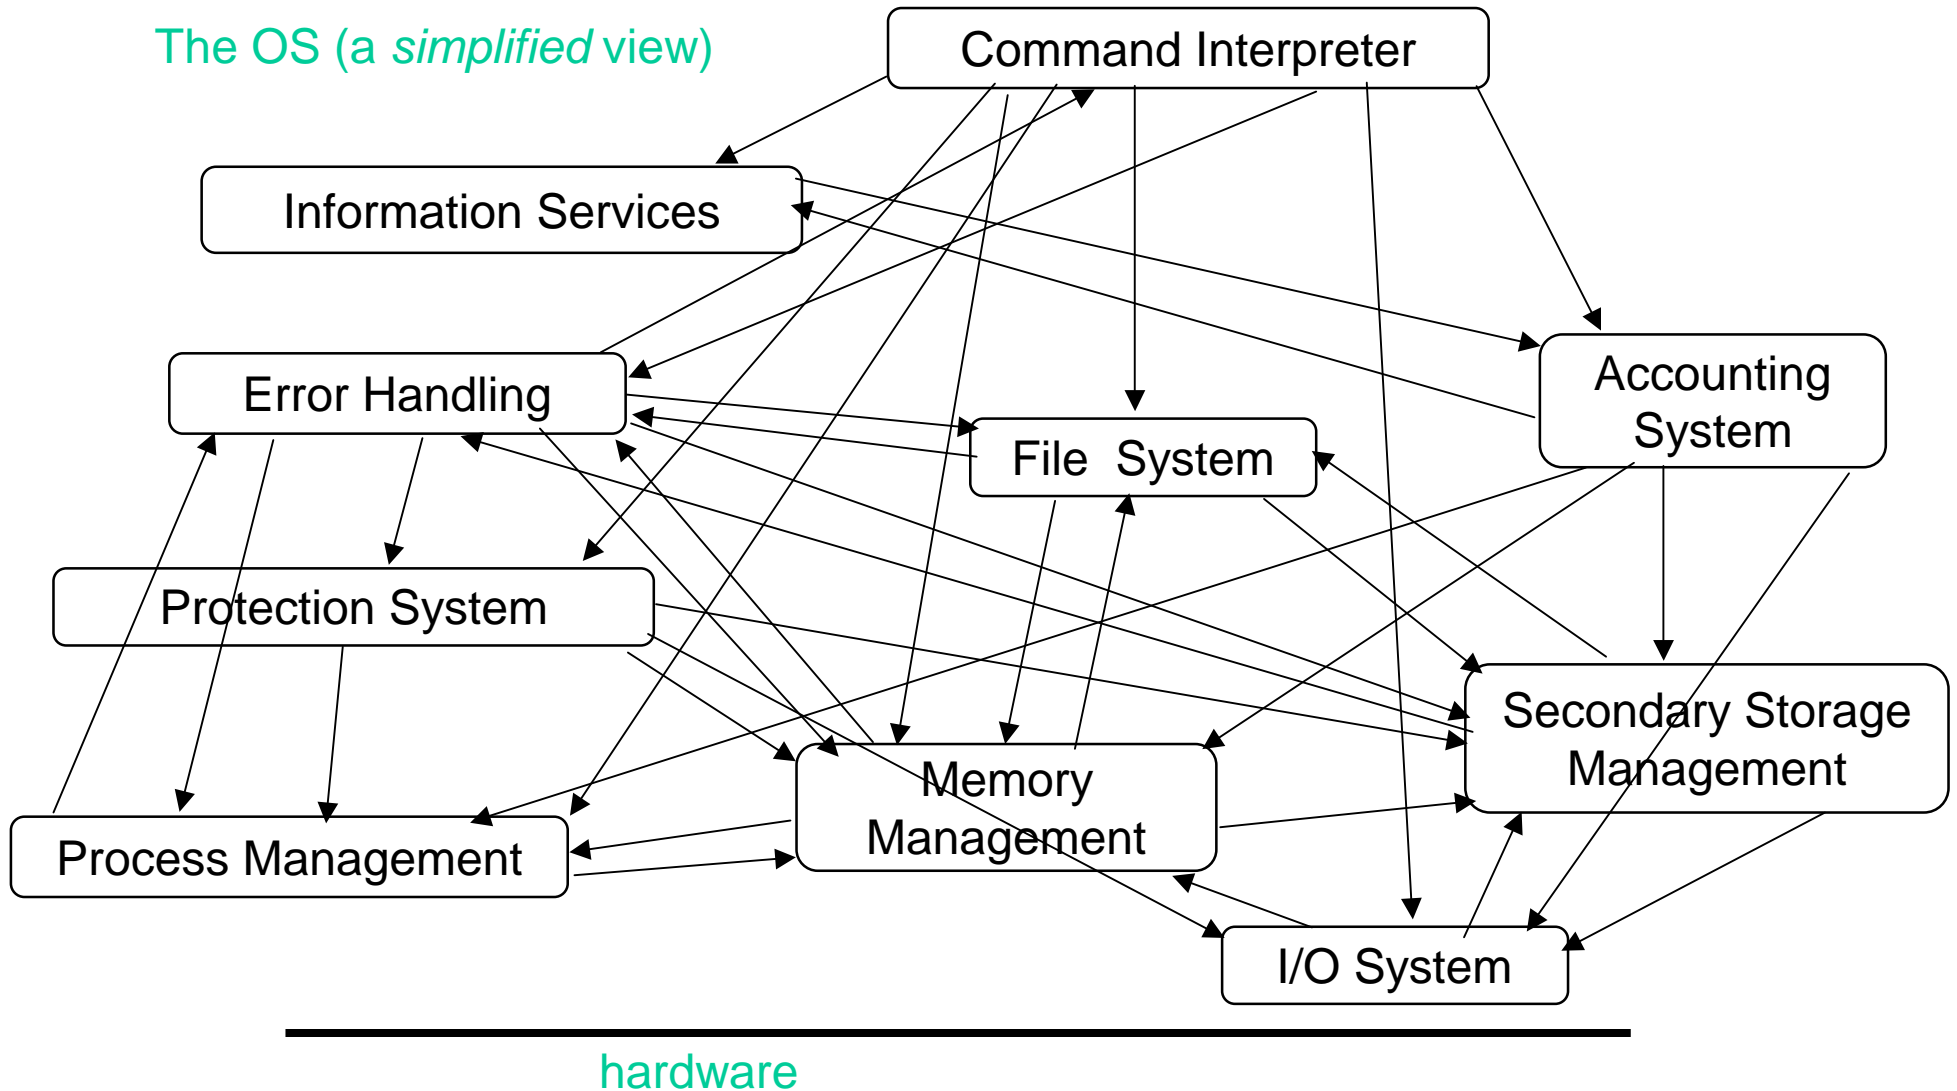
OS Structure (Cont'd)

