

Operating Systems

CPU Scheduling Policies

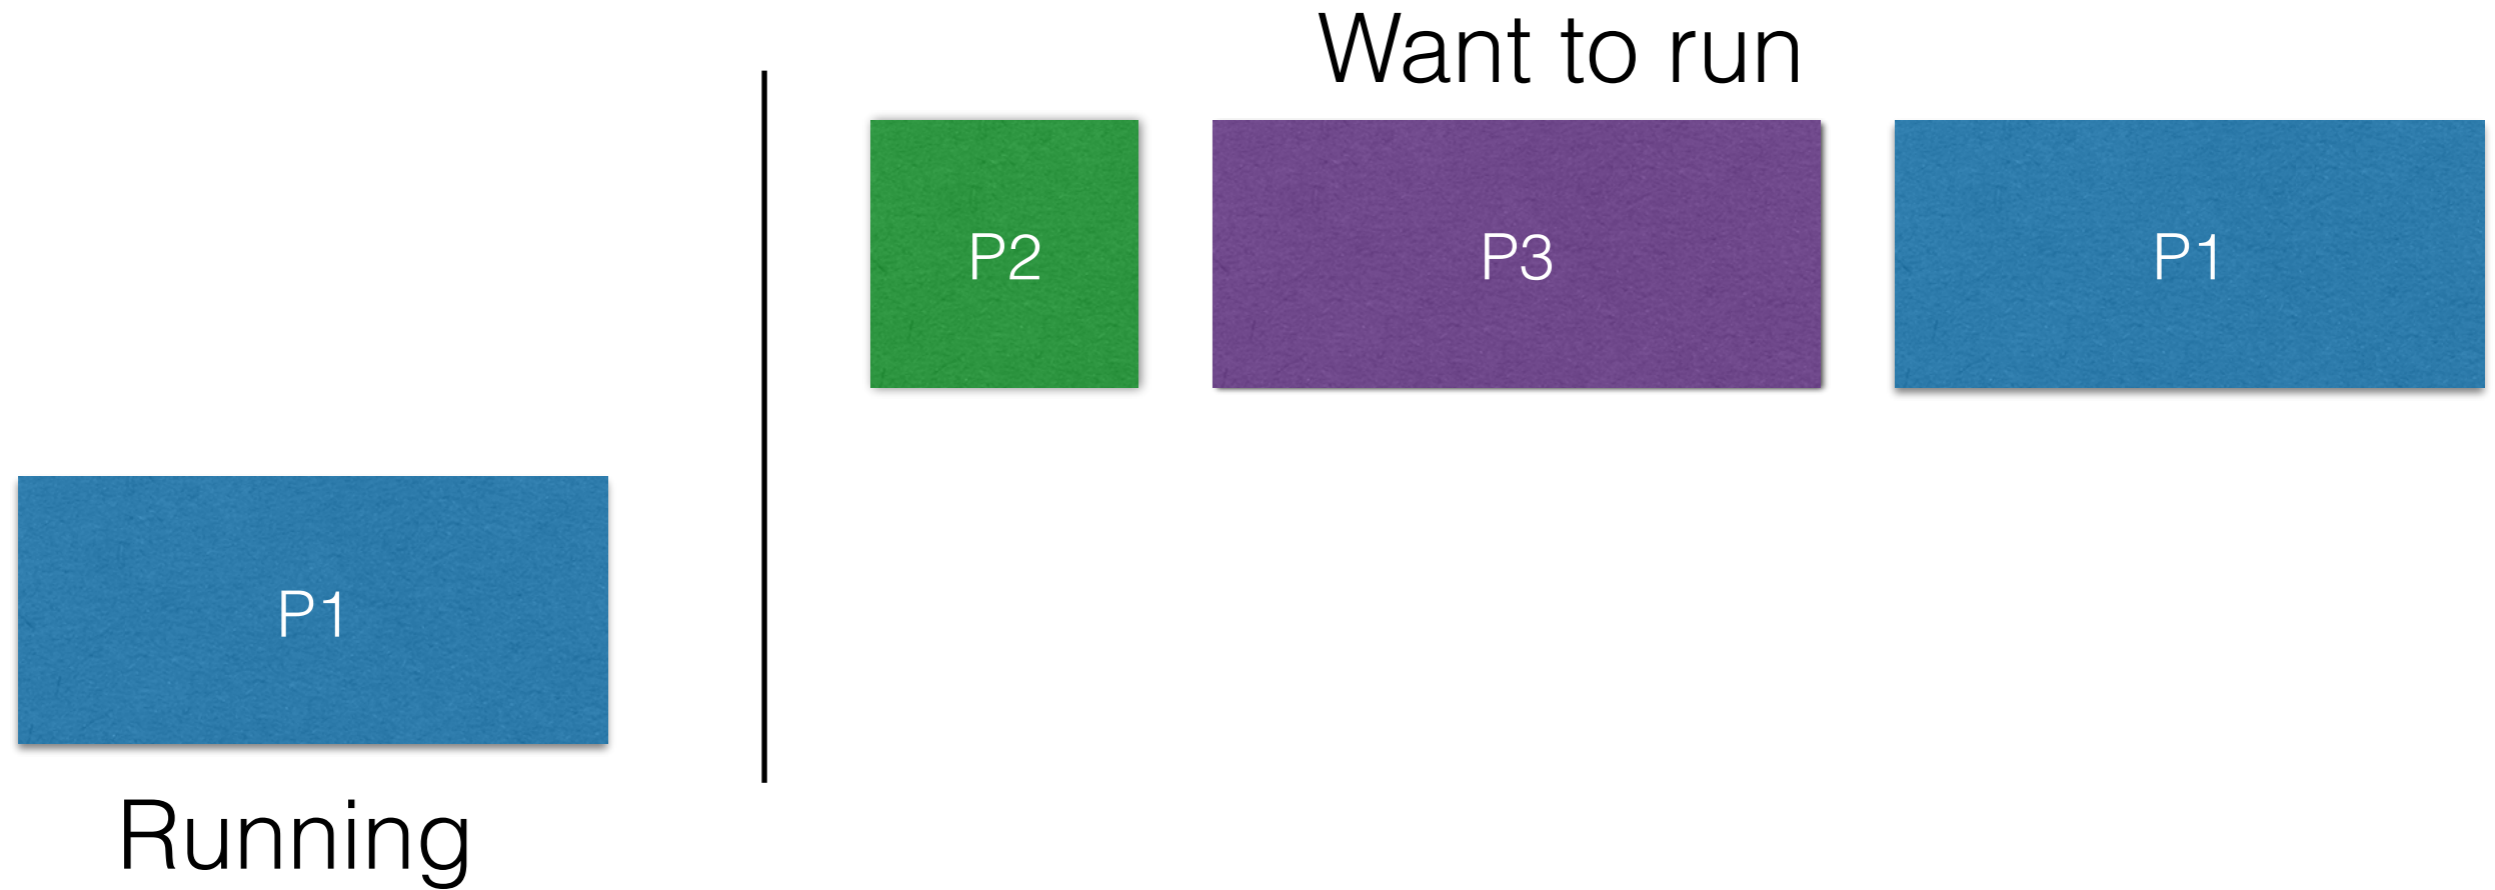
Nipun Batra

