

*Operating
Systems:
Internals
and Design
Principles*

Chapter 3 Process Description and Control

Seventh Edition
By William Stallings

OS Management of Application Execution

- Resources are made available to multiple applications
- The processor is switched among multiple applications so all will appear to be progressing
- The processor and I/O devices can be used efficiently

Process Elements

- Two essential elements of a process are:

Program code

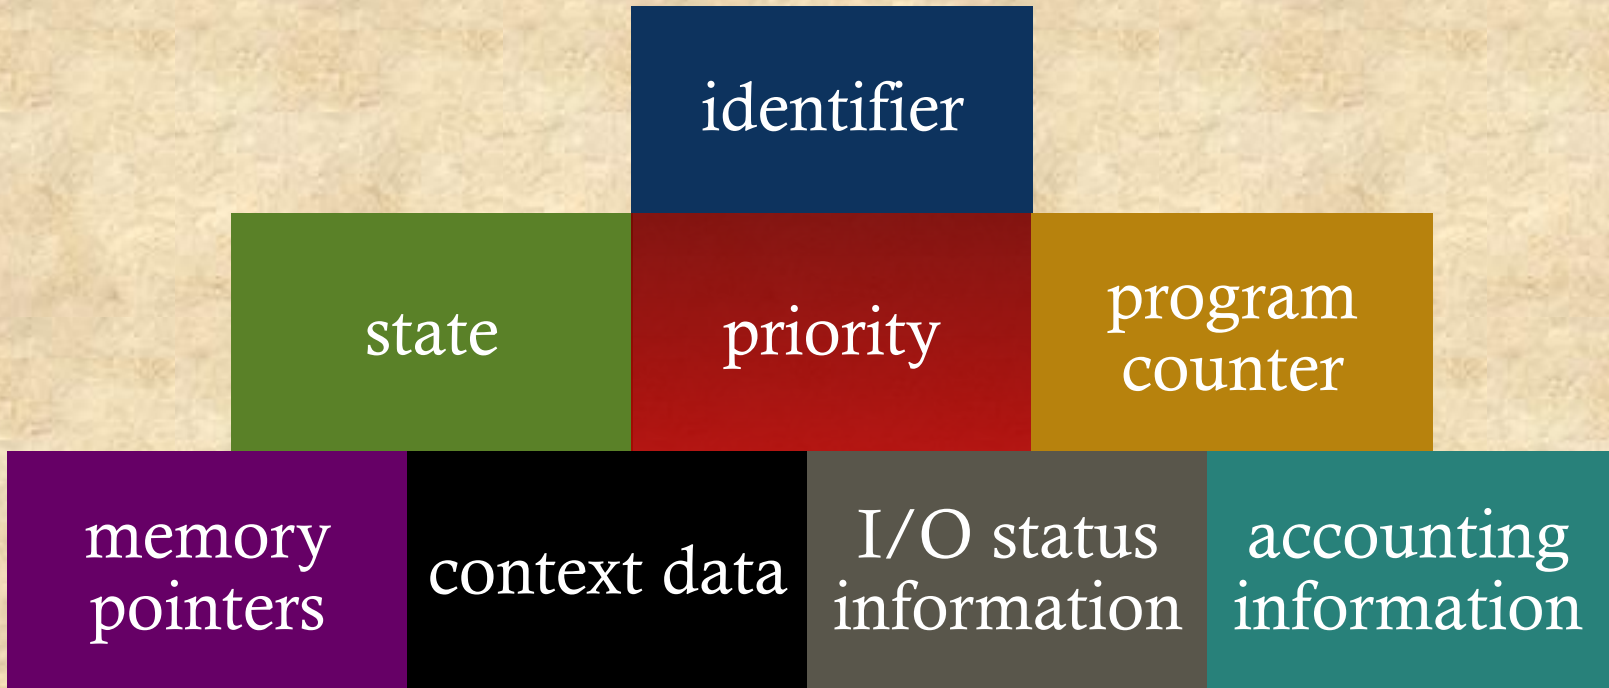
- which may be shared with other processes that are executing the same program

A set of data associated with that code

- When the processor begins to execute the program code, we refer to this executing entity as a *process*

Process Elements

- While the program is executing, this process can be uniquely characterized by a number of elements, including:



Process Control Block

- Contains the process elements
- Makes it possible to interrupt a running process and later resume execution as if the interruption had not occurred
- Created and managed by the operating system
- Key tool that allows support for multiple processes (multiprogramming)

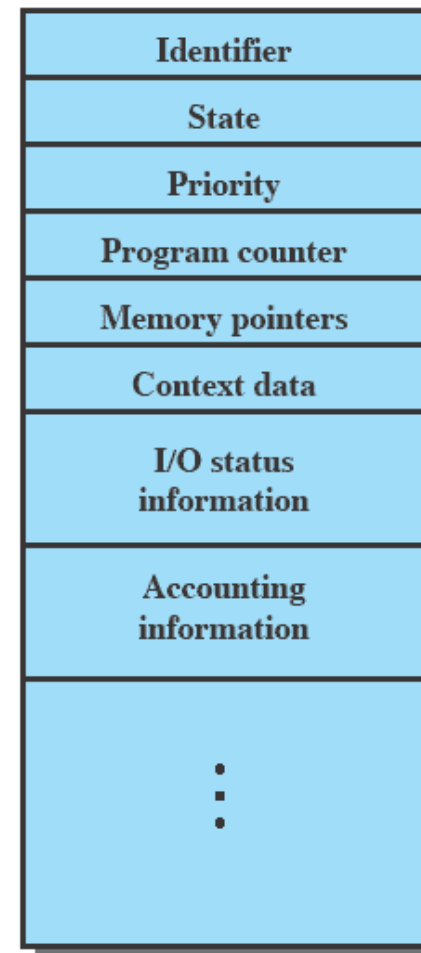
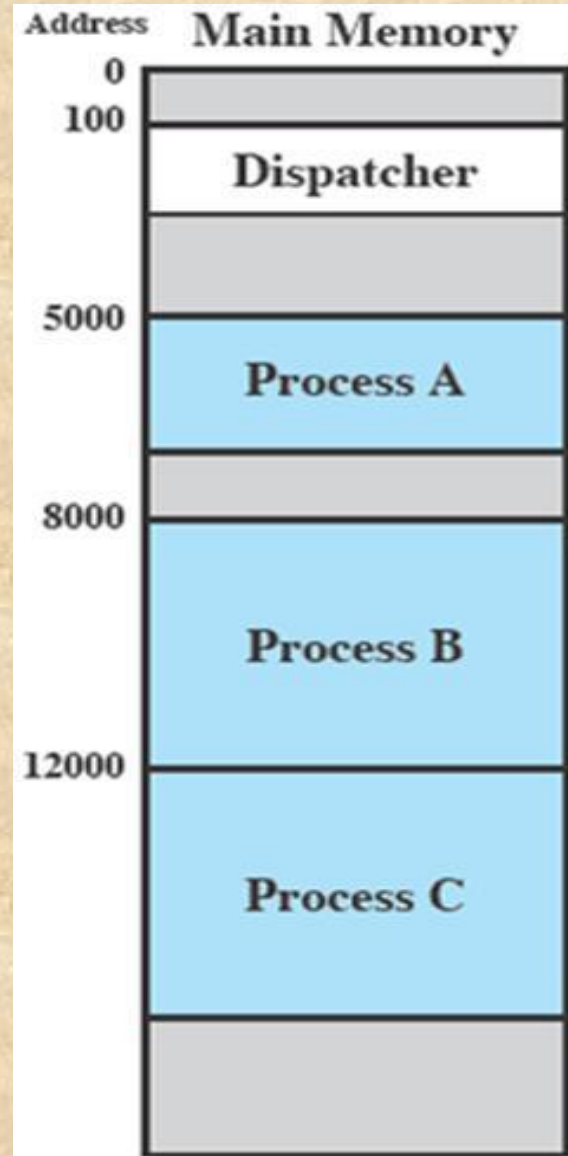


