CSC 4103 - Operating Systems
Spring 2007

LECTURE - XI
MIDTERM REVIEW
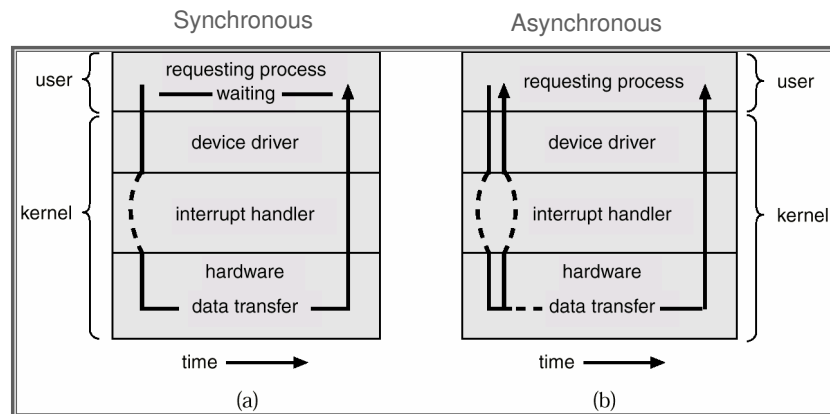
Tevfik Koşar

Louisiana State University
March 1st, 2007

I/O Structure

- After I/O starts, control returns to user program only upon I/O completion ➔ synchronous
  - Wait instruction idles the CPU until the next interrupt
  - Wait loop (contention for memory access).
  - At most one I/O request is outstanding at a time, no simultaneous I/O processing.

- After I/O starts, control returns to user program without waiting for I/O completion ➔ asynchronous
  - *System call* – request to the operating system to allow user to wait for I/O completion.
  - *Device-status table* contains entry for each I/O device indicating its type, address, and state.
  - Operating system indexes into I/O device table to determine device status and to modify table entry to include interrupt.

2

Two I/O Methods

Synchronous   Asynchronous

user	requesting process
	waiting
	device driver
kernel	interrupt handler
	hardware
	data transfer

time ⟶

(a)

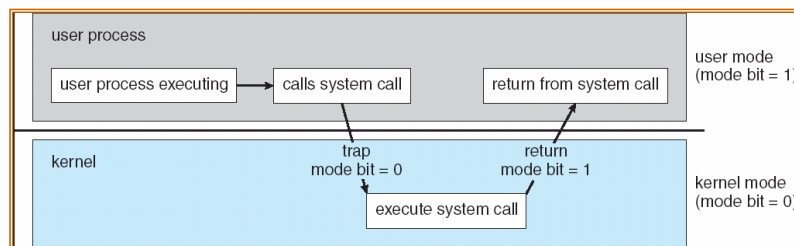| | requesting process | user |
	device driver	
	interrupt handler	kernel
	hardware	
	data transfer	

time ⟶

(b)

3

Dual-Mode Operation

- **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
        - Protects OS from errant users, and errant users from each other
        - Some instructions designated as **privileged**, only executable in kernel mode
        - System call changes mode to kernel, return from call resets it to user
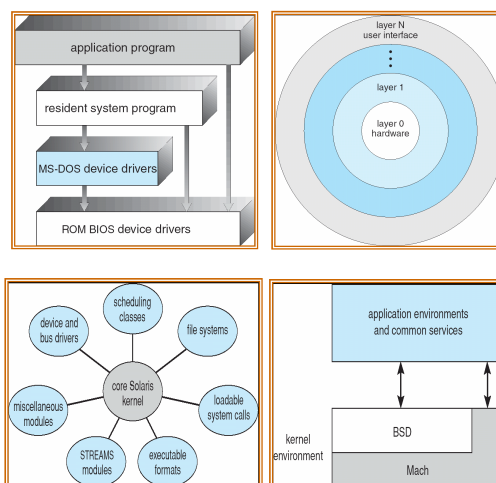
4

