



## Chapter 2: OS Structures

---



## Operating System Services (1)

---

- ❑ **One set of operating-system services provides functions that are helpful to the user (or user processes):**
  - **User interface** - almost all operating systems have a user interface (UI).
    - ***Command-line interface*** (CLI): require a program to allow entering and editing of text commands.
    - ***Graphics user interface*** (GUI): a window system with a pointing device and a keyboard to enter commands.
    - ***Batch***: commands and directives are entered into files to be executed.
  - **Program execution** - the system must be able to load a program into memory and to run that program, end execution, either normally or abnormally.

## Operating System Services (2)

---

- **I/O operations** - a user program may require I/O.
  - For efficiency and protection, users cannot control I/O devices directly.
  - The operating system must provide a means to do I/O.
- **File-system manipulation** - user programs need to read/write/create/delete/search files and directories.
  - The operating system provides permission management to allow or deny access to files or directories.
- **Communications** – user processes may exchange information, on the same computer or between computers over a network.
  - Communications may be via *shared memory* or through *message passing*.
- **Error detection** – the operating system needs to be constantly aware of possible errors.
  - And fix errors generated from hardware (disk fail) or software (arithmetic error).
  - Debugging facilities can enhance the user's and programmer's abilities to efficiently use the system.

3

## Operating System Services (3)

---

- **For systems with multiple users (processes), another set of operating-system functions exists for ensuring the efficient operation of the system itself.**
- **Resource allocation** - when multiple users or multiple jobs running concurrently, resources must be allocated to each of them.
  - CPU, memory, file storage ...
  - Operating systems have *CPU-scheduling routines* to determine the best way to use the CPU.
- **Accounting** - To keep track of which users use how much and what kinds of computer resources.
  - Usage statistics may be a valuable tool for researchers who wish to reconfigure the system to improve computing services.

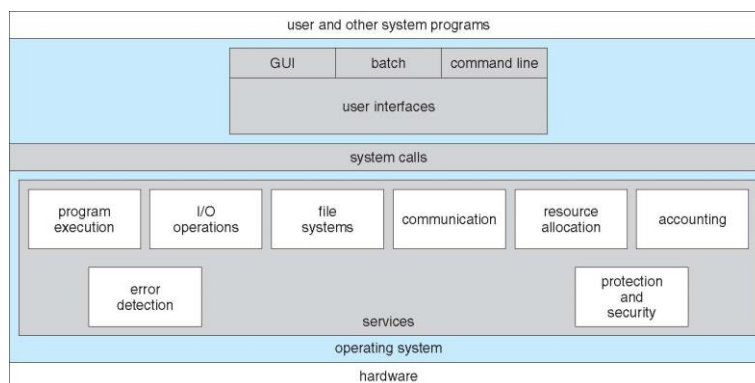
4

## Operating System Services (4)

- **Protection and security** - a multiuser or networked computer system may want to control use of user information.
  - **Concurrent** processes should not interfere with each other.
  - **Protection** involves ensuring that all access to system resources is controlled.
  - **Security** of the system from outsiders requires user authentication, extends to defending external I/O devices (e.g., network adapters) from invalid access attempts.

5

## A View of Operating System Services



## User Operating-System Interface – CLI

---

### □ Two ways to implement commands and command interpreters:

- The command interpreter contains the code to execute the command.
  - For example, a command to delete a file may cause the interpreter to jump to a section of its code that makes the appropriate system call.
  - The number of commands that can be given determines the size of the interpreter.
- An alternative approach implements **most** commands through system programs.
  - The interpreter does not understand the command.
  - It **identify** the command file and **load** it into memory for **execution**.

7

## CLI (cont'd)

---

### □ UNIX example: `rm file.txt`

- The interpreter search for a file called `rm (/bin/rm)`.
- **Load** it into memory and **execute** it with the parameter `file.txt`.
- The function `rm` is completely defined by the code in the file `/bin/rm`.
- Programmers can add new commands to the system easily.
- The interpreter program can be **small**, and does not have to be changed for new commands to be added.
- Used mostly among operating system, e.g., UNIX.

8

## CLI (cont'd)

- ❑ **CLI is sometimes implemented in kernel, sometimes by systems program.**
- ❑ **An operating system can have multiple interpreters to choose from, known as *shells*.**
  - For example, on UNIX and Linux systems, there are *Bourne/C/Korn ...shell*.
  - The name 'shell' originates from shells being an outer layer of interface between the user and the innards of the operating system (the kernel).
  - Most shells provide similar functionality with only minor differences; most users choose a shell based upon personal preference.
    - E.g., the syntax of *shell script*.

9

## GUI

- ❑ **A GUI provides a *desktop metaphor interface*.**
  - **Icons** represent files, programs, actions, etc.
  - A **mouse click** can invoke a program, select a file ...
- ❑ **GUI Timeline:**
  - Experimentally appeared in the early 1970s.
  - Became widespread by Apple Macintosh computer (Mac OS) in the 1980s.
  - Dominated by Microsoft Windows (3.1, NT, 95, 98, 2000, XP, Vista).
- ❑ **UNIX systems have been dominated by command-line interface.**
  - Although there are various GUI interface available.
    - X-Windows systems, K Desktop Environment (KDE) by GNU project (open source – source code is in the public domain).
- ❑ **Many systems now include both CLI and GUI interfaces**
  - Microsoft Windows is GUI with CLI “command” shell
  - Apple Mac OS X is “Aqua” GUI interface with UNIX kernel underneath and shells available
  - Unix and Linux have CLI with optional GUI interfaces (CDE, KDE, GNOME)

10

## GUI (cont'd)

- ❑ **Preference:**
  - Many UNIX users often prefer a command-line interface.
  - Most Windows user are pleased to use the Windows GUI and seldom use the MS-DOS shell interface.
- ❑ **Nevertheless, many systems now include both CLI and GUI interfaces.**

11

## Touchscreen Interfaces

- ❑ **Touchscreen devices require new interfaces.**
  - Mouse not possible or not desired
  - Actions and selection based on gestures
  - Virtual keyboard for text entry
- ❑ **Voice commands.**



