

Process Scheduling

Notes prepared by Matt Bishop

Goal

What characterizes a “fair internal policy?” Which process is given the CPU next? This is the province of *schedulers*.

Schedulers

Three kinds:

- *long-term scheduler* determines which jobs are admitted to the system for processing
example: in a batch system, often more jobs are submitted than can be done at once, so some are spooled out to a mass storage device; the long-term scheduler selects the next one to be loaded into memory. So it controls the degree of multiprogramming, *i.e.*, the number of processes in memory.
- *short-term scheduler* determines which job in memory (*i.e.*, in the ready queue) goes next
- *medium-term scheduler:* at times jobs may have to be removed from the system temporarily; that is, too many jobs may be competing for memory. The removed process will be restarted where it left off later; called *swapping*. This scheduler decides who gets swapped out and in.

The long term scheduler is invoked relatively infrequently, but the short term one is invoked often — whenever *any* process returns control to the operating system. Hence the short-term scheduler must be very fast. (Context switching also must be very fast; typically, 10 μ s to 100 μ s. Many machines have special-purpose instructions, like the VAX LDCTX, for just this reason.)

The system should try to balance CPU-bound and I/O-bound jobs.

Scheduling Considerations and Overview

These choose which process goes next. Which one is used depends on what is wanted from the system; possible measures are:

- *throughput*; get the most work done in a given time
- *turnaround*; complete jobs as soon as possible after submission
- *response*; minimize the amount of time from submission to the first response (called the *response time*); this interval does *not* include the time to output the response
- *resource use*; keep each type of resource assigned to some process as much as possible, but avoid waiting too long for certain resources.
- *waiting time*; minimize the amount of time the process sits in the ready queue
- *consistency*; treat processes with given characteristics in a predictable manner that doesn't vary greatly over time.

In the process of scheduling, the processes being considered must be distinguished upon many parameters, among them

priority

anticipated resource need (including running time)

running time, resources used so far

interactive/non-interactive

frequency of I/O requests

time spent waiting for service

To demonstrate how algorithms work, we'll use this set of jobs:

	Arrival Time	Service Time
A	0	10
B	1	29
C	2	3
D	3	7
E	4	12

and measure 3 quantities:

- turnaround time: time the process is present in the system

$$T = \text{finish time} - \text{arrival time}$$

- waiting time: time the process is present and not running

$$W = T - \text{service time}$$

- response ratio (sometimes called the "penalty ratio"): the factor by which the processing rate is reduced, from the user's point of view:

$$R = \frac{T}{\text{service time}}$$

Characterization of Scheduling Algorithms

decision mode

This is *non-preemptive* if a process runs until it blocks or completes; at no time during its run will the operating system replace it with another job. It is *preemptive* if the operating system can interrupt the currently running process to start another one.

priority function

This is a mathematical function which assigns a priority to the process; the process with the highest (numerical) priority goes next. The function usually involves the service time so far a , the real time spent in the system so far r , and the total required service time t .

arbitration rule

If two processes have the same priority, this rule states how one of them is selected to run.

The Scheduling Algorithms

First Come, First Served (FCFS)

decision mode: non-preemptive

priority function: $p(a, r, t) = r$

arbitration rule: random

	service time	arrival time	start	finish	T	W	R
A	10	0	0	10	10	0	1.0
B	29	1	10	39	38	9	1.3
C	3	2	39	42	40	37	13.3
D	7	3	42	49	46	39	6.6
E	12	4	49	61	57	45	4.8
<i>mean</i>					38.2	26	5.4

A potential problem is when a short job follows a long one:

	service time	arrival time	start	finish	T	W	R
A'	1000	0	0	1000	1000	0	1.0
B'	1	1	1000	1001	1000	999	1000.0

Gantt Chart:



Basically, long processes love FCFS, but short ones seem to be much slower.

Shortest Job Next (SJN), Shortest Job First (SJF), Shortest Process Next (SPN)

As an estimate of the total service time needed is required, this algorithm is usually used in batch systems.

decision mode: non-preemptive
 priority function: $p(a, r, t) = -t$
 arbitration rule: chronological or random

	service time	arrival time	start	finish	T	W	R
A	10	0	0	10	10	0	1.0
B	29	1	32	61	60	31	2.1
C	3	2	10	13	11	8	3.7
D	7	3	13	20	17	39	2.4
E	12	4	20	32	28	10	2.3
<i>mean</i>					25.2	17.6	2.3

Claim: Shortest Job First gives the smallest average turnaround time T out of all non-preemptive priority functions.

