



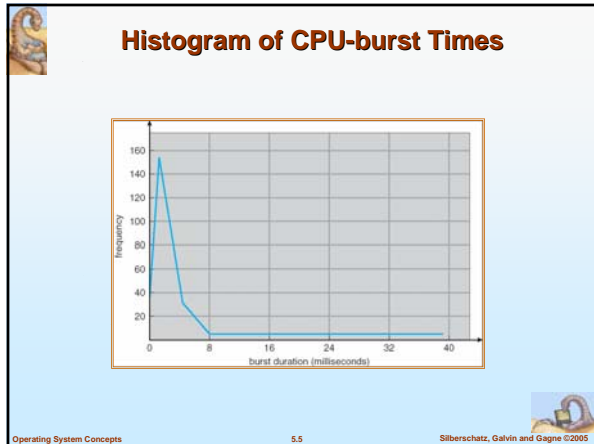
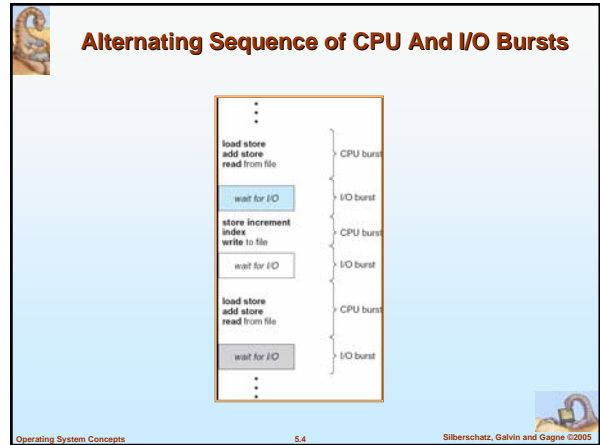
Chapter 5: CPU Scheduling

Chapter 5: CPU Scheduling

- Basic Concepts
- Scheduling Criteria
- Scheduling Algorithms
- Multiple-Processor Scheduling
- Real-Time Scheduling
- Thread Scheduling
- Operating Systems Examples
- Java Thread Scheduling
- Algorithm Evaluation

- ## 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



- ## 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*

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 5.13 Silberschatz, Galvin and Gagne ©2005

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 5.14 Silberschatz, Galvin and Gagne ©2005

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

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

Operating System Concepts 5.15 Silberschatz, Galvin and Gagne ©2005

Prediction of the Length of the Next CPU Burst

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

Operating System Concepts 5.16 Silberschatz, Galvin and Gagne ©2005

Examples of Exponential Averaging

- $\alpha = 0$
 - $\tau_{n+1} = \tau_n$
 - Recent history does not count
- $\alpha = 1$
 - $\tau_{n+1} = \alpha t_n$
 - Only the actual last CPU burst counts
- If we expand the formula, we get:

$$\tau_{n+1} = \alpha t_n + (1 - \alpha)\alpha t_{n-1} + \dots + (1 - \alpha)^j \alpha t_{n-j} + \dots + (1 - \alpha)^{n+1} \tau_0$$
- Since both α and $(1 - \alpha)$ are less than or equal to 1, each successive term has less weight than its predecessor

Operating System Concepts 5.17 Silberschatz, Galvin and Gagne ©2005

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

Operating System Concepts 5.18 Silberschatz, Galvin and Gagne ©2005

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 5.19 Silberschatz, Galvin and Gagne ©2005

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:

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

Operating System Concepts 5.20 Silberschatz, Galvin and Gagne ©2005

Time Quantum and Context Switch Time

process time = 10	quantum	context switches
	12	0
	6	1
	1	9

Operating System Concepts 5.21 Silberschatz, Galvin and Gagne ©2005

Turnaround Time Varies With The Time Quantum

process	time
P_1	6
P_2	3
P_3	1
P_4	7

Operating System Concepts 5.22 Silberschatz, Galvin and Gagne ©2005

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 5.23 Silberschatz, Galvin and Gagne ©2005

Multilevel Queue Scheduling

Operating System Concepts 5.24 Silberschatz, Galvin and Gagne ©2005

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 5.25 Silberschatz, Galvin and Gagne ©2005

Example of Multilevel Feedback Queue

- Three queues:
 - Q_0 – RR with time quantum 8 milliseconds
 - Q_1 – RR 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 5.26 Silberschatz, Galvin and Gagne ©2005

