

Priority Scheduling and Memory Mapping

2025 Winter ECE353: Systems Software
Jon Eyolfson

Lecture 10
2.0.1

Let's Explore a Dynamic Priority Scheduling

This may also be called: Feedback Scheduling

We let the algorithm manage the priorities

- We use set time slices, and measure CPU usage

Increase the priority of processes that don't use their time slice

Decrease the priority of processes that use their full time slice

We Pick the Lowest Number as Highest Priority

Each process gets assigned a priority when started, P_n

Pick the lowest priority number to schedule, if it yields, pick the next lowest number

- Break ties with arrival order

- If a lower priority number becomes ready, switch to it

Record how much time each process executes for in this priority interval, C_n

- Timer interrupts still occur

At the end of the priority interval, update the priority of each process with:

$$P_n = \frac{P_n}{2} + C_n$$

- and reset the value of C_n back to 0

Assume All Processes Have an Initial Priority of 0


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



Assume All Processes Have an Initial Priority of 0


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



Assume All Processes Have an Initial Priority of 0


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



Assume All Processes Have an Initial Priority of 0


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



Assume All Processes Have an Initial Priority of 0


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



Assume All Processes Have an Initial Priority of 0


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



Assume All Processes Have an Initial Priority of 0


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



	$P_X = \frac{0}{2} + 1 = 1$
X	$P_Y = \frac{0}{2} + 1 = 1$
ready	$P_A = \frac{0}{2} + 8 = 8$
Y	$P_B = \frac{0}{2} + 0 = 0$
ready	

Assume All Processes Have an Initial Priority of 0


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



Now, Processes A and B have a Priority of 6 (others are 0)

Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



Now, Processes A and B have a Priority of 6 (others are 0)


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



Now, Processes A and B have a Priority of 6 (others are 0)


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



Now, Processes A and B have a Priority of 6 (others are 0)


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



Now, Processes A and B have a Priority of 6 (others are 0)


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



Now, Processes A and B have a Priority of 6 (others are 0)


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



Now, Processes A and B have a Priority of 6 (others are 0)


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



Now, Processes A and B have a Priority of 6 (others are 0)


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



$$P_X = \frac{0}{2} + 2 = 2$$

$$P_Y = \frac{0}{2} + 2 = 2$$

$$P_A = \frac{6}{2} + 6 = 9$$

$$P_B = \frac{6}{2} + 0 = 3$$

Now, Processes A and B have a Priority of 6 (others are 0)


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



$$P_X = \frac{0}{2} + 2 = 2$$

$$P_Y = \frac{0}{2} + 2 = 2$$

$$P_A = \frac{6}{2} + 6 = 9$$

$$P_B = \frac{6}{2} + 0 = 3$$

Now, Processes A and B have a Priority of 6 (others are 0)


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



$$P_X = \frac{0}{2} + 2 = 2$$

$$P_Y = \frac{0}{2} + 2 = 2$$

$$P_A = \frac{6}{2} + 6 = 9$$

$$P_B = \frac{6}{2} + 0 = 3$$

Now, Processes A and B have a Priority of 6 (others are 0)


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



$$P_X = \frac{0}{2} + 2 = 2$$

$$P_Y = \frac{0}{2} + 2 = 2$$

$$P_A = \frac{6}{2} + 6 = 9$$

$$P_B = \frac{6}{2} + 0 = 3$$

Now, Processes A and B have a Priority of 6 (others are 0)


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



$$P_X = \frac{0}{2} + 2 = 2$$

$$P_Y = \frac{0}{2} + 2 = 2$$

$$P_A = \frac{6}{2} + 6 = 9$$

$$P_B = \frac{6}{2} + 0 = 3$$

Now, Processes A and B have a Priority of 6 (others are 0)


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



$$P_X = \frac{0}{2} + 2 = 2$$

$$P_Y = \frac{0}{2} + 2 = 2$$

$$P_A = \frac{6}{2} + 6 = 9$$

$$P_B = \frac{6}{2} + 0 = 3$$

Now, Processes A and B have a Priority of 6 (others are 0)


Assume we have 4 processes ready to execute arriving in order: X, Y, A, B

A and B are CPU bound processes

X and Y are I/O bound processes that execute for 1 and block for 5

Timer interrupts occur every 1, and each time slice is 10, priority interval is 10

What is the scheduling? (each box is a timer interrupt)



$$P_X = \frac{0}{2} + 2 = 2$$

$$P_Y = \frac{0}{2} + 2 = 2$$

$$P_A = \frac{6}{2} + 6 = 9$$

$$P_B = \frac{6}{2} + 0 = 3$$

A 30B LLM on with Only 6.8 GB of RAM, how?

LLaMA is Meta's large language model, and this C++ implementation is efficient

It's so efficient people are discussing how this is possible!

<https://github.com/ggerganov/llama.cpp/discussions/638>

We Can Control Our Processes' Virtual Memory

Memory map, or `mmap` is used to map files to a processes' virtual address space

The pointer (virtual address) returned will allow you to access the file directly
There's no need for `read` and `write` system calls

The `mmap` API

`mmap` takes 6 arguments:

1. `void *addr`: suggested starting address (NULL means you don't care)
2. `size_t length`: number of bytes to map
3. `int prot`: protection flags (read/write/execute)
4. `int flags`: mapping flags (shared/private/anonymous)
anonymous means the mapping isn't backed by a file
5. `int fd`: file descriptor to map (ignored for anonymous)
6. `off_t offset`: offset to start the mapping (must be a multiple of page size)

Let's See a `mmap` Example

Check out `lectures/14-priority-scheduling-and-memory-mapping` in the materials repository

mmap Is Lazy

It just sets up the page tables, it doesn't actually read from the file

It would create an invalid PTE during the mmap call

The kernel uses the remaining bits of the PTE for bookkeeping

Where on the disk is this entry

The first access to the page would generate a page fault

The kernel would then read from disk into memory

This ensures only the used parts of the file get read

Back to the Question

<https://github.com/ggerganov/llama.cpp/discussions/638>

How does an approximately 20 GB file only use 6.8 GB of real memory?

Hint: when you do model inference, the models are sparse (you don't use all of it)

How Much Space Would the Kernel Need for Page Tables?

Someone posted you'd need 40 MB of page tables:

$(20 * (1024 * 1024 * 1024) / 4096 * 8) / (1024 * 1024)$

Someone clarified it's:

$(20GB / 4KB \text{ Page size} * 8 \text{ bytes per PTE}) / 1KB$

(the 1KB at the end should be 1MB)

Is this correct? Why or why not?

How Much Space Do Our Page Tables Need In the Best Case?

$$\frac{20 \times 2^{30}}{2^{12}} = 20 \times 2^{18} \text{ PTEs}$$

However, these are how many PTEs we need across only the L0 page tables!

$$\frac{20 \times 2^{18}}{2^9} = 20 \times 2^9 = 10240 \text{ full L0 page tables (40 MB)}$$

Each L1 page table can point to 512 L0 page tables

$$\frac{10240}{512} = 20 \text{ full L1 page tables}$$

$$\text{So we'd need } \mathbf{10260} \text{ full page tables} = \frac{10260 \times 4096}{2^{20}} = 40.078125 \text{ MiB}$$