

Types of Schedulers

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

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

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.

<i>job name</i>	<i>arrival time</i>	<i>service time</i>	<i>start time</i>	<i>finish time</i>	<i>turnaround time</i>	<i>waiting time</i>	<i>response ratio</i>
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.

<i>job name</i>	<i>arrival time</i>	<i>service time</i>	<i>start time</i>	<i>finish time</i>	<i>turnaround time</i>	<i>waiting time</i>	<i>response ratio</i>
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.

<i>job name</i>	<i>arrival time</i>	<i>service time</i>	<i>start time</i>	<i>finish time</i>	<i>turnaround time</i>	<i>waiting time</i>	<i>response ratio</i>
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

Highest Response Ratio Next (HRN)

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

<i>job name</i>	<i>arrival time</i>	<i>service time</i>	<i>start time</i>	<i>finish time</i>	<i>turnaround time</i>	<i>waiting time</i>	<i>response ratio</i>
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

Round Robin (RR)

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

<i>job name</i>	<i>arrival time</i>	<i>service time</i>	<i>start time</i>	<i>finish time</i>	<i>turnaround time</i>	<i>waiting time</i>	<i>response ratio</i>
A	0	10	0	5		<i>end of quantum; B starts</i>	
B	1	29	5	10	28	18	2.8
		24	28	33		<i>end of quantum; C starts</i>	
		19	40	45		<i>end of quantum; E starts</i>	
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
<i>mean</i>					35	23	3.3

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.

<i>time</i>	<i>level 1</i>	<i>level 2</i>	<i>level 3</i>	<i>comments</i>
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)

<i>time</i>	<i>level 1</i>	<i>level 2</i>	<i>level 3</i>	<i>comments</i>
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:

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

Fair Share Scheduler

Introduction

A *fair share scheduler* is used when CPU time is to be divided equally between groups of processes. For this scheduling algorithm, processes are divided into groups based upon external factors. Such factors include the organizational divisions of the owners of the computer, or classes of customers, or other criteria.

For example, suppose group A has 1 process, group B has 2 processes, group C has 3 processes, and group D has 4 processes. Under a regular scheduler, each of the 10 processes would get 10% of the CPU. Under a fair share scheduler, each of the 4 groups would get 25% of the CPU.

Example

Suppose there are 3 processes. Process p_1 is in group A, and processes p_2 and p_3 are in group B. The following formula assigns process p_i a priority P_i :

$$P_i = (p_i\text{'s recent CPU usage})/2 + (p_i\text{'s group CPU usage})/2$$

In addition, a decay function decrements the current CPU usage of all processes. This “spreads out” the priority of the processes in the ready queue. The decay D_i for p_i is:

$$D_i = (p_i\text{'s recent CPU usage})/2$$

In this system, the lower the numerical value of P_i , the higher the priority of process p_i .

The following shows how processes execute, given a quantum of 60 ticks. All arithmetic is *integer* arithmetic, and the decay function is applied after the most recent CPU time is added in, but *before* the priorities are computed.

First 60-Tick Interval

At the beginning of this interval, all priorities are equal, so the process to run is chosen randomly. Say p_1 is selected to run. It runs, and at the end of the interval, its CPU usage is updated to 60. The group CPU usage for group A, to which p_1 belongs, also is updated to 60. The decay function is then applied, cutting both to 30. The CPU usage for p_2 and p_3 , and for group B, are 0, so the decay function does not change them. The priority P_1 of p_1 becomes

$$P_1 = (p_1\text{'s recent CPU usage})/2 + (p_1\text{'s group CPU usage})/2 = 30/2 + 30/2 = 15 + 15 = 30$$

Second 60-Tick Interval

At the beginning of this interval, P_2 and P_3 are equal, and both are less than P_1 , so either p_2 or p_3 will run. Say p_2 is selected to run. It runs, and at the end of the interval, its CPU usage is updated to 60. The group CPU usage for group B, to which p_2 belongs, also is updated to 60. The decay function is then applied, cutting both to 30. It also cuts the CPU usage of p_1 to 15, and the group CPU usage of group A to 15. The CPU usage for p_3 is 0, so the decay function does not change it. The priorities become

$$P_1 = (p_1\text{'s recent CPU usage})/2 + (p_1\text{'s group CPU usage})/2 = 15/2 + 15/2 = 7 + 7 = 14$$

$$P_2 = (p_2\text{'s recent CPU usage})/2 + (p_2\text{'s group CPU usage})/2 = 30/2 + 30/2 = 15 + 15 = 30$$

$$P_3 = (p_3\text{'s recent CPU usage})/2 + (p_3\text{'s group CPU usage})/2 = 0/2 + 30/2 = 0 + 15 = 15$$

Third 60-Tick Interval

At the beginning of this interval, P_1 is less than P_2 or P_3 , so p_1 runs. At the end of the interval, its CPU usage is updated to $15 + 60 = 75$. The group CPU usage for group A, to which p_1 belongs, is similarly updated to $15 + 60 = 75$. The decay function is then applied, cutting both to 37. It also cuts the CPU usage of p_2 to 15, and the group CPU usage of group B to 15. The CPU usage for p_3 is 0, so the decay function does not change it. The priorities become

$$P_1 = (p_1\text{'s recent CPU usage})/2 + (p_1\text{'s group CPU usage})/2 = 37/2 + 37/2 = 18 + 18 = 36$$

$$P_2 = (p_2\text{'s recent CPU usage})/2 + (p_2\text{'s group CPU usage})/2 = 15/2 + 15/2 = 7 + 7 = 14$$

$$P_3 = (p_3\text{'s recent CPU usage})/2 + (p_3\text{'s group CPU usage})/2 = 0/2 + 15/2 = 0 + 7 = 7$$

Fourth 60-Tick Interval

At the beginning of this interval, P_3 is less than P_1 or P_2 , so p_3 runs. At the end of the interval, its CPU usage is updated to $0 + 60 = 60$. The group CPU usage for group B, to which p_2 belongs, is similarly updated to $15 + 60 = 75$. The decay function is then applied, cutting p_3 's CPU usage to 30 and the group CPU usage to 37. It also cuts the CPU usage of p_1 to 18, the CPU usage of p_2 to 7, and the group CPU usage of group A to 18. The priorities become

$$P_1 = (p_1\text{'s recent CPU usage})/2 + (p_1\text{'s group CPU usage})/2 = 18/2 + 18/2 = 9 + 9 = 18$$

$$P_2 = (p_2\text{'s recent CPU usage})/2 + (p_2\text{'s group CPU usage})/2 = 7/2 + 37/2 = 3 + 18 = 21$$

$$P_3 = (p_3\text{'s recent CPU usage})/2 + (p_3\text{'s group CPU usage})/2 = 30/2 + 37/2 = 15 + 18 = 33$$

Summary Table

This table summarizes the first 8 seconds. The figures shown are for after the ticks, and after the calculations of priorities.

<i>ticks</i>	<i>Priorities</i>			<i>CPU Usage After Decay</i>			<i>Group CPU Usage After Decay</i>		<i>runs</i>
	p_1	p_2	p_3	p_1	p_2	p_3	<i>A</i>	<i>B</i>	
0	0	0	0	0	0	0	0	0	A
60	30	0	0	30	0	0	30	0	B
120	14	30	15	15	30	0	15	30	A
180	36	14	7	37	15	0	37	15	C
240	18	21	33	18	7	30	18	37	A
300	38	10	16	39	3	15	39	18	B
360	18	34	22	19	31	7	19	39	A
420	38	16	10	39	15	3	39	19	C
...