Proof: Suppose n jobs arrive at the same time, with $t_1 \leq t_2 \leq \dots \leq t_n$. Then $T(t_1) = t_1$, $T(t_2) = t_1 + t_2$, ..., hence the average turnaround time is

$$T_{av} = \sum_i i t_i$$

Now suppose t_a and t_b , $a < b$, are swapped. The new average turnaround time is:

$$T'_{av} = \frac{1}{n} (n t_1 + (n-1) t_2 + \dots + (n-a+1) t_b + \dots + (n-b+1) t_a + \dots + t_n)$$

so

$$T'_{av} - T_{av} = \frac{1}{n} ((n-a+1) t_b - (n-b+1) t_a + (n-a+1) t_a - (n-b+1) t_b)$$

and

$$T'_{av} - T_{av} = \frac{1}{n} (b-a)(t_b - t_a) \geq 0 \text{ because } b \geq a \text{ implies } t_b \geq t_a.$$

Problem: need to know service times into the future so you can run the process with the shortest next CPU burst. How does the short-term scheduler choose the next process to run? It can use a number of different ways:

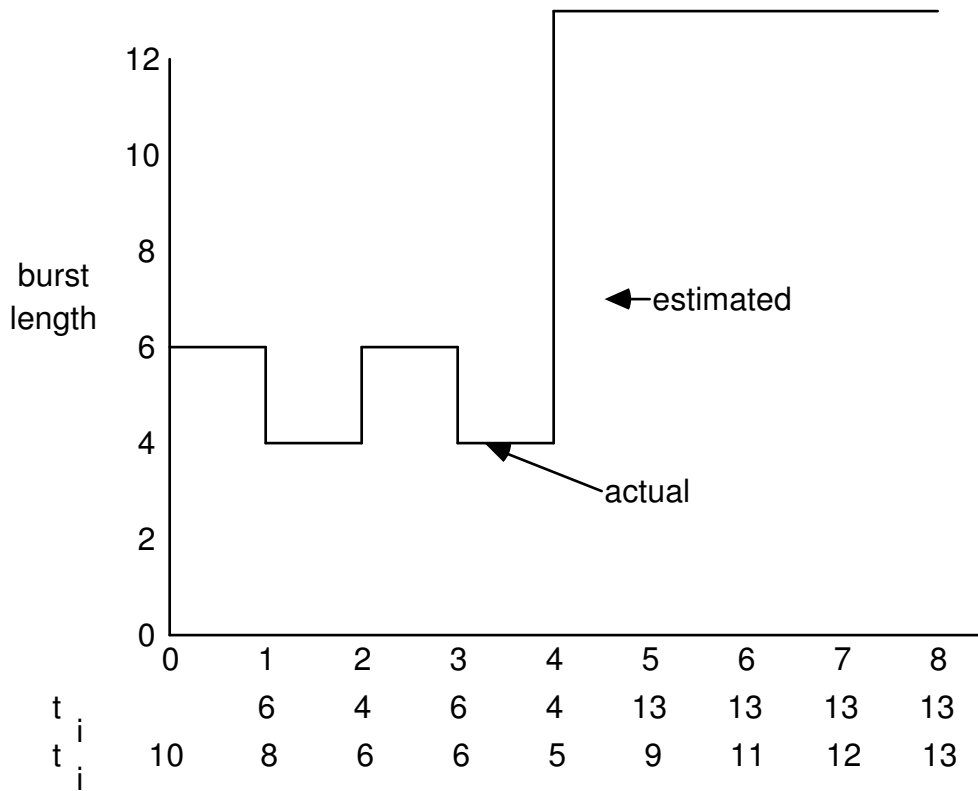
- Most accurate is to run all ready processes, time the CPU bursts, and then schedule them (*snicker*)
- Characterize each process as CPU-bound or I/O-bound, and specify for each an "average service time needed" based upon timing processes over a period of time and averaging. Note that characteristics might change over a period of time; that is, a process might be CPU-bound for a time, then I/O-bound, then CPU-bound, etc.
- Compute the expected time of the next CPU-burst as an exponential average of previous CPU-bursts of the process. Let t_n be the length of the n-th CPU burst, and t_{n+1} the expected length of the next burst; then

$$t_{n+1} = a t_n + (1-a) t_n$$

where a is a parameter indicating how much to count past history (usually chosen around $\frac{1}{2}$)

$a = 1$ the estimate is simply the length of the last CPU burst

$a = 0$ the estimate is the initial estimate holds



Comparing exponential estimation with actual values: $a = \frac{1}{2}$

SPN is better than FCFS for short jobs, but long jobs may have to wait for some time for service.

