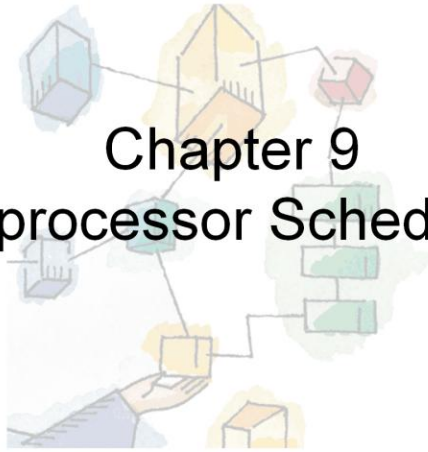


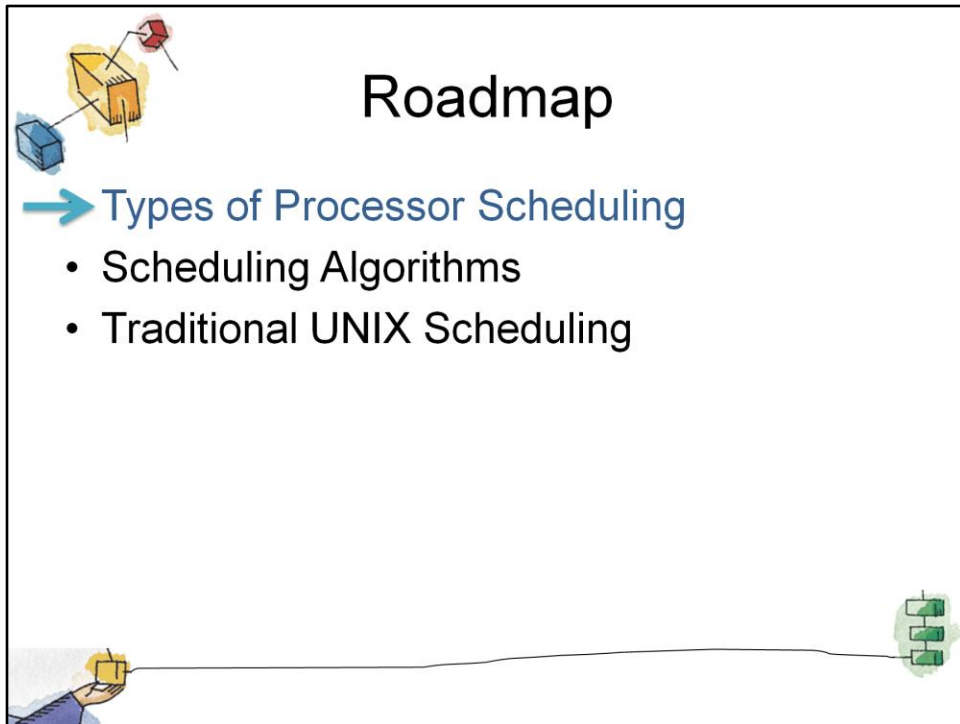
*Operating Systems:  
Internals and Design Principles, 6/E*  
William Stallings

## Chapter 9 Uniprocessor Scheduling



Dave Bremer  
Otago Polytechnic, N.Z.  
©2008, Prentice Hall

These slides are intended to help a teacher develop a presentation. This PowerPoint covers the entire chapter and includes too many slides for a single delivery. Professors are encouraged to adapt this presentation in ways which are best suited for their students and environment.



This chapter is concerned with scheduling on a system with a single processor.

It is possible to define and clarify many design issues related to scheduling in this context.

We examine the three types of processor scheduling:

- long term,
- medium term, and
- short term.

The bulk of the chapter focuses on short-term scheduling issues.

# Scheduling

- An OS must allocate resources amongst competing processes.
- The resource provided by a processor is execution time
  - The resource is allocated by means of a schedule

An operating system must allocate computer resources among the potentially competing requirements of multiple processes.

In the case of the processor, the resource to be allocated is execution time on the processor and the means of allocation is scheduling.



## Overall Aim of Scheduling

- The aim of processor scheduling is to assign processes to be executed by the processor over time,
  - in a way that meets system objectives, such as response time, throughput, and processor efficiency.




## Scheduling Objectives

- The scheduling function should
  - Share time *fairly* among processes
  - Prevent starvation of a process
  - Use the processor efficiently
  - Have low overhead
  - Prioritise processes when necessary (e.g. real time deadlines)

The scheduling function must be designed to satisfy a number of objectives, including fairness, lack of starvation of any particular process, efficient use of processor time, and low overhead.


In addition, the scheduling function may need to take into account different levels of priority or real-time deadlines for the start or completion of certain processes.



# Types of Scheduling

**Table 9.1 Types of Scheduling**

<b>Long-term scheduling</b>	The decision to add to the pool of processes to be executed
<b>Medium-term scheduling</b>	The decision to add to the number of processes that are partially or fully in main memory
<b>Short-term scheduling</b>	The decision as to which available process will be executed by the processor
<b>I/O scheduling</b>	The decision as to which process's pending I/O request shall be handled by an available I/O device



**Long-term scheduling** is performed when a new process is created.

- This is a decision whether to add a new process to the set of processes that are currently active.

**Medium-term scheduling** is a part of the swapping function.

- This is a decision whether to add a process to those that are at least partially in main memory and therefore available for execution.

**Short-term scheduling** is the actual decision of which ready process to execute next.

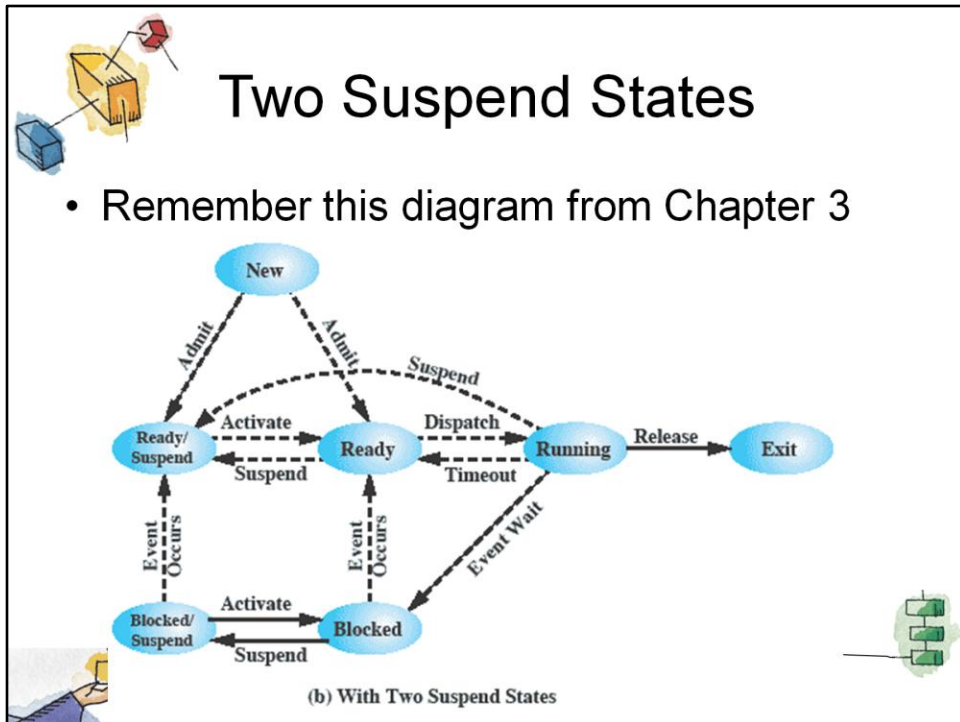
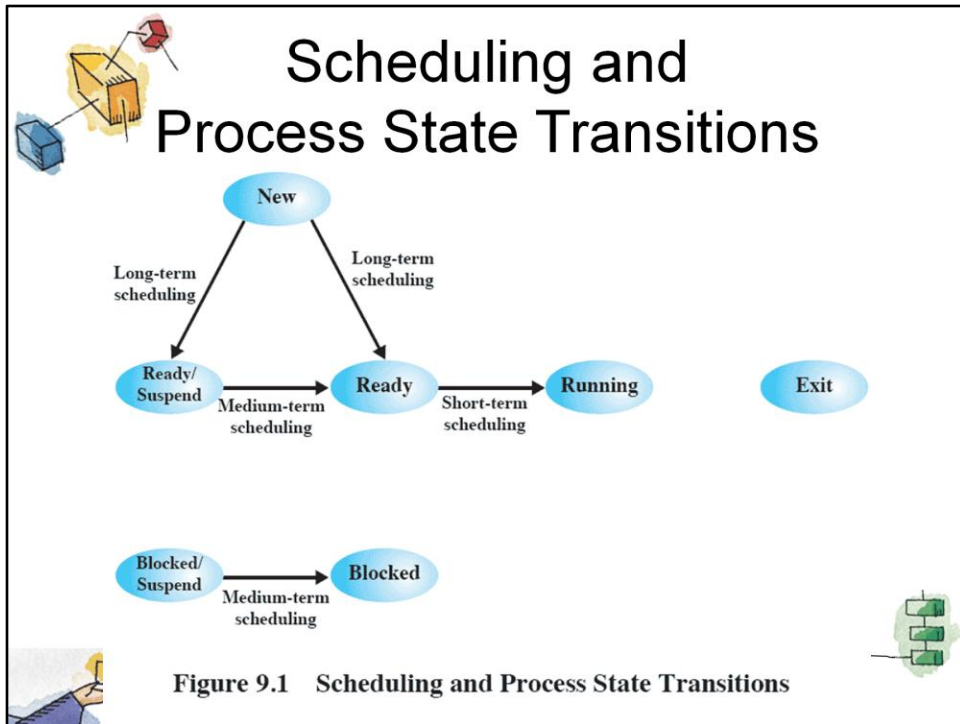


Figure 3.9b

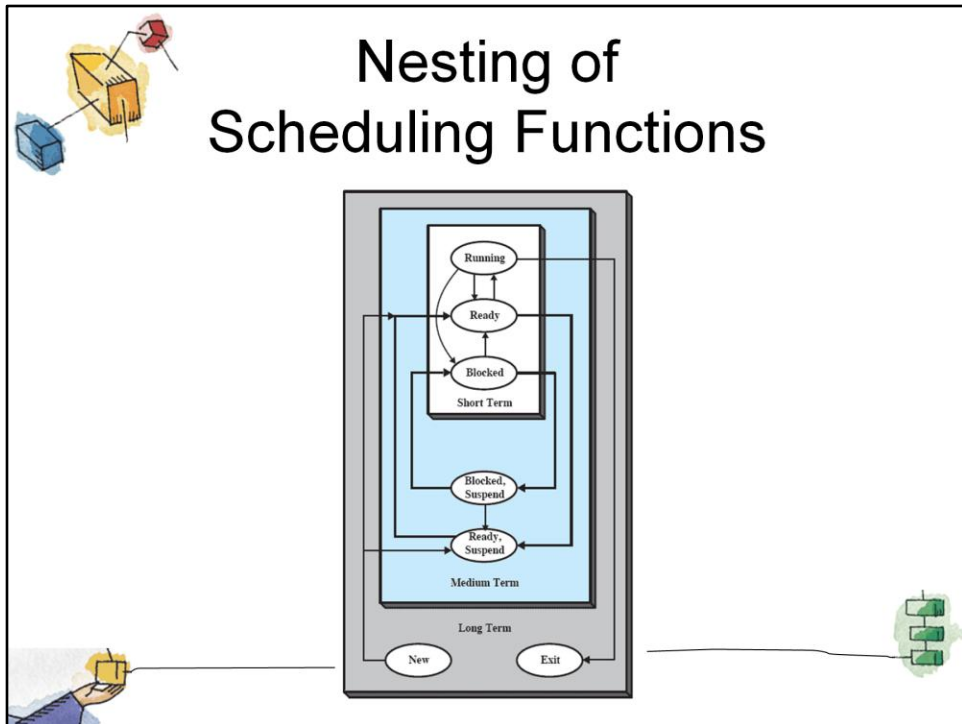
Two suspend states allow all processes which are not actually running to be swapped.