Transition from User to Kernel Mode

- How to prevent user program getting stuck in an infinite loop / process hogging resources
  - → Timer: Set interrupt after specific period (1ms to 1sec)
  - – Operating system decrements counter
  - – When counter zero generate an interrupt
  - – Set up before scheduling process to regain control or terminate program that exceeds allotted time

| user process			user mode (mode bit = 1)
user process executing	calls system call	return from system call	

kernel — trap mode bit = 0 — return mode bit = 1 — kernel mode (mode bit = 0)

execute system call

5

OS Design Approaches

- Simple Structure
- Layered Approach
- Microkernels
- Modules

application program

resident system program

MS-DOS device drivers

ROM BIOS device drivers

layer N user interface
⋮
layer 1
layer 0 hardware

scheduling classes
device and bus drivers
file systems
miscellaneous modules
core Solaris kernel
loadable system calls
STREAMS modules
executable formats

application environments and common services
kernel environment
BSD
Mach

6

3

Process State

- As a process executes, it changes *state*
  - **new**: The process is being created
  - **running**: Instructions are being executed
  - **waiting**: The process is waiting for some event to occur
  - **ready**: The process is waiting to be assigned to a process
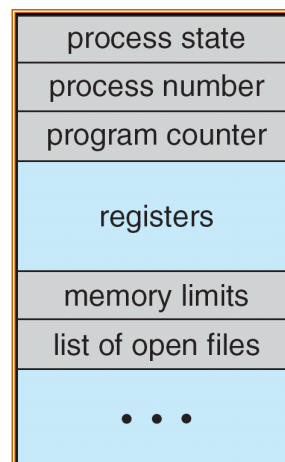 - **terminated**: The process has finished execution
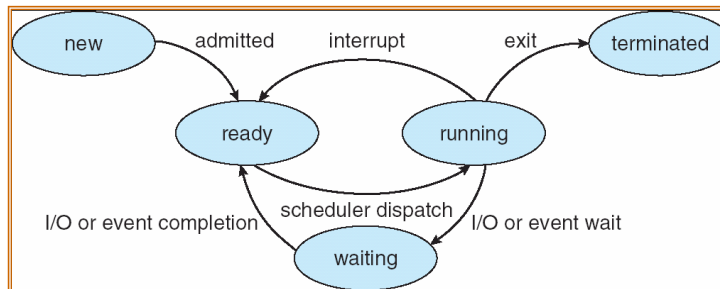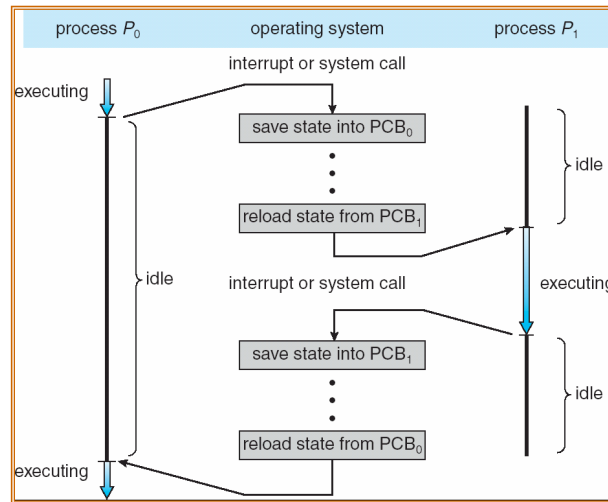
Process Control Block (PCB)

Information associated with each process

- Process state
- Program counter
- CPU registers
- CPU scheduling information
- Memory-management information
- Accounting information
- I/O status information

process state

process number

program counter

registers

memory limits

list of open files

• • •

CPU Switch From Process to Process

process P_0 operating system   process P_1

interrupt or system call

executing

save state into PCB_0

•
•
•

reload state from PCB_1

idle

idle

interrupt or system call   executing

save state into PCB_1

•
•
•

reload state from PCB_0

idle

executing

9

Schedulers (Cont.)

