# Operating Systems — Introduction and Processes

INF107

Stefano Zacchiroli
2023

# INF107 — Where are we in our Trip

INF107: “from the logic gate to the operating system”

- Part 1: *logic gate* → *processor*
- Part 2: *processor* → *system programs* (C programming language)
- Part 3: *system programs* → *operating system* ⇐ we are here

Goals of Part 3:

- provide an overview of what **Operating Systems (OS)** *do*,
- *how* OS work internally and how to implement one.

# What Operating Systems Do

# Warm-up Quiz

- Q: what’s an operating system (in your own words)?
- A: ⟨*your answer here*⟩

# What is an Operating System? (Intuition)

- A program that acts as an **intermediary** between a *user* of a computer and the computer *hardware*

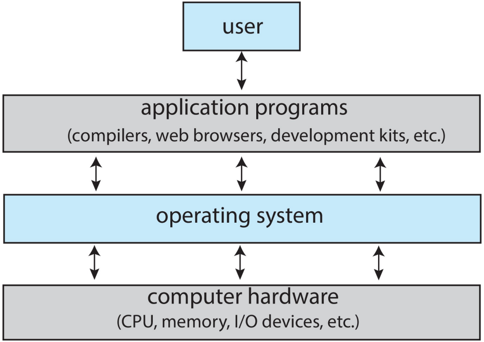

- Operating system goals:
  - *Execute user programs* and make solving user problems easier
  - Make the computer system *convenient to use*
  - Use the computer hardware in an *efficient* manner

# What Operating Systems Do

- Depends on the **point of view**
  - **Users** want convenience, ease of use and good performance
  - Don't care about resource utilization
- But shared computer such as mainframe or minicomputer must *keep all users happy*
  - Operating system is a **resource allocator** and control program *making efficient use of HW* and managing execution of user programs
- **Resources are scarce** in many contexts for different reasons
  - Servers: many users, need to share resources between them
  - Mobile devices: optimize for battery life
  - Embedded devices: limited hardware

Operating systems arbiter the allocation of scarce hardware resources to demanding users, in the best possible way.[1]

---

[1] For some precise measure of "best".

# What is an Operating System? (Definitions)

- No universally accepted definition
- *"Everything a software vendor ships when you order an OS"* is a good approximation
  - But varies wildly
- *"The one program running at all times on the computer"* is the **kernel**, *part* of the OS
- Everything else is either:
  - A **system program**[2] (ships with the OS, but is not part of the kernel), or
  - An **application program**, all programs not associated with the OS
- Today's OSes for general purpose and mobile computing also include *middleware* — a set of software frameworks that provide additional services to application developers such as databases, multimedia, graphics

---

[2] cf. INF107, part 2

# Basics of Computer System Structure

# Computer System Organization and the Bus

- One or more **CPUs** and **device controllers** connect through a common **system bus** providing access to shared memory
- **Concurrent execution** of CPUs and devices competing for memory cycles

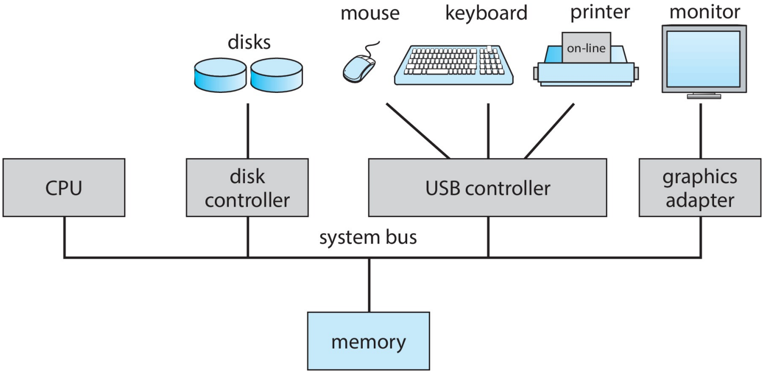

# Device Controllers and Interrupts

- Each device controller is in charge of a particular device type
- Each device controller has a **local buffer**
- Each device controller type has an operating system **device driver** (= software) to manage it
- CPU moves data: main memory ↔ local buffers
- I/O is from the device to local buffer of controller
- Device controller informs CPU that it has finished its operation by causing an **interrupt**

# Interrupts

- Interrupt transfers control to the **interrupt service routine** generally, through the interrupt vector, which contains the addresses of all the service routines
- Interrupt architecture must **save the address of the interrupted instruction** (to return to it later)
- A **trap or exception** is a software-generated interrupt caused either by an error or a user request
  - They are handled the same way than I/O interrupt
- Modern operating systems are mostly *interrupt-driven*

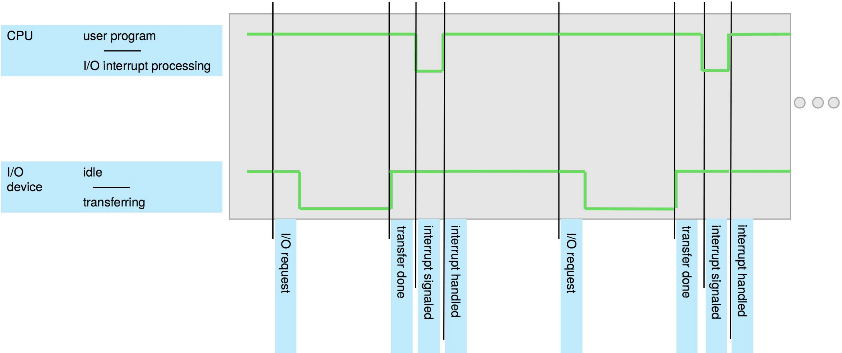

# Interrupts (cont.)

- Upon receiving an interrupt, the OS preserves the **state of the CPU** by storing the *registers and the program counter (PC)*
- Determines which type of interrupt has occurred
- Separate segments of code determine what action should be taken for each type of interrupt

A typical I/O scenario hence corresponds to the workflow shown on the right.