- **Ready:** The process is in main memory and available for execution.
- **Blocked:** The process is in main memory and awaiting an event.
- **Blocked/Suspend:** The process is in secondary memory and awaiting an event.
- **Ready/Suspend:** The process is in secondary memory but is available for execution as soon as it is loaded into main memory.

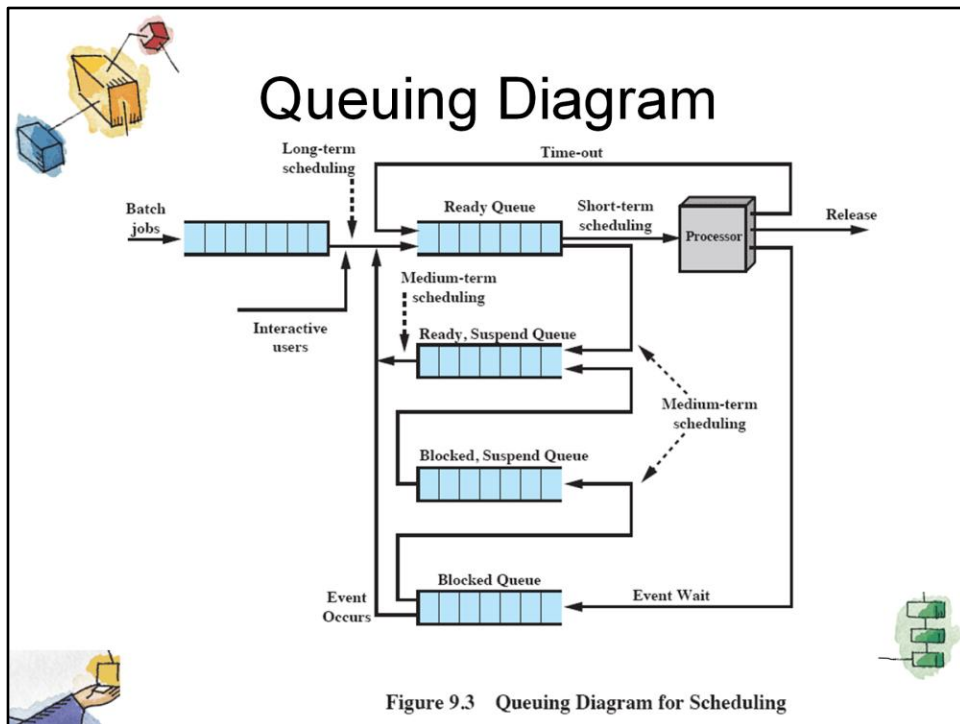


This figure relates the scheduling functions to the process state transition diagram shown in chapter 3 (fig 3.9b)





This figure reorganizes the state transition diagram to suggest the nesting of scheduling functions.

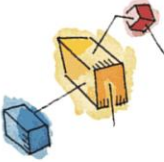


This figure shows that scheduling affects the performance of the system because it determines which processes will wait and which will progress.

The figure shows the queues involved in the state transitions of a process.



- For simplicity, new processes are shown going directly to the Ready state, whereas the earlier figures (9.1 and 9.2) show the option of either the Ready state or the Ready/Suspend state.

Fundamentally, scheduling is a matter of managing queues to minimize queuing delay and to optimize performance in a queuing environment.



# Long-Term Scheduling

- Determines which programs are admitted to the system for processing
  - May be first-come-first-served
  - Or according to criteria such as priority, I/O requirements or expected execution time
- Controls the degree of multiprogramming
- More processes, smaller percentage of time each process is executed



The long-term scheduler determines which programs are admitted to the system for processing.

- Thus, it controls the degree of multiprogramming.

In terms of frequency of execution, the long-term scheduler executes relatively infrequently and makes the coarse-grained decision of whether or not to take on a new process and which one to take.

Once admitted, a **job** or user program becomes a process and is added to the queue for the short-term scheduler.

- Or (as on some systems) a newly created process begins in a swapped-out condition, in which case it is added to a queue for the medium-term scheduler.

The decision as to which job to admit next can be on a simple first-come-first-served basis,

- or it can be a tool to manage system performance.

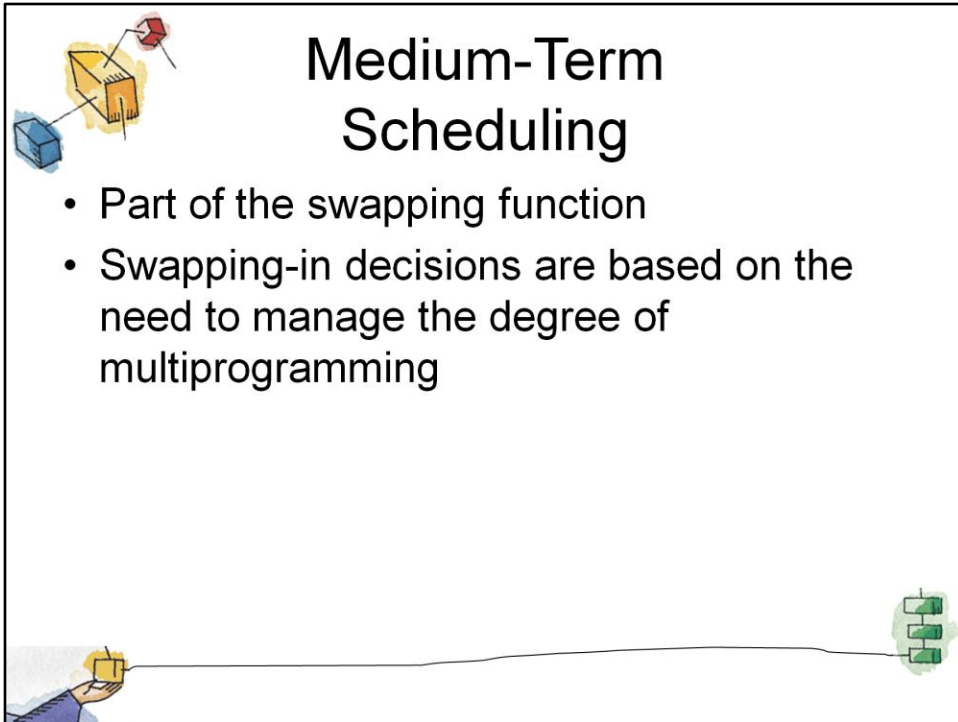
The criteria used may include

- priority,
- expected execution time, and
- I/O requirements.

The decision as to when to create a new process is generally driven by the desired degree of multiprogramming.

The more processes that are created, the smaller is the percentage of time that each process can be executed

- i.e. more processes are competing for the same amount of processor time.

An illustration within a black border. In the top-left corner, there is a yellow box with a red dot and a blue box. In the bottom-left corner, a hand in a blue sleeve holds a yellow box. In the bottom-right corner, there is a green server rack. A thin black line connects the yellow box in the bottom-left to the server rack in the bottom-right.

## Medium-Term Scheduling

- Part of the swapping function
- Swapping-in decisions are based on the need to manage the degree of multiprogramming

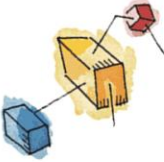
The medium-term scheduler is executed somewhat more frequently.

Medium-term scheduling is part of the swapping function.

Typically, the swapping-in decision is based on the need to manage the degree of multiprogramming.


On a system that does not use virtual memory, memory management is also an issue.

- Thus, the swapping-in decision will consider the memory requirements of the swapped-out processes.



## Short-Term Scheduling

- Known as the dispatcher
- Executes most frequently
- Invoked when an event occurs
  - Clock interrupts
  - I/O interrupts
  - Operating system calls
  - Signals



Also known as the dispatcher, executes most frequently and makes the fine-grained decision of which process to execute next.

The short-term scheduler is invoked whenever an event occurs that may lead to the blocking of the current process or that may provide an opportunity to preempt a currently running process in favor of another.

Examples of such events include

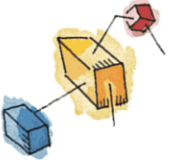
- Clock interrupts
- I/O interrupts
- Operating system calls
- Signals (e.g., semaphores)



# Roadmap


- Types of Processor Scheduling
- Scheduling Algorithms
- Traditional UNIX Scheduling





## Aim of Short Term Scheduling

- Main objective is to allocate processor time to optimize certain aspects of system behaviour.
- A set of criteria is needed to evaluate the scheduling policy.



The main objective of short-term scheduling is to allocate processor time in such a way as to optimize one or more aspects of system behavior.

Generally, a set of criteria is established against which various scheduling policies may be evaluated.

## Short-Term Scheduling Criteria: User vs System

- We can differentiate between user and system criteria
- User-oriented
  - Response Time
    - Elapsed time between the submission of a request until there is output.
- System-oriented
  - Effective and efficient utilization of the processor

The slide features several decorative icons: a blue cube, a yellow cube with a red dot, a red cube, a hand holding a yellow cube, and a green circuit-like icon.

The commonly used criteria can be categorized along two dimensions.

We can make a distinction between user-oriented and system-oriented criteria.

**User** oriented criteria relate to the behavior of the system as perceived by the individual user or process.

- E.g. response time in an interactive system which is the elapsed time between the submission of a request until the response begins to appear as output.
- We would like a scheduling policy that provides “good” service to various users. A threshold may be defined, say 2 seconds. Then a goal of the scheduling mechanism should be to maximize the number of users who experience an average response time of 2 seconds or less.

**System oriented** criteria is the focused on effective and efficient utilization of the processor.

- An example is throughput, i.e. the rate at which processes are completed. Thus, throughput is of concern to a system administrator but not to the user population.

On a single-user system, it probably is not important to achieve high processor utilization or high throughput as long as the responsiveness of the system to user applications is acceptable.



