

CSE 4/521

Introduction to Operating Systems

Lecture 3 – Operating Systems Structures

(Operating-System Services, User and Operating-System Interface, System Calls, Types of System Calls, System Programs, Operating-System Structure, Operating-System Debugging)

Summer 2018

Overview

Objective:

1. To describe the **services** an operating system provides to users, processes, and other systems
2. To discuss the various ways of **structuring an operating system**

- Operating-System Services
- User and Operating-System Interface
- System Calls
- Types of System Calls
- System Programs
- Operating-System Structure
- Operating-System Debugging

Recap

- Storage Management
 - Migration of Data from Disk -> Registers and associated tradeoffs.
- Protection and Security
 - Implemented through userIDs and groupIDs, privilege escalation
- Kernel Data Structures
 - Linked lists, binary (search) trees, hash maps
- Computing Environments
 - Traditional, Mobile, Distributed (Client-Server, Peer-to-peer), Virtualization, Cloud Computing
- Open-Source Operating Systems
 - GNU/Linux, BSD UNIX

Question

1. Distinguish between **Client-Server** and **Peer-to-Peer** models of distributed systems. (Easy)
2. Why is **Cache Coherency** important for **Multi-processor environments**? Why must **multi-tasking environments** be careful to use **most recent value**? (Medium)
3. In your analysis of Open-source Oses, did you find any **other data structures** apart from linked lists, binary trees or hash maps? (Hard and open-ended)

Table of Content

- Operating-System Services
- User and Operating-System Interface
- System Calls
- Types of System Calls
- System Programs
- Operating-System Structure
- Operating-System Debugging

Table of Content

- **Operating-System Services**
- User and Operating-System Interface
- System Calls
- Types of System Calls
- System Programs
- Operating-System Structure
- Operating-System Debugging

