
Process

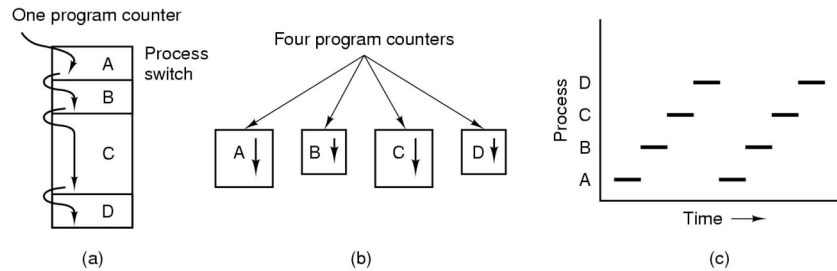
CSCE 351: Operating System
Kernels

Major Requirements of an OS

- Interleave the execution of several processes to maximize processor utilization while providing reasonable response time
- Allocate resources to processes
- Support interprocess communication and user creation of processes

Processes

The Process Model



- Multiprogramming of four programs
- Conceptual model of 4 independent, sequential processes
- Only one process active at any instant

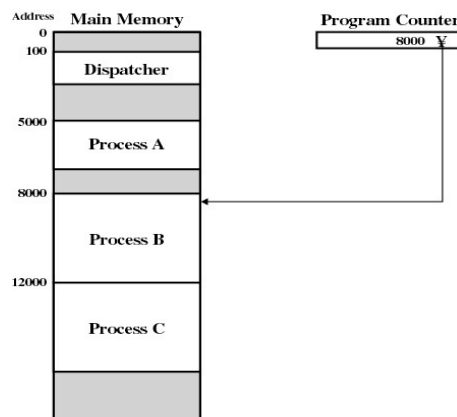



Figure 3.1 Snapshot of Example Execution (Figure 3.3) at Instruction Cycle 13

Two-State Process Model

- Process may be in one of two states
 - Running
 - Not-running




(a) State transition diagram

Two-State Process Model (2)

- Not-running
 - ready to execute
- Blocked
 - waiting for I/O
- Dispatcher cannot just select the process that has been in the queue the longest because it may be blocked

Not-Running Process in a Queue



(b) Queuing diagram

Process Creation

Principal events that cause process creation

1. System initialization
 - Execution of a process creation system
2. User request to create a new process
3. Initiation of a batch job

Process Termination


Conditions which terminate processes

1. Normal exit (voluntary)
2. Error exit (voluntary)
3. Fatal error (involuntary)
4. Killed by another process (involuntary)

Process Hierarchies

- Parent creates a child process, child processes can create its own process
- Forms a hierarchy
 - UNIX calls this a "process group"
- Windows has no concept of process hierarchy
 - all processes are created equal


A Three-State Model (1)



1. Process blocks for input
2. Scheduler picks another process
3. Scheduler picks this process
4. Input becomes available

- Possible process states
 - running
 - blocked
 - ready
- Transitions between states shown

A Three-State Model (2)



- Lowest layer of process-structured OS
 - handles interrupts, scheduling
- Above that layer are sequential processes

A Five-State Model

- Running
- Ready
- Blocked
- New
- Exit

A Five-State Model



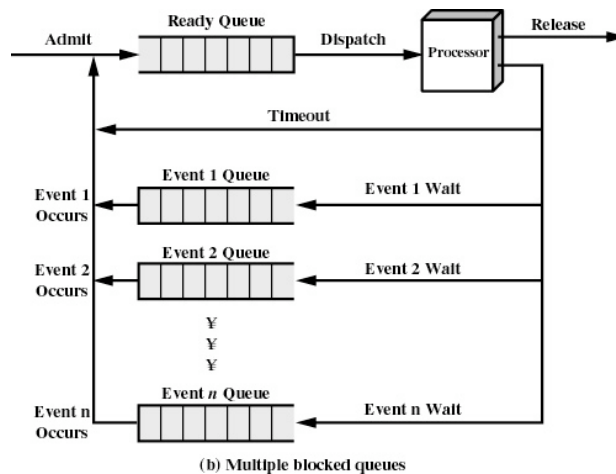


Figure 3.5 Five-State Process Model

Using Two Queues




Using Multiple Queues



Suspended Processes


- Processor is faster than I/O so all processes could be waiting for I/O
- Swap these processes to disk to free up more memory
- Blocked state becomes suspend state when swapped to disk
- Two new states
 - Blocked, suspend
 - Ready, suspend

