

Chapter 3: Operating-System Structures

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

Common System Components

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

Process Management

- A *process* is a program in execution. A process needs certain resources, including CPU time, memory, files, and I/O devices, to accomplish its task.
- The operating system is responsible for the following activities in connection with process management.
 - ☞ Process creation and deletion.
 - ☞ Process suspension and resumption.
 - ☞ Mechanisms for:
 - 📄 process synchronization
 - 📄 process communication

Main-Memory Management

- Memory is a large array of words or bytes, each with its own address. It is a repository of quickly accessible data shared by the CPU and I/O devices.
- Main memory is a volatile storage device. It loses its contents in the case of system failure.
- The operating system is responsible for the following activities in connections with memory management:
 - ☞ Keep track of which parts of memory are currently being used and by whom.
 - ☞ Decide which processes to load when memory space becomes available.
 - ☞ Allocate and deallocate memory space as needed.

File Management

- A file is a collection of related information defined by its creator. Commonly, files represent programs (both source and object forms) and data.
- The operating system is responsible for the following activities in connections with file management:
 - ☞ File creation and deletion.
 - ☞ File organization (directory creation and deletion).
 - ☞ Support of primitives for manipulating files and directories (read/write).
 - ☞ Mapping files onto secondary storage.

I/O System Management

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

Secondary-Storage Management

- 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.
- Most modern computer systems use disks as the principle on-line storage medium, for both programs and data.
- The operating system is responsible for the following activities in connection with disk management:
 - ☞ Free space management
 - ☞ Storage allocation
 - ☞ Reliability (RAID)
 - ☞ Disk scheduling

Networking (Distributed Systems)

- A *distributed* system is a collection of processors that do not share memory or a clock.
- Communication takes place using a *protocol*:
 - ☞ FTP: examine secondary storage (list files, read contents) and remotely alter it (add, delete files).
 - ☞ HTTP: like FTP but low-setup overhead, ideal for quick transfer of small content.
 - ☞ POP/IMAP: like FTP but partial contents of mailbox file are transferred and changed. POP has single mailbox file; IMAP supports folder-oriented organization.

Protection System

- *Protection* refers to a mechanism for controlling access by programs, processes, or users to both system and user resources.
- The protection mechanism must:
 - ☞ allow user to specify the controls to be imposed (e.g. user can set file access permissions).
 - ☞ provide a means of enforcement (e.g. ensure user is identified before an attempt is made to access a file).
 - ☞ distinguish between authorized and unauthorized usage.
- **Security:**
 - ☞ security vs. protection: external vs. internal.
 - ☞ covers user authentication, logging/auditing trails, encrypted communications

Command-Interpreter System

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

Command-Interpreter System (Cont.)

- The program that reads and interprets control statements (typed text) is called variously:

- ☞ command-line interpreter
- ☞ shell (in UNIX: csh, tcsh, bash)

Gets and executes the next command statement.

- Statement can be built-in command (MS-DOS) or a regular program that calls system calls and formats result for user (UNIX).
- Alternative is GUI: control defined via user actions (mouse or joystick clicks, touch screen). Most such systems still have command-line interpreter for administration (often a security loophole, e.g. voting machine tampering).

Operating System Services

- Program execution – system capability to load a program into memory and to run it.
- I/O operations – since user programs cannot execute I/O operations directly, the operating system must provide some means to perform I/O.
- File-system manipulation – program capability to read, write, create, and delete files.
- Communications – 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*.
- Error detection – ensure correct computing by detecting errors in the CPU and memory hardware, in I/O devices, or in user programs.