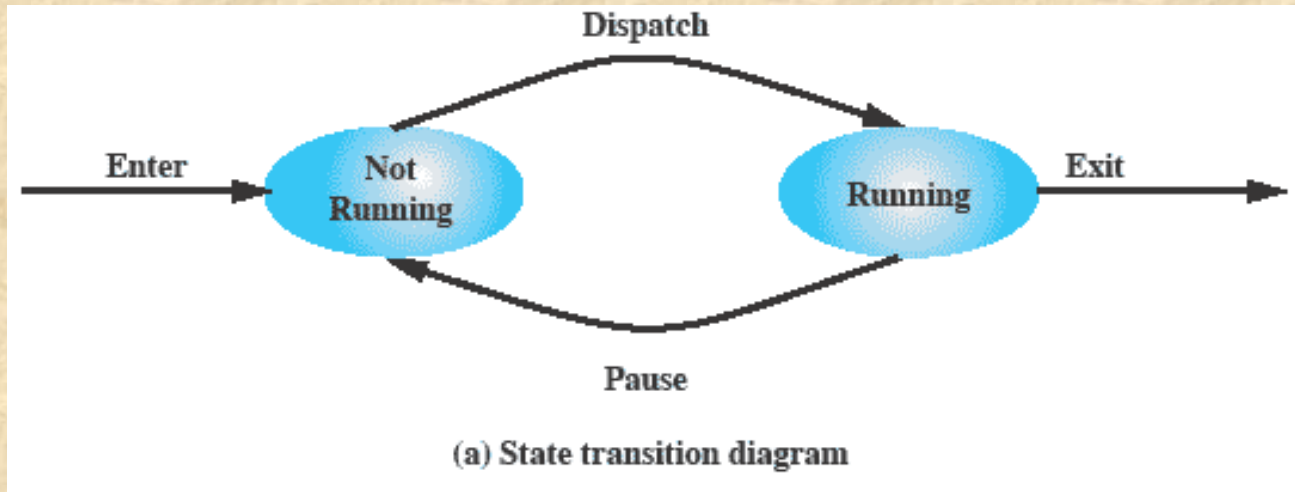
Figure 3.1 Simplified Process Control Block

Process Execution



Two-State Process Model

- A process may be in one of two states:
 - running
 - not-running



Queuing Diagram

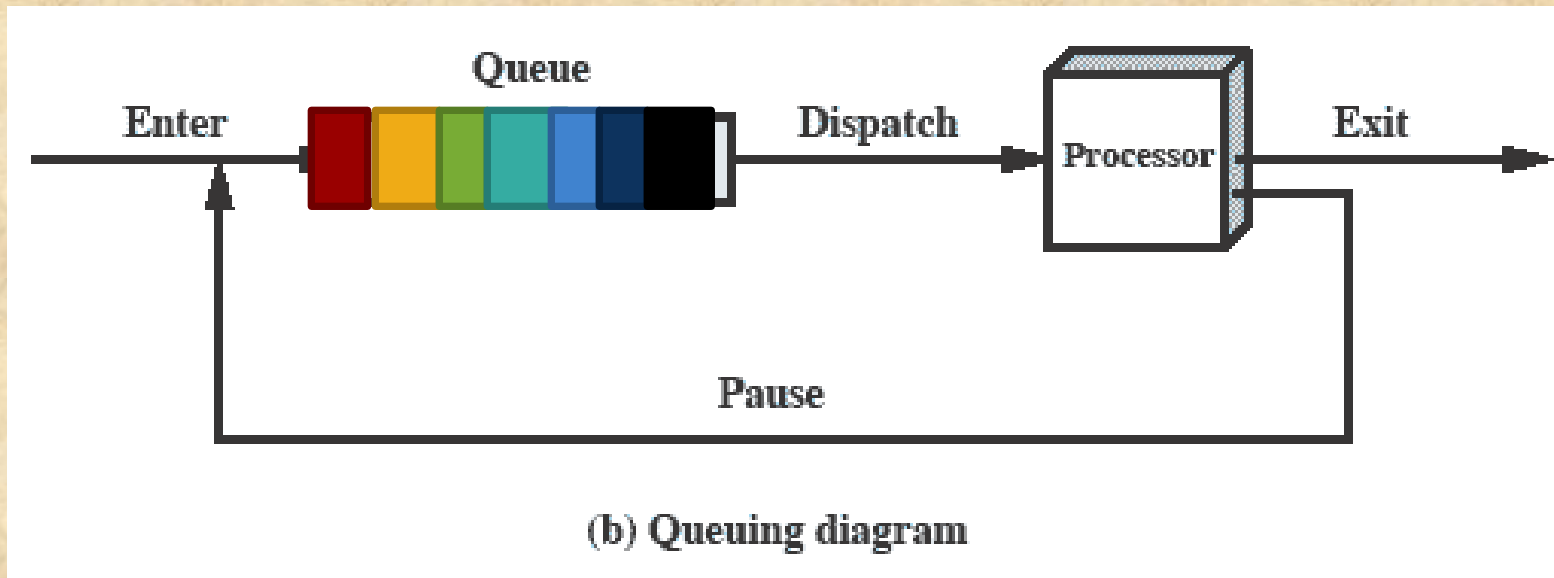
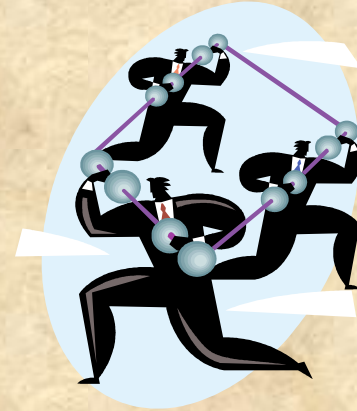


Table 3.1 Reasons for Process Creation

New batch job	The OS is provided with a batch job control stream, usually on tape or disk. When the OS is prepared to take on new work, it will read the next sequence of job control commands.
Interactive logon	A user at a terminal logs on to the system.
Created by OS to provide a service	The OS can create a process to perform a function on behalf of a user program, without the user having to wait (e.g., a process to control printing).
Spawned by existing process	For purposes of modularity or to exploit parallelism, a user program can dictate the creation of a number of processes.

Process Creation

Process spawning

- when the OS creates a process at the explicit request of another process

Parent process

- is the original, creating, process

Child process

- is the new process

Process Termination

- There must be a means for a process to indicate its completion
- A batch job should include a HALT instruction or an explicit OS service call for termination
- For an interactive application, the action of the user will indicate when the process is completed (e.g. log off, quitting an application)



Table 3.2

Reasons for Process Termination



Normal completion	The process executes an OS service call to indicate that it has completed running.
Time limit exceeded	The process has run longer than the specified total time limit. There are a number of possibilities for the type of time that is measured. These include total elapsed time ("wall clock time"), amount of time spent executing, and, in the case of an interactive process, the amount of time since the user last provided any input.
Memory unavailable	The process requires more memory than the system can provide.
Bounds violation	The process tries to access a memory location that it is not allowed to access.
Protection error	The process attempts to use a resource such as a file that it is not allowed to use, or it tries to use it in an improper fashion, such as writing to a read-only file.
Arithmetic error	The process tries a prohibited computation, such as division by zero, or tries to store numbers larger than the hardware can accommodate.
Time overrun	The process has waited longer than a specified maximum for a certain event to occur.
I/O failure	An error occurs during input or output, such as inability to find a file, failure to read or write after a specified maximum number of tries (when, for example, a defective area is encountered on a tape), or invalid operation (such as reading from the line printer).
Invalid instruction	The process attempts to execute a nonexistent instruction (often a result of branching into a data area and attempting to execute the data).
Privileged instruction	The process attempts to use an instruction reserved for the operating system.
Data misuse	A piece of data is of the wrong type or is not initialized.
Operator or OS intervention	For some reason, the operator or the operating system has terminated the process (e.g., if a deadlock exists).
Parent termination	When a parent terminates, the operating system may automatically terminate all of the offspring of that parent.
Parent request	A parent process typically has the authority to terminate any of its offspring.