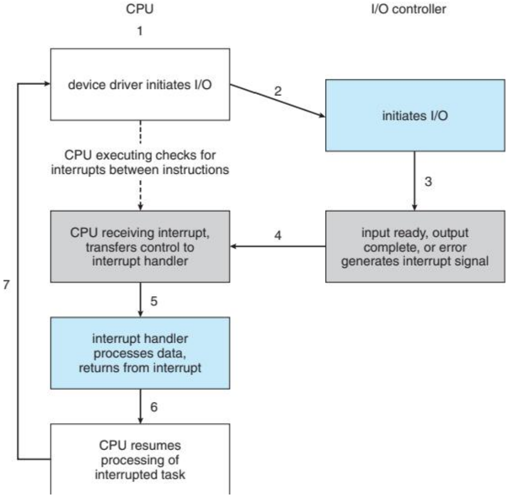

# Storage Hierarchy

- Storage is organized in a **hierarchy** with varying: speed, cost, volatility
  - **Main memory**: only large storage that *CPU can access directly*
  - **Secondary storage**: large nonvolatile storage capacity. Main types: hard disk drives (*HDD*), non-volatile memory (*NVM*)
  - **Tertiary storage**: even larger and slower (e.g., for backup purposes)

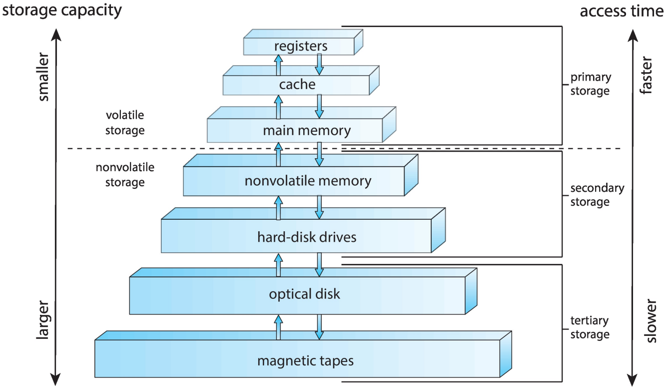

# Storage Characteristics

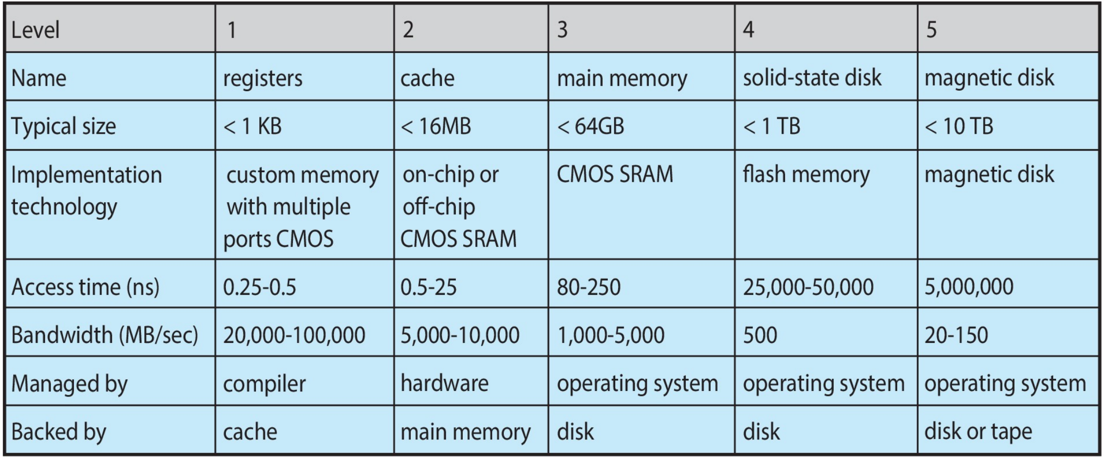

| Level | 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|---|
| Name | registers | cache | main memory | solid-state disk | magnetic disk |
| Typical size | < 1 KB | < 16MB | < 64GB | < 1 TB | < 10 TB |
| Implementation technology | custom memory with multiple ports CMOS | on-chip or off-chip CMOS SRAM | CMOS SRAM | flash memory | magnetic disk |
| Access time (ns) | 0.25-0.5 | 0.5-25 | 80-250 | 25,000-50,000 | 5,000,000 |
| Bandwidth (MB/sec) | 20,000-100,000 | 5,000-10,000 | 1,000-5,000 | 500 | 20-150 |
| Managed by | compiler | hardware | operating system | operating system | operating system |
| Backed by | cache | main memory | disk | disk | disk or tape |

# Caching

- Important principle, performed at many levels in a computer (hardware, OS, software)
- Information in use **copied from slower to faster storage** temporarily
- Faster storage **(cache) checked first** to determine if information is there
  - If it is (*cache "hit"*), information used directly from the cache (fast)
  - If not (*cache "miss"*), data copied to cache and used there
- Cache smaller than storage being cached
  - Cache management important design problem
  - Cache size and replacement policy

## Example

The path of an integer $x$ from disk to register (where the CPU can actually do something with it):