## Short-Term Scheduling Criteria: Performance

- We could differentiate between performance related criteria, and those unrelated to performance
- Performance-related
  - Quantitative, easily measured
  - E.g. response time and throughput
- Non-performance related
  - Qualitative
  - Hard to measure

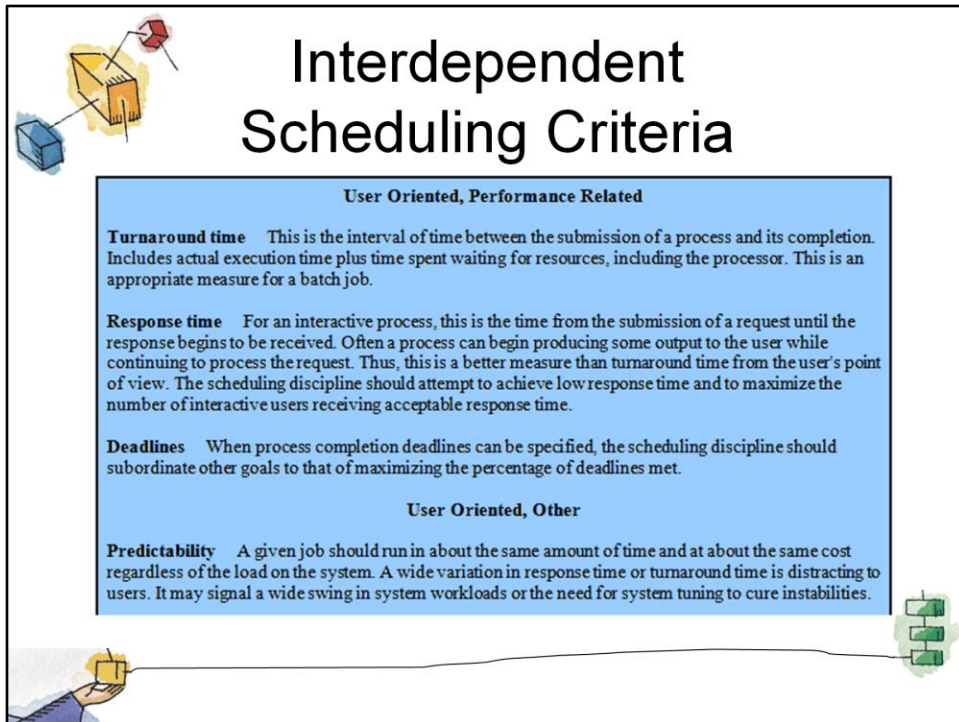
Another dimension along which criteria can be classified is those that are performance related and those that are not directly performance related.

Performance-related criteria are quantitative and generally can be readily measured.

- e.g. include response time and throughput.

Criteria that are not performance related are either qualitative in nature or do not lend themselves readily to measurement and analysis.

- e.g. predictability. We would like for the service provided to users to exhibit the same characteristics over time, independent of other work being performed by the system.
- To some extent, this criterion can be measured, by calculating variances as a function of workload. However, this is not nearly as straightforward as measuring throughput or response time as a function of workload.



This table summarizes key scheduling criteria.

These are interdependent, and it is impossible to optimize all of them simultaneously.

- Thus, the design of a scheduling policy involves compromising among competing requirements;
- The relative weights given the various requirements will depend on the nature and intended use of the system.

In most interactive operating systems, whether single user or time shared, adequate response time is the critical requirement.



# Interdependent Scheduling Criteria cont.

## System Oriented, Performance Related

**Throughput** The scheduling policy should attempt to maximize the number of processes completed per unit of time. This is a measure of how much work is being performed. This clearly depends on the average length of a process but is also influenced by the scheduling policy, which may affect utilization.

**Processor utilization** This is the percentage of time that the processor is busy. For an expensive shared system, this is a significant criterion. In single-user systems and in some other systems, such as real-time systems, this criterion is less important than some of the others.

## System Oriented, Other

**Fairness** In the absence of guidance from the user or other system-supplied guidance, processes should be treated the same, and no process should suffer starvation.

**Enforcing priorities** When processes are assigned priorities, the scheduling policy should favor higher-priority processes.

**Balancing resources** The scheduling policy should keep the resources of the system busy. Processes that will underutilize stressed resources should be favored. This criterion also involves medium-term and long-term scheduling.



# Priorities

- Scheduler will always choose a process of higher priority over one of lower priority
- Have multiple ready queues to represent each level of priority

In many systems, each process is assigned a priority and the scheduler will always choose a process of higher priority over one of lower priority.

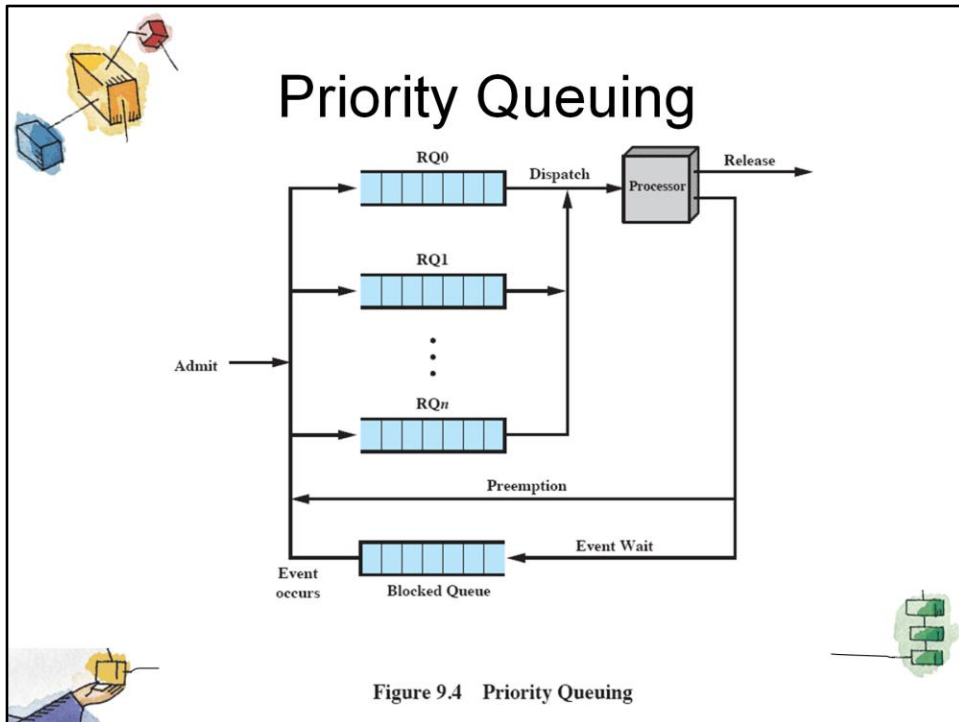


Figure 9.4 Priority Queuing

This figure illustrates the use of priorities. This figure is simplified and ignores the blocked and suspend queues of fig3.8a (chapter 3)

Instead of a single ready queue, we provide a set of queues, in descending order of priority:

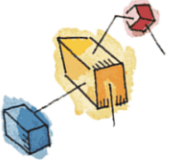
- RQ0, RQ1, . . . RQn, with  $\text{priority}[\text{RQ}_i] > \text{priority}[\text{RQ}_j]$  for  $i < j$ .

**NB:** In UNIX and many other systems, larger priority values represent lower priority processes;

- unless other-wise stated we follow that convention.
- Some systems, such as Windows, use the opposite convention: a higher number means a higher priority.


When a scheduling selection is to be made, the scheduler will start at the highest-priority ready queue (RQ0).

- If there are one or more processes in the queue, a process is selected using some scheduling policy.
- If RQ0 is empty, then RQ1 is examined, and so on.




# Starvation

- Problem:
  - Lower-priority may suffer starvation if there is a steady supply of high priority processes.
- Solution
  - Allow a process to change its priority based on its age or execution history



One problem with a pure priority scheduling scheme is that lower-priority processes may suffer starvation.



- This will happen if there is always a steady supply of higher-priority ready processes.
- If this behaviour is not desirable, the priority of a process can change with its age or execution history.



# Alternative Scheduling Policies

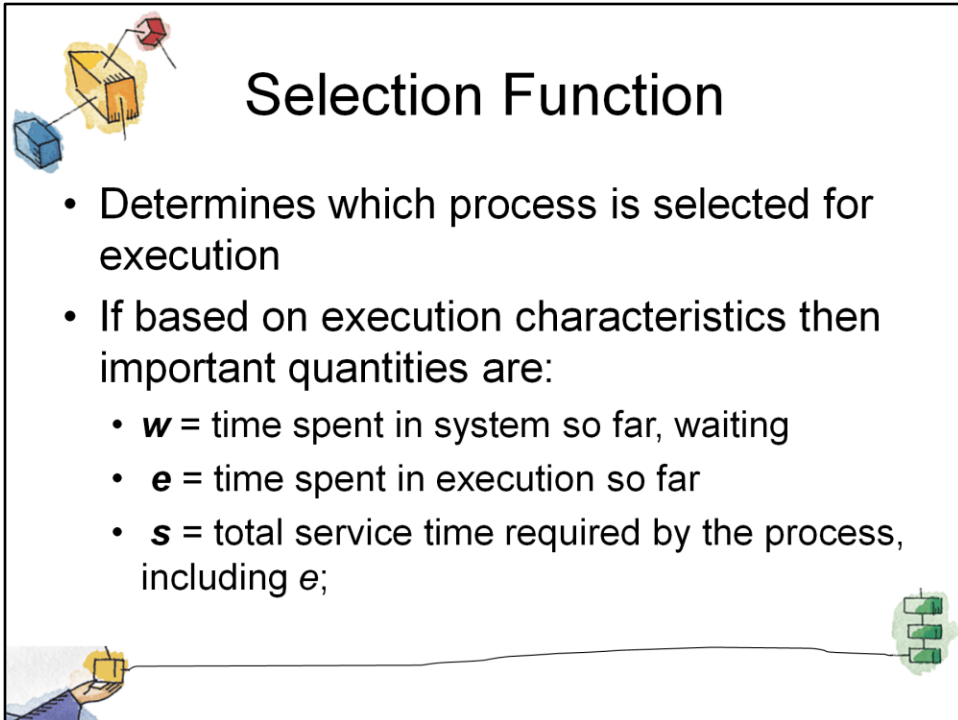
**Table 9.3** Characteristics of Various Scheduling Policies

	FCFS	Round robin	SPN	SRT	HRRN	Feedback
<b>Selection function</b>	max[w]	constant	min[s]	min[s - e]	$\max\left(\frac{w + s}{s}\right)$	(see text)
<b>Decision mode</b>	Non-preemptive	Preemptive (at time quantum)	Non-preemptive	Preemptive (at arrival)	Non-preemptive	Preemptive (at time quantum)
<b>Throughput</b>	Not emphasized	May be low if quantum is too small	High	High	High	Not emphasized
<b>Response time</b>	May be high, especially if there is a large variance in process execution times	Provides good response time for short processes	Provides good response time for short processes	Provides good response time	Provides good response time	Not emphasized
<b>Overhead</b>	Minimum	Minimum	Can be high	Can be high	Can be high	Can be high
<b>Effect on processes</b>	Penalizes short processes; penalizes I/O bound processes	Fair treatment	Penalizes long processes	Penalizes long processes	Good balance	May favor I/O bound processes
<b>Starvation</b>	No	No	Possible	Possible	No	Possible