CPU Virtualisation Revisited

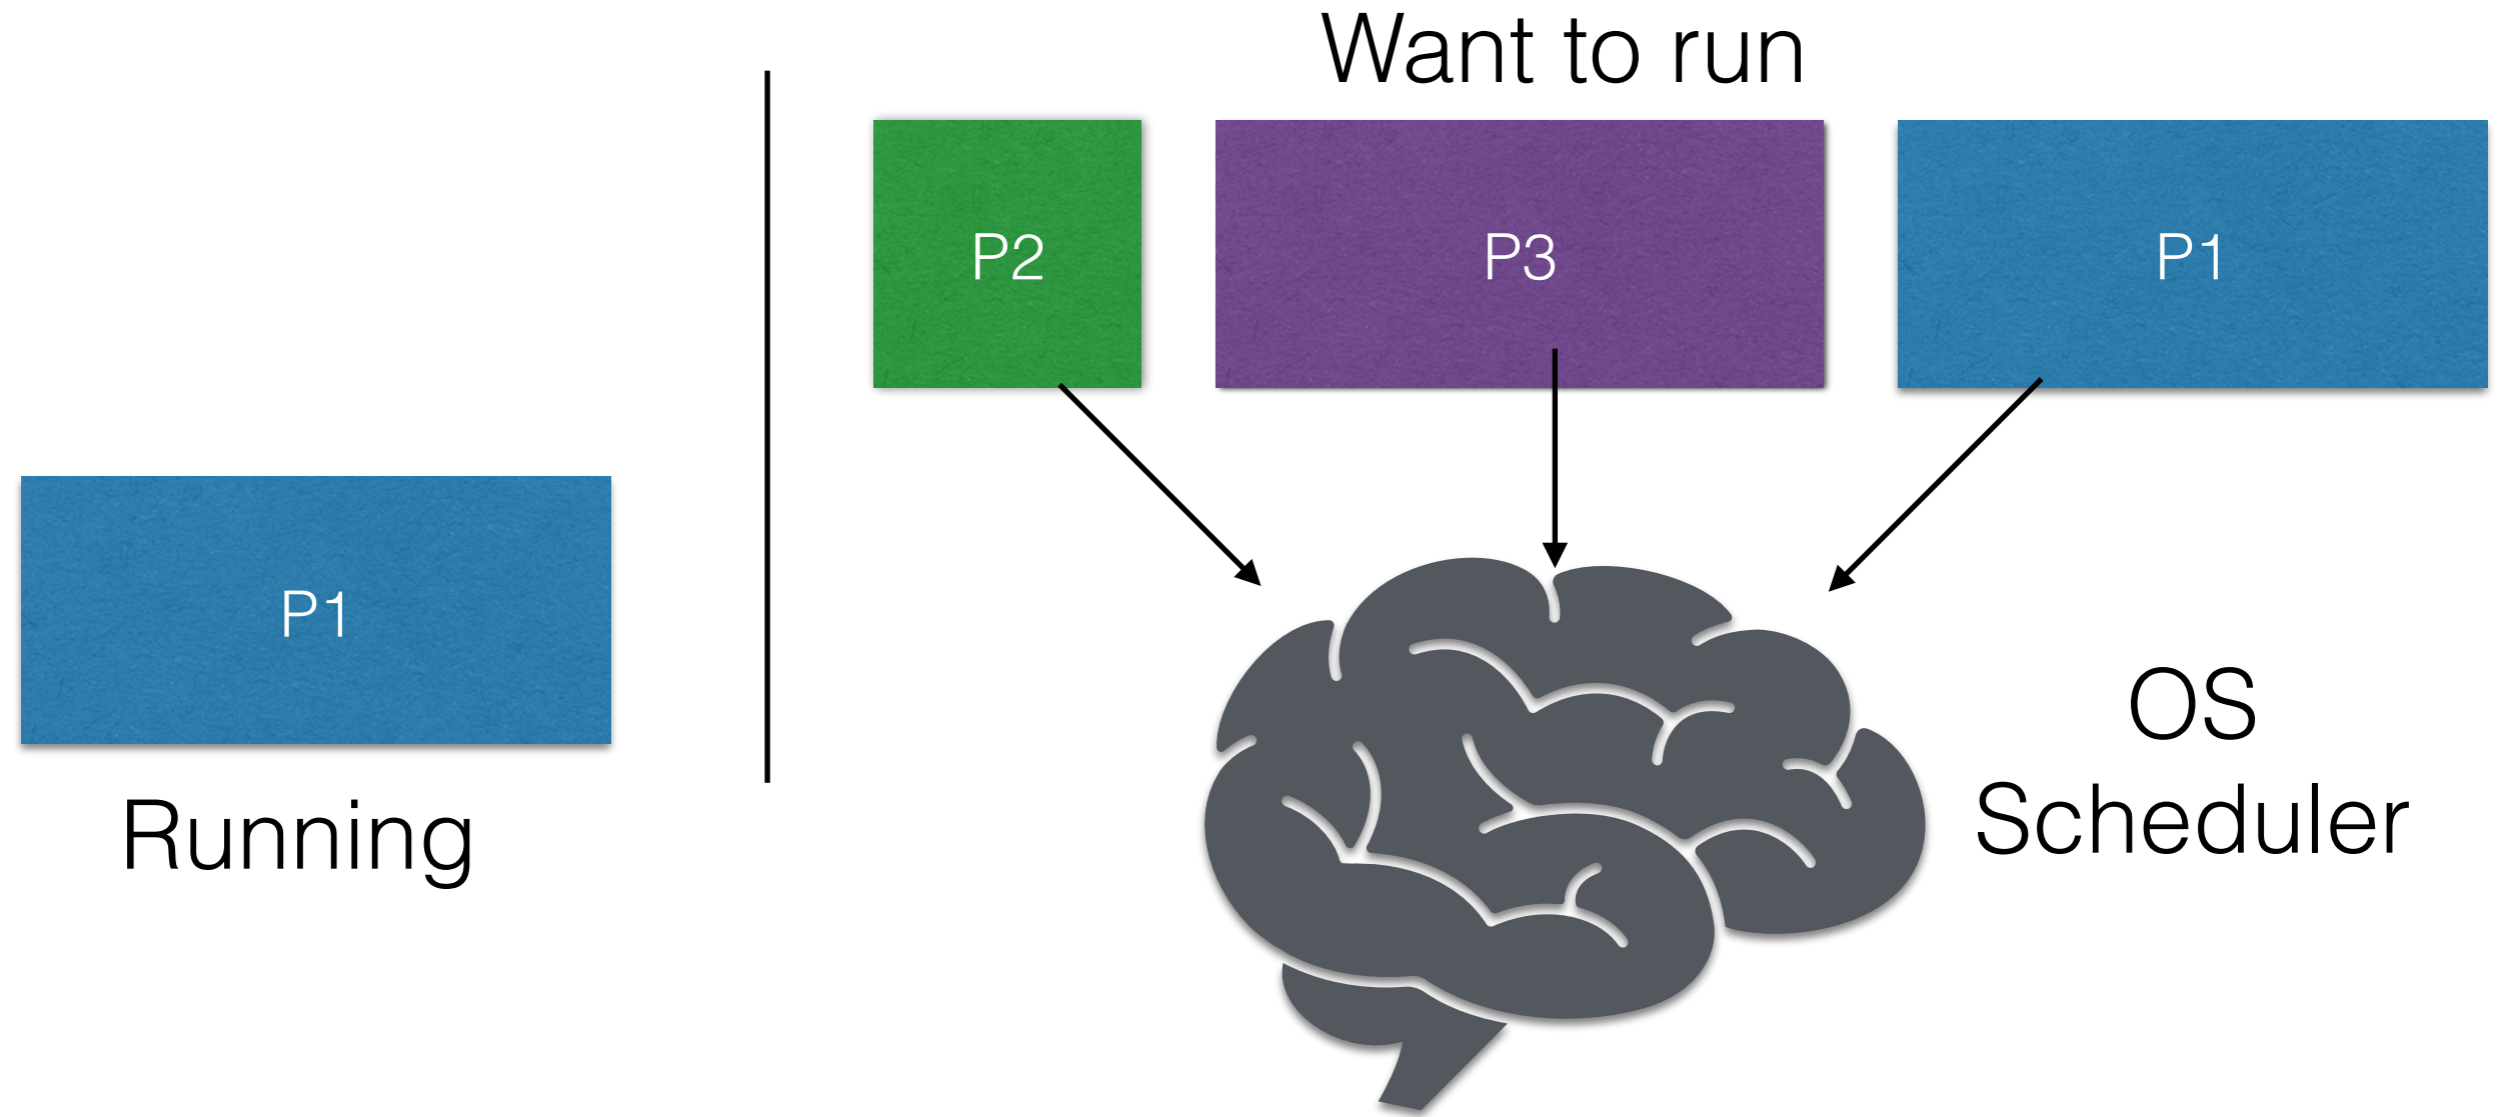


Running

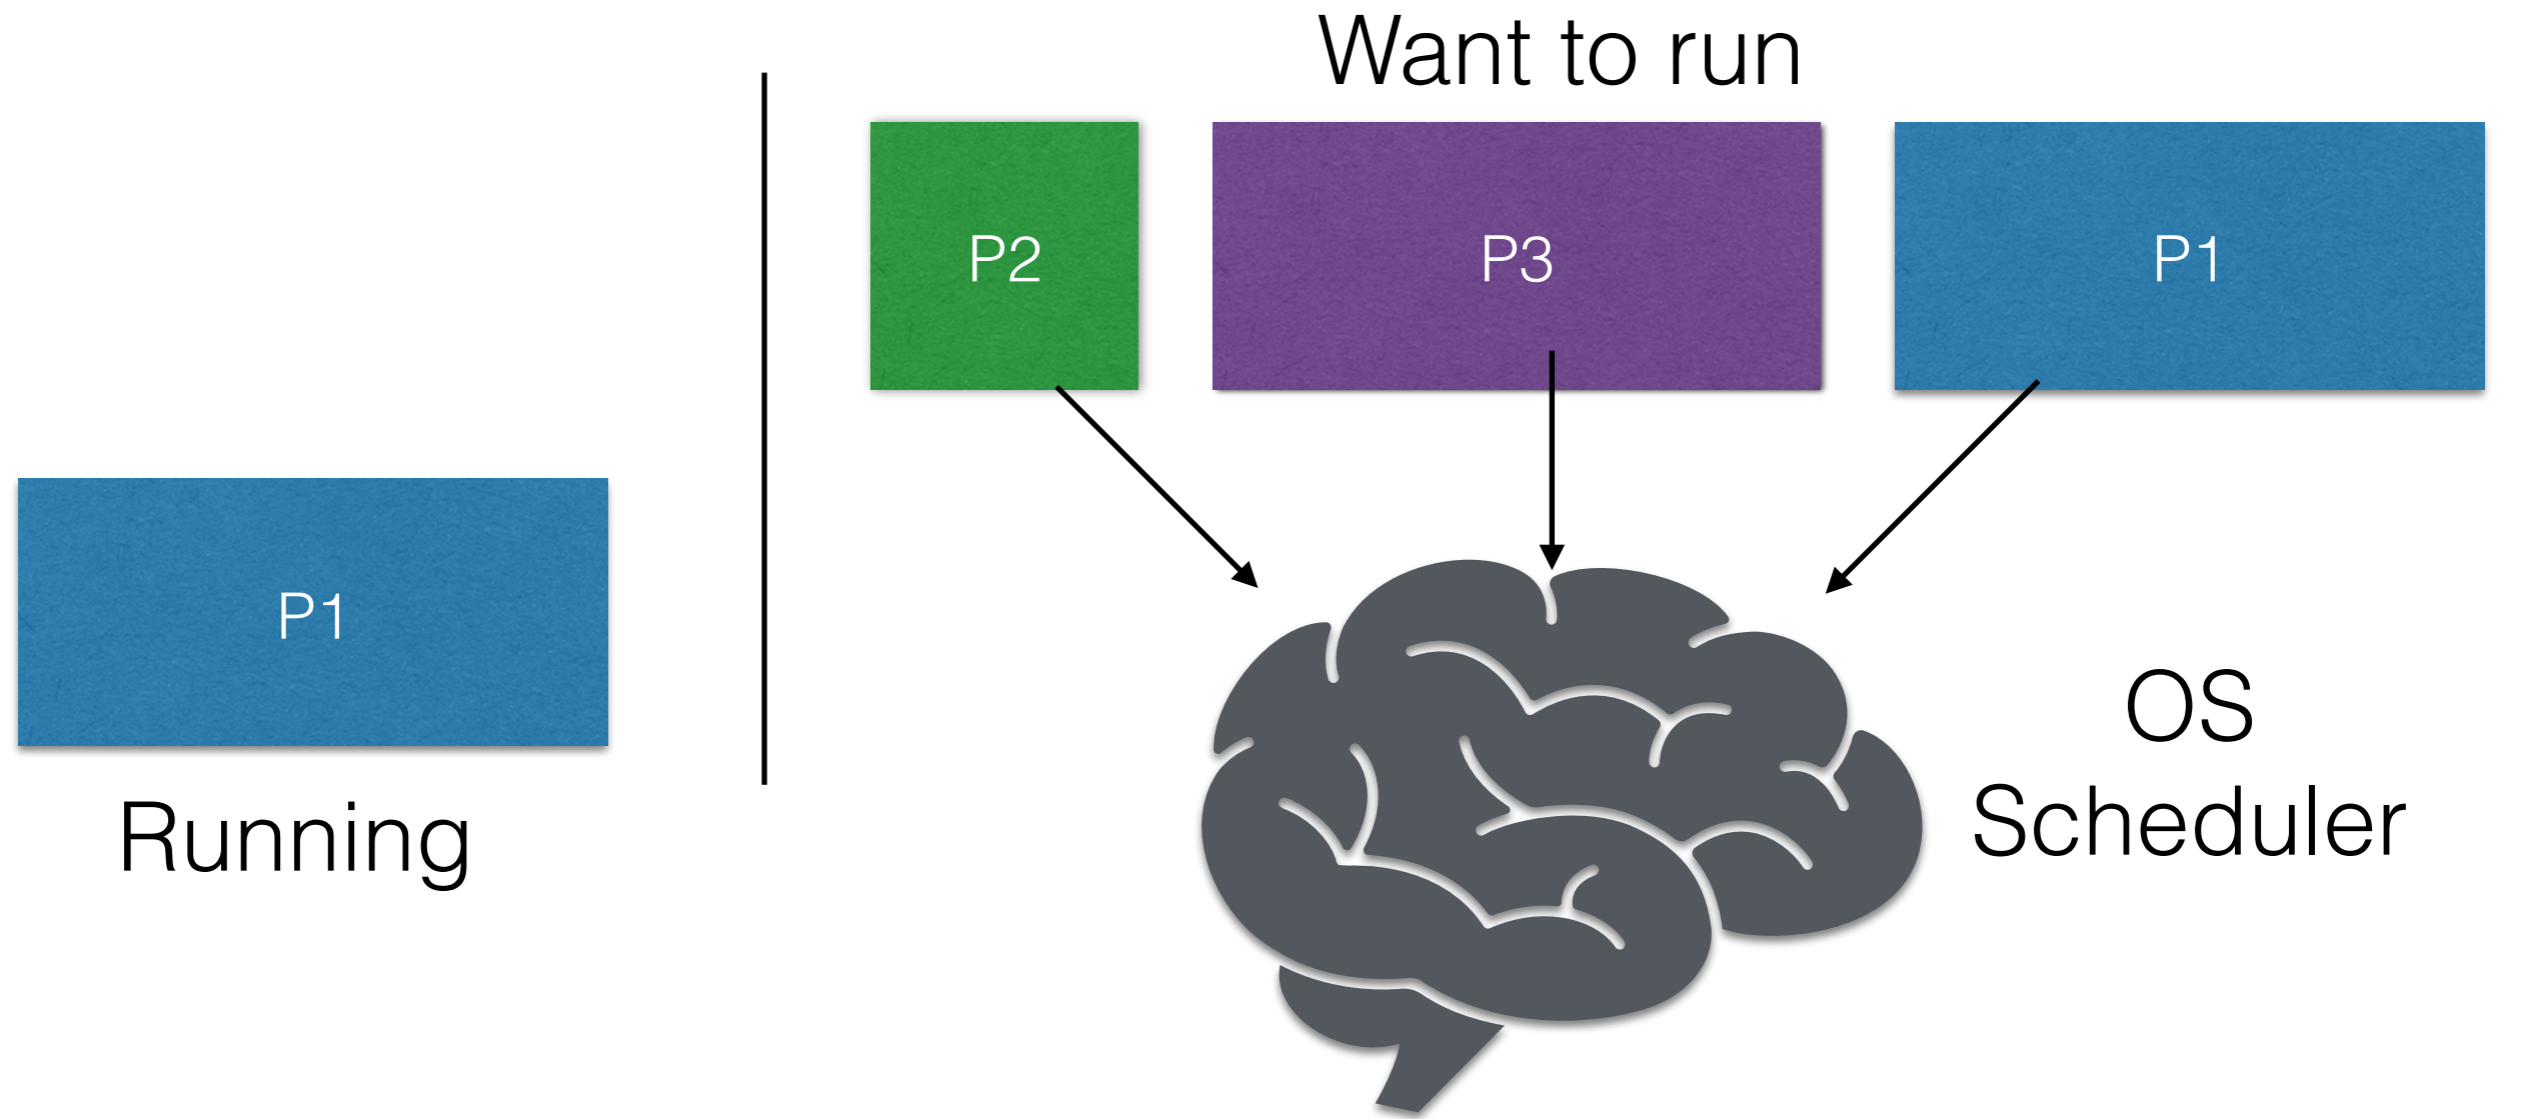