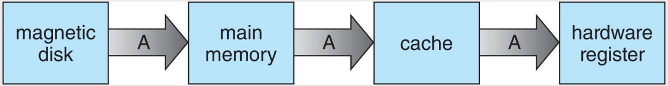

magnetic disk → A → main memory → A → cache → A → hardware register

# How a Modern Computer Works

Modern general-purpose computers:

- Implement the **von Neumann architecture**, where memory contains both data and instructions, interpreted one way or another by the CPU (if the PC points to it → it's an instruction)
- Allow devices to read/write memory directly (**Direct Memory Access**, or DMA) to reduce bus contention

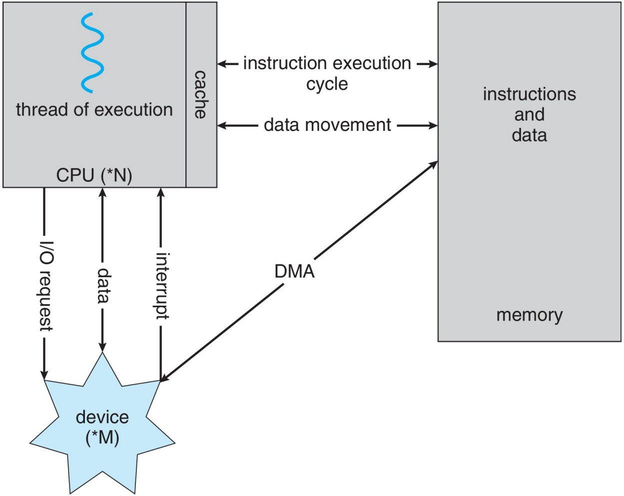

# Multiprocessor Systems

- Most systems use a single general-purpose processor
  - Plus several special-purpose processors, e.g., in device controllers
- **Multiprocessors** systems growing in use and importance
  - Also known as parallel systems, tightly-coupled systems
  - Advantages:
    1. Increased throughput
    2. Economy of scale
    3. Increased reliability (e.g., fault tolerance)
- Two types:
  - **Asymmetric** Multiprocessing – each processor is assigned a special task
  - **Symmetric** Multiprocessing – each processor performs all tasks

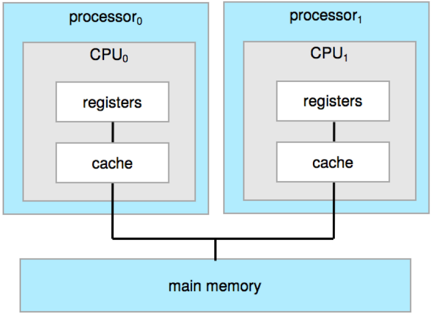

Figure: a symmetric multiprocessing architecture

# Multicore Systems

- Each physical processor *chip* (sometime confused with the term “CPU”) can host one or more units capable of executing CPU instructions at a time, called **core**
- A chip containing more than one core is called **multicore**
- Can mix and match multiprocessor and multicore in the same system
    - E.g., current high-end laptop: 1 processor, 14 cores
    - E.g., current high-end server: 4 processors, 24 cores each
- Note: only with more than one core (no matter if on the same chip or different ones) there can be parallelism, i.e., more than one CPU instructions executed *at the same time*

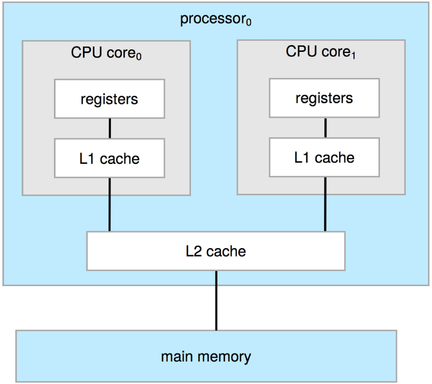

Figure: a single-processor, dual-core (= two cores) architecture

# Multiprogramming and Multitasking

## Multiprogramming (batch systems)

- Single user cannot always keep CPU and I/O devices busy
- Multiprogramming organizes jobs (code and data) so CPU always has one to execute
- A subset of total jobs in system is kept in memory
- One job selected and run via **job scheduling**
- When job has to wait (for I/O for example), OS switches to another job

## Multitasking (timesharing)

A logical extension of Batch systems — the CPU switches jobs so frequently that users can interact with each job while it is running, creating **interactive computing**

- Response time should be short (<< 1 second)
- Each user has at least one program executing in memory → **process**
- If several jobs ready to run at the same time → CPU scheduling
- If processes don't fit in memory, swapping moves them in and out to run
- Virtual memory allows execution of processes not completely in memory

Memory layout: code+data of OS and all executing programs is in memory.

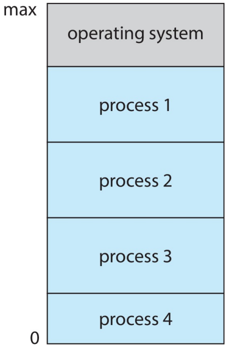

# Dual-mode Operation

- Dual-mode operation allows OS to *protect itself and other system components*
  - **User mode** and **kernel mode**
- **Mode bit** provided by *hardware*
  - Provides ability to distinguish when system is running **user code** or **kernel code**
  - When user code is running → mode bit is “user”
  - When kernel code is executing → mode bit is “kernel”
- How do we guarantee that user does not explicitly set the mode bit to “kernel”?
  - User code can requests system services by invoking **system calls** (more on this later); system calls change mode to kernel, return from call resets it to user
- Some **instructions designated as “privileged”** are only executable in kernel mode

