

COP 4600 – Summer 2012

Introduction To Operating Systems

Uni-processor Scheduling

Instructor : Dr. Mark Llewellyn
markl@cs.ucf.edu
HEC 236, 407-823-2790
<http://www.cs.ucf.edu/courses/cop4600/sum2012>

Department of Electrical Engineering and Computer Science
Computer Science Division
University of Central Florida



Uni-processor Scheduling

- In a multiprogramming system, multiple processes exist concurrently in main memory. Each process alternates between using the processor and waiting for some event to occur, such as the completion of I/O.
- The key to multiprogramming is scheduling.
- The goals of scheduling are:
 1. Assign processes to be executed by the processor(s)
 2. Improve response time
 3. Improve throughput
 4. Increase processor efficiency



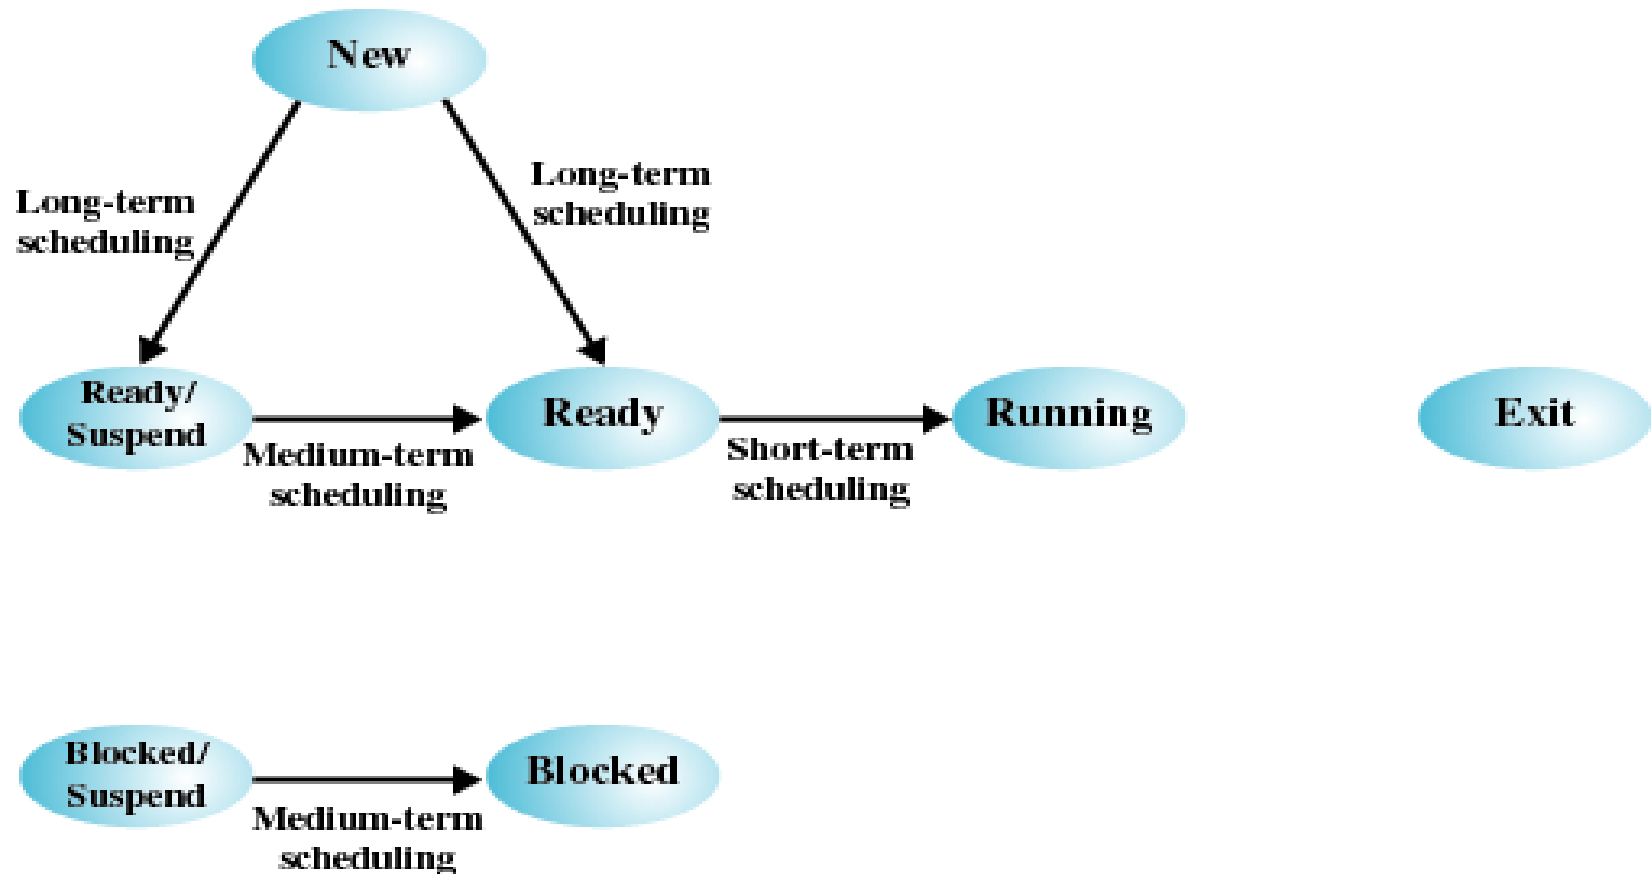
Types of Scheduling

There are typically four different types of scheduling involved.

- **Long-term scheduling:** The decision to add to the pool of processes to be executed.
- **Medium-term scheduling:** The decision to add to the number of processes that are partially or fully in main memory.
- **Short-term scheduling (dispatcher):** The decision as to which available process will be executed by the processor.
- **I/O scheduling:** The decision as to which process's pending I/O request will be handled by an available I/O device. (We'll defer this type of scheduling until we discuss I/O management later in the course.)



Scheduling and Process State Transitions



The diagram illustrates a process state transition model. It is organized into three main horizontal sections: Short Term, Medium Term, and Long Term.

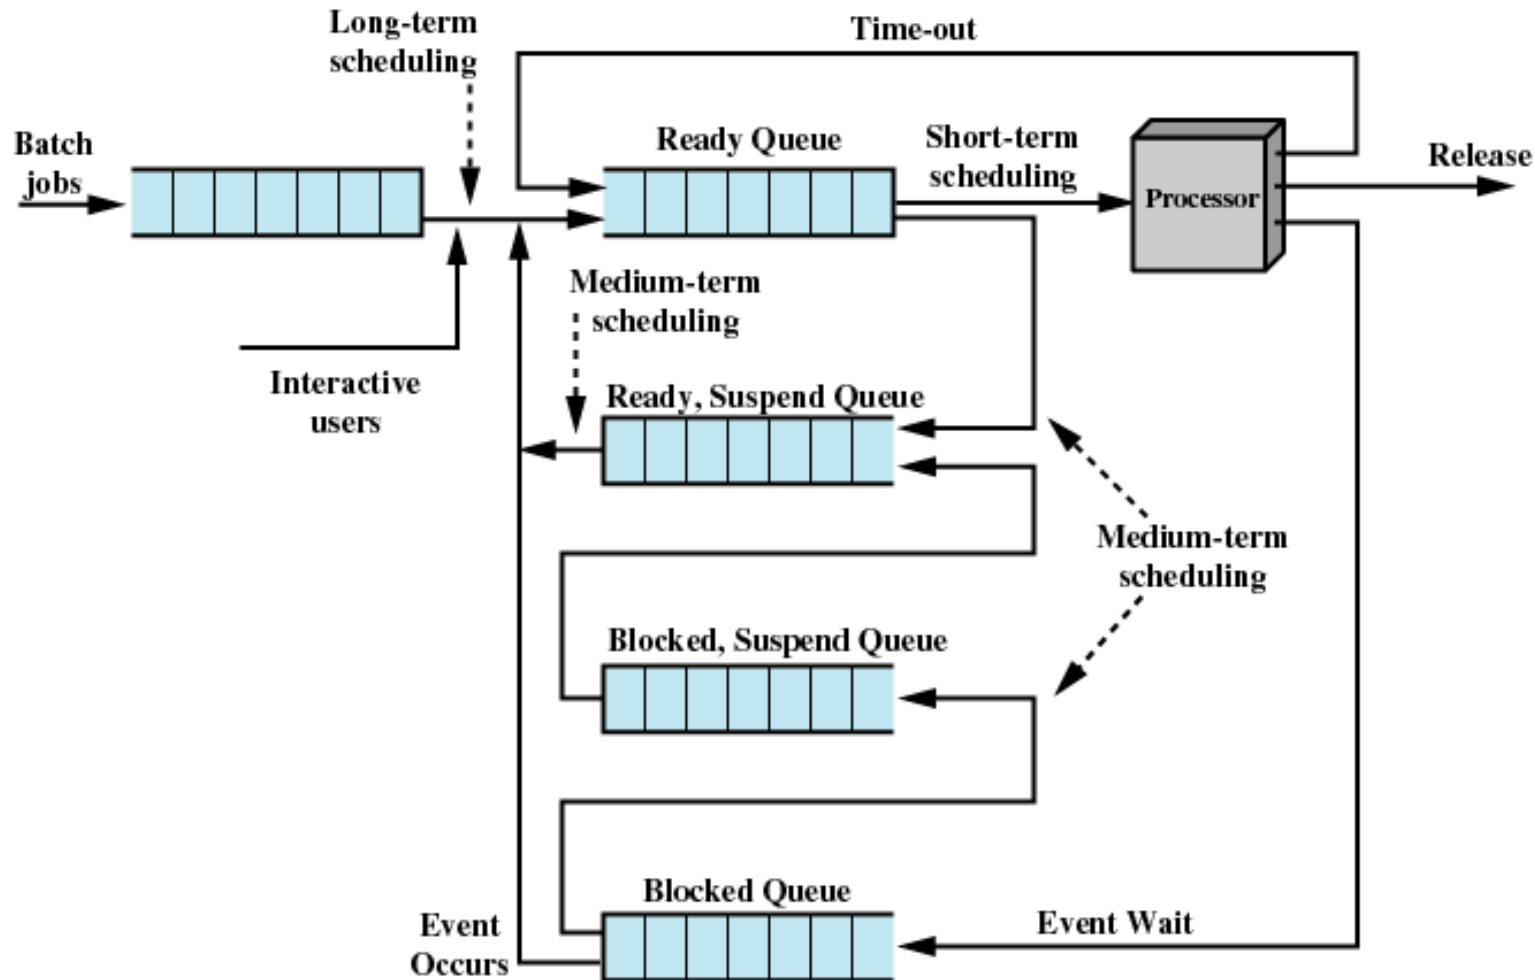
- Short Term:** Contains three states: **Running**, **Ready**, and **Blocked**. Transitions include:
 - Running to Ready (downward arrow)
 - Ready to Running (upward arrow)
 - Ready to Blocked (downward arrow)
 - Blocked to Ready (upward arrow)
 - A curved arrow from Running to Blocked.
- Medium Term:** Contains two states: **Blocked, Suspend** and **Ready, Suspend**. Transitions include:
 - Blocked, Suspend to Ready, Suspend (downward arrow)
- Long Term:** Contains two states: **New** and **Exit**.

Transitions between sections:

- New** (Long Term) transitions to **Ready** (Short Term).
- Running** (Short Term) transitions to **Ready, Suspend** (Medium Term).
- Ready** (Short Term) transitions to **Ready, Suspend** (Medium Term).
- Blocked** (Short Term) transitions to **Blocked, Suspend** (Medium Term).
- Blocked, Suspend** (Medium Term) transitions back to **Ready** (Short Term).
- Ready, Suspend** (Medium Term) transitions back to **Ready** (Short Term).
- Ready, Suspend** (Medium Term) transitions to **Exit** (Long Term).
- Blocked, Suspend** (Medium Term) transitions to **Exit** (Long Term).
- Exit** (Long Term) is the final state.

Scheduling affects the performance of the system because it determines which processes will wait and which will progress. This is illustrated by the next diagram.

Queuing Diagram For Scheduling



Long-Term Scheduling

- Determines which programs are admitted to the system for processing
- Long term scheduling controls the degree of multiprogramming
- The decision as to when to create a new process is general driven by the desired degree of multiprogramming. The more processes that are created, the smaller is the percentage of time each process can be executed (i.e., more processes are competing for the same amount of processor time).
- Thus, the long term scheduler may limit the degree of multiprogramming to provide satisfactory service to the current set of processes.



Long-Term Scheduling (cont.)

- The decision as to which job to admit next can be based on a simple first-come-first-served basis, or it can be based on a much more elaborate protocol to assist in the management of system performance.
- Many different criteria can be used including:
 - Priority
 - Expected execution time
 - I/O requirements
 - Overall system balance (CPU bound versus I/O bound processes)
- Note: for time sharing systems, process creation will occur when a user attempts to connect to the system. Time sharing users are not queued up and kept waiting, rather all comers are accepted until the system reaches some saturation point.



Medium-Term Scheduling

- 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 also becomes an issue that must be addressed by the medium-term scheduler. This means that the swapping-in decision must consider the memory requirements of the swapped-out process.



Short-Term Scheduling

- 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.
- The medium-term scheduler is executed somewhat more frequently to make a swapping decision.
- The short-term scheduler is also known as the **dispatcher**, executes the most frequently and makes the fine-grained decision of which process to execute next.



Short-Term Scheduling (cont.)

- The short-term scheduler is invoked when 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.
- Example of such events include:
 - Clock interrupts
 - I/O interrupts
 - Operating system calls
 - Signals (semaphores)



Short-Term Scheduling (cont.)