Multilevel Feedback Queues

Operating System Concepts 5.27 Silberschatz, Galvin and Gagne ©2005

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 5.28 Silberschatz, Galvin and Gagne ©2005

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 5.29 Silberschatz, Galvin and Gagne ©2005

Thread Scheduling

- Local Scheduling – How the threads library decides which thread to put onto an available LWP
- Global Scheduling – How the kernel decides which kernel thread to run next

Operating System Concepts 5.30 Silberschatz, Galvin and Gagne ©2005

Pthread Scheduling API

```

#include <pthread.h>
#include <stdio.h>
#define NUM_THREADS 5
int main(int argc, char *argv[])
{
    int i;
    pthread_t tid[NUM_THREADS];
    pthread_attr_t attr;
    /* get the default attributes */
    pthread_attr_t attr;
    /* set the scheduling algorithm to PROCESS or SYSTEM */
    pthread_attr_t attr;
    /* set the scheduling policy - FIFO, RT, or OTHER */
    pthread_attr_t attr;
    /* create the threads */
    for (i = 0; i < NUM_THREADS; i++)
        pthread_create(&tid[i], &attr, runner, NULL);
}

```

Operating System Concepts 5.31 Silberschatz, Galvin and Gagne ©2005

Pthread Scheduling API

```

/* now join on each thread */
for (i = 0; i < NUM_THREADS; i++)
    pthread_join(tid[i], NULL);
}
/* Each thread will begin control in this function */
void *runner(void *param)
{
    printf("I am a thread\n");
    pthread_exit(0);
}

```

Operating System Concepts 5.32 Silberschatz, Galvin and Gagne ©2005

Operating System Examples

- Solaris scheduling
- Windows XP scheduling
- Linux scheduling

Operating System Concepts 5.33 Silberschatz, Galvin and Gagne ©2005

Solaris 2 Scheduling

Operating System Concepts 5.34 Silberschatz, Galvin and Gagne ©2005

Solaris Dispatch Table

priority	time quantum	time quantum expired	return from sleep
0	200	0	50
5	200	0	50
10	160	0	51
15	160	5	51
20	120	10	52
25	120	15	52
30	80	20	53
35	80	25	54
40	40	30	55
45	40	35	56
50	40	40	58
55	40	45	58
59	20	49	59

Operating System Concepts 5.35 Silberschatz, Galvin and Gagne ©2005

Windows XP 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 5.36 Silberschatz, Galvin and Gagne ©2005

Linux Scheduling

- Two algorithms: time-sharing and real-time
- Time-sharing
 - Prioritized credit-based – process with most credits is scheduled next
 - Credit subtracted when timer interrupt occurs
 - When credit = 0, another process chosen
 - When all processes have credit = 0, recrediting occurs
 - ▶ Based on factors including priority and history
- Real-time
 - Soft real-time
 - Posix.1b compliant – two classes
 - ▶ FCFS and RR
 - ▶ Highest priority process always runs first

The Relationship Between Priorities and Time-slice length

numeric priority	relative priority	time quantum
0	highest	real-time tasks
•		
•		
•		
99		
100	other tasks	10 ms
•		
•		
•		
140		
	lowest	

