


7 Operating Systems



Objectives

After studying this chapter, the student 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 NT.

A computer is a system composed of two major components:

- ◆ hardware
- ◆ software

1. **Computer hardware** is the **physical equipment**. Software is the **collection of programs** that allows the hardware to do its job.
2. **Computer software** is divided into two broad categories: the **operating system** and **application programs** (Figure 7.1).
 - a. The operating system, controls the access to hardware by users.
 - b. Application programs use the computer hardware to solve users' problems.




Figure 7.1 A computer system

7-1 INTRODUCTION

An **operating system** 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 a **program** (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.



An operating system is an interface between the hardware of a computer and the user (programs or humans) that facilitates the execution of other programs and the access to hardware and software resources.

Two major design goals of an operating system are:

- Efficient use of hardware.**
- Ease of use of resources.**

Bootstrap process

The operating system, based on the above definitions, provides supports for other programs. For example, it is responsible for loading other programs into memory for execution. However, **the operating system itself is a program that needs to be loaded into the memory and be run.**

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 RAM.

...

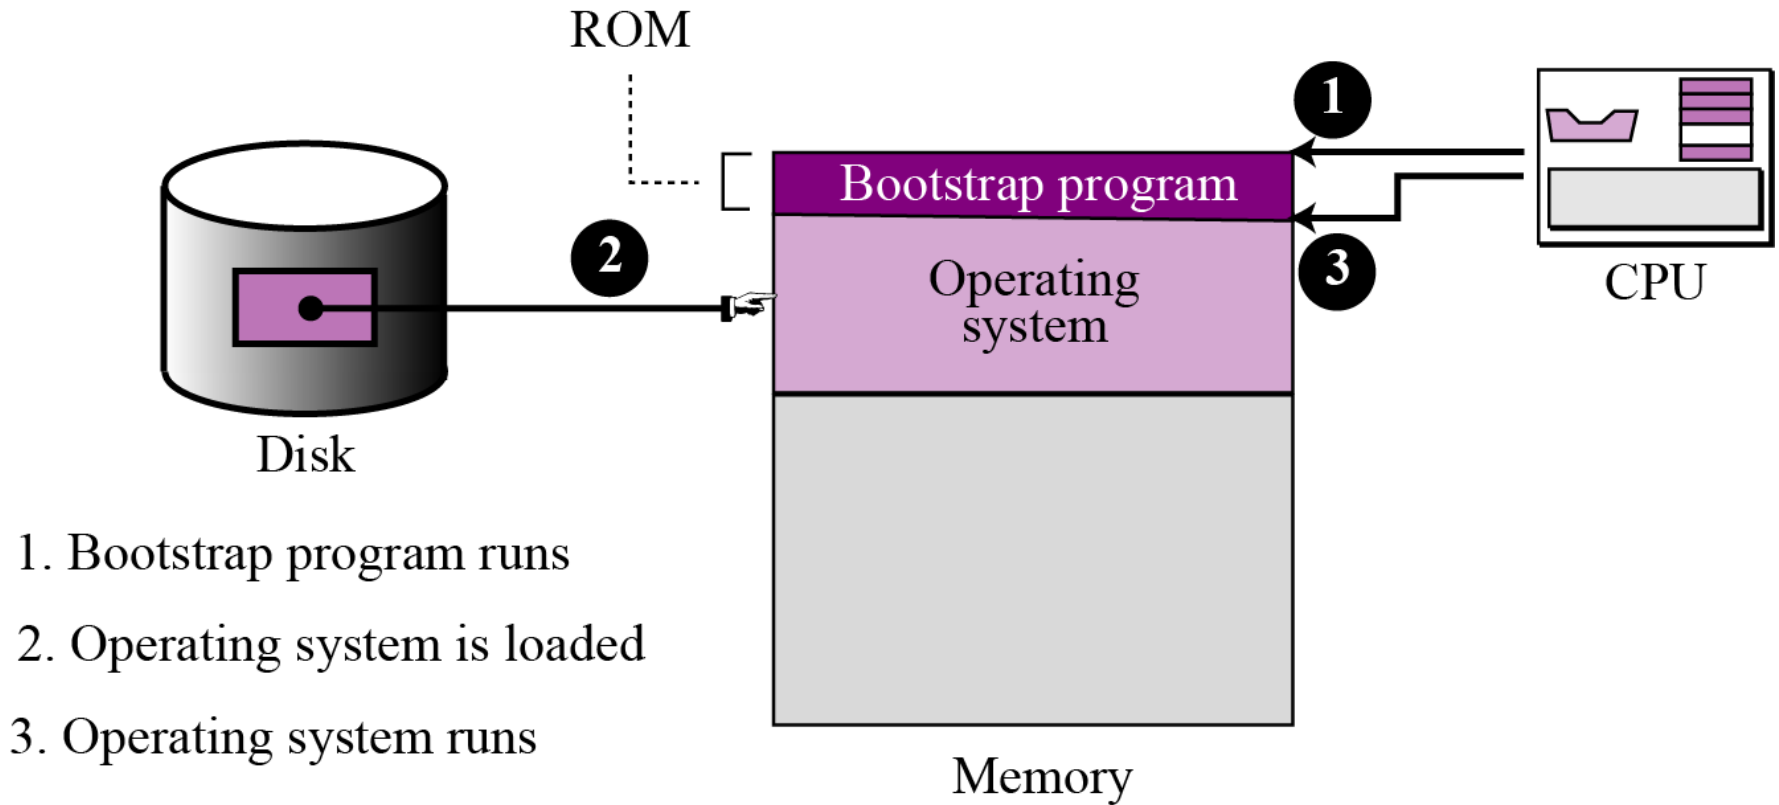


Figure 7.2 The bootstrap process

7-2 EVOLUTION

Batch systems

Batch operating systems were designed in the 1950s to control mainframe computers. At that time, computers were large machines that used punched cards for input, line printers for output and tape drives for secondary storage media.

Each program to be executed was called a job. A 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.

Multiprogramming 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.

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

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

A job that was previously done on one computer can now be shared between computers (**using network**) that may be thousands of miles apart.

Distributed systems combine features of the previous generation with new duties such as controlling security.

Real-time systems

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.

7-3 COMPONENTS

A modern operating system has at least four duties:

1. **Memory manager**
2. **Process manager**
3. **Device manager**
4. **File manager**