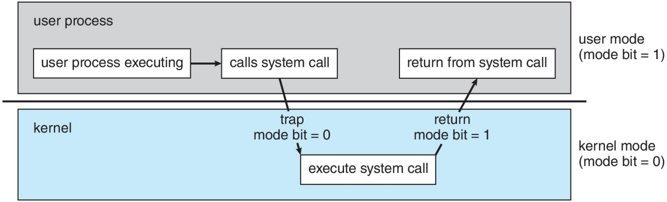

# Timer

An example of how privileged instructions are used (as well as a useful service): the **system timer** to prevent infinite loops (or, more generally, process hogging resources):

- Timer is set to interrupt the computer after some time period
- Keep a counter that is decremented by the physical clock
- Operating system set the counter (privileged instruction)
- When counter zero generate an interrupt
- Set up before scheduling process to regain control or terminate program that exceeds allotted time

# Operating System Responsibilities

# Operating System Responsibilities

An operating system has several **responsibilities**, which we briefly present in the following; we will expand upon most of them later in the course of INF107.

Several of OS responsibilities belong to the general area of **managing resources** that executing programs need to run:

- CPU, memory, file-system, mass-storage, I/O

Other OS responsibilities are more general and **cross-cutting**, such as:

- Protection and security
- Virtualization

# Process Management

- A **process** is *a program in execution* (more on this later). It is a unit of work within the system.
- Process needs resources to accomplish its task
  - CPU, memory, I/O, files
  - Initialization data
- Process termination requires reclaim of any reusable resources

- **Single-threaded** process has one program counter specifying location of next instruction to execute
  - Process executes instructions sequentially, one at a time, until completion
- **Multi-threaded** process has one program counter per thread[3]
- Typically system has many processes, some user, some operating system running concurrently on one or more CPUs
  - Concurrency by multiplexing the CPUs among the processes / threads

---

[3] More on threads in next lecture.

# Process Management (cont.)

## OS activities for process management

- Creating and terminating processes
- Suspending and resuming processes
- Providing mechanisms for:
  - Process synchronization
  - Process communication
  - Deadlock handling (more on this later)

# Memory Management

- To execute a program all (or part) of the instructions must be in memory
- All(or part) of the data that is needed by the program must be in memory
- Memory management determines what is in memory and when

## OS activities for memory management

- Keeping track of which parts of memory are currently being used and by whom
- Deciding which processes (or parts thereof) and data to move into and out of memory
- Allocating and deallocating memory space as needed

## File-system Management

OS provides uniform, logical view of information storage:

- Abstracts physical properties to logical storage unit: file
- Each medium is controlled by device (i.e., disk drive, tape drive)
- Files usually organized into directories
- Access control to determine who can access what

### OS activities for file-system management

- Creating and deleting files and directories
- Primitives to manipulate files and directories
- Mapping files onto secondary storage
- Backup files onto stable (non-volatile) storage media

# Mass-storage Management

- Usually disks used to store data that does not fit in main memory or data that must be kept for a “long” period of time
- Proper management is of central importance
- Entire speed of computer operation hinges on disk subsystem and its algorithms

### OS activities related to mass-storage management

- Mounting and unmounting
- Free-space management
- Storage allocation
- Disk scheduling
- Partitioning
- Protection

# I/O Management

- One purpose of OS is to hide peculiarities of hardware devices from the user
- I/O subsystem responsible for
    - Memory management of I/O including buffering (storing data temporarily while it is being transferred), caching (storing parts of data in faster storage for performance), spooling (the overlapping of output of one job with input of other jobs)
    - General device-driver interface
    - Drivers for specific hardware devices

# Protection and Security

- **Protection:** any mechanism for controlling access of processes or users to resources defined by the OS
- **Security:** defense of the system against internal and external attacks
  - Huge range, including: denial-of-service, worms, viruses, identity theft, theft of service
- Systems generally first distinguish among users, to determine who can do what
  - User identities (**user IDs**, security IDs) include name and associated number, one per user
  - User ID then associated with all files, processes of that user to determine access control
  - Group identifier (**group ID**) allows set of users to be defined and controls managed, then also associated with each process, file
  - **Privilege escalation** (controlled) allows user to change to effective ID with more rights

# Virtualization

- Allows operating systems to run applications within other OSes
    - Vast and growing industry
- **Emulation** used when source CPU type different from target type (i.e., PowerPC to Intel x86)
    - Generally slowest method
    - When computer language not compiled to native code — **Interpretation**
- **Virtualization** — OS natively compiled for CPU, running **guest OS** also natively compiled
    - E.g., VMware running WinXP guests, each running applications, all on native WinXP **host OS**
    - **VMM** (Virtual Machine Manager, part of the OS) provides virtualization services

# Operating System Services

# Operating System Services

Operating systems provide a number of **services** to users and running programs

Services for users:

- User interfaces: CLI, GUI, touch screen
- Program execution

Services for running programs:

- I/O, file-system ops.
- Communication between programs (locally or via the network)
- Resource allocation, error detection
- Accounting, protection, security

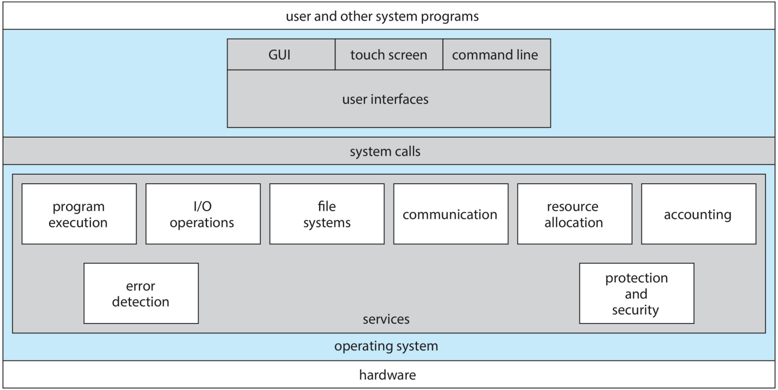

## System Calls

Running programs request OS services by invoking **system calls** (or *syscalls*, for short).