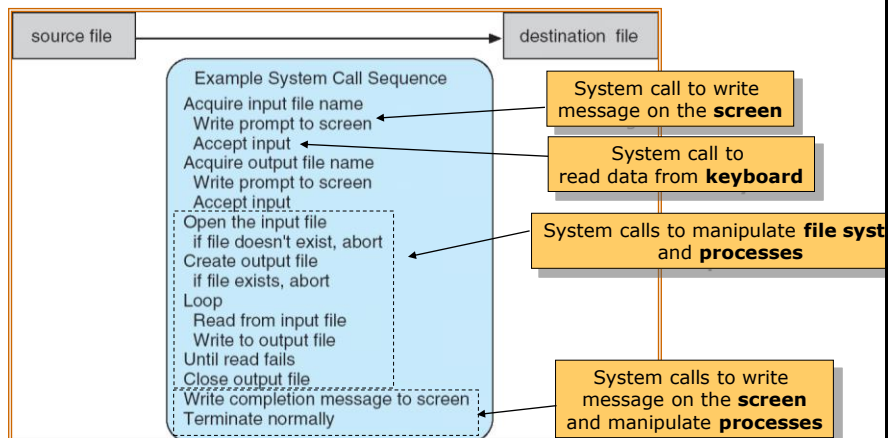
## System Calls

- ❑ Can be regarded as a ***programming interface to the services provided by the OS.***
  - Called by user applications.
- ❑ Are generally available as  ***routines,*** typically written in a high-level language (C or C++).
  - Also in low-level assembly language (for accessing hardware).

13

## System Calls (cont'd)

- ❑ System call sequence to copy the contents of one file to another file.



**Even a simple program makes heavy use of system calls!!**

## System Calls (cont'd)

- ❑ Programs mostly access system services via a high-level *application program interface* (API) rather than using system call directly.
  - API specifies a set of functions (specifications) that are available to programmers.
  - **The functions of the API invoke the actual system calls on behalf of the programmer.**
- ❑ Three most common APIs:
  - **Win32 (Win) API** for Windows.
  - **POSIX API** for POSIX-based systems (including virtually all versions of UNIX, Linux, and Mac OS X).
  - **Java API** for the Java virtual machine (JVM).

Portable Operating System Interface

15

## System Calls (cont'd)

- ❑ Consider the ReadFile() function in the
- ❑ Win API—a function for reading from a file

```

return value
  ↓
BOOL ReadFile c (HANDLE file,
                LPVOID buffer,
                DWORD bytes To Read,
                LPDWORD bytes Read,
                LPOVERLAPPED ovl);
                ↑
                function name
    
```

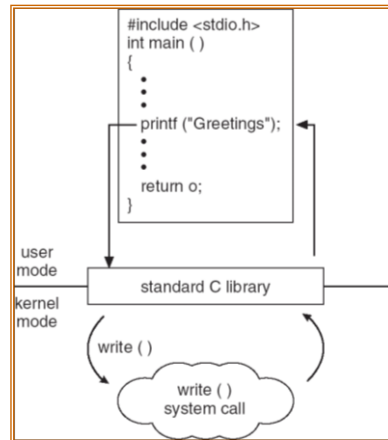
] parameters

- ❑ A description of the parameters passed to ReadFile()
  - HANDLE file—the file to be read
  - LPVOID buffer—a buffer where the data will be read into and written from
  - DWORD bytesToRead—the number of bytes to be read into the buffer
  - LPDWORD bytesRead—the number of bytes read during the last read
  - LPOVERLAPPED ovl—indicates if overlapped I/O is being used



## System Calls (cont'd)

- C program invoking `printf()` library call, which calls `write()` system call.



17

## System Calls (cont'd)

### EXAMPLE OF STANDARD API

As an example of a standard API, consider the `read()` function that is available in UNIX and Linux systems. The API for this function is obtained from the `man` page by invoking the command

```
man read
```

on the command line. A description of this API appears below:

```
#include <unistd.h>
ssize_t read(int fd, void *buf, size_t count)
```

return value	function name	parameters
--------------	---------------	------------

A program that uses the `read()` function must include the `unistd.h` header file, as this file defines the `ssize_t` and `size_t` data types (among other things). The parameters passed to `read()` are as follows:

- `int fd`—the file descriptor to be read
- `void *buf`—a buffer where the data will be read into
- `size_t count`—the maximum number of bytes to be read into the buffer

On a successful read, the number of bytes read is returned. A return value of 0 indicates end of file. If an error occurs, `read()` returns `-1`.

## System Calls (cont'd)


- ❑ **Why use APIs rather than system calls?**
  - **Program portability** — a program using an API can be expected to compile and run on any system that supports the same API.
    - E.g., your Windows programs on various versions of Windows.
  - Programming with API is **simpler** than using system calls.

19

## System Calls (cont'd)

- ❑ **As system calls are routines in kernel space, using it causes a change in privileges.**
- ❑ **But before that ...**
  - Similar to hardware interrupt, we need a number (index) to indicate the required system call, which is store in the `EAX` register.
- ❑ **How ?**
  - Via software interrupt (e.g., `INT 0x80` assembly instruction on Intel 386 arch of Linux system).
- ❑ **System contains a table of code pointers.**
  - Using the system call number, we jump to the address of the system call for execution.

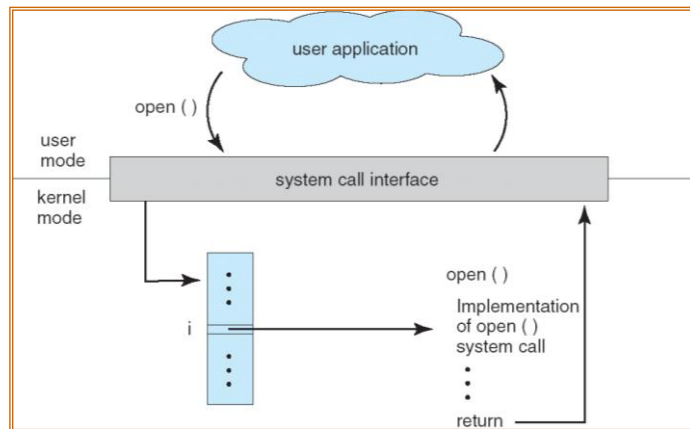
API helps us wrap all the details by simply invoking a library function.



20

## System Calls (cont'd)

- How the operating system handles a user application invoking the `open ()` system call through API.