- The main objective of short-term scheduling is to allocate processor time in such a way as to optimize one or more aspects of the systems behavior.
- The commonly used criteria can be categorized into two broad dimensions.
 1. We can make the distinction between user-oriented and system-oriented criteria.
 2. We can also make the distinction between criteria which are performance related and those that are not directly performance related.



Short-Term Scheduling Criteria

- User-oriented (perceived by the user or process)
 - Response Time in an interactive system
 - Elapsed time between the submission of a request until there is output.
 - For example, a threshold of 2 seconds may be defined such that the goal of the scheduling is to maximize the number of users who experience an average response time of 2 seconds or less.
- System-oriented
 - Effective and efficient utilization of the processor
 - An example is throughput, which is the rate at which processes are completed. Focus is clearly on system performance rather than service provided to the user, although the users may also benefit from increased throughput.



Short-Term Scheduling Criteria

- Performance-related
 - Quantitative
 - Readily measurable and analyzable.
 - Examples: response time and throughput.
- Non-performance related
 - Qualitative
 - Not readily measurable.
 - Example is predictability. Service provided to users exhibits the same characteristics over time independent of other work being performed by the system.



Summary of Scheduling Criteria

User Oriented, Performance Related

Turnaround time This is the interval of time between the submission of a process and its completion. Includes actual execution time plus time spent waiting for resources, including the processor. This is an appropriate measure for a batch job.

Response time For an interactive process, this is the time from the submission of a request until the response begins to be received. Often a process can begin producing some output to the user while continuing to process the request. Thus, this is a better measure than turnaround time from the user's point of view. The scheduling discipline should attempt to achieve low response time and to maximize the number of interactive users receiving acceptable response time.

Deadlines When process completion deadlines can be specified, the scheduling discipline should subordinate other goals to that of maximizing the percentage of deadlines met.

User Oriented, Other

Predictability A given job should run in about the same amount of time and at about the same cost regardless of the load on the system. A wide variation in response time or turnaround time is distracting to users. It may signal a wide swing in system workloads or the need for system tuning to cure instabilities.



Summary of 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.

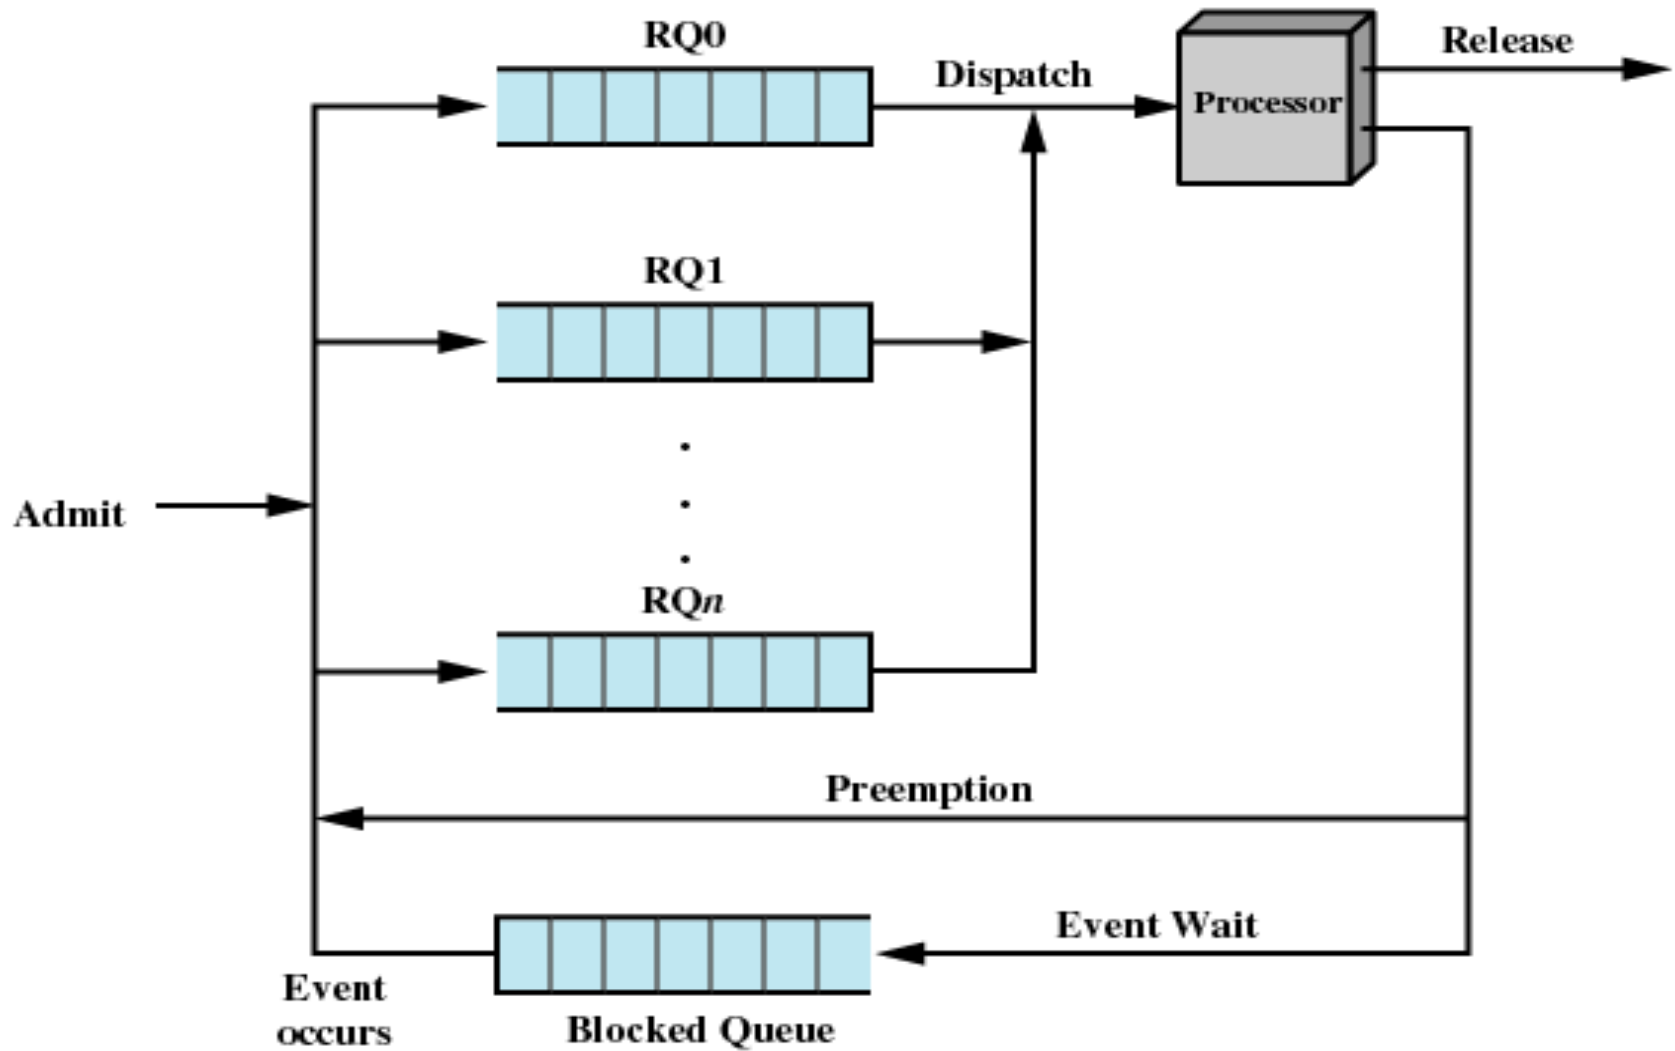


The Use Of Priorities

- 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
- Have multiple ready queues (RQ #) to represent each level of priority
- One problem with a pure priority scheduling scheme is that lower-priority processes may suffer starvation. This happens when there is always a steady supply of higher-priority processes.
 - To prevent this it is possible to allow a process to change its priority based on its age or execution history.



Priority Queuing



Alternative Scheduling Protocols

- The table on the following page illustrates some of the possible scheduling protocols.
- The selection function determines which process, among ready processes, is selected next for execution. This function 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 and executing
 - e = time spent in execution so far
 - s = total service time required by the process, including e : generally this quantity is estimated.
- For example, the selection function $\max[w]$ indicates a first-come-first-served protocol.



Characteristics of Various Scheduling Protocols

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

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

FCFS = first come first served SPN = shortest process next

SRT = shortest remaining time

HRRN = highest response ratio next



Decision Mode

- The decision mode specifies the instants in time at which the selection function is applied. There are two general categories:
- **Non-preemptive**
 - Once a process is in the running state, it will continue until (a) it terminates or (b) blocks itself to wait for I/O or request some operating system service.
- **Preemptive**
 - 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.



Decision Mode (cont.)

- Preemptive protocols incur greater overhead than non-preemptive ones but will in general provide better service to the total population of processes, because they prevent any one process from monopolizing the processor for very long.
- In addition, the cost of preemption may be kept relatively low by using efficient process-switching mechanisms (with hardware support) and by providing a large main memory to key a high percentage of programs in main memory.



Process Scheduling Example

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

As we examine the various scheduling protocols we'll use this set of processes as a running example.

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

Alternatively, we can think of these as ongoing processes that require alternate use of the processor and I/O in repetitive fashion. In this case, the service time represents the processor time required in one cycle.

In either case, in terms of a queuing model, this quantity corresponds to the service time.



First-Come-First-Served (FCFS)

- The FCFS scheduling policy is the simplest scheduling algorithm we will examine.
- The FCFS protocol specifies that the first process to request the CPU is allocated to the CPU first.
- The FCFS protocol maintains the ready list as a straight queue (i.e., not a priority queue but a FIFO structure).
- The FCFS protocol is non-preemptive. Once a process is allocated to the CPU it keeps the CPU until it terminates or requests I/O (interrupt).
- While the FCFS protocol is easy to implement and oversee – it does not lead to a minimization of the average waiting time. The following example illustrates how the average waiting time is computed.

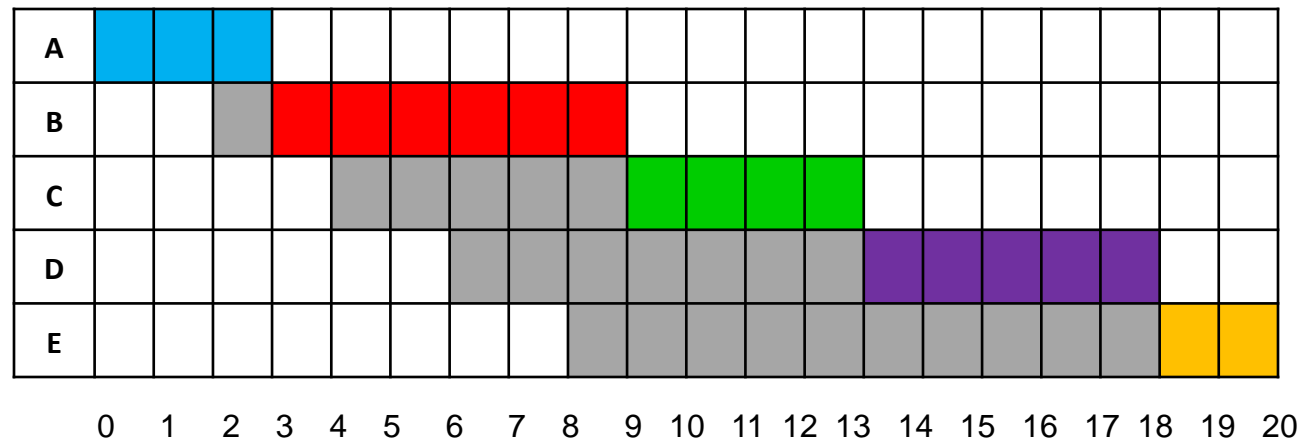


First-Come-First-Served (FCFS)

The Gantt chart notation:

Gray shading means process is waiting

Bright color means process is executing



- The waiting time (w) for process A = 0, for B = 1, C = 5, D = 7 and E = 10
- The average waiting time is then: $(0 + 1 + 5 + 7 + 10) / 5 = 23/5 = 4.6$
- The turnaround time (T_r) for process A = 3, B = 7, C = 9, D = 12, and E = 12
- The average turnaround time is then $(3 + 7 + 9 + 12 + 12) / 5 = 43/5 = 8.6$
- T_r/T_s : A = $3/3 = 1$, B = $7/6 = 1.17$, C = $9/4 = 2.25$, D = $12/5 = 2.4$, E = $12/2 = 6$
- The average for T_r/T_s : $(1 + 1.17 + 2.25 + 2.4 + 6) / 5 = 2.56$



First-Come-First-Served (FCFS)

- The average waiting time under a FCFS protocol is generally not minimal. Further, if the variance in CPU burst time is large, then the average waiting time will vary drastically depending upon the order in which the processes arrive for service in the ready queue. The following example illustrates the variance in the average waiting time of this protocol.
- Suppose the processes arrive in the order B, D, C, A, E. This causes their waiting times to become: $B = 0$, $D = 4$, $C = 7$, $A = 9$, $E = 10$. The average waiting time is then: $(0 + 4 + 7 + 9 + 10)/5 = 30/5 = 6$. Similarly the turnaround times become: $B = 6$, $D = 11$, $A = 15$, $C = 18$, and $E = 20$, with the average turnaround time being $(6 + 11 + 14 + 18 + 20)/5 = 65/5 = 13.8$.



