

Lecture 6

CPU scheduling II & Process Synchronization I

Lecture Information

Ceng328 *Operating Systems* at March 23, 2010

Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Dr. Cem Özdoğan
Computer Engineering Department
Çankaya University



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Contents

1 Scheduling Algorithms

First-Come, First-Served Scheduling

Shortest-Job-First Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue Scheduling

Multilevel Feedback-Queue Scheduling

2 Multiple-Processor Scheduling

Approaches to Multiple-Processor Scheduling

Load Balancing

3 Operating System Examples

Example: Linux Scheduling

4 Process Synchronization

5 Race Condition

6 The Critical-Section Problem

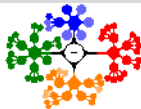
Disabling Interrupts: (Systems approach)

Lock Variables: (Software approach)

Lock Variables: (Software approach)

7 Mutual Exclusion with Busy Waiting

Strict Alternation (Software approach)



Scheduling Algorithms

First-Come, First-Served
SchedulingShortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
SchedulingMultilevel Feedback-Queue
SchedulingMultiple-Processor
SchedulingApproaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System
Examples

Example: Linux Scheduling

Process
Synchronization

Race Condition

The Critical-Section
Problem

Disabling Interrupts:

First-Come, First-Served Scheduling I

- By far the simplest CPU-scheduling algorithm is the first-come, first-served (FCFS) scheduling algorithm.
- When a process enters the ready queue, its PCB is linked onto the tail of the queue.
- The average waiting time under the FCFS policy, however, is often quite long. Consider the following set of processes that arrive at time 0, with the length of the CPU burst given in milliseconds:

Process	Burst Time	Waiting Time	Turnaround Time
P_1	24	0	24
P_2	3	24	27
P_3	3	27	30
Average	-	17	27

- If the processes arrive in the order P_1, P_2, P_3 , and are served in FCFS order, we get the result shown in the following chart:



First-Come, First-Served Scheduling II



- If the processes arrive in the order P_2, P_3, P_1 , the results will be as shown in the following chart:



- The average waiting time is now $(6 + 0 + 3)/3 = 3$ msecs.
- The average waiting time under an FCFS policy is generally not minimal.
- The FCFS scheduling algorithm is nonpreemptive.
- Once the CPU has been allocated to a process, that process keeps the CPU until it releases the CPU, either by terminating or by requesting I/O.

Scheduling Algorithms

First-Come, First-Served Scheduling

Shortest-Job-First Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue Scheduling

Multilevel Feedback-Queue Scheduling

Multiple-Processor Scheduling

Approaches to Multiple-Processor Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process

Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

First-Come, First-Served Scheduling III

- Assume we have one CPU-bound process and many I/O-bound processes.
 - The CPU-bound process will get and hold the CPU. During this time, all the other processes will finish their I/O and will move into the ready queue, waiting for the CPU.
 - While the processes wait in the ready queue, the I/O devices are idle. Eventually, the CPU-bound process finishes its CPU burst and moves to an I/O device.
 - All the I/O-bound processes, which have short CPU bursts, execute quickly and move back to the I/O queues.
 - At this point, the CPU sits idle. The CPU-bound process will then move back to the ready queue and be allocated the CPU.
 - Again, all the I/O processes end up waiting in the ready queue until the CPU-bound process is done.
 - There is a **convoy effect** as all the other processes wait for the one big process to get off the CPU. A long CPU-bound job may take the CPU and may force shorter (or I/O-bound) jobs to wait prolonged periods.
- This effect results in lower CPU and device utilization than might be possible if the shorter processes were allowed to go first.



Scheduling Algorithms

First-Come, First-Served Scheduling

Shortest-Job-First Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue Scheduling

Multilevel Feedback-Queue Scheduling

Multiple-Processor Scheduling

Approaches to Multiple-Processor Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling


Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

First-Come, First-Served Scheduling IV



Scheduling Algorithms

First-Come, First-Served Scheduling

Shortest-Job-First Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue Scheduling

Multilevel Feedback-Queue Scheduling

Multiple-Processor Scheduling

Approaches to Multiple-Processor Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:




Figure: An example to First-Come First-Served.