Five-State Process Model

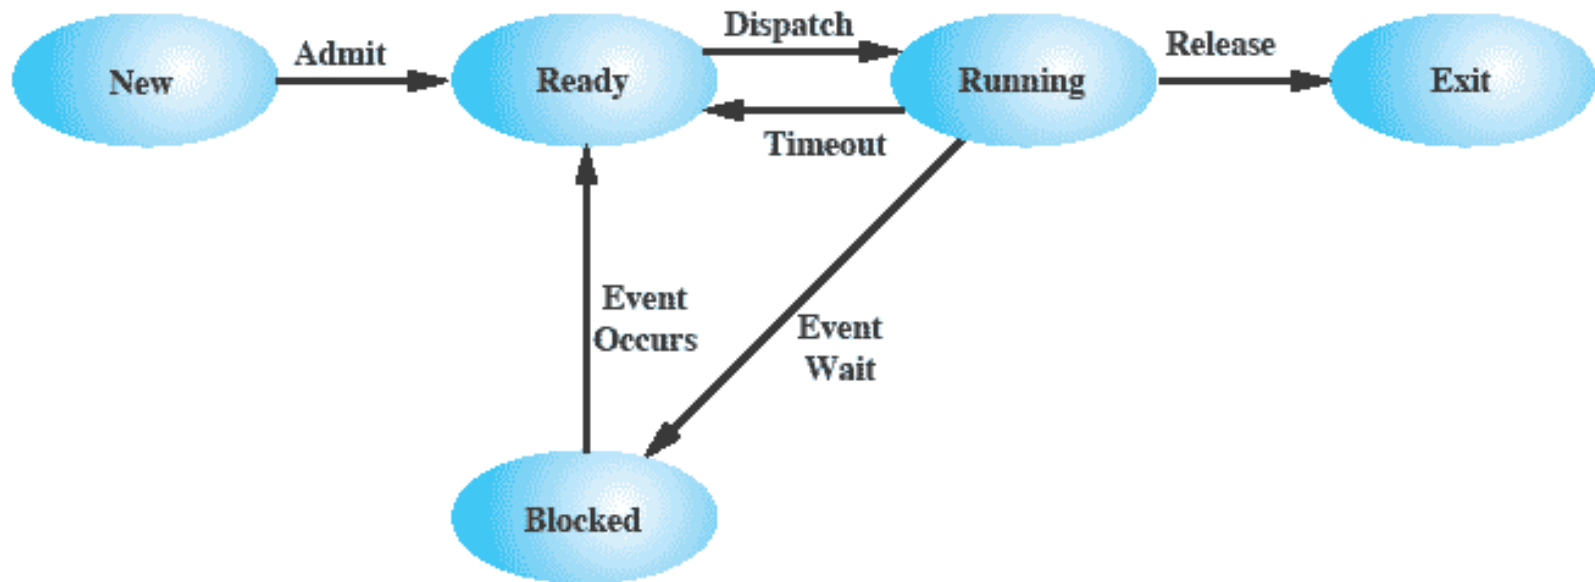


Figure 3.6 Five-State Process Model

Process States example

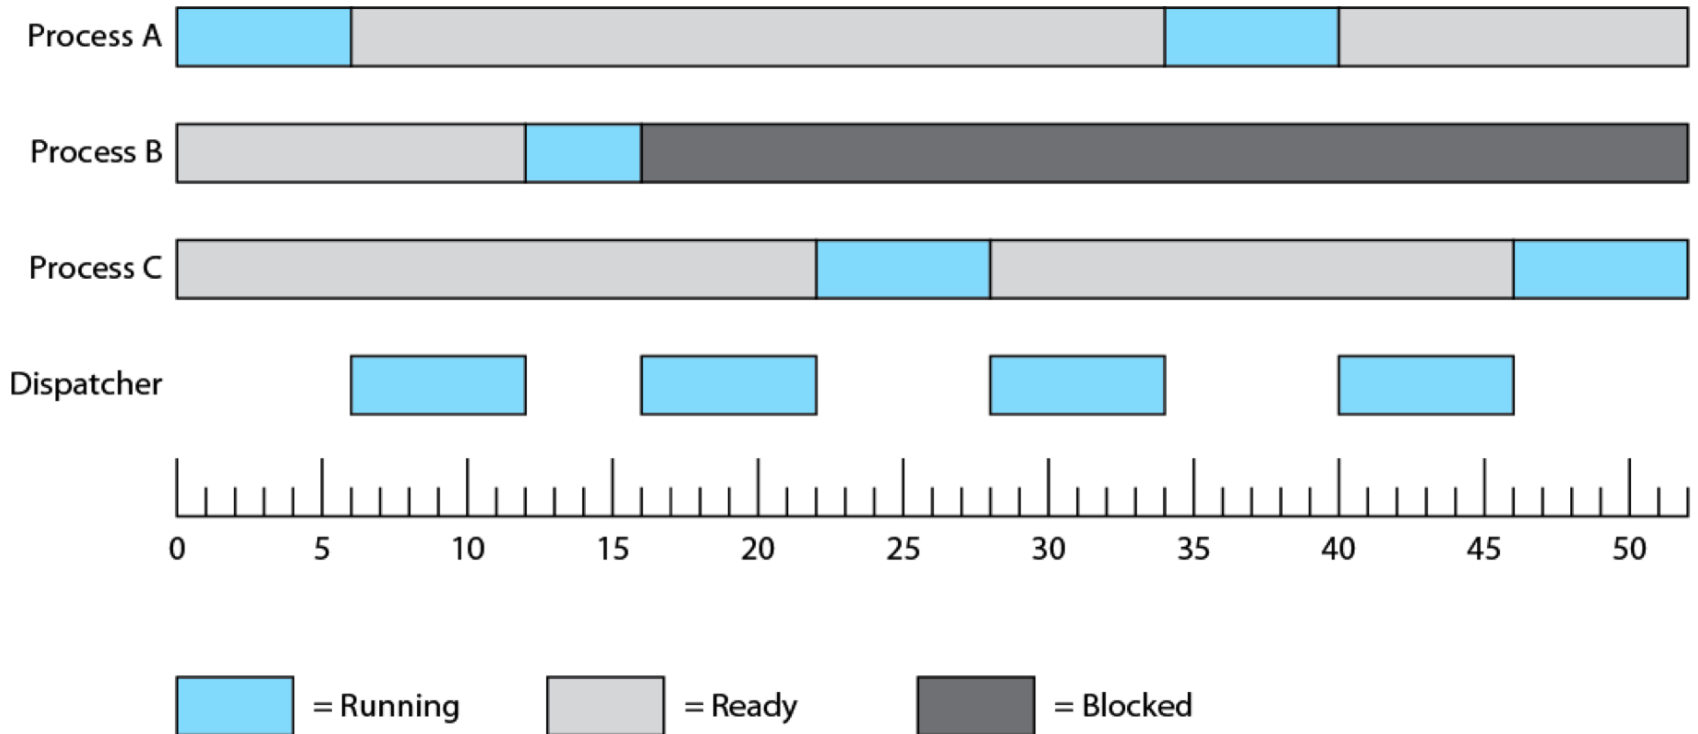
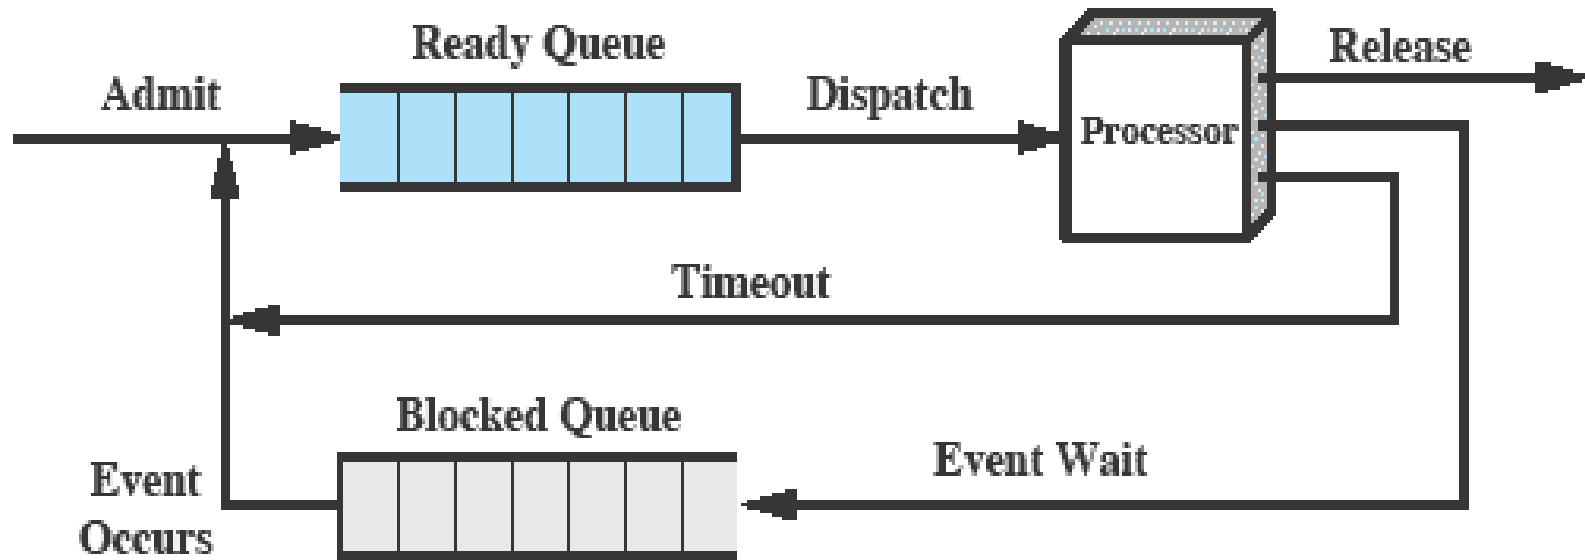