The long-term scheduler can simply use the job's time limit as specified by the user; this motivates users to be realistic in their limits, as:

- limits too low: job aborts with a "time limit exceeded".
- limits too high: the turnaround time may be very long.

Shortest Remaining Time (SRT), Preemptive Shortest Process Next (PSPN)

This is like SPN, but preemptive.

decision mode: preemptive (at arrival)
 priority function: $p(a, r, t) = a-t$
 arbitration rule: chronological or random

	service time	arrival time	start	finish	T	W	R
A	10	0	0, 12	2, 20	20	10	2.0
B	29	1	32	61	60	31	2.1
C	3	2	2	5	3	0	1.0
D	7	3	5	12	9	2	1.3
E	12	4	20	32	28	16	2.3
<i>mean</i>					24	11.8	1.74

Miscellaneous:

- Whenever a new job comes in, check the remaining service time on the current job.
- For all but the longest jobs, SRT better than SJF
- The response ratio is good (low)
- Waiting time is also quite low for most processes.

Highest Response Ratio Next (HRRN, HRN)

This tries to level out bias towards long or short jobs

decision mode: non-preemptive
 priority function: $p(a, r, t) = a/c$
 arbitration rule: random or FIFO

	service time	arrival time	start	finish	T	W	R
A	10	0	0	10	20	10	2.0
B	29	1	32	61	60	31	2.1
C	3	2	2	5	3	0	1.0
D	7	3	5	12	9	2	1.3
E	12	4	20	32	28	16	2.3
<i>mean</i>					25.2	13	2.3

Why? Here are the response ratios as each process completes:

time	A	B	C	D	E
10		$\frac{29+9}{29} = 1.3$	$\frac{3+8}{3} = 3.7$	$\frac{7+7}{7} = 2.0$	$\frac{12+6}{12} = 1.5$
13		$\frac{29+12}{29} = 1.4$		$\frac{7+10}{7} = 2.4$	$\frac{12+9}{12} = 1.8$
20		$\frac{29+19}{29} = 1.7$			$\frac{12+16}{12} = 2.3$
32		$\frac{29+31}{29} = 2.1$			

The ratio used is actually

$$\frac{\text{estimated service time} + \text{waiting time so far}}{\text{estimated service time}}$$

The idea behind this method is to get the mean response ratio low, so if a job has a high response ratio, it should be run at once to reduce the mean.

Round Robin (RR) with Quantum q

This is especially designed for time sharing; the quantum is typically $\frac{1}{60} \leq q \leq 1$ seconds.

decision mode: preemptive (at quantum)
 priority function: $p(a, r, t) = c$
 arbitration rule: cyclic

In this example the quantum is 5:

	service time	arrival time	start	finish	T	W	R
A	10	0	...	28	28	18	2.8
B	29	1	...	61	60	31	2.1
C	3	2	...	13	11	8	3.7
D	7	3	...	35	32	25	4.6
E	12	4	...	47	43	31	3.5
<i>mean</i>					34.8	22.6	3.3

Why? Here is what things look like:

time	0	5	10	13	18	23	28	33	35	40	45	47	52	57	61
proc.	A	B	C	D	E	A	B	D	E	B	E	B	B	B	
rem	5	24	0	2	7	0	19	0	2	14	0	9	4	0	

(here, "proc" is the process starting at the indicated time, and "rem" the remaining time after the quantum is complete.)

- As each process is preempted, it moves to the rear of the queue
- All new arrivals come in at the rear of the queue
- As $q \rightarrow 0$, every process thinks it is getting constant service from a processor that is slower in proportion to the number of competing processes; this is called *processor sharing*. This scheme is used in hardware in CDC6600 to implement 10 peripheral processors with one set of hardware (*i.e.*, processor) and 10 sets of registers; the processor does 1 instruction for one set of registers, then goes on to the next set. (This turns out to be not much slower than a real processor.)