CPU Virtualisation Revisited



CPU Virtualisation Revisited

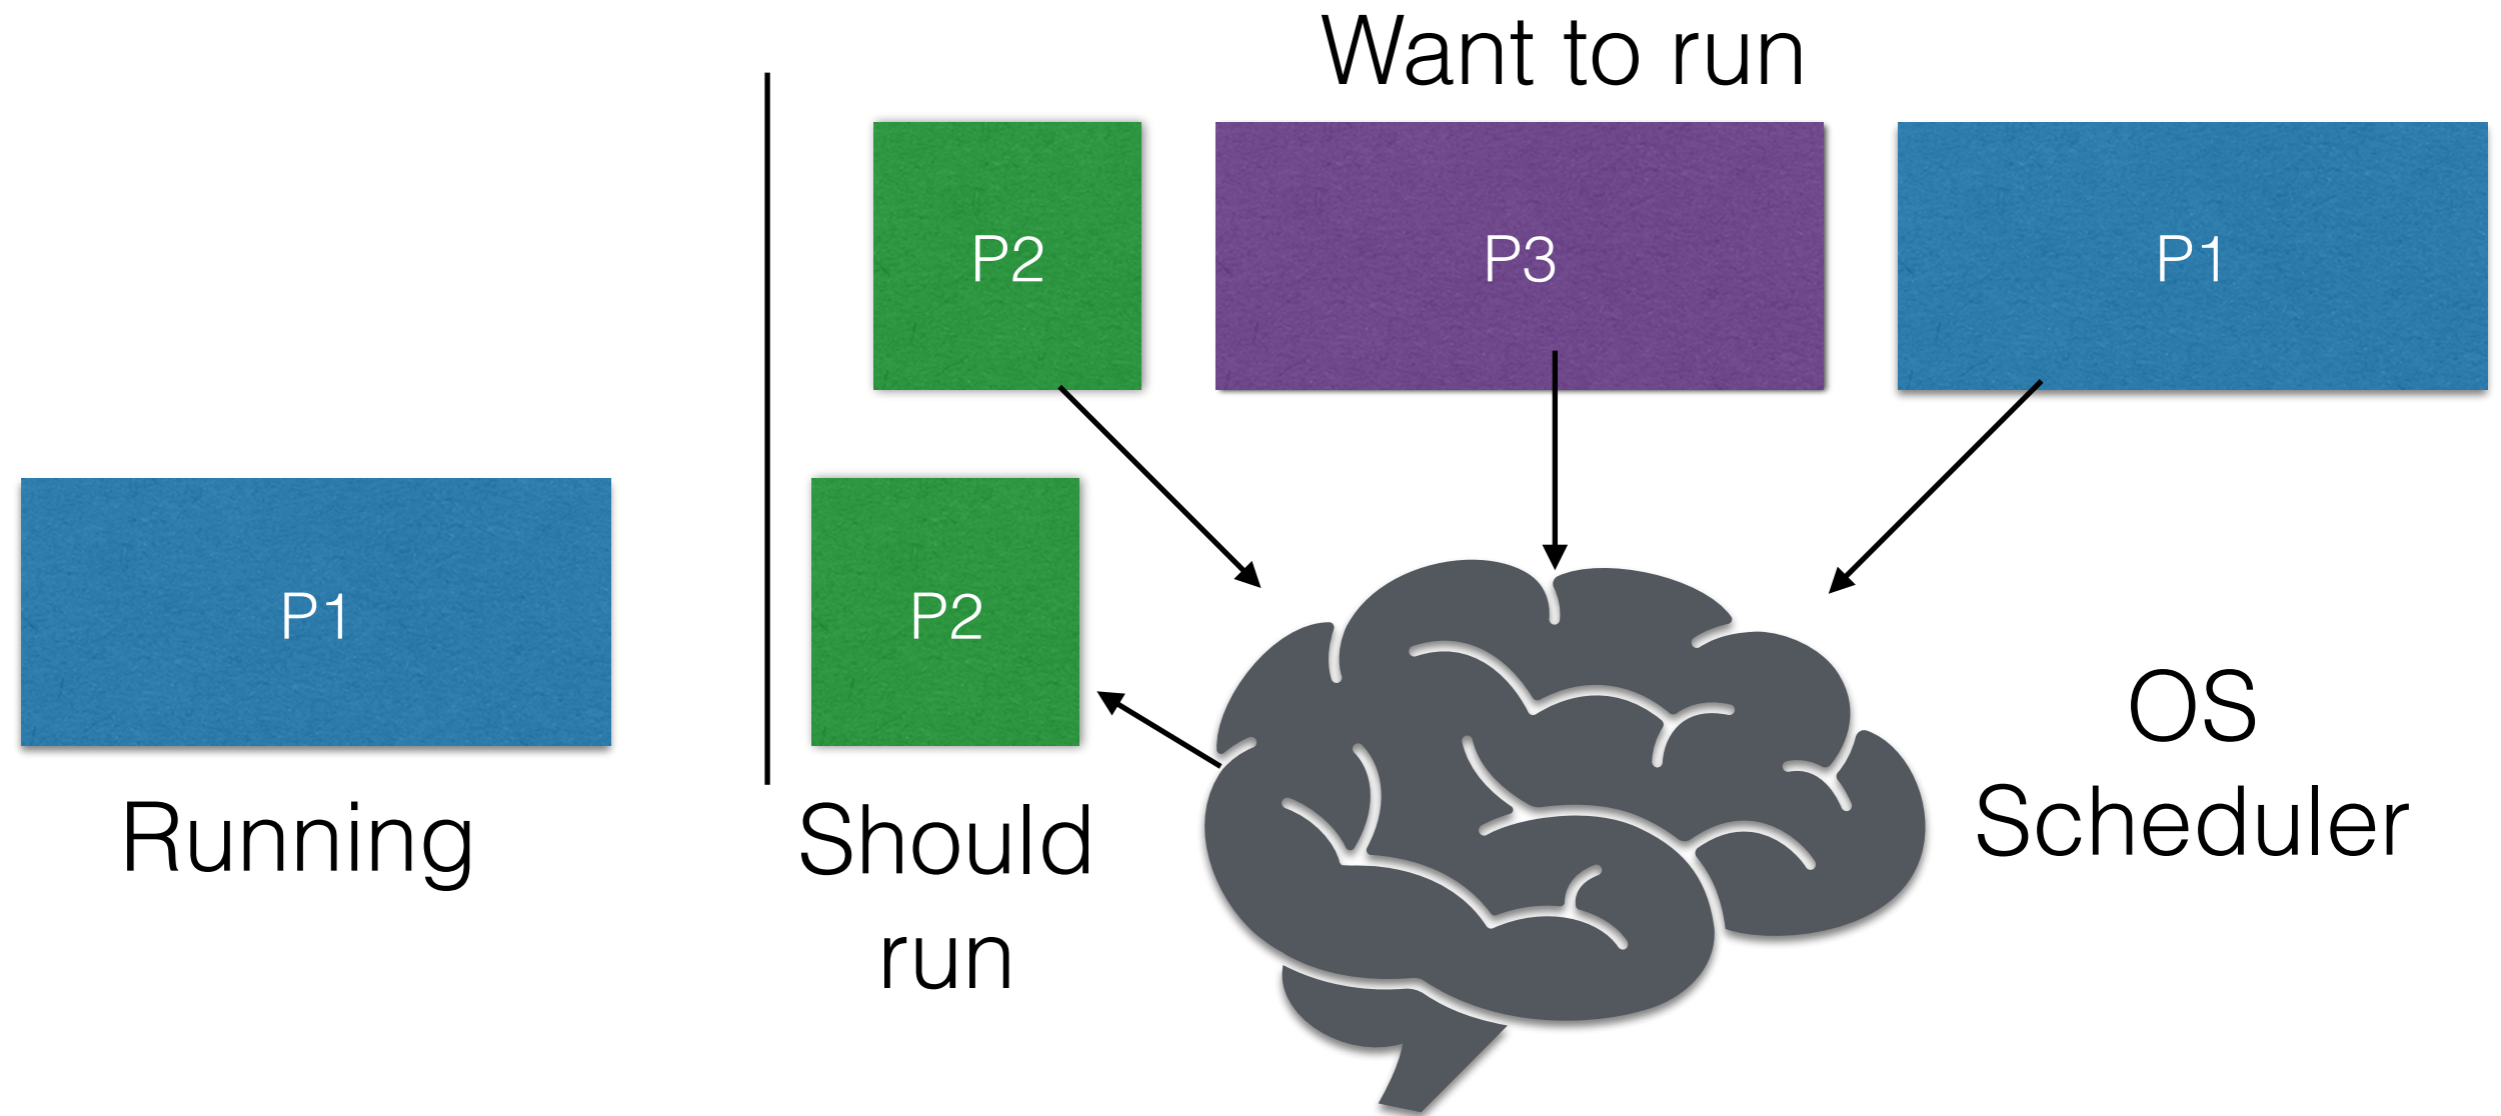