n Traditionally, systems such as Unix were built as a *monolithic* kernel:



OS Structure (Cont'd)

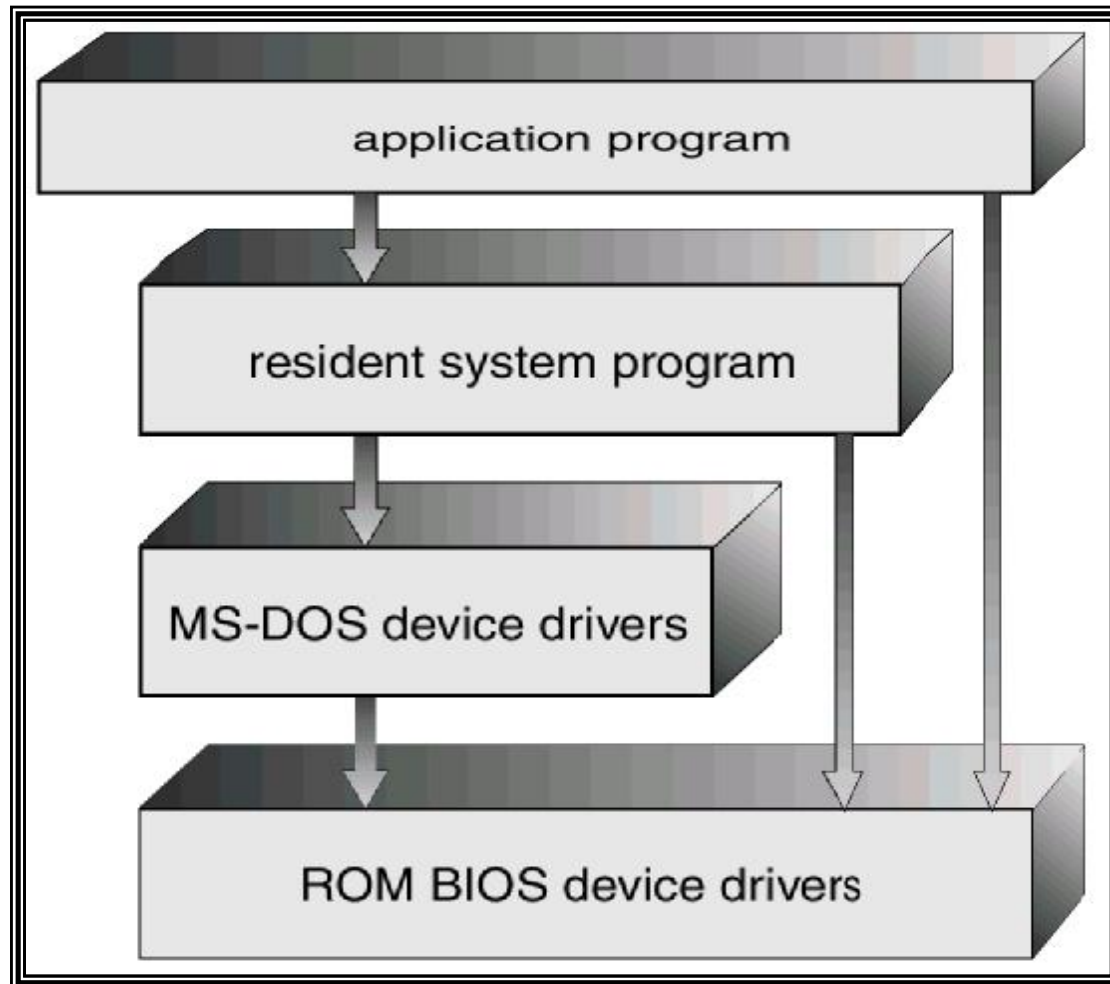
The OS (a *simplified* view)



MS-DOS System Structure

- n MS-DOS – written to provide the most functionality in the least space
 - ü not divided into modules
 - ü Although MS-DOS has some structure, its interfaces and levels of functionality are not well separated