Additional Operating System Functions

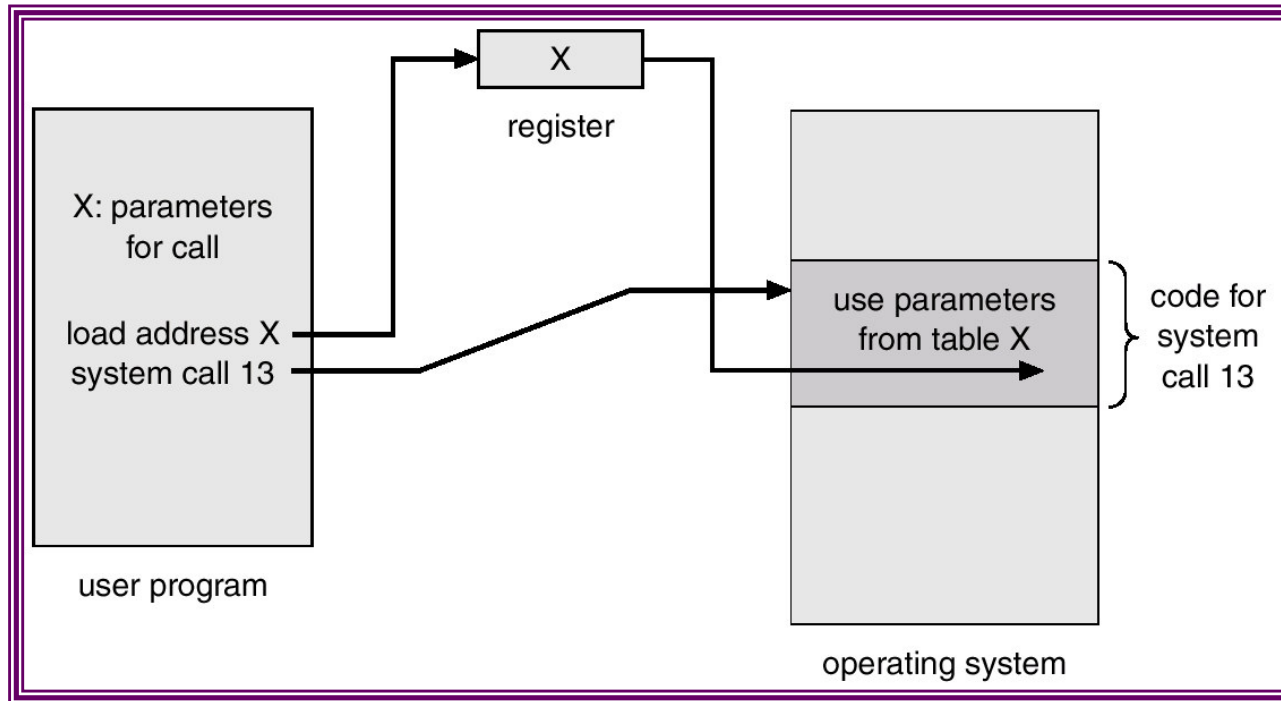
Additional functions exist not for helping the user, but rather for ensuring efficient system operations:

- Resource allocation – allocating resources to multiple users or multiple jobs running at the same time.
- 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.
- Protection – ensuring that all access to system resources is controlled.