CPU Virtualisation Revisited

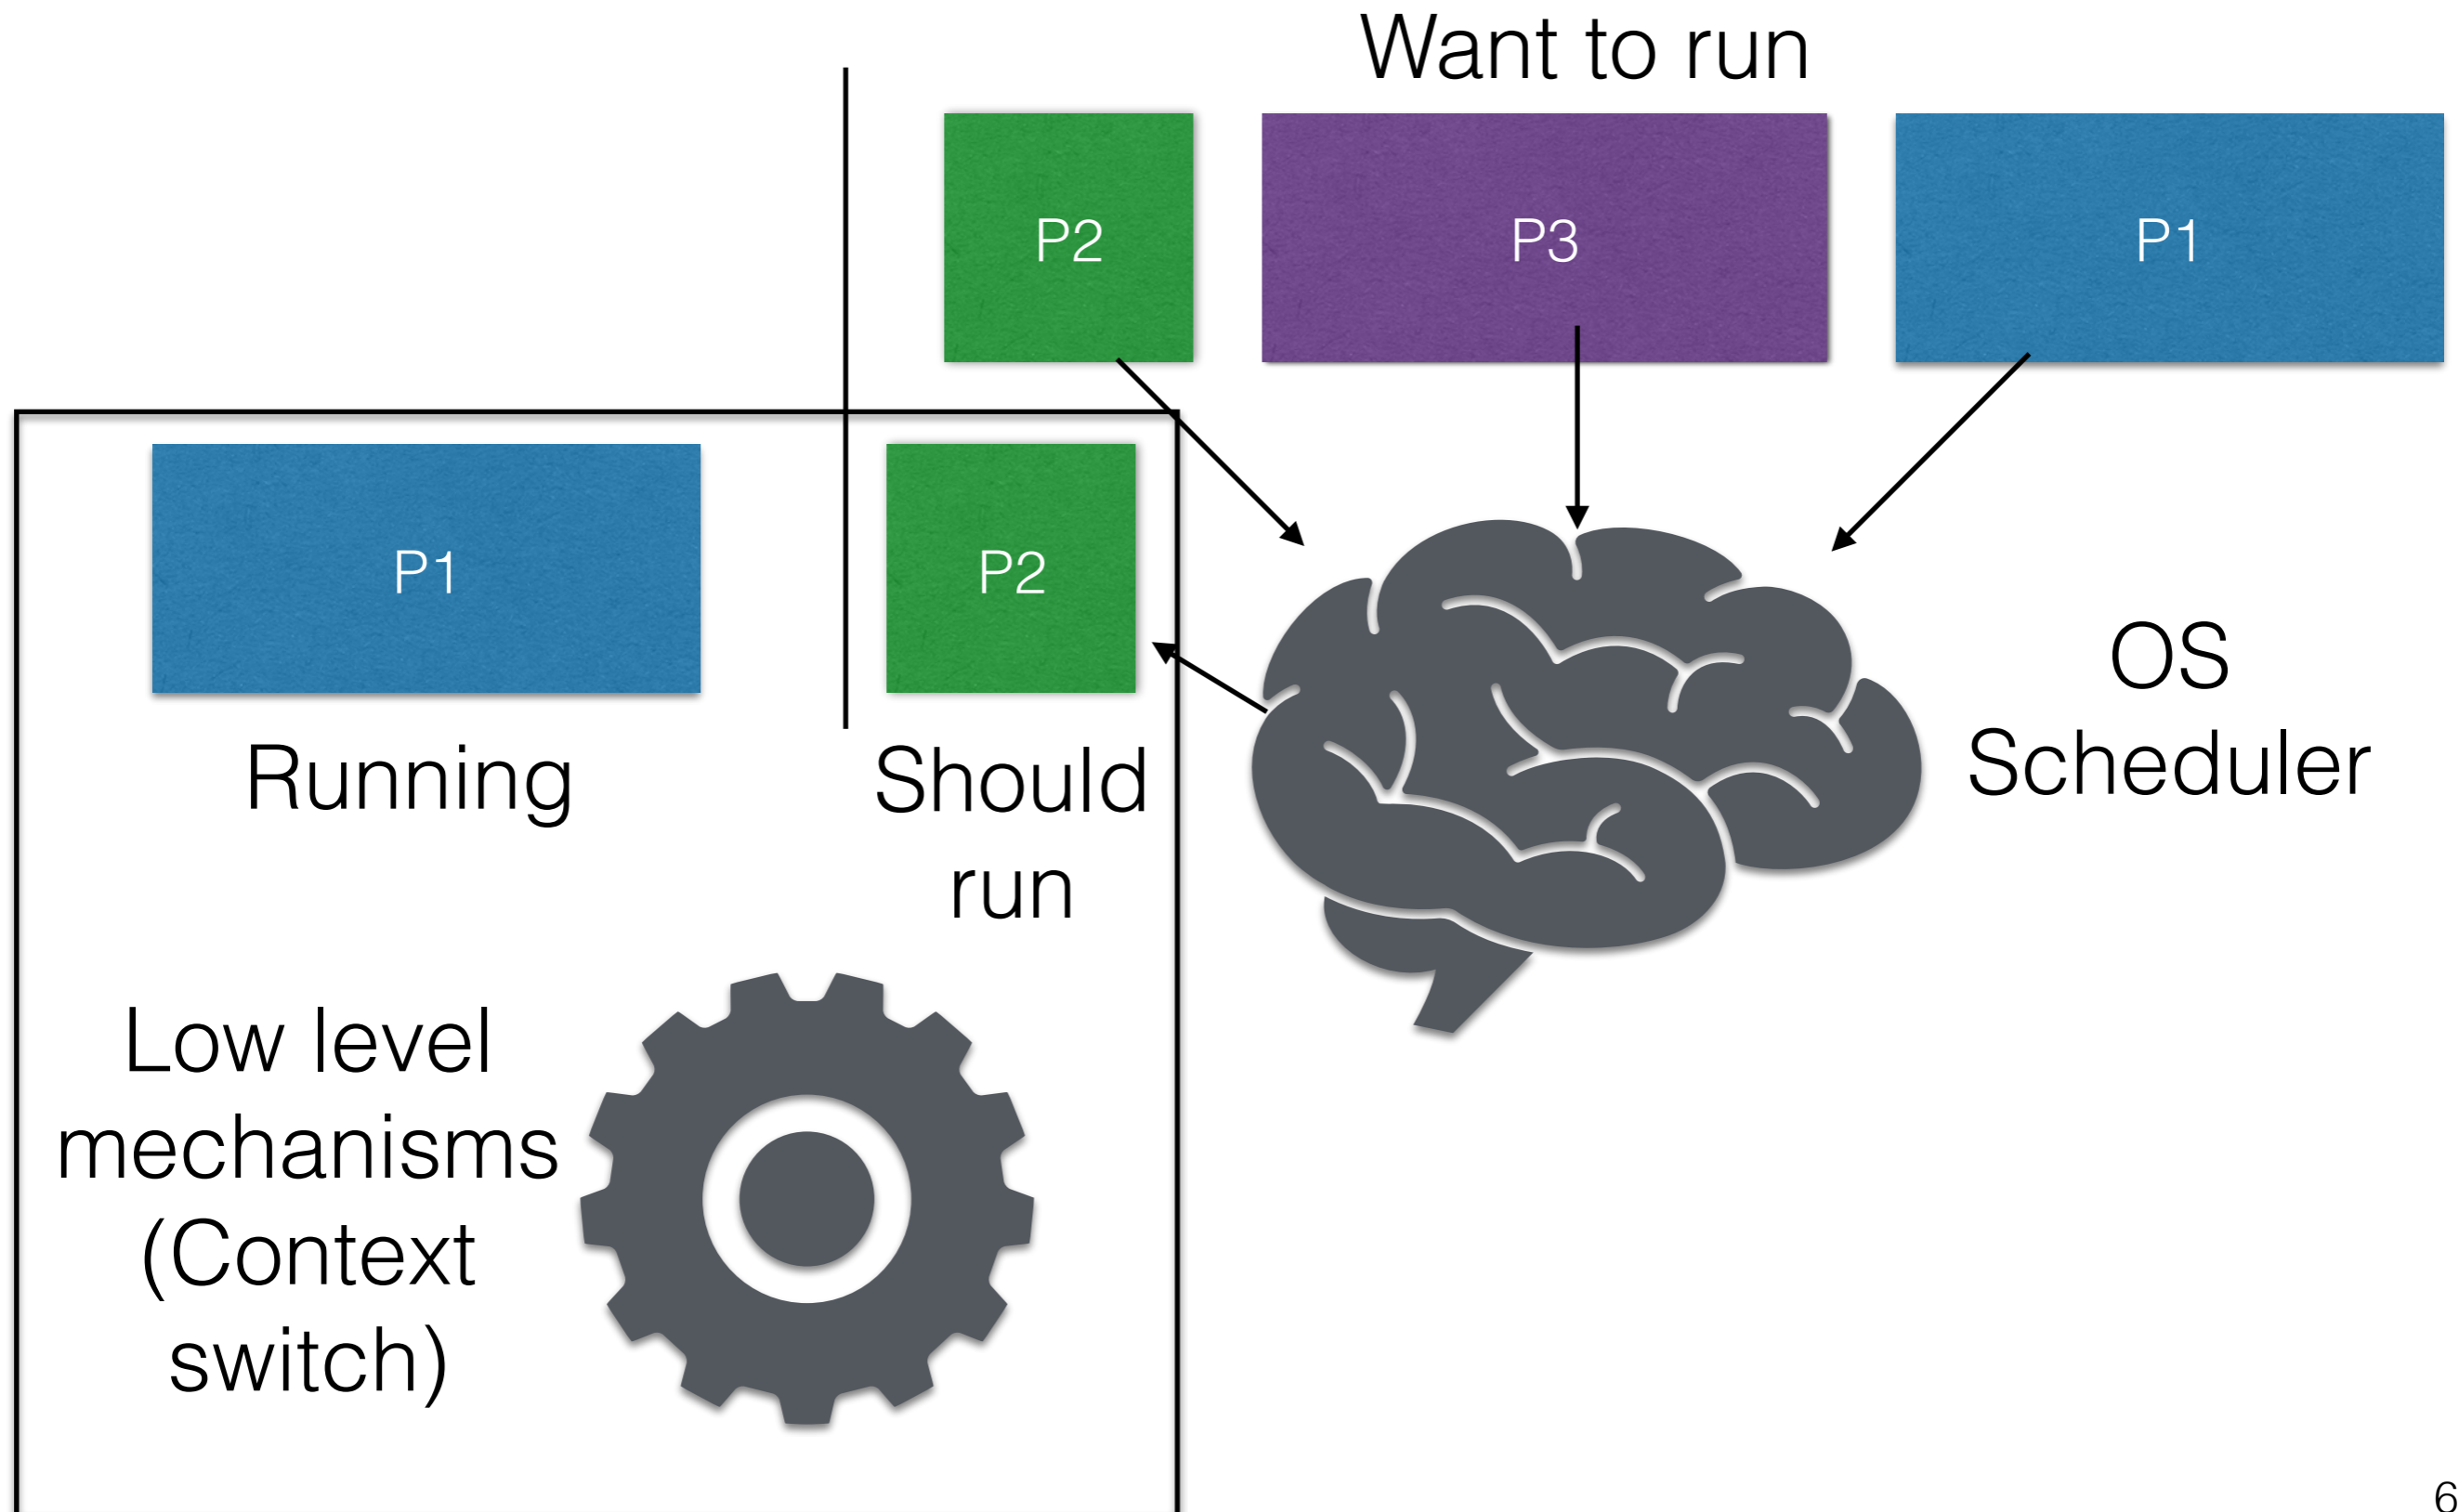


Next P = f(run time,
metric, type of process, ...)

