

## Chapter 6: CPU Scheduling

- Basic Concepts
- Scheduling Criteria
- Scheduling Algorithms
- Multiple-Processor Scheduling
- Real-Time Scheduling
- Algorithm Evaluation

Operating System Concepts 6.1 Silberschatz, Galvin and Gagne ©2002

## Basic Concepts

- Maximum CPU utilization obtained with multiprogramming
- CPU I/O Burst Cycle Process execution consists of a cycle of CPU execution and I/O wait.
- CPU burst distribution

Operating System Concepts 6.2 Silberschatz, Galvin and Gagne ©2002

## Alternating Sequence of CPU And I/O Bursts

⋮

load store  
add store  
read from file

wait for I/O

store increment  
index  
write to file

wait for I/O

load store  
add store  
read from file

wait for I/O

⋮

} CPU burst  
} I/O burst  
} CPU burst  
} I/O burst  
} CPU burst  
} I/O burst

Operating System Concepts 6.3 Silberschatz, Galvin and Gagne ©2002

## Histogram of CPU-burst Times

Burst Duration (ms)	Frequency
0-4	150
4-8	35
8-12	10
12-16	5
16-20	2
20-24	1
24-28	1
28-32	1
32-36	1
36-40	1

Operating System Concepts 6.4 Silberschatz, Galvin and Gagne ©2002

## CPU Scheduler

- Selects from among the processes in memory that are ready to execute, and allocates the CPU to one of them.

CPU scheduling decisions may take place when a process:

1. Switches from running to waiting state.
2. Switches from running to ready state.
3. Switches from waiting to ready.
4. Terminates.

Scheduling under 1 and 4 is *nonpreemptive*.  
All other scheduling is *preemptive*.

Operating System Concepts 6.5 Silberschatz, Galvin and Gagne ©2002

## Dispatcher

- Dispatcher module gives control of the CPU to the process selected by the short-term scheduler; this involves:
  - switching context
  - switching to user mode
  - jumping to the proper location in the user program to restart that program

*Dispatch latency* time it takes for the dispatcher to stop one process and start another running.

Operating System Concepts 6.6 Silberschatz, Galvin and Gagne ©2002

## Scheduling Criteria

- CPU utilization keep the CPU as busy as possible
- Throughput # of processes that complete their execution per time unit
- Turnaround time amount of time to execute a particular process
- Waiting time amount of time a process has been waiting in the ready queue
- Response time amount of time it takes from when a request was submitted until the first response is produced, **not** output (for time-sharing environment)

Operating System Concepts 6.7 Silberschatz, Galvin and Gagne ©2002

## Optimization Criteria

- Max CPU utilization
- Max throughput
- Min turnaround time
- Min waiting time
- Min response time

Operating System Concepts 6.8 Silberschatz, Galvin and Gagne ©2002

## 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:

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

Operating System Concepts 6.9 Silberschatz, Galvin and Gagne ©2002

## FCFS Scheduling (Cont.)

Suppose that the processes arrive in the order  $P_2, P_3, P_1$ .  
The Gantt chart for the schedule is:

Waiting time for  $P_1 = 6; P_2 = 0, P_3 = 3$   
Average waiting time:  $(6 + 0 + 3)/3 = 3$   
Much better than previous case.  
*Convoy effect* short process behind long process

Operating System Concepts 6.10 Silberschatz, Galvin and Gagne ©2002

## 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 known as the Shortest-Remaining-Time-First (SRTF).

SJF is optimal gives minimum average waiting time for a given set of processes.

Operating System Concepts 6.11 Silberschatz, Galvin and Gagne ©2002

## 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)

Average waiting time =  $(0 + 6 + 3 + 7)/4 = 4$

Operating System Concepts 6.12 Silberschatz, Galvin and Gagne ©2002

### 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)

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

Operating System Concepts 6.13 Silberschatz, Galvin and Gagne ©2002

### Determining Length of Next CPU Burst

- Can only estimate the length.
- Can be done by using the length of previous CPU bursts, using exponential averaging.

- $t_n$  = actual length of  $n^{th}$  CPU burst
- $\tau_{n+1}$  = predicted value for the next CPU burst
- $\alpha, 0 \leq \alpha \leq 1$
- Define:
 
$$\tau_{n+1} = \alpha t_n + (1 - \alpha)\tau_n$$

Operating System Concepts 6.14 Silberschatz, Galvin and Gagne ©2002

### Prediction of the Length of the Next CPU Burst

CPU burst ( $t_i$ )	6	4	6	4	13	13	13	...	
"guess" ( $\tau_i$ )	10	8	6	6	5	9	11	12	...

Operating System Concepts 6.15 Silberschatz, Galvin and Gagne ©2002

### Examples of Exponential Averaging

- $\alpha = 0$ 
  - $\tau_{n+1} = \tau_n$
  - Recent history does not count.
- $\alpha = 1$ 
  - $\tau_{n+1} = t_n$
  - Only the actual last CPU burst counts.

If we expand the formula, we get:

$$\tau_{n+1} = \alpha t_n + (1 - \alpha) \alpha t_{n-1} + (1 - \alpha)^2 \alpha t_{n-2} + \dots + (1 - \alpha)^{n-1} \alpha t_1 + (1 - \alpha)^n \tau_0$$

Since both  $\alpha$  and  $(1 - \alpha)$  are less than or equal to 1, each successive term has less weight than its predecessor.