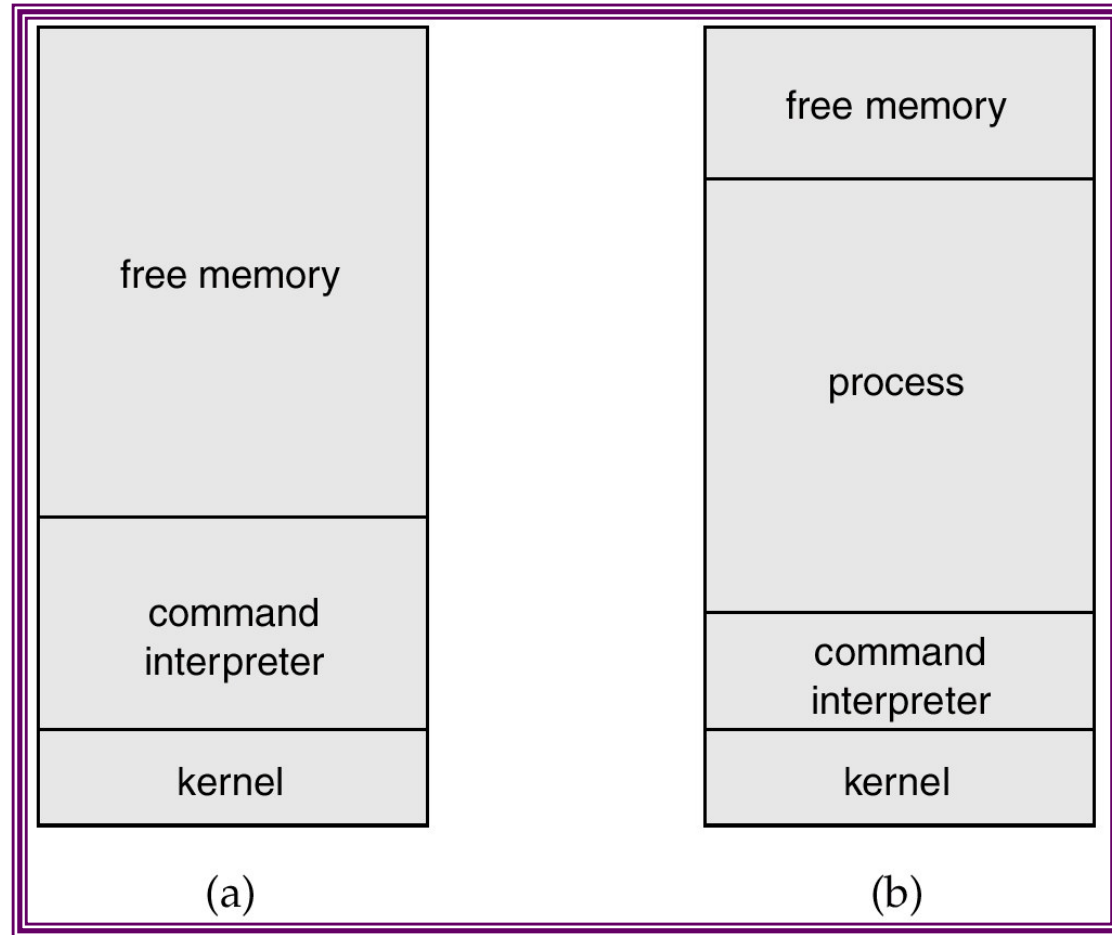
System Calls

- System calls provide the interface between a running program and the operating system.
 - ☞ Generally available as assembly-language instructions.
 - ☞ Languages defined to replace assembly language for systems programming allow system calls to be made directly (e.g., C, C++)
- Three general methods are used to pass parameters between a running program and the operating system.
 - ☞ Pass parameters in *registers*.
 - ☞ Store the parameters in a table in memory, and the table address is passed as a parameter in a register.
 - ☞ *Push* (store) the parameters onto the *stack* by the program, and *pop* off the stack by operating system.