Variants:

- Round Robin, but adjust quantum periodically.
example: after every process switch, the quantum becomes q/n , where n is the number of processes in the ready list
 - few ready processes means that each gets a long quantum, minimizing process switches.
 - a lot of ready processes means that this algorithm gives more processes a shot at the CPU over a fixed period of time, at the price of more process switching
 - processes needing a small amount of CPU time get a quantum fairly soon, and hence *may* finish sooner.
- Round Robin, but give the current process an extra quantum when a new process arrives
 This reduces process switching in proportion to the number of processes arriving.

Selfish Round Robin (SRR) with Parameters a and b and Quantum q

This is like round robin, except newly-arrived jobs move into the *arrived queue* (or *new queue*), and assigned an initial priority of 0. Associated with the *accepted* queue is a single priority. The priority of all processes in the new queue increases at rate *a*; that of all processes in the accepted queue increases at rate *b*. When the priority of a process in the new queue is equal to the priority of the accepted queue, the process moves to the accepted queue.

decision mode: preemptive (at quantum)
 priority function: here, *W* is the time that a process must wait before entering the accepted queue.

$$p(r, W) = \begin{cases} br & r \leq W \\ bW + a(r - W) & r > W \end{cases}$$

arbitration rule: first in, first out

In this example the quantum is 1, and *a* = 3, *b* = 2:

	service time	arrival time	start	finish	T	W	R
A	10	0	27	17	2.7
B	29	1	60	31	2.1
C	3	2	15	12	5.0
D	7	3	33	26	4.7
E	12	4	44	32	3.7
<i>mean</i>					35.8	23.6	3.6

Why? Assuming that new process promotions from the new queue to the ready queue precede quanta expirations, we have:

time	job running	ready queue (at end of interval)	new queue (at end of interval)
0-1	A	A(2)	B(0)
1-2	A	A(4)	B(3), C(0)
2-3	A	B(6), A(6)	C(3), D(0)
3-4	B	A(8), B(8)	C(6), D(3), E(0)
4-5	A	B(10), A(10)	C(9), D(6), E(3)
5-6	B	A(12), C(12), B(12)	D(9), E(6)
6-7	A	C(14), B(14), A(14)	D(12), E(9)
7-8	C	B(16), A(16), C(16)	D(15), E(12)
8-9	B	A(18), C(18), D(18), B(18)	E(15)
9-10	A	B(20), C(20), D(20), A(20)	E(18)
10-11	B	C(22), D(22), A(22), B(22)	E(21)
11-12	C	D(24), A(24), B(24), E(24), C(24)	

... round robin from here on ...

How *a* and *b* are chosen greatly affects the way the algorithm works:

- *a* = *b* £ *p*(*r*,*W*) = *ar*, which is the same as the FCFS algorithm.
- *b* = 0 £ *p*(*r*,*W*) = 0, which is the same as Round Robin.
- *b* ≥ *a* £ accept no new processes until all accepted processes complete

This algorithm gives better service to processes executing for a while as opposed to new ones.

Multilevel Feedback Queues (MLF, MLFB) with n different priority levels each of priority T_p

Processes start out in the uppermost level. After getting T_0 units of CPU time, it drops to the next lower level, and after units of CPU time at that level, it drops down again ..., until it reaches the lowest level. If it blocks or otherwise leaves the scheduling system, and later returns, it may reenter the feedback queues at another queue (for example, the top one).

- decision mode: preemptive (at quantum)
- priority function: $p(a) = n - i$, where i satisfies both $0 \leq i < n$ and $T_0(2^i - 1) \leq a < T_0(2^{i+1} - 1)$, assuming that $T_p = 2^p T_0$
- arbitration rule: cyclic or chronological within queues

In this example the quantum is 1, $n = 3$, $T_0 = 2$, and $T_p = 2^p T_0$:

	service time	arrival time	start	finish	T	W	R
A	10	0	38	28	3.8
B	29	1	60	31	2.1
C	3	2	11	8	3.7
D	7	3	27	20	3.9
E	12	4	40	28	3.3
<i>mean</i>					35.2	23.3	3.4

This algorithm favors short processes by giving them more of the CPU.

It is also adaptive, in that it responds to the changing behavior of the system it controls.

Variants

- MLFB with round robin for all but the lowest level, and thatr first come first serve (but preemption possible, of course):

	service time	arrival time	start	finish	T	W	R
A	10	0	25	15	2.5
B	29	1	49	20	2.5
C	3	2	11	8	1.4
D	7	3	50	43	1.2
E	12	4	57	45	1.3
<i>mean</i>					38.4	26.2	1.8

Comparison of Scheduling Methods

These policies were simulated assuming an exponential distribution of arrival and service rates (with parameters of $a = 0.8$, $b = 1.0$), and using 50,000 processes (the first 100 of which were ignored, to get to a steady state). The processes were categorized into percentiles based on service time required; each percentile had about 500 processes.