First-Come-First-Served (FCFS)

- The FCFS protocol performs poorly in terms of maximizing the utilization of the CPU and the various I/O devices.
 - Consider the following scenario of one CPU bound process and many I/O bound processes currently in the system. Once the CPU bound process is allocated to the CPU it will keep it. During this time all of the I/O bound jobs will finish their I/O and reenter the ready queue to await their next turn on the CPU. While the I/O bound processes wait in the ready queue all of the I/O devices are idle. Eventually, the CPU bound process will finish its current CPU burst and requests I/O. Now all of the I/O bound processes in the ready queue will execute their CPU burst very quickly and move back into their I/O queues. At this point the CPU remains idle (as all processes are currently awaiting I/O completions. At some point the CPU bound process will reenter the CPU and the process will repeat as the I/O bound jobs will finish and arrive back in the ready queue. This is a *convoy effect* as all the I/O bound and short CPU processes wait for one CPU bound job to complete.
 - The overall effect is to lower both CPU utilization and I/O device utilization while increasing the average waiting time in the system for all processes (except perhaps for the one CPU bound process).



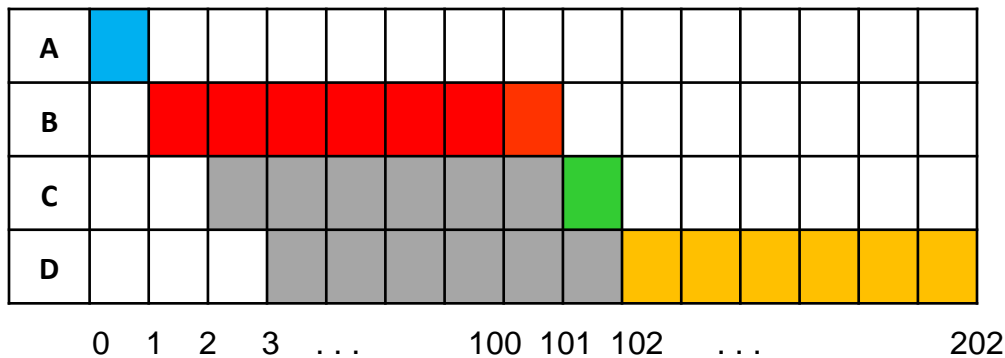
First-Come-First-Served (FCFS)

- 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 FCFS protocol is particularly unsuited to time-shared systems where the average response time begins to skyrocket if a single process is allowed to control the CPU for an extended period.
- In general, FCFS performs much better for long processes than short processes. This is illustrated by the example on the following page.



First-Come-First-Served (FCFS)

Process	Arrival Time	Service Time (T_s)	Start Time	Finish Time	Turnaround Time (T_r)	T_r/T_s
A	0	1	0	1	1	1
B	1	100	1	101	100	1
C	2	1	101	102	100	100
D	3	100	102	202	199	1.99



The normalized turnaround time for C is way out of line compared to the other processes: the total time it is in the system is 100 times the required processing time. This will happen whenever a short process arrives just after a long process. On the other hand, even in this extreme case, long processes do not do too badly. Process D has a turnaround time that is almost double that of C, but its normalized residence time is under 2.0.



First-Come-First-Served (FCFS)

- FCFS is not an attractive alternative on its own for a uni-processor system.
- It is sometimes combined with a priority scheme to provide an effective scheduler. In this case, the scheduler maintains a number of queues, one for each priority level, and dispatch within each queue on a FCFS basis.
- This is a common technique employed with feedback systems.



Round-Robin

- The round-robin protocol is a straightforward way to reduce the penalty that short jobs suffer under FCFS.
- Round-robin uses preemption based on a clock. A clock interrupt signal 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.
- This technique is also known as time-slicing, because each process is given a slice of time before being preempted.

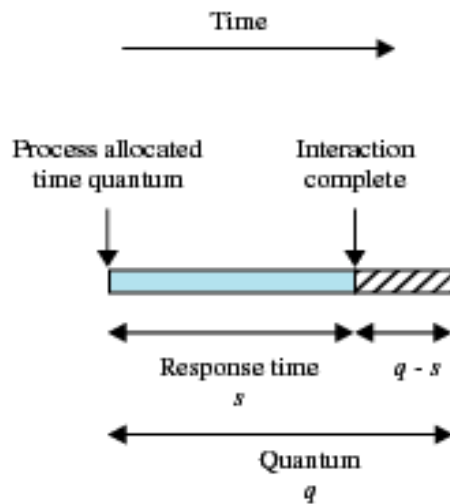


Round-Robin

- With round-robin, 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.
- On the other hand, there is processing overhead involved in handling the clock interrupt and performing the scheduling and dispatching functions.
- This implies that very short time quantum should be avoided.
- One useful guideline 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 quanta. (See next slide.)

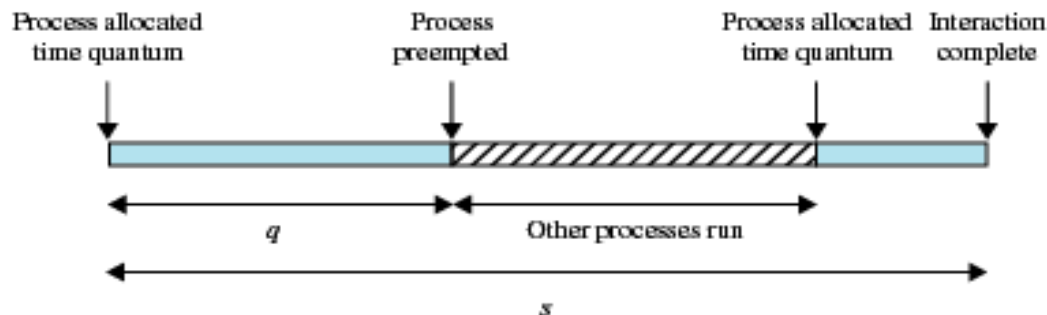


Effect of Size on Preemption Time Quantum



(a) Time quantum greater than typical interaction

Figure (a) shows the effect when the time quantum is larger than the typical interaction time. Typical processes complete in one time quantum.



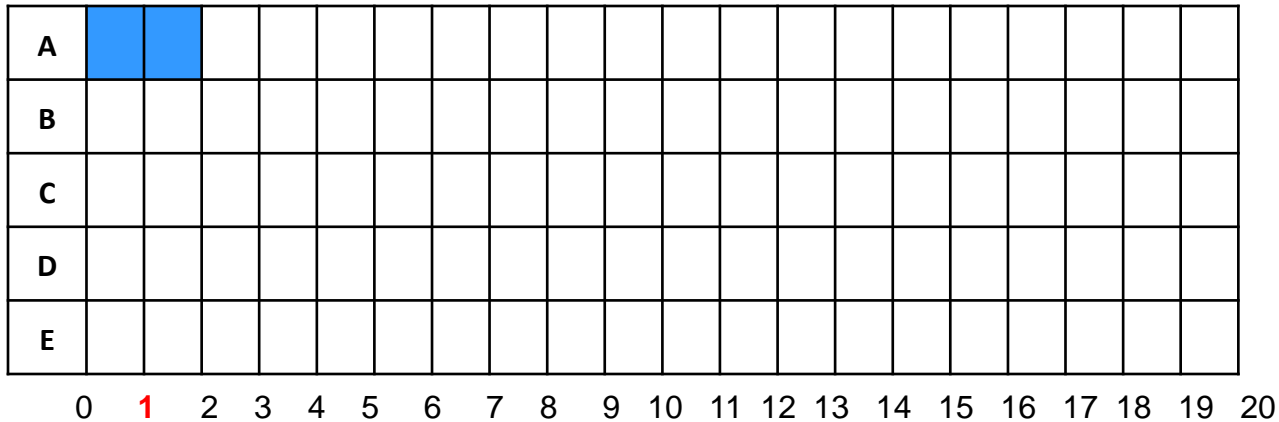
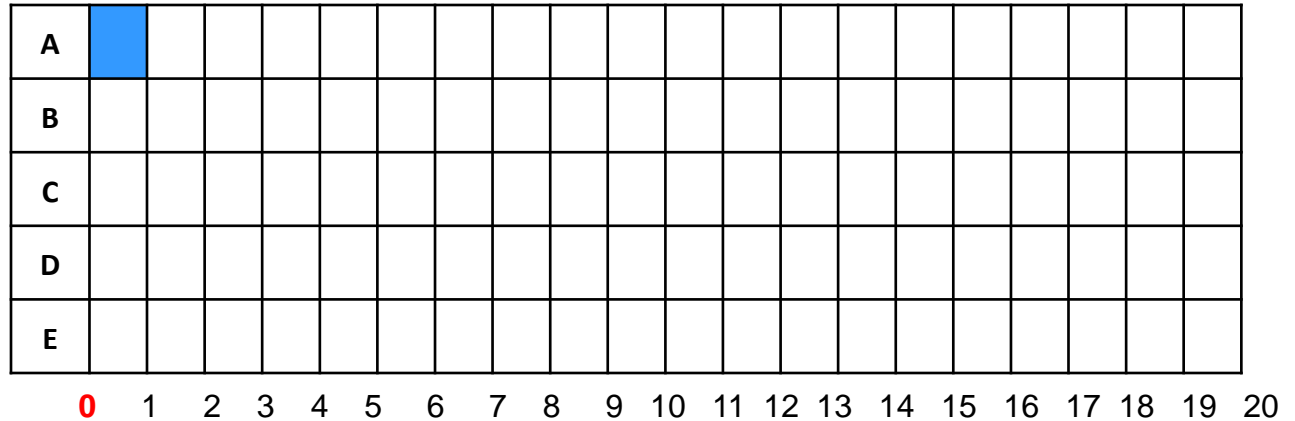
(b) Time quantum less than typical interaction

Figure (b) illustrates the case when the time quantum is smaller than the typical interaction time. Typical processes require at least two time quantum.

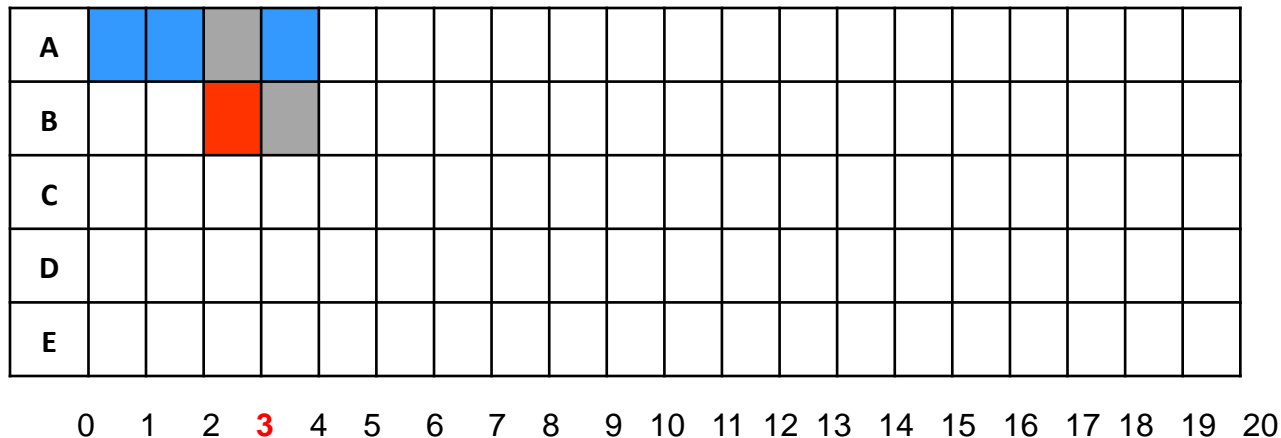
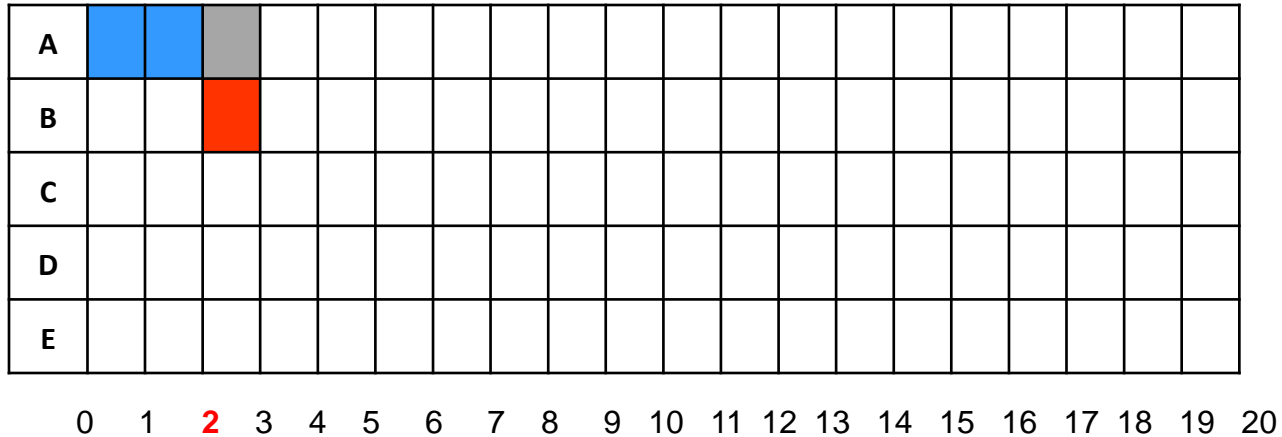