Passing of Parameters As A Table



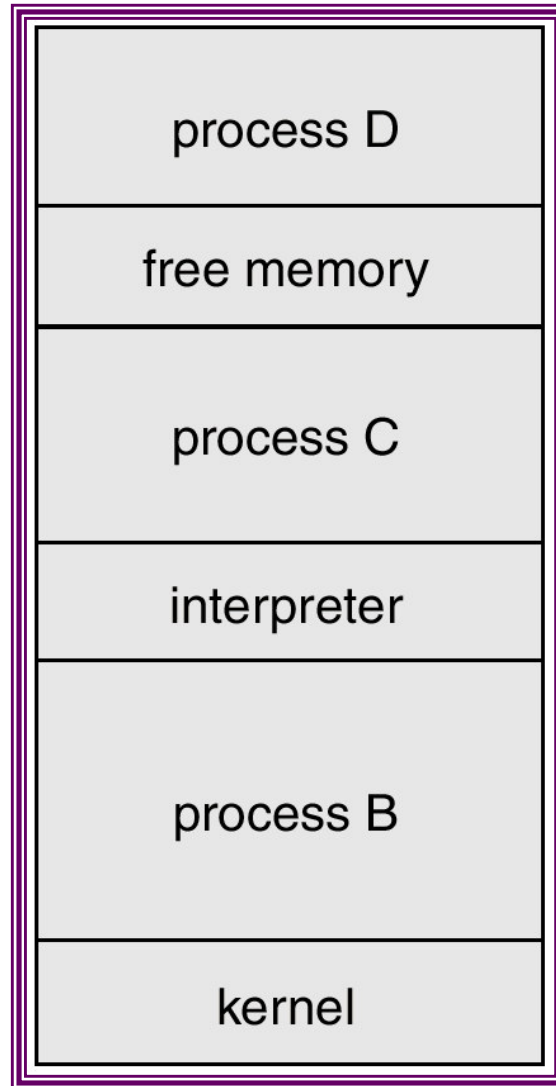
Process Management: MS-DOS



At System Start-up

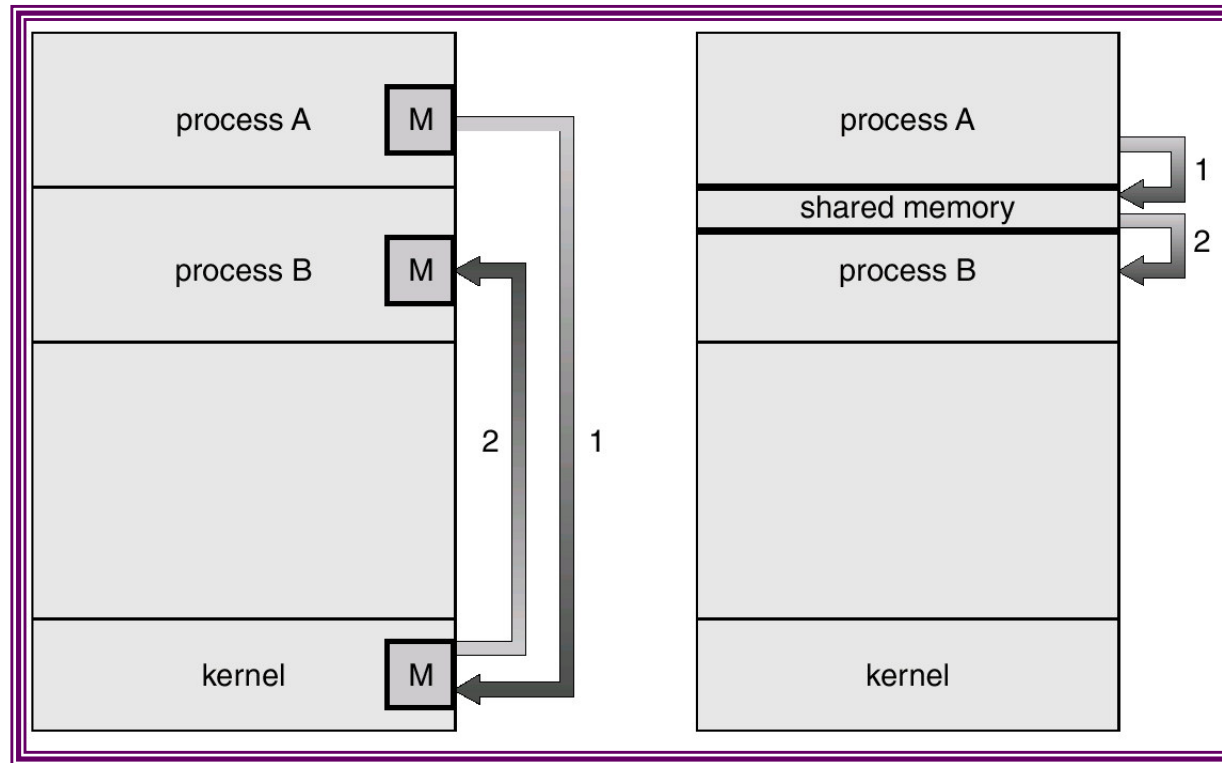
Running a Program

Process Management: UNIX



Communication Models

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



Message Passing

Shared Memory

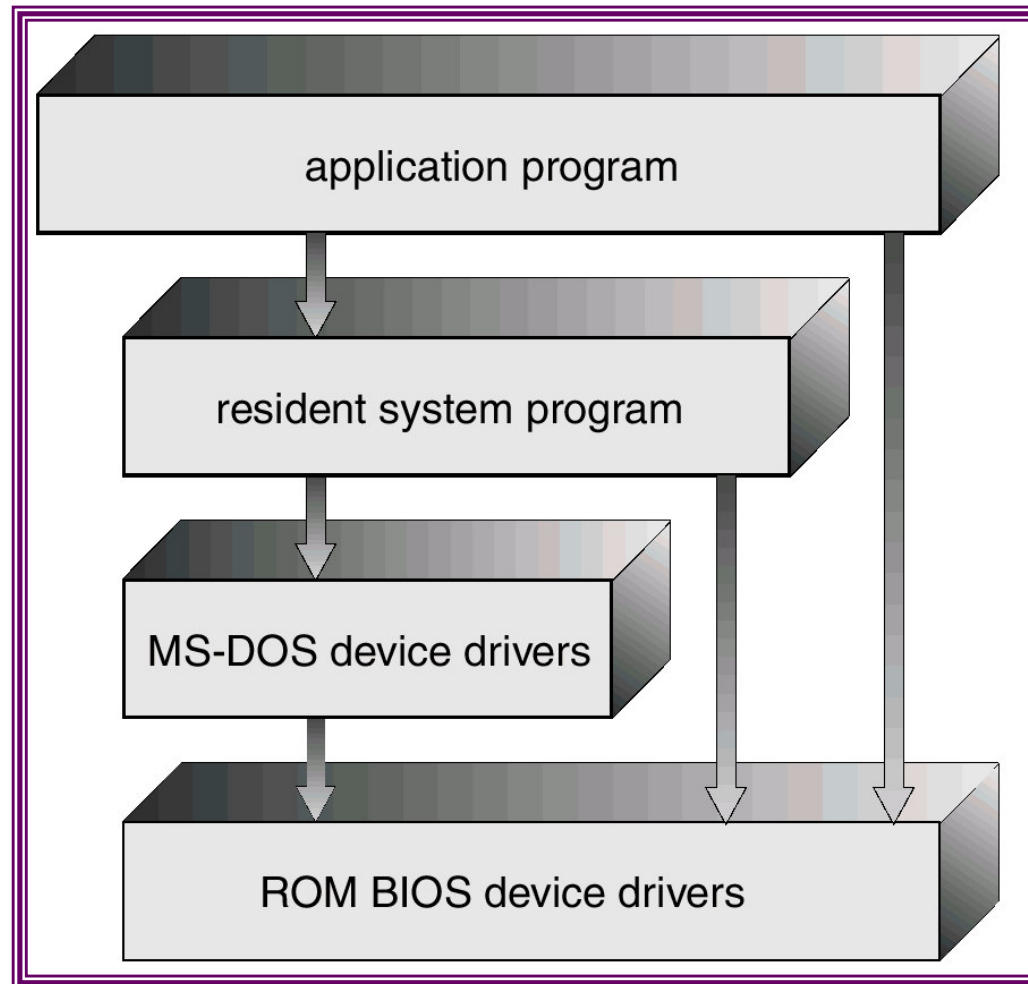
System Programs

- System programs provide a convenient environment for program development and execution. They can be divided into:
 - ☞ File manipulation: Unix 'cp'
 - ☞ Status information: Unix 'ps'
 - ☞ File modification: Unix 'cat'
 - ☞ Programming language support: Unix 'cc'