External Priority Methods

These adjust priority based on some external factors, and are quite common when users pay based upon their computer use.

Examples:

- round robin, where the quantum is set independently for each process, based on the external priority of process (*i.e.*, the more you pay, the bigger the quantum.)
- *Worst Service Next*: after each quantum, compute a “suffering function” (based on how long the process had to wait, how many times it has been preempted, how much the user is paying, and/or the amount of time and resources used). The process with the greatest suffering gets the next quantum.
- The user buys a response ratio guarantee; the suffering function used takes into account the difference between the guaranteed response ratio and the actual response ratio at the moment.
- *Deadline Scheduling*: each process specifies how much service it needs and by what real time it must be finished. The algorithm tries not to run jobs that cannot meet their deadline.
- *Fair-Share Scheduling*: allocate blocks of CPU time to a particular set of processes, usually by splitting user processes into groups; within each group, use a standard schedule, but allocate the CPU proportionately to each group
example: All processes are infinite loops; 1 process in group 1, 2 in group 2, 3 in group 3, and 4 in group 4
regular scheduler: each process gets 10%

fair share scheduler: each group gets 25%; processes in sgroup share equally

example: This uses UNIX internal, not external, priorities. Here, 3 processes: process A in one group; processes B and C in another group.

The internal priority function is:

$$\text{priority} = \frac{\text{recent CPU usage}}{2} + \frac{\text{group CPU usage}}{2} + \text{threshold}$$

(with the threshold being 60 for user processes). A decay function decrements the current CPU usage of processes not run; this has the effect of raising their priority. The function is:

$$\text{decay of CPU usage} = \frac{\text{CPU usage}}{2}$$

example of the UNIX Fair Share Scheduler. Here, the quantum is 1 second. Note that the higher the priority, the lower the integer representing that priority.

A runs for 1 second

decay applied to CPU and group CPU usage; A's new priority is $60 + \frac{30}{2} +$

$\frac{30}{2} = 90$. As B and C now have higher priority, one of them (say, B) goes next.

B runs for 1 second

decay applied to CPU and group CPU usage; A's new priority is $60 + \frac{15}{2} + \frac{15}{2} = 74$, B's new priority is $60 + \frac{30}{2} + \frac{30}{2} = 90$, and C's new priority is $60 + \frac{0}{2} + \frac{30}{2} = 75$; A has the highest priority, so it runs next.

A runs for 1 second.

decay applied to CPU and group CPU usage; A's new priority is $60 + \frac{37}{2} + \frac{37}{2} = 96$, B's new priority is $60 + \frac{15}{2} + \frac{15}{2} = 74$, and C's new priority is $60 + \frac{0}{2} + \frac{15}{2} = 67$; C has the highest priority, so it runs next.

C runs for 1 second.

decay applied to CPU and group CPU usage; A's new priority is $60 + \frac{18}{2} + \frac{18}{2} = 96$, B's new priority is $60 + \frac{7}{2} + \frac{37}{2} = 81$, and C's new priority is $60 + \frac{30}{2} + \frac{37}{2} = 93$; A has the highest priority, so it runs next.

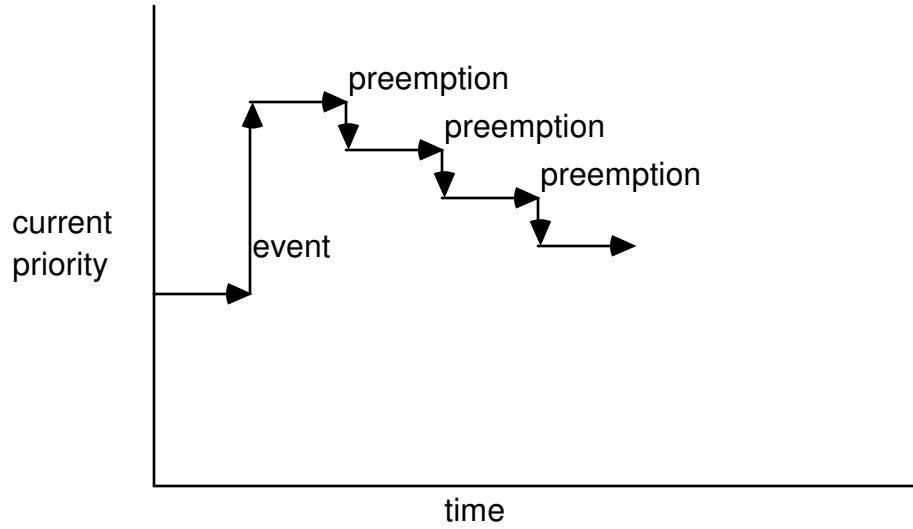
Hence the order of running is A B A C A B A C ..., with A getting 50% of the CPU and B and C together getting 50%.