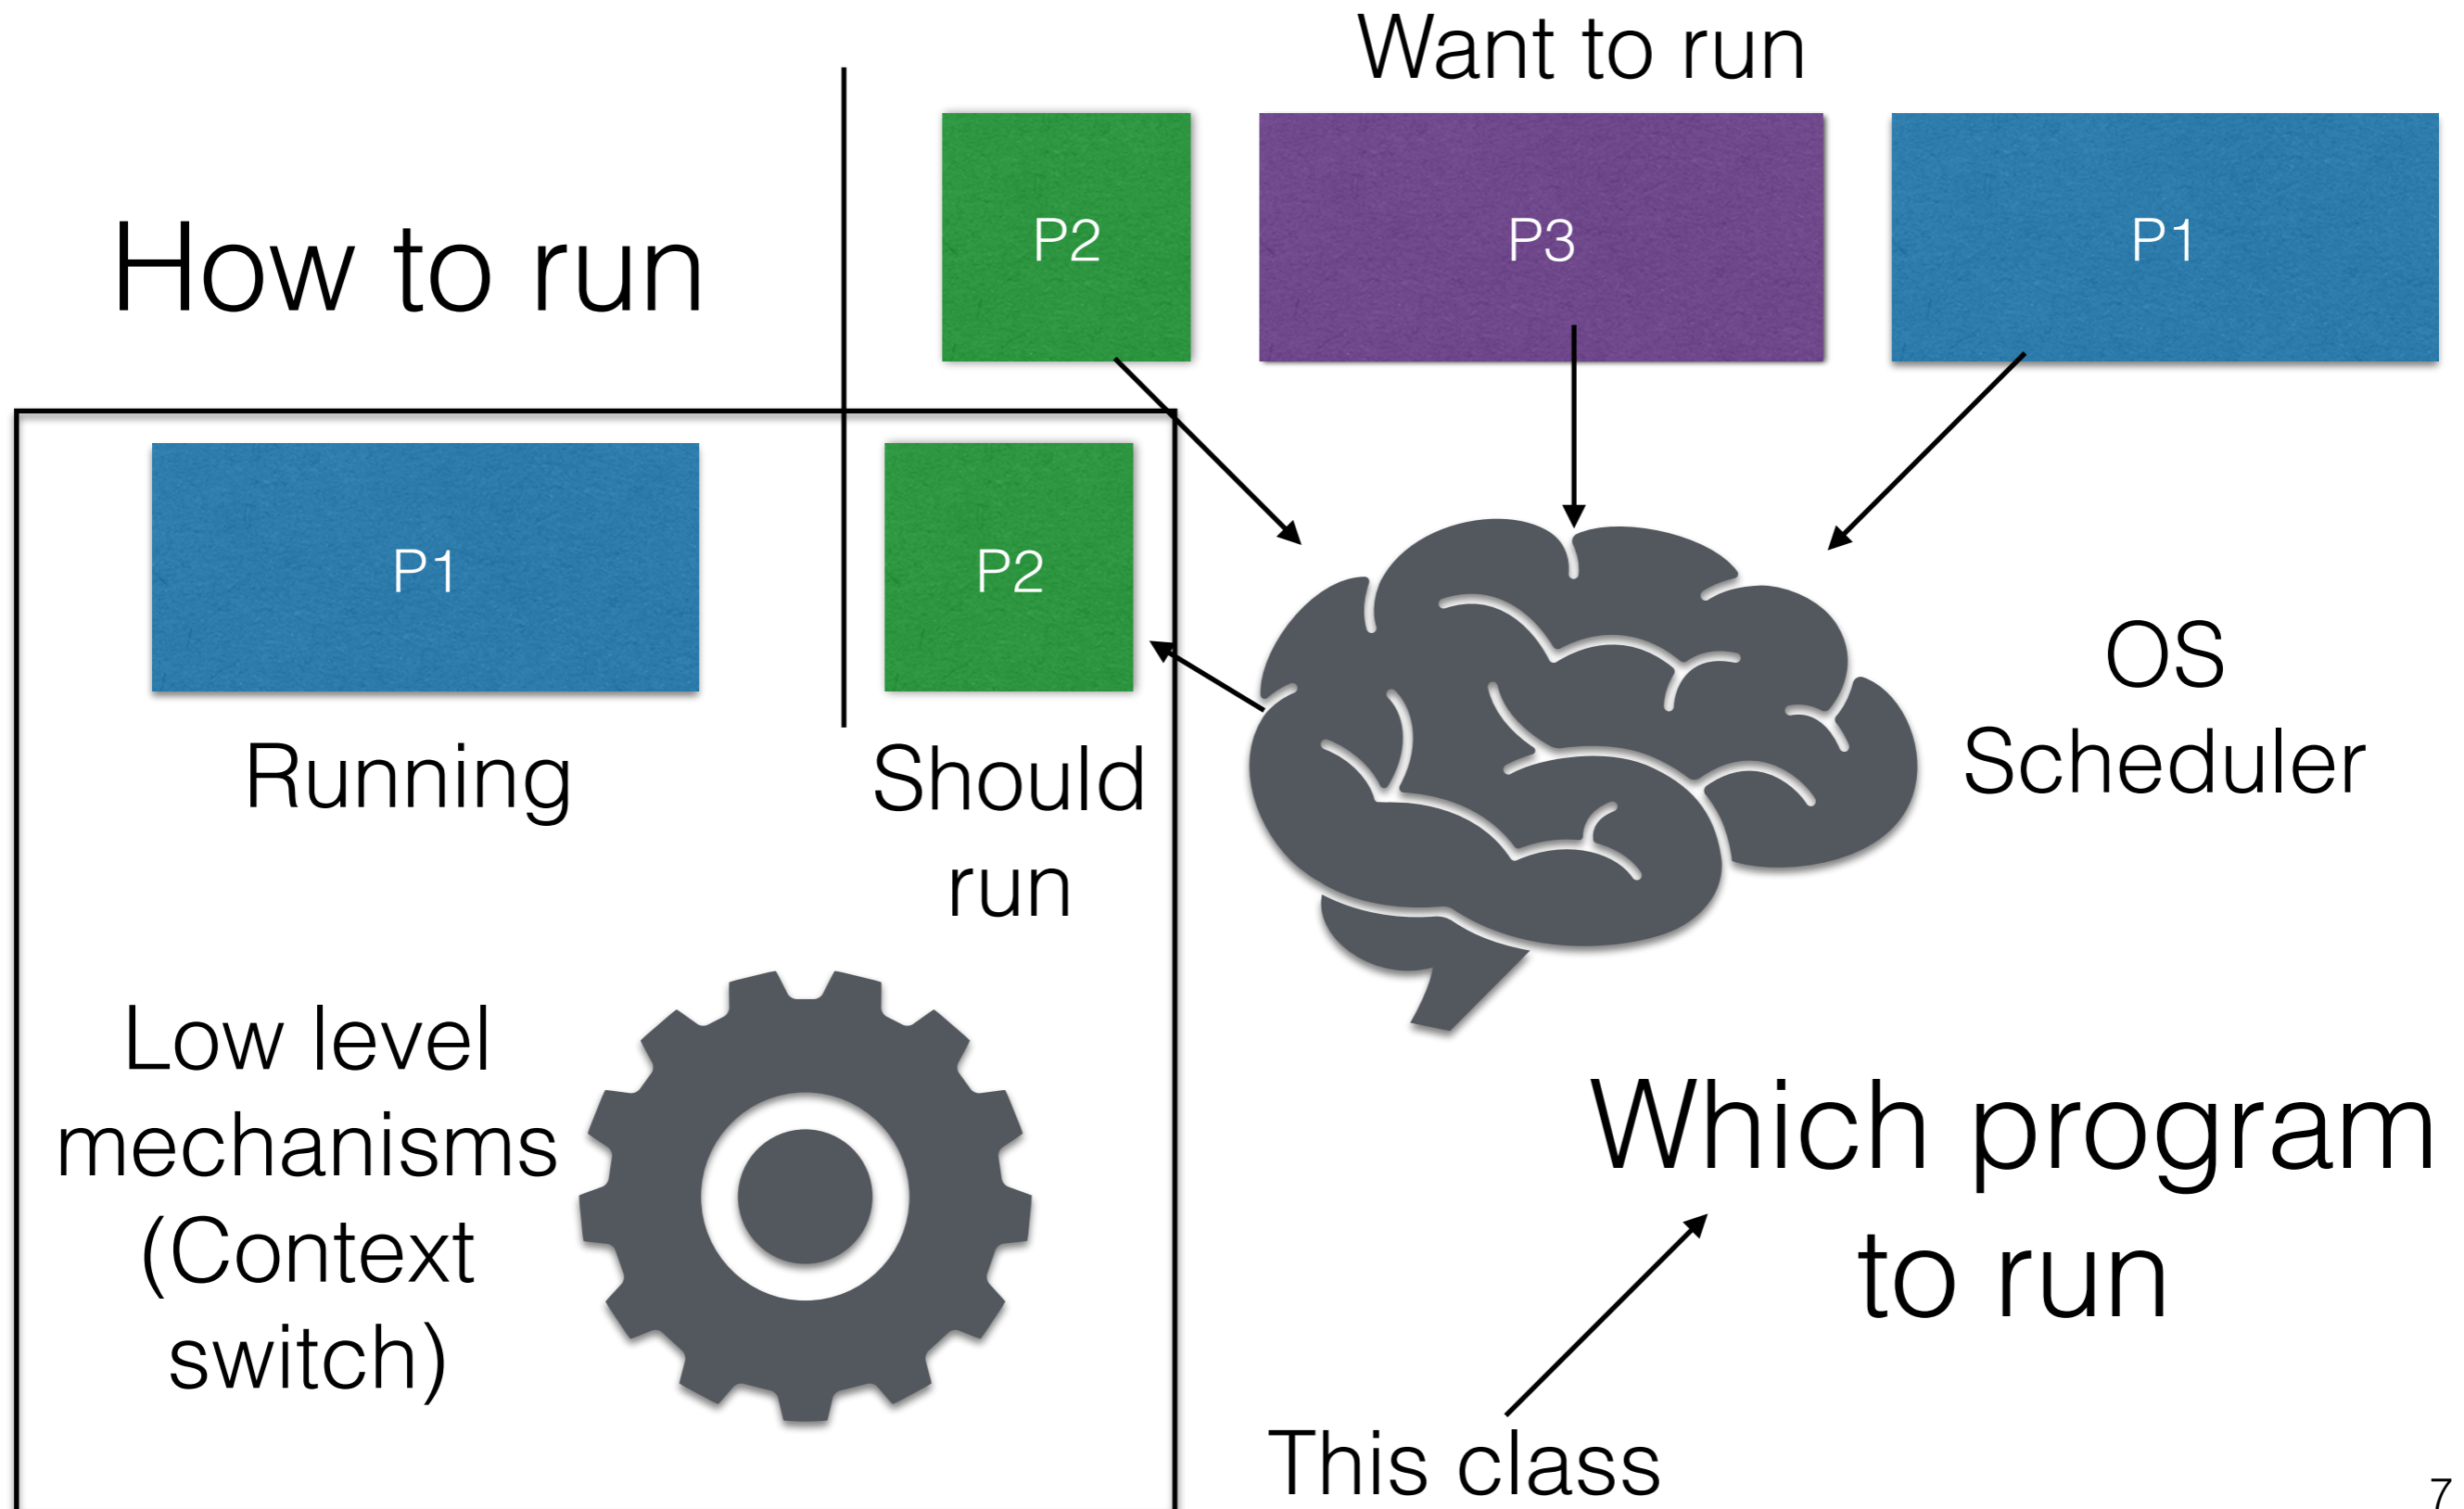
CPU Virtualisation Revisited



CPU Virtualisation Revisited



CPU Virtualisation Revisited






Workload






Set of
processes

Workload





Set of
processes

Workload





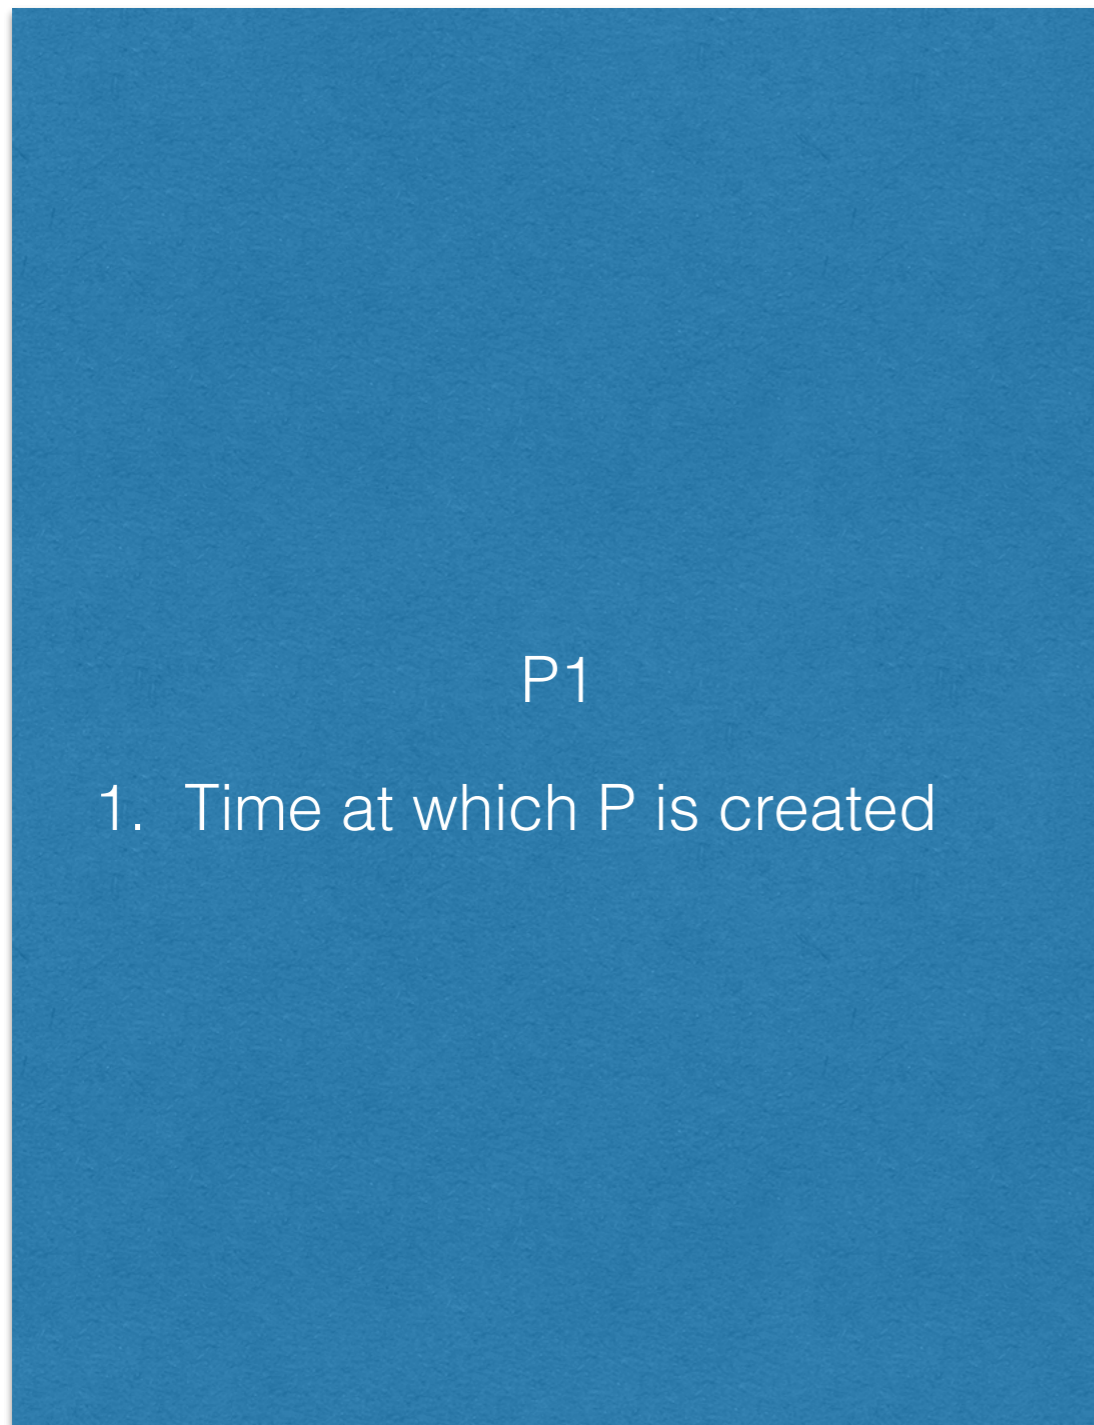
Set of
processes

Workload




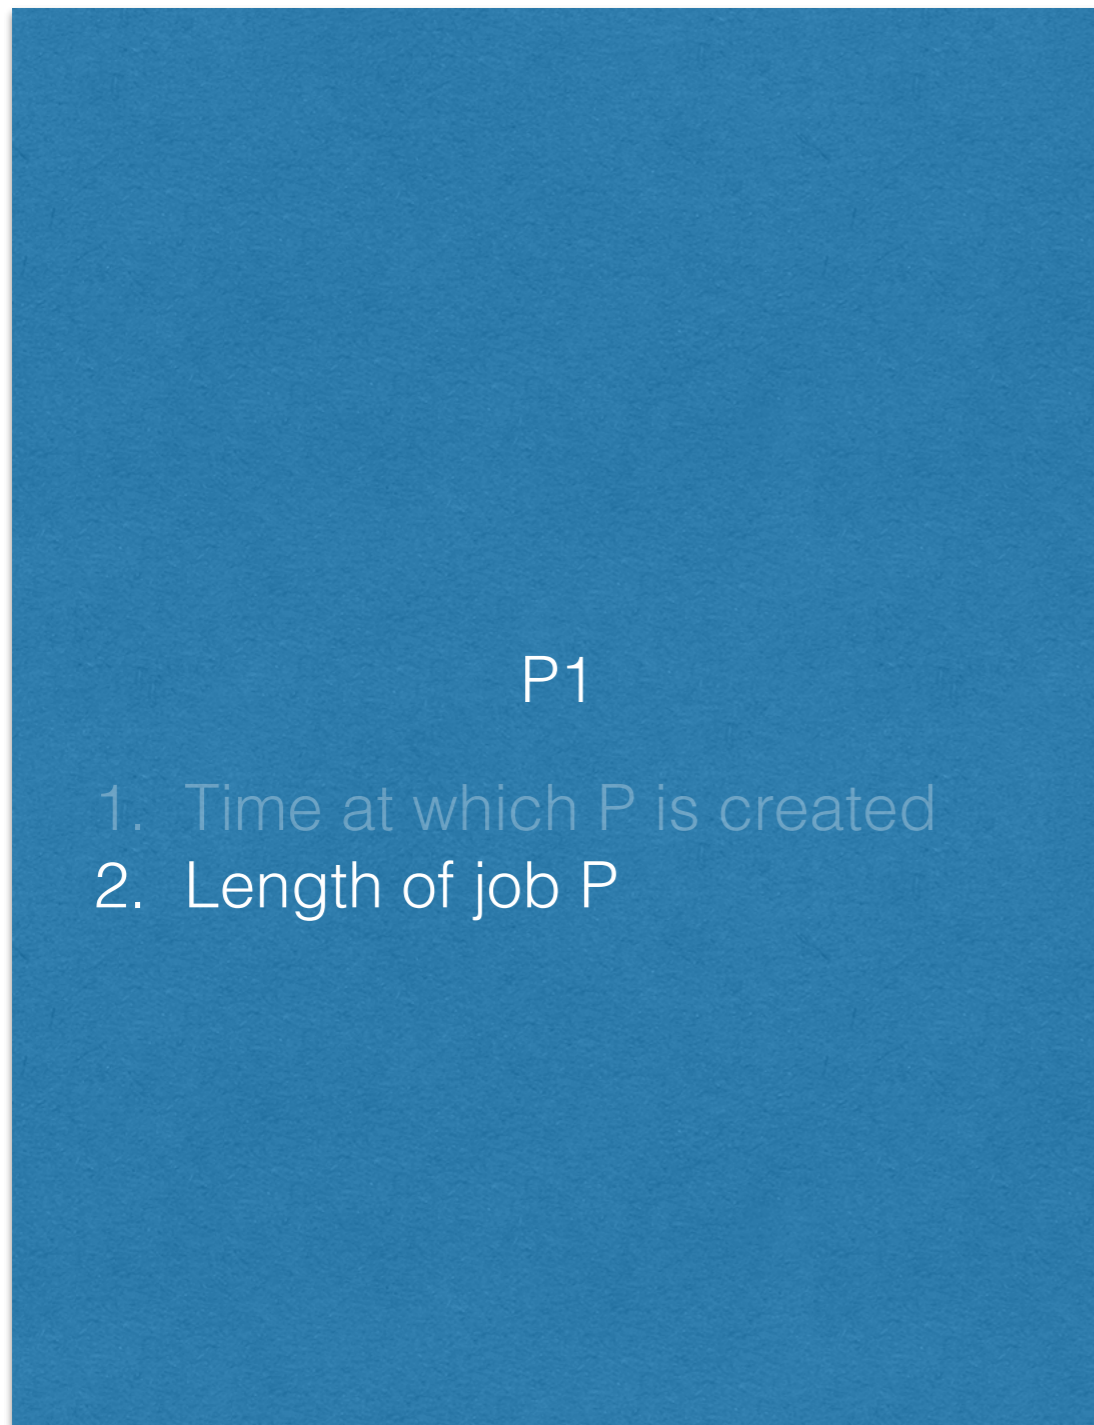
Set of
processes

Workload




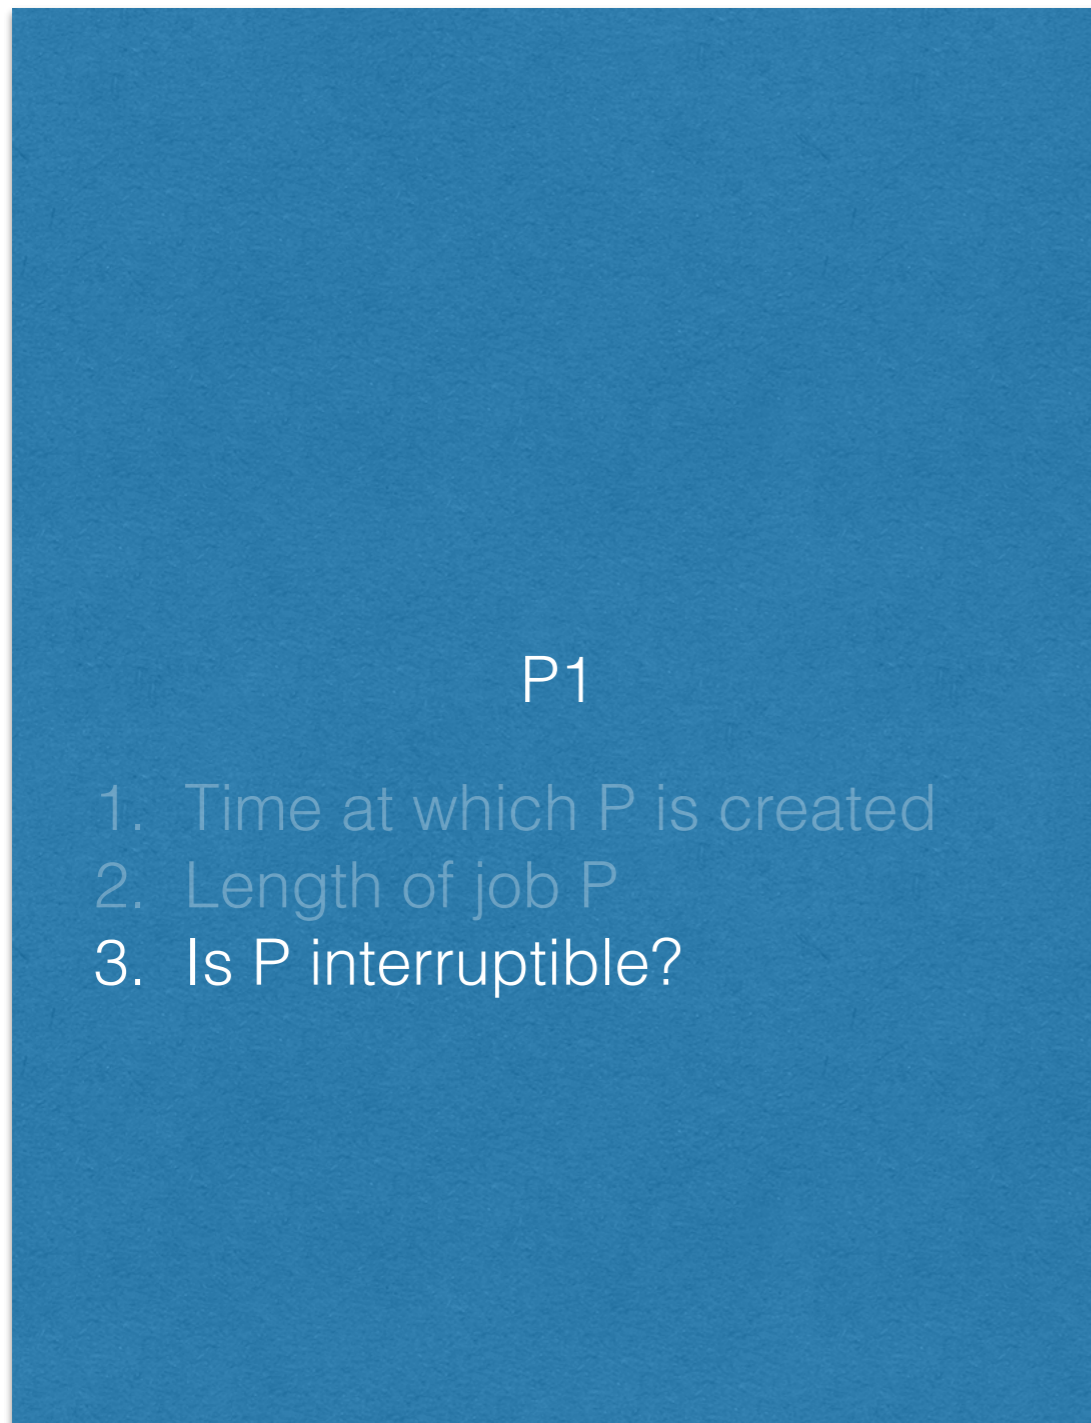
Set of
processes

Workload



Set of
processes

Workload

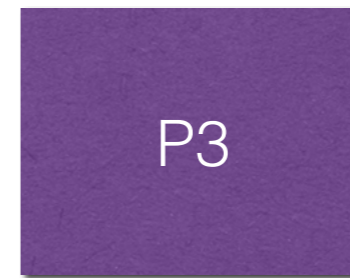


Set of
processes

Workload

P1

1. Time at which P is created
2. Length of job P
3. Is P interruptible?
4. Is it only CPU? Or CPU & IO?



Set of
processes

Workload Assumptions

1. Each job runs for the same time
2. All jobs arrive at the same time
3. Once started, each job runs to completion
4. All jobs use only the CPU
5. Run time of each job is known

Scheduling Metrics


Scheduling Metrics



P1

Scheduling Metrics


Arrived at t_1



Scheduling Metrics

Arrived at t_1


Completed at t_2



Scheduling Metrics

Arrived at t_1


Completed at t_2



Turnaround time (P1) = $t_2 - t_1$

Scheduling Metrics


Arrived at t_1 Completed at t_2



Turnaround time (P1) = $t_2 - t_1$

Scheduling Metrics

Arrived at t_1 Completed at t_2




Turnaround time (P1) = $t_2 - t_1$

Scheduling Metrics

Arrived at t_1

Completed at t_2

Completed at t_3



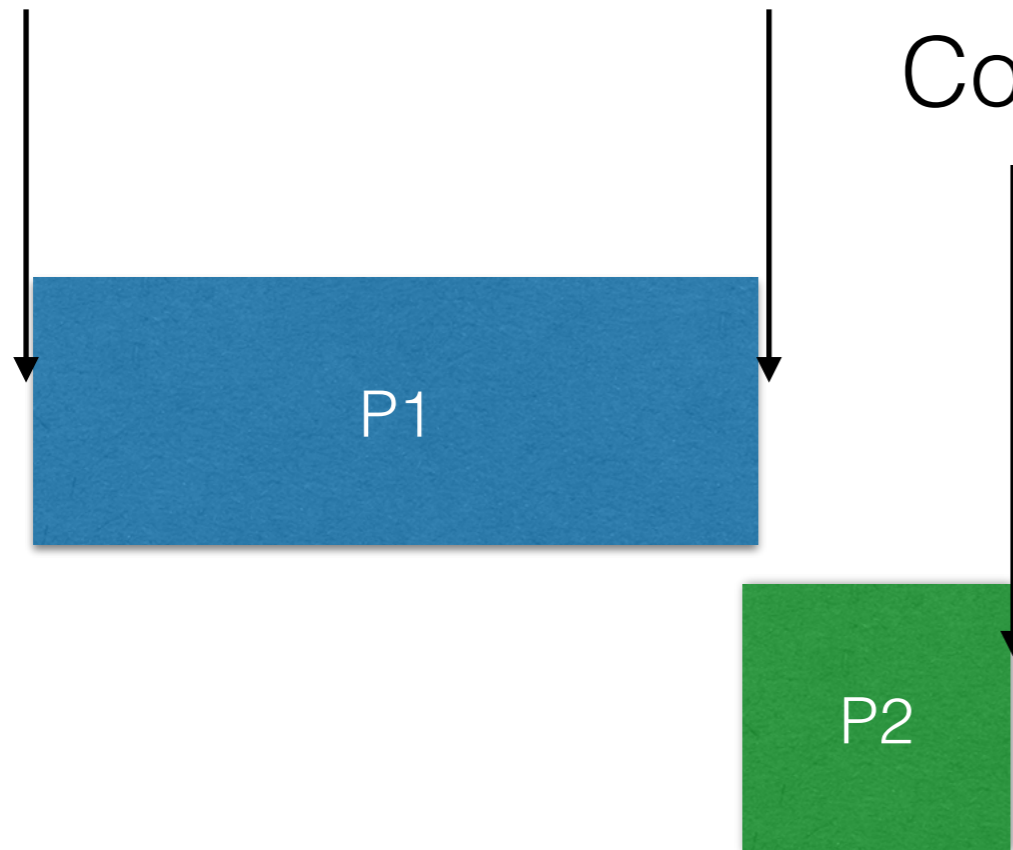
Turnaround time (P1) = $t_2 - t_1$

Scheduling Metrics

Arrived at t_1

Completed at t_2

Completed at t_3



Turnaround time (P1) = $t_2 - t_1$


Turnaround time (P2) = $t_3 - t_1$

FIFO scheduling

A, B and C come at $T = 0$, $0 + \Delta$, $0 + 2\Delta$

FIFO scheduling


A, B and C come at $T = 0$, $0 + \Delta$, $0 + 2\Delta$



FIFO scheduling


A, B and C come at $T = 0$, $0 + \Delta$, $0 + 2\Delta$

- $T_{\text{completion}}(A) = 10$
- $\text{Turnaround}(A) = 10 - 0 = 0$



FIFO scheduling


A, B and C come at $T = 0$, $0 + \Delta$, $0 + 2\Delta$



FIFO scheduling


A, B and C come at $T = 0$, $0 + \Delta$, $0 + 2\Delta$