List of Tasks Indexed According to Priorities

active array

priority	task lists
[0]	● ○ ○ ○ ○
[1]	● ○ ○ ○ ○
•	•
•	•
•	•
[140]	● ○ ○ ○ ○

expired array

priority	task lists
[0]	● ○ ○ ○ ○
[1]	● ○ ○ ○ ○
•	•
•	•
•	•
[140]	● ○ ○ ○ ○

Algorithm Evaluation

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

5.15

actual process execution

trace tape

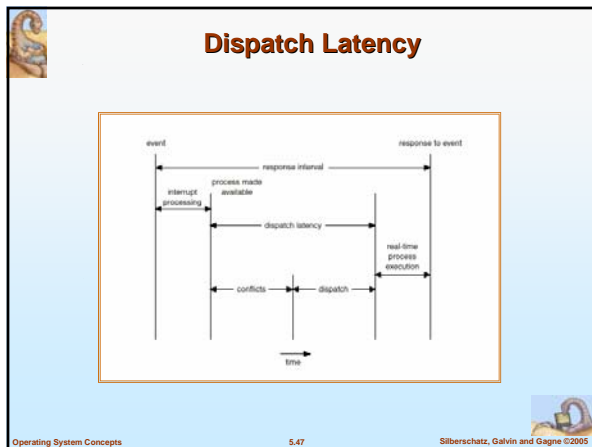
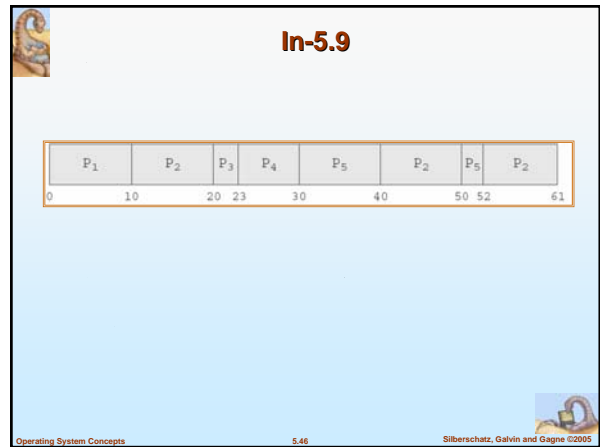

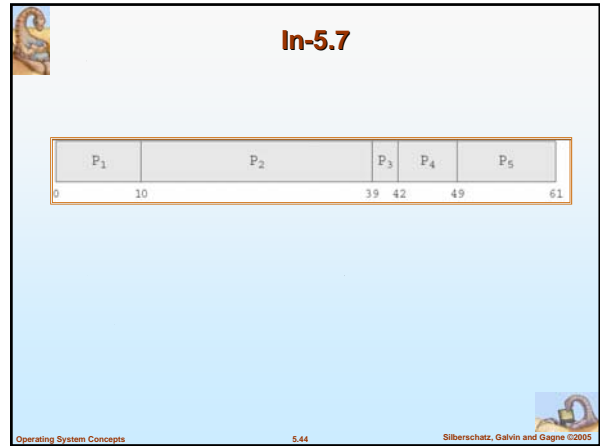
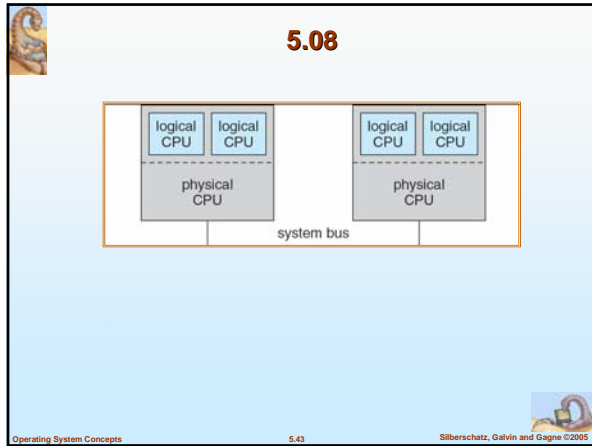
... CPU 10 IO 213 CPU 12 IO 112 CPU 2 IO 147 CPU 173 ...

simulation for FCFS → performance statistics for FCFS


simulation for SJF → performance statistics for SJF

simulation for RR (q = 14) → performance statistics for RR (q = 14)

End of Chapter 5



- ### Java Thread Scheduling
- JVM Uses a Preemptive, Priority-Based Scheduling Algorithm
 - FIFO Queue is Used if There Are Multiple Threads With the Same Priority
- Operating System Concepts 5.48 Silberschatz, Galvin and Gagne ©2005




Java Thread Scheduling (cont)


JVM Schedules a Thread to Run When:

1. The Currently Running Thread Exits the Runnable State
2. A Higher Priority Thread Enters the Runnable State

* Note – the JVM Does Not Specify Whether Threads are Time-Sliced or Not



Operating System Concepts 5.49 Silberschatz, Galvin and Gagne ©2005




Time-Slicing


Since the JVM Doesn't Ensure Time-Slicing, the yield() Method May Be Used:

```
while (true) {  
    // perform CPU-intensive task  
    ...  
    Thread.yield();  
}
```

This Yields Control to Another Thread of Equal Priority




Operating System Concepts 5.50 Silberschatz, Galvin and Gagne ©2005



Thread Priorities

<u>Priority</u>	<u>Comment</u>
Thread.MIN_PRIORITY	Minimum Thread Priority
Thread.MAX_PRIORITY	Maximum Thread Priority
Thread.NORM_PRIORITY	Default Thread Priority

Priorities May Be Set Using setPriority() method:
setPriority(Thread.NORM_PRIORITY + 2);



Operating System Concepts 5.51 Silberschatz, Galvin and Gagne ©2005