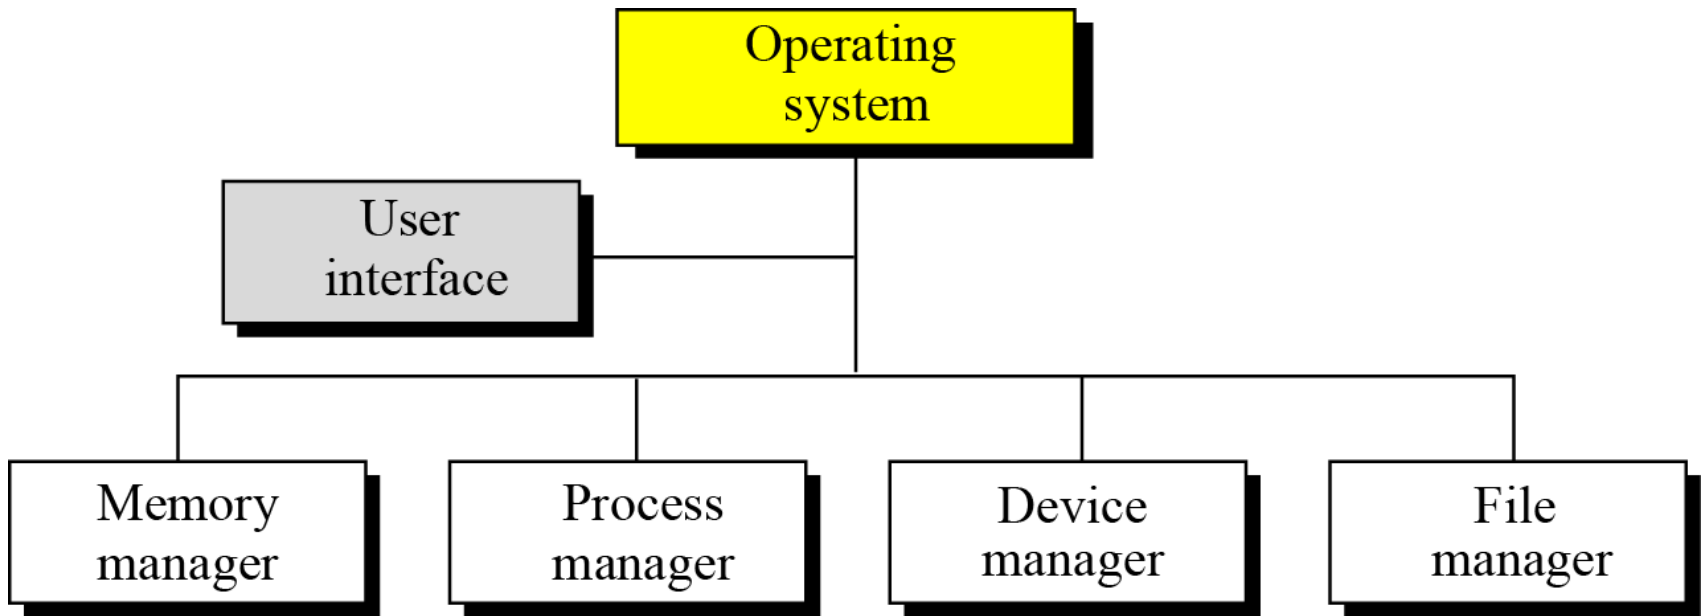


Figure 7.3 Components of an operating system

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 **window** to denote that it is menu driven and has a GUI (graphical user interface) 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, the sizes of the programs and data increase more tremendous.

Memory allocation must be managed to prevent applications from running out of memory. Operating systems can be divided into two broad categories of memory management: monoprogramming and multiprogramming.

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.




Figure 7.4 Monoprogramming

Multiprogramming

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

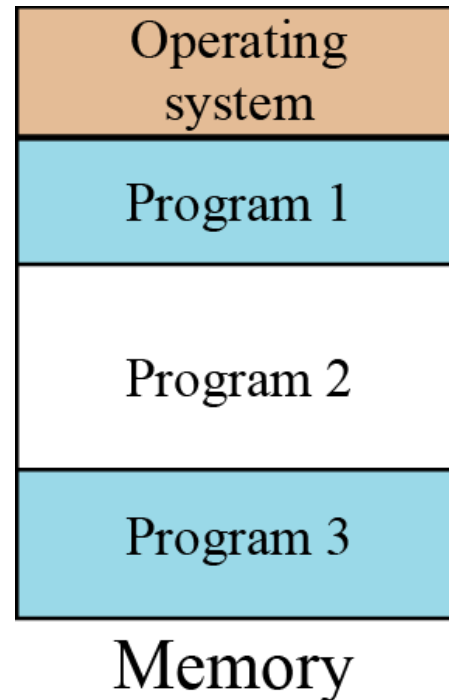


Figure 7.5 Multiprogramming

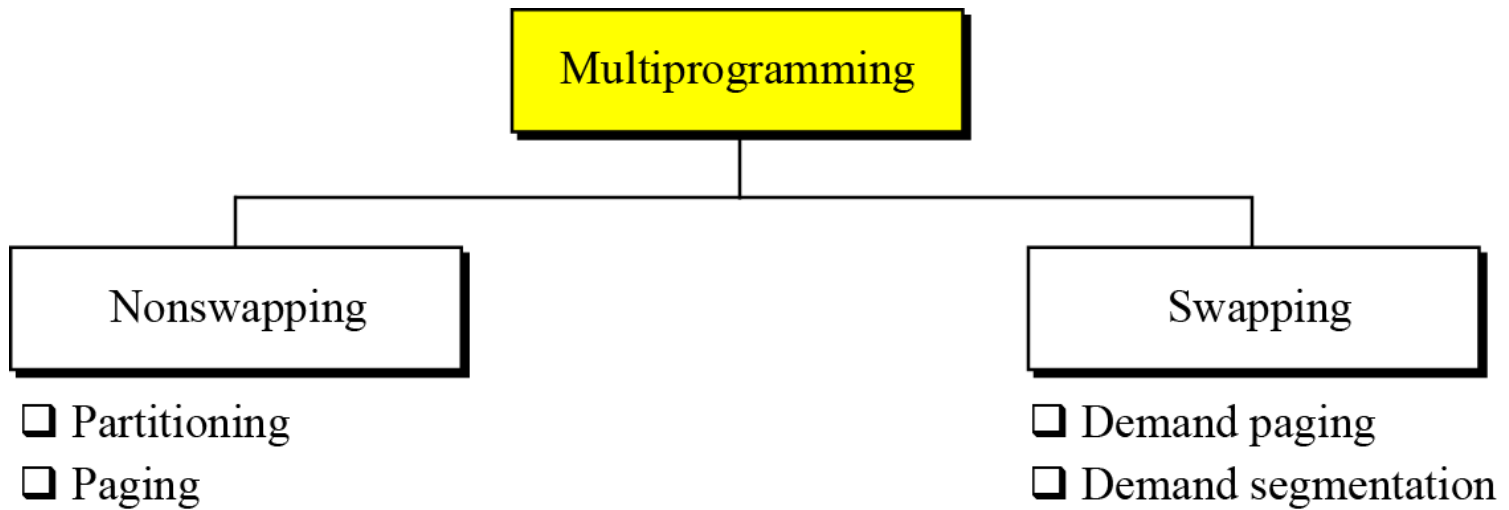
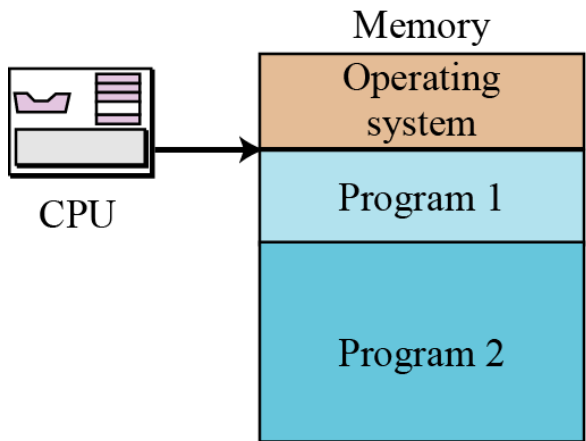