- Short-term scheduler is invoked very frequently (milliseconds) $\Rightarrow$ (must be fast)
- Long-term scheduler is invoked very infrequently (seconds, minutes) $\Rightarrow$ (may be slow)
- The long-term scheduler controls the *degree of multiprogramming*
- Processes can be described as either:
  - **I/O-bound process** – spends more time doing I/O than computations, many short CPU bursts
  - **CPU-bound process** – spends more time doing computations; few very long CPU bursts

  ➔long-term schedulers need to make careful decision

10

5

Addition of Medium Term Scheduling

- In time-sharing systems: remove processes from memory "temporarily" to reduce degree of multiprogramming.
- Later, these processes are resumed ➜ Swaping
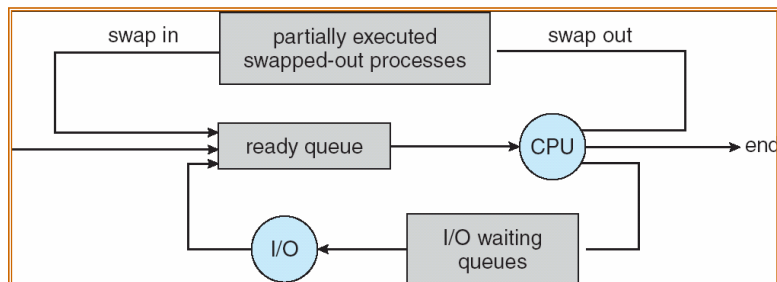
Process Creation (Cont.)

- Address space
  - Child duplicate of parent
  - Child has a program loaded into it
- UNIX examples
  - **fork** system call creates new process
  - **exec** system call used after a **fork** to replace the process' memory space with a new program
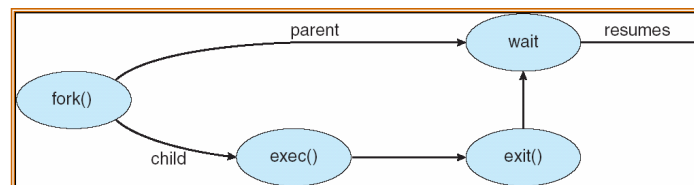
C Program Forking Separate Process

```
int main()
{
Pid_t  pid;
    /* fork another process */
    pid = fork();
    if (pid < 0) { /* error occurred */
        fprintf(stderr, "Fork Failed");
        exit(-1);
    }
    else if (pid == 0) { /* child process */
        execlp("/bin/ls", "ls", NULL);
    }
    else { /* parent process */
        /* parent will wait for the child to complete */
        wait (NULL);
        printf ("Child Complete");
        exit(0);
    }
}
```

13

Synchronization

- Message passing may be either blocking or non-blocking
- **Blocking** is considered **synchronous**
  - **Blocking send** has the sender block until the message is received
  - **Blocking receive** has the receiver block until a message is available
- **Non-blocking** is considered **asynchronous**
  - **Non-blocking** send has the sender send the message and continue
  - **Non-blocking** receive has the receiver receive a valid message or null

14

Threads vs Processes

- Heavyweight Process = Process
- Lightweight Process = Thread

Advantages (Thread vs. Process):
- Much quicker to create a thread than a process
- Much quicker to switch between threads than to switch between processes
- Threads share data easily

Disadvantages (Thread vs. Process):
- Processes are more flexible
  - They don't have to run on the same processor
- No security between threads: One thread can stomp on another thread's data
- For threads which are supported by user thread package instead of the kernel:
  - If one thread blocks, all threads in task block.

15

Different Multi-threading Models
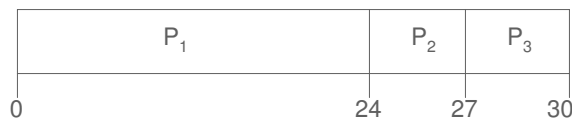
- Many-to-One
- One-to-One
- Many-to-Many
- Two-level Model
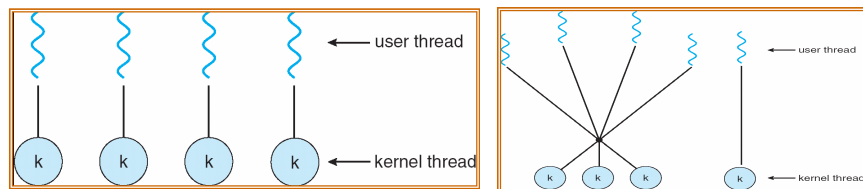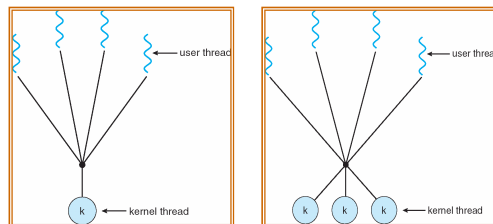
16