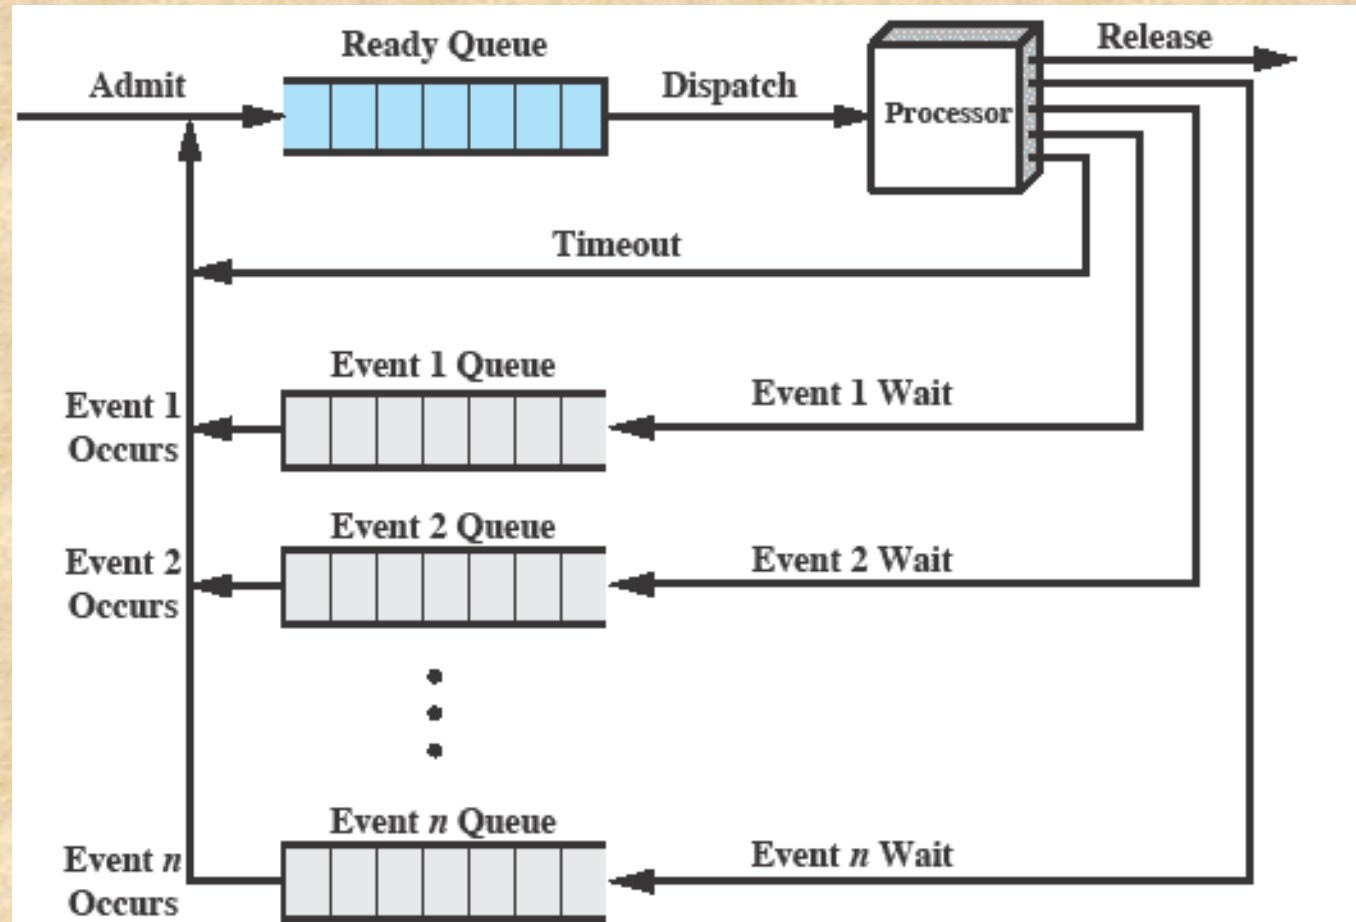

Figure 3.7 Process States for Trace of Figure 3.4

Using Two Queues



(a) Single blocked queue

Multiple Blocked Queues



(b) Multiple blocked queues

Processes and Resources

What does the OS need to know to manage processes & resources?

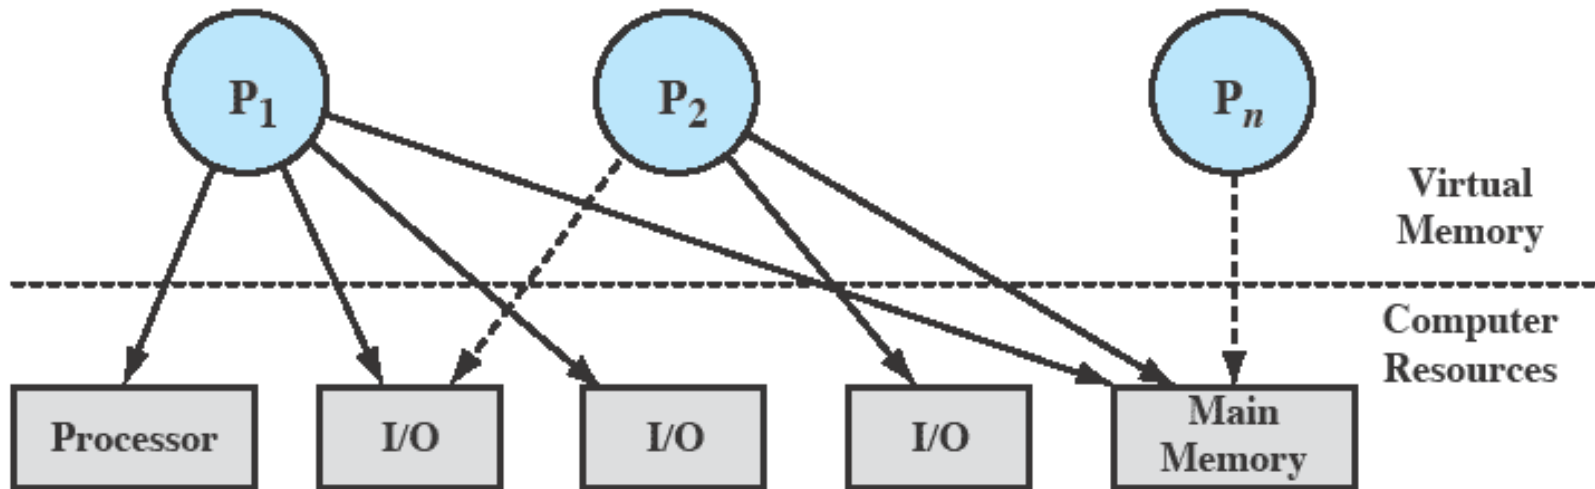


Figure 3.10 Processes and Resources (resource allocation at one snapshot in time)