- $T_{\text{completion}}(B) = 20$
- $\text{Turnaround}(B) = 20 - 0 = 20$



FIFO scheduling


A, B and C come at $T = 0$, $0 + \Delta$, $0 + 2\Delta$



FIFO scheduling


A, B and C come at $T=0$, $0 + \Delta$, $0 + 2\Delta$

- $T_{\text{completion}}(C) = 30$
- $\text{Turnaround}(C) = 30 - 0 = 30$




FIFO scheduling

A, B and C come at $T = 0$, $0 + \Delta$, $0 + 2\Delta$



FIFO scheduling

A, B and C come at $T = 0$, $0 + \Delta$, $0 + 2\Delta$



$$\text{Avg. Turnaround Time} = \frac{(10+20+30)}{3} = 20$$

Workload Assumptions

1. Each job runs for the same time
2. All jobs arrive at the same time
3. Once started, each job runs to completion
4. All jobs use only the CPU
5. Run time of each job is known

Workload Assumptions


1. ~~Each job runs for the same time~~
2. All jobs arrive at the same time
3. Once started, each job runs to completion
4. All jobs use only the CPU
5. Run time of each job is known

FIFO scheduling




$$\text{Avg. Turnaround Time} = \frac{(30+40+50)}{3} = 40$$


FIFO scheduling




$$\text{Avg. Turnaround Time} = (30+40+50)/3 = 40$$



FIFO scheduling




$$\text{Avg. Turnaround Time} = (30+40+50)/3 = 40$$




$$\text{Avg. Turnaround Time} = (10+40+50)/3 = 33.3$$


FIFO scheduling



$$\text{Avg. Turnaround Time} = (30+40+50)/3 = 40$$




$$\text{Avg. Turnaround Time} = (10+40+50)/3 = 33.3$$




FIFO scheduling



$$\text{Avg. Turnaround Time} = (30+40+50)/3 = 40$$



$$\text{Avg. Turnaround Time} = (10+40+50)/3 = 33.3$$




$$\text{Avg. Turnaround Time} = (10+20+50)/3 \sim 27$$

FIFO scheduling



$$\text{Avg. Turnaround Time} = (30+40+50)/3 = 40$$





$$\text{Avg. Turnaround Time} = (10+40+50)/3 = 33.3$$




$$\text{Avg. Turnaround Time} = (10+20+50)/3 \sim 27$$

FIFO scheduling



Shortest Job First (SJF) scheduling



Shortest Job First (SJF) scheduling



$$\text{Avg. Turnaround Time} = \frac{(10+20+50)}{3} \sim 27$$

Workload Assumptions

1. ~~Each job runs for the same time~~
2. All jobs arrive at the same time
3. Once started, each job runs to completion