Operating System Concepts 6.16 Silberschatz, Galvin and Gagne ©2002

### 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  $\equiv$  Starvation low priority processes may never execute.
- Solution  $\equiv$  Aging as time progresses increase the priority of the process.

Operating System Concepts 6.17 Silberschatz, Galvin and Gagne ©2002

### 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  $(n-1)q$  time units.
- Performance
  - $q$  large  $\Rightarrow$  FIFO
  - $q$  small  $\Rightarrow q$  must be large with respect to context switch, otherwise overhead is too high.

Operating System Concepts 6.18 Silberschatz, Galvin and Gagne ©2002

### 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.

Operating System Concepts 6.19 Silberschatz, Galvin and Gagne ©2002

### Time Quantum and Context Switch Time

process time = 10	quantum	context switches
0 10	12	0
0 10	6	1
0 1 2 3 4 5 6 7 8 9 10	1	9

Operating System Concepts 6.20 Silberschatz, Galvin and Gagne ©2002

### Turnaround Time Varies With The Time Quantum

process	time
$P_1$	6
$P_2$	3
$P_3$	1
$P_4$	7

Operating System Concepts 6.21 Silberschatz, Galvin and Gagne ©2002

### Multilevel Queue

- Ready queue is partitioned into separate queues:
  - foreground (interactive)
  - background (batch)
- Each queue has its own scheduling algorithm,
  - foreground RR
  - background FCFS
- Scheduling must be done between the queues.
  - Fixed priority scheduling; (i.e., serve all from foreground then from background). Possibility of starvation.
  - Time slice each queue gets a certain amount of CPU time which it can schedule amongst its processes; i.e., 80% to foreground in RR
  - 20% to background in FCFS

Operating System Concepts 6.22 Silberschatz, Galvin and Gagne ©2002

### Multilevel Queue Scheduling

Operating System Concepts 6.23 Silberschatz, Galvin and Gagne ©2002

### 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

Operating System Concepts 6.24 Silberschatz, Galvin and Gagne ©2002

### Example of Multilevel Feedback Queue

- Three queues:
  - $Q_0$  time quantum 8 milliseconds
  - $Q_1$  time quantum 16 milliseconds
  - $Q_2$  FCFS
- Scheduling
 

A new job enters queue  $Q_0$  which is served FCFS. When it gains CPU, job receives 8 milliseconds. If it does not finish in 8 milliseconds, job is moved to queue  $Q_1$ .

At  $Q_1$  job is again served FCFS and receives 16 additional milliseconds. If it still does not complete, it is preempted and moved to queue  $Q_2$ .

Operating System Concepts 6.25 Silberschatz, Galvin and Gagne ©2002

### Multilevel Feedback Queues

Operating System Concepts 6.26 Silberschatz, Galvin and Gagne ©2002

### Multiple-Processor Scheduling

- CPU scheduling more complex when multiple CPUs are available.
  - Homogeneous processors* within a multiprocessor.
    - Load sharing*
    - Asymmetric multiprocessing* only one processor accesses the system data structures, alleviating the need for data sharing.

Operating System Concepts 6.27 Silberschatz, Galvin and Gagne ©2002

### Real-Time Scheduling

- Hard real-time systems* required to complete a critical task within a guaranteed amount of time.
  - Soft real-time computing* requires that critical processes receive priority over less fortunate ones.

Operating System Concepts 6.28 Silberschatz, Galvin and Gagne ©2002


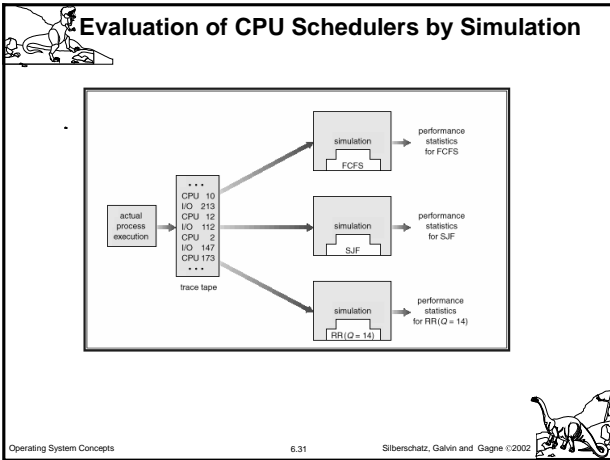
### Dispatch Latency

Operating System Concepts 6.29 Silberschatz, Galvin and Gagne ©2002

### Algorithm Evaluation

- Deterministic modeling takes a particular predetermined workload and defines the performance of each algorithm for that workload.
  - Queueing models
  - Implementation

Operating System Concepts 6.30 Silberschatz, Galvin and Gagne ©2002



### Windows 2000 Priorities

	real-time	high	above normal	normal	below normal	idle priority
time-critical	31	15	15	15	15	15
highest	26	15	12	10	8	6
above normal	25	14	11	9	7	5
normal	24	13	10	8	6	4
below normal	23	12	9	7	5	3
lowest	22	11	8	6	4	2
idle	16	1	1	1	1	1

Operating System Concepts 6.33 Silberschatz, Galvin and Gagne ©2002

This document was created with Win2PDF available at <http://www.daneprairie.com>.  
The unregistered version of Win2PDF is for evaluation or non-commercial use only.