First-Come, First-Served (FCFS) Scheduling

Process	Burst Time
P_1	24
P_2	3
P_3	3

- Suppose that the processes arrive in the order: P_1, P_2, P_3
  The Gantt Chart for the schedule is:

P_1		P_2	P_3
0	24	27	30

- Waiting time for P_1 = 0; P_2 = 24; P_3 = 27
- Average waiting time: (0 + 24 + 27)/3 = 17

17

## Shortest-Job-First (SJF) Scheduling

- Associate with each process the length of its next CPU burst. Use these lengths to schedule the process with the shortest time
- Two schemes:
  - nonpreemptive – once CPU given to the process it cannot be preempted until completes its CPU burst
  - preemptive – if a new process arrives with CPU burst length less than remaining time of current executing process, preempt. This scheme is know as the Shortest-Remaining-Time-First (SRTF)
- SJF is optimal – gives minimum average waiting time for a given set of processes

18

Example of Non-Preemptive SJF

Process	Arrival Time	Burst Time
P_1	0.0	7
P_2	2.0	4
P_3	4.0	1
P_4	5.0	4

- SJF (non-preemptive)

P_1					P_3	P_2		P_4	

```
0       3           7  8       12          16
```
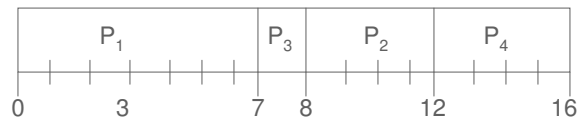
- Average waiting time = $(0 + 6 + 3 + 7)/4 = 4$

19

Example of Preemptive SJF

Process	Arrival Time	Burst Time
P_1	0.0	7
P_2	2.0	4
P_3	4.0	1
P_4	5.0	4

- SJF (preemptive)

P_1	P_2	P_3	P_2	P_4	P_1

```
0   2   4  5     7       11          16
```

- Average waiting time = $(9 + 1 + 0 + 2)/4 = 3$

20

Priority Scheduling

- A priority number (integer) is associated with each process
- The CPU is allocated to the process with the highest priority (smallest integer ≡ highest priority)
  - Preemptive
  - nonpreemptive
- SJF is a priority scheduling where priority is the predicted next CPU burst time
- Problem ≡ Starvation – low priority processes may never execute
- Solution ≡ Aging – as time progresses increase the priority of the process

21

Round Robin (RR)

- Each process gets a small unit of CPU time (*time quantum*), usually 10-100 milliseconds. After this time has elapsed, the process is preempted and added to the end of the ready queue.
- If there are *n* processes in the ready queue and the time quantum is *q*, then each process gets $1/n$ of the CPU time in chunks of at most *q* time units at once. No process waits more than $(n\text{-}1)q$ time units.
- Performance
  - *q* large $\Rightarrow$ FIFO
  - *q* small $\Rightarrow$ *q* must be large with respect to context switch, otherwise overhead is too high

22

Example of RR with Time Quantum = 20

Process	Burst Time
P_1	53
P_2	17
P_3	68
P_4	24

- The Gantt chart is:

P_1	P_2	P_3	P_4	P_1	P_3	P_4	P_1	P_3	P_3

0   20   37   57   77   97   117   121   134   154   162

- Typically, higher average turnaround than SJF, but better *response*

23

Multilevel Feedback Queue

- A process can move between the various queues; aging can be implemented this way
- Multilevel-feedback-queue scheduler defined by the following parameters:
  - number of queues
  - scheduling algorithms for each queue
  - method used to determine when to upgrade a process
  - method used to determine when to demote a process
  - method used to determine which queue a process will enter when that process needs service

24

12

Solution to Critical-Section Problem

A solution to the critical-section problem must satisfy the following requirements:

1. Mutual Exclusion - If process P_i is executing in its critical section, then no other processes can be executing in their critical sections