OS Control Tables

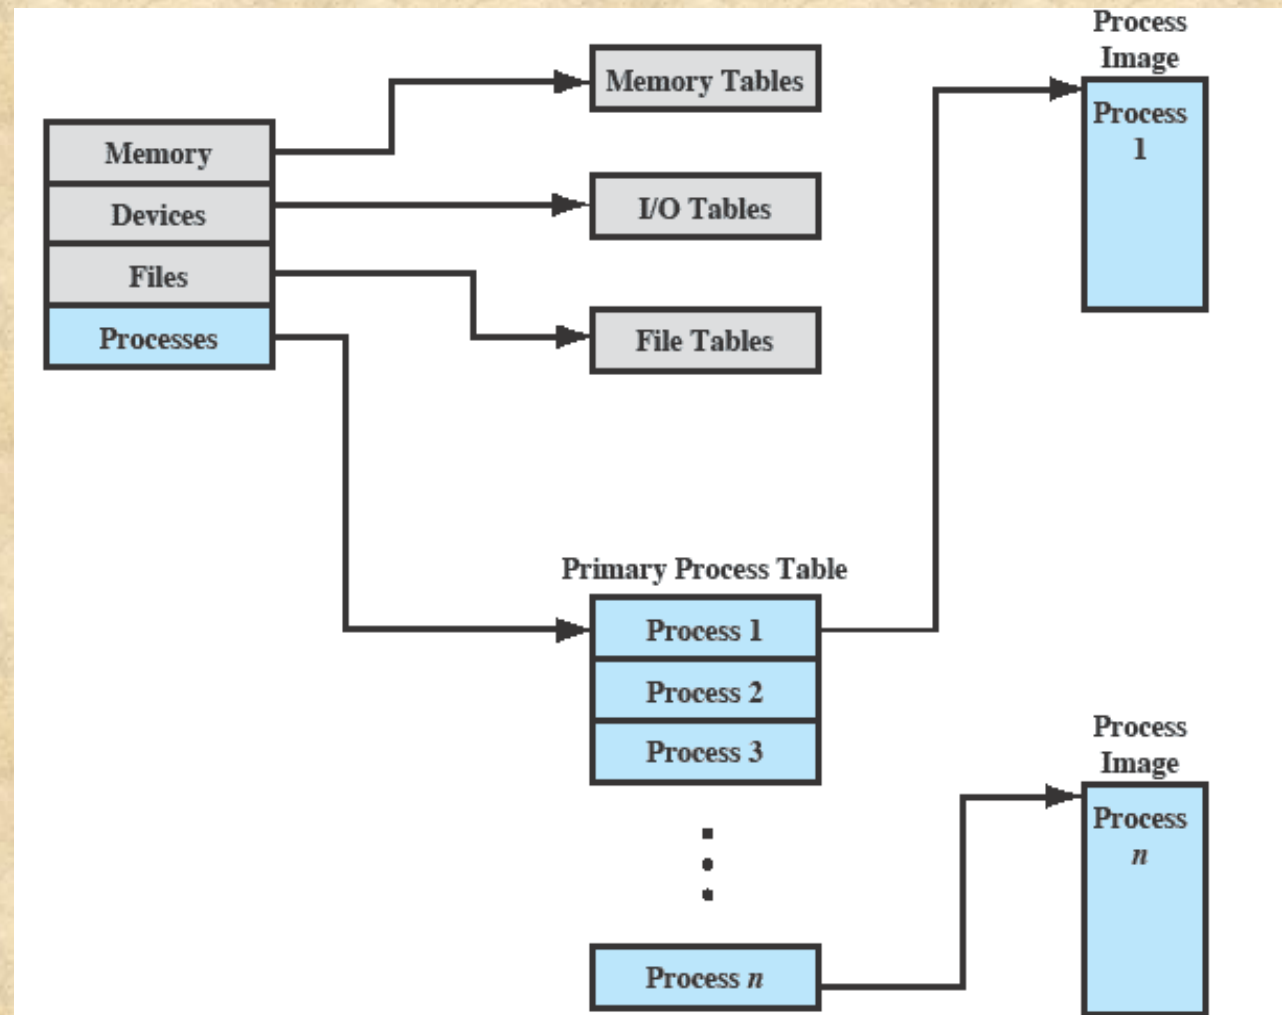


Figure 3.11 General Structure of Operating System Control Tables

Memory Tables

- Used to keep track of both main (real) and secondary (virtual) memory
- Processes are maintained on secondary memory using some sort of virtual memory or simple swapping mechanism

Must include:

allocation of main memory to processes

allocation of secondary memory to processes

protection attributes of blocks of main or virtual memory

information needed to manage virtual memory

I/O Tables

- Used by the OS to manage the I/O devices and channels of the computer system
- At any given time, an I/O device may be available or assigned to a particular process

If an I/O operation is in progress, the OS needs to know:

the status of the I/O operation

the location in main memory being used as the source or destination of the I/O transfer

File Tables

These tables provide information about:

- Information may be maintained and used by a file management system, or the OS itself
- existence of files
- location on secondary memory
- current status
- other attributes

Process Tables

- Must be maintained to manage processes
- Process tables must have some reference to memory, I/O, and file tables
 - e.g., pointers from each individual process entry to its data in the other three sets of tables.
- The tables themselves must be accessible by the OS and therefore are subject to memory management
- Processes are represented by a *process image*

Process Control Structures - PCB

Process Components

- A process must include a program or set of programs to be executed
- A process will consist of at least sufficient memory to hold the program(s) and data of that process
- The execution of a program typically involves a stack that is used to keep track of procedure calls and parameter passing between procedures

Process Attributes

- Each process has associated with it a number of attributes that are used by the OS for process control –maintained in a Process Control Block (PCB)
- The collection of program, data, stack, and attributes is referred to as the process image
- Process image location will depend on the memory management scheme being used

PCB elements

Process Identification

Identifiers

Numeric identifiers that may be stored with the process control block include

- Identifier of this process
- Identifier of the process that created this process (parent process)
- User identifier

Processor State Information

User-Visible Registers

A user-visible register is one that may be referenced by means of the machine language that the processor executes while in user mode. Typically, there are from 8 to 32 of these registers, although some RISC implementations have over 100.

Control and Status Registers

These are a variety of processor registers that are employed to control the operation of the processor. These include

- Program counter:** Contains the address of the next instruction to be fetched
- Condition codes:** Result of the most recent arithmetic or logical operation (e.g., sign, zero, carry, equal, overflow)
- Status information:** Includes interrupt enabled/disabled flags, execution mode

Stack Pointers

Each process has one or more last-in-first-out (LIFO) system stacks associated with it. A stack is used to store parameters and calling addresses for procedure and system calls. The stack pointer points to the top of the stack.

PCB elements (continuation)



Process Control Information

Scheduling and State Information

This is information that is needed by the operating system to perform its scheduling function. Typical items of information:

- Process state:** Defines the readiness of the process to be scheduled for execution (e.g., running, ready, waiting, halted).
- Priority:** One or more fields may be used to describe the scheduling priority of the process. In some systems, several values are required (e.g., default, current, highest-allowable)
- Scheduling-related information:** This will depend on the scheduling algorithm used. Examples are the amount of time that the process has been waiting and the amount of time that the process executed the last time it was running.
- Event:** Identity of event the process is awaiting before it can be resumed.

Data Structuring

A process may be linked to other process in a queue, ring, or some other structure. For example, all processes in a waiting state for a particular priority level may be linked in a queue. A process may exhibit a parent-child (creator-created) relationship with another process. The process control block may contain pointers to other processes to support these structures.

Interprocess Communication

Various flags, signals, and messages may be associated with communication between two independent processes. Some or all of this information may be maintained in the process control block.

Process Privileges

Processes are granted privileges in terms of the memory that may be accessed and the types of instructions that may be executed. In addition, privileges may apply to the use of system utilities and services.

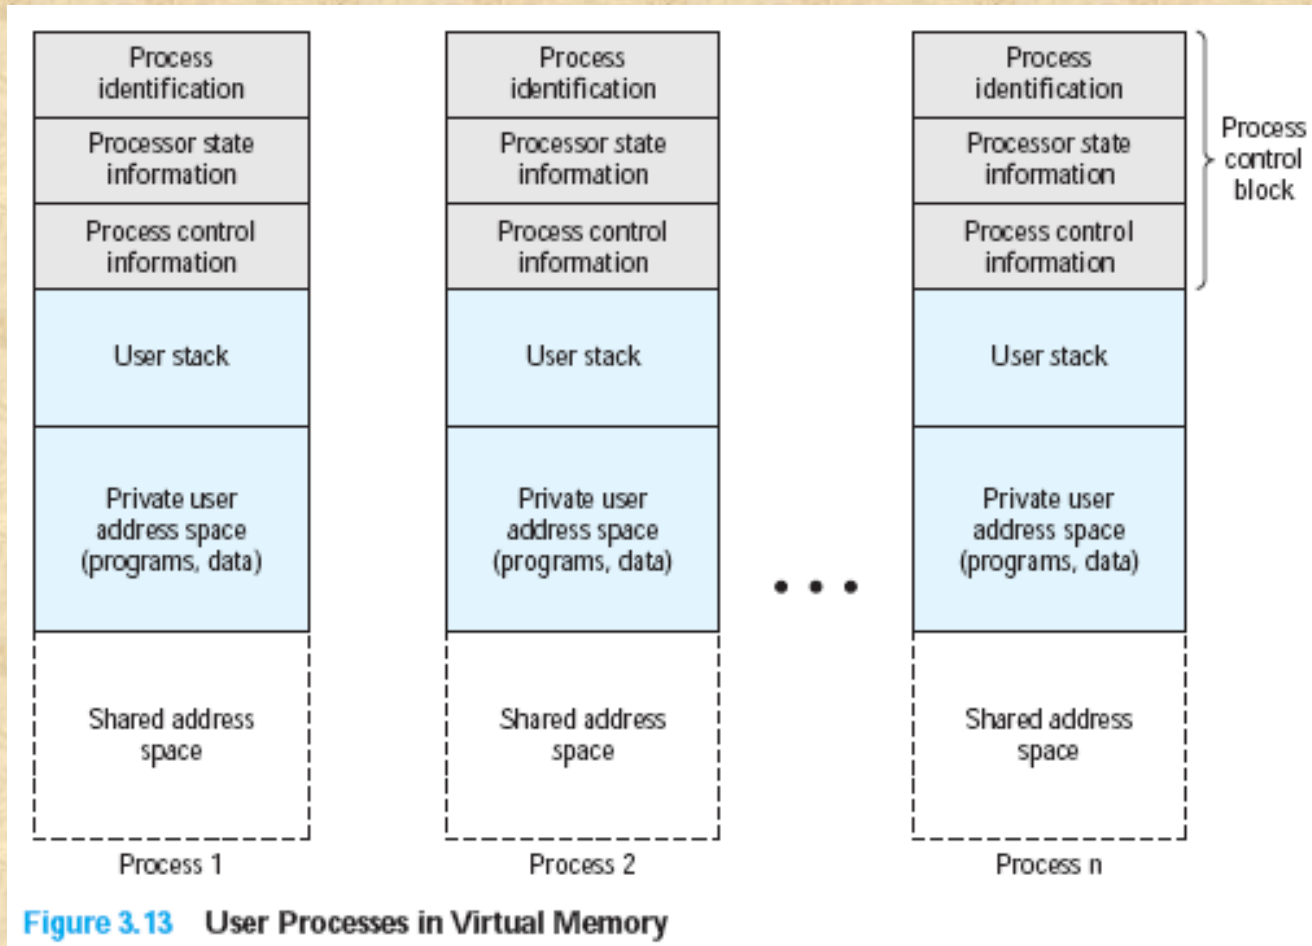

Memory Management

This section may include pointers to segment and/or page tables that describe the virtual memory assigned to this process.

Resource Ownership and Utilization

Resources controlled by the process may be indicated, such as opened files. A history of utilization of the processor or other resources may also be included; this information may be needed by the scheduler.

Structure of Process Images in Virtual Memory



Process List Structures

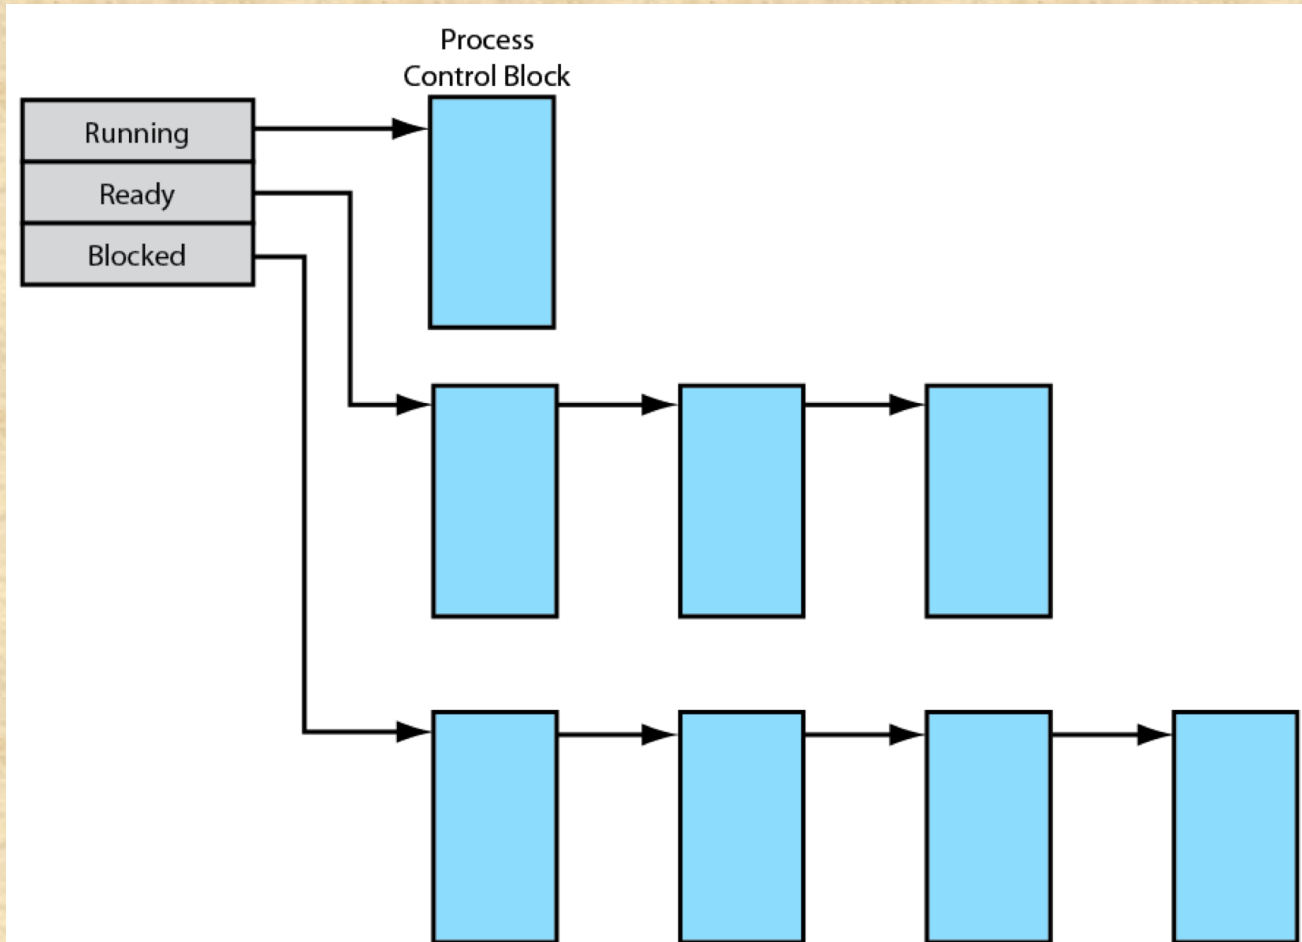


Figure 3.14 Process List Structures

Role of the Process Control Block

- The most important data structure in an OS
 - contains all of the information about a process that is needed by the OS
 - blocks are read and/or modified by virtually every module in the OS
 - defines the state of the OS
- Difficulty is not access, but protection
 - a bug in a single routine could damage process control blocks, which could destroy the system's ability to manage the affected processes
 - a design change in the structure or semantics of the process control block could affect a number of modules in the OS

Modes of Execution

User Mode

- less-privileged mode
- user programs typically execute in this mode



System Mode

- more-privileged mode
- also referred to as control mode or *kernel mode*
- kernel of the operating system




Table 3.7

Typical

Functions

of an

Operating

System

Kernel

Process Management

- Process creation and termination
- Process scheduling and dispatching
- Process switching
- Process synchronization and support for interprocess communication
- Management of process control blocks

Memory Management

- Allocation of address space to processes
- Swapping
- Page and segment management

I/O Management

- Buffer management
- Allocation of I/O channels and devices to processes

Support Functions

- Interrupt handling
- Accounting
- Monitoring

Process Creation

- Once the OS decides to create a new process it:

assigns a unique process identifier to the new process

allocates space for the process

initializes the process control block

sets the appropriate linkages

creates or expands other data structures

Process (Context) Switching

A process switch may occur any time that the OS has gained control from the currently running process. Possible events giving OS control are:

Mechanism	Cause	Use
Interrupt	External to the execution of the current instruction	Reaction to an asynchronous external event
Trap	Associated with the execution of the current instruction	Handling of an error or an exception condition
Supervisor call	Explicit request	Call to an operating system function

System Interrupts

Interrupt

- Due to some sort of event that is external to and independent of the currently running process
 - clock interrupt
 - I/O interrupt
 - memory fault
- Time slice
 - the maximum amount of time that a process can execute before being interrupted

Trap

- An error or exception condition generated within the currently running process
- OS determines if the condition is fatal
 - moved to the Exit state and a process switch occurs
 - action will depend on the nature of the error

Change of Process State

- The steps in a full process switch are:

save the context of the processor



update the process control block of the process currently in the Running state



move the process control block of this process to the appropriate queue



select another process for execution



update the process control block of the process selected



update memory management data structures



restore the context of the processor to that which existed at the time the selected process was last switched out

If the currently running process is to be moved to another state (Ready, Blocked, etc.), then the OS must make substantial changes in its environment