This table may be difficult to read depending on the display – refer to table 9.3 in textbook

***Animation on-click enlarges the image***



## Selection Function

- Determines which process is selected for execution
- If based on execution characteristics then important quantities are:
  - $w$  = time spent in system so far, waiting
  - $e$  = time spent in execution so far
  - $s$  = total service time required by the process, including  $e$ ;

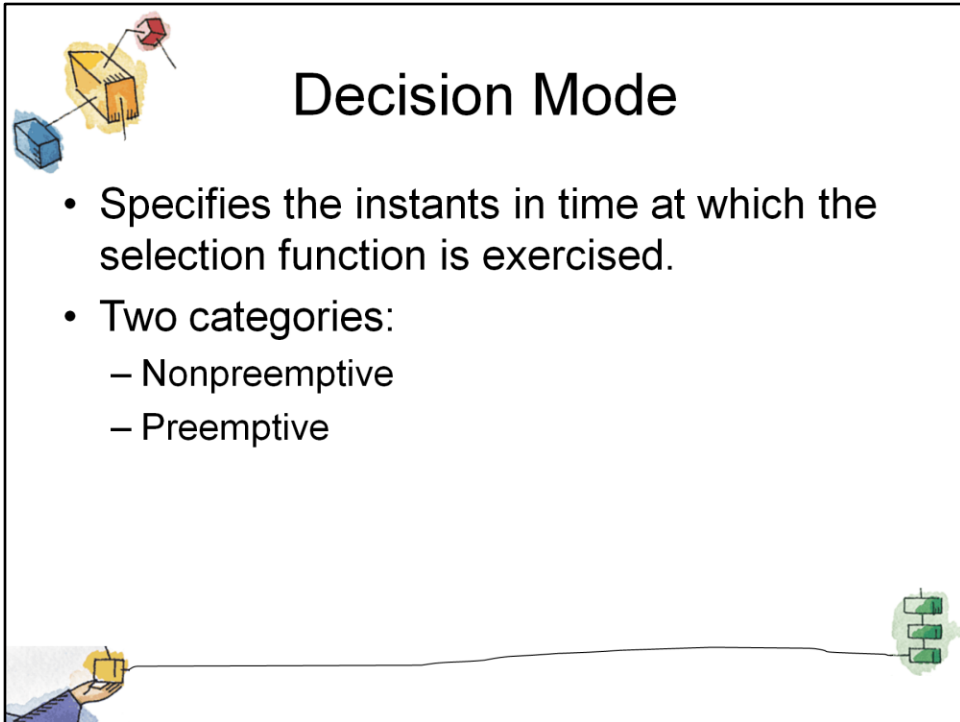
**Selection function** determines which process, among ready processes, is selected next for execution.

- May be based on priority, resource requirements, or the execution characteristics of the process.

In the latter case, three quantities are significant:

- $w$  = time spent in system so far, waiting
- $e$  = time spent in execution so far
- $s$  = total service time required by the process, including  $e$ ;






## Decision Mode

- Specifies the instants in time at which the selection function is exercised.
- Two categories:
  - Nonpreemptive
  - Preemptive

The slide features several decorative illustrations: a blue cube and a yellow cube with a red dot on top in the top-left corner; a hand holding a yellow cube in the bottom-left corner; and a green cube with a white dot on top in the bottom-right corner. A thin black line connects the hand and the green cube.



**Decision mode** specifies the instants in time at which the selection function is exercised.

There are two general categories: non-preemptive and preemptive



# Nonpreemptive vs Preemptive

- Non-preemptive
  - Once a process is in the running state, it will continue until it terminates or blocks itself for I/O
- Preemptive
  - Currently running process may be interrupted and moved to ready state by the OS
  - Preemption may occur when new process arrives, on an interrupt, or periodically.



## Non-preemptive

In this case, once a process is in the Running state, it continues to execute until

- (a) it terminates or
- (b) it blocks itself to wait for I/O or to request some operating system service.

## Preemptive:

The currently running process may be interrupted and moved to the Ready state by the operating system.

The decision to preempt may be performed

- when a new process arrives;
- when an interrupt occurs that places a blocked process in the Ready state;
- or periodically, based on a clock interrupt.



# Process Scheduling Example

- Example set of processes, consider each a batch job

Table 9.4 Process Scheduling Example

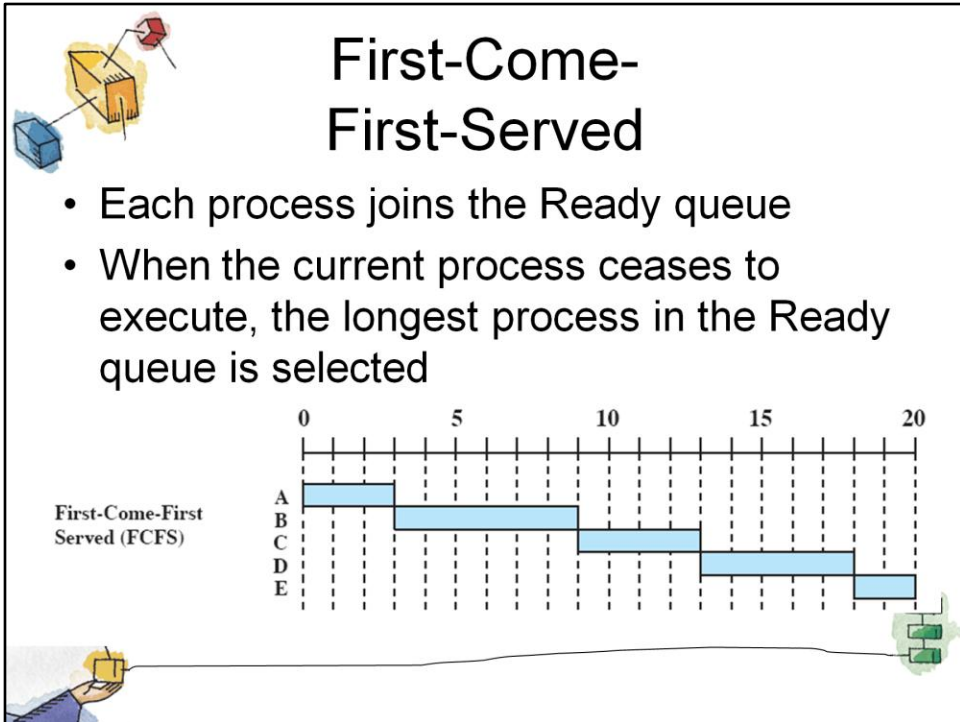
Process	Arrival Time	Service Time
A	0	3
B	2	6
C	4	4
D	6	5
E	8	2

– Service time represents total execution time



We can think of these as batch jobs, with the service time being the total execution time required.

**Movie button** links to animation of *Process Scheduling Algorithms* at <http://gaia.ecs.csus.edu/%7ezhangd/oscal/pscheduling.html>

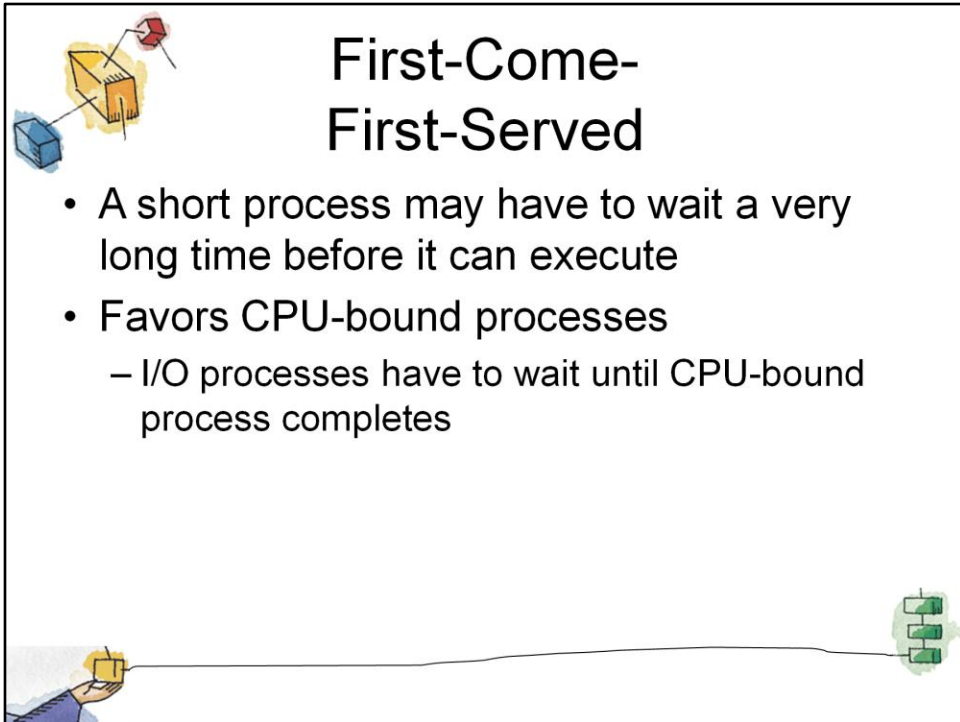


The simplest scheduling policy is first-come-first-served (FCFS),

- aka. first-in-first-out (FIFO) or a strict queuing scheme.

As each process becomes ready, it joins the ready queue.

When the currently running process ceases to execute, the process that has been in the ready queue the longest is selected for running.



## First-Come-First-Served

- A short process may have to wait a very long time before it can execute
- Favors CPU-bound processes
  - I/O processes have to wait until CPU-bound process completes

The slide features several illustrations: a yellow cube with a red dot and a blue cube in the top left; a hand holding a yellow cube in the bottom left; and a green server rack in the bottom right. A thin black line connects the hand to the server rack.


FCFS performs much better for long processes than short ones.

Another difficulty with FCFS is that it tends to favor processor-bound processes over I/O-bound processes.

FCFS is not an attractive alternative on its own for a uniprocessor system.

However, it is often combined with a priority scheme to provide an effective scheduler.



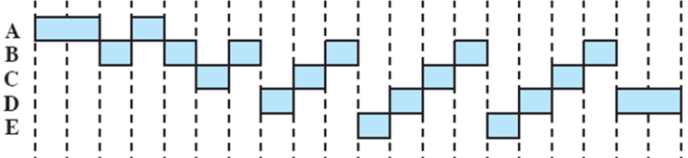
Thus, the scheduler may maintain a number of queues, one for each priority level, and dispatch within each queue on a first-come-first-served basis.



# Round Robin

- Uses preemption based on a clock
  - also known as time slicing, because each process is given a slice of time before being preempted.

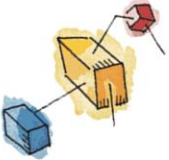
Round-Robin  
(RR),  $q = 1$



A straightforward way to reduce the penalty that short jobs suffer with FCFS is to use preemption based on a clock.