4. All jobs use only the CPU
5. Run time of each job is known


Workload Assumptions

1. ~~Each job runs for the same time~~
2. ~~All jobs arrive at the same time~~
3. Once started, each job runs to completion
4. All jobs use only the CPU
5. Run time of each job is known


Shortest Job First (SJF) scheduling

0 10 20 30 40 50


Shortest Job First (SJF) scheduling




Shortest Job First (SJF) scheduling




Shortest Job First (SJF) scheduling



Shortest Job First (SJF) scheduling



Shortest Job First (SJF) scheduling



Shortest Job First (SJF) scheduling



$$\text{Avg. Turnaround Time} = \frac{(30 + (40 - 10) + (50 - 10))}{3} \sim 33$$

Workload Assumptions

1. ~~Each job runs for the same time~~
2. ~~All jobs arrive at the same time~~
3. Once started, each job runs to completion (Pre-emptible)
4. All jobs use only the CPU
5. Run time of each job is known

Workload Assumptions

1. ~~Each job runs for the same time~~
2. ~~All jobs arrive at the same time~~
3. ~~Once started, each job runs to completion (Pre-emptible)~~
4. All jobs use only the CPU
5. Run time of each job is known

500 Mile Email Problem

Subject: The case of the 500-mile email (was RE:
[SAGE] Favorite impossible
task?)

Here's a problem that *sounded* impossible... I almost regret posting the story to a wide audience, because it makes a great tale over drinks at a conference. :-) The story is slightly altered in order to protect the guilty, elide over irrelevant and boring details, and generally make the whole thing more entertaining.

500 Mile Email Problem

I was working in a job running the campus email system some years ago when I got a call from the chairman of the statistics department.

"We're having a problem sending email out of the department."

"What's the problem?" I asked.

"We can't send mail more than 500 miles," the chairman explained.

I choked on my latte. "Come again?"

500 Mile Email Problem

"We can't send mail farther than 500 miles from here," he repeated. "A little bit more, actually. Call it 520 miles. But no farther."