Figure 7.6 Categories of multiprogramming



a. CPU starts executing program 1



b. CPU starts executing program 2

Figure 7.7 Partitioning

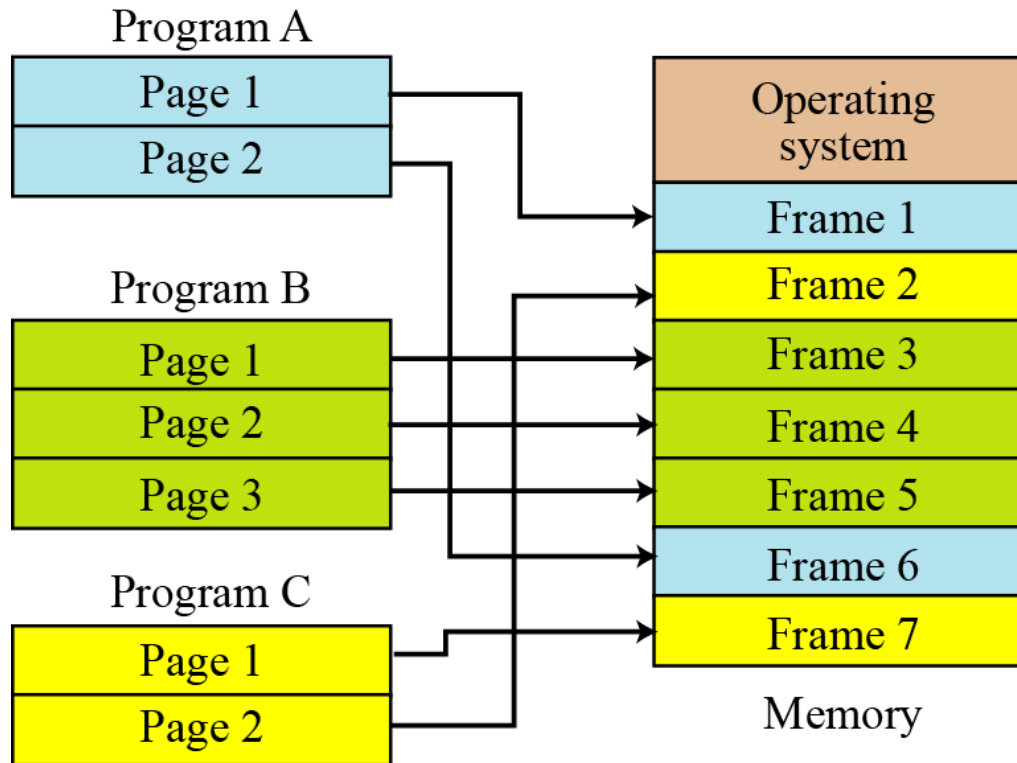


Figure 7.8 Paging

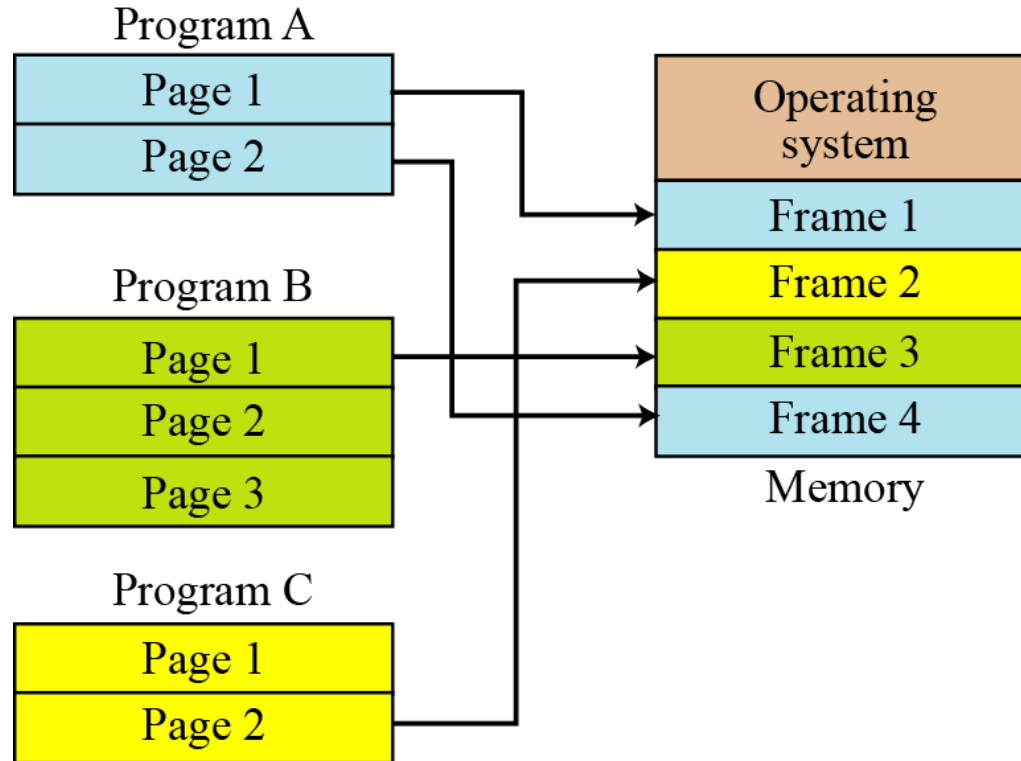


Figure 7.9 Demand paging

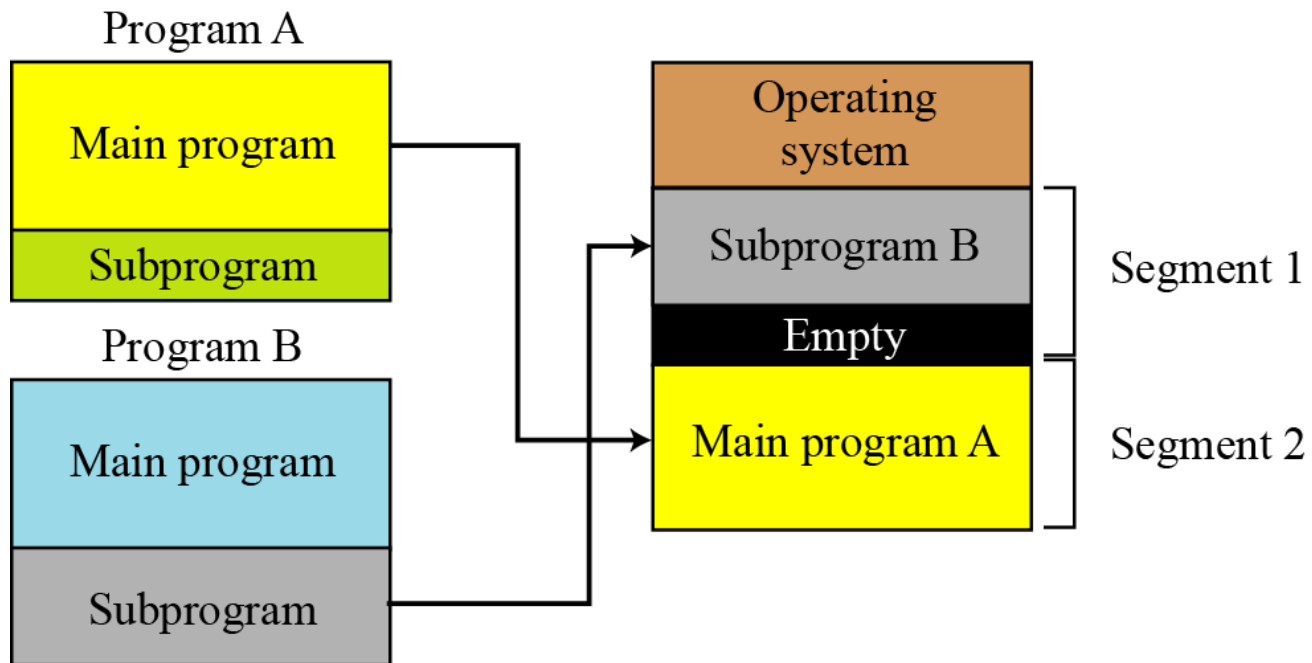


Figure 7.10 Demand segmentation

Virtual Memory

- Virtual Memory
 - A technique that allows the execution of a process that may not be completely in memory.
- Motivation:
 - An entire program in execution may not all be needed at the same time!
 - e.g. error handling routines, a large array, certain program features, etc

Virtual Memory

- Potential Benefits