- The simplest such policy is round robin.

Also known as time slicing, because each process is given a slice of time before being preempted.




# Round Robin

- Clock interrupt is generated at periodic intervals
- When an interrupt occurs, the currently running process is placed in the ready queue
  - Next ready job is selected



A clock interrupt is generated at periodic intervals.

- When the interrupt occurs, the currently running process is placed in the ready queue, and the next ready job is selected on a FCFS basis.

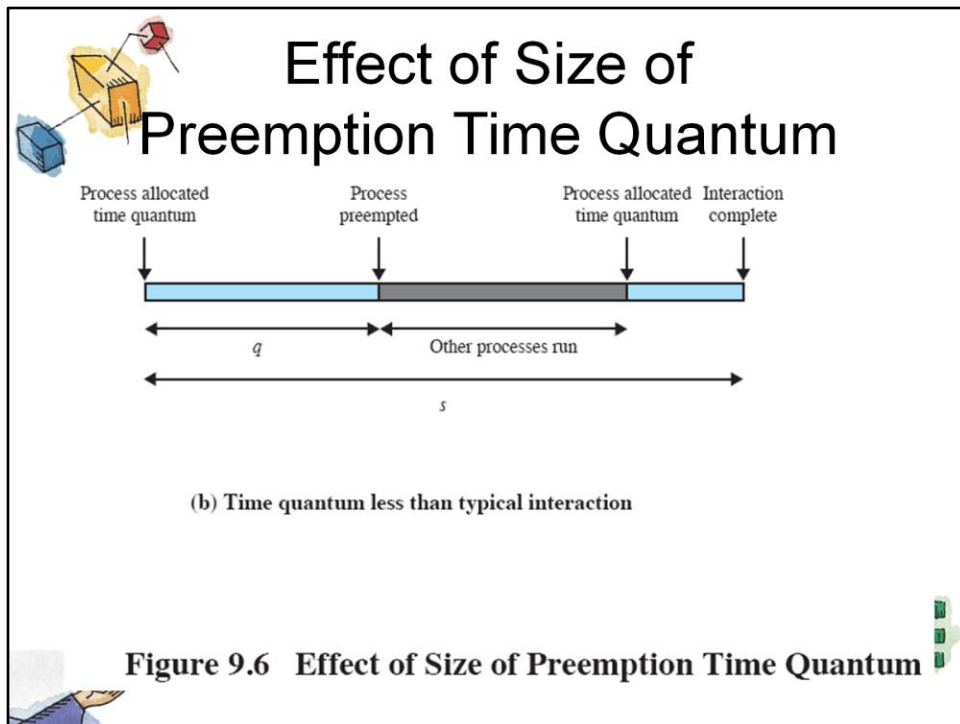


The principal design issue is the length of the time quantum, or slice, to be used.

If the quantum is very short, then short processes will move through the system relatively quickly.

- BUT there is processing over-head involved in handling the clock interrupt and performing the scheduling and dispatching function.
- Thus, very short time quanta should be avoided.



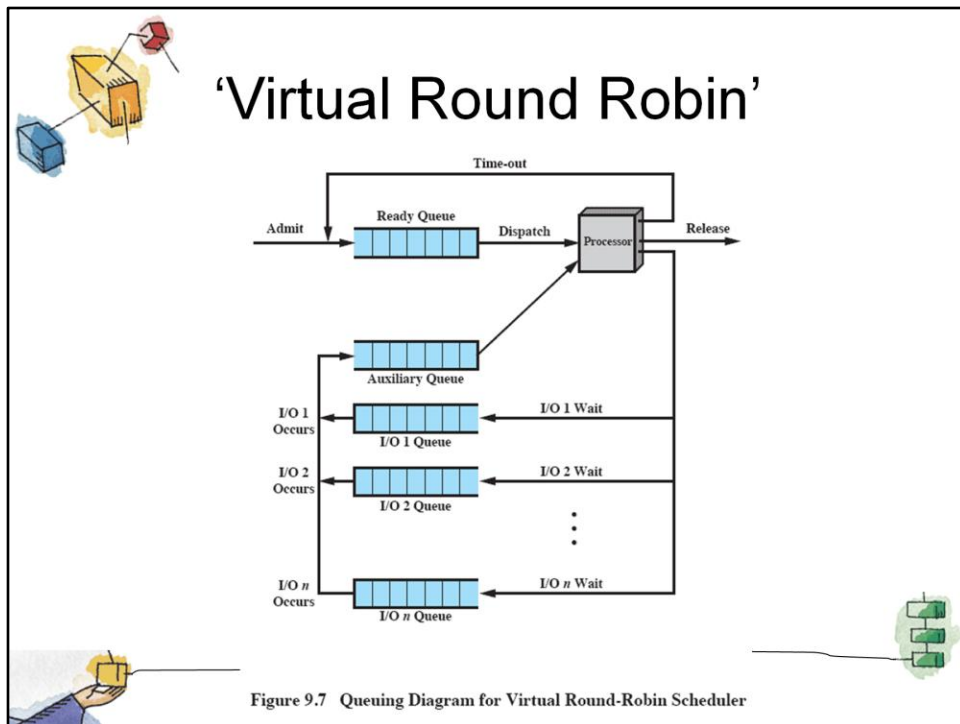


One useful guide is that the time quantum should be slightly greater than the time required for a typical interaction or process function.

- If it is less, then most processes will require at least two time quanta.

Processor-bound processes tend to receive an unfair portion of processor time, which results

- in poor performance for I/O-bound processes,
- inefficient use of I/O devices,
- and an increase in the variance of response time.



As per usual:


- New processes arrive and join the ready queue, which is managed on an FCFS basis.
- When a running process times out, it is returned to the ready queue.
- When a process is blocked for I/O, it joins an I/O queue.

The new feature is an FCFS auxiliary queue to which processes are moved after being released from an I/O block.

When a dispatching decision is to be made, processes in the auxiliary queue get preference over those in the main ready queue.

When a process is dispatched from the auxiliary queue, it runs no longer than a time equal to the basic time quantum minus the total time spent running since it was last selected from the main ready queue.

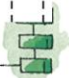

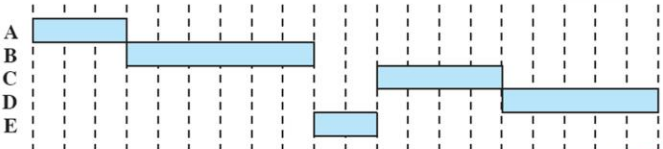
Performance studies indicate that this approach is indeed superior to round robin in terms of fairness.



# Shortest Process Next

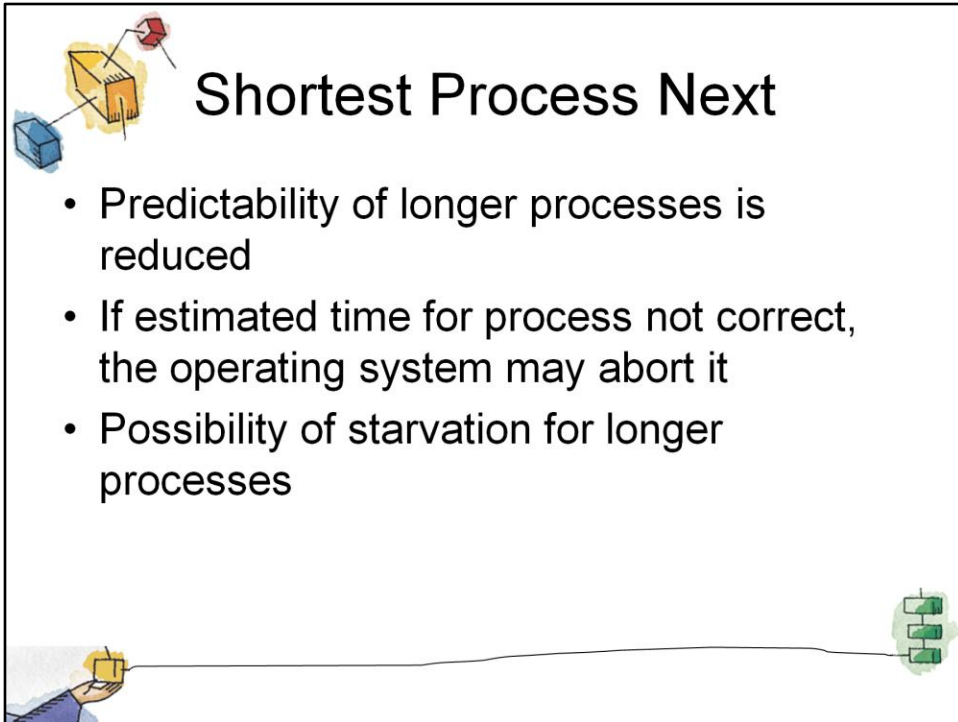
- Nonpreemptive policy
- Process with shortest expected processing time is selected next
- Short process jumps ahead of longer processes

Shortest Process Next (SPN)



A non-preemptive policy in which the process with the shortest expected processing time is selected next.

- Thus a short process will jump to the head of the queue past longer jobs.



## Shortest Process Next

- Predictability of longer processes is reduced
- If estimated time for process not correct, the operating system may abort it
- Possibility of starvation for longer processes

Overall performance is significantly improved in terms of response time.

- However, the variability of response times is increased, especially for longer processes, and thus predictability is reduced.

One difficulty with the SPN policy is the need to know or at least estimate the required processing time of each process.

- For batch jobs, the system may require the programmer to estimate the value and supply it to the operating system.
- If the programmer's estimate is substantially under the actual running time, the system may abort the job.




## Calculating Program 'Burst'

- Where:
  - $T_i$  = processor execution time for the  $i$ th instance of this process
  - $S_i$  = predicted value for the  $i$ th instance
  - $S_1$  = predicted value for first instance; not calculated

$$S_{n+1} = \frac{1}{n} \sum_{i=1}^n T_i$$






# Exponential Averaging

- A common technique for predicting a future value on the basis of a time series of past values is exponential averaging

$$S_{n+1} = \alpha T_n + (1 - \alpha)S_n$$





The size of the coefficient as a function of its position in the expansion is shown here.

The larger the value of  $\alpha$ , the greater the weight given to the more recent observations.

For  $\alpha = 0.8$ , virtually all of the weight is given to the four most recent observations, whereas for  $\alpha = 0.2$ , the averaging is effectively spread out over the eight or so most recent observations.

The advantage of using a value of  $\alpha$  close to 1 is that the average will quickly reflect a rapid change in the observed quantity.

The disadvantage is that if there is a brief surge in the value of the observed quantity and it then settles back to some average value, the use of a large value of  $\alpha$  will result in jerky changes in the average.