2. Progress - If no process is executing in its critical section and there exist some processes that wish to enter their critical section, then the selection of the processes that will enter the critical section next cannot be postponed indefinitely

25

Solution to Critical-Section Problem

3. Bounded Waiting - A bound must exist on the number of times that other processes are allowed to enter their critical sections after a process has made a request to enter its critical section and before that request is granted

- Assume that each process executes at a nonzero speed
- No assumption concerning relative speed of the N processes

26

Semaphores as Synchronization Tool

- Counting semaphore – integer value can range over an unrestricted domain
- Binary semaphore – integer value can range only between 0 and 1; can be simpler to implement
  - Also known as mutex locks

- Provides mutual exclusion
  - Semaphore S;   // initialized to 1
  - wait (S);
          Critical Section
      signal (S);

27

Deadlock and Starvation

- Deadlock – two or more processes are waiting indefinitely for an event that can be caused by only one of the waiting processes
- Let S and Q be two semaphores initialized to 1

P_0	P_1
wait (S);	wait (Q);
.	.
wait (Q);	wait (S);
.	.
.	.
signal (S);	signal (Q);
signal (Q);	signal (S);

- Starvation – indefinite blocking. A process may never be removed from the semaphore queue in which it is suspended.

28

14

Classical Problems of Synchronization

- Bounded-Buffer Problem
- Readers and Writers Problem
- Dining-Philosophers Problem
- Sleeping Barber Problem

- You should be able to solve them using semaphores as well as monitors.

Monitors

- A high-level abstraction that provides a convenient and effective mechanism for process synchronization
- Only one process may be active within the monitor at a time

```
monitor monitor-name
{
    // shared variable declarations
    procedure P1 (…) { …. }
            …
    procedure Pn (…) {……}

     Initialization code ( ….) { … }
            …
    }
}
```

- A monitor procedure takes the lock before doing anything else, and holds it until it either finishes or waits for a condition

Deadlock Characterization

Deadlock can arise if four conditions hold simultaneously.

1. Mutual exclusion: nonshared resources; only one process at a time can use a specific resource

2. Hold and wait: a process holding at least one resource is waiting to acquire additional resources held by other processes

3. No preemption: a resource can be released only voluntarily by the process holding it, after that process has completed its task

31

Deadlock Characterization (cont.)

Deadlock can arise if four conditions hold simultaneously.

4. Circular wait: there exists a set $\{P_0, P_1, ..., P_0\}$ of waiting processes such that P_0 is waiting for a resource that is held by P_1, P_1 is waiting for a resource that is held by P_2, ..., P_{n-1} is waiting for a resource that is held by P_n, and P_n is waiting for a resource that is held by P_0.

32

Resource-Allocation Graph

• Used to describe deadlocks

• Consists of a set of vertices *V* and a set of edges *E*.

• V is partitioned into two types:
  – $P = \{P_1, P_2, ..., P_n\}$, the set consisting of all the processes in the system.

  – $R = \{R_1, R_2, ..., R_m\}$, the set consisting of all resource types in the system.

• P requests R – directed edge $P_1 \rightarrow R_j$
• R is assigned to P – directed edge $R_j \rightarrow P_i$

33

Example of a Resource Allocation Graph
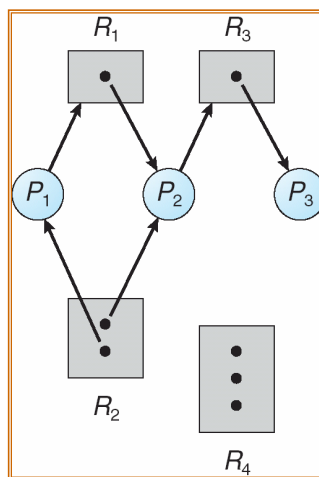
34

Basic Facts

- If graph contains no cycles $\Rightarrow$ no deadlock.

- If graph contains a cycle $\Rightarrow$ there may be a deadlock
  - if only one instance per resource type, then deadlock.
  - if several instances per resource type, possibility of deadlock.

35

Deadlock Prevention

➔ Ensure one of the deadlock conditions cannot hold

➔ Restrain the ways request can be made.