21

## System Calls (cont'd)

- To link system call made available by the operating system:
  - The **run-time library** for most programming languages provides a **system-call interface**.
  - Typically, a **number** associated with each system call.
  - System-call interface (or kernel) maintains a **table** indexed according to these numbers.
  - The system call interface invokes intended system call in operating system kernel and returns status of the system call and any return values.

22

## System Calls (cont'd)

---

- ❑ **With the help of API:**
  - The caller need know nothing about how the system call is implemented.
  - Just needs to obey API and understand what the operating system will do as a result of that system call.
  - Details of the operating system are hidden from programmer.

23

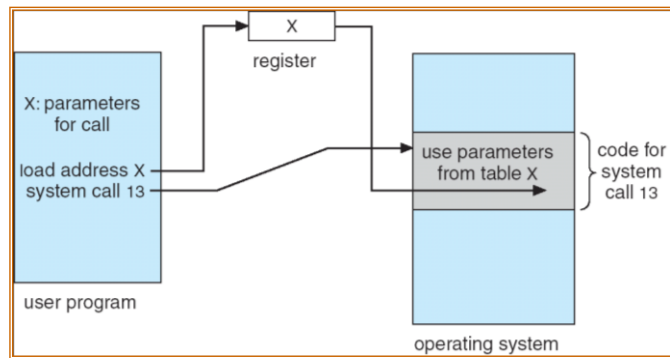
## System Calls (cont'd)

---

- ❑ **Often, more information is required than simply identity of desired system call**
  - E.g., system call parameters.
- ❑ **Three general methods used to pass parameters to the OS:**
  - Simplest: pass the parameters in *registers*
    - In some cases, may be more parameters than registers.
  - Parameters are stored in a *block* in memory, and address of block passed as a parameter in a register.
    - This approach taken by Linux and Solaris.
  - Parameters are *pushed* onto a ***stack*** by the program and *popped* off the stack by the operating system.
- Block and stack methods are popular because they do not limit the number or length of parameters being passed.

24

## System Calls (cont'd)



25

## Types of System Calls

- ❑ **Many of today's operating system have hundreds of system calls.**
  - Linux/ has more than 328 different system calls.
    - [http://blog.rchapman.org/posts/Linux\\_System\\_Call\\_Table\\_for\\_x86\\_64/](http://blog.rchapman.org/posts/Linux_System_Call_Table_for_x86_64/)
- ❑ **Those system calls can be grouped roughly into five major categories:**
  - Process control.
  - File management.
  - Device management.
  - Information maintenance.
  - Communications.
  - Protection

26

## Examples of Windows and Unix System Calls

	Windows	Unix
<b>Process Control</b>	CreateProcess() ExitProcess() WaitForSingleObject()	fork() exit() wait()
<b>File Manipulation</b>	CreateFile() ReadFile() WriteFile() CloseHandle()	open() read() write() close()
<b>Device Manipulation</b>	SetConsoleMode() ReadConsole() WriteConsole()	ioctl() read() write()
<b>Information Maintenance</b>	GetCurrentProcessID() SetTimer() Sleep()	getpid() alarm() sleep()
<b>Communication</b>	CreatePipe() CreateFileMapping() MapViewOfFile()	pipe() shmget() mmap()
<b>Protection</b>	SetFileSecurity() InitializeSecurityDescriptor() SetSecurityDescriptorGroup()	chmod() umask() chown()

## Process Control

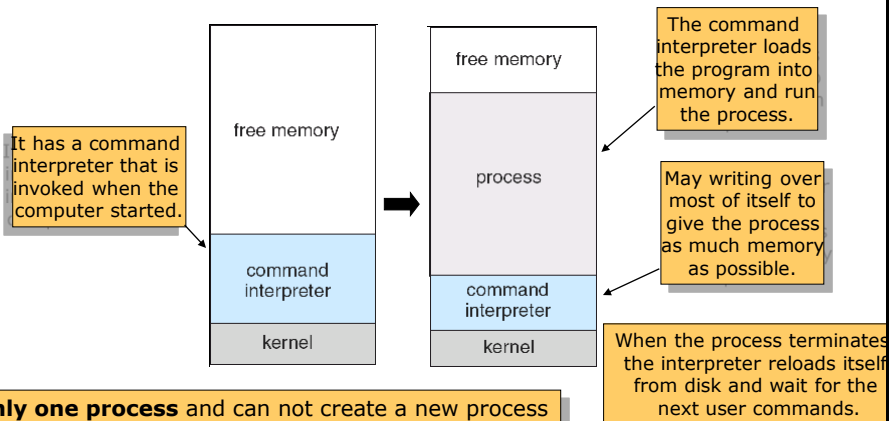
- ❑ **end and abort:**
  - A running program needs to be able to halt its execution either normally or abnormally.
  - The operating system then transfers control to the invoking command interpreter to read the next command.
- ❑ **load and execute:**
  - A process executing one program may want to load and execute another program.
  - The existing process can be lost, saved, or allowed to continue execution concurrently with the new process.
  - Chapter 6 (synchronization) discusses coordination of concurrent processes in great detail.
- ❑ **wait time/event:**
  - Having created new processes, we may need to wait for them to finish their execution.
  - We may want to wait for a certain amount of time to pass.
  - We may want to wait for a specific event to occur.

28

## Process Control (cont'd)

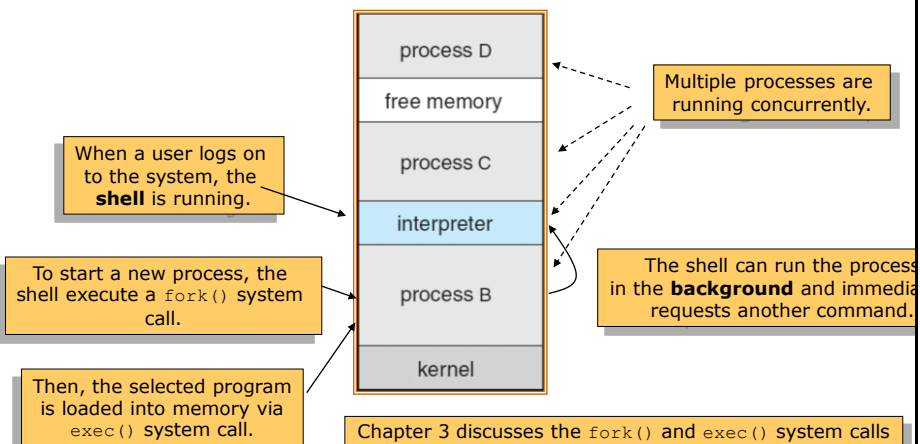
❑ **Two popular variations in process control:**

- **Single-tasking system:** the MS-DOS operating system.