Process	Arrival Time	Service Time
A	0	3
B	2	6
C	4	4
D	6	5
E	8	2

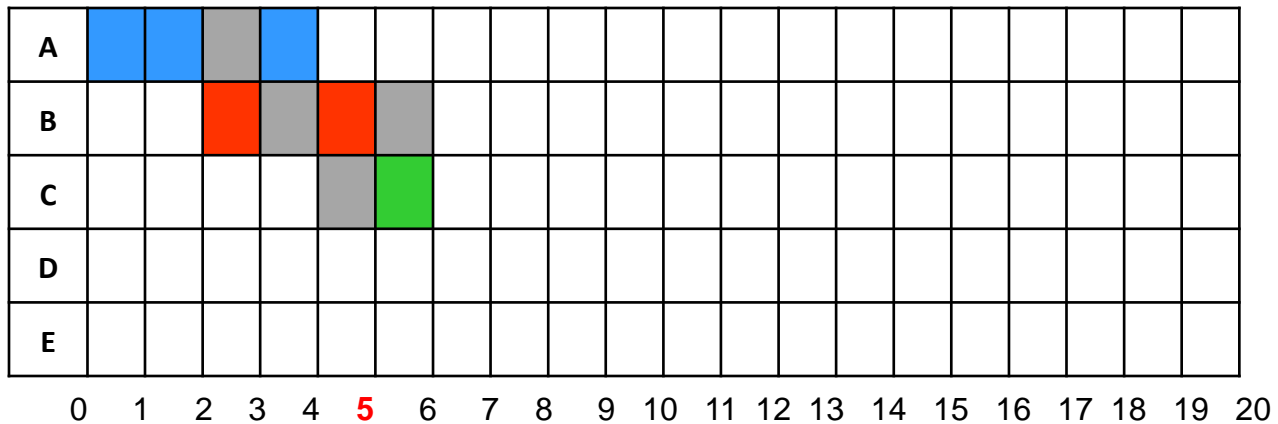
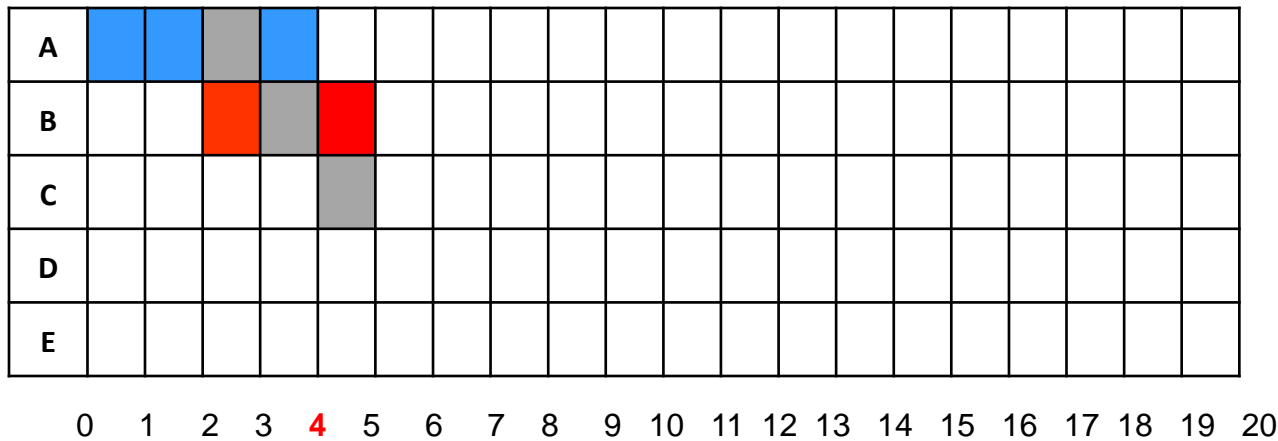
Round-Robin (quantum = 1)



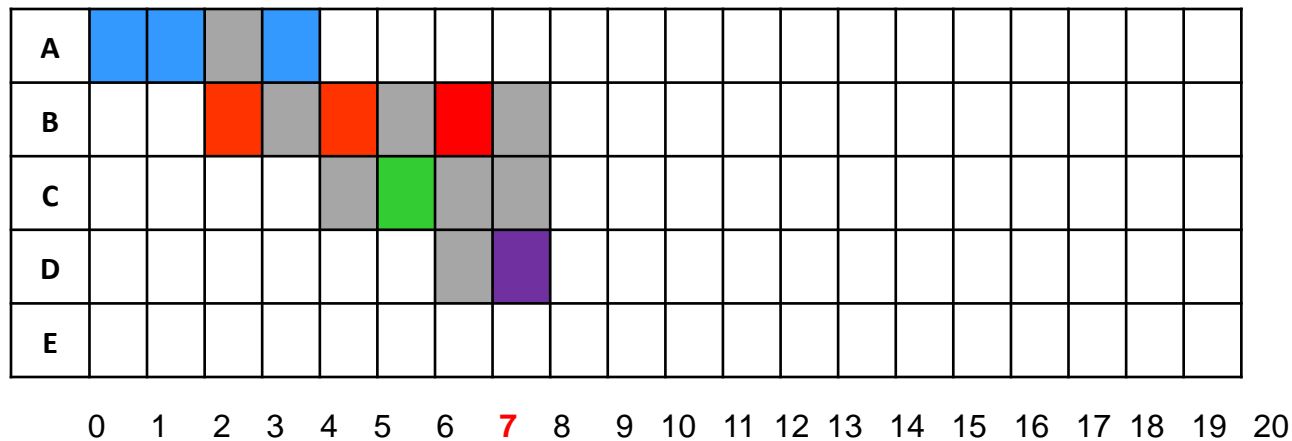
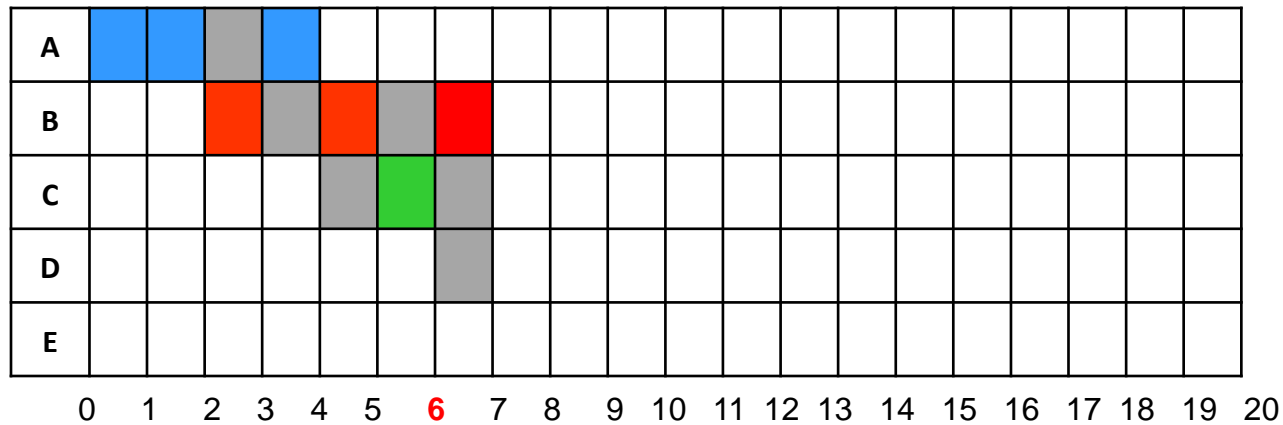
Round-Robin (quantum = 1)



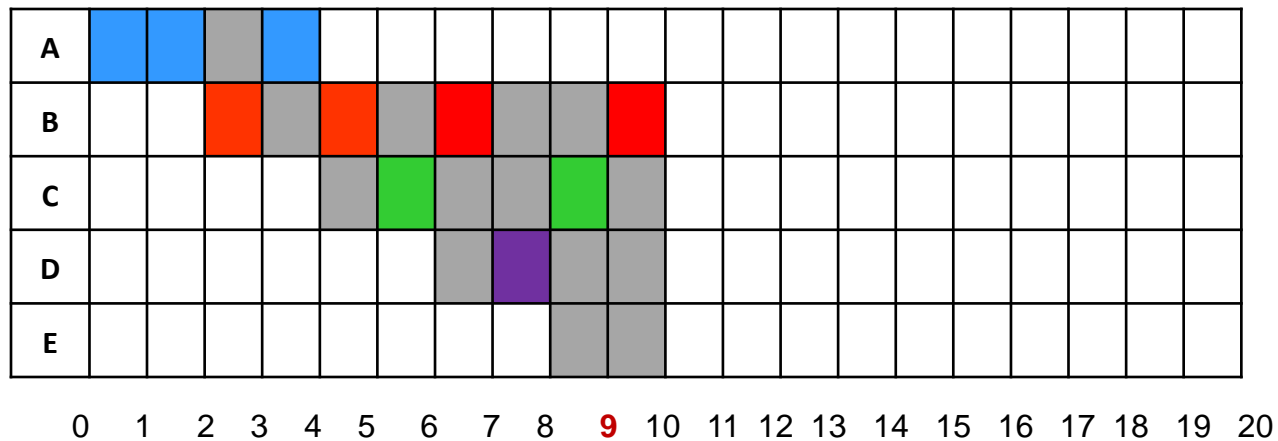
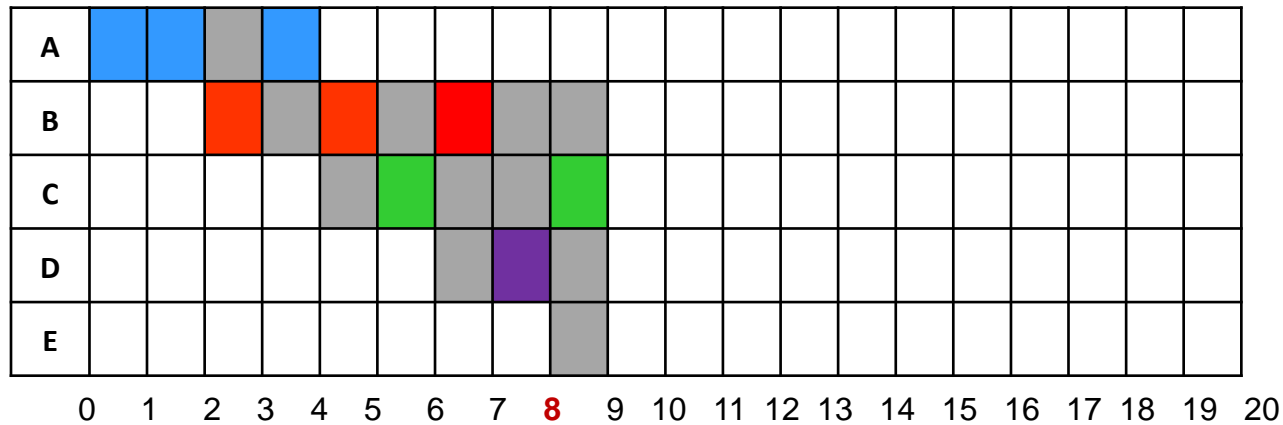
Round-Robin (quantum = 1)



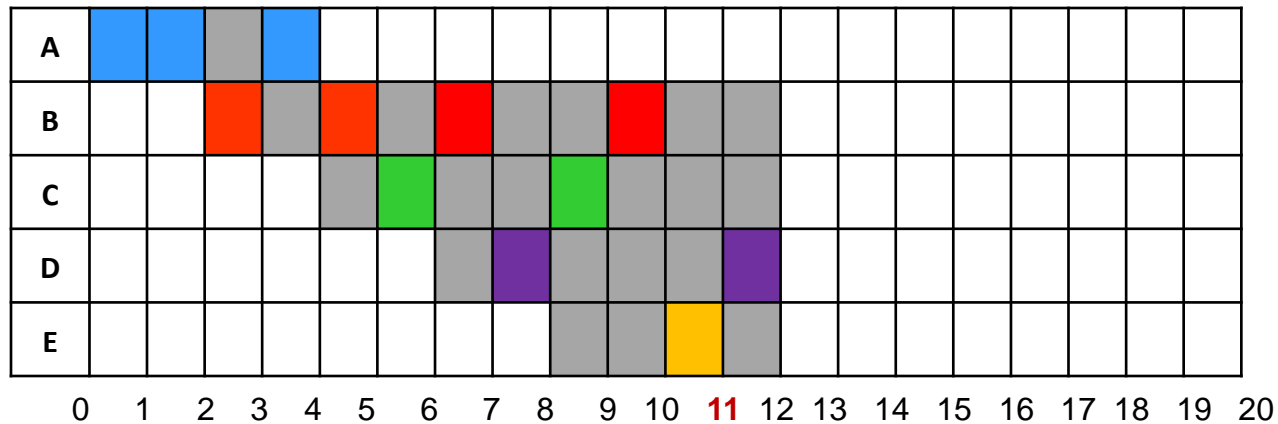
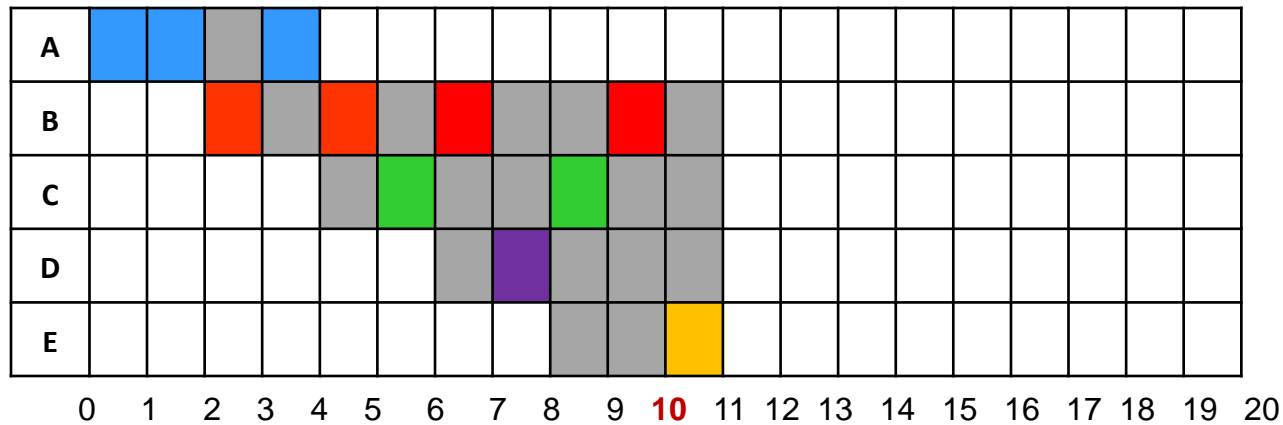
Round-Robin (quantum = 1)