 - ☞ Program loading and execution: Windows 'start'
 - ☞ Communications: Windows 'ipconfig'
 - ☞ Application programs: Windows 'ie'
- Most users' view of the operating system is defined by system programs, not the actual system calls.

MS-DOS System Structure

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

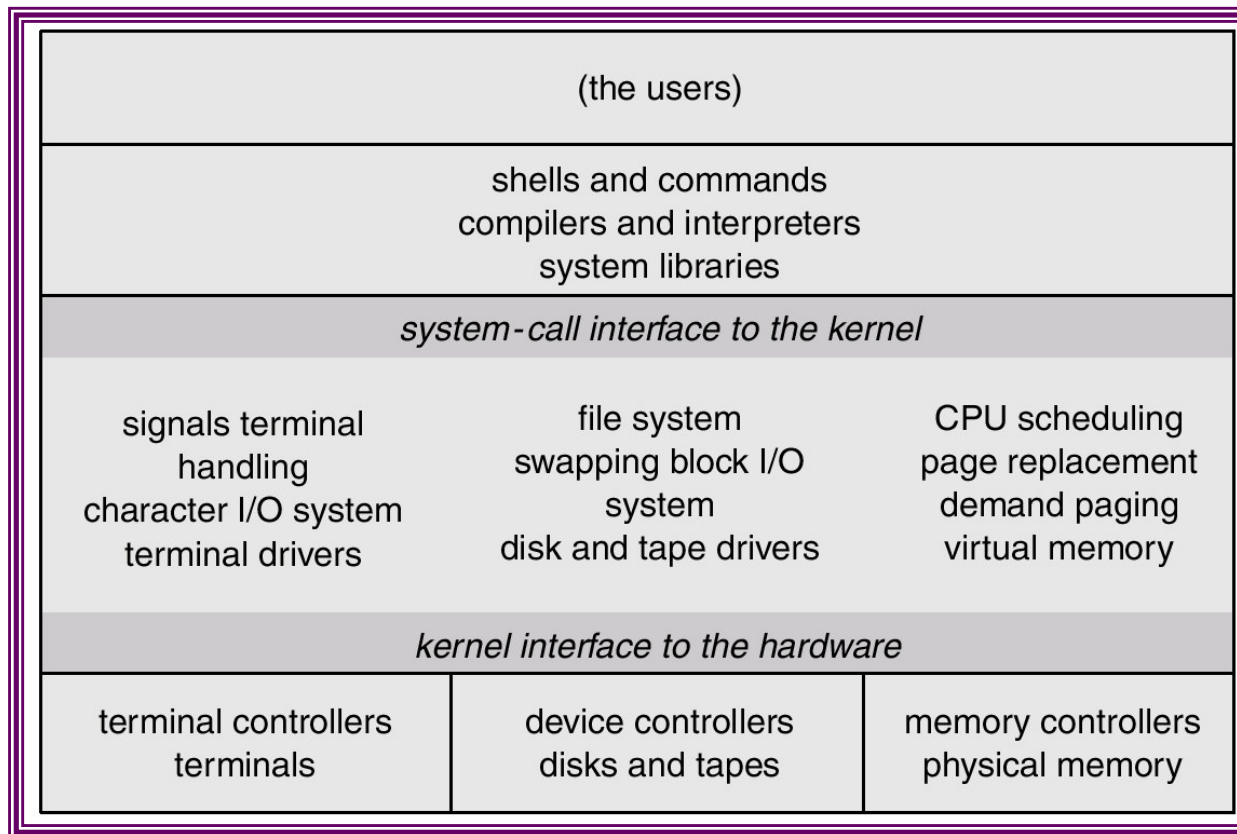
MS-DOS Layer Structure



UNIX System Structure

- UNIX – limited by hardware functionality, the original UNIX operating system had limited structuring. The UNIX OS consists of two separable parts.
 - ☞ Systems programs
 - ☞ 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.