Operating Systems Services

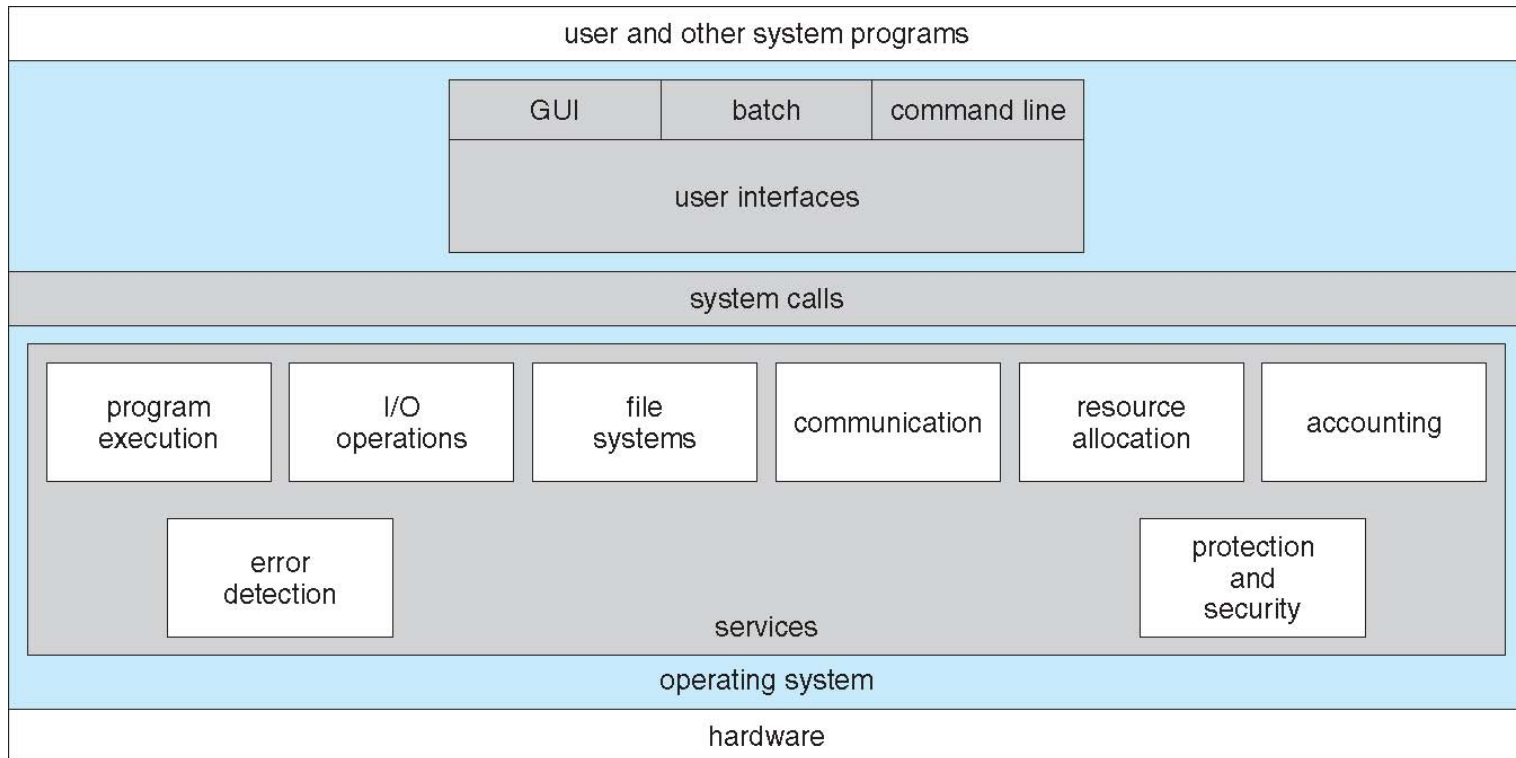


Table of Content

- Operating-System Services
- **User and Operating-System Interface**
- System Calls
- Types of System Calls
- System Programs
- Operating-System Structure
- Operating-System Debugging

User and Operating-System Interface

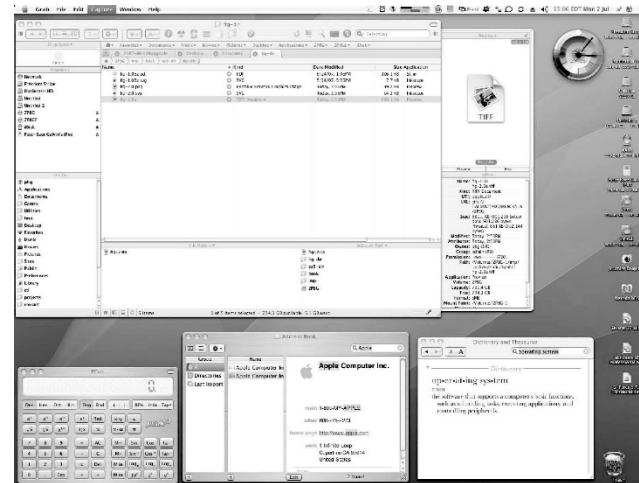
- Three predominant types of User and Operating-System Interface:

1. Command line interface

```
Default
New Info Close Execute Bookmarks
Default Default
PBG-Mac-Pro:~ pbg$ w
15:24 up 56 mins, 2 users, load averages: 1.51 1.53 1.65
USER      TTY      FROM          LOGIN@  IDLE WHAT
pbg       console  14:34         50      - w
pbg       s000    -             15:05   - w
PBG-Mac-Pro:~ pbg$ iostat 5
disk0      disk1      disk10      cpu      load average
KB/t tps MB/s  KB/t tps MB/s  KB/t tps MB/s  us sy id 1m 5m 15m
33.75 343 11.30  64.31 14  0.88  39.67 0 0.02 11 5 84 1.51 1.53 1.65
5.27 320 1.65  0.00 0  0.00  0.00 0 0.00  4 2 94 1.39 1.51 1.65
4.28 329 1.37  0.00 0  0.00  0.00 0 0.00  5 3 92 1.44 1.51 1.65
AC
PBG-Mac-Pro:~ pbg$ ls
Applications          Music                  WebEx
Applications (Parallels)  Pando Packages      config.log
Desktop               Pictures              getsmartdata.txt
Documents             Public                imp
Downloads              Sites                 log
Dropbox                Thumbs.db            pando-dist
Library                Virtual Machines     prob.txt
Movies                 Volumes              scripts
PBG-Mac-Pro:~ pbg$ pwd
/Users/pbg
PBG-Mac-Pro:~ pbg$ ping 192.168.1.1
PING 192.168.1.1 (192.168.1.1): 56 data bytes
64 bytes from 192.168.1.1: icmp_seq=0 ttl=64 time=2.257 ms
64 bytes from 192.168.1.1: icmp_seq=1 ttl=64 time=1.262 ms
AC
--- 192.168.1.1 ping statistics ---
2 packets transmitted, 2 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 1.262/1.766/2.257/0.498 ms
PBG-Mac-Pro:~ pbg$
```