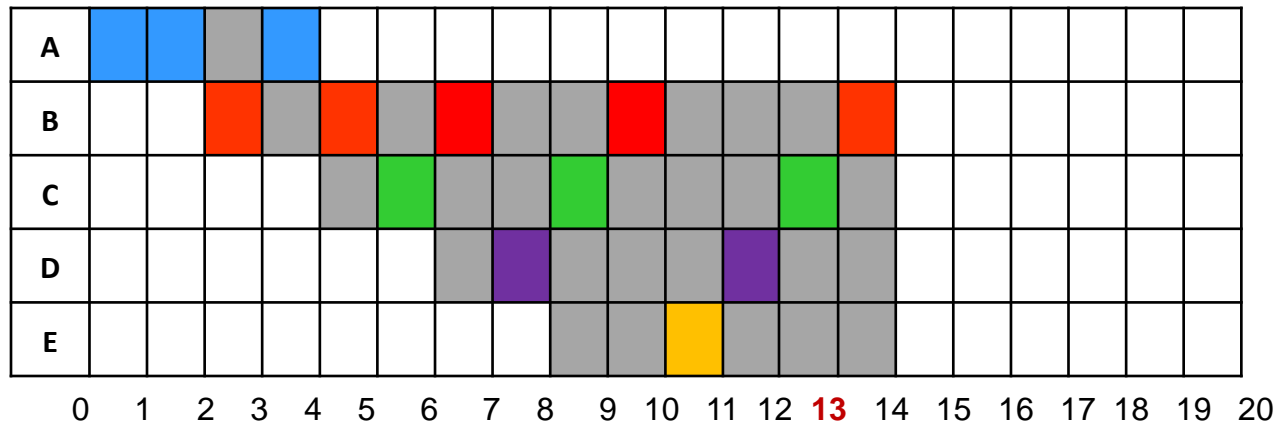
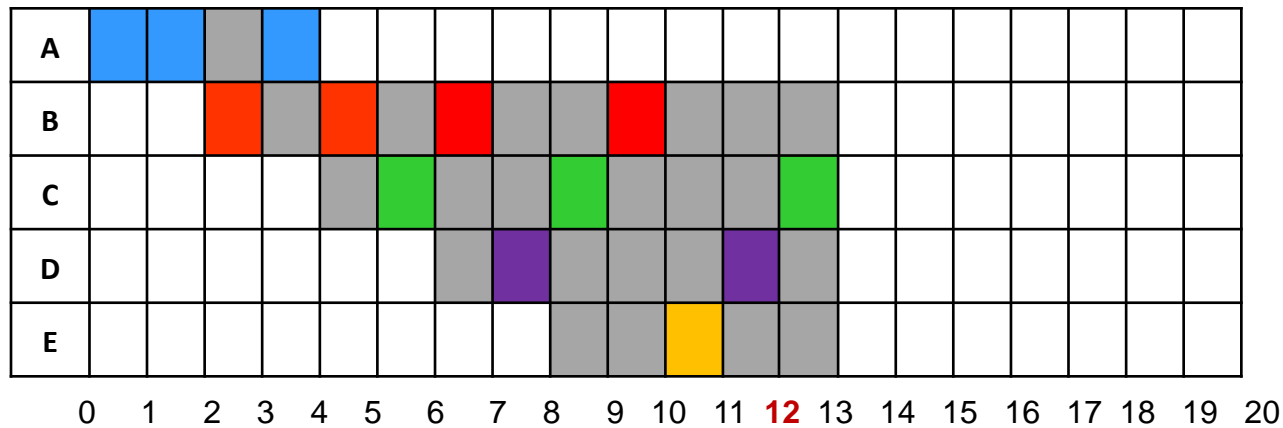
Round-Robin (quantum = 1)



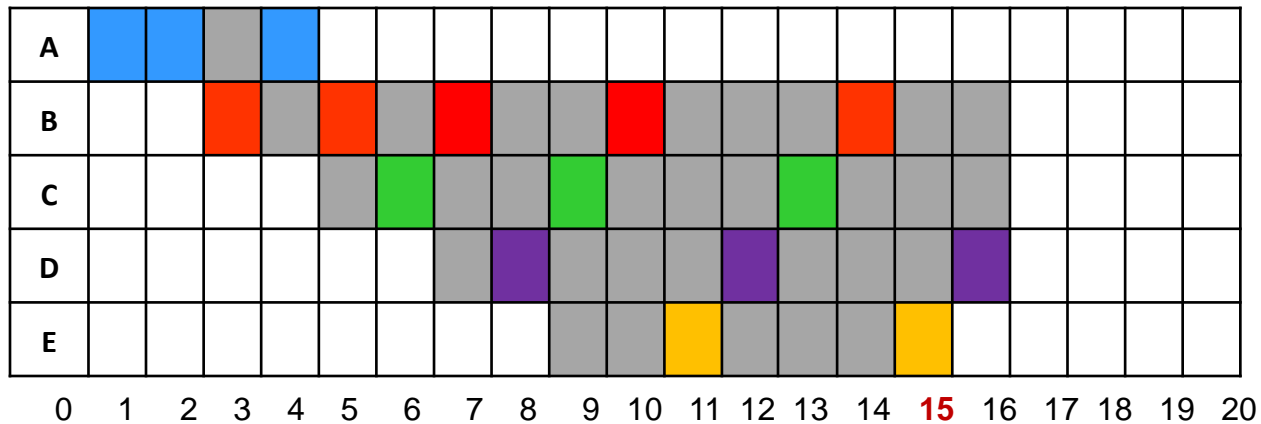
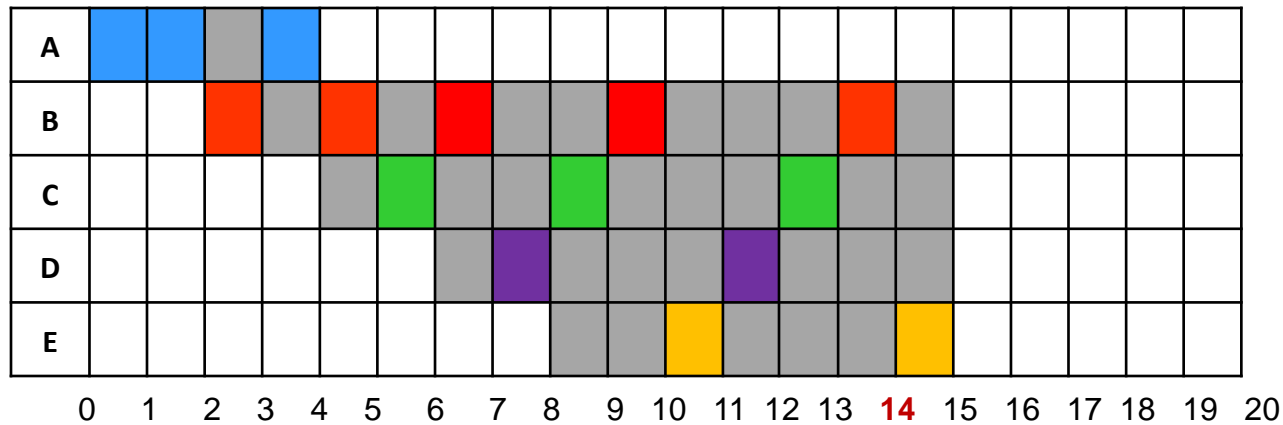
Round-Robin (quantum = 1)



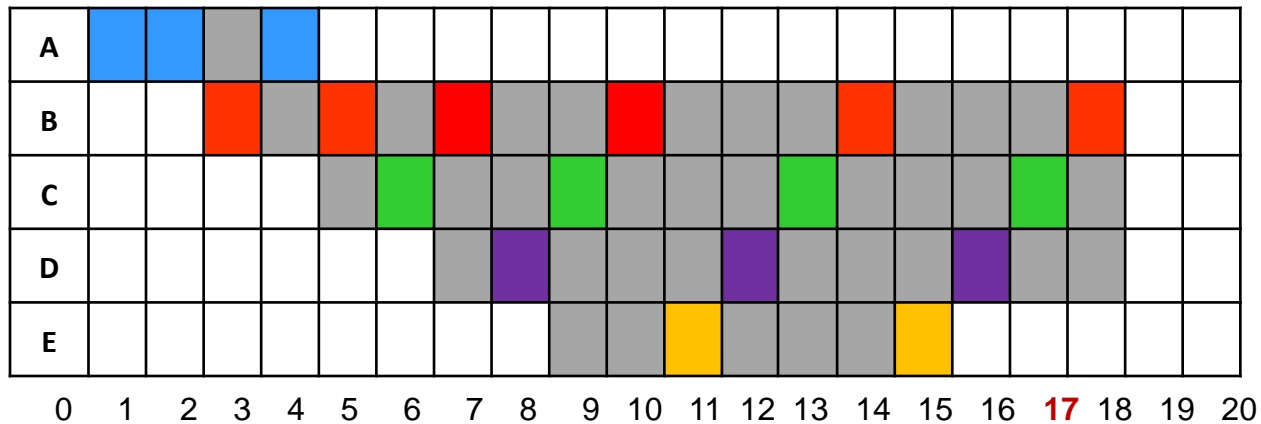
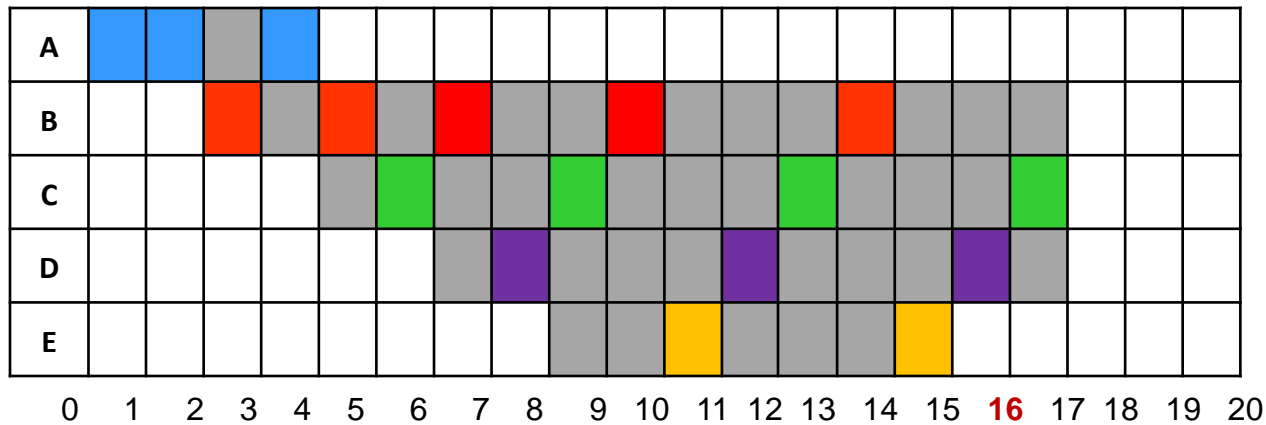
Round-Robin (quantum = 1)



Round-Robin (quantum = 1)



Round-Robin (quantum = 1)



Round-Robin (quantum = 1)