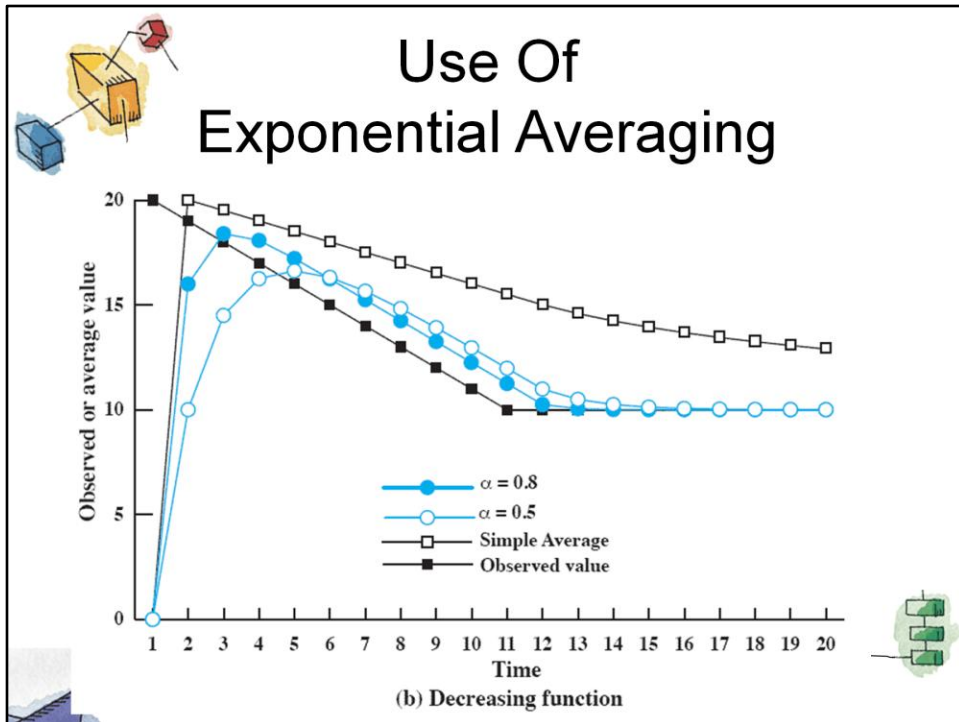


Figure 9.9 compares simple averaging with exponential averaging (for two different values of  $\alpha$ ).

Here (Figure 9.9a), the observed value begins at 1, grows gradually to a value of 10, and then stays there.





While here (Figure 9.9b), the observed value begins at 20, declines gradually to 10, and then stays there.

In both cases, we start out with an estimate of  $S_1 = 0$ .

- This gives greater priority to new processes.

Note that exponential averaging tracks changes in process behavior faster than does simple averaging and that the larger value of  $\alpha$  results in a more rapid reaction to the change in the observed value.

# Shortest Remaining Time

- Preemptive version of shortest process next policy
- Must estimate processing time and choose the shortest

Shortest Remaining Time (SRT)

A B C D E

A pre-emptive version of SPN.

In this case, the scheduler always chooses the process that has the shortest expected remaining processing time.

SRT does not have the bias in favor of long processes found in FCFS.

- Unlike round robin, no additional interrupts are generated, reducing overhead.
- On the other hand, elapsed service times must be recorded, contributing to overhead.

SRT should also give superior turnaround time performance to SPN, because a short job is given immediate preference to a running longer job.

## Highest Response Ratio Next

- Choose next process with the greatest ratio

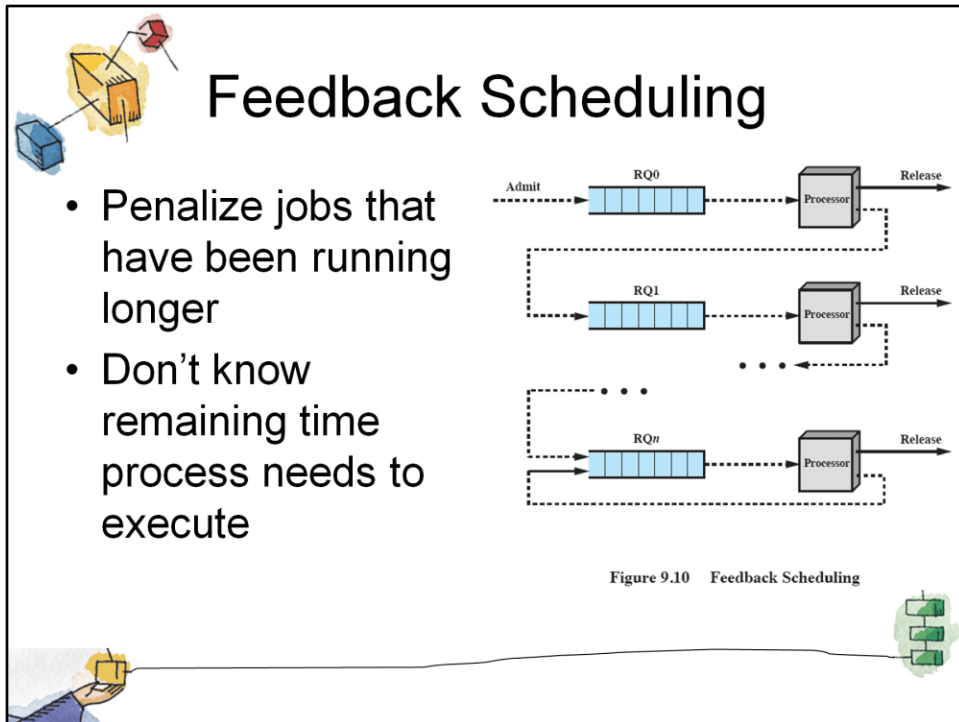
$$\text{Ratio} = \frac{\text{time spent waiting} + \text{expected service time}}{\text{expected service time}}$$

Highest Response Ratio Next (HRRN)

The Gantt chart illustrates the execution order of processes A, B, C, D, and E. The processes are represented by horizontal bars of different lengths. The execution order is A, C, B, E, and D. This order is determined by the Highest Response Ratio Next (HRRN) scheduling algorithm, which selects the process with the highest ratio at each step. The ratio is calculated as the sum of the time spent waiting and the expected service time, divided by the expected service time. As processes wait, their ratio increases, eventually allowing longer processes like D to be scheduled before shorter ones like B.

A smaller denominator yields a larger ratio so that shorter jobs are favored, but aging without service increases the ratio so that a longer process will eventually get past competing shorter jobs.

As with SRT and SPN, the expected service time must be estimated to use highest response ratio next (HRRN).



If we have no indication of the relative length of various processes, then none of SPN, SRT, and HRRN can be used.

Another way of establishing a preference for shorter jobs is to penalize jobs that have been running longer.

- IE if we cannot focus on the time remaining to execute, let's focus on the time spent in execution so far.

Scheduling is done on a preemptive (at time quantum) basis, and a dynamic priority mechanism is used.

When a process first enters the system, it is placed in RQ0 (refer to Figure 9.4).

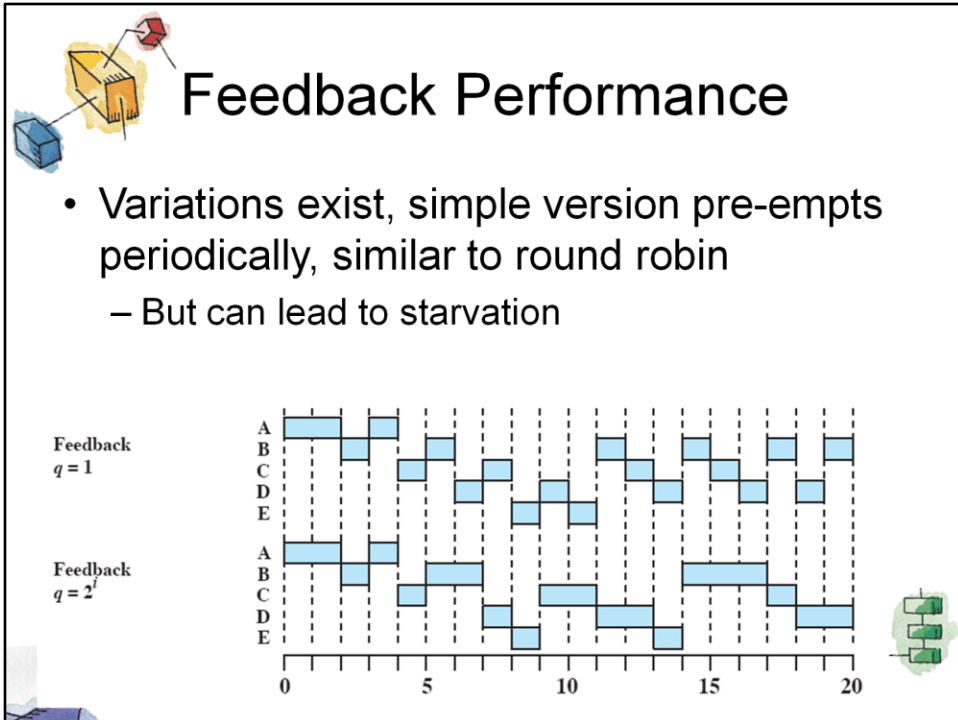
- After its first preemption, when it returns to the Ready state, it is placed in RQ1.
- Each subsequent time that it is preempted, it is demoted to the next lower-priority queue.

A short process will complete quickly, without migrating very far down the hierarchy of ready queues.

- A longer process will gradually drift downward.
- Thus, newer, shorter processes are favored over older, longer processes.

Within each queue, except the lowest-priority queue, a simple FCFS mechanism is used.

- Once in the lowest-priority queue, a process cannot go lower, but is returned to this queue repeatedly until it completes execution.
- Thus, this queue is treated in round-robin fashion.



A number of variations of this scheme exist.

A simple version is to perform preemption in the same fashion as for round robin: at periodic intervals.

Our example shows this (Figure 9.5 and Table 9.5) for a quantum of one time unit.

Note that in this case, the behavior is similar to round robin with a time quantum of 1.

But the turnaround time of longer processes can stretch out alarmingly.

- it is possible for starvation to occur if new jobs are entering the system frequently.


To compensate for this, we can vary the preemption times according to the queue:

- A process scheduled from RQ0 is allowed to execute for one time unit and then is preempted;
- a process scheduled from RQ1 is allowed to execute two time units, and so on.

In general, a process scheduled from RQ $i$  is allowed to execute  $2^i$  time units before preemption.

- Even with the allowance for greater time allocation at lower priority, a longer process may still suffer starvation.

A possible remedy is to promote a process to a higher-priority queue after it spends a certain amount of time waiting for service in its current queue.