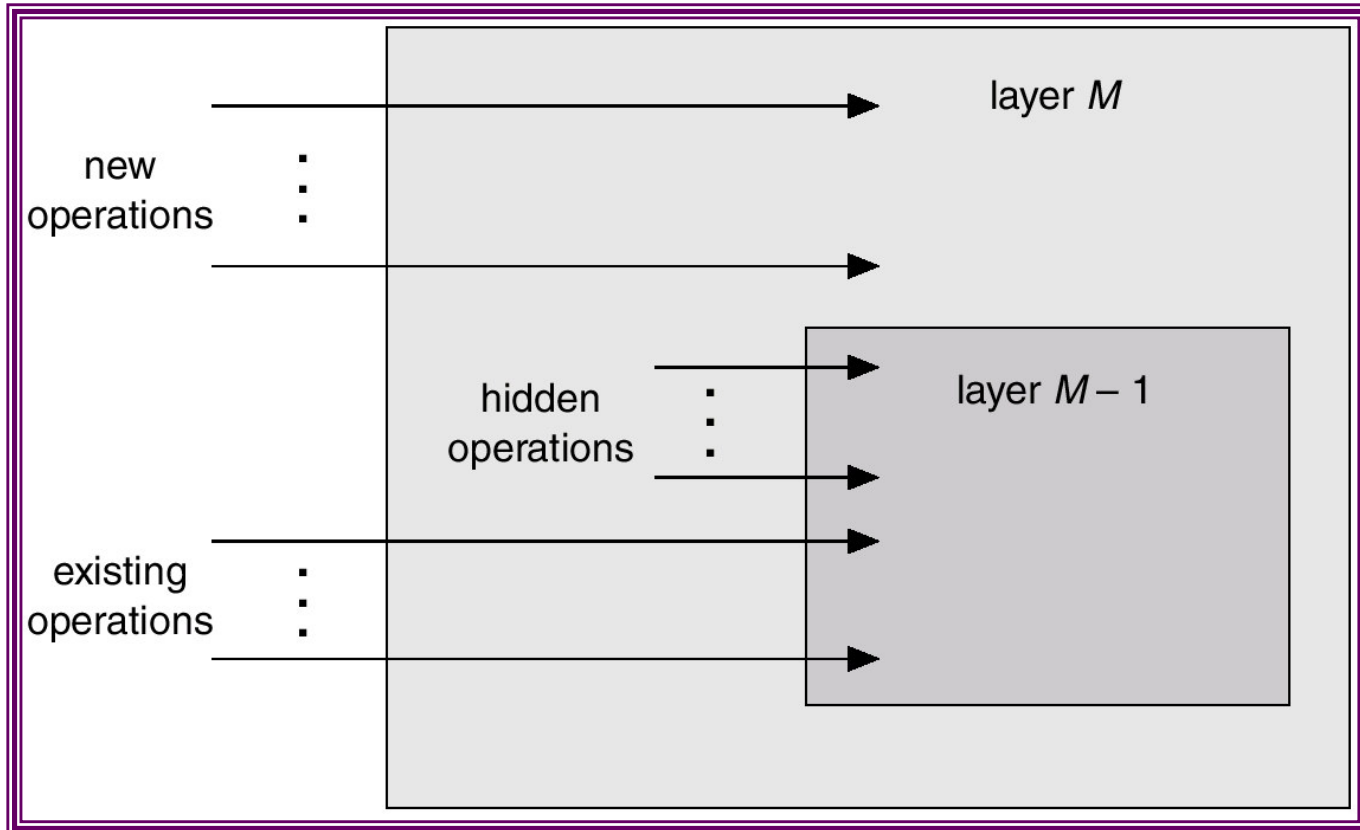
UNIX System Structure



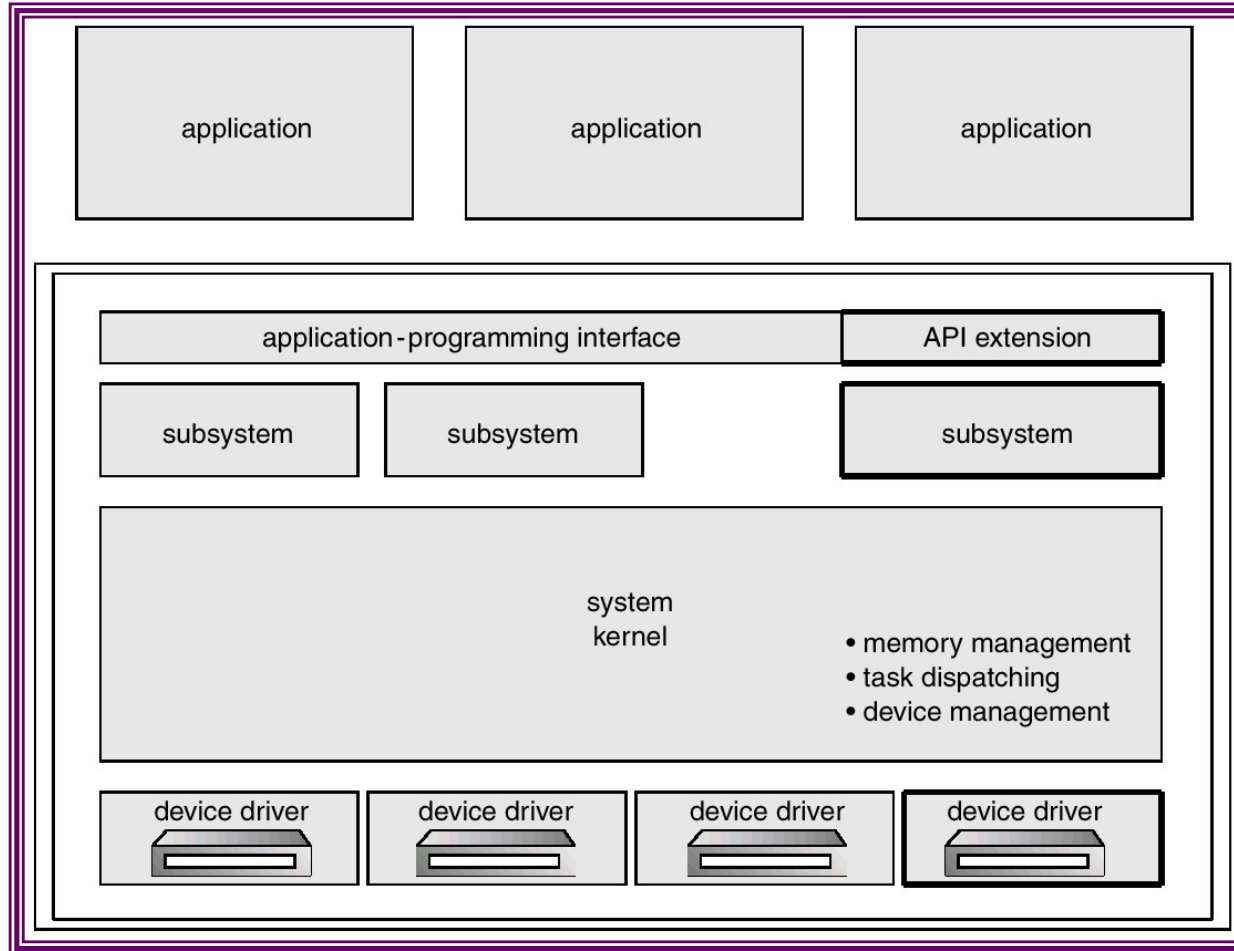
Layered Approach

- The operating system is divided into a number of layers (levels), each built on top of lower layers. The bottom layer (layer 0), is the hardware; the highest (layer N) is the user interface.
- 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

- Moves as much from the kernel into “*user*” space.
- Communication takes place between user modules using message passing.
- Benefits:
 - easier to extend a microkernel via user-space extensions.
 - easier to port the operating system to new architectures (less code to port)
 - more reliable (less code is running in kernel mode)
 - more secure (less to protect)
- Windows NT: Win32, POSIX subsystem are user-level.