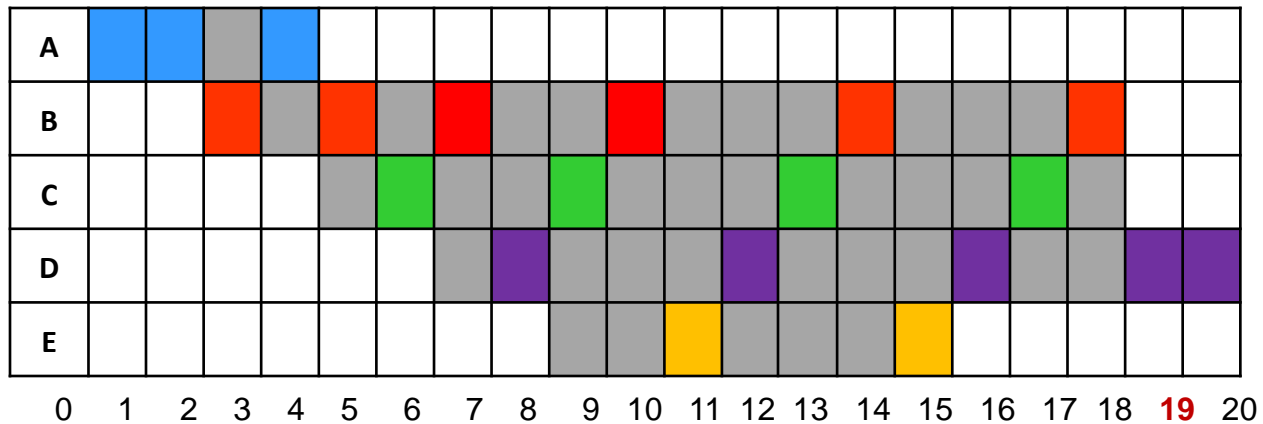
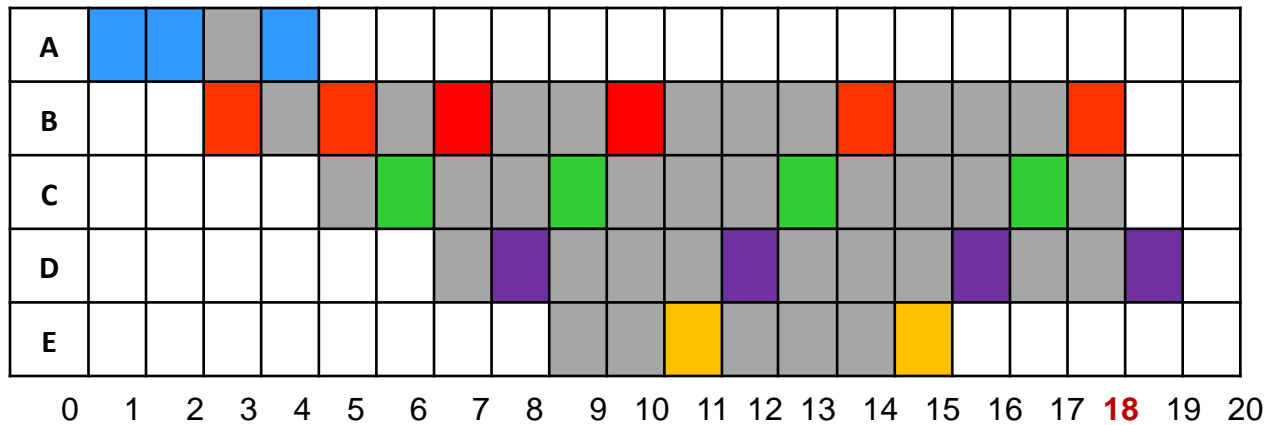
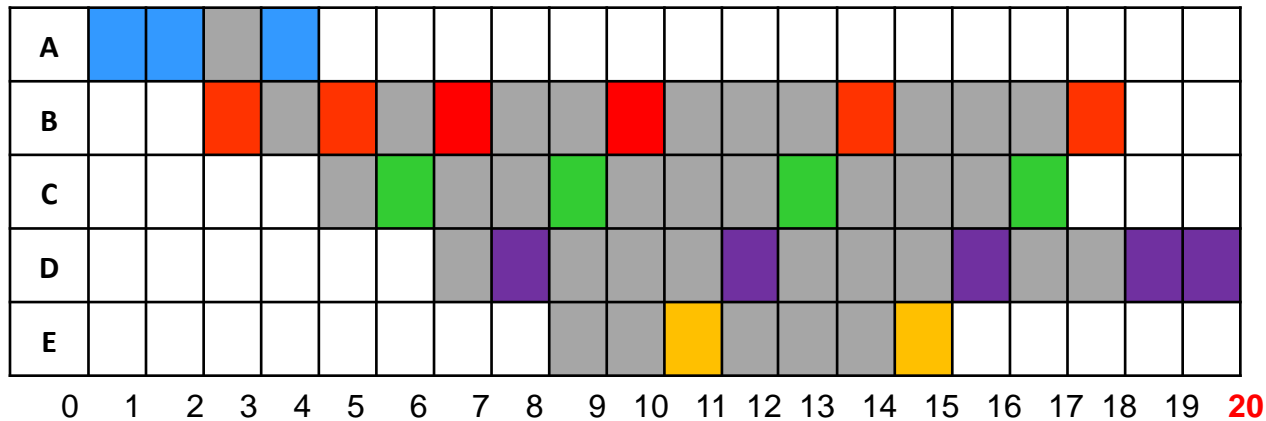


Diagram Illustrating Round-Robin queue and
Processor allocation for quantum = 1

Time	Queue (R - F)	Process In CPU	Action
0	A	A	A arrives and gets CPU
1		A	
2	A, B	B	B arrives, A loses CPU
3	B, A	A	A completes ($T_r = 4$)
4	C, B	B	C arrives, B gets CPU
5	B, C	C	
6	C, D, B	B	D arrives
7	B, C, D	D	
8	D, E, B, C	C	E arrives
9	C, D, E, B	B	
10	B, C, D, E	E	
11	E, B, C, D	D	
12	D, E, B, C	C	
13	C, D, E, B	B	
14	B, C, D, E	E	E completes ($T_r = 7$)
15	B, C, D	D	
16	D, B, C	C	C completes ($T_r = 13$)
17	D, B	B	B completes ($T_r = 16$)
18	D	D	
19		D	D completes ($T_r = 14$)



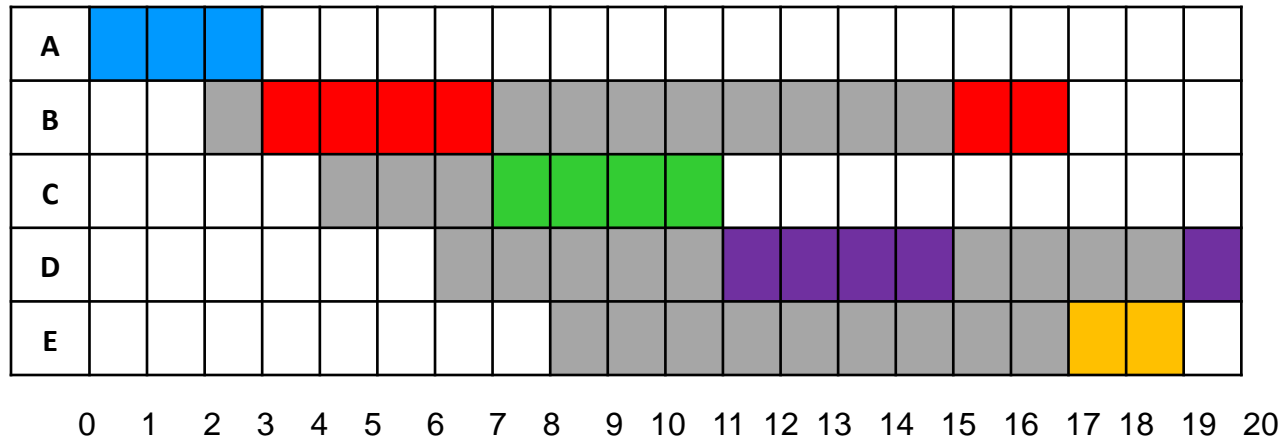
Round-Robin (quantum = 1)



- Waiting time of process A = 1, B = 10, C = 9, D = 9, and E = 5
- Average waiting time: $(1 + 10 + 9 + 9 + 5)/5 = 34/5 = 6.8$
- Turnaround time (T_r): A = 4, B = 16, C = 13, D = 14, and E = 7
- Average turnaround time: $(4 + 16 + 13 + 14 + 7)/5 = 54/5 = 10.8$
- T_r/T_s : A = $4/3=1.33$, B = $16/6=2.66$, C = $13/4=3.25$, D = $14/5=2.8$, E = $7/2 = 3.5$
- Average T_r/T_s : $(1.33 + 2.66 + 3.25 + 2.8 + 3.5)/5 = 13.54/5 = 2.71$



Round-Robin (quantum = 4)



- Waiting time of process A = 0, B = 9, C = 3, D = 9, and E = 9
- Average waiting time: $(0 + 10 + 7 + 9 + 7)/5 = 33/5 = 6.6$
- Turnaround time (T_r): A = 3, B = 15, C = 7, D = 14, and E = 11
- Average turnaround time: $(3 + 15 + 7 + 14 + 11)/5 = 50/5 = 10.0$
- T_r/T_s : A = $3/3=1.0$, B = $15/6=2.5$, C = $7/4=1.75$, D = $14/5=2.8$, E = $11/2 = 5.5$
- Average T_r/T_s : $(1.0 + 2.5 + 1.75 + 2.8 + 5.5)/5 = 13.55/5 = 2.71$



Round-Robin

- Round-robin is particularly effective in a general-purpose time-sharing system or transaction processing system.
- One drawback to round-robin is its relative treatment of CPU-bound and I/O-bound processes. Generally, an I/O bound process has a shorter processor burst (the amount of time spent executing between I/O operations) than a CPU-bound process.
- With a mix of CPU and I/O bound processes the following will happen: An I/O bound process uses the CPU for a short period of time and is then blocked for I/O; it waits for the I/O to complete then joins the ready queue. On the other hand, a CPU bound process generally uses its entire quantum while executing and immediately returns to the ready queue. Thus, CPU 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.

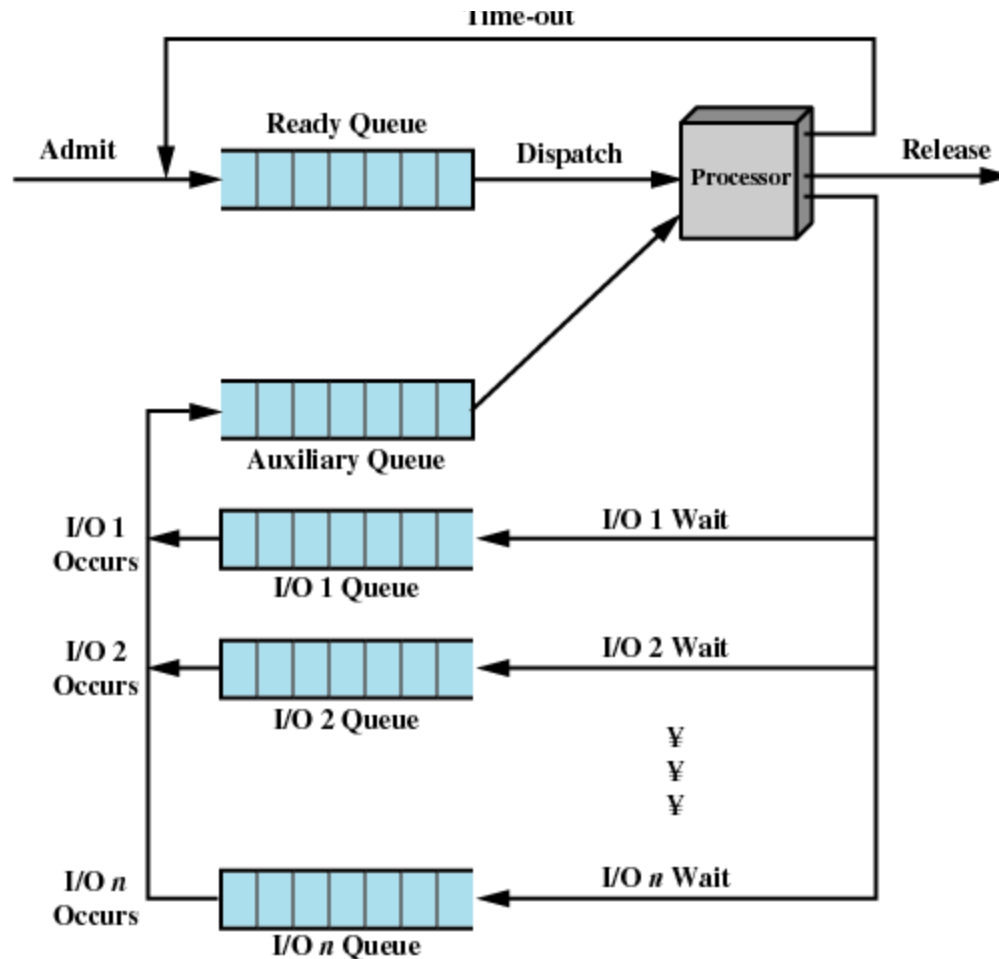


Round-Robin

- One possible solution to this problem that has been developed is referred to as a virtual round-robin (VRR) which avoids this unfairness to I/O bound processes.
- In VRR, new processes arrive and join the ready queue, which is managed on a 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. (So far, this method is no different from what we've seen previously).
- 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 are given 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. This method is illustrated on the next slide.



Set-up For Virtual Round-Robin

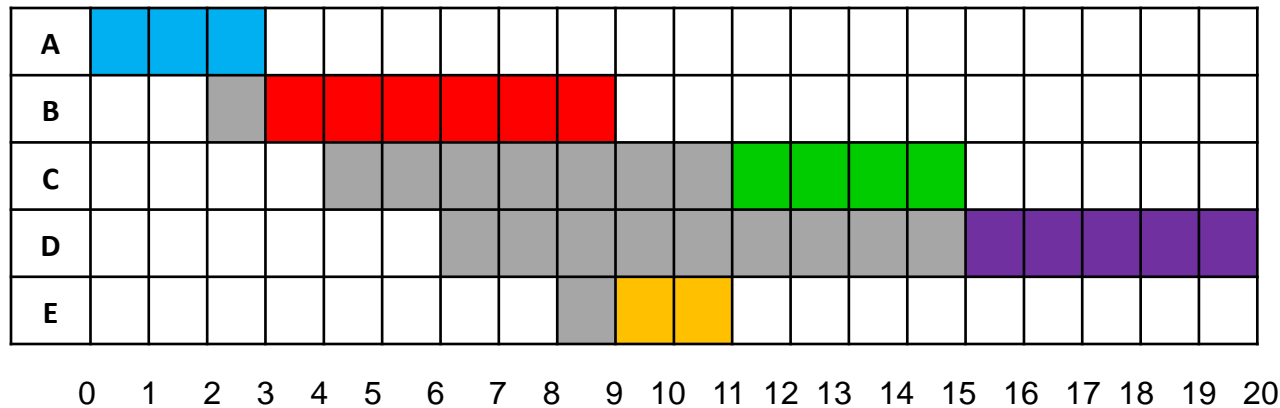


Shortest Process Next (SPN)

- Shortest Process Next (SPN) is another approach to reduce the bias in favor of long processes that is inherent with FCFS.
- SPN is non-preemptive.
- The process with the shortest expected processing time is selected next by the scheduler. Thus, a short process job will jump to the head of the queue passing longer jobs.