# Performance Comparison

- Any scheduling discipline that chooses the next item to be served independent of service time obeys the relationship:

$$\frac{T_r}{T_s} = \frac{1}{1 - \rho}$$

where

- $T_r$  = turnaround time or residence time; total time in system, waiting plus execution
- $T_s$  = average service time; average time spent in Running state
- $\rho$  = processor utilization

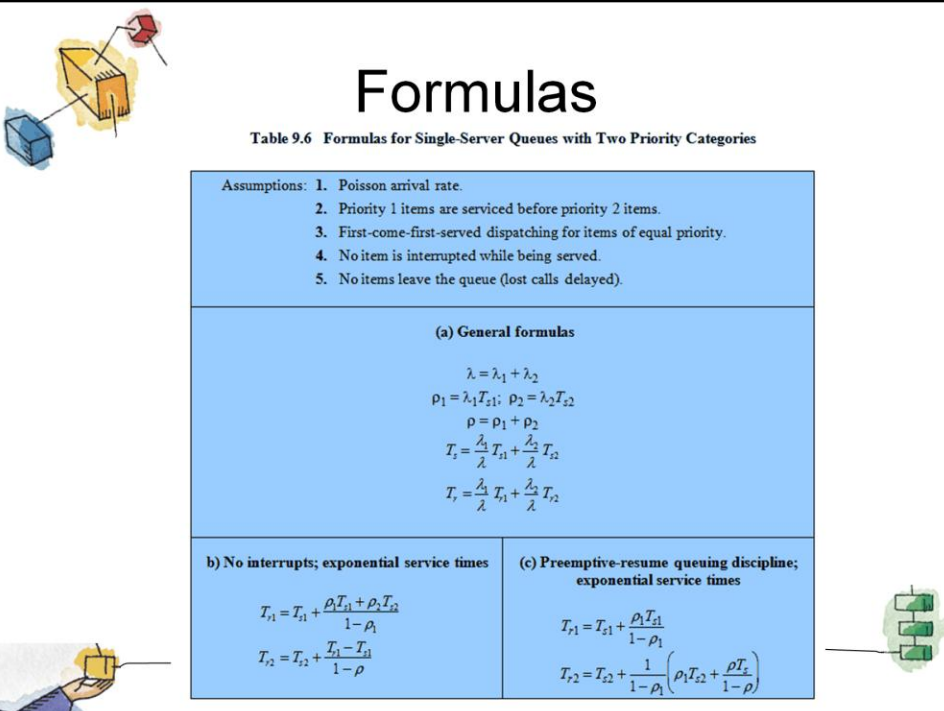


The performance of various scheduling policies is a critical factor in the choice of a scheduling policy.

But, it is impossible to make definitive comparisons because relative performance will depend on a variety of factors,

- But it is possible to make some general conclusions

For any scheduling discipline that chooses the next item to be served independent of service time obeys the above relationship



# Formulas

**Table 9.6 Formulas for Single-Server Queues with Two Priority Categories**

Assumptions: 1. Poisson arrival rate. 2. Priority 1 items are serviced before priority 2 items. 3. First-come-first-served dispatching for items of equal priority. 4. No item is interrupted while being served. 5. No items leave the queue (lost calls delayed).	
<b>(a) General formulas</b>  $\lambda = \lambda_1 + \lambda_2$ $\rho_1 = \lambda_1 T_{s1}; \quad \rho_2 = \lambda_2 T_{s2}$ $\rho = \rho_1 + \rho_2$ $T_s = \frac{\lambda_1}{\lambda} T_{s1} + \frac{\lambda_2}{\lambda} T_{s2}$ $T_r = \frac{\lambda_1}{\lambda} T_{r1} + \frac{\lambda_2}{\lambda} T_{r2}$	
<b>b) No interrupts; exponential service times</b>  $T_{r1} = T_{s1} + \frac{\rho_1 T_{s1} + \rho_2 T_{s2}}{1 - \rho_1}$ $T_{r2} = T_{s2} + \frac{T_{s1} - T_{s2}}{1 - \rho}$	<b>(c) Preemptive-resume queuing discipline; exponential service times</b>  $T_{r1} = T_{s1} + \frac{\rho_1 T_{s1}}{1 - \rho_1}$ $T_{r2} = T_{s2} + \frac{1}{1 - \rho_1} \left( \rho_1 T_{s2} + \frac{\rho T_{s2}}{1 - \rho} \right)$

If scheduling is done on the basis of priority and if processes are assigned to a priority class on the basis of service time, then differences do emerge.


This slide shows the formulas that result when we assume two priority classes, with different service times for each class.

$\lambda$  refers to the arrival rate.

These results can be generalized to any number of priority classes.

**Note:** the formulas differ for nonpreemptive versus preemptive scheduling.

In the preemptive, it is assumed that a lower-priority process is immediately interrupted when a higher-priority process becomes ready.




This figure shows the overall result.

By giving preference to shorter jobs, the average normalized turnaround time is improved at higher levels of utilization.

The improvement is greatest with the use of preemption.

Notice, however, that overall performance is not much affected.





Significant differences emerge when we consider the two priority classes separately.


This figure shows the results for the higher-priority, shorter processes.

For comparison, the upper line on the graph assumes that priorities are not used but that we are simply looking at the relative performance of half of all processes that have the shorter processing time.

- The other two lines assume that these processes are assigned a higher priority.

When the system is run using priority scheduling without preemption, the improvements are significant.

- They are even more significant when preemption is used.



This shows the same analysis for the lower-priority, longer processes.

As expected, such processes suffer a performance degradation under priority scheduling.




Figure 9.14 Simulation Results for Normalized Turnaround Time

Figure 9.15 Simulation Results for Waiting Time

Some of the difficulties of analytic modelling are over-come by using discrete-event simulation, which allows a wide range of policies to be modelled.

The disadvantage of simulation is that the results for a given “run” only apply to that particular collection of processes under that particular set of assumptions. Nevertheless, useful insights can be gained.

This figure shows the result of one study. Figure 9.14 shows the normalized turnaround time, and Figure 9.15 shows the average waiting time.

Looking at the turnaround time, we can see that the performance of FCFS is very unfavorable,

- with one-third of the processes having a normalized turnaround time greater than 10 times the service time (these were the shortest)
- On the other hand, the absolute waiting time is uniform, as is to be expected because scheduling is independent of service time.

The figures show round robin using a quantum of one time unit.

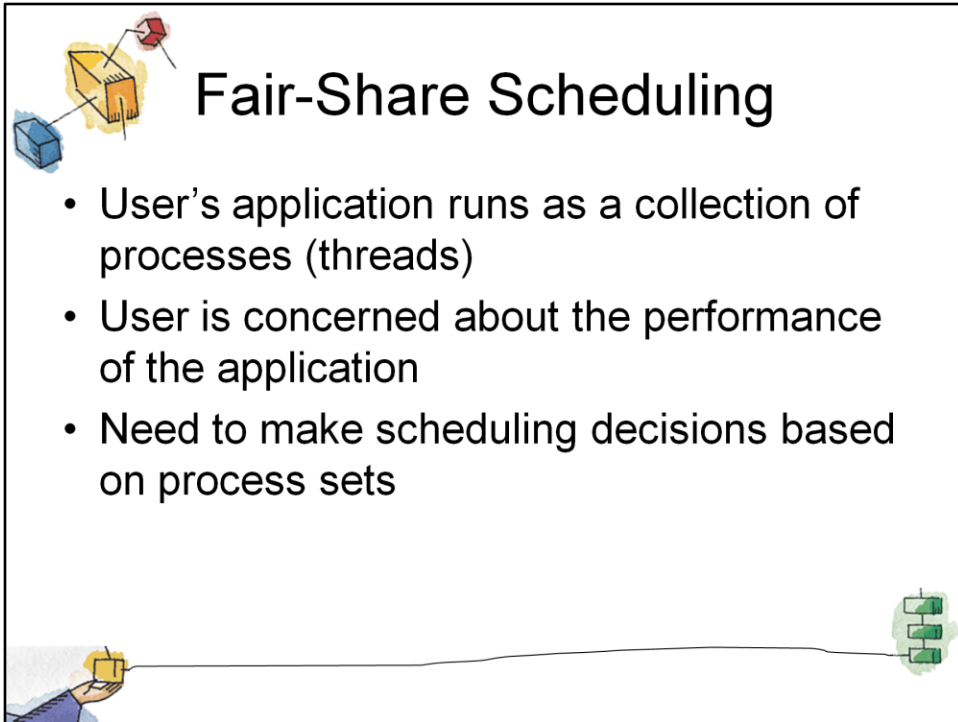
- Round robin yields a normalized turnaround time of about 5 for all processes, treating all fairly, except for the shortest processes, which execute in less than one quantum,
- Shortest process next performs better than round robin, except for the shortest processes.
- Shortest remaining time, the preemptive version of SPN, performs better than SPN except for the longest 7% of all processes.

We have seen that, among nonpreemptive policies, FCFS favors long processes and SPN favors short ones.

Highest response ratio next is intended to be a compromise between these two effects, and this is indeed confirmed in the figures.

Finally, the figure shows feedback scheduling with fixed, uniform quanta in each priority queue.

- As expected, FB performs quite well for short processes.

The diagram is enclosed in a black rectangular border. In the top-left corner, there is a cluster of icons: a blue cube, a yellow cube with a red dot on top, and a red cube. In the top-right corner, the title "Fair-Share Scheduling" is written in a large, black, sans-serif font. Below the title, there is a bulleted list of three items. In the bottom-left corner, there is an illustration of a hand holding a yellow cube. In the bottom-right corner, there is a green icon consisting of three stacked rectangular blocks. A thin black line runs horizontally across the bottom of the diagram, connecting the hand icon to the green icon.

## Fair-Share Scheduling

- User's application runs as a collection of processes (threads)
- User is concerned about the performance of the application
- Need to make scheduling decisions based on process sets

From the user's point of view, the concern is not how a particular process performs but rather how his or her set of processes, which constitute a single application, performs.

It would be attractive to make scheduling decisions on the basis of these process sets.

- This approach is generally known as fair-share scheduling.
- The concept can be extended to groups of users, even if each user is represented by a single process.

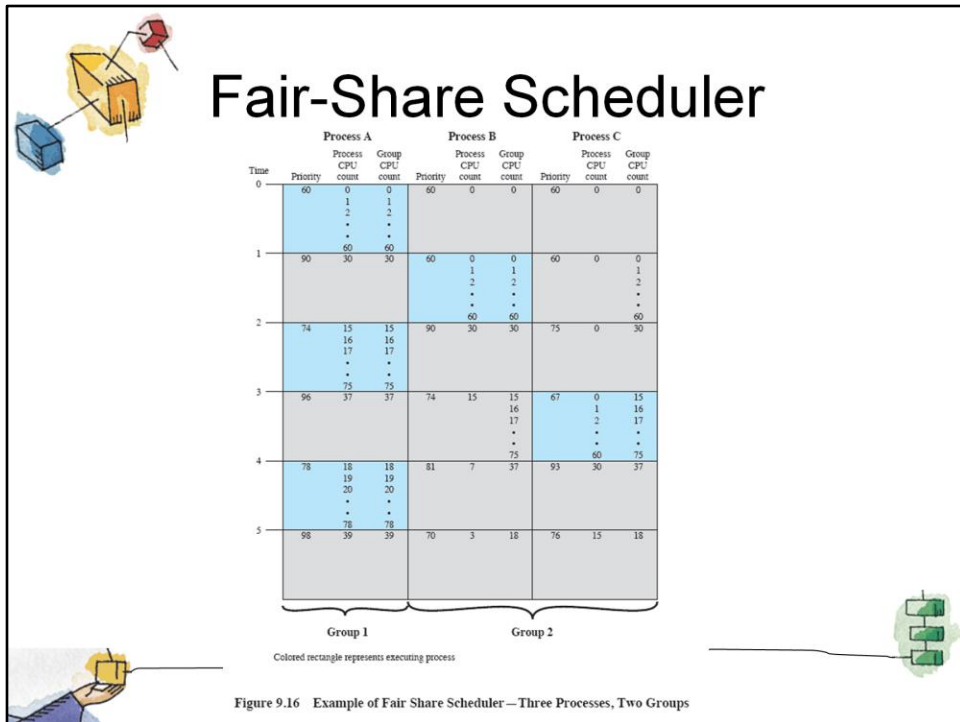


Figure 9.16 Example of Fair Share Scheduler—Three Processes, Two Groups

This figure is an example in which process A is in one group and process B and process C are in a second group,

- with each group having a weighting of 0.5.