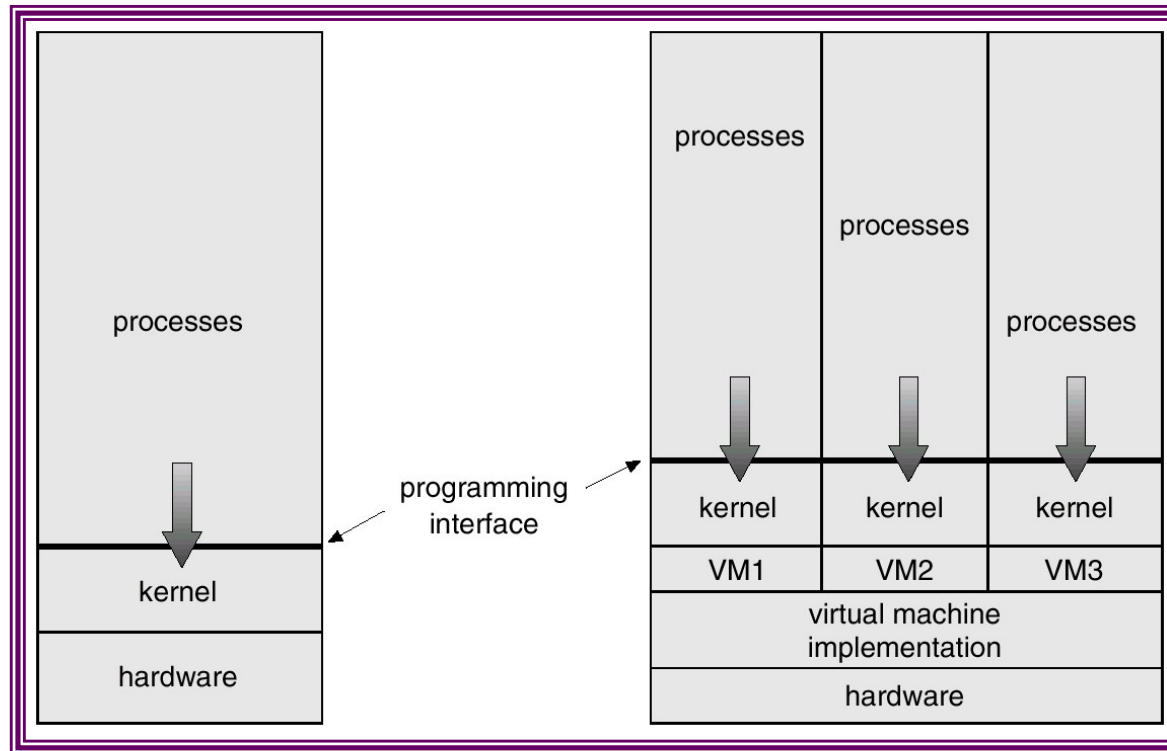
Virtual Machines

- A *virtual machine* is a program that acts as a hardware simulator. Run N copies of this simulator, the one physical machine becomes N virtual machines.
- Each machine can run:
 - ☞ a single process under a simple OS
 - ☞ all processes of a single user under a moderate OS
 - ☞ a complex time-sharing OS (e.g. for debugging)
- “OS” has three parts:
 - ☞ hardware simulator,
 - ☞ resource (processor, memory) sharing between simulators,
 - ☞ OS running inside each simulator.

Virtual Machines (Cont.)

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

System Models



Non-virtual Machine

Virtual Machine

Advantages of Virtual Machines

- Complete protection of system resources since each virtual machine is isolated from all other virtual machines.
- Ideal for operating-systems research and development. System development does not disrupt normal system operation.


Disadvantage of Virtual Machines

- No direct sharing of resources.
- Difficult to implement. For efficiency, we run non-privileged instruction on hardware. But then, what if...
 - ☞ User program under simulator makes system call in *real user* mode, triggers *real* interrupt
 - ☞ Real OS, in *real kernel* mode, sets simulator to *simulated kernel* mode, restarts simulator's implementation of system call in *real user* mode
 - ☞ Simulator runs privileged instruction (e.g. I/O) in *real user* mode, triggers *real* trap
 - ☞ Real OS simulates I/O in *real kernel* mode, restarts simulator in *simulated kernel* mode and *real user* mode.
- If this was hard to understand, imagine how hard it is to code it correctly...

Java Virtual Machine

- Compiled Java programs are platform-neutral bytecodes executed by a Java Virtual Machine (JVM).
- JVM consists of
 - class loader
 - class verifier (no pointers, no stack over/underflow)
 - runtime interpreter
- Just-In-Time (JIT) compilers increase performance

Java Virtual Machine



VMWare

- Commercial product
- Simulates a basic PC
- Persistent state of machine stored in two real files:
 - ☞ NVRAM (non-volatile RAM)
 - ☞ Disk contents (1 file per 1GB of simulated disk)
- Simulated disk when simulator runs
 - ☞ Discard changes (disk loaded in memory)
 - ☞ Keep changes (real file was modified)
 - ☞ Choose to commit/revert: database-like journal kept on disk, can issue commit or abort
- Different than SoftWindows: simulates Win32 API, not generic hardware allowing any OS.

System Design Goals

- User goals – operating system should be convenient to use, easy to learn, reliable, safe, and fast.
- System goals – operating system should be easy to design, implement, and maintain, as well as flexible, reliable, error-free, and efficient.

Mechanisms and Policies

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

System Implementation

- Traditionally written in assembly language, operating systems can now be written in higher-level languages.
- Code written in a high-level language:
 - ☞ can be written faster.
 - ☞ is more compact.
 - ☞ is easier to understand and debug.
- 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)

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