 - Programs can be much larger than the amount of physical memory. Users can concentrate on their problem programming.
 - The level of multiprogramming increases because processes occupy less physical memory.
 - Each user program may run faster because less I/O is needed for loading or swapping user programs.
- Implementation: demand paging, demand segmentation (more difficult), etc.

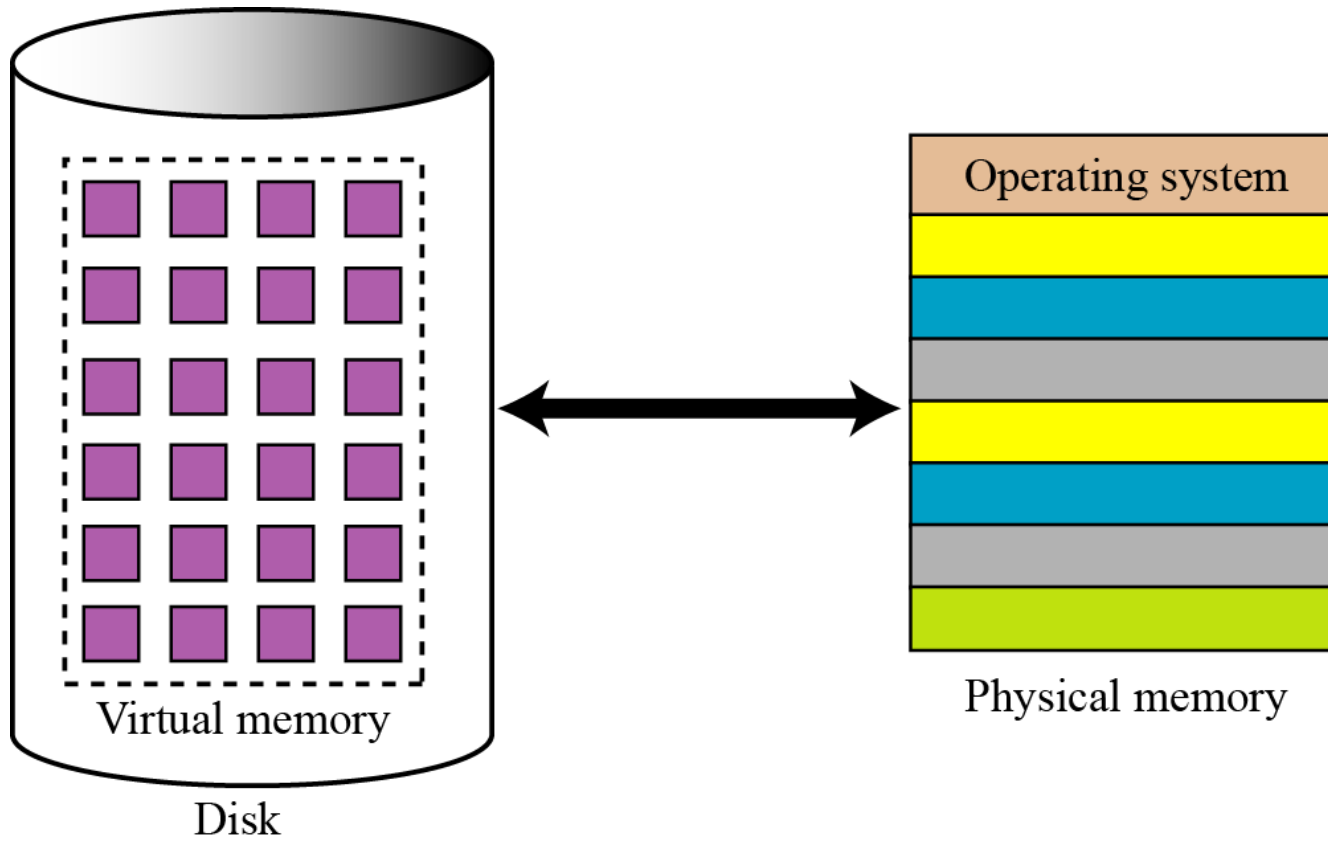


Figure 7.11 Virtual memory

Process manager

A second function of an operating system is process management.

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.

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.

This can be illustrated with a state diagram that shows the different states of each of these entities.