MS-DOS Layer Structure

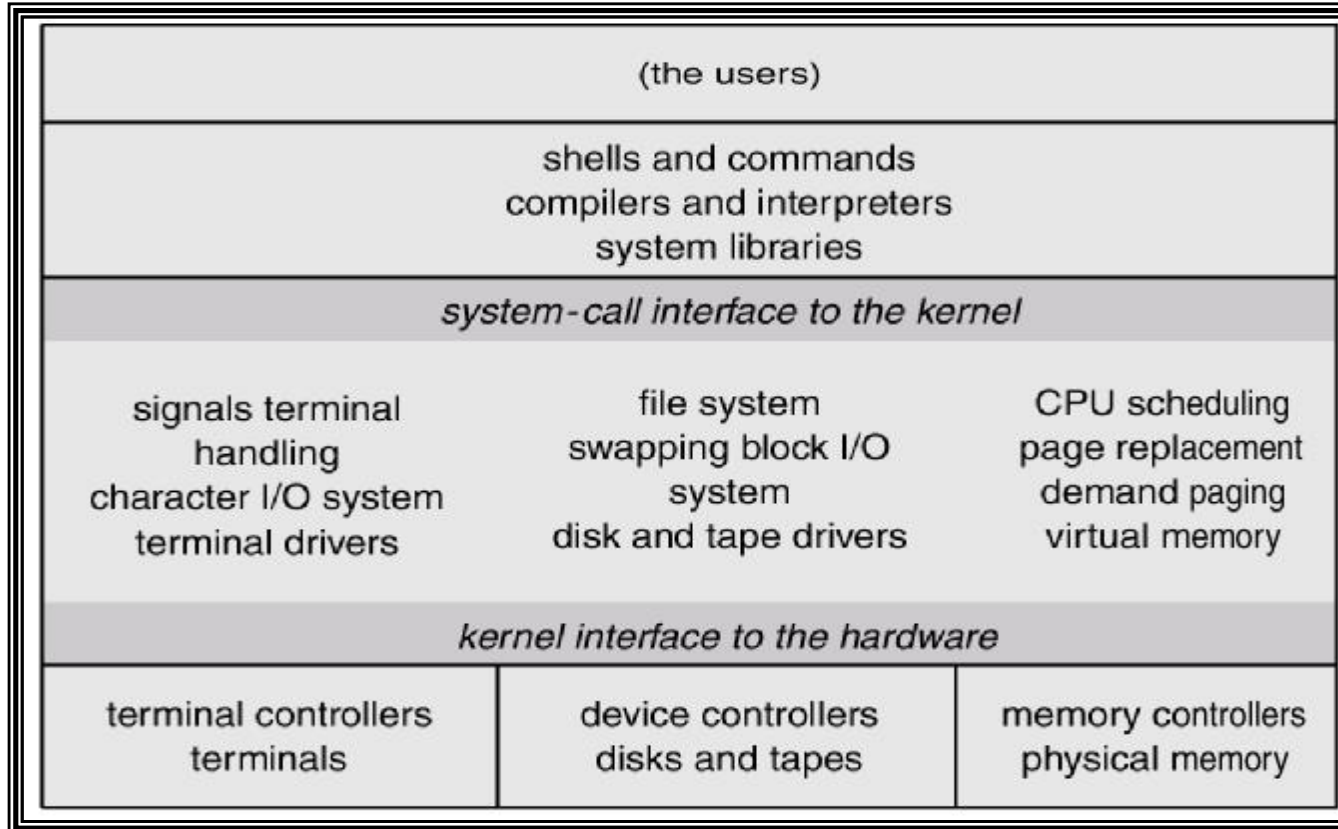


UNIX System Structure

- n UNIX – limited by hardware functionality, the original UNIX operating system had limited structuring

- n The UNIX OS consists of two separable parts
 - ü Systems programs
 - ü The kernel
 - § Consists of everything below the system-call interface and above the physical hardware
 - § Provides the file system, CPU scheduling, memory management, and other operating-system functions; a large number of functions for one level