- 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)** implemented by system libraries (e.g., the **C standard library**, or `libc`) rather than direct syscall invocation
- Common high-level APIs for syscalls:
  - Win32 API for Windows
  - POSIX API for UNIX systems (including Linux and Mac OS)
  - (subset of) Java API for the Java Virtual Machine (JVM)

# System Calls — Example

Consider a program that interactively asks the user for two file names and copies the content of one file to the other. How many system call (invocations) are involved in such a task?

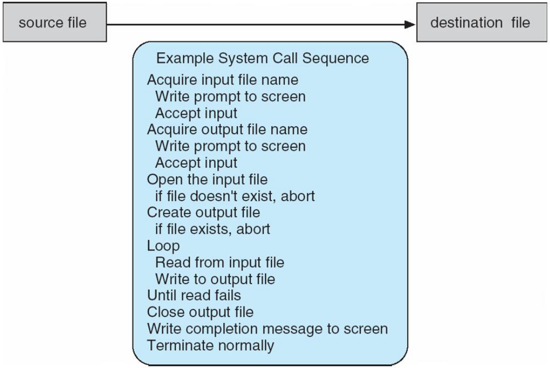

Example System Call Sequence

```
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
```

Try it out for yourself by running `strace cp input_file output_file` in a terminal.

Bottom line: *a lot* of what running programs do is using OS services.

# System Calls — Program, libc, OS — Example (1)

- `read` is a standard POSIX system call that requests the service of reading content from a file into a memory buffer of the requesting program (e.g., a byte array)
  - (You will learn about `read` details later.)
- The `read` *system call* can be invoked by C programs by calling the `read` *function* implemented in the libc
- `read` is a *blocking* system call; calling program suspends its execution, waiting for completion

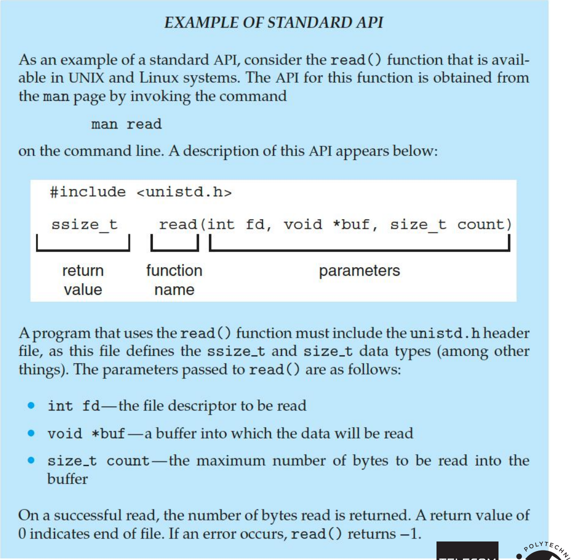

*EXAMPLE OF STANDARD API*

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 into which the data will be read
- `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 — Program, libc, OS — Example (2)

- Other, higher-level functions of the C standard library (and other libraries) are not 1-1 mappings to system calls, but call into system calls nonetheless
- For example, `printf` uses `write` (the complementary system call of `read`) to write formatted output to standard output (usually connected to your terminal)
- Note that user code (program and libc) executes in **user mode** whereas system call code executes in **kernel mode**

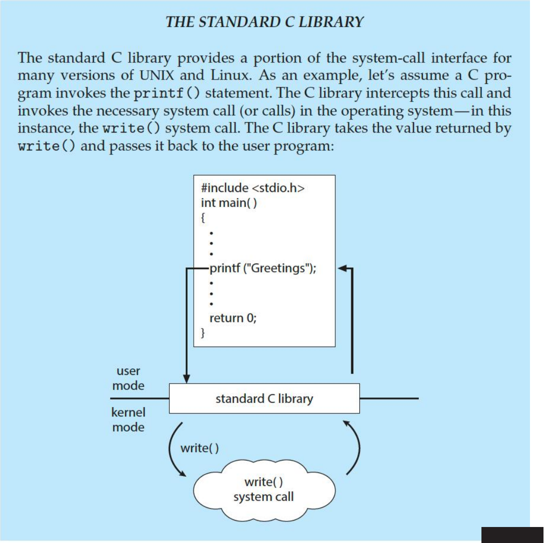

*THE STANDARD C LIBRARY*

The standard C library provides a portion of the system-call interface for many versions of UNIX and Linux. As an example, let's assume a C program invokes the `printf()` statement. The C library intercepts this call and invokes the necessary system call (or calls) in the operating system—in this instance, the `write()` system call. The C library takes the value returned by `write()` and passes it back to the user program:

```
#include <stdio.h>
int main( )
{
  •
  •
  •
  printf ("Greetings");
  •
  •
  •
  return 0;
}
```