One Suspend State



(a) With One Suspend State

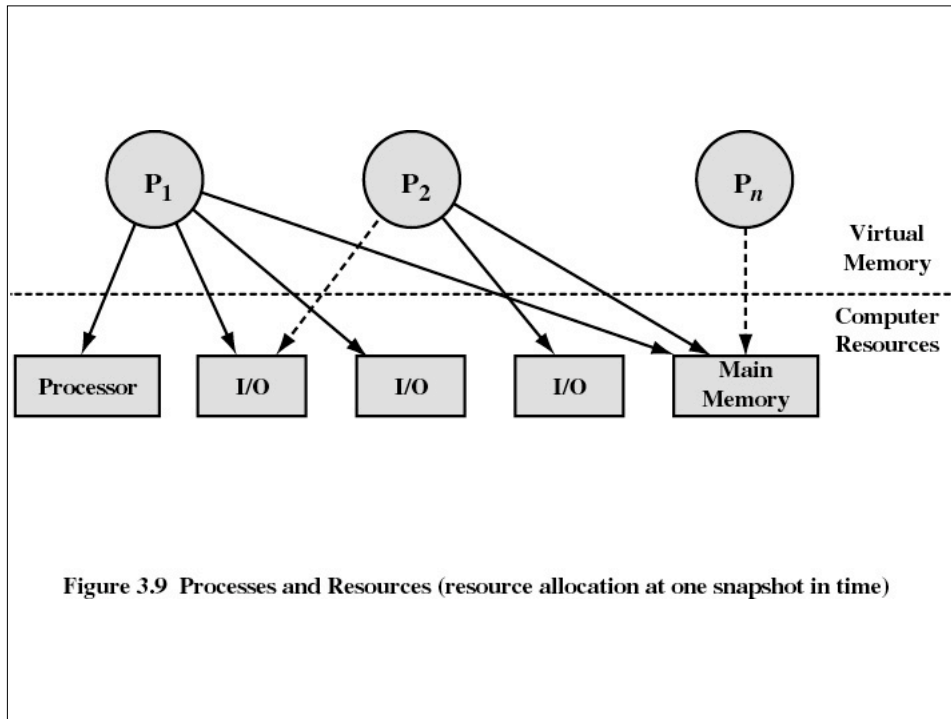
Two Suspend States



(b) With Two Suspend States

Reasons for Process Suspension

Swapping	The operating system needs to release sufficient main memory to bring in a process that is ready to execute.
Other OS reason	The operating system may suspend a background or utility process or a process that is suspected of causing a problem.
Interactive user request	A user may wish to suspend execution of a program for purposes of debugging or in connection with the use of a resource.
Timing	A process may be executed periodically (e.g., an accounting or system monitoring process) and may be suspended while waiting for the next time interval.
Parent process request	A parent process may wish to suspend execution of a descendent to examine or modify the suspended process, or to coordinate the activity of various descendents.



Operating System Control Structures

- Information about the current status of each process and resource
- Tables are constructed for each entity the operating system manages

Memory Tables

- Allocation of main memory to processes
- Allocation of secondary memory to processes
- Protection attributes for access to shared memory regions
- Information needed to manage virtual memory

I/O Tables

- I/O device is available or assigned
- Status of I/O operation
- Location in main memory being used as the source or destination of the I/O transfer

File Tables

- Existence of files
- Location on secondary memory
- Current Status
- Attributes
- Sometimes this information is maintained by a file-management system

Process Table

- Where process is located
- Attributes necessary for its management
 - Process ID
 - Process state
 - Location in memory

Process Table (2)

Process management	Memory management	File management
Registers Program counter Program status word Stack pointer Process state Priority Scheduling parameters Process ID Parent process Process group Signals Time when process started CPU time used Children's CPU time Time of next alarm	Pointer to text segment Pointer to data segment Pointer to stack segment	Root directory Working directory File descriptors User ID Group ID

Fields of a process table entry

Process Location

- Process includes set of programs to be executed
 - Data locations for local and global variables
 - Any defined constants
 - Stack
- Process control block
 - Collection of attributes
- Process image
 - Collection of program, data, stack, and attributes