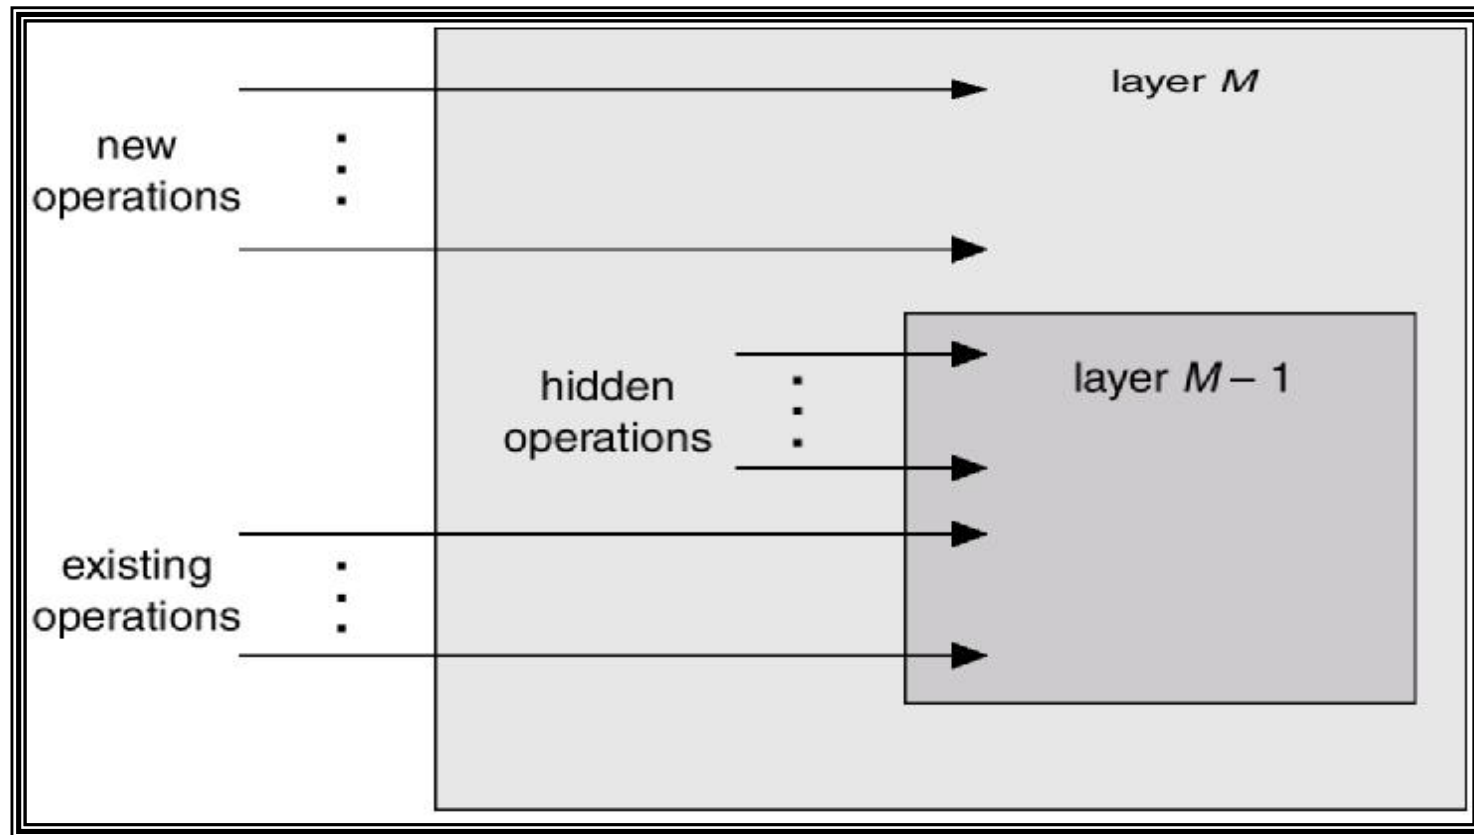
UNIX System Structure (Cont'd)



