

Operating Systems

Chapter 7

Source: Foundations of Computer Science © Cengage Learning

Objectives

After studying this chapter, students should be able to:

- □ Understand the role of the operating system.
- Understand the process of bootstrapping to load the operating system into memory.
- □ List the components of an operating system.
- Discuss the role of the memory manager.
- □ Discuss the role of the process manager.
- □ Discuss the role of the device manager.
- Discuss the role of the file manager in an operating system.
- □ Understand the main features of three common operating systems: UNIX, Linux and Windows.

7.2

7-1 INTRODUCTION

7.3

A computer is a system composed of two major components: hardware and software. Computer hardware is the physical equipment. Software is the collection of programs that allows the hardware to do its job.

<u>Computer software is divided into two broad categories:</u> the operating system and application programs. Application programs use the computer hardware to solve users' problems. The operating system, on the other hand, controls the access to hardware by users.



Operating System

An **<u>operating system</u>** is complex, so it is difficult to give a simple universal definition. Instead, here are some common definitions:

- □ An operating system is an interface between the hardware of a computer and the user (programs or humans).
- □ An operating system is <u>a program</u> (or a set of programs) that facilitates the execution of other programs.
- □ An operating system acts as a general manager supervising the activity of each component in the computer system.

Two major design goals of an operating system are: **D** Efficient use of hardware.

Ease of use of resources.

Bootstrap Process

The operating system provides supports for other programs. For example, it is responsible for loading other programs into memory for execution. However, <u>the operating system</u> itself is a program that needs to be loaded into the memory and be run. How is this dilemma solved?

The solution is a two-stage process. A very small section of memory is made of **ROM** and holds a small program called the **bootstrap program**. When the computer is turned on, the CPU counter is set to the first instruction of this bootstrap program and executes the instructions in this program. When loading is done, the program counter is set to the first instruction of the operating system in <u>RAM</u>.

The Bootstrap Process



7.6

7-2 EVOLUTION

Operating systems have gone through a long history of evolution, which we summarize here.

Batch systems

Batch operating systems were designed in the 1950s to control mainframe computers. At that time, <u>computers were</u> <u>large machines</u> that used <u>punched cards for input</u>, <u>line</u> printers for output and <u>tape drives for secondary storage</u> media. Each program to be executed was called a job. <u>A</u> programmer who wished to execute a job sends a request to the operating system.

Time-sharing systems

To use computer system resources efficiently, *multiprogramming* was introduced. The idea is to hold several jobs in memory at a time, and only assign a resource to a job that needs it on the condition that the resource is available.

<u>Multiprogramming</u> brought the idea of time sharing: resources could be shared between different jobs, with each job being allocated a portion of time to use a resource. Because a computer is much faster than a human, time sharing is hidden from the user—each user has the impression that the whole system is serving them exclusively.

7.8

Personal systems

When personal computers were introduced, there was a need for an operating system for this new type of computer. During this era, **single-user operating systems** such as **DOS** (**Disk Operating System**) were introduced.

Parallel systems

7.9

The need for more speed and efficiency led to the design of **parallel systems: multiple CPUs on the same machine**. Each CPU can be used to serve one program or a part of a program, which means that many tasks can be accomplished in parallel instead of serially. The operating systems required for this are more complex than those that support single CPUs.

Distributed systems

Networking and internetworking, as we saw in Chapter 6, have created a new dimension in operating systems. A job that was previously done on one computer can now <u>be</u> shared between computers that may be thousands of <u>miles apart</u>. Distributed systems combine features of the previous generation with new duties such as controlling security.

Real-time systems (RTOS)

A real-time system is expected to do a task within a specific time constraint. They are used with real-time applications, which monitor, respond to or control external processes or environments.

7.10

7-3 COMPONENTS

Today's operating systems are very complex. An operating system needs to manage different resources in a computer system. It resembles an organization with several managers at the top level. Each manager is responsible for managing their department, but also needs to cooperate with others and coordinate activities. A modern operating system has at least four duties: memory manager, process manager, device manager and file manager.

Components of an operating system



User interface

Each operating system has a **user interface**, a program that accepts requests from users (processes) and interprets them for the rest of the operating system. A user interface in some operating systems, such as UNIX, is called a **shell**.

In others, it is called a <u>window</u> to denote that it is menu driven and has a <u>GUI (graphical user interface)</u> component.

Memory manager

One of the responsibilities of a modern computer system is **memory management**. Although the memory size of computers has increased tremendously in recent years, so has the size of the programs and data to be processed. <u>Memory allocation must be managed to prevent applications from running out of memory</u>. Operating systems can be divided into two broad categories of memory management: monoprogramming and multiprogramming.