Assume that all processes are processor bound and are usually ready to run.

All processes have a base priority of 60.

Processor utilization is measured as follows:

- The processor is interrupted 60 times per second;
- during each interrupt, the processor usage field of the currently running process is incremented, as is the corresponding group processor field.
- Once per second, priorities are recalculated.

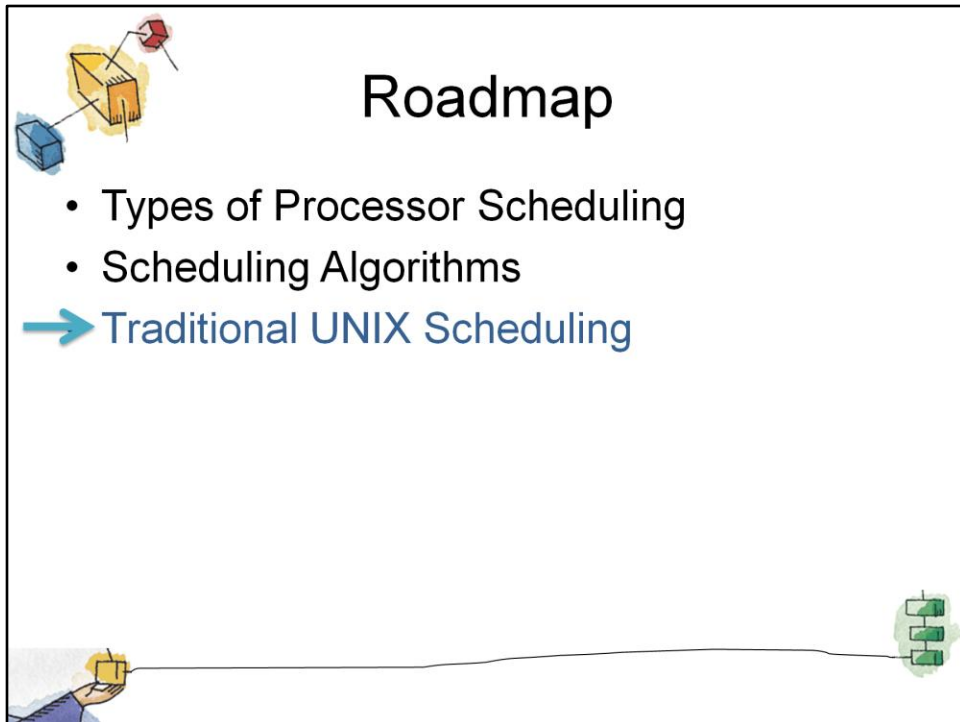
In the figure, process A is scheduled first.

- At the end of one second, it is pre-empted.
- Processes B and C now have the higher priority, and process B is scheduled.

At the end of the second time unit, process A has the highest priority.

Note that the pattern repeats:

- the kernel schedules the processes in order: A, B, A, C, A, B, and so on.
- Thus, 50% of the processor is allocated to process A, which constitutes one group, and 50% to processes B and C, which constitute another group.

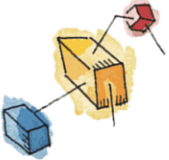


In this section we examine traditional UNIX scheduling, which is used in both SVR3 and 4.3 BSD UNIX.

- Primarily targeted at the time-sharing interactive environment.



The scheduling algorithm is designed to provide good response time for interactive users while ensuring that low-priority background jobs do not starve.

- Although this algorithm has been replaced in modern UNIX systems, the approach because is representative of practical time-sharing scheduling algorithms.



## Traditional UNIX Scheduling

- Multilevel feedback using round robin within each of the priority queues
- If a running process does not block or complete within 1 second, it is preempted
- Priority is based on process type and execution history.



The traditional UNIX scheduler employs multilevel feedback using round robin within each of the priority queues.

The system makes use of 1-second preemption.

- ie if a running process does not block or complete within 1 second, it is preempted.

Priority is based on process type and execution history.



# Scheduling Formula

$$CPU_j(i) = \frac{CPU_j(i-1)}{2}$$

$$P_j(i) = Base_j + \frac{CPU_j(i)}{2} + nice_j$$

where

$CPU_j(i)$  = measure of processor utilization by process  $j$  through interval  $i$

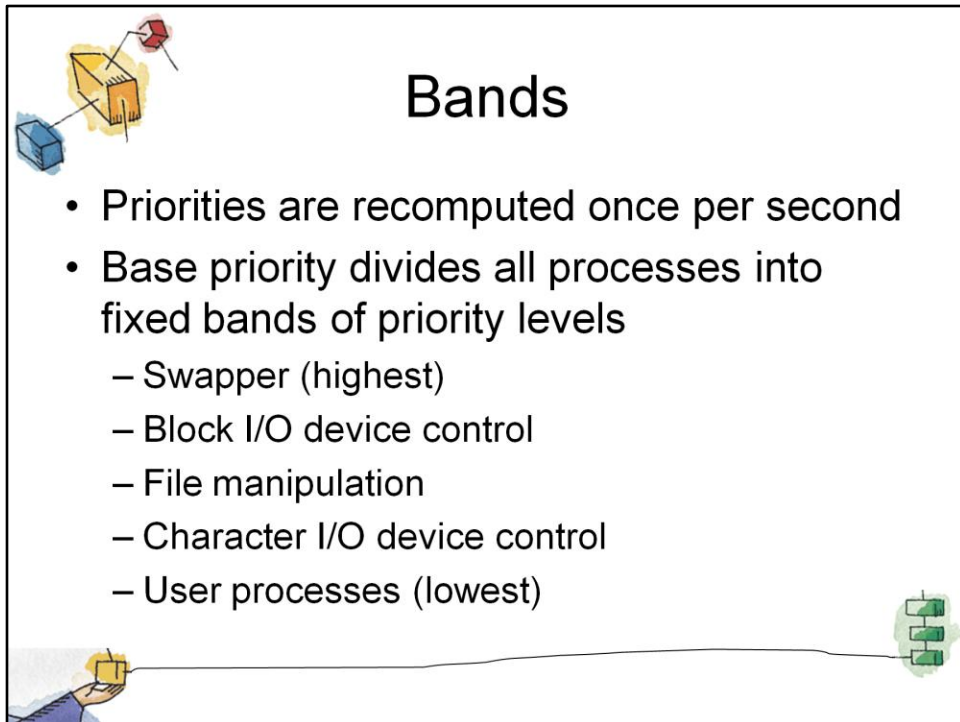
$P_j(i)$  = priority of process  $j$  at beginning of interval  $i$ ; lower values equal higher priorities

$Base_j$  = base priority of process  $j$

$nice_j$  = user-controllable adjustment factor







# Bands

- Priorities are recomputed once per second
- Base priority divides all processes into fixed bands of priority levels
  - Swapper (highest)
  - Block I/O device control
  - File manipulation
  - Character I/O device control
  - User processes (lowest)

The slide features several decorative illustrations: a blue cube and a yellow cube with a red dot on top in the upper left; a hand holding a yellow cube in the lower left; and a green circuit-like diagram in the lower right.

The priority of each process is recomputed once per second, at which time a new scheduling decision is made.

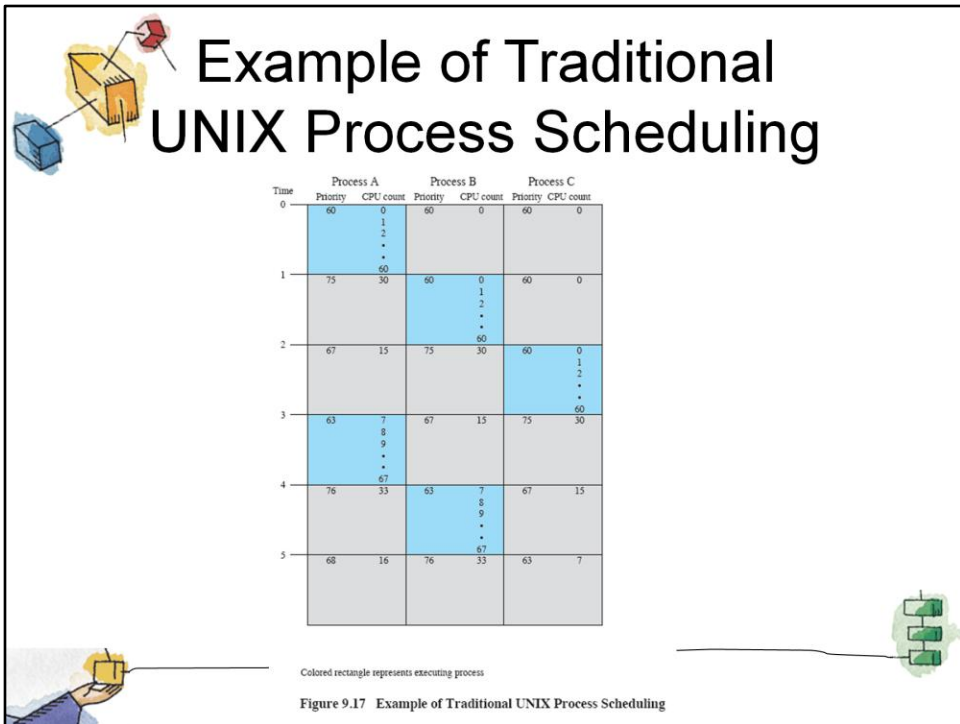
- The base priority is to divide all processes into fixed bands of priority levels.

The CPU and nice components are restricted to prevent a process from migrating out of its assigned band (assigned by the base priority level).

These bands are used to optimize access to block devices (e.g., disk) and to allow the operating system to respond quickly to system calls.

In decreasing order of priority, the bands are

- Swapper
- Block I/O device control
- File manipulation
- Character I/O device control
- User processes



This figure shows an example of process scheduling

Processes A, B, and C are created at the same time with base priorities of 60 (ignoring the nice value).

The clock interrupts the system 60 times per second and increments a counter for the running process.

The example assumes that none of the processes block themselves and that no other processes are ready to run.