Shortest-Job-First Scheduling I

- This algorithm associates with each process the length of the process's next CPU burst.
- When the CPU is available, it is assigned to the process that has the smallest next CPU burst.
- As an example of SJF scheduling, consider the following set of processes, with the length of the CPU burst given in milliseconds:

Process	Burst Time	Waiting Time	Turnaround Time
P_1	6	3	9
P_2	8	16	24
P_3	7	9	16
P_4	3	0	3
Average	-	7	13



- By comparison, if we were using the FCFS scheduling scheme, the average waiting time would be 10.25 milliseconds.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization


Race Condition

The Critical-Section Problem

Disabling Interrupts:

Shortest-Job-First Scheduling II

- The SJF scheduling algorithm gives the minimum average waiting time for a given set of processes.
- The real difficulty with the SJF algorithm is knowing the length of the next CPU request.
- Although the SJF algorithm is optimal, it can not be implemented at the level of short-term CPU scheduling. There is no way to know the length of the next CPU burst.
- Also, long running jobs may starve for the CPU when there is a steady supply of short jobs.
- Example: In Fig. 2a, the average turnaround time is 14 minutes. Consider running these four jobs using SJF, as shown in Fig. 2b, the average turnaround time now becomes 11 minutes.



(a)



(b)

Figure: An example of shortest job first scheduling. (a) Running four jobs in the original order. (b) Running them in shortest job first order.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Shortest-Job-First Scheduling III





Figure: An example to Shortest Job First.

Scheduling Algorithms

First-Come, First-Served Scheduling

Shortest-Job-First Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue Scheduling

Multilevel Feedback-Queue Scheduling

Multiple-Processor Scheduling

Approaches to Multiple-Processor Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Shortest-Job-First Scheduling IV

- The SJF algorithm can be either **pre-emptive** or **nonpreemptive**.
- A pre-emptive SJF algorithm will preempt the currently executing process, whereas a nonpreemptive SJF algorithm will allow the currently running process to finish its CPU burst.

Process	Arrival Time	Burst Time	Waiting Time	Turnaround Time
P_1	0	8	9	17
P_2	1	4	0	4
P_3	2	9	15	24
P_4	3	5	2	7
Average	-	-	6.5	13

- If the processes arrive at the ready queue at the times shown and need the indicated burst times, then the resulting pre-emptive SJF schedule is as depicted in the following chart:



- Nonpreemptive SJF scheduling would result in an average waiting time of 7.75 milliseconds.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling


Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Shortest-Job-First Scheduling V



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition


The Critical-Section Problem

Disabling Interrupts:

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



$$\text{average waiting time} = (0 + 6 + 3 + 7)/4 = 4$$


COMP3300

13

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



$$\text{average waiting time} = (9 + 1 + 0 + 2)/4 = 3$$

COMP3300

14

Figure: Example of non-pre-emptive SJF and pre-emptive SJF.

Priority Scheduling I

- The SJF algorithm is a special case of the general **priority scheduling algorithm**.
- A priority is associated with each process, and the CPU is allocated to the process with the highest priority.
- Priorities are generally indicated by some fixed range of numbers, such as 0 to 7 or 0 to 4095. We use as low numbers represent high priority.

Process	Burst Time	Priority	Waiting Time	Turnaround Time
P_1	10	3	6	16
P_2	1	1	0	1
P_3	2	4	16	18
P_4	1	5	18	19
P_5	5	2	1	6
Average	-	-	8.2	12



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Priority Scheduling II

- Priorities can be defined either **internally** or **externally**.
 - Internally defined priorities use some measurable quantity or quantities to compute the priority of a process.
 - External priorities are set by criteria outside the OS (such as the importance of the process).
- Priority scheduling can be either **pre-emptive** or **nonpreemptive**.
 - A pre-emptive priority scheduling algorithm will preempt the CPU if the priority of the newly arrived process is higher than the priority of the currently running process.
 - A nonpreemptive priority scheduling algorithm will simply put the new process at the head of the ready queue.
- A major problem with priority scheduling algorithms is **indefinite blocking**, or **starvation**. A process that is ready to run but waiting for the CPU can be considered blocked.
 - A priority scheduling algorithm can leave some low priority processes waiting indefinitely.
 - In a heavily loaded computer system, a steady stream of higher-priority processes can prevent a low-priority process from ever getting the CPU.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Priority Scheduling III



- It is often convenient to group processes into priority classes and use priority scheduling among the classes but round-robin scheduling within each class. Figure 5 shows a system with four priority classes.