Shortest Process Next



Notice that Process E receives service much sooner under SPN than it did under FCFS.

- Waiting times: $A = 0$, $B = 1$, $C = 7$, $D = 9$, $E = 1$
- Average waiting time = $(0 + 1 + 7 + 9 + 1)/5 = 3.6$
- Turnaround times (T_r): $A = 3$, $B = 7$, $C = 11$, $D = 14$, $E = 3$
- Average turnaround time = $(3 + 7 + 11 + 14 + 3)/5 = 38/5 = 7.6$
- T_r/T_s : $A = 3/3 = 1$, $B = 7/6 = 1.17$, $C = 11/4 = 2.75$, $D = 14/5 = 2.8$, $E = 3/2 = 1.5$
- Average $T_r/T_s = (1 + 1.17 + 2.75 + 2.8 + 1.5)/5 = 9.22/5 = 1.84$



Shortest Process Next

- In terms of response time, overall performance has improved under this protocol. However, the variability of response times has also increased, especially for longer processes, and thus predictability of longer processes is reduced.
- One difficulty with the SPN protocol is the need to know or accurately predict the required processing time for each process.
- If the estimated time for a process is not correct, the operating system may abort it.
- Possibility of starvation for longer processes occurs if there is a steady supply of short processes.
- Not favored for time-sharing or transaction processing environments due to the lack of preemption.



Shortest Remaining Time (SRT)

- The Shortest Remaining Time (SRT) protocol is a preemptive version of SPN.
- In SRT, the scheduler always chooses the process that has the shortest expected remaining processing time.
- When a new process joins the ready queue, it may have a shorter remaining time than the currently running process. If this occurs, the scheduler may preempt the current process when the new process arrives.
- As with SPN, the SRT scheduler must have an estimate of the processing time in order to perform the selection function.
- Again, there is the possibility of starvation for longer processes.

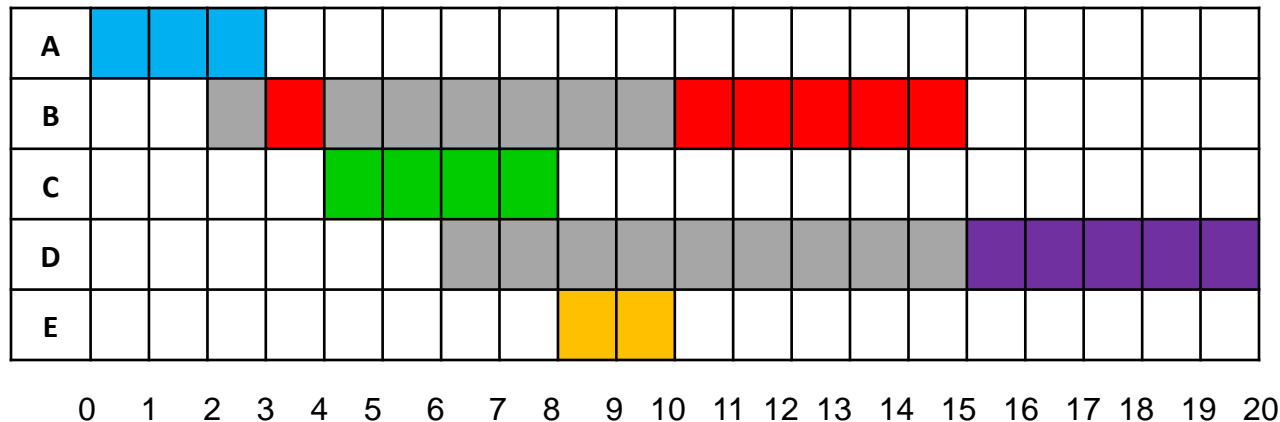


Shortest Remaining Time (SRT)

- SRT does not have the bias in favor of long processes that we saw with FCFS.
- Unlike RR, no additional interrupts are generated, which reduces the overhead.
- On the other hand, elapsed service time must be recorded which contributes to overhead.
- SRT typically gives superior turnaround time performance when compared to SPN, because a short job is given immediate preference to a running longer job.
- Note in the example on the next page that the three shortest processes (A, C, and E) all receive immediate service, which produces a normalized turnaround time of 1.0 for each.



Shortest Remaining Time



- Waiting times: A = 0, B = 7, C = 0, D = 9, E = 0
- Average waiting time = $(0 + 7 + 0 + 9 + 0)/5 = 16/5 = 3.2$
- Turnaround times (T_r): A = 3, B = 13, C = 4, D = 14, E = 2
- Average turnaround time = $(3 + 13 + 4 + 14 + 2)/5 = 36/5 = 7.2$
- T_r/T_s : A = $3/3 = 1$, B = $13/6 = 2.17$, C = $4/4 = 1$, D = $14/5 = 2.8$, E = $2/2 = 1$
- Average $T_r/T_s = (1 + 2.17 + 1 + 2.8 + 1)/5 = 7.97/5 = 1.59$



Highest Response Ratio Next (HRRN)

- Highest Response Ratio Next (HRRN) uses the normalized turnaround time, which is the ratio T_r/T_s calculated as:

$$\text{response ratio} = \frac{\text{time waiting} + T_s}{T_s}$$

- HRRN attempts to minimize the average of this ratio over all processes.
- In general, it is not possible to know in advance what the exact service time will be, but it can be approximated, based either on past history or some input from the user or a configuration manager.
- The scheduler's decision is now determined as follows: when the current process completes or is blocked, choose the ready process with the greatest value of this ratio.



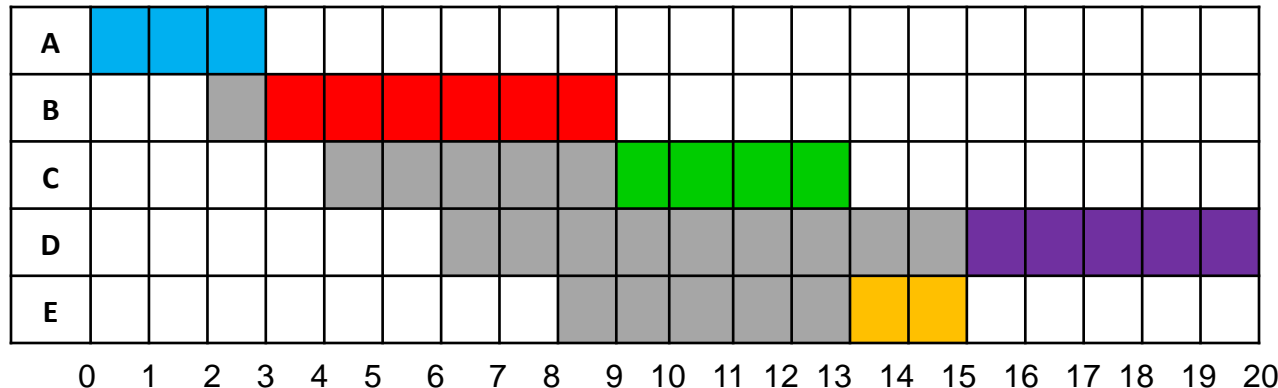
Highest Response Ratio Next (HRRN)

- This approach is attractive because it accounts for the age of the process.
- While shorter jobs are favored (a smaller denominator results in a larger ratio), aging without service increases the ratio (since T_r gets larger) so that a longer process will eventually get past competing shorter jobs.



Highest Response Ratio Next (HRRN)

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



At time = 3, only process B is dispatchable, so HRRN is not calculated.

At time 9, processes C, D, and E are all dispatchable, so HRRNs are calculated as:

$$C = (5+4)/4 = 2.25$$

$$D = (3+5)/5 = 1.6$$

$$E = (1+2)/2 = 1.5$$

So C is dispatched.

At time 13, processes D and E are dispatchable, so HRRNs are calculated as:

$$D = (7+5)/5 = 2.4$$

$$E = (5+2)/2 = 3.5$$

So E is dispatched.

- Waiting times: A = 0, B = 1, C = 5, D = 9, E = 5
- Average waiting time = $(0 + 1 + 5 + 9 + 5)/5 = 19/5 = 3.8$
- Turnaround times (T_r): A = 3, B = 7, C = 9, D = 14, E = 7
- Average turnaround time = $(3 + 7 + 9 + 14 + 7)/5 = 40/5 = 8.0$
- T_r/T_s : A = $3/3 = 1$, B = $7/6 = 1.17$, C = $9/4 = 2.25$, D = $14/5 = 2.8$, E = $7/2 = 3.5$
- Average $T_r/T_s = (1 + 1.17 + 2.25 + 2.8 + 3.5)/5 = 10.72/5 = 2.14$



Feedback Techniques (FB)

- If the scheduler has no way of knowing or estimating with any degree of accuracy the relative length of the various processes it may schedule, then none of SPN, SRT, or HRRN can be used.
- Another technique for establishing a preference for shorter jobs is to penalize jobs that have been running longer. In other words, if the scheduler cannot focus on the time remaining to execute, then let it focus on the time spent in execution so far.
- The mechanism for doing this is as follows:

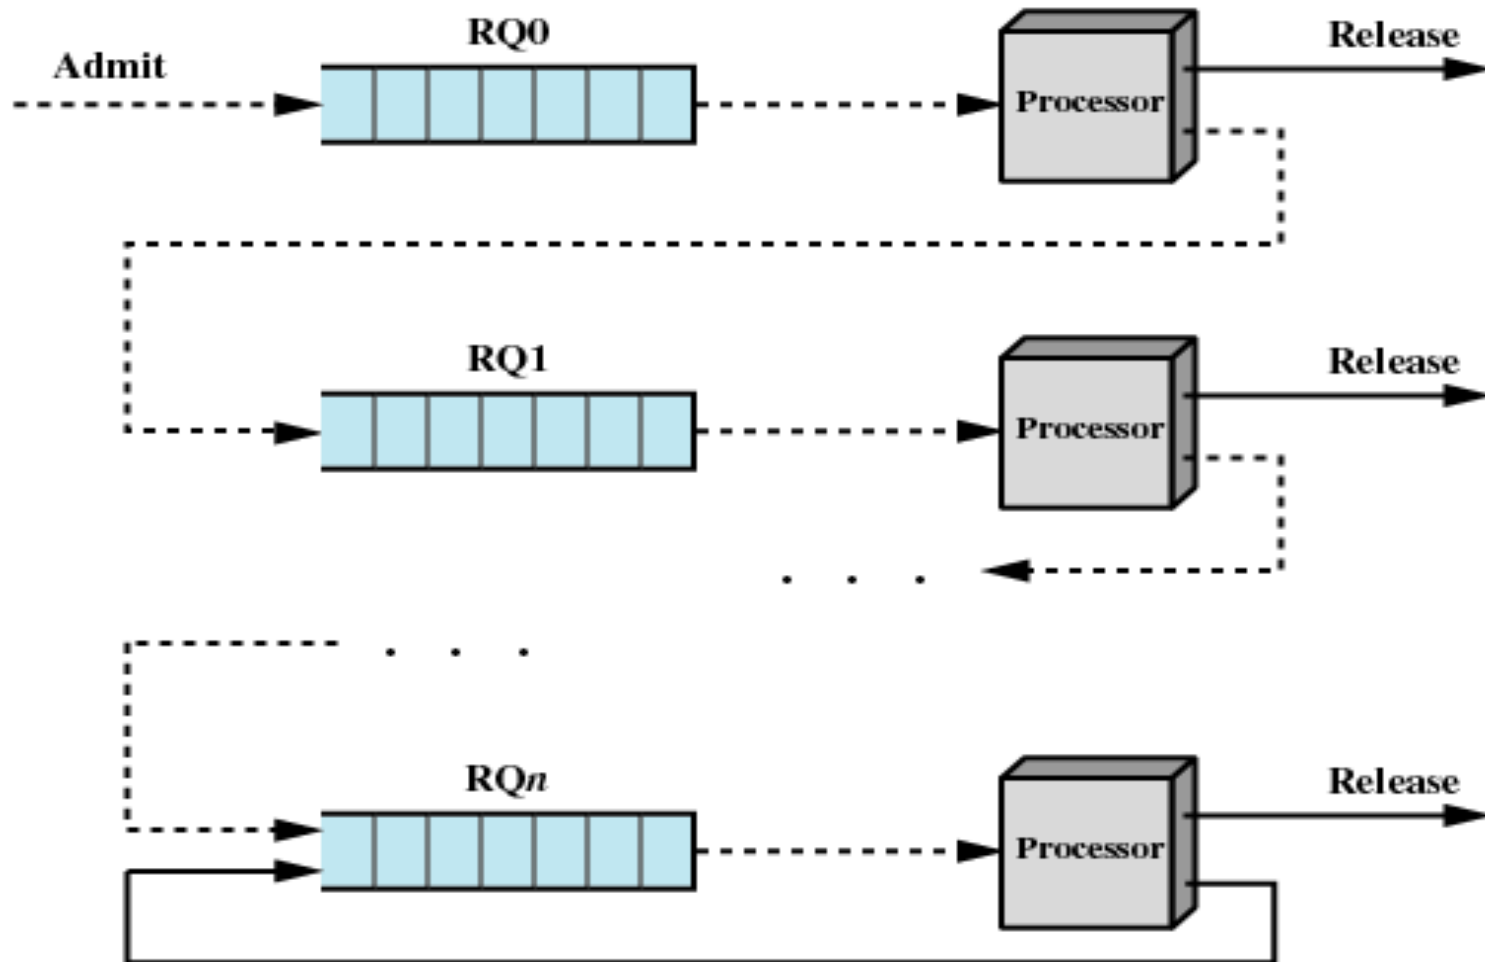


Feedback Techniques

- Scheduling is done on a preemptive basis (assuming some time quantum), and a dynamic priority mechanism is applied.
 - When a process first enters the system, it is placed in RQ0 (see diagram on next page).
 - After its first preemption, when it returns to the ready state, it is placed in RQ1 (next lowest priority).
 - Each subsequent time that it is preempted, it is demoted to the next lower priority queue.
 - Within each priority queue, except for the lowest level queue, a simple FCFS mechanism is used. Once in the lowest level queue a process cannot have a lower priority so it is repeatedly returned to this queue until it completes execution. So this queue is handled in round-robin fashion.