example: VAX/VMS scheduler

This scheduler has 32 priority levels: levels 31 to 16 are for real-time processes, and levels 15 to 0 for regular processes. Real-time processes have fixed priority throughout their lifetime; but the priority of regular processes is dynamic:

- at process creation, a *base* priority assigned; this is the process' minimum priority
- the current priority of the process is altered by a system events, each of which has an associated *increment*, *i.e.*, terminal read increment > terminal write increment > disk I/O

When awakened due to a system event, the appropriate increment is added to the current priority value; on preemption due to quantum expiration, the current priority drops by 1.



- Processes are dispatched by their current priority.
This scheme is like a MLFB scheme, with two differences:
- processes need not start at the highest level; and
 - quanta are associated with each *process*, not level

Evaluation

How well does the algorithm meet the desired goal(s)? 4 ways to do this:

- (1) *deterministic modelling* - we did this when looking at the algorithms when we assumed a given workload.
 - It gives exact answers for the given workload;
 - Those answers are only applicable to the specific case, and many cases must be analyzed to discern trends; *eg.*, that SJF has the minimum average turnaround time.
- (2) *simulation* - program a model of the system, using a variable representing a clock, and as the value of that variable increases, the simulator modifies the state of the system to reflect the activities of devices, jobs, schedulers; it gathers statistics and prints them to indicate the algorithm's performance. To get the data for the simulation:
 - Use a RNG to generate it using a distribution based upon the distribution of jobs in the system
 - Monitor the real system, recording sequence of actual events, and use the resulting observations.

Using this method, you can compare algorithms exactly for the given site data.

Advantages: can be tailored for site's actual job mix

Problems: It can be expensive to do in time and/or space (storing all the observations).

- (3) *implementation* - implement the scheduling algorithm and put it into the system. Then watch the results.

Problems: It can be quite expensive to modify the system and mollify the users; also, users can quickly learn how to work around the scheduling algorithm, *eg.*, if short jobs get higher priority, they will split longer jobs up.

Advantage: it is completely accurate.
- (4) *queueing theory* – represent the system mathematically. Usually, you assume that processes arrival times are exponentially distributed. Let a be

the arrival rate, so the average time between arrivals is $\frac{1}{a}$; this means:

- when a process arrives, the next one will arrive within time t with probability $1 - e^{-at}$; in other words,
- the probability that k processes will arrive within 1 time unit is $e^{-a} \frac{a^k}{k!}$.

Why assume an exponential distribution? Two reasons:

- It's easy to deal with mathematically.
- It reflects the way processes arrive in practice.

We can also assume the processes' service times are exponentially distributed with parameter b , so:

- the probability that k processes will be serviced within 1 time unit is $e^{-b} \frac{b^k}{k!}$.

The *saturation level* $r = \frac{a}{b}$ indicates how busy the computer is on the average:

$r = 0$ $\&$ new processes never arrive, so the machine is idle

$r = 1$ $\&$ processes arrive on the average at exactly the same rate as they can be finished

$r > 1$ $\&$ processes come in faster than they can be served; hence the ready list is expected to get longer and longer.

example (Little's Law): Let:

L be the mean queue length (excluding the job currently being serviced);

W be the mean waiting time in the queue;

a be the mean arrival rate for new jobs into the queue (*ie*, the units might be jobs/second)

Hence you expect that during time **W**, **aW** new jobs arrive. Now assume the system is in a "steady state;" then the number of jobs leaving the queue is the same as the number of jobs arriving, so

$$\mathbf{L} = \mathbf{aW}$$

This is Little's Law, and it applies to any arrival distribution and scheduling algorithm. For example, say jobs arrive at the mean rate of 7 jobs per second, and there are normally 14 jobs in the queue; hence

$$\mathbf{a} = 7 \frac{\text{jobs}}{\text{sec}}, \quad \mathbf{L} = 14 \text{ jobs} \quad \& \quad \mathbf{W} = \frac{\mathbf{L}}{\mathbf{a}} = \frac{14}{7} \text{ sec} = 2 \text{ seconds.}$$

So the mean waiting time in the queue is 2 seconds per job.