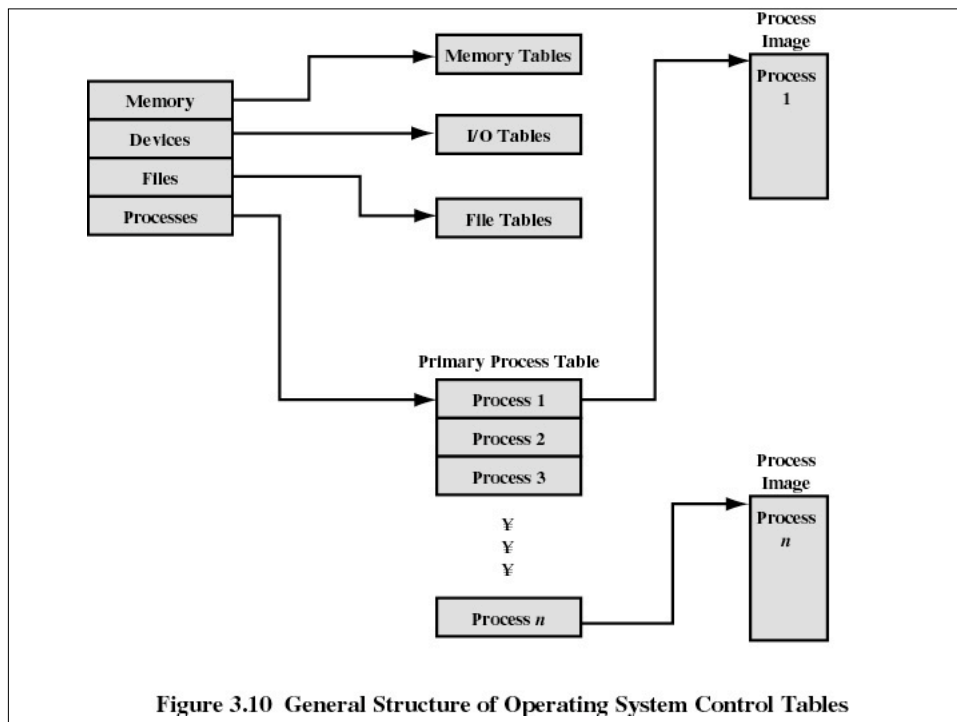


Figure 3.10 General Structure of Operating System Control Tables

Process Control Block (1)

- 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

Process Control Block (2)

- 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. Typically, there are from 8 to 32 of these registers, although some RISC implementations have over 100.

Process Control Block (3)

- Processor State Information
 - 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

Process Control Block (4)

- Processor State Information
 - 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.

Process Control Block (5)

- Process Control Information
 - Scheduling and State
 - *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

Process Control Block (6)

- Process Control Information
 - 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.

Process Control Block (7)

- Process Control Information
 - 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.

Process Control Block (8)

- Process Control Information
 - 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.

Process Control Block (9)

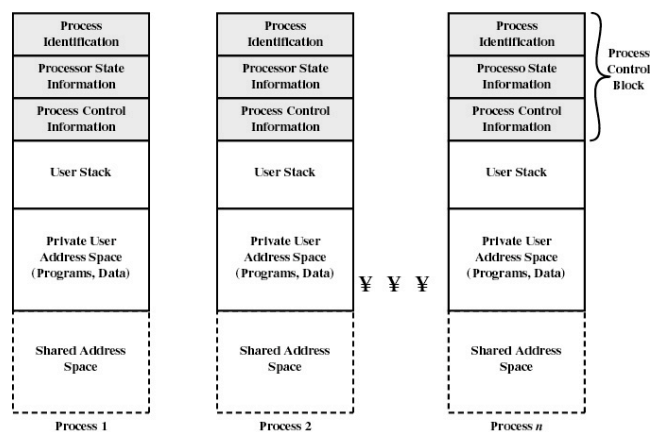



Figure 3.12 User Processes in Virtual Memory

Processor State Information

- Contents of processor registers
 - User-visible registers
 - Control and status registers
 - Stack pointers
- Program status word (PSW)
 - contains status information
 - Example: the EFLAGS register on Pentium machines

Pentium II EFLAGS Register



ID	=	Identification flag	DF	=	Direction flag
VIP	=	Virtual interrupt pending	IF	=	Interrupt enable flag
VIF	=	Virtual interrupt flag	TF	=	Trap flag
AC	=	Alignment check	SF	=	Sign flag
VM	=	Virtual 8086 mode	ZF	=	Zero flag
RF	=	Resume flag	AF	=	Auxiliary carry flag
NT	=	Nested task flag	PF	=	Parity flag
IOPL	=	I/O privilege level	CF	=	Carry flag
OF	=	Overflow flag			

Figure 3.11 Pentium II EFLAGS Register

Modes of Execution

- User mode
 - Less-privileged mode
 - User programs typically execute in this mode
- System mode, control mode, or kernel mode
 - More-privileged mode
 - Kernel of the operating system

Process Creation

- Assign a unique process identifier
- Allocate space for the process
- Initialize process control block
- Set up appropriate linkages
 - Ex: add new process to linked list used for scheduling queue
- Create or expand other data structures
 - Ex: maintain an accounting file

