

Concurrency

Introduction to Computer Systems, Fall 2024

Instructors: Joel Ramirez Travis McGaha

Head TAs: Adam Gorka Daniel Gearhardt
Ash Fujiyama Emily Shen

TAs:

Ahmed Abdellah

Ethan Weisberg

Maya Huizar

Angie Cao

Garrett O'Malley Kirsch

Meghana Vasireddy

August Fu

Hassan Rizwan

Perrie Quek

Caroline Begg

Iain Li

Sidharth Roy

Cathy Cao

Jerry Wang

Sydney-Shea Cohen

Claire Lu

Juan Lopez

Vivi Li

Eric Sungwon Lee

Keith Mathe

Yousef AlRabiah



pollev.com/cis2400

- ❖ I hope you were able to enjoy your breaks!
- ❖ Anything exciting happen? Any really good food made?

What is a Program?

- ❖ A "program" is a set of instructions, essentially a *static* file containing code.

```
void answer_emails() {  
    // I'm Jeff Besos  
    // My Inbox has 1,000,000 emails  
    for (auto& email : inbox) {  
        email.send("Sorry, I'm on my yacht...");  
    }  
}  
  
int main(int argc, char* argv[]) {  
    answer_emails();  
    return 0;  
}                                     auto_responder.c
```

It's just text...
nothing special about it.

What is a Program *in execution*?

- ❖ A "program" in *execution* is called a *process*.

```

void answer_emails() {
    // I'm Jeff Besos
    // My Inbox has 1,000,000 emails
    for (auto& email : inbox) {
        email.send("Sorry, I'm on my yacht...");
    }
}

int main(int argc, char* argv[]) {
    answer_emails();
    return 0;
}
    
```

auto_responder.c



The compiled Instructions executed on the processor

What does the process consist of?

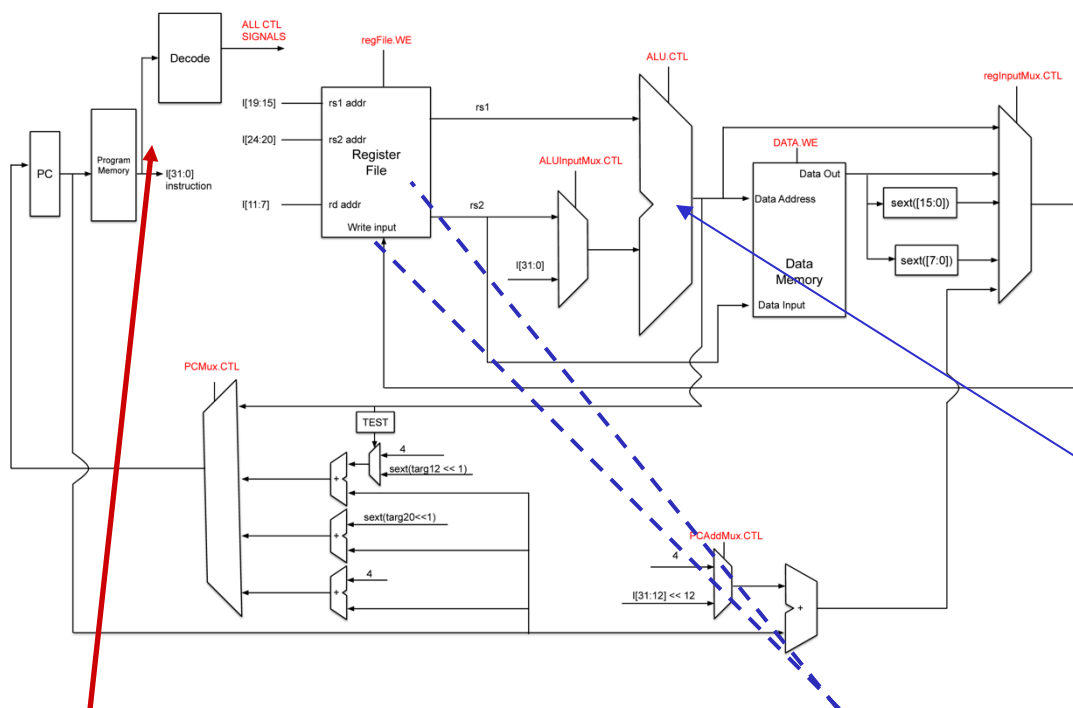
memory, instructions, registers, and other state involved in program execution.

What is *necessary* to run a process?



- ❖ You need a CPU with at least one core!
- ❖ What's a core?

ITS THIS! (ESSENTIALLY)



Execution Unit (ALU)

Fetch/Decode Instructions

Register State/Memory State

What is *necessary* to run a process?



- ❖ You need a CPU with a single core!
- ❖ Single Core
 - Fetch/Decode, Register/Memory, Execution Unit (ALU)
 - *Fundamental unit of systems hardware*

Truth: most things aren't singly cored anymore

- ❖ Does anyone know the # of cores in an Intel i9 CPU?
 - (the ceo just quit yesterday btw. Company isn't doing well I hear)
 - We'll leave that to the Wharton people

What is *necessary* to run a process?

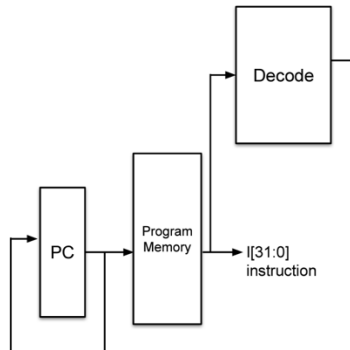


- ❖ You need a CPU with a single core!
- ❖ Single Core
 - Fetch/Decode, Register/Memory, Execution Unit (ALU)
 - *Fundamental unit of systems hardware*