## Process Control (cont'd)

- **A multitasking system:** the FreeBSD (derived from Berkeley UNIX).



## File Management

---

- ❑ **create and delete:**
  - Able to create and delete files/directories.
- ❑ **open and close:**
  - Able to open and close a file.
- ❑ **read, write, and reposition:**
  - Able to read, write, or skipping to the end/head of the file.
- ❑ **Other system calls for obtaining/setting file/directory attributes.**

31

## Device Management

---

- ❑ **The various resources (memory, disks, file, ...) controlled by the operating system can be thought of as devices.**
- ❑ **To access a resource, a process has to:**
  - First `request` the device, to ensure **exclusive** use of it.
  - Then we can `read`, `write`, and `reposition` the device.
  - After we are finished with the device, we `release` it.
- The similarity between I/O devices and files is so great that many operating systems (UNIX) merge the two into a combined file-device structure.
  - A set of system calls is used on files and devices.
  - Sometimes, I/O devices are identified by special file names.

32



## Information Maintenance

---

- ❑ **Many system calls exist simply for the purpose of transferring information between the user program and the operating system.**
  - `time` and `date` return the current time and date of the system.
  - Other system calls can return the number of current users, the amount of free memory or disk space, ...
  - Get and set processes attributes.

33

## Communication

---

- ❑ **There are two common models of interprocess communication:**
  - **Message-passing model:**
    - The communicating processes (may be on different computers) exchange messages with one another to transfer information.
    - Client and server (daemon) architecture.
      - Client: ask for connecting communication.
      - Server: wait for connection.
    - Require system calls to `build up/terminate` connection, `read`, and `write` messages.
  - **Shared-memory model:**
    - Processes use `shared memory` `create/attach` system calls to create and gain access to regions of memory owned by other processes.
    - They can then exchange information by reading and writing data in the shared areas.

34

## Communication (cont'd)

---

- ❑ **Message passing:**
  - Is useful for exchanging smaller amounts of data, because no conflicts need be avoided.
  - Is easy to implement.
- ❑ **Shared memory:**
  - Allows maximum speed of communication, since it can be done at memory speeds when it takes place within a computer.
  - However, problems exist in the areas of protection and synchronization between the processes sharing memory.

35

## System Programs

---

- ❑ **A perspective of operating systems is a collection of system programs.**
  - System programs provide a convenient environment for program execution and development.
  - Most users' view of the operation system is defined by system programs, not the actual system calls.
  - **Some of them are just user interfaces to system calls!!**
- ❑ **Categories of system programs:**
  - **File manipulation:**
    - Create, delete, copy, rename, print, dump, list, and generally manipulate files and directories

36

## System Programs (cont'd)

---

- **File modification:**
  - Text editors to create and modify files.
- **Status information:**
  - Some ask the system for info - date, time, amount of available memory, disk space, number of users
  - Typically, these programs format and print the output to the terminal or other output devices
  - Some systems implement a **registry** - used to store and retrieve configuration information
- **Programming-language support:**
  - Compilers, assemblers, debuggers and interpreters sometimes provided.
- **Program loading and execution:**
  - Loaders to load assembled or compiled programs into memory for execution.

37

## System Programs (cont'd)

---

- **Communications:**
  - Provide the mechanism for creating virtual connections among processes, users, and computer systems.
- **Background Services**
- ❑ **In addition to system programs, application programs are supplied to solve common problems or perform common operations.**
  - Web browsers, word processors, database systems, games ...
- ❑ **The view of the operating system seen by most users is defined by the application and system programs, rather than the actual system calls.**

38

## System Programs (Cont.)

---

- ❑ **Background Services**
  - Launch at boot time
    - Some for system startup, then terminate
    - Some from system boot to shutdown
  - Provide facilities like disk checking, process scheduling, error logging, printing
  - Run in user context not kernel context
  - Known as **services, subsystems, daemons**
  
- ❑ **Application programs**
  - Don't pertain to system
  - Run by users
  - Not typically considered part of OS
  - Launched by command line, mouse click, finger poke

## Operating System Design

---

- ❑ **Design and Implementation of OS not “solvable”, but some approaches have proven successful**
- ❑ **Internal structure of different Operating Systems can vary widely**
- ❑ **Start the design by defining goals and specifications**
- ❑ **Affected by choice of hardware, type of system**
  - There is no unique solution to the problem of defining the requirements for an operating system.
  - Requirements can be affected by choice of hardware, type of system.
    - Handheld devices vs. PCs.
    - Single process vs. multitasking.
- ❑ **The requirement can be divided into user goals and system 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.

40

## Operating System Design (cont'd)

- ❑ One important principle of system design is the separation of *policy* from *mechanism*.
  - **Policy: What** will be done?
  - **Mechanism: How** to do it?
  - Example: *timer* is a **mechanism** for ensuring CPU protection, but deciding how long the timer is to be set is a **policy** decision.
- ❑ **Flexibility of the separation:**
  - Policies are likely to change across places or over time.
  - The separation enables a change in policy to redefine certain policy parameters rather than changing the underlying mechanism (example – timer).
- ❑ **Most of Windows services mix mechanisms with policies to enforce a global look and feel.**
- ❑ **Specifying and designing an OS is highly creative task of software engineering**

41

## Operating System Implementation

- ❑ Traditionally, operating systems have been written in assembly language, then Algol, PL/1.
- ❑ Now, they are most written in higher-level languages such as C or C++.
  - The code can be written faster and is easier to understand and debug.
  - System is easier to port (to move to some other hardware).
    - The Linux operating system is written mostly in C and is available on a number of different CPUs.
    - Modern C compiler techniques can perform complex analysis and optimizations that produce excellent code.
- ❑ **Actually usually a mix of languages**
  - Lowest levels in assembly
  - Main body in C
  - Systems programs in C, C++, scripting languages like PERL, Python, shell scripts
- ❑ **More high-level language easier to port to other hardware**
  - But slower
- ❑ **Emulation can allow an OS to run on non-native hardware**

42

## Operating System Implementation (cont'd)

---

- ❑ Moreover, major performance improvements in operating systems (and other systems) are more likely to be the result of better data structures and algorithms than of excellent assembly-language code.
- ❑ Should pay more attentions on the memory manager and the CPU scheduler.
  - They are probably the most critical routines.