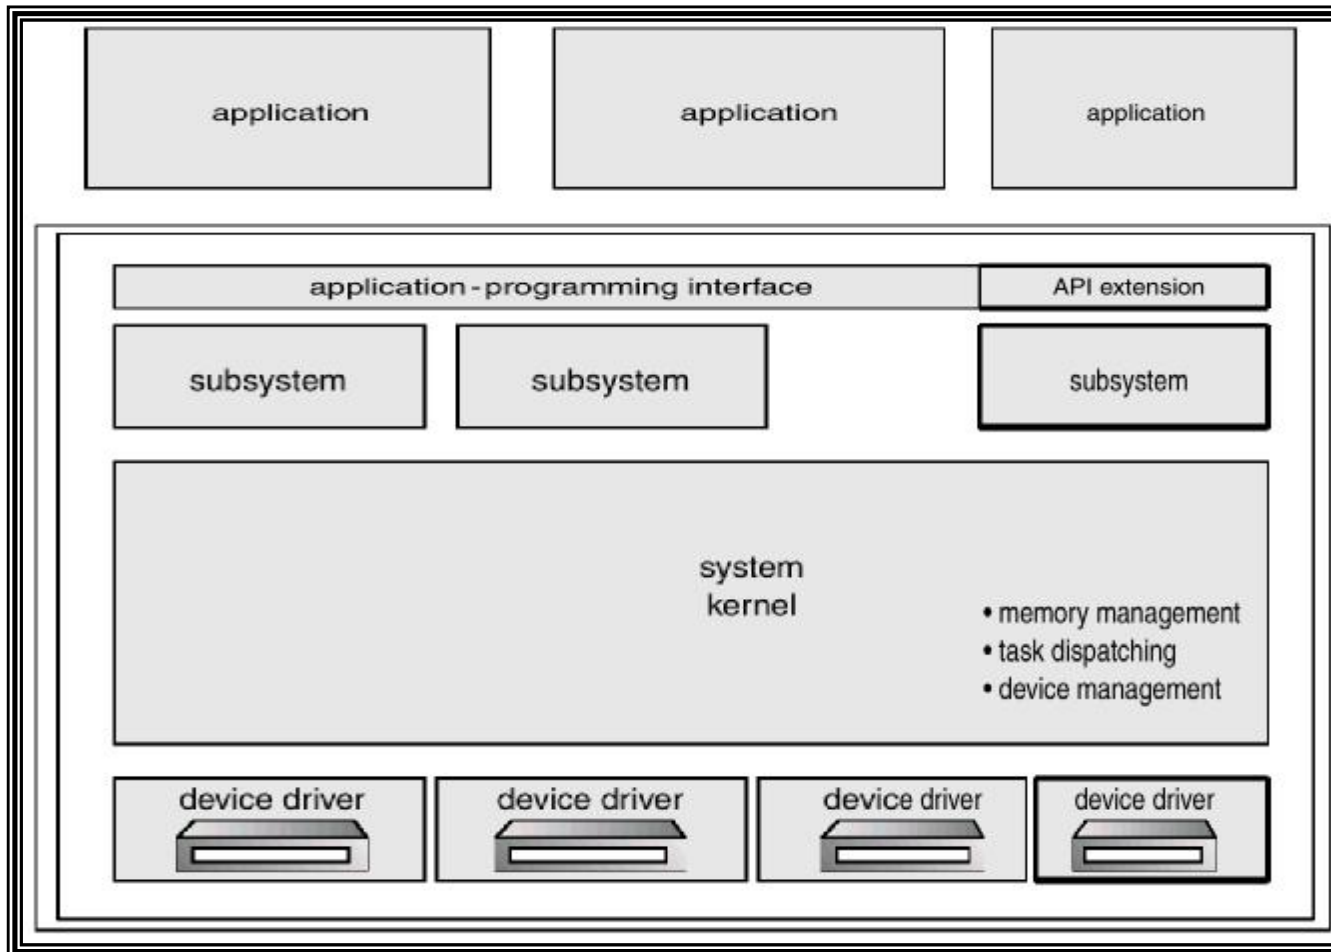
Layered Approach

- n The operating system is divided into a number of layers (levels), each built on top of lower layers
- n The bottom layer (layer 0), is the hardware; the highest (layer N) is the user interface
- n With modularity, layers are selected such that each uses functions (operations) and services of only lower-level layers

