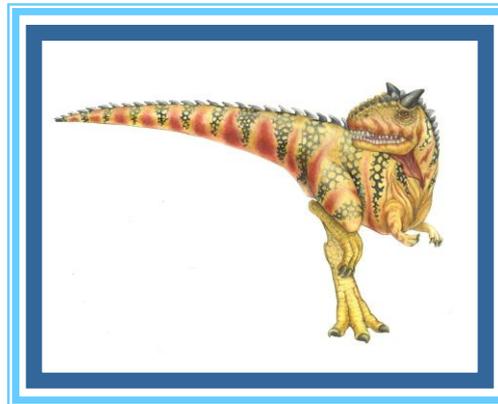



Chapter 1: Introduction



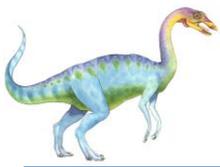


Chapter 1: Introduction

- What Operating Systems Do
- Computer-System Organization
- Computer-System Architecture
- Operating-System Structure (overview of Chapter 2)
- Operating-System Operations
- Process Management (overview of Chapters 3-7)
- Memory Management (overview of Chapters 8-9)
- Storage Management (overview of Chapters 10-13)
- Protection and Security (overview of Chapters 14-15)
- Kernel Data Structures
- Computing Environments
- Open-Source Operating Systems

Overview in this chapter, study in detail later on






Objectives

- To describe the **basic organization** of computer systems
 - Quick recap of previous Lecture
- To provide a grand tour of the **major components** of operating systems
- To give an overview of the many types of computing environments
 - Most were overviewed in Lecture 1
- To explore several open-source operating systems

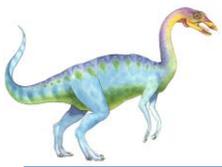





Recap: What is an Operating System?


- OS is 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





Four Components of a Computer System

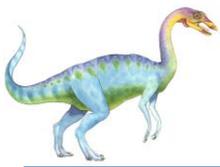




Computer Startup


- **Bootstrap program** is loaded at power-up or reboot
 - Typically stored in ROM or EPROM, generally known as **firmware**
 - Initializes all aspects of system
 - **Loads operating system** kernel and starts execution
 - ▶ We will come back to it

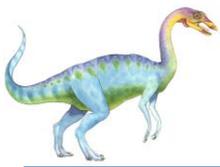




Computer System Organization

- Computer-system operation
 - One or more CPUs, device controllers connect through common bus providing access to shared memory

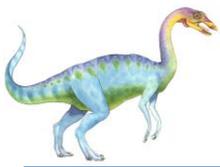




Computer-System Operation

- I/O devices and the CPU can execute **concurrently**
- Each device controller is in charge of a particular device type
- Each device controller has a **local buffer**
- CPU moves data from/to main memory to/from 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**






Common Functions of 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
- A **trap** or **exception** is a software-generated interrupt caused either by an error or a user request
- An operating system is **interrupt driven**

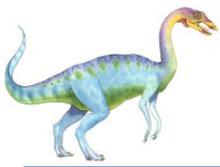





Interrupt Handling


- The OS preserves the state of the CPU by storing registers and the program counter
- Determines which type of interrupt has occurred:
 - **polling**
 - ▶ The interrupt controller polls (send a signal out to) each device to determine which one made the request
 - **vectored** interrupt system
- Separate segments of code determine what action should be taken for each type of interrupt





Interrupt Timeline






I/O Structure

- Synchronous (blocking) I/O
 - Waiting for I/O to complete
 - Easy to program, not always efficient
 - **Wait** instruction idles the CPU until the next interrupt
 - At most one I/O request is outstanding at a time
 - ▶ no simultaneous I/O processing
- Asynchronous (nonblocking) I/O
 - After I/O starts, control returns to user program without waiting for I/O completion
 - Harder to program, more efficient
 - **System call** – request to the OS to allow user to wait for I/O completion (polling periodically to check busy/done)
 - **Device-status table** contains entry for each I/O device indicating its type, address, and state