Goals: Batch processing

2. Graphic user interface



Goals: Intuitive user experience

3. Touchscreen interface



Goals: Maximizing efficiency of touch

Table of Content

- Operating-System Services
- User and Operating-System Interface
- **System Calls**
- Types of System Calls
- System Programs
- Operating-System Structure
- Operating-System Debugging

System Calls

- **Programming interface** to the **services** provided by the OS
- Typically **written in a high-level language** (C or C++)
- Mostly accessed by programs via a high-level **Application Programming Interface (API)**.
- Three most common APIs are
 - **Win32 API** for Windows,
 - **POSIX API** for UNIX, Linux and Mac OS X systems
 - **Java API** for the Java virtual machine (JVM)

System Calls - Example

Collection of Systems Call APIs to copy source file to destination file



Example System Call Sequence

- Acquire input file name
- Write prompt to screen
- Accept input
- Acquire output file name
- Write prompt to screen
- Accept input
- Open the input file
 - if file doesn't exist, abort
- Create output file
 - if file exists, abort
- Loop
 - Read from input file
 - Write to output file
- Until read fails
- Close output file
- Write completion message to screen
- Terminate normally

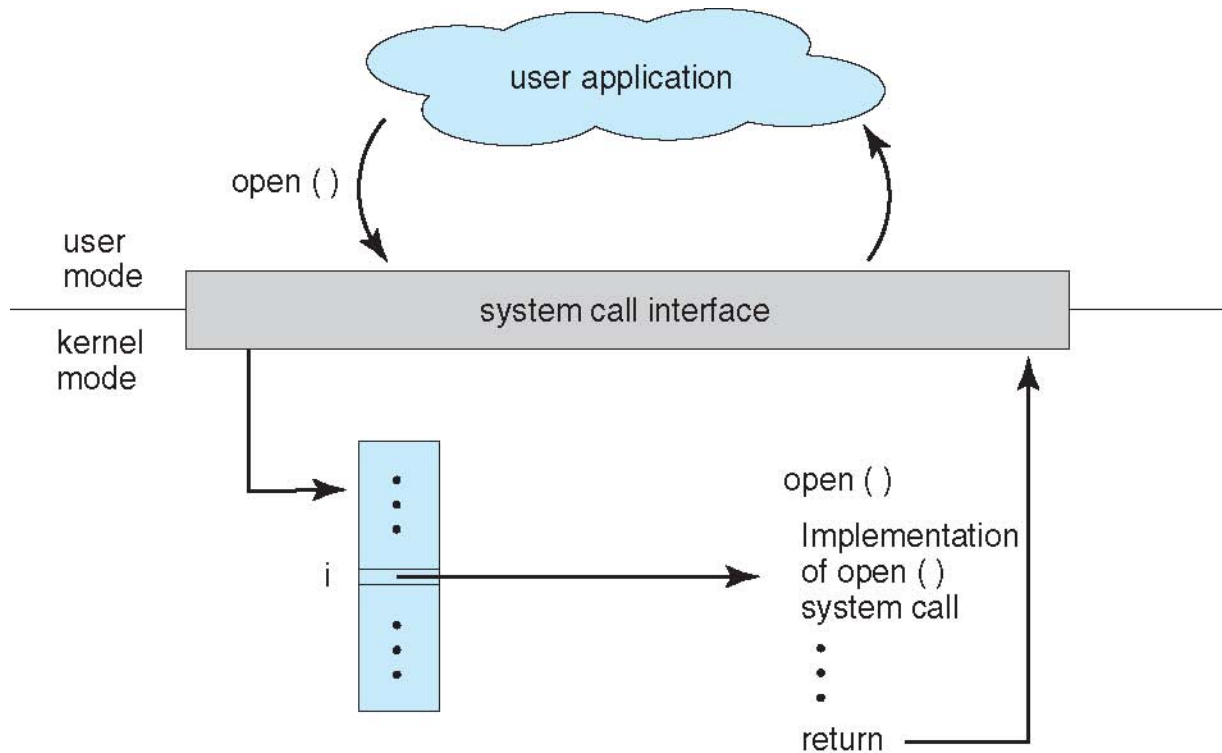
POSIX API for read system call

```
#include <unistd.h>

ssize_t read(int fd, void *buf, size_t count)
```

return value function name parameters

System Calls - Invocation



System Calls – Passing Parameters

- **Three general methods** used to pass parameters to the OS
 - Pass the parameters in **registers**
 - In some cases, may be more parameters than registers
 - Parameters stored in a **block**, or **table**, and **address of block** passed as a parameter in a register
 - This approach taken by Linux and Solaris
 - Parameters placed, or pushed, onto the **stack** by the program and popped off the stack by the operating system