An Operating System Layer



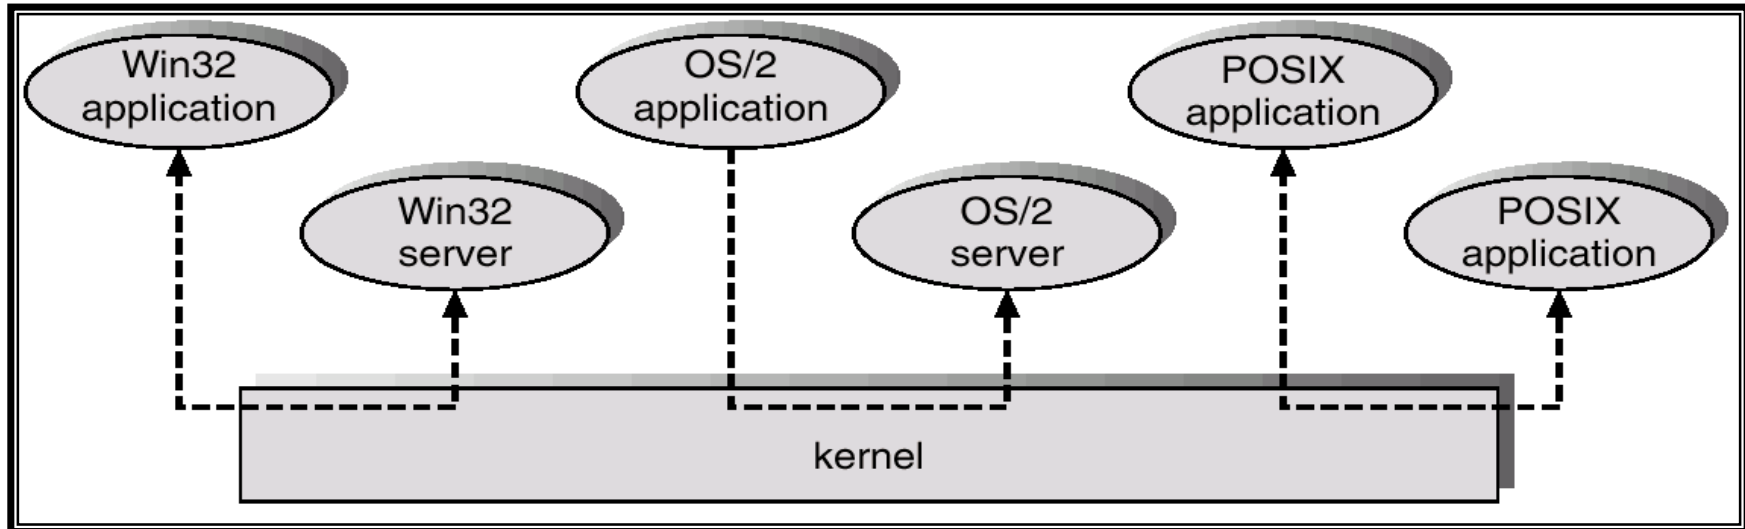
OS/2 Layer Structure



Microkernel System Structure

- n Moves as much from the kernel into “*user*” space
- n Communication takes place between user modules using message passing
- n Benefits:
 - ü easier to extend a microkernel
 - ü easier to port the operating system to new architectures
 - ü more reliable (less code is running in kernel mode)
 - ü more secure