# System Calls — User and Kernel Mode

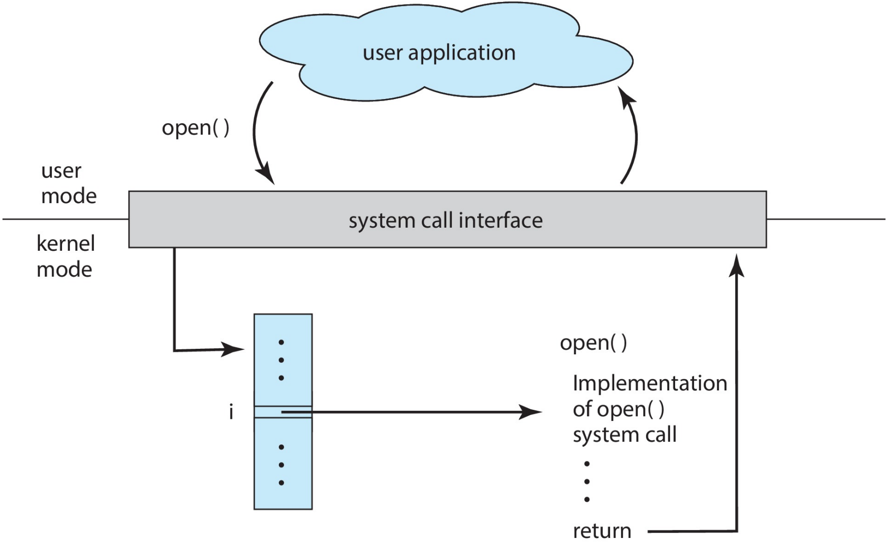

# Types of System Calls

- Many classes of system call exist, depending on the type of service requested
  - Process control
  - File management
  - Device management
  - Information maintenance
  - Communications
  - Protection
- The sets of available system calls vary across OS
- You will learn about several (UNIX) system calls in the lab sessions of INF107

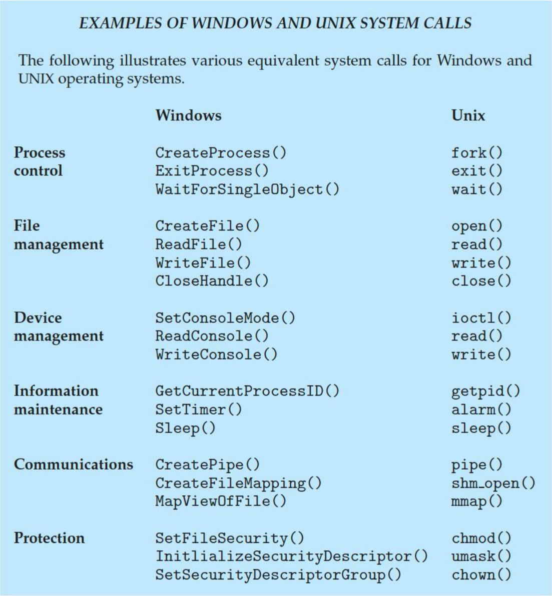

*EXAMPLES OF WINDOWS AND UNIX SYSTEM CALLS*

The following illustrates various equivalent system calls for Windows and UNIX operating systems.

| | Windows | Unix |
|---|---|---|
| **Process control** | `CreateProcess()`<br>`ExitProcess()`<br>`WaitForSingleObject()` | `fork()`<br>`exit()`<br>`wait()` |
| **File management** | `CreateFile()`<br>`ReadFile()`<br>`WriteFile()`<br>`CloseHandle()` | `open()`<br>`read()`<br>`write()`<br>`close()` |
| **Device management** | `SetConsoleMode()`<br>`ReadConsole()`<br>`WriteConsole()` | `ioctl()`<br>`read()`<br>`write()` |
| **Information maintenance** | `GetCurrentProcessID()`<br>`SetTimer()`<br>`Sleep()` | `getpid()`<br>`alarm()`<br>`sleep()` |
| **Communications** | `CreatePipe()`<br>`CreateFileMapping()`<br>`MapViewOfFile()` | `pipe()`<br>`shm_open()`<br>`mmap()` |
| **Protection** | `SetFileSecurity()`<br>`InitlializeSecurityDescriptor()`<br>`SetSecurityDescriptorGroup()` | `chmod()`<br>`umask()`<br>`chown()` |

# Processes

# Process — Concept

- An operating system executes a variety of programs that run as a process
- A **process** is a *a program in execution*
- Process execution must progress in sequential fashion
- Multiple parts (seen in INF 107, Part 2):
  - The **program code**, also called text section
  - Current activity including program counter and other processor **registers**
  - **Stack** containing temporary data
    - Function parameters, return addresses, local variables
  - **Data section** containing global variables
  - **Heap** containing memory dynamically allocated during run time

## Program vs Process

- A program is a **passive entity** stored on disk (executable file)
- A process is an **active entity**
- Program turns into a process when an executable file is loaded into memory for execution
- One program can correspond to several processes (e.g., multiple users executing same program)