Table of Content

- Operating-System Services
- User and Operating-System Interface
- System Calls
- **Types of System Calls**
- System Programs
- Operating-System Structure
- Operating-System Debugging

Types of System Calls

1. Process control

- Create process, terminate process, wait for time

2. File Management

- Open file, read file, write file, reposition file, close file

3. Device Management

- Request device, read device, write device, release device

4. Information Maintenance

- Get and set process, get and set time

5. Communications

- Create and delete communication connection, shared memory model create and gain access to memory regions

6. Protection

- Control access, get and set permissions, allow and deny users
-

Types of System Calls - Examples

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

Total LINUX
system calls:
362

Table of Content

- Operating-System Services
- User and Operating-System Interface
- System Calls
- Types of System Calls
- **System Programs**
- Operating-System Structure
- Operating-System Debugging

System Programs

- Most **users' view** of the operation system is **defined by system programs, not the actual system calls**
- Mostly divided into:
 - **File management** – (create, delete and generally manipulate files and directories)
 - **Status information** – (date, time, logging, amount of memory)
 - **Programming language support** – (compilers, assemblers, debuggers)
 - **Program loading and execution** – (absolute loaders, relocatable loaders, linkage editors)
 - **Communications** – (provide mechanism for creating virtual connections among processes, users and computer systems)
 - **Background services** – (disk checking, error logging)
 - **Application programs** – (launched by command line, click)

Table of Content

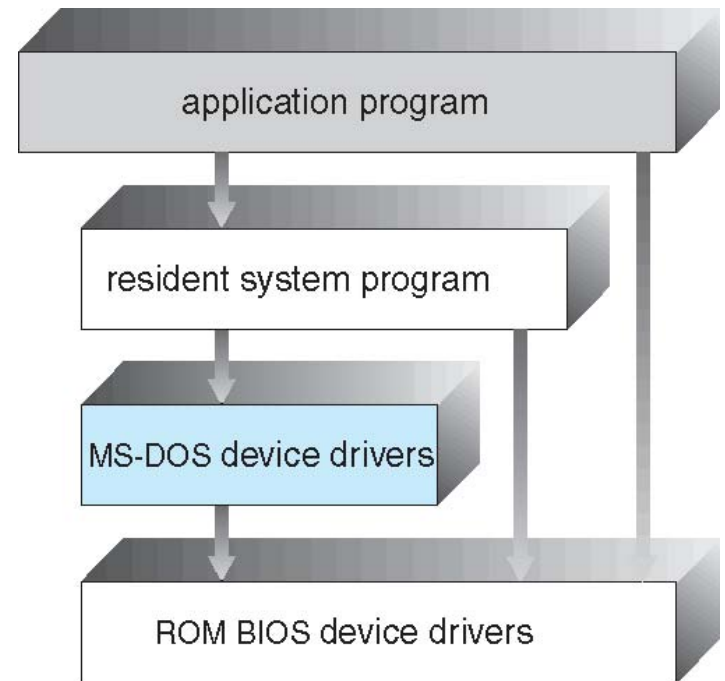
- Operating-System Services
- User and Operating-System Interface
- System Calls
- Types of System Calls
- System Programs
- **Operating-System Structure**
- Operating-System Debugging

Operating-System Structure

- General-purpose OS is very large program
- 4 ways to structure an OS
 - Simple structure – MS-DOS
 - Complex -- UNIX
 - Layered – an abstraction
 - Microkernel -Mach