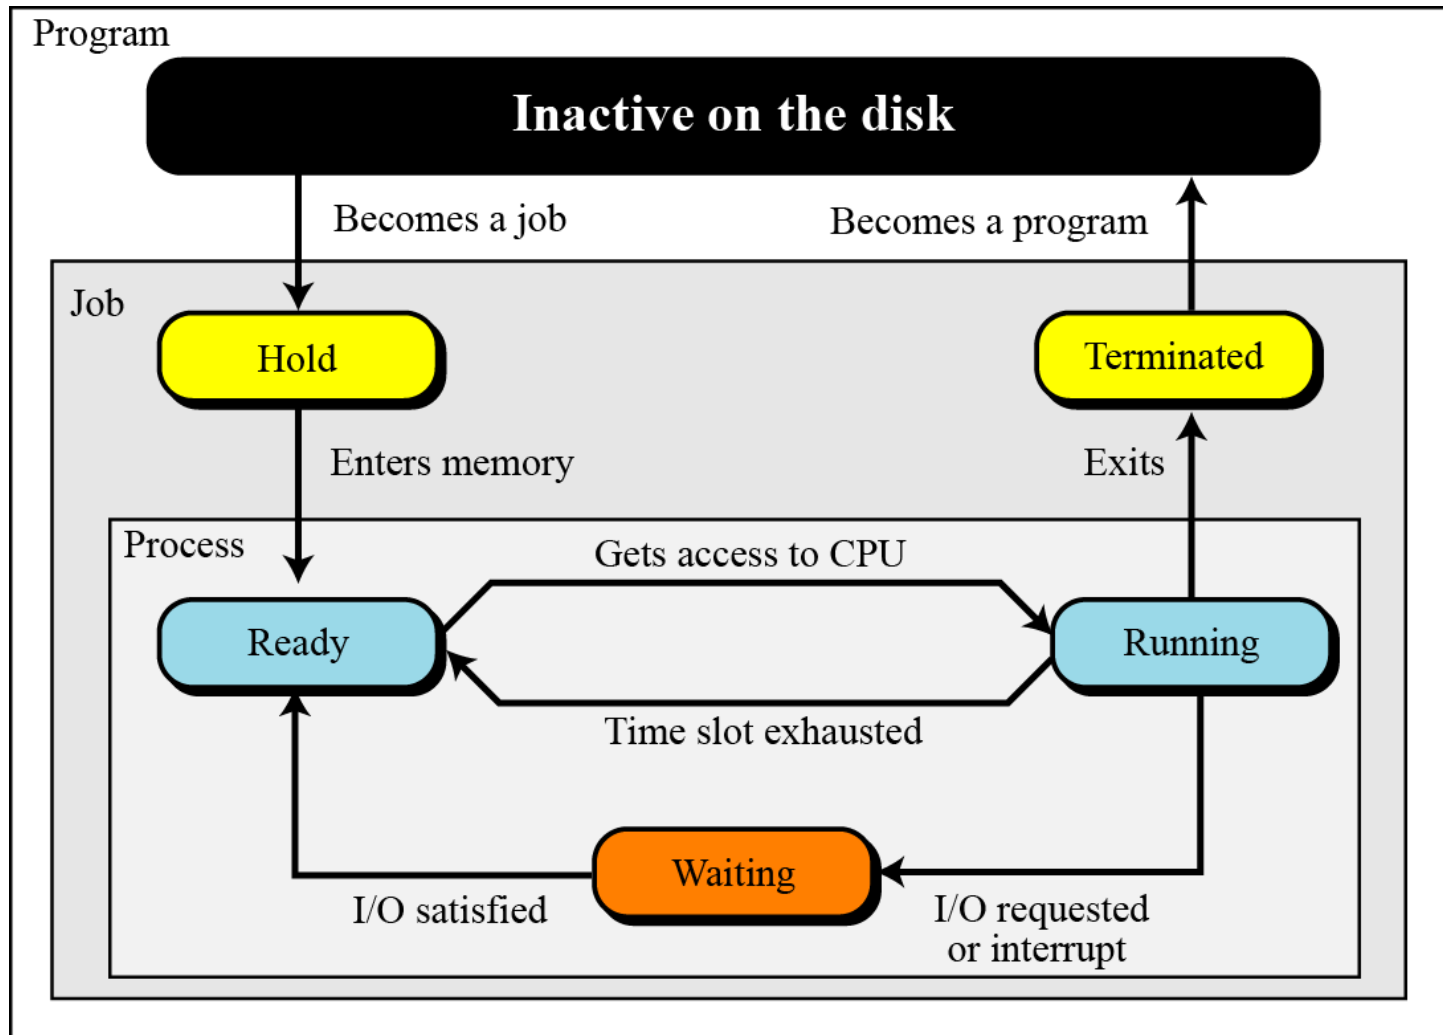


Figure 7.12 State diagram with boundaries between program, job and process

Schedulers

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

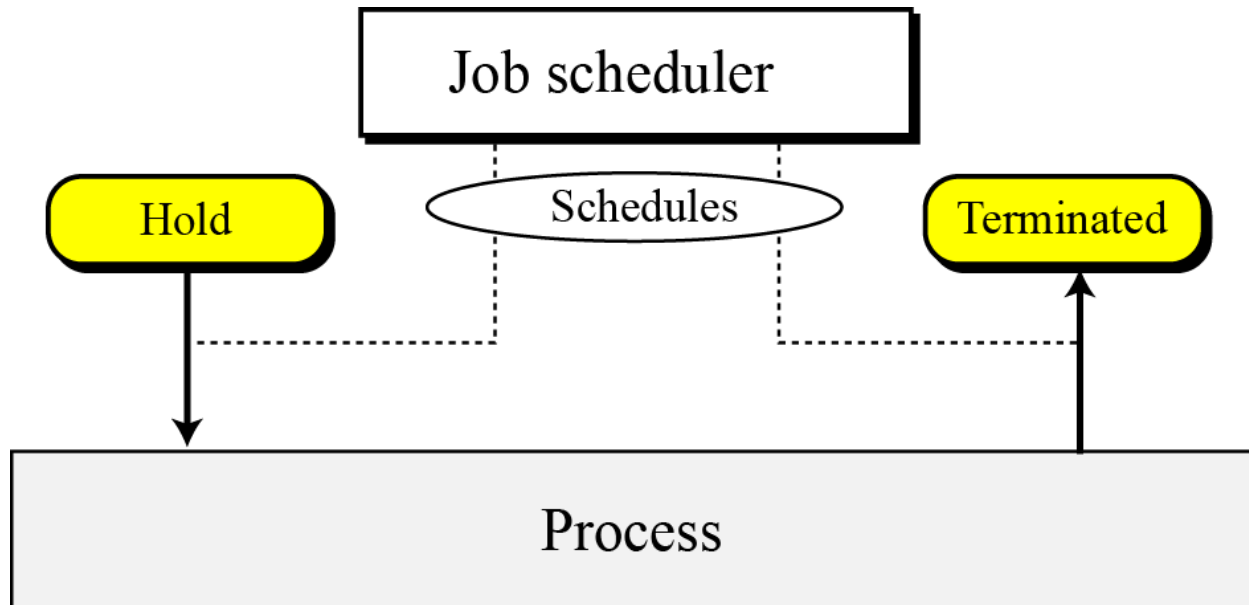


Figure 7.13 Job scheduler

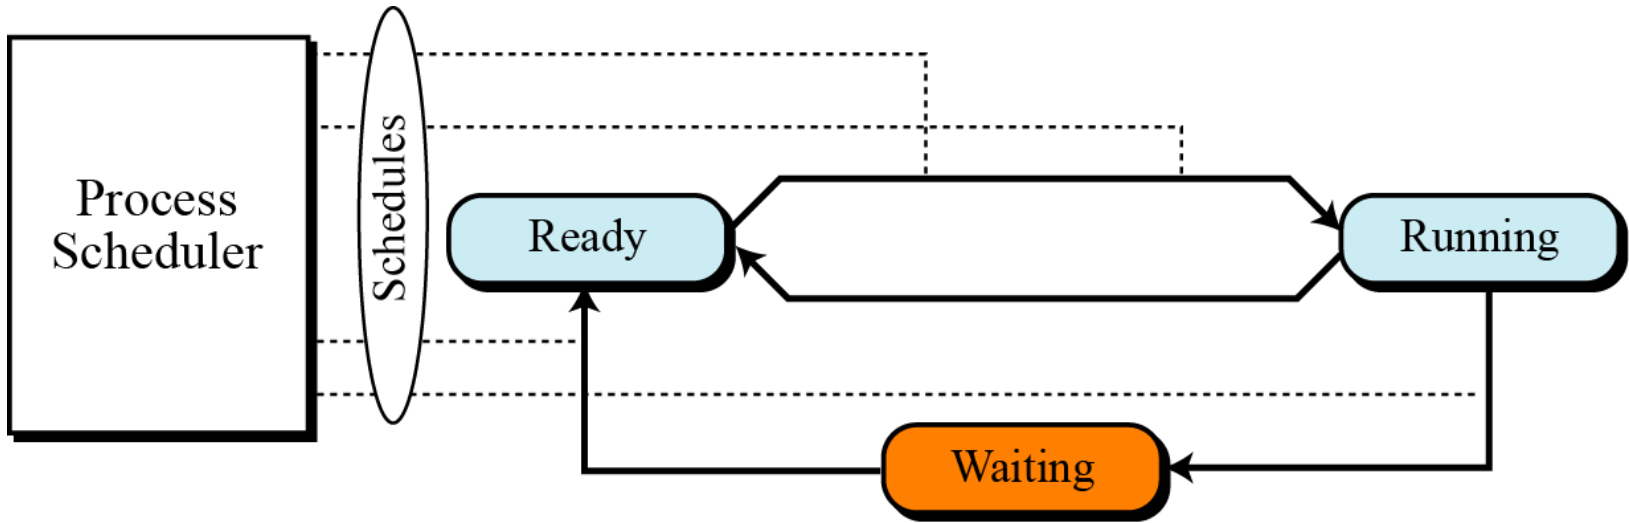


Figure 7.14 Process scheduler

Queuing

In reality, there are many jobs and many processes competing with each other for computer resources.

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.

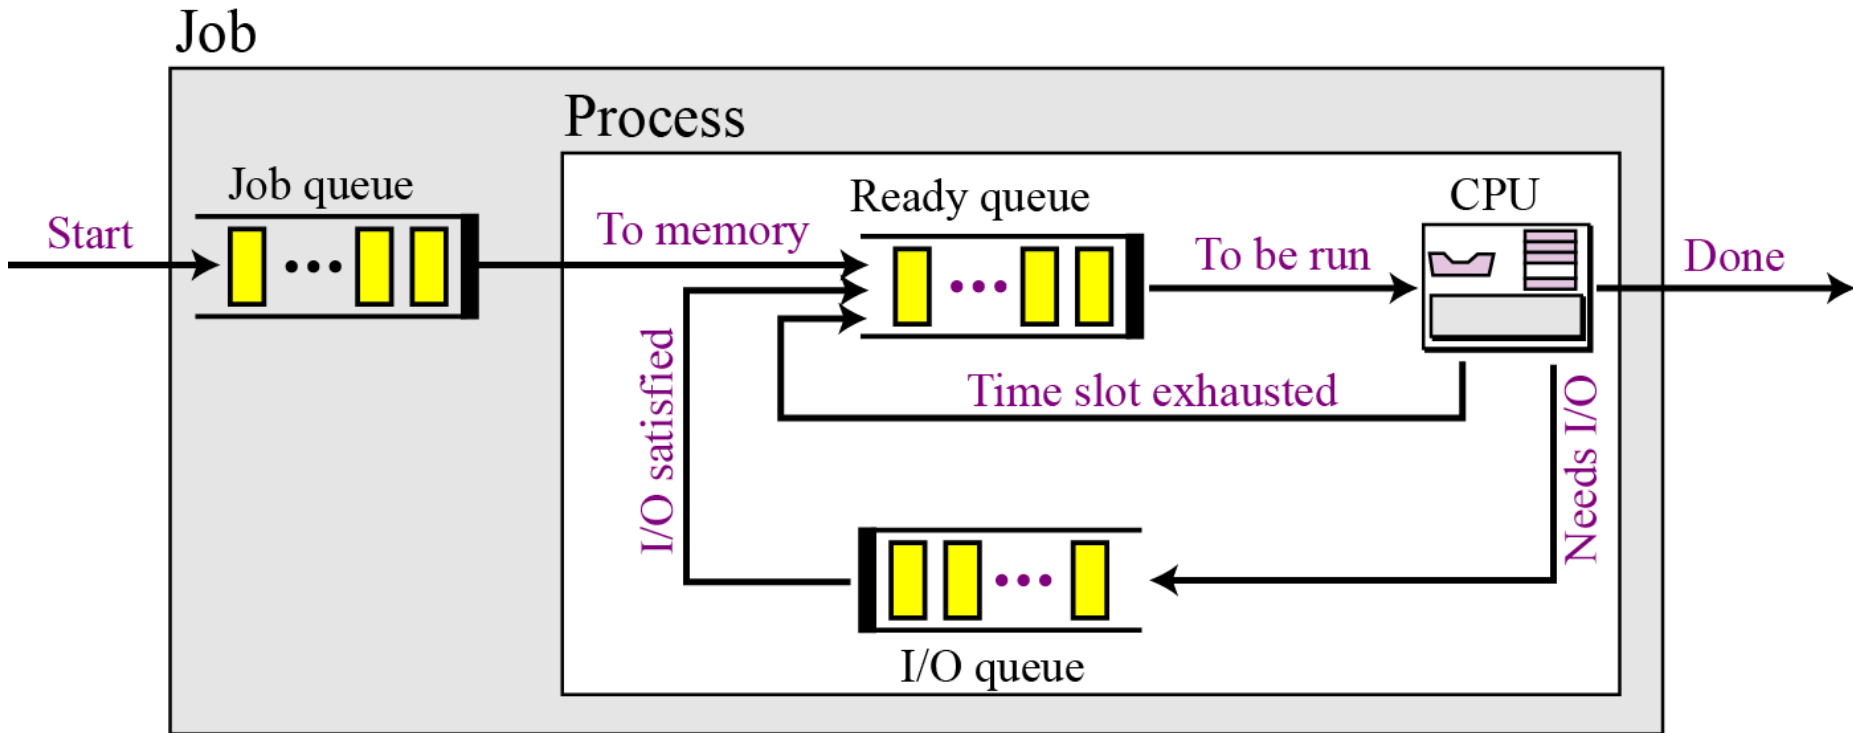


Figure 7.15 Queues for process management

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:

1. Deadlock
2. Starvation

Deadlocks

- A set of process is in a deadlock state when every process in the set is waiting for an event that can be caused by only another process in the set.
- A System Model
 - Resources:
 - Physical Resources, e.g., CPU, printers, memory, etc.
 - Logical Resources, e.g., files, semaphores, etc.




Figure 7.16 Deadlock



Deadlock occurs when the operating system does not put resource restrictions on processes.