# Memory Layout of a C Program (Redux)

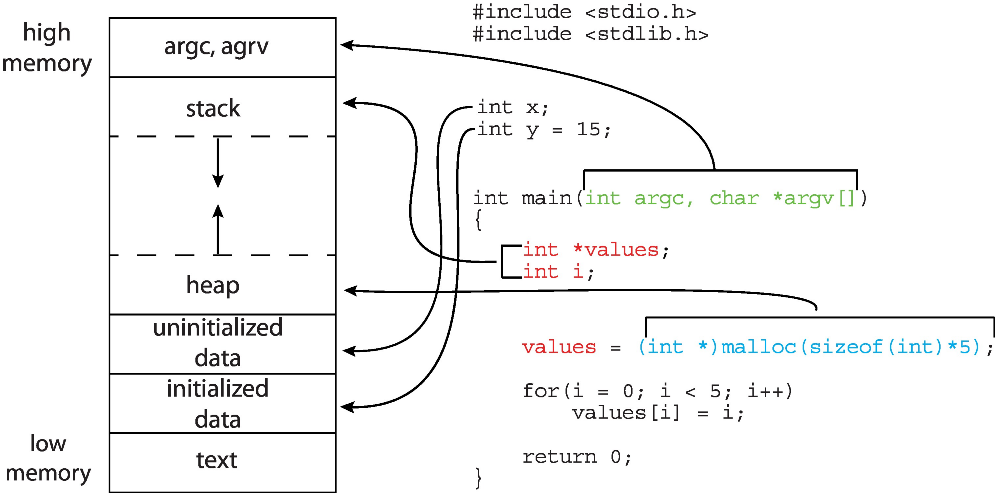

```
#include <stdio.h>
#include <stdlib.h>

int x;
int y = 15;

int main(int argc, char *argv[])
{
    int *values;
    int i;

    values = (int *)malloc(sizeof(int)*5);

    for(i = 0; i < 5; i++)
        values[i] = i;

    return 0;
}
```

(Just a reminder, you’ve seen this before.)

# Process State

As a process executes it changes **state:**

- **New:** being created
- **Running:** its instructions are being executed on a processor
- **Waiting:** waiting for some event to occur (cannot be executed, temporarily)
- **Ready:** waiting to be assigned to a processor
- **Terminated:** finished execution

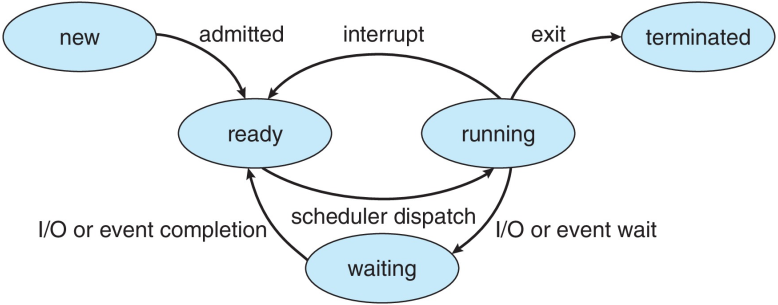

Figure: process state machine

# Process Control Block (PCB)

The full status of a process is captured by its **Process Control Block** (PCB), which contains:

- Process state — running, waiting, etc.
- Program counter — location of instruction to next execute
- CPU registers — contents of all process-centric registers
- CPU scheduling information — priorities, scheduling queue pointers
- Memory-management information — memory allocated to the process
- Accounting information — CPU used, clock time elapsed since start, time limits
- I/O status information — I/O devices allocated to process, list of open files

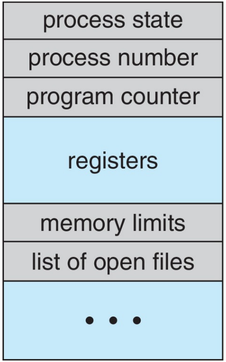
process state
process number
program counter
registers
memory limits
list of open files
• • •

All process PCBs are maintained by the OS using dedicated data structures.

# Process Scheduling

- **Process scheduler** selects among available processes for next execution on CPU core
- Goal: maximize CPU use; Implementation: quickly switch processes on/off CPU core(s)
- Maintains **scheduling queues** of processes
  - **Ready queue:** set of all processes residing in main memory, ready and waiting to execute
  - **Wait queues:** (plural!) set of processes waiting for an event (e.g., I/O, process termination, etc.)
  - Processes migrate among the various queues as they change state

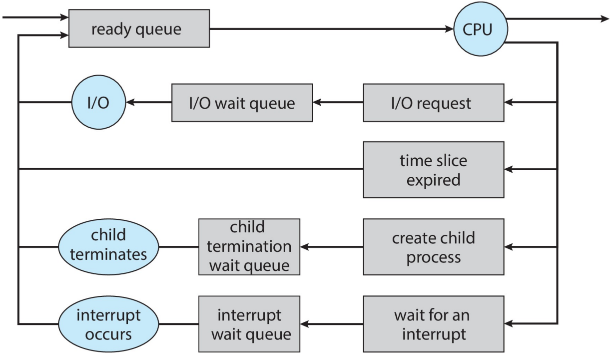

# Context Switch

A **context switch** occurs when the CPU switches from executing one process to another.

- To execute a context switch, the OS must **save the state** (or “context”) of the *old process* and load the **saved state** for the *new process*
- Full context of a process represented in the PCB
- Context-switch time is **pure overhead**; the system does no useful work while switching
  - The more complex the OS and the PCB → the longer the context switch

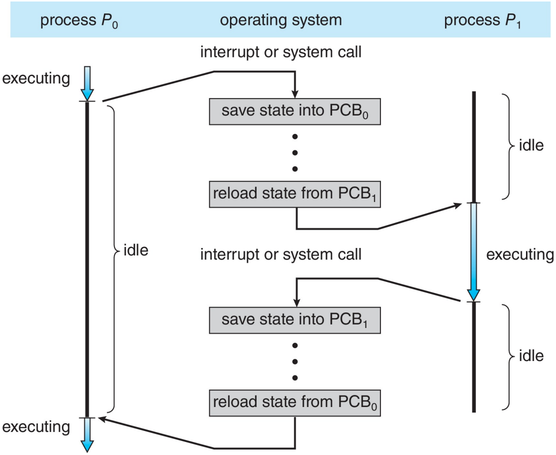

Figure: context switch timeline (from top to bottom)