Windows NT Client-Server Structure



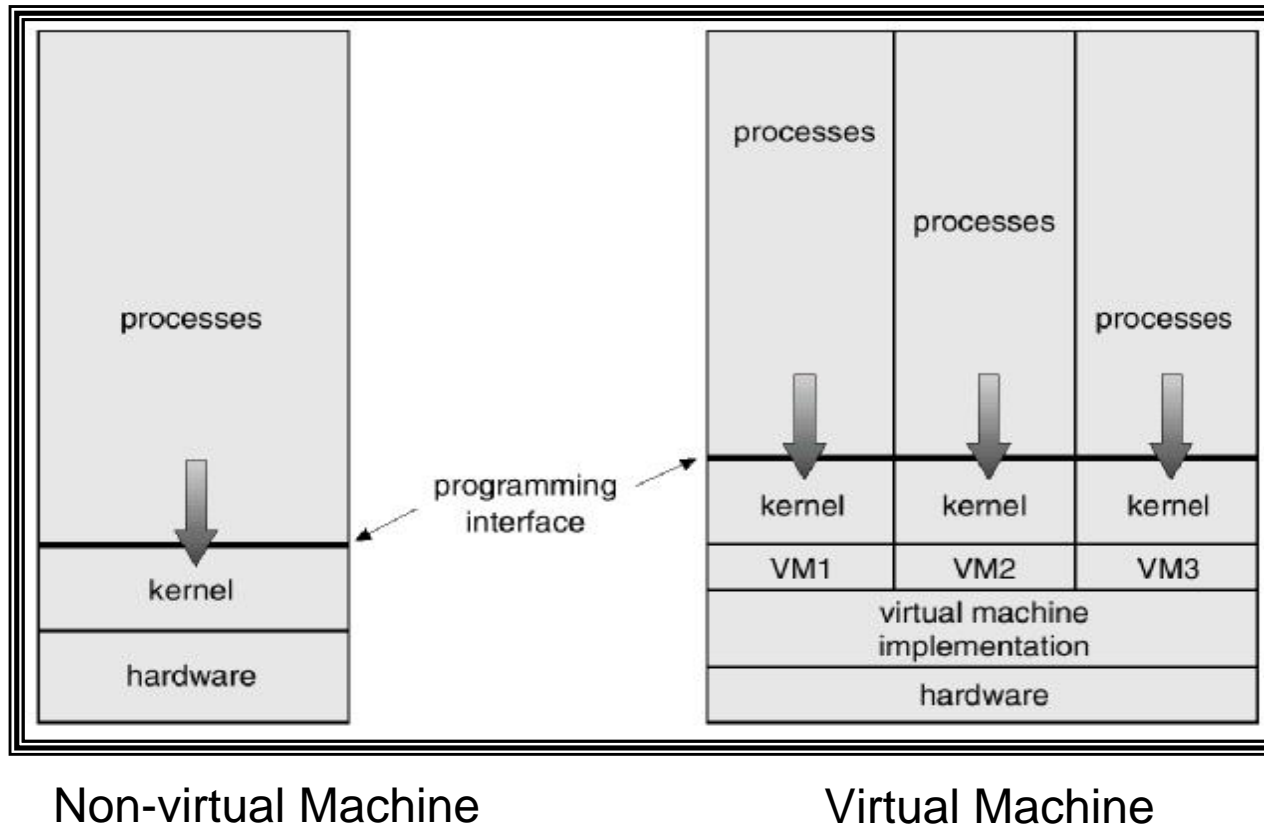
Virtual Machines

- n A *virtual machine* takes the layered approach to its logical conclusion
 - ü It treats hardware and the operating system kernel as though they were all hardware
- n A virtual machine provides an interface *identical* to the underlying bare hardware
- n The operating system creates the illusion of multiple processes, each executing on its own processor with its own (virtual) memory

Virtual Machines (Cont'd)

- n The resources of the physical computer are shared to create the virtual machines
 - ü CPU scheduling can create the appearance that users have their own processor
 - ü Spooling and a file system can provide virtual card readers and virtual line printers
 - ü A normal user time-sharing terminal serves as the virtual machine operator's console

System Models




Advantages/Disadvantages of Virtual Machines

- n The virtual-machine concept provides complete protection of system resources since each virtual machine is isolated from all other virtual machines
- n This isolation, however, permits no direct sharing of resources
- n A virtual-machine system is a perfect vehicle for operating-systems research and development
- n System development is done on the virtual machine, instead of on a physical machine and so does not disrupt normal system operation
- n The virtual machine concept is difficult to implement due to the effort required to provide an *exact* duplicate to the underlying machine

Java Virtual Machine

- n Compiled Java programs are platform-neutral byte-codes executed by a Java Virtual Machine (JVM)
- n JVM consists of
 - class loader
 - class verifier
 - runtime interpreter
- n Just-In-Time (JIT) compilers increase performance

Java Virtual Machine



System Design Goals

n User goals

- ü operating system should be convenient to use, easy to learn, reliable, safe, and fast

n System goals

- ü operating system should be easy to design, implement, and maintain, as well as flexible, reliable, error-free, and efficient

Mechanisms and Policies

- n Mechanisms determine how to do something, policies decide what will be done
- n The separation of policy from mechanism is a very important principle, it allows maximum flexibility if policy decisions are to be changed later

Mechanisms and Policies

n Policy

- ü *What* should be done?

n Mechanism

- ü *How* to do something?

- ü Policies are likely to change across places or over time.

- ü A general mechanism is desirable.

- ü A change in policy would then require redefinition of only certain parameters of the system instead of resulting in a change in the mechanism.

System Implementation

- n Traditionally written in assembly language, operating systems can now be written in higher-level languages
- n Code written in a high-level language:
 - ü can be written faster
 - ü is more compact
 - ü is easier to understand and debug
- n An operating system is far easier to *port* (move to some other hardware) if it is written in a high-level language

System Generation (SYSGEN)

- n Operating systems are designed to run on any of a class of machines; the system must be configured for each specific computer site
- n SYSGEN program obtains information concerning the specific configuration of the hardware system
- n *Booting*
 - ü starting a computer by loading the kernel
- n *Bootstrap program*
 - ü code stored in ROM that is able to locate the kernel, load it into memory, and start its execution

n Linux Booting Process

- ü The CPU initializes itself and then execute an instruction at a fixed location (0xffffffff).
- ü This instruction jumps into the BIOS.
- ü The BIOS finds a boot device and fetches its MBR (Master Boot Record), which points to LILO (Linux Loader).
- ü The BIOS loads and transfers control to LILO.
- ü LILO loads the compressed kernel.
- ü The compressed kernel decompresses itself and transfers control to the uncompressed kernel.