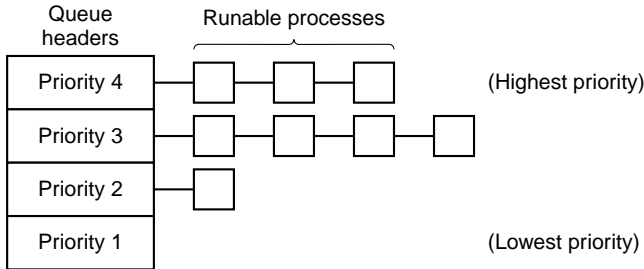


Figure: A scheduling algorithm with four priority classes.

Scheduling Algorithms

First-Come, First-Served
Scheduling
Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling
Multilevel Queue
Scheduling
Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling
Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization


Race Condition

The Critical-Section Problem

Disabling Interrupts:

Priority Scheduling IV

- A solution to the problem of indefinite blockage of low-priority processes is **aging**.
- Aging is a technique of gradually increasing the priority of processes that wait in the system for a long time.
 - For example, if priorities range from 127 (low) to 0 (high), we could increase the priority of a waiting process by 1 every 15 minutes.
 - Eventually, even a process with an initial priority of 127 would have the highest priority in the system and would be executed.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Priority Scheduling V




Figure: An example to Priority-based Scheduling.

Scheduling Algorithms

First-Come, First-Served Scheduling

Shortest-Job-First Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue Scheduling

Multilevel Feedback-Queue Scheduling

Multiple-Processor Scheduling

Approaches to Multiple-Processor Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Round-Robin Scheduling I

- The **round-robin (RR) scheduling algorithm** is designed especially for time-sharing systems.
- It is similar to FCFS scheduling, but pre-emption is added to switch between processes.
- A small unit of time, called a **time quantum** or **time slice**, is defined.
- A time quantum is generally from 10 to 100 milliseconds.
- The ready queue is treated as a circular queue.


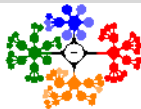


Figure: Round-robin scheduling. (a) The list of runnable processes. (b) The list of runnable processes after B uses up its quantum.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Round-Robin Scheduling II

- To implement RR scheduling,
 - we keep the ready queue as a FIFO queue of processes.
 - New processes are added to the tail of the ready queue.
 - The CPU scheduler picks the first process from the ready queue, sets a timer to interrupt after 1 time quantum, and dispatches the process.
 - The process may have a CPU burst of less than 1 time quantum.
 - In this case, the process itself will release the CPU voluntarily.
 - The scheduler will then proceed to the next process in the ready queue.
 - Otherwise, if the CPU burst of the currently running process is longer than 1 time quantum,
 - the timer will go off and will cause an interrupt to the OS.
 - A context switch will be executed, and the process will be put at the tail of the ready queue.
 - The CPU scheduler will then select the next process in the ready queue.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling


Process Synchronization

Race Condition

The Critical-Section Problem


Disabling Interrupts:

Round-Robin Scheduling II



- The average waiting time under the RR policy is often long.
- Consider the following set of processes that arrive at time 0, with the length of the CPU burst given in milliseconds: (a time quantum of 4 milliseconds)

Process	Burst Time	Waiting Time	Turnaround Time
P_1	24	6	30
P_2	3	4	7
P_3	3	7	10
Average	-	5.66	15.66



Scheduling Algorithms

First-Come, First-Served Scheduling

Shortest-Job-First Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue Scheduling

Multilevel Feedback-Queue Scheduling

Multiple-Processor Scheduling

Approaches to Multiple-Processor Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Round-Robin Scheduling IV

- In the RR scheduling algorithm, no process is allocated the CPU for more than 1 time quantum in a row (unless it is the only runnable process).
- The performance of the RR algorithm depends heavily on the size of the time quantum.
 - If the time quantum is extremely large, the RR policy is the same as the FCFS policy.
 - If the time quantum is extremely small (say, 1 millisecond), the RR approach is called **processor sharing** and (in theory) creates the appearance that each of n processes has its own processor running at $\frac{1}{n}$ the speed of the real processor.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process

Synchronization

Race Condition

The Critical-Section
Problem

Disabling Interrupts:

Round-Robin Scheduling V

- We need also to consider the effect of context switching on the performance of RR scheduling (see Fig. 8).
 - Let us assume that we have only one process of 10 time units.
 - If the quantum is 12 time units, the process finishes in less than 1 time quantum, with no overhead.
 - If the quantum is 6 time units, however, the process requires 2 quanta, resulting in a context switch.
 - If the time quantum is 1 time unit, then nine context switches will occur, slowing the execution of the process accordingly





Figure: The way in which a smaller time quantum increases context switches.