# Process Creation

- **Parent process create children processes**, which, in turn create other processes, forming a **tree of processes**
- Generally, process identified and managed via a **process identifier (pid)**
- Resource sharing options (depending on the OS):
  - Parent and children share all resources
  - Children share subset of parent’s resources
  - Parent and child share no resources
- Execution options (ditto):
  - Parent and children execute concurrently
  - Parent waits until children terminate
- Address space options (ditto):
  - Child duplicate parent’s address space
  - Child has a (new) program loaded into it

# Process Creation on UNIX — Example

- `fork()` system call creates new process
  - Child shares some parent's resources (e.g., open files)
  - Parent and child execute concurrently
  - Child duplicates parent's address space
- (optional) `exec()` system call used after a fork() to replace the process address space with a new program
- (optional, for coordination) Parent process calls `wait()` system call to wait for the child to terminate

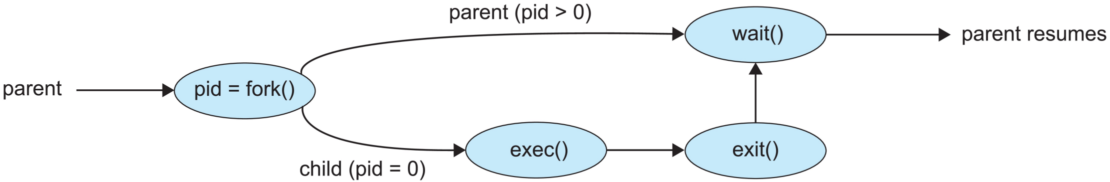

(More on this in the upcoming INF107 lab session.)

# Process Creation on UNIX — Example (cont.)

```c
#include <stdio.h>
#include <sys/types.h>
#include <sys/wait.h>
#include <unistd.h>

int main() {
    pid_t pid;
    pid = fork();  /* create a child process */
    if (pid < 0) {  /* fork() syscall failed */
        fprintf(stderr, "E: Fork failed.\n");
        return 1;
    } else if (pid == 0) { /* child process */
        execlp("/bin/ls", "ls", NULL);
    } else { /* parent process */
        wait(NULL);  /* parent will wait for the child to complete */
        printf("I: Child completed.\n");
    }
    return 0;
}
```

# Process Representation in Linux — Example

In the Linux kernel, the full status of a process (PCB) is captured in a `task_struct` structure (defined in `include/linux/sched.h`).

```c
pid t_pid;                      /* process identifier */
long state;                     /* state of the process */
unsigned int time_slice;        /* scheduling information */
struct task_struct *parent;     /* this process's parent */
struct list_head children;      /* this process's children */
struct files_struct *files;     /* list of open files */
struct mm_struct *mm;           /* address space of this process */
/* ... */
```

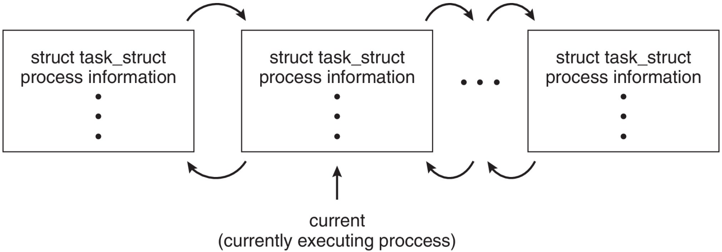

struct task_struct
process information
⋮

struct task_struct
process information
⋮

. . .

struct task_struct
process information
⋮

current
(currently executing proccess)

# Interprocess Communication

- Processes within a system may be **independent** or **cooperating**
- Cooperating process can affect or be affected by other processes, including sharing data
- Reasons for cooperating processes: information sharing, computation speedup, modularity
- Cooperating processes need **interprocess communication (IPC)**
- Two models of IPC: (a) **shared memory**, (b) **message passing**

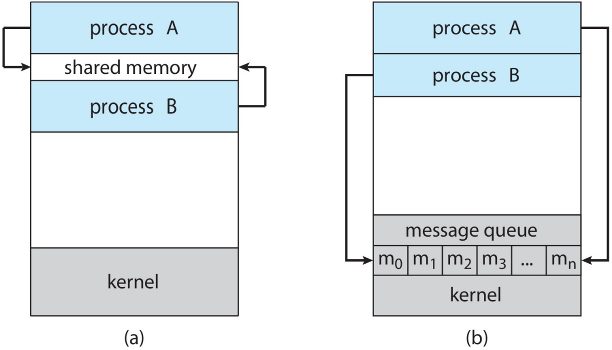

# Interprocess Communication — Examples

- IPC cannot happen without OS intervention
  - OS services must be requested either for each communication (often the case in message passing),
  - or initially to setup the communication mechanism (often the case for shared memory)
- Example of UNIX / POSIX IPC mechanisms:
  - Message passing: `pipe`, `mkfifo`, `mq_open`/`mq_send`/`mq_receive`/`mq_close`, `socket`
  - Shared memory: `mmap`, `shm_open`/`shm_unlink`

(More on some of these in later INF107 lectures and lab sessions.)

# Reading List

You should study on books, not slides! Reading material for this lecture is:

- Silberschatz, Galvin, Gagne. Operating System Concepts, Tenth Edition:
  - Chapter 1: Introduction
  - Chapter 2: Operating-System Structures
  - Chapter 3: Processes

Credits:

- Some of the material in these slides is reused (with modifications) from the official slides of the book Operating System Concepts, Tenth Edition, as permitted by their copyright note.