Types of Schedulers Chart

Introduction

This chart shows the function of each of the three types of schedulers (long-term, short-term, and medium-term) for each of three types of operating systems (batch, interactive, and real-time).

Chart

	batch	interactive	real-time
long-term	job admission, based on characteristics and resource needs	sessions and processes normally accepted unless capacity reached	processes either permanent or accepted at once
medium-term	usually none - jobs remain in storage until done	processes swapped when necessary	processes never swapped
short-term	processes scheduled by priority; continue until wait voluntarily, request service, are terminated	processes scheduled on rotating basis; continue until service requested, time quantum expires, preempted	scheduling based on strict priority with immediate preemption; may time-share among equal priorities

From

Malcolm G. Lane and James D. Mooney, *A Practical Approach to Operating Systems*, Boyd and Fraser Publishing Company, Boston, MA ©1988, page 168.

Job Scheduling Algorithms

Introduction

This handout shows how the various job scheduling algorithms work.

First Come, First Serve (FCFS)

This policy services jobs in the order they arrive.

job name	arrival time	service time	start time	finish time	turnaround time	waiting time	response ratio
A	0	10	0	10	10	0	1.0
B	1	29	10	39	38	9	1.3
C	2	3	39	42	40	37	13.3
D	3	7	42	49	46	39	6.6
E	4	12	49	61	57	45	4.8
<i>mean</i>					38	26	5.4

Shortest Job Next (SJN)

This policy services the shortest job next.

job name	arrival time	service time	start time	finish time	turnaround time	waiting time	response ratio
A	0	10	0	10	10	0	1.0
B	1	29	32	61	60	31	2.1
C	2	3	10	13	11	8	3.7
D	3	7	13	20	17	10	2.4
E	4	12	20	32	28	16	2.3
<i>mean</i>					25	13	2.3

Pre-emptive Shortest Job Next (PSJN)

This policy services the shortest job next, pre-emptively.

job name	arrival time	service time	start time	finish time	turnaround time	waiting time	response ratio
A	0	10	0	2			
		8	12	20	20	10	2.0
B	1	29	32	61	60	31	2.1
C	2	3	2	5	3	0	1.0
D	3	7	5	12	9	2	1.3
E	4	12	20	32	28	16	2.3
<i>mean</i>					24	12	1.7

pre-empted by C

Highest Response Ratio Next (HRN)

This policy services the job with the highest response ratio next.

job name	arrival time	service time	start time	finish time	turnaround time	waiting time	response ratio
A	0	10	0	10	10	0	1.0
B	1	29	32	61	60	31	2.1
C	2	3	10	13	11	8	3.7
D	3	7	13	20	17	10	2.4
E	4	12	20	32	28	16	2.3
mean					25	13	2.3

Round Robin (RR)

This policy services jobs for a fixed quantum (here, 5).

job name	arrival time	service time	start time	finish time	turnaround time	waiting time	response ratio
A	0	10	0	5	<i>end of quantum; B starts</i>		
		5	23	28	28	18	2.8
B	1	29	5	10	<i>end of quantum; C starts</i>		
		24	28	33	<i>end of quantum; D starts</i>		
		19	40	45	<i>end of quantum; E starts</i>		
		14	47	61	60	31	2.1
C	2	3	10	13	11	8	3.7
D	3	7	13	18	<i>end of quantum; E starts</i>		
		2	33	35	32	25	4.6
E	4	12	18	23	<i>end of quantum; A starts</i>		
		7	35	40	<i>end of quantum; B starts</i>		
		2	45	47	43	31	3.5
mean					35	23	3.3

Selfish Round Robin (SRR)

This policy services jobs for a fixed quantum (here, 1). The priority of new processes increases at a rate of 3 per quantum, of accepted processes at a rate of 2 per quantum. Note that new process promotions from the new queue to the ready queue precede quanta expirations.

time	job running	ready queue (at end of interval)	new queue (at end of interval)
0-1	A	A(2)	B(0)
1-2	A	A(4)	B(3), C(0)
2-3	A	B(6), A(6)	C(3), D(0)
3-4	B	A(8), B(8)	C(6), D(3), E(0)
4-5	A	B(10), A(10)	C(9), D(6), E(3)
5-6	B	A(12), C(12), B(12)	D(9), E(6)
6-7	A	C(14), B(14), A(14)	D(12), E(9)
7-8	C	B(16), A(16), C(16)	D(15), E(12)
8-9	B	A(18), C(18), D(18), B(18)	E(15)
9-10	A	C(20), D(20), B(20), A(20)	E(18)
10-11	C	D(22), B(22), A(22), C(22)	E(21)
11-12	D	B(24), A(24), C(24), E(24), D(24)	

... round robin from here on ...