Feedback Techniques



Feedback Techniques

- Short processes will complete quickly, without moving very far down the hierarchy of ready queues.
- A longer process will gradually drift down the priority queue hierarchy.
- Thus newer shorter processes are favored over older longer processes.
- There are a number of variations on the feedback protocol. In the simplest case, preemption is performed in the same fashion as for round-robin, i.e., at periodic intervals. (This is shown in the example on page 63 for a time quantum of 1.)



Feedback Techniques

- For the example set of processes assuming a time quantum of 1 and 3 levels of feedback queues, we would have the following situation:

time	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
RQ0	A		B		C		D		E											
RQ1		A		B,A	B	C,B	C	D,C	D	E,D	E									
RQ2							B	B	C,B	C,B	D,C,B	D, C	B, D	C, B	D, C	B, D	B	D	B	
Dispatched	A	A	B	A	C	B	D	C	E	D	E	B	C	D	B	C	D	B	D	B
Finished – leaves queue				A							E					C			D	B



Feedback Techniques

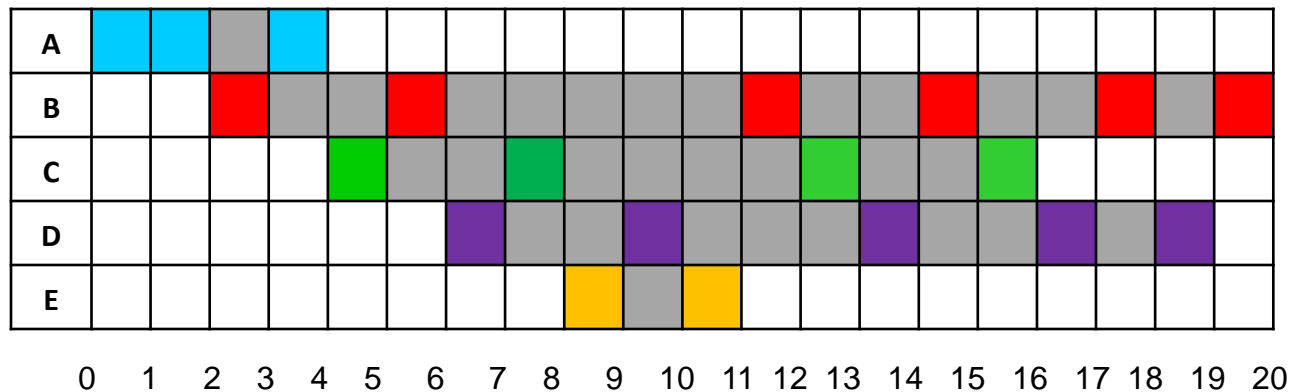
- For the example set of processes assuming a time quantum of 1 and 3 queue levels with no feedback if only single ready process, we would have the following situation (green CPU cell indicates process ends):

Time interval	0-1	1-2	2-3	3-4	4-5	5-6	6-7	7-8	8-9	9-10	10-11	11-12	12-13	13-14	14-15
RQ0	A	A	B		C		D		E						
RQ1			A	B, A	B	C,B	C	D,C	D	E,D	E				
RQ2							B	B	C,B	C,B	D,C,B	D,C	B,D	C,B	D,C
CPU	A	A	B	A	C	B	D	C	E	D	E	B	C	D	B

Time interval	15-16	16-17	17-18	18-19	19-20
RQ0					
RQ1					
RQ2	B,D	B	D	B	
CPU	C	D	B	D	B



Feedback Techniques



Time quantum = 1 for all queue levels – assume 3 queue levels

Waiting time: A = 1, B = 12, C = 8, D = 8, E = 1

Average waiting time = $(8 + 12 + 8 + 8 + 0)/5 = 36/5 = 7.2$

Turnaround times (T_r): A = 4, B = 18, C = 12, D = 13, E = 3

Average turnaround time = $(4 + 18 + 12 + 13 + 3)/5 = 50/5 = 10.0$

T_r/T_s : A = $4/3 = 1.33$, B = $18/6 = 3$, C = $12/4 = 3$, D = $13/5 = 2.6$, E = $3/2 = 1.5$

Average $T_r/T_s = (1.33 + 3 + 3 + 2.6 + 1.5)/5 = 11.43/5 = 2.286 = 2.29$



Feedback Techniques

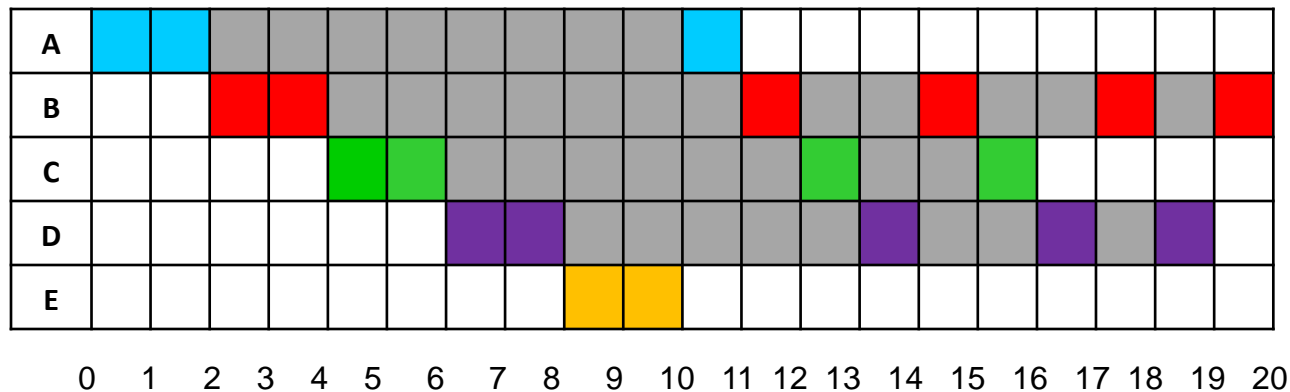
- For the example set of processes assuming a time quantum of 1 and 3 queue levels with mandatory feedback, we would have the following situation (green CPU cell indicates process ends):

Time interval	0-1	1-2	2-3	3-4	4-5	5-6	6-7	7-8	8-9	9-10	10-11	11-12	12-13	13-14
RQ0	A		B		C		D		E					
RQ1		A		B		C		D		E				
RQ2			A	A	B,A	B,A	C,B,A	C,B,A	D,C,B,A	D,C,B,A	D,C,B	A,D,C	B,A,D	C,B
CPU	A	A	B	B	C	C	D	D	E	E	A	B	C	D

Time interval	14-15	15-16	16-17	17-18	18-19	19-20
RQ0						
RQ1						
RQ2	D,C	B,D	B	D	B	
CPU	B	C	D	B	D	B



Feedback Techniques



Time quantum = 1 for all queue levels – assume 3 queue levels

Waiting time: A = 8, B = 12, C = 8, D = 8, E = 0

Average waiting time = $(8 + 12 + 8 + 8 + 0)/5 = 36/5 = 7.2$

Turnaround times (T_r): A = 11, B = 18, C = 12, D = 13, E = 2

Average turnaround time = $(11 + 18 + 12 + 13 + 2)/5 = 56/5 = 11.2$

T_r/T_s : A = $8/3 = 2.67$, B = $18/6 = 3$, C = $12/4 = 3$, D = $13/5 = 2.6$, E = $2/2 = 1.0$

Average $T_r/T_s = (2.67 + 3 + 3 + 2.6 + 1.0)/5 = 12.27/5 = 2.45$



Feedback Techniques

- There is a problem with this simple mechanism however, in that the turnaround time of longer processes can grow at an alarming rate.
- Starvation is possible, if new jobs are entering the system frequently.
- To compensate for this, the preemption times can be varying depending on the queue (feedback level) in the following fashion:
 - A process scheduled from queue RQ0 is allowed to execute for 1 time quantum and then is preempted.
 - A process scheduled from queue RQ1 is allowed to execute for 2 time quanta.
 - In general, a process scheduled from queue RQ_i is allowed to execute for 2^i time quanta. This is shown on the next two pages.



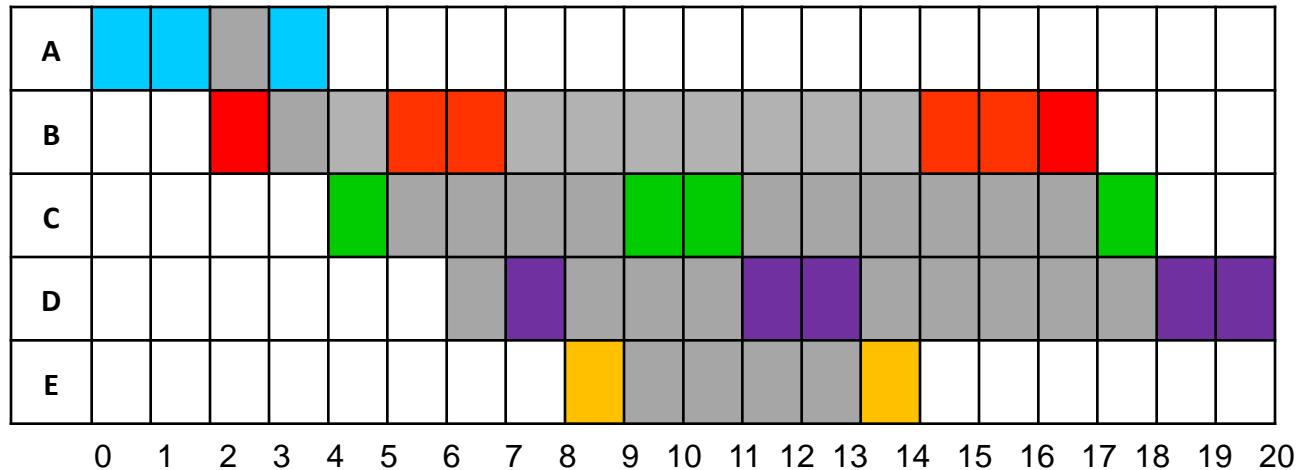
Feedback Techniques

- For the example set of processes assuming a time quantum of 2^i for each level of feedback and 3 levels of feedback queues, we would have the following situation:

time	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
RQ0 (q = 1)	A		B		C		D		E											
RQ1 (q = 2)		A		B, A	B	C, B	C	D, C	D, C	E, D, C	E, D	E	E							
RQ2 (q = 4)								B	B	B	C, B	C, B	C, B	D, C, B	D, C	D, C	D, C	D		
Dispatched	A	A	B	A	C	B	B	D	E	C	C	D	D	E	B	B	B	C	D	D
Finished – leaves queue				A										E			B	C		D



Feedback Techniques



Time quantum = 2^i for all queue levels

Waiting times: A = 1, B = 9, C = 10, D = 9, E = 4

Average waiting time = $(1 + 9 + 10 + 9 + 4)/5 = 33/5 = 6.6$

Turnaround times (Tr): A = 4, B = 15, C = 14, D = 14, E = 6

Average turnaround time = $(4 + 15 + 14 + 14 + 6)/5 = 53/5 = 10.7$

Tr/Ts: A = $4/3 = 1.33$, B = $15/6 = 2.5$, C = $14/4 = 3.5$, D = $14/5 = 2.8$, E = $6/2 = 3$

Average Tr/Ts = $(1.33 + 2.5 + 3.5 + 2.8 + 3)/5 = 13.13/5 = 2.626 = 2.63$



Overall Comparison Chart

	Process Arrival Time Service Time (T_S)	A 0 3	B 2 6	C 4 4	D 6 5	E 8 2	Mean
FCFS	Finish Time Turnaround - T_R T_R/T_S	3 3 1.00	9 7 1.17	13 9 2.25	18 12 2.40	20 12 6.00	8.60 2.56
RR ($q = 1$)	Finish Time Turnaround - T_R T_R/T_S	4 4 1.33	18 16 2.67	17 13 3.25	20 14 2.8	15 7 3.5	10.8 2.71
RR ($q = 4$)	Finish Time Turnaround - T_R T_R/T_S	3 3 1.00	17 15 2.50	11 7 1.75	20 14 2.80	19 11 5.50	10.0 2.71
SPN	Finish Time Turnaround - T_R T_R/T_S	3 3 1.00	9 7 1.17	15 11 2.75	20 14 2.80	11 3 1.50	7.60 1.84
SRT	Finish Time Turnaround - T_R T_R/T_S	3 3 1.00	15 13 2.17	8 4 1.00	20 14 2.80	10 2 1.00	7.20 1.59
HRRN	Finish Time Turnaround - T_R T_R/T_S	3 3 1.00	9 7 1.17	13 9 2.25	20 14 2.80	15 7 3.5	8.00 2.14
Feedback ($q = 1$)	Finish Time Turnaround - T_R T_R/T_S	4 4 1.33	20 18 3.00	16 12 3.00	19 13 2.60	11 3 1.5	10.0 2.29
Feedback ($q = 2^i$)	Finish Time Turnaround - T_R T_R/T_S	4 4 1.33	17 15 2.50	18 14 3.50	20 14 2.80	14 6 3.00	10.6 2.63