Round-Robin Scheduling VI

- Thus, we want the time quantum to be large with respect to the context-switch time.
 - If the context-switch time is approximately 10 percent of the time quantum, then about 10 percent of the CPU time will be spent in context switching.
 - In practice, most modern systems have time quanta ranging from 10 to 100 milliseconds.
 - The time required for a context switch is typically less than 10 microseconds; thus, the context-switch time is a small fraction of the time quantum.
- Setting the quantum too short causes too many process switches and lowers the CPU efficiency, but setting it too long may cause poor response to short interactive requests.
- Although the time quantum should be large compared with the context-switch time, it should not be too large. If the time quantum is too large, RR scheduling degenerates to FCFS policy.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Round-Robin Scheduling VII





Figure: An example to Round Robin.

Scheduling Algorithms

First-Come, First-Served Scheduling

Shortest-Job-First Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue Scheduling

Multilevel Feedback-Queue Scheduling

Multiple-Processor Scheduling

Approaches to Multiple-Processor Scheduling

Load Balancing

Operating System Examples


Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:



Scheduling Algorithms

First-Come, First-Served
SchedulingShortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
SchedulingMultilevel Feedback-Queue
SchedulingMultiple-Processor
SchedulingApproaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System
Examples

Example: Linux Scheduling

Process
Synchronization

Race Condition

The Critical-Section
Problem

Disabling Interrupts:

Multilevel Queue Scheduling I

- Another class of scheduling algorithms has been created for situations in which processes are easily classified into different groups.
- A common division is made between **foreground (interactive)** processes and **background (batch)** processes.
- A multilevel queue scheduling algorithm partitions the ready queue into several separate queues (see Fig. 10).

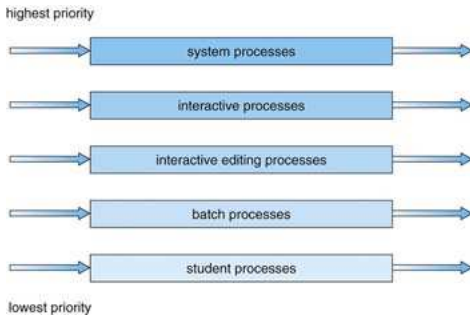


Figure: Multilevel queue scheduling.

Multilevel Queue Scheduling II

- The processes are permanently assigned to one queue, generally based on some property of the process, such as memory size, process priority, or process type.
- Each queue has absolute priority over lower-priority queues and also each queue has its own scheduling algorithm.
- The foreground queue might be scheduled by an RR algorithm, while the background queue is scheduled by an FCFS algorithm.
- In addition, there must be scheduling among the queues, which is commonly implemented as fixed-priority pre-emptive scheduling.
- For example, the foreground queue may have absolute priority over the background queue.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Multilevel Feedback-Queue Scheduling I

- If there are separate queues for foreground and background processes, processes do not move from one queue to the other.
- This setup has the advantage of low scheduling overhead, but it is inflexible.
- The **multilevel feedback-queue scheduling algorithm**, in contrast, allows a process to move between queues.
- The idea is to separate processes according to the characteristics of their CPU bursts.
 - If a process uses too much CPU time, it will be moved to a lower-priority queue.
 - This scheme leaves I/O-bound and interactive processes in the higher-priority queues.
 - In addition, a process that waits too long in a lower-priority queue may be moved to a higher-priority queue.
- The definition of a multilevel feedback-queue scheduler makes it the most general CPU-scheduling algorithm. Unfortunately, it is also the most complex algorithm.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Multilevel Feedback-Queue Scheduling II

- This form of **aging** prevents starvation (see Fig. 11).




Figure: Multilevel feedback queues.