The relevant numbers (ignoring start and finish time) are:

job name	arrival time	service time	start time	finish time	turnaround time	waiting time	response ratio
A	0	10	27	17	2.7
B	1	29	60	31	2.1
C	2	3	15	12	5.0
D	3	7	33	26	4.7
E	4	12	44	32	3.7
<i>mean</i>					35.8	23.6	3.6

Multilevel Feedback (MLFB)

The variant of this class of scheduling algorithms uses three levels:

- processes at level 1 are scheduled round robin; the relevant quantum is 2, and when a quantum expires the job is moved to level 2.
- processes at level 2 are scheduled round robin; the quantum is 4, and processes are allowed 2 quanta before being moved to level 3.
- processes at level 3 are serviced first come first serve.

The jobs A, B, C, D, and E have been augmented by F, a 1-second job arriving at time 13, and G, an 11-second job arriving at time 50. These are to demonstrate that quanta are usually **not** interrupted.

In what follows, the number in parentheses in the comment field is the remaining service time for the job.

time	level 1	level 2	level 3	comments
0	A	—	—	A(10) arrives, runs
1	AB	—	—	B(29) arrives, A continues quantum
2	BC	A	—	C(3) arrives, A's quantum expires (8), moves to level 2, B runs
3	BCD	A	—	D(7) arrives, B continues quantum
4	CDE	AB	—	E(12) arrives, B's quantum expires (27), moves down, C runs
6	DE	ABC	—	C's quantum expires (1), moves down, D runs
8	E	ABCD	—	D's quantum expires (5), moves down, E runs
10	—	ABCDE	—	E's quantum expires (10), moves down, A runs from level 2 (level 1 is empty)
13	F	ABCDE	—	F(1) arrives, A's quantum continues
14	F	ABCDE	—	A's quantum expires (4), F runs (at level 1)
15	—	ABCDE	—	F finishes, B runs from level 2 (level 1 is empty)
19	—	ABCDE	—	B's quantum expires (23), C runs
20	—	ABDE	—	C finishes, D runs
24	—	ABDE	—	D's quantum expires (1), E runs
28	—	ABDE	—	E's quantum expires (6), A runs
32	—	BDE	—	A finishes, B runs
36	—	DE	B	B's quantum expires (19), moves down, D runs
37	—	E	B	D finishes, E runs
41	—	—	BE	E's quantum expires (2), moves down, B runs from level 3 (since there is nothing in higher levels)
50	G	—	BE	G arrives(11), B continues to run
60	G	—	E	B finishes, G runs (since it is in the highest level)
62	—	G	E	G's quantum expires (9), moves down, G runs from level 2
66	—	G	E	G's quantum expires (5), G runs
70	—	—	EG	G's quantum expires (1), moves down, E runs
72	—	—	G	E finishes, G runs
73	—	—	—	G finishes

The relevant numbers (ignoring start and finish time) are:

job name	arrival time	service time	start time	finish time	turnaround time	waiting time	response ratio
A	0	10	0	2	<i>preempted by B</i>		3.2
		8	10	14	<i>preempted by F</i>		
		4	28	32	32	22	
B	1	29	2	4	<i>preempted by C</i>		2.0
		27	15	19	<i>preempted by C</i>		
		23	32	36	<i>preempted by D</i>		
C	2	19	41	60	59	30	6.0
		3	4	6	<i>preempted by D</i>		
D	3	1	19	20	18	15	4.9
		7	6	8	<i>preempted by E</i>		
E	4	5	20	24	<i>preempted by E</i>		5.7
		1	36	37	34	27	
		12	8	10	<i>preempted by A</i>		
F	13	10	24	28	<i>preempted by A</i>		2.0
		6	37	41	<i>preempted by B</i>		
		2	70	72	68	56	
G	50	11	60	70	2	1	2.1
		1	72	73	23	12	3.7
<i>mean</i>					33.7	23.3	