7.14

Monoprogramming

In **monoprogramming**, most of the memory capacity is dedicated to a single program; only a small part is needed to hold the operating system. In this configuration, the whole **program is in memory for execution**. When the program finishes running, the program area is occupied by another program.



Multiprogramming

programs.

7.16

In multiprogramming, <u>more than one program is in</u> <u>memory at the same time, and they are executed</u> <u>concurrently</u>, with the CPU switching rapidly between the

Operating system
Program 1
Program 2
Program 3
Memory



Demand segmentation



Virtual memory

7.20

Demand paging and **demand segmentation** mean that, when a program is being executed, part of the program is in memory and part is on disk. This means that, for example, a memory size of 10 MB can execute 10 programs (each of size 3 MB). At any moment, <u>10 MB</u> of the 10 programs are in memory and <u>20 MB are on disk</u>. Therefore a virtual memory size of <u>30 MB</u>. Virtual memory, which implies demand paging, demand segmentation or both, is used in almost all operating systems today.



Process manager

A second function of an operating system is process management, but before discussing this concept, we need to define some terms.

Program, job, and process

- A program is a non-active set of instructions stored on disk.
- □ A program becomes a job from the moment it is selected for execution until it has finished running and becomes a program again.
- □ A process is a program in execution. It is a program that has started but has not finished.

7.21

State diagrams

The relationship between a program, a job and a process becomes clearer if we consider how a program becomes a job and how a job becomes a process.



Schedulers

To move a job or process from one state to another, the process manager uses two schedulers: job scheduler and process scheduler.



Queuing

To handle multiple processes and jobs, the process manager uses queues (waiting lists). A job control block or process control block is associated with each job or process. This is a block of memory that stores information about that job or process. The process manager stores the job or process control block in the queues instead of the job or process itself.



Process synchronization

The whole idea behind process management is to **synchronize different processes with different resources**. Whenever resources can be used by more than one user (or process, in this case), we can have two problematic situations: **deadlock** and **starvation**.





Device manager

The device manager, or input/output manager, is responsible for access to input/ output devices. There are limitations on the number and speed of input/output devices in a computer system.

□ The device manager <u>monitors every input/output</u> <u>device</u> constantly to ensure that the device is functioning properly.

□ The device manager <u>maintains a queue for each</u> <u>input/output device</u> or one or more queues for similar input/output devices.

□ The device manager <u>controls the different policies for</u> <u>accessing input/output devices</u>.

File manager

Operating systems today use a file manager to control access to files. A detailed discussion of the file manager also requires advanced knowledge of operating system principles and file access concepts that are beyond the scope of this book. The file manager:

□ controls access to files.

□ supervises the <u>creation</u>, <u>deletion</u>, <u>and modification of</u> <u>files</u>.

- □ controls the <u>naming of files</u>.
- □ supervises the <u>storage of files</u>.
- □ is responsible for <u>archiving and backups</u>.

7-4 A SURVEY OF OPERATING SYSTEMS

In this section we introduce some popular operating systems and encourage you to study them further. We have chosen three operating systems that are familiar to most computer users: **UNIX**, **Linux** and **Windows**.

UNIX was originally developed in 1969 by Thomson and Ritchie of the Computer Science Research Group at **Bell Laboratories**. UNIX has been a popular operating system among computer programmers and computer scientists. UNIX is a <u>multiuser</u>, <u>multiprocessing</u>, <u>portable operating</u> <u>system</u>. It is designed to <u>facilitate programming</u>, <u>text</u> <u>processing</u> and <u>communication</u>. 729



Linux

In 1991, Linus Torvalds, a Finish student at the University of Helsinki at the time, developed a new operating system that is known today as Linux. The initial kernel, which was similar to a small subset of UNIX, has grown into a fullscale operating system today. The Linux 2.0 kernel, released in 1997, was accepted as a commercial operating system: it has all features traditionally attributed to UNIX.

Windows

In the late 1980s Microsoft, under the leadership of Dave Cutler, started development of a new single-user operating system to replace MS-DOS (Microsoft Disk Operating System). Several versions of windows are followed, such as Windows NT (NT standing for New Technology), Windows 2000, Windows XP (XP stands for eXPerience in 2001), Windows 7 (supports the functions for touch-controlled screen and has released this year of 2009), Windows 8/8.1, Windows 10, and so on.



The architecture of Windows



Review Questions

> What is **<u>the definition</u>** of operating system (OS)?

> Please describe <u>the bootstrap process</u> when PC is turned on.

> Explain the following <u>terminologies</u>: RTOS, GUI, deadlock, and starvation.

Please descript <u>their functions</u> of memory manager, process manager, device manager, and file manager.

> Please give three latest kind of operating systems for PCs, iPad, and cellphones.