- **Mutual Exclusion** – not required for sharable resources; must hold for nonsharable resources.
  - Eg. read-only files

- **Hold and Wait** – must guarantee that whenever a process requests a resource, it does not hold any other resources.
  1. Require process to request and be allocated all its resources before it begins execution
  2. or allow process to request resources only when the process has none.

  Example: Write from DVD to disk, then print.
  1. holds printer unnecessarily for the entire execution
     - Low resource utilization
  2. may never get the printer later
     - starvation possible

36

18

Deadlock Prevention (Cont.)

- **No Preemption** –
  - If a process that is holding some resources requests another resource that cannot be immediately allocated to it, then all resources currently being held are released.
  - Preempted resources are added to the list of resources for which the process is waiting.
  - Process will be restarted only when it can regain its old resources, as well as the new ones that it is requesting.

- **Circular Wait** – impose a total ordering of all resource types, and require that each process requests resources in an increasing order of enumeration.
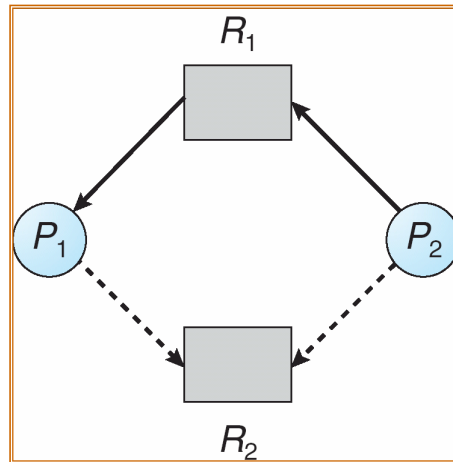
37

Deadlock Avoidance

- If a system is in safe state $\Rightarrow$ no deadlocks.

- If a system is in unsafe state $\Rightarrow$ possibility of deadlock.

- Avoidance $\Rightarrow$ ensure that a system will never enter an unsafe state.

38

Resource-Allocation Graph For Deadlock Avoidance
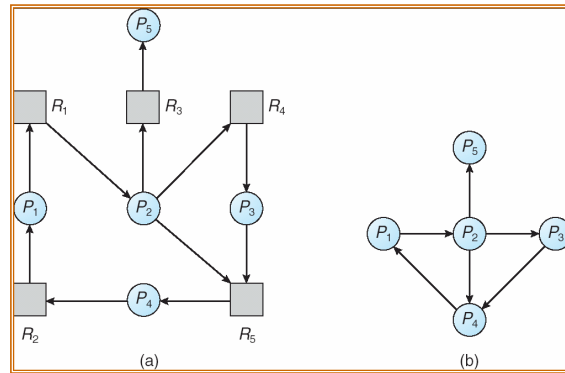
R_1

P_1 P_2

R_2

Banker's Algorithm

- Multiple instances.

- Each process must a priori claim maximum use.

- When a process requests a resource it may have to wait.

- When a process gets all its resources it must return them in a finite amount of time.

Resource-Allocation Graph and Wait-for Graph

(a)

(b)

Resource-Allocation Graph   Corresponding wait-for graph

Detection Algorithm

1. Let *Work* and *Finish* be vectors of length *m* and *n*, respectively Initialize:
   - (a) *Work = Available*
   - (b) For $i = 1, 2, ..., n$, if *Allocation$_i$* $\neq$ 0, then *Finish*[i] = false; otherwise, *Finish*[i] = *true*.
2. Find an index *i* such that both:
   - (a) *Finish*[i] == *false*
   - (b) *Request$_i$* $\leq$ *Work*

   If no such *i* exists, go to step 4.

Recovery from Deadlock: Process Termination

- Abort all deadlocked processes.

- Abort one process at a time until the deadlock cycle is eliminated.

- In which order should we choose to abort?
  - Priority of the process.
  - How long process has computed, and how much longer to completion.
  - Resources the process has used.
  - Resources process needs to complete.
  - How many processes will need to be terminated.
  - Is process interactive or batch?

43

Recovery from Deadlock: Resource Preemption

- Selecting a victim – minimize cost.

- Rollback – return to some safe state, restart process for that state.

- Starvation – same process may always be picked as victim, include number of rollback in cost factor.

44