43

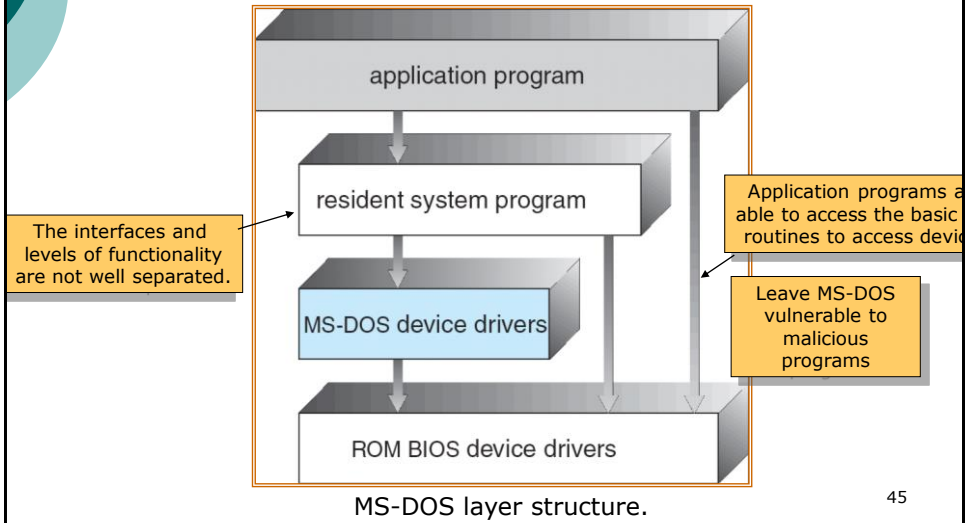
## OS Structure: Simple Structure

---

- ❑ A common approach to implement an operating system is to partition the task into small components.
  - **Rather than have one monolithic system!!**
  - These components are interconnected and meld into a kernel.
  - But...
- ❑ **Simple structure:**
  - Many commercial system do not have well-defined structures initially.
    - Started as small, simple, and limited systems and then grew beyond their original scope.
    - For example, MS-DOS.

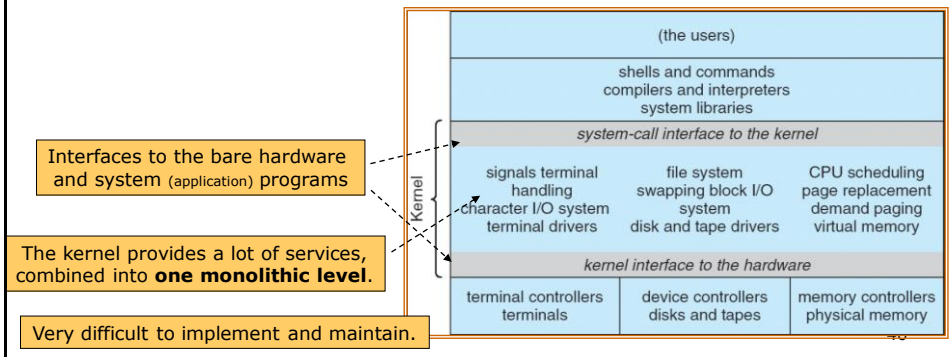
44

## OS Structure: Simple Structure (cont'd)



## OS Structure: Non Simple Structure

- ❑ **UNIX – limited by hardware functionality, the original UNIX operating system had limited structuring.**
- ❑ **The UNIX OS consists of two separable parts:**
  - **System programs.**
  - **The kernel.**



## OS Structure: Layered Approach


- ❑ **With the improvements of hardware and programming techniques, operating systems can be broken into pieces of components.**
  - That is **modular** operating systems.
    - Information hiding: hide the internal implementation detail of modules and provide external access **interfaces**.
- ❑ **One way of modular system: 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.

47

## OS Structure: Layered Approach (cont'd)

Layer  $M$  consists of data structures and a set of routines that can be invoked by higher-level layers.

Layer  $M$ , in turn, can invoke operations on lower-level layers.



48



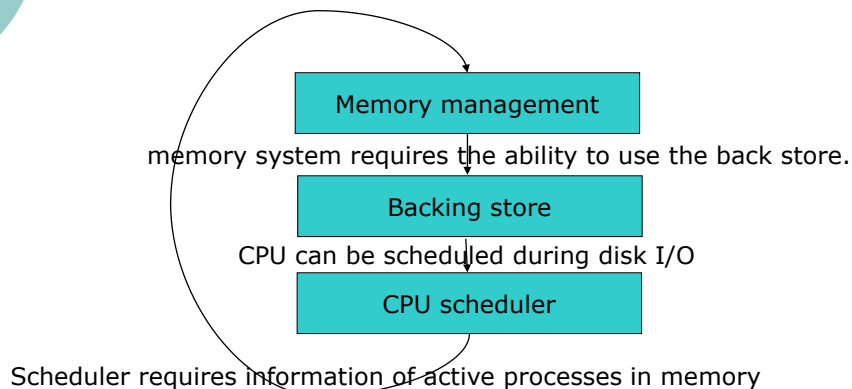
## OS Structure: Layered Approach (cont'd)

- ❑ **The main advantage of the layered approach:**
  - Simplicity of construction and **debugging**.
  - Layer-by-layer debugging, starting from layer 0.
  - If an error is found during the debugging of a particular layer, the error must be on that layer.
- ❑ **The major difficulty of the layered approach:**
  - Because only lower layers operations can be invoked, appropriately defining the various layers is difficult.
    - **System services usually tangle together.**
  - Layered implementation tend to be less efficient.
    - A function call on the top layer can lead to many lower-layer calls.
    - Function calls need to pass (redundant) parameters.
- ❑ **Recently, fewer layers with more functionality are being designed.**
  - Providing the advantages of modularization.
  - Avoiding the difficulties of layer definition and interaction.