- A process entering the ready queue is put in queue 0. A process in queue 0 is given a time quantum of 8 milliseconds.
- If it does not finish within this time, it is moved to the tail of queue 1.
- If queue 0 is empty, the process at the head of queue 1 is given a quantum of 16 milliseconds.
- If it does not complete, it is preempted and is put into queue 2.
- Processes in queue 2 are run on an FCFS basis but are run only when queues 0 and 1 are empty.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Approaches to Multiple-Processor Scheduling

- One approach to CPU scheduling in a multiprocessor system has all scheduling decisions, I/O processing, and other system activities handled by a single processor - the master server.
- This **asymmetric multiprocessing** is simple because only one processor accesses the system data structures, reducing the need for data sharing.
- A second approach uses **symmetric multiprocessing (SMP)**, where each processor is self-scheduling.
 - All processes may be in a common ready queue, or each processor may have its own private queue of ready processes.
 - Regardless, scheduling proceeds by having the scheduler for each processor examine the ready queue and select a process to execute.
- if we have multiple processors trying to access and update a common data structure, we must ensure that two processors do not choose the same process and that processes are not lost from the queue.



Scheduling Algorithms

First-Come, First-Served Scheduling
Shortest-Job-First Scheduling
Priority Scheduling
Round-Robin Scheduling
Multilevel Queue Scheduling
Multilevel Feedback-Queue Scheduling

Multiple-Processor Scheduling

Approaches to Multiple-Processor Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:



- Load balancing attempts to keep the workload evenly distributed across all processors in an SMP system.
 - Load balancing is typically only necessary on systems where each processor has its own private queue of eligible processes to execute.
 - On systems with a common run queue, load balancing is often unnecessary, because once a processor becomes idle, it immediately extracts a runnable process from the common run queue.
- It is important to note that in most contemporary OSs supporting SMP, each processor does have a private queue of eligible processes.

Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples


Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:



- There are two general approaches to load balancing: **push migration** and **pull migration**.
 - With push migration, a specific task periodically checks the load on each processor and -if it finds an imbalance- evenly distributes the load by moving (or pushing) processes from overloaded to idle or less-busy processors.
 - Pull migration occurs when an idle processor pulls a waiting task from a busy processor.
- Linux runs its load balancing algorithm every 200 milliseconds (push migration) or whenever the run queue for a processor is empty (pull migration).

Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:



Scheduling Algorithms

First-Come, First-Served
SchedulingShortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
SchedulingMultilevel Feedback-Queue
SchedulingMultiple-Processor
SchedulingApproaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System
Examples

Example: Linux Scheduling

Process

Synchronization

Race Condition

The Critical-Section
Problem

Disabling Interrupts:


Example: Linux Scheduling I

- The Linux scheduler is a pre-emptive, priority-based algorithm with two separate priority ranges:
 - a real-time range from 0 to 99
 - and a nice value ranging from 100 to 140.
- The relationship between priorities and time-slice length is shown in Fig. 12.

<u>numeric priority</u>	<u>relative priority</u>		<u>time quantum</u>
0	highest	realtime tasks	200 ms
·			
·			
99			
100		other tasks	10 ms
·			
·			
140	lowest		

Figure: The relationship between priorities and time-slice length.

- A runnable task is considered eligible for execution on the CPU as long as it has time remaining in its time slice.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor
Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System
Examples

Example: Linux Scheduling

Process

Synchronization

Race Condition

The Critical-Section
Problem

Disabling Interrupts:

Example: Linux Scheduling II

- The kernel maintains a list of all runnable tasks in a **runqueue** data structure.
- Each runqueue contains two priority arrays -**active** and **expired**.
 - The active array contains all tasks with time remaining in their time slices,
 - and the expired array contains all expired tasks.


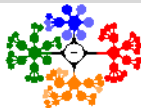


Figure: List of tasks indexed according to priority.

- Each of these priority arrays contains a list of tasks indexed according to priority (see Fig. 13).

Race Condition I



- A potential problem; the order of instructions of cooperating processes (see Table 1).

Table: Race Condition

Process A	Process B	concurrent access
$X = 1;$	$Y = 2;$	<i>does not matter</i>
$X = Y + 1;$	$Y = Y * 2;$	<i>important!</i>

- A **race condition** is a situation where two or more processes access shared data concurrently and final value of shared data depends on *timing* (*race* to access and modify data)
- To guard against the race condition above, we need to ensure that only one process at a time can be manipulating the variable counter (**process synchronization**).