When to Switch a Process

- Clock interrupt
 - process has executed for the maximum allowable time slice
- I/O interrupt
- Memory fault
 - memory address is in virtual memory so it must be brought into main memory

When to Switch a Process

- Trap
 - error occurred
 - may cause process to be moved to Exit state
- Supervisor call
 - such as file open

Change of Process State (1)

- Save context of processor including program counter and other registers
- Update the process control block of the process that is currently running
- Move process control block to appropriate queue - ready, blocked
- Select another process for execution

Change of Process State (2)

- Update the process control block of the process selected
- Update memory-management data structures
- Restore context of the selected process

Change of Process State (3)

1. Hardware stacks program counter, etc.
2. Hardware loads new program counter from interrupt vector.
3. Assembly language procedure saves registers.
4. Assembly language procedure sets up new stack.
5. C interrupt service runs (typically reads and buffers input).
6. Scheduler decides which process is to run next.
7. C procedure returns to the assembly code.
8. Assembly language procedure starts up new current process.

Skeleton of what lowest level of OS does when an interrupt occurs

Execution of the Operating System

- Non-process Kernel
 - Execute kernel outside of any process
 - Operating system code is executed as a separate entity that operates in privileged mode
- Execution Within User Processes
 - Operating system software within context of a user process
 - Process executes in privileged mode when executing operating system code

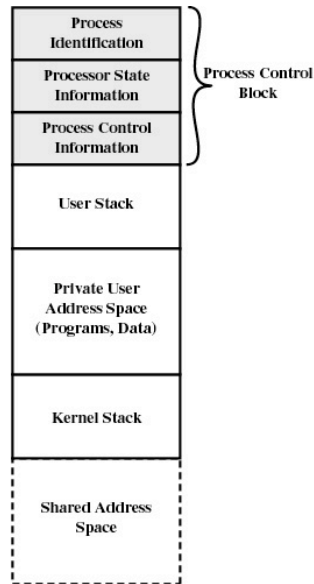



Figure 3.15 Process Image: Operating System Executes Within User Space

Execution of the Operating System

- Process-Based Operating System
 - Major kernel functions are separate processes
 - Useful in multi-processor or multi-computer environment

UNIX SVR4 Process Management

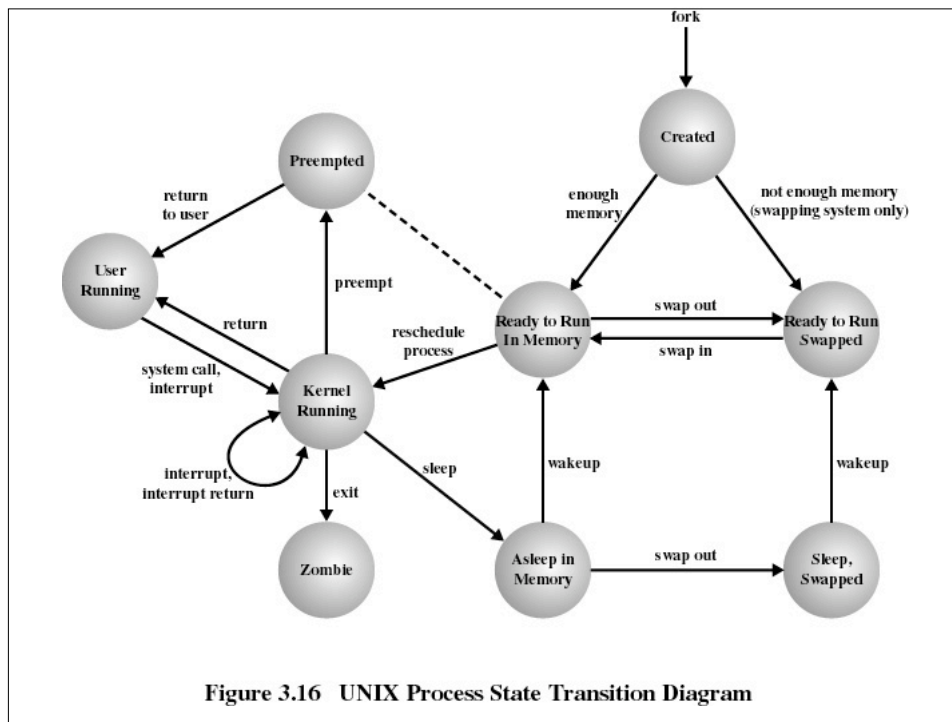
- Most of the operating system executes within the environment of a user process



(b) OS functions execute within user processes

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.



Summary

- In this presentation we discussed
 - Process life-cycle
 - Process swapping mechanism
 - Basic data structure that maintain process
- At this point, we just finished chapter 3 in the book