49

## OS Structure: Layered Approach (cont'd)

- ❑ **Example of tangled layers:**

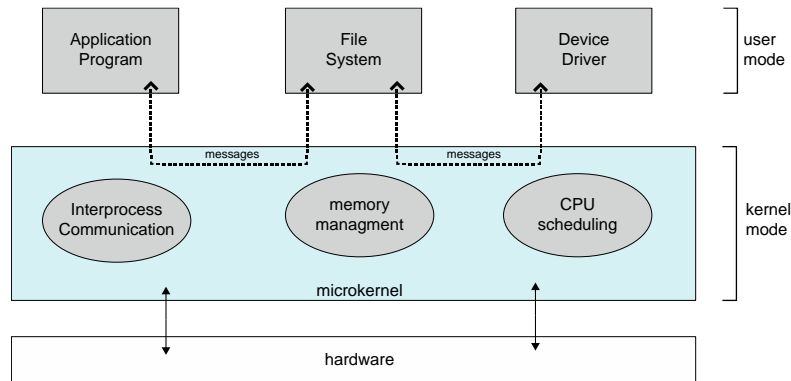


50

## Microkernel System Structure (1)

- ❑ Moves as much from the kernel into user space.
- ❑ Mach example of microkernel.
  - Mac OS X kernel (Darwin) partly based on Mach
- ❑ Communication takes place between user modules using message passing.
- ❑ Benefits:
  - Easier to extend a microkernel
  - Easier to port the operating system to new architectures
  - More reliable (less code is running in kernel mode)
  - More secure
- ❑ Detriments:
  - Performance overhead of user space to kernel space communication

## Microkernel System Structure (2)



## OS Structure: Microkernel

---

- ❑ As operating systems expanded, the kernel became large and difficult to manage.
- ❑ In the mid-1980s, CMU developed an operating system called Mach that modularized the kernel using the microkernel approach.
  - Micro → removing all nonessential components from the kernel and implementing them as system and user-level programs (servers).
  - Typically, microkernels provide **process** and **memory** management, and a **communication** facility.
  - The client program and services communicate indirectly by exchanging message with the microkernel.

53

## OS Structure: Microkernel (cont'd)

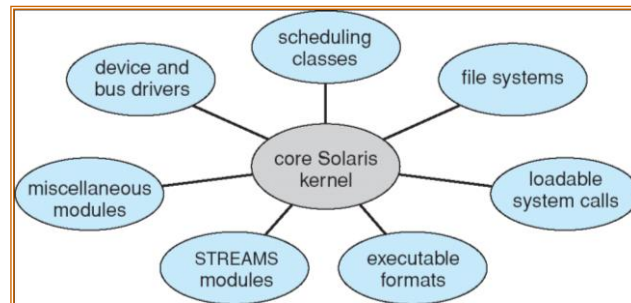
---

- ❑ **Benefits:**
  - Easier to include new operating system services to a microkernel.
    - Do not require modification of the kernel.
  - The small kernel makes it easier to port to new hardware architectures.
  - More reliable and secure (less code is running in kernel mode).
- ❑ **Problems:**
  - Performance overhead of user space to kernel space communication.
  - Initial Windows NT (a micorkernel organization) → Windows NT 4.0 (moving layers from user space to kernel space).

54

## OS Structure: Module

- ❑ A better methodology for operating-system design involves using object-oriented programming techniques to create a modular kernel.
  - Consists of a **core kernel**, and system service as **kernel modules (Linux, Solaris, etc.)**.
    - Each module talks to the others over known **interfaces**



55

## OS Structure: Module (cont'd)

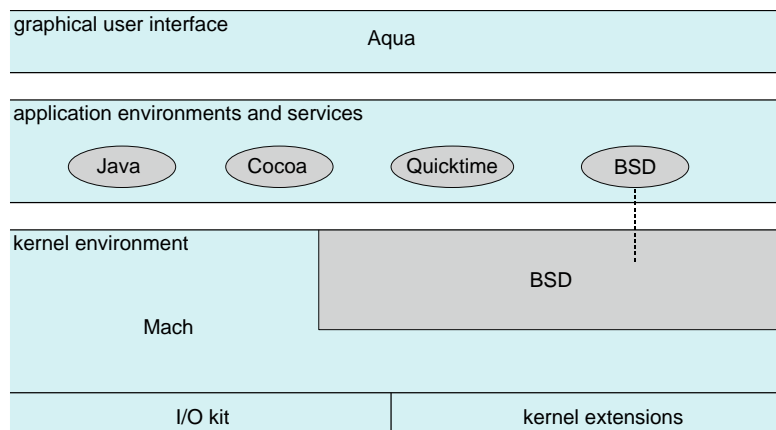
- ❑ Moreover, modules (system services) can be linked into the system either during boot time or during run time (that is, loaded as needed within the kernel).
  - Load different file system (ext2/3/4fs, FAT32 or NTFS) as needed, to save main memory.
- ❑ The module structure is similar to layered (communicate with interfaces) and microkernel approaches (a core), but with more flexible.
  - Any module can call any other module, but the layered approach can not.
  - Is efficient than microkernel approach because modules are in the kernel space and do not need to invoke message passing to communicate.
- ❑ The strategy of dynamically loadable modules is very popular in modern UNIX-based operating systems, such as Linux.

56

## Hybrid Systems

- ❑ **Most modern operating systems are actually not one pure model.**
  - Hybrid combines multiple approaches to address performance, security, usability needs
  - Linux and Solaris kernels in kernel address space, so monolithic, plus modular for dynamic loading of functionality
  - Windows mostly monolithic, plus microkernel for different subsystem *personalities*
- ❑ **Apple Mac OS X hybrid, layered, Aqua UI plus Cocoa programming environment.**
  - Below is kernel consisting of Mach microkernel and BSD Unix parts, plus I/O kit and dynamically loadable modules (called **kernel extensions**)

## Mac OS X Structure



# iOS

- ❑ **Apple mobile OS for iPhone, iPad**
  - Structured on Mac OS X, added functionality
  - Does not run OS X applications natively
    - Also runs on different CPU architecture (ARM vs. Intel)
  - **Cocoa Touch** Objective-C API for developing apps
  - **Media services** layer for graphics, audio, video
  - **Core services** provides cloud computing, databases
  - Core operating system, based on Mac OS X kernel