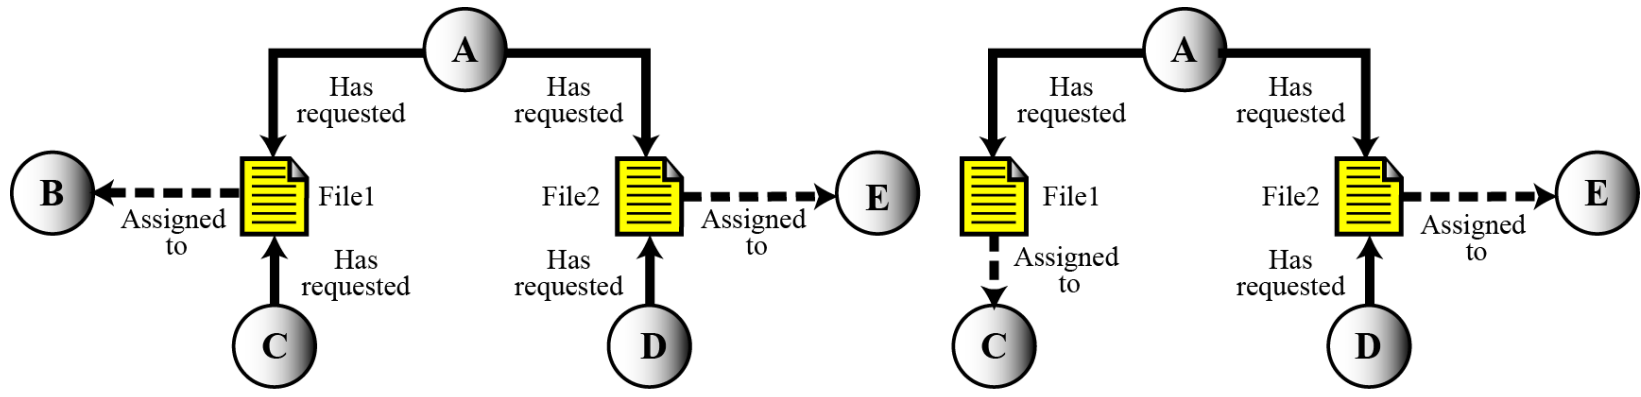



Figure 7.17 Deadlock on a bridge



a. Process A needs both File1 and File2

b. Process A still needs both File1 and File2

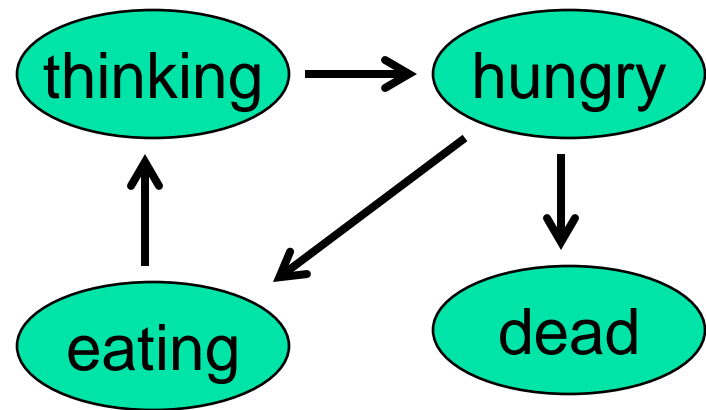
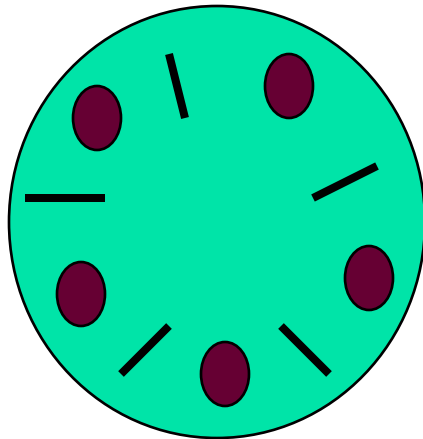


c. Process A still needs both File1 and File2 (starving)

Figure 7.18 Starvation

Classical Synchronization Problems – Dining-Philosophers

- Each philosopher must pick up one chopstick beside him/her at a time
- When two chopsticks are picked up, the philosopher can eat.



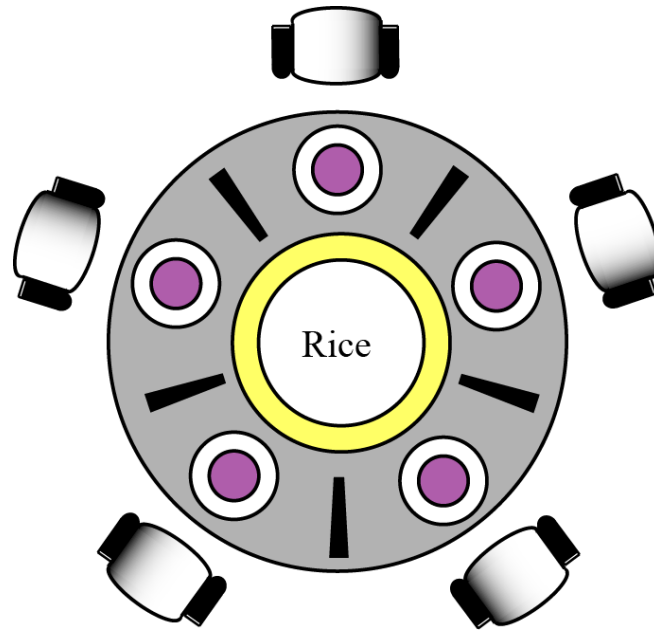


Figure 7.19 The dining philosophers problem



Starvation is the opposite of deadlock. It can happen when the operating system puts too many resource restrictions on a process.

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 monitors every input/output device constantly to ensure that the device is functioning properly.**
- ❑ The device manager maintains a queue for each input/output device or one or more queues for similar input/output devices.**
- ❑ The device manager controls the different policies for accessing input/output devices.**

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 creation, deletion, and modification of files.**
- controls the naming of files.**
- supervises the storage of files.**
- is responsible for archiving and backups.**

7-4 A SURVEY OF OPERATING SYSTEMS

We have chosen three operating systems that are familiar to most computer users:

1. UNIX
2. Linux
3. Windows

UNIX

UNIX was originally developed in 1969 by Thomson and Ritchie of the Computer Science Research Group at Bell Laboratories.

UNIX has gone through many versions since then. It has been a popular operating system among computer programmers and computer scientists.



UNIX is a multiuser, multiprocessing, portable operating system.

It is designed to facilitate programming, text processing and communication.