"Um... Email really doesn't work that way, generally," I said, trying to keep panic out of my voice. One doesn't display panic when speaking to a department chairman, even of a relatively impoverished department like statistics. "What makes you think you can't send mail more than 500 miles?"

500 Mile Email Problem

"It's not what I *think*," the chairman replied testily. "You see, when we first noticed this happening, a few days ago--"

"You waited a few DAYS?" I interrupted, a tremor tingling my voice. "And you couldn't send email this whole time?"

"We could send email. Just not more than--"

--500 miles, yes," I finished for him, "I got that. But why didn't you call earlier?"

"Well, we hadn't collected enough data to be sure of what was going on until just now." Right. This is the chairman of *statistics*. "Anyway, I asked one of the geostatisticians to look into it--" 24

500 Mile Email Problem

"Geostatisticians..."

"--yes, and she's produced a map showing the radius within which we can send email to be slightly more than 500 miles. There are a number of destinations within that radius that we can't reach, either, or reach sporadically, but we can never email farther than this radius."

"I see," I said, and put my head in my hands. "When did this start? A few days ago, you said, but did anything change in your systems at that time?"

500 Mile Email Problem

"Well, the consultant came in and patched our server and rebooted it.

But I called him, and he said he didn't touch the mail system."

"Okay, let me take a look, and I'll call you back," I said, scarcely believing that I was playing along. It wasn't April Fool's Day. I tried to remember if someone owed me a practical joke.

I logged into their department's server, and sent a few test mails. This was in the Research Triangle of North Carolina, and a test mail to my own account was delivered without a hitch. Ditto for one sent to Richmond, and Atlanta, and Washington. Another to Princeton (400 miles) worked.

500 Mile Email Problem

But then I tried to send an email to Memphis (600 miles). It failed.

Boston, failed. Detroit, failed. I got out my address book and started trying to narrow this down. New York (420 miles) worked, but Providence (580 miles) failed.

I was beginning to wonder if I had lost my sanity. I tried emailing a friend who lived in North Carolina, but whose ISP was in Seattle.

Thankfully, it failed. If the problem had had to do with the geography of the human recipient and not his mail server, I think I would have broken down in tears.

500 Mile Email Problem

Having established that--unbelievably--the problem as reported was true, and repeatable, I took a look at the `sendmail.cf` file. It looked fairly normal. In fact, it looked familiar.

I diffed it against the `sendmail.cf` in my home directory. It hadn't been altered--it was a `sendmail.cf` I had written. And I was fairly certain I hadn't enabled the `"FAIL_MAIL_OVER_500_MILES"` option. At a loss, I telnetted into the SMTP port. The server happily responded with a SunOS `sendmail` banner.

500 Mile Email Problem

Wait a minute... a SunOS sendmail banner? At the time, Sun was still shipping Sendmail 5 with its operating system, even though Sendmail 8 was fairly mature. Being a good system administrator, I had standardized on Sendmail 8. And also being a good system administrator, I had written a `sendmail.cf` that used the nice long self-documenting option and variable names available in Sendmail 8 rather than the cryptic punctuation-mark codes that had been used in Sendmail 5.

500 Mile Email Problem

The pieces fell into place, all at once, and I again choked on the dregs of my now-cold latte. When the consultant had "patched the server," he had apparently upgraded the version of SunOS, and in so doing *downgraded* Sendmail. The upgrade helpfully left the `sendmail.cf` alone, even though it was now the wrong version.

It so happens that Sendmail 5--at least, the version that Sun shipped, which had some tweaks--could deal with the Sendmail 8 `sendmail.cf`, as most of the rules had at that point remained unaltered. But the new long configuration options--those it saw as junk, and skipped. And the sendmail binary had no defaults compiled in for most of these, so, finding no suitable settings in the `sendmail.cf` file, they were set to zero.

500 Mile Email Problem

One of the settings that was set to zero was the timeout to connect to the remote SMTP server. Some experimentation established that on this particular machine with its typical load, a zero timeout would abort a connect call in slightly over three milliseconds.

An odd feature of our campus network at the time was that it was 100% switched. An outgoing packet wouldn't incur a router delay until hitting the POP and reaching a router on the far side. So time to connect to a lightly-loaded remote host on a nearby network would actually largely be governed by the speed of light distance to the destination rather than by incidental router delays.

500 Mile Email Problem

Feeling slightly giddy, I typed into my shell:

```
$ units
```

```
1311 units, 63 prefixes
```

```
You have: 3 millilightseconds
```

```
You want: miles
```

```
    * 558.84719
```


```
    / 0.0017893979
```

"500 miles, or a little bit more."


Shortest Time to Completion (SCTF) scheduling

0 10 20 30 40 50


Shortest Time to Completion (SCTF) scheduling




Shortest Time to Completion (SCTF) scheduling




Shortest Time to Completion (SCTF) scheduling




Shortest Time to Completion (SCTF) scheduling




Shortest Time to Completion (SCTF) scheduling




Shortest Time to Completion (SCTF) scheduling



Shortest Time to Completion (SCTF) scheduling




Shortest Time to Completion (SCTF) scheduling



$$\text{Avg. Turnaround Time} = \frac{(50 + (20 - 10) + (30 - 10))}{3} \sim 27$$

Metric - Response Time




Metric - Response Time




$$\text{Avg. Turnaround Time} = (10+20+50)/3 \sim 27$$


Metric - Response Time



Metric - Response Time




Metric - Response Time




- Arrives at 0

Metric - Response Time




- Arrives at 0
- Runs for first time at 0

Metric - Response Time




- Arrives at 0
- Runs for first time at 0
- Response time = 0


Metric - Response Time



Metric - Response Time




Metric - Response Time




- Arrives at 0

Metric - Response Time




- Arrives at 0
- Runs for first time at 10

Metric - Response Time




- Arrives at 0
- Runs for first time at 10
- Response time = 10


Metric - Response Time



Metric - Response Time




Metric - Response Time




- Arrives at 0

Metric - Response Time




- Arrives at 0
- Runs for first time at 20

Metric - Response Time




- Arrives at 0
- Runs for first time at 20
- Response time = 20

Metric - Response Time




Metric - Response Time




$$\text{Avg. Response Time} = (0+10+20)/3 \sim 10$$


Metric - Response Time



Metric - Response Time




Metric - Response Time




$$\text{Avg. Response Time} = \frac{(0+10+40)}{3} \sim 17$$


Metric - Response Time



Metric - Response Time




Metric - Response Time




$$\text{Avg. Response Time} = \frac{(0+30+40)}{3} \sim 23$$

Metric - Response Time




Metric - Response Time



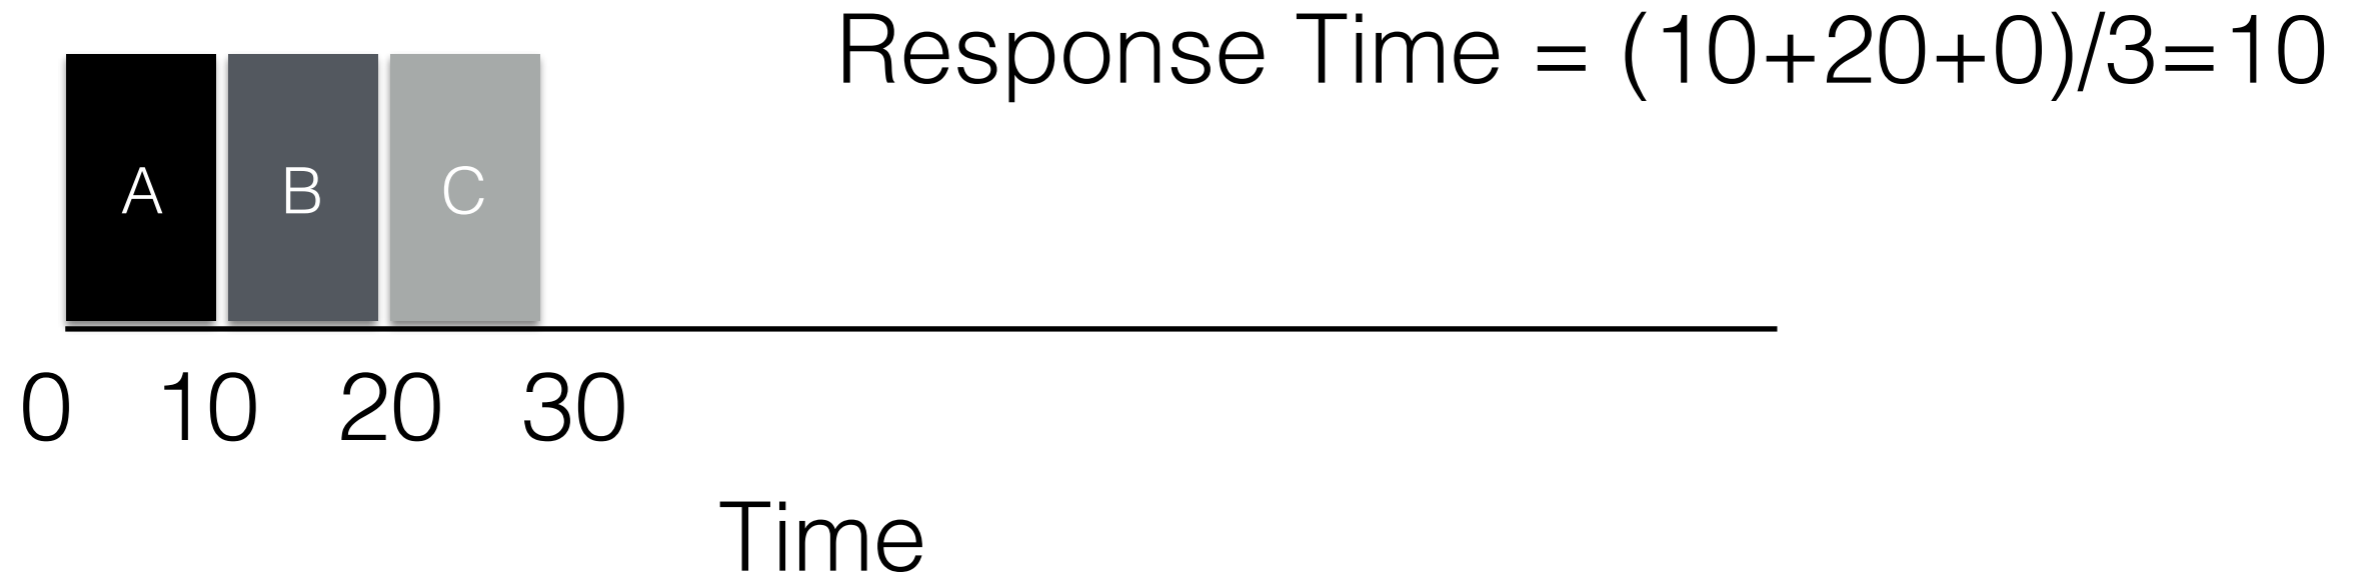
None of the previous policies particularly good for response time

Round Robin


Round Robin



Round Robin



Round Robin




$$\text{Response Time} = (10+20+0)/3=10$$

$$\text{Turnaround Time} = (10+20+30)/3=20$$


Time

Round Robin




$$\text{Response Time} = (10+20+0)/3=10$$

$$\text{Turnaround Time} = (10+20+30)/3=20$$




Time

Round Robin




$$\text{Response Time} = (10+20+0)/3=10$$

$$\text{Turnaround Time} = (10+20+30)/3=20$$




$$\text{Response Time} = (0+5+10)/3=5$$

Round Robin



$$\text{Response Time} = (10+20+0)/3=10$$
$$\text{Turnaround Time} = (10+20+30)/3=20$$



$$\text{Response Time} = (0+5+10)/3=5$$
$$\text{Turnaround Time} = (20+25+30)/3=25$$