Cocoa Touch

Media Services

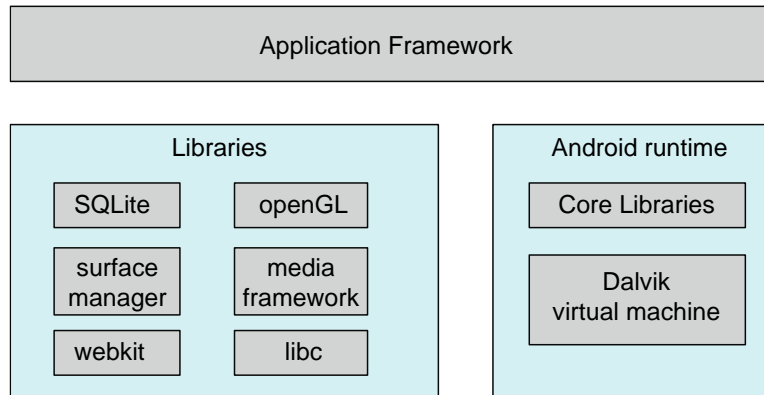
Core Services

Core OS

# Android

- ❑ **Developed by Open Handset Alliance (mostly Google).**
  - Open Source
- ❑ **Similar stack to IOS.**
- ❑ **Based on Linux kernel but modified.**
  - Provides process, memory, device-driver management
  - Adds power management
- ❑ **Runtime environment includes core set of libraries and Dalvik virtual machine.**
  - Apps developed in Java plus Android API
    - Java class files compiled to Java bytecode then translated to executable than runs in Dalvik VM
- ❑ **Libraries include frameworks for web browser (webkit), database (SQLite), multimedia, smaller libc.**


# Android Architecture



# Virtual Machines

□ **Virtual machine software** (such as Vmware, Virtual Box) **abstracts the hardware of a single computer into several different execution environments (virtual machines).**

- Creating the illusion that each virtual machine is running its own private computer.
- Being able to share the same hardware, yet run several different operating systems **concurrently**.
- Is difficult to implement due to the effort required to provide an *exact* duplicate to the underlying machine.
  - E.g., simulate dual mode and disk system.



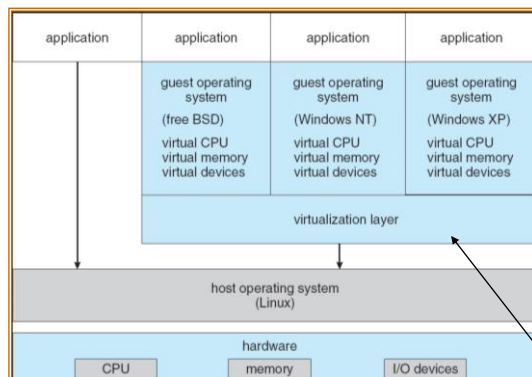
## Virtual Machines (cont'd)

### Why virtual machines?

- Today, virtual machines are frequently used as a means of solving system **compatibility** problems.
  - E.g., Java virtual machine (JVM).
- A virtual-machine system is a perfect vehicle for **operating-systems research and development**.
  - System development is done on the virtual machine, instead of on a physical machine and so does not disrupt normal system operation or crash the whole system.
- **Save money and time** for application development.
  - Virtual machines can help system developers develop an application on different operating system.

63

## VM Example – VMware



To concurrently run several different **guest operating systems** as virtual machines


VMware runs as an application on a host operating system

64



## VM Example – JVM

- ❑ JVM for system compatibility.
- ❑ To make java programs compatible to different systems ...
  - The compiler produces an **architecture-neutral** bytecode output (.class) file that will run on any JVM.
- ❑ JVM is a software on a host operating system.
- ❑ It abstracts computer and manages memory (garbage collection) to increase the performance of Java programs.



65

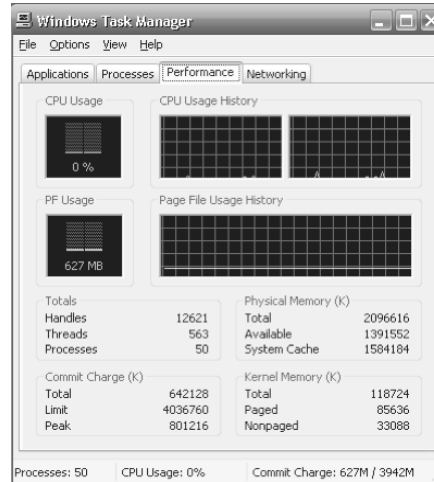
## Operating-System Debugging

- ❑ Debugging is finding and fixing errors, or bugs.
- ❑ OS generate log files containing error information.
- ❑ Failure of an application can generate core dump file capturing memory of the process.
- ❑ Operating system failure can generate crash dump file containing kernel memory.
- ❑ Beyond crashes, performance tuning can optimize system performance.
  - Sometimes using **trace listings** of activities, recorded for analysis
  - **Profiling** is periodic sampling of instruction pointer to look for statistical trends

**Kernighan's Law:** "Debugging is twice as hard as writing the code in the first place. Therefore, if you write the code as cleverly as possible, you are, by definition, not smart enough to debug it."

# Performance Tuning

- ❑ Improve performance by removing bottlenecks.
- ❑ OS must provide means of computing and displaying measures of system behavior.
- ❑ For example, “top” program or Windows Task Manager.



# DTrace

- DTrace tool in Solaris, FreeBSD, Mac OS X allows live instrumentation on production systems.
- **Probes** fire when code is executed within a **provider**, capturing state data and sending it to **consumers** of those probes .
- Example of following XEventsQueued system call move from libc library to kernel and back.

```
# ./all.d 'pgrep xclock' XEventsQueued
dtrace: script './all.d' matched 52377 probes
CPU FUNCTION
0 -> XEventsQueued U
0 -> _XEventsQueued U
0 -> _XllTransBytesReadable U
0 <- _XllTransBytesReadable U
0 -> _XllTransSocketBytesReadable U
0 <- _XllTransSocketBytesreadable U
0 -> ioctl U
0 -> ioctl K
0 -> getf K
0 -> set_active_fd K
0 <- set_active_fd K
0 <- getf K
0 -> get_udatamodel K
0 <- get_udatamodel K
...
0 -> releasef K
0 -> clear_active_fd K
0 <- clear_active_fd K
0 -> cv_broadcast K
0 <- cv_broadcast K
0 <- releasef K
0 <- ioctl K
0 <- ioctl U
0 <- _XEventsQueued U
0 <- XEventsQueued U
```