Security Issues

- An OS associates a set of privileges with each process
- Typically a process that executes on behalf of a user has the privileges that the OS recognizes for that user
- Highest level of privilege is referred to as administrator, supervisor, or root access
- A key security issue in the design of any OS is to prevent, or at least detect, attempts by a user or malware from gaining unauthorized privileges on the system and from gaining root access



System Access Threats

Intruders

- Often referred to as a hacker or cracker

Classes:

- Masquerader
- Mifeasor
- Clandestine user



- Objective is to gain access to a system or to increase the range of privileges accessible on a system
- Attempts to acquire information that should have been protected

Malicious Software

- Most sophisticated types of threats to computer systems

Categories:

- those that need a host program (parasitic)
 - viruses, logic bombs, backdoors
- those that are independent
 - worms, bots

- Can be relatively harmless or very damaging



Countermeasures: Intrusion Detection

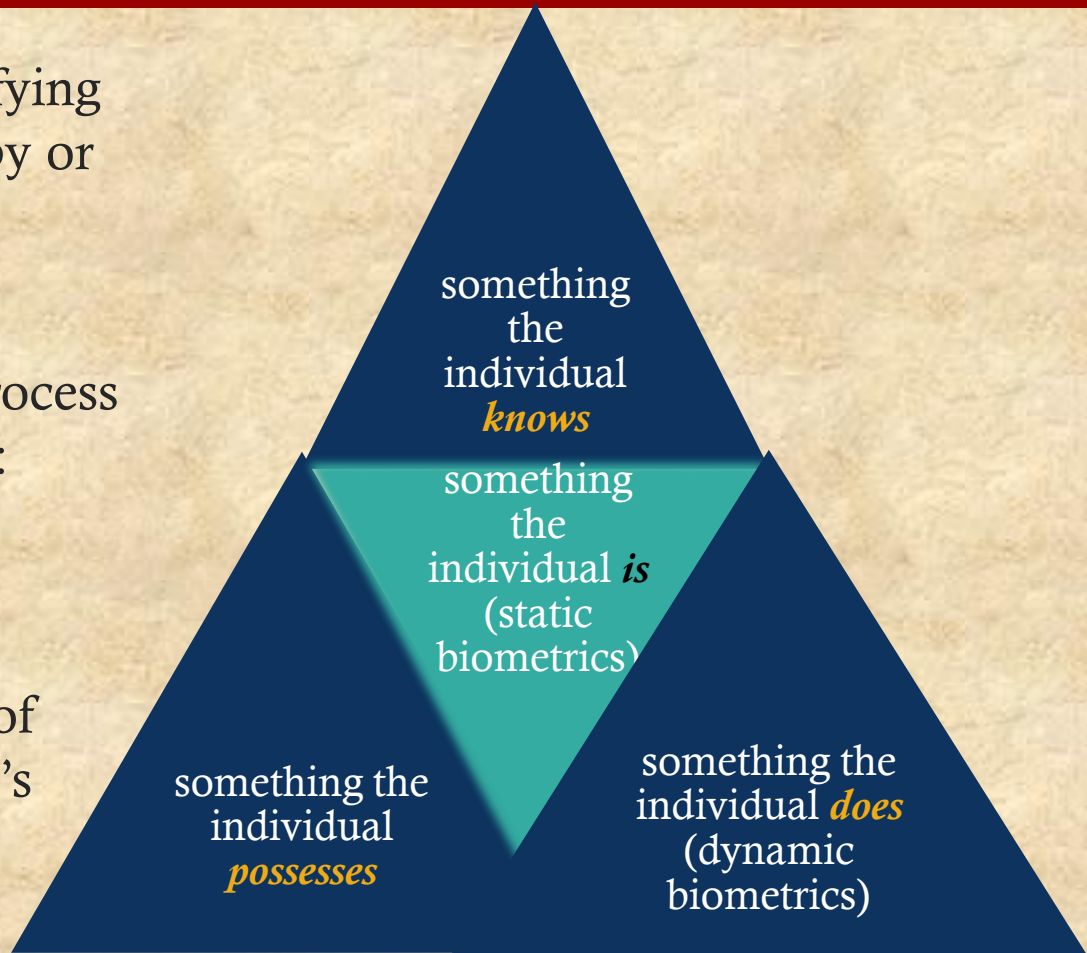
- “A security service that monitors and analyzes system events for the purpose of finding, and providing real-time or near real-time warning of, attempts to access system resources in an unauthorized manner” (RFC 2828)
- May be host or network based
- An intrusion detection system (IDS) comprises three logical components:



- IDSs are typically designed to detect human intruder behavior as well as malicious software behavior

Countermeasures: Authentication

- “The process of verifying an identity claimed by or for a system entity.” (RFC2828)
- An authentication process consists of two steps:
 - Identification
 - Verification
- Four general means of authenticating a user’s identity:



Countermeasures: Access Control

- Implements a security policy that specifies who or what may have access to each specific system resource and the type of access that is permitted in each instance
- Mediates between a user and system resources
- A security administrator maintains an authorization database
- An auditing function monitors and keeps a record of user accesses to system resources



Countermeasures: Firewalls

A dedicated
computer
that:

- interfaces with computers outside a network
- has special security precautions built into it to protect sensitive files on computers within the network

Design
goals of a
firewall:

- all traffic must pass through the firewall
- only authorized traffic will be allowed to pass
- immune to penetration



UNIX Process States

User Running	Executing in user mode.
Kernel Running	Executing in kernel mode.
Ready to Run, in Memory	Ready to run as soon as the kernel schedules it.
Asleep in Memory	Unable to execute until an event occurs; process is in main memory (a blocked state).
Ready to Run, Swapped	Process is ready to run, but the swapper must swap the process into main memory before the kernel can schedule it to execute.
Sleeping, Swapped	The process is awaiting an event and has been swapped to secondary storage (a blocked state).
Preempted	Process is returning from kernel to user mode, but the kernel preempts it and does a process switch to schedule another process.
Created	Process is newly created and not yet ready to run.
Zombie	Process no longer exists, but it leaves a record for its parent process to collect.