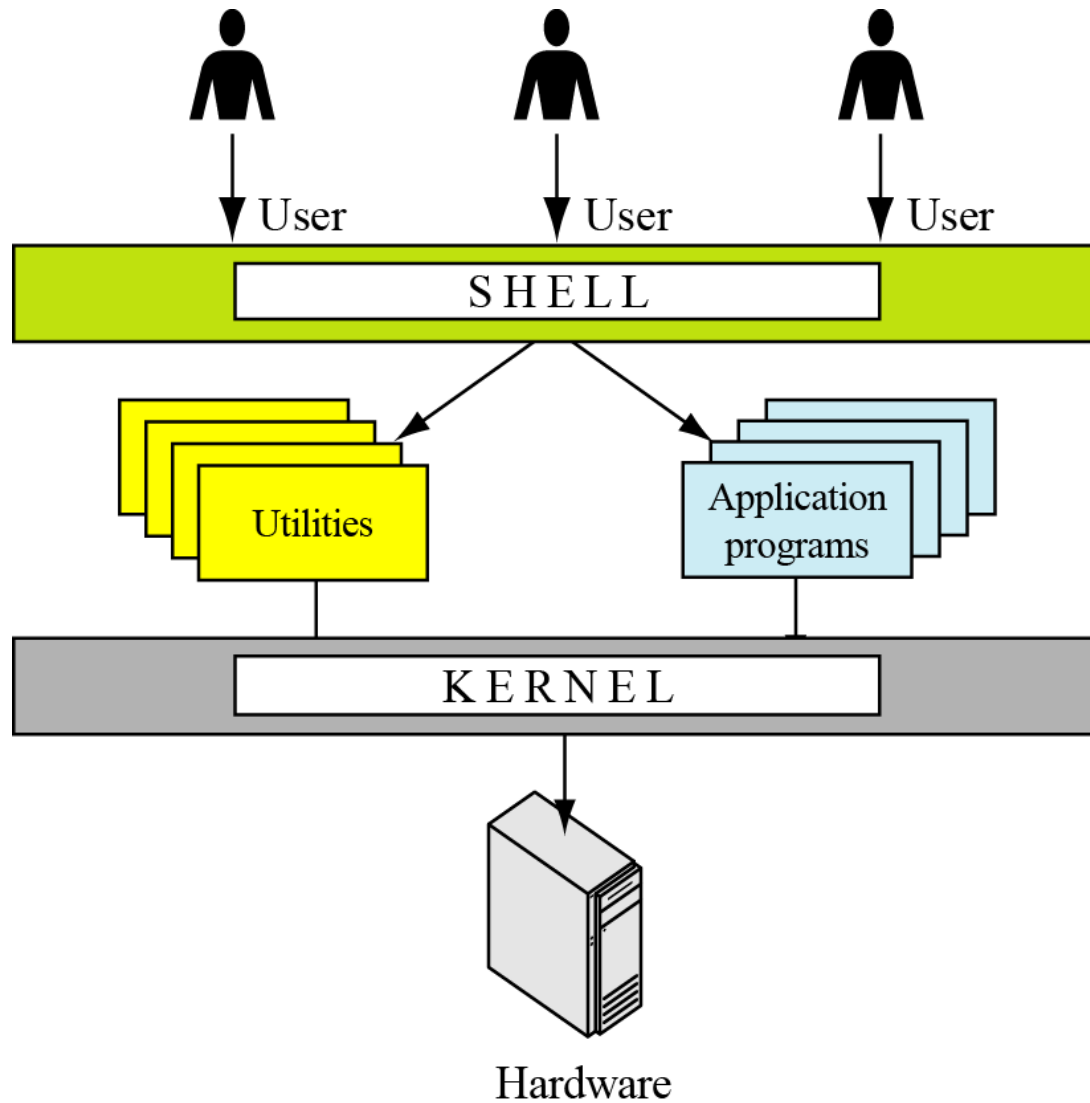


Figure 7.20 Components of the UNIX operating system

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 full-scale 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). **Windows** was the result. Several versions of Windows followed. We refer to all of these versions as Windows.

The architecture of Windows

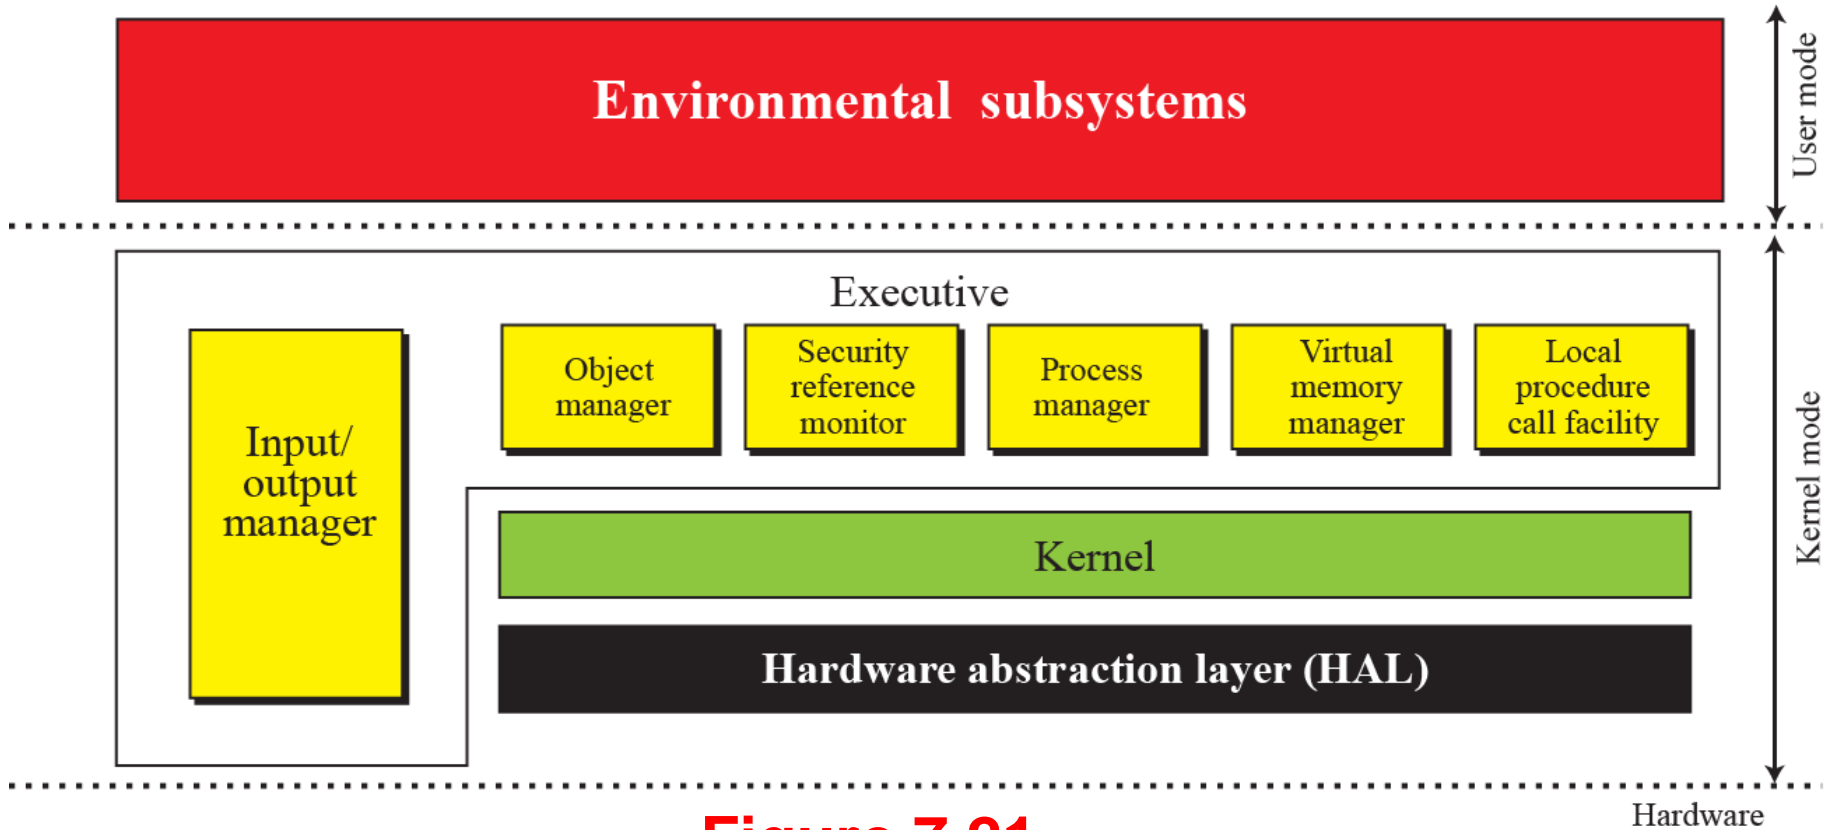


Figure 7.21