## Dtrace (Cont.)

- DTrace code to record amount of time each process with UserID 101 is in running mode (on CPU) in nanoseconds.

```

sched:::on-cpu
uid == 101
{
    self->ts = timestamp;
}

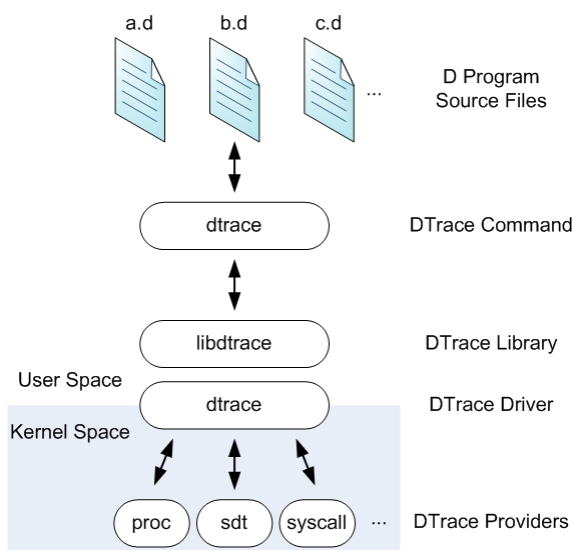
sched:::off-cpu
self->ts
{
    @time[execname] = sum(timestamp - self->ts);
    self->ts = 0;
}
    
```

```

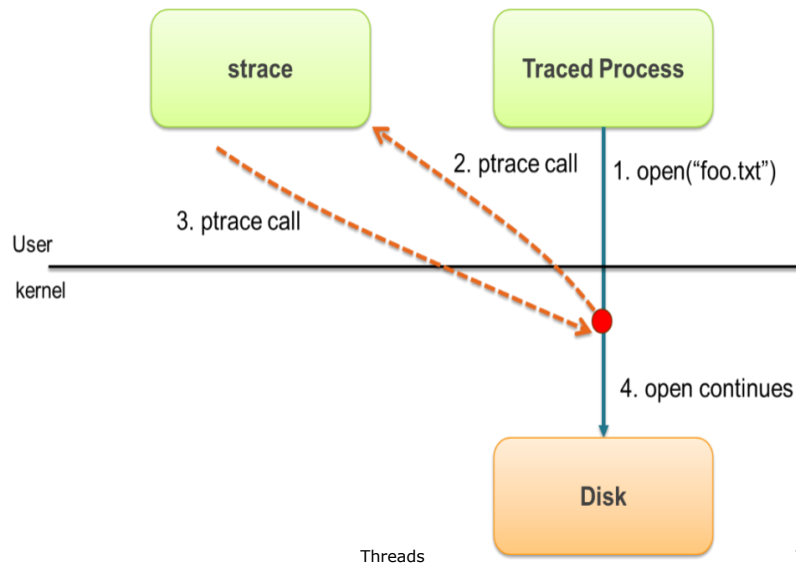
# dtrace -s sched.d
dtrace: script 'sched.d' matched 6 probes
^C
gnome-settings-d          142354
gnome-vfs-daemon         158243
dsdm                      189804
wnck-applet              200030
gnome-panel              277864
clock-applet             374916
mapping-daemon           385475
xscreensaver             514177
metacity                 539281
Xorg                     2579646
gnome-terminal           5007269
mixer-applet2            7388447
java                     10769137
    
```

Figure 2.21 Output of the D code.

## Dtrace (Cont.)



## strace (Linux)



## strace (Linux)

- ❑ Strace is based on a facility called *ptrace* that is exported by Linux and other operating systems. *Ptrace* can do many complex things and is used, for example, by debuggers like *gdb* to look into a running process. Strace uses it to instrument a target process and “listen” to that process’s system calls.
- ❑ The *ptrace* mechanism makes it possible for strace to interrupt the traced process every time a system call is invoked, capture the call, decode it, and then resume the execution of the traced process.

Threads

72

## Operating System Generation

---

- Operating systems are designed to run on any of a class of machines; the system must be configured for each specific computer site.
- **SYSGEN program obtains information concerning the specific configuration of the hardware system.**
  - Used to build system-specific compiled kernel or system-tuned
  - Can generate more efficient code than one general kernel

## Operating System Generation (cont'd)

---

- **The information must be determined:**
  - **CPU:**
    - What CPU is to be used?
    - Number of CPUs.
    - Has extended instruction sets or floating point arithmetic.
  - **Memory:**
    - Size.
  - **Devices:**
    - Type and model.
    - Interrupt number.
  - **Operating-system options:**
    - Maximum number of processes to be supported.
    - CPU-scheduling algorithm.

74

## Operating System Generation (cont'd)

- ❑ **Once the information is determined ...**
  - Source code of the operating system can be **modified** and completely **compiled** to produce a tailored operating system.
    - System generation is slower.
    - But more specific to the underlying hardware.
  - Or, the description can cause the selection of **modules** from a **precompiled library**, which are linked together to form the operating system.
    - Because the system is not recompiled, system generation is faster.
    - The resulting system may be general.
    - Easy to modify the generated system as the hardware configuration changes (such as, add a new hardware).

75

## System Boot

- ❑ **The generated operating system must be made available by the hardware.**
- ❑ **How does the hardware know where the kernel is and how to load that kernel??**
- ❑ **Booting** – the procedure of starting a computer by loading the kernel.
  - Power up or reset.
- ❑ **Need a *bootstrap program* to:**
  - Locate the kernel on the disk.
  - Load it into memory.
  - Start its execution.
  - A simple code stored in ROM or EPROM.
    - At a fixed location so that can be loaded and executed when computer is on.
  - But before that, it first initializes all aspects of the system:
    - CPU registers, device controller, the contents of main memory ...

76

## System Boot (cont'd)

---

- ❑ **Some computer systems (such as PCs) use a two-step booting process:**
  - A simple bootstrap loader fetches a more complex boot program from disk.
  - Which in turn loads the kernel.
  
- ❑ **The boot program stored in the *boot block* (a fixed location on disk) is usually sophisticated and modifiable and is to load an (or different) operating system into memory and begin its execution.**
  - Then the operating system is said to be **running**.
- ❑ **Common bootstrap loader, GRUB, allows selection of kernel from multiple disks, versions, kernel options.**
- ❑ **Kernel loads and system is then running.**

77