Operating System Structure - Simple

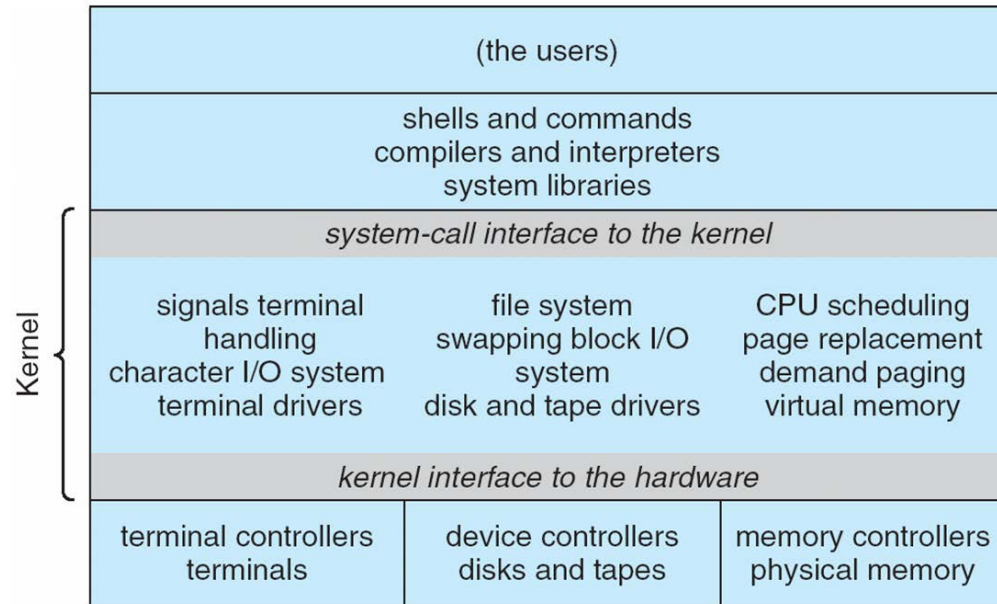
- MS-DOS – written to provide the most functionality in the least space
 - Not divided into modules
 - Although MS-DOS has some structure, its **interfaces and levels of functionality are not well separated**



Operating System Structure - Complex


The **UNIX OS** consists of 2 separate 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



Operating System Structure - Layered

- 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**



Operating System Structure - Microkernel

- 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
- Drawbacks:
 - **Performance overhead** of user space to kernel space communication

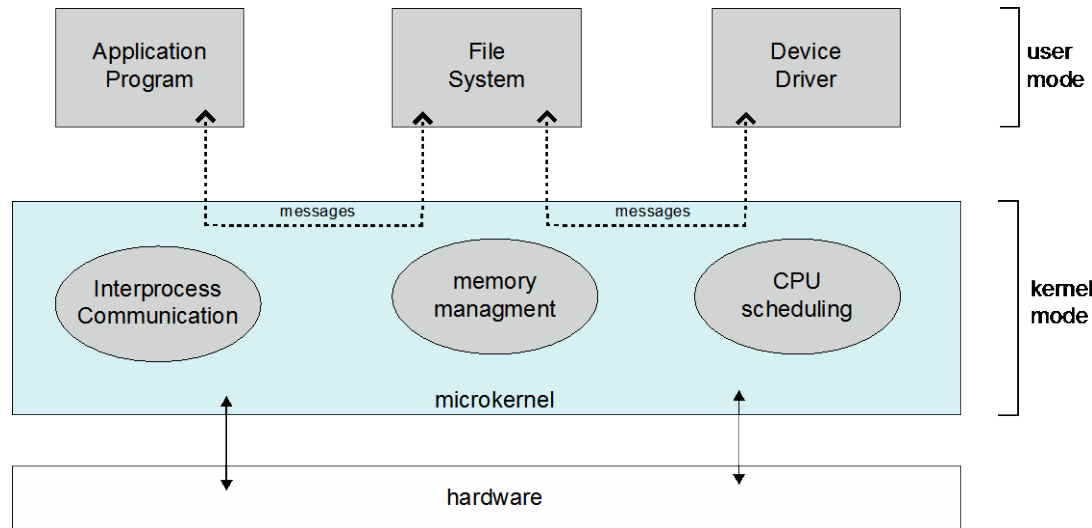


Table of Content

- Operating-System Services
- User and Operating-System Interface
- System Calls
- Types of System Calls
- System Programs
- Operating-System Structure
- **Operating-System Debugging**

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
 - **Profiling** is periodic sampling of instruction pointer to look for statistical trends

Operating System Debugging - 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

```
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.

Credential for slides

Silberschatz, Galvin and Gagne