UNIX Process State Transition Diagram

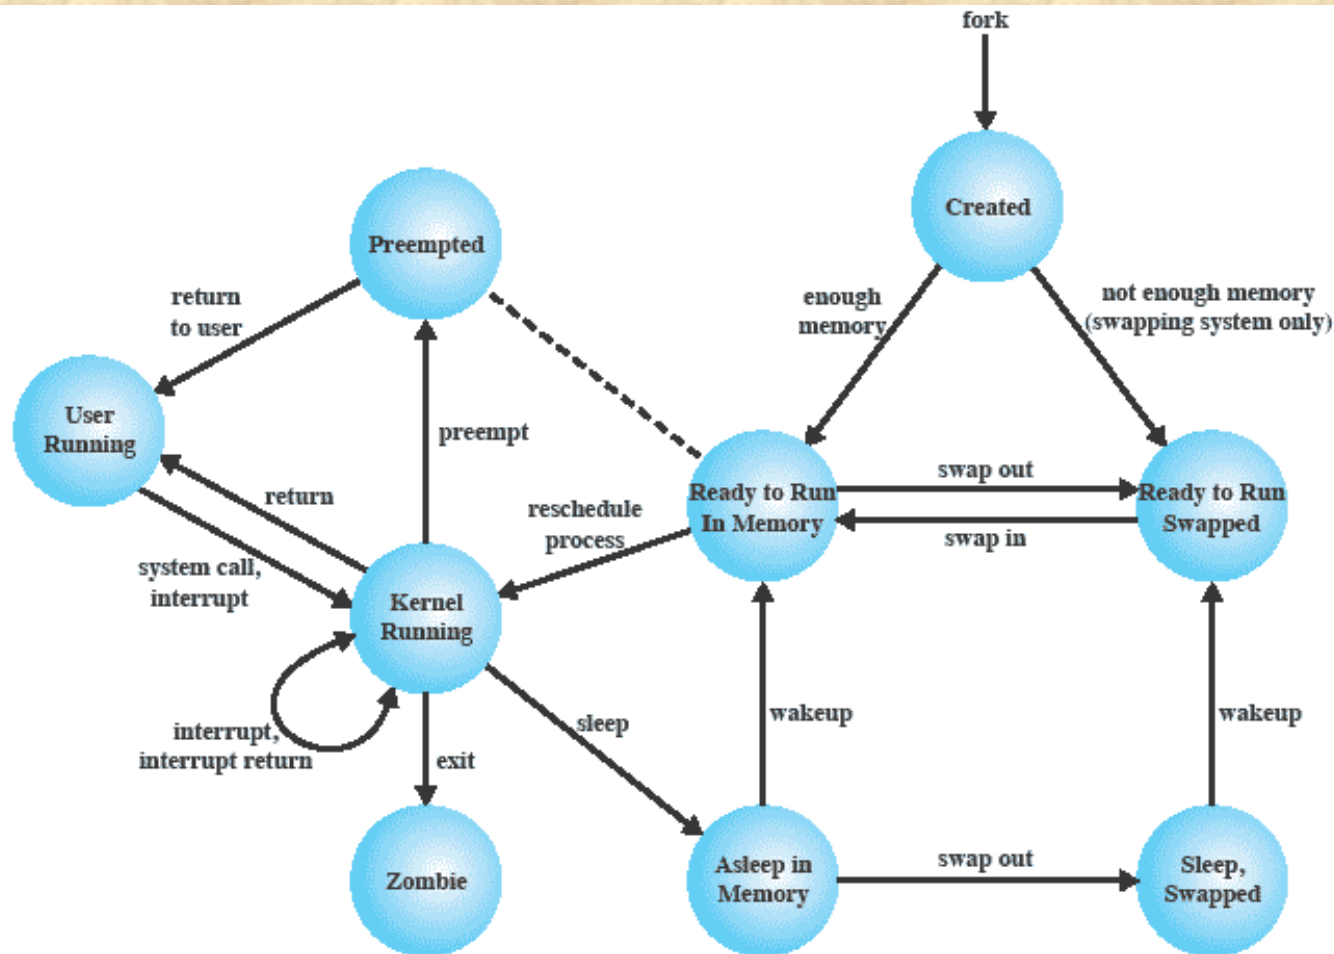



Figure 3.17 UNIX Process State Transition Diagram

A Unix Process



User-Level Context

Process text	Executable machine instructions of the program
Process data	Data accessible by the program of this process
User stack	Contains the arguments, local variables, and pointers for functions executing in user mode
Shared memory	Memory shared with other processes, used for interprocess communication

Register Context

Program counter	Address of next instruction to be executed; may be in kernel or user memory space of this process
Processor status register	Contains the hardware status at the time of preemption; contents and format are hardware dependent
Stack pointer	Points to the top of the kernel or user stack, depending on the mode of operation at the time or preemption
General-purpose registers	Hardware dependent

System-Level Context

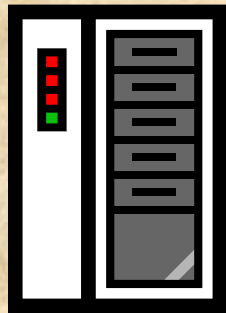
Process table entry	Defines state of a process; this information is always accessible to the operating system
U (user) area	Process control information that needs to be accessed only in the context of the process
Per process region table	Defines the mapping from virtual to physical addresses; also contains a permission field that indicates the type of access allowed the process: read-only, read-write, or read-execute
Kernel stack	Contains the stack frame of kernel procedures as the process executes in kernel mode

Table 3.11

UNIX

Process

Table Entry

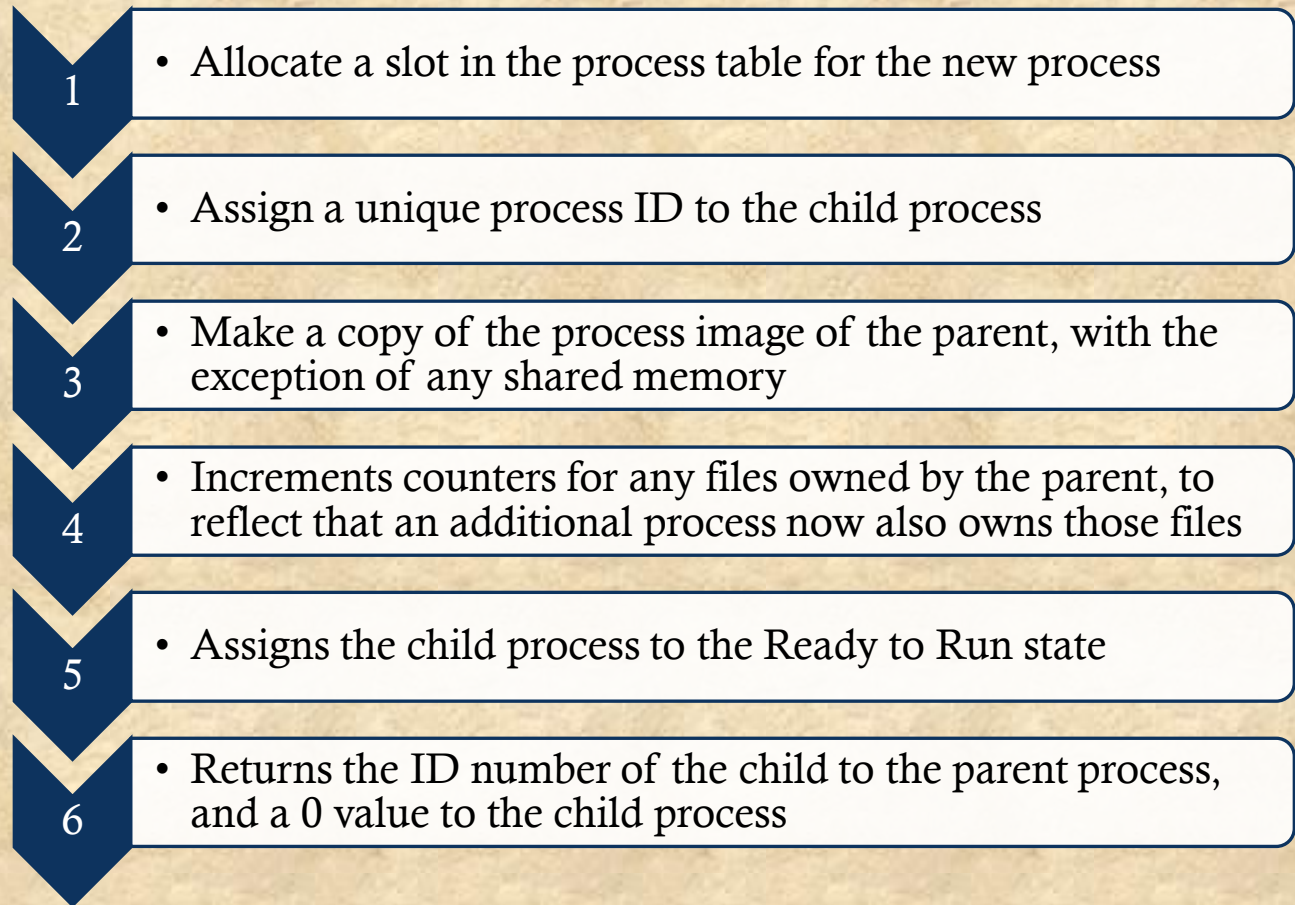


Process status	Current state of process.
Pointers	To U area and process memory area (text, data, stack).
Process size	Enables the operating system to know how much space to allocate the process.
User identifiers	The real user ID identifies the user who is responsible for the running process. The effective user ID may be used by a process to gain temporary privileges associated with a particular program; while that program is being executed as part of the process, the process operates with the effective user ID.
Process identifiers	ID of this process; ID of parent process. These are set up when the process enters the Created state during the fork system call.
Event descriptor	Valid when a process is in a sleeping state; when the event occurs, the process is transferred to a ready-to-run state.
Priority	Used for process scheduling.
Signal	Enumerates signals sent to a process but not yet handled.
Timers	Include process execution time, kernel resource utilization, and user-set timer used to send alarm signal to a process.
P_link	Pointer to the next link in the ready queue (valid if process is ready to execute).
Memory status	Indicates whether process image is in main memory or swapped out. If it is in memory, this field also indicates whether it may be swapped out or is temporarily locked into main memory.

Process Creation

■ Process creation is by means of the kernel system call, `fork()`

■ This causes the OS, in Kernel Mode, to:



After Creation

- After creating the process the Kernel can do one of the following, as part of the dispatcher routine:
 - stay in the parent process
 - transfer control to the child process
 - transfer control to another process



Summary

- The most fundamental concept in a modern OS is the process
- The principal function of the OS is to create, manage, and terminate processes
- Process control block contains all of the information that is required for the OS to manage the process, including its current state, resources allocated to it, priority, and other relevant data
- The most important states are Ready, Running and Blocked
- The running process is the one that is currently being executed by the processor
- A blocked process is waiting for the completion of some event
- A running process is interrupted either by an interrupt or by executing a supervisor call to the OS