Storage Definitions and Notation Review

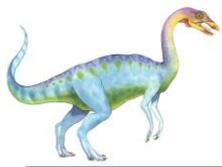
The basic unit of computer storage is the **bit**. A bit can contain one of two values, 0 and 1. All other storage in a computer is based on collections of bits. Given enough bits, it is amazing how many things a computer can represent: numbers, letters, images, movies, sounds, documents, and programs, to name a few. A **byte** is 8 bits, and on most computers it is the smallest convenient chunk of storage. For example, most computers don't have an instruction to move a bit but do have one to move a byte. A less common term is **word**, which is a given computer architecture's **native unit of data**. A word is made up of one or more bytes. For example, a computer that has 64-bit registers and 64-bit memory addressing typically has 64-bit (8-byte) words. A computer executes many operations in its native word size rather than a byte at a time.

Computer storage, along with most computer throughput, is generally measured and manipulated in bytes and collections of bytes.

A **kilobyte**, or **KB**, is $1,024$ bytes
a **megabyte**, or **MB**, is $1,024^2$ bytes
a **gigabyte**, or **GB**, is $1,024^3$ bytes
a **terabyte**, or **TB**, is $1,024^4$ bytes
a **petabyte**, or **PB**, is $1,024^5$ bytes


Computer manufacturers often round off these numbers and say that a megabyte is 1 million bytes and a gigabyte is 1 billion bytes. Networking measurements are an exception to this general rule; they are given in bits (because networks move data a bit at a time).






Storage Structure


- Main memory – only large storage media that the CPU can access directly
 - **Random access**
 - Typically **volatile**
- Secondary storage – extension of main memory that provides large **nonvolatile** storage capacity
- Hard disks – rigid metal or glass platters covered with magnetic recording material
 - Disk surface is logically divided into **tracks**, which are subdivided into **sectors**
 - The **disk controller** determines the logical interaction between the device and the computer
- **Solid-state disks** – faster than hard disks, nonvolatile
 - Various technologies
 - Becoming more popular






Storage Hierarchy

- Storage systems organized in hierarchy
 - Speed
 - Cost (per byte of storage)
 - Volatility
- **Device Driver** for each device controller to manage I/O
 - Provides uniform interface between controller and kernel






Performance of Various Levels of Storage

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
 - in hardware,
 - operating system,
 - software
- Information in use copied from slower to faster storage temporarily
 - Efficiency
- Faster storage (cache) checked first to determine if information is there
 - If it is, information used directly from the cache (fast)
 - If not, data copied to cache and used there
- Cache smaller than storage being cached
 - Cache management important design problem
 - Cache size and replacement policy

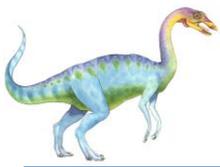





Direct Memory Access Structure

- Typically used for I/O devices that generate data in blocks, or generate data fast
- Device controller transfers **blocks** of data from buffer storage directly to main memory **without CPU intervention**
- Only **one interrupt** is generated **per block**, rather than the one interrupt per byte






How a Modern Computer Works



A von Neumann architecture







Computer-System Architecture


- Most systems use a single general-purpose processor
 - Most systems have special-purpose processors as well
- **Multiprocessors** systems growing in use and importance
 - Also known as **parallel systems**, **tightly-coupled systems**
 - Advantages include:
 1. **Increased throughput**
 2. **Economy of scale**
 3. **Increased reliability** – graceful degradation or fault tolerance
 - Two types:
 1. **Asymmetric Multiprocessing** – each processor is assigned a specific task
 2. **Symmetric Multiprocessing** – each processor performs all tasks






Symmetric Multiprocessing Architecture







A Dual-Core Design

- **Multicore**
 - Several cores on a single chip
 - On chip communication is faster than between-chip
 - Less power used






Clustered Systems



- Like multiprocessor systems, but multiple systems working together
 - Provides a **high-availability** service which survives failures
 - ▶ **Asymmetric clustering** has one machine in hot-standby mode
 - ▶ **Symmetric clustering** has multiple nodes running applications, monitoring each other
 - Some clusters are for **high-performance computing (HPC)**
 - ▶ Applications must be written to use **parallelization**






Operating System Structure


- **Multiprogramming (Batch system)** needed for efficiency
 - Single user cannot keep CPU and I/O devices busy at all times
 - 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 it has to wait (for I/O for example), OS switches to another job


- **Timesharing (multitasking)** is logical extension in which CPU switches jobs so frequently that users can interact with each job while it is running, creating **interactive** computing
 - **Response time** should be < 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 for Multiprogrammed System







Operating-System Operations

- **Interrupt driven** (hardware and software)
 - Hardware interrupt by one of the devices
 - Software interrupt (**exception** or **trap**):
 - ▶ Software error (e.g., division by zero)
 - ▶ Request for operating system service
 - ▶ Other process problems include infinite loop, processes modifying each other or the operating system







Operating-System Operations (cont.)



- **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
 - ▶ Some instructions designated as **privileged**, only executable in kernel mode
 - ▶ System call changes mode to kernel, return from call resets it to user






Using timer for preventing certain events

- **Timer** to prevent infinite loop / 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






Process Management

- A process is a program in execution. It is a unit of work within the system. Program is a ***passive entity***, process is an ***active entity***.
- Process needs resources to accomplish its task
 - CPU, memory, I/O, files
 - Initialization data
- 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






Process Management Activities

The operating system is responsible for the following activities in connection with process management:

- **Creating** and **deleting** both user and system processes
- **Suspending** and **resuming** processes
- Providing mechanisms for **process synchronization**
- Providing mechanisms for **process communication**
- Providing mechanisms for **deadlock handling**






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**
 - Optimizing CPU utilization and computer response to users
- Memory management activities
 - 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






Storage Management

- OS provides **uniform, logical view of information storage**
 - Different devices, same view
 - Abstracts physical properties to logical storage unit - **file**
 - Each medium is controlled by device (i.e., disk drive, tape drive)
 - ▶ Varying properties include access speed, capacity, data-transfer rate, access method (sequential or random)
- File-System management
 - Files usually organized into **directories**
 - **Access control** on most systems to determine who can access what
 - OS activities include
 - ▶ 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
 - Disk is slow, its I/O is often a bottleneck
- OS activities
 - Free-space management
 - Storage allocation
 - Disk scheduling






I/O Subsystem

- 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),
 - 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
 - Access control for users and groups






Computing Environments


- Recall from Lecture 1:
- Stand-alone general purpose machines
- Mobile:
 - Handheld smartphones, tablets, etc
- Distributed computing
- Client-server computing
- Peer-to-Peer computing
- Cloud-computing
- Real-Time systems
- Embedded Systems







Computing Environments - Virtualization

- **Virtualization** (more details in [Chapter 16](#), which will not be tested)
 - **Host** OS, natively compiled for CPU
 - **VMM** - virtual machine manager
 - ▶ Creates and runs **virtual machines**
 - VMM runs **guest** OSes, also natively compiled for CPU
 - Applications run within these guest OSes
 - Example: Parallels for OS X running Win and/or Linux and their apps
 - Some VMM'es run within a host OS
 - But, **some act as a specialized OS**
 - ▶ Example. VMware ESX: installed on hardware, runs when hardware boots, provides services to apps, runs guest OSes
- Vast and growing industry
- Use cases
 - Developing apps for multiple different OSes on 1 PC
 - Very important for **cloud computing**
 - ▶ Executing and managing **compute environments** in data centers





Computing Environments - Virtualization



End of Chapter 1