Scheduling Algorithms

First-Come, First-Served Scheduling

Shortest-Job-First Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue Scheduling

Multilevel Feedback-Queue Scheduling

Multiple-Processor Scheduling

Approaches to Multiple-Processor Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

- **Producer-consumer problem.** It is described that how a bounded buffer could be used to enable processes to share memory
 - **Bounded buffer problem.** The solution allows at most $BUFFER_SIZE - 1$ items in the buffer at the same time.
 - An integer variable *counter*, initialized to 0. *counter* is incremented every time we add a new item to the buffer and is decremented every time we remove one item from the buffer.
- Although both the producer and consumer routines are correct separately, *they may not function correctly when executed concurrently.*
- We would arrive at incorrect state because we allowed both processes to manipulate the variable *counter* concurrently.



Scheduling Algorithms

First-Come, First-Served Scheduling
Shortest-Job-First Scheduling
Priority Scheduling
Round-Robin Scheduling
Multilevel Queue Scheduling
Multilevel Feedback-Queue Scheduling

Multiple-Processor Scheduling

Approaches to Multiple-Processor Scheduling
Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Race Condition III

- The code for the producer process:

```
while (true)
{
    /* produce an item in next Produced */
    while (counter == BUFFER_SIZE)
        ; /* do nothing */
    buffer [in] = nextProduced;
    in = (in + 1) % BUFFER_SIZE;
    counter++;
}
```

- The code for the consumer process:

```
while (true)
{
    while (counter == 0)
        ; /* do nothing */
    nextConsumed = buffer [out] ;
    out = (out + 1) % BUFFER_SIZE;
    counter--;
    /* consume the item in nextConsumed */
}
```



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

The Critical-Section Problem I

- How do we avoid *race conditions*? What we need is **mutual exclusion** (see Fig. 14).

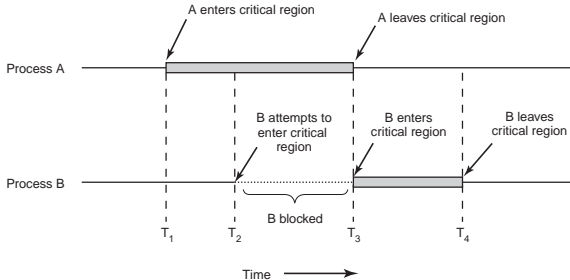


Figure: Mutual exclusion using critical regions.

- Consider a system consisting of n processes. Each process has a segment of code, called a **critical section (CS)**, in which the process may be changing common variables, updating a table, writing a file, and so on.
- The important feature of the system is that, when one process is executing in its CS, no other process is to be allowed to execute in its CS.



Scheduling Algorithms

First-Come, First-Served Scheduling
Shortest-Job-First Scheduling
Priority Scheduling
Round-Robin Scheduling
Multilevel Queue Scheduling
Multilevel Feedback-Queue Scheduling

Multiple-Processor Scheduling

Approaches to Multiple-Processor Scheduling
Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

The Critical-Section Problem II

- That is, no two processes are executing in their CSs at the same time.
- Each process must request permission to enter its CS. The section of code implementing this request is the **entry section**.
- The CS may be followed by an **exit section**.
- The remaining code is the **remainder section** (see Fig. 15).

```
do {  
    entry section  
    critical section  
    exit section  
    remainder section  
} while (TRUE);
```

Figure: General structure of a typical process P_i .



Scheduling Algorithms

First-Come, First-Served
Scheduling
Shortest-Job-First
Scheduling
Priority Scheduling
Round-Robin Scheduling
Multilevel Queue
Scheduling
Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling
Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

The Critical-Section Problem III



- A solution to the CS problem must satisfy the following requirements:
 - 1 **Mutual exclusion.** If process P_i is executing in its CS, then no other processes can be executing in their CSs.
 - 2 **Progress.** If no process is executing in its CS and some processes wish to enter their CSs, then only those processes that are not executing in their remainder sections can participate in the decision on which will enter its CS next, and this selection cannot be postponed indefinitely. (No process should have to wait forever to enter its CS.)
 - 3 **Bounded waiting.** There exists a bound, or limit, on the number of times that other processes are allowed to enter their CSs after a process has made a request to enter its CS and before that request is granted.
 - 4 **Fault tolerance.** Process running outside its CR should not block other processes accessing the CR.
 - 5 **No assumptions** may be made about speeds or the number of CPUs.

Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

The Critical-Section Problem IV



- **Atomic operation.** Atomic means either an operation happens in its entirely or NOT at all (it cannot be interrupted in the middle).
- Machine instructions are atomic, high level instructions are not (count++; this is actually 3 machine level instructions, an interrupt can occur in the middle of instructions).
- Various proposals for achieving mutual exclusion, so that while one process is busy updating shared memory in its CS, no other process will enter its CS and cause trouble.
 - Disabling Interrupts
 - Lock Variables
 - Strict Alternation
 - Peterson's Solution
 - The TSL instructions (Hardware approach)

Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Disabling Interrupts: (Systems approach)

- The simplest solution is to have each process disable all interrupts just after entering its CS and re-enable them just before leaving it.
- With interrupts disabled, no clock interrupts can occur (The CPU is only switched from process to process as a result of clock or other interrupts)
- With interrupts turned off the CPU will not be switched to another process!! Thus, once a process has disabled interrupts, it can examine and update the shared memory without fear that any other process will intervene.
- This approach is generally unattractive because it is unwise to give user processes the power to turn off interrupts. Suppose that one of them did it and never turned them on again? That could be the end of the system.
- On the other hand, it is frequently convenient for the kernel itself to disable interrupts for a few instructions while it is updating variables or lists.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

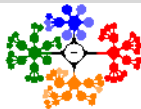
Disabling Interrupts:

Disabling Interrupts: (Systems approach)

```
do {  
    acquire lock  
    critical section  
    release lock  
    remainder section  
} while (TRUE);
```

Figure: Solution to the critical-section problem using locks.

- Consider having a single, shared (lock) variable, initially 0.
 - When a process wants to enter its CS, it first tests the lock.
 - If the lock is 0, the process sets it to 1 and enters the CS.
 - If the lock is already 1, the process just waits until it becomes 0.
 - Thus, a 0 means that no process is in its CS, and a 1 means that some process is in its CS.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Disabling Interrupts: (Systems approach)



- Unfortunately, this idea contains a fatal flaw;
 - Suppose that one process reads the lock and sees that it is 0.
 - Before it can set the lock to 1, another process is scheduled, runs, and sets the lock to 1.
 - When the first process runs again, it will also set the lock to 1, and two processes will be in their CSs at the same time.

Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling


Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts:

Strict Alternation (Software approach) I



- **Busy waiting** (notice the semicolons terminating the while statements in Fig. 17); continuously testing a variable until some value appears, a lock that uses busy waiting is called a **spin lock**.
- It should usually be avoided, since it wastes CPU time.

```
while (TRUE) {  
    while (turn != 0)    /* loop */ ;  
    critical_region();  
    turn = 1;  
    noncritical_region();  
}
```

(a)

```
while (TRUE) {  
    while (turn != 1)    /* loop */ ;  
    critical_region();  
    turn = 0;  
    noncritical_region();  
}
```

(b)

Figure: A proposed solution to the critical region problem. (a) Process 0. (b) Process 1.

Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization


Race Condition

The Critical-Section Problem

Disabling Interrupts:

Strict Alternation (Software approach) II

- the integer variable **turn** (keeps track of whose turn it is to enter the CR),
- initially, process 0 inspects turn, finds it to be 0, and enters its CR,
- process 1 also finds it to be 0 and therefore sits in a tight loop continually testing turn to see when it becomes,
- when process 0 leaves the CR, it sets turn to 1, to allow process 1 to enter its CR,
- suppose that process 1 finishes its CR quickly, so both processes are in their nonCR (with turn set to 0)
- process 0 finishes its nonCR and goes back to the top of its loop. Process 0 executes its whole loop quickly, exiting its CR and setting turn to 1.
- at this point turn is 1 and both processes are executing in their nonCR,
- process 0 finishes its nonCR and goes back to the top of its loop,
- unfortunately, it is not permitted to enter its CR, **turn** is 1 and process 1 is busy with its nonCR,
- it hangs in its while loop until process 1 sets turn to 0,
- this algorithm does avoid all races. But violates condition **Fault tolerance**.



Scheduling Algorithms

First-Come, First-Served
Scheduling

Shortest-Job-First
Scheduling

Priority Scheduling

Round-Robin Scheduling

Multilevel Queue
Scheduling

Multilevel Feedback-Queue
Scheduling

Multiple-Processor Scheduling

Approaches to
Multiple-Processor
Scheduling

Load Balancing

Operating System Examples

Example: Linux Scheduling

Process Synchronization

Race Condition

The Critical-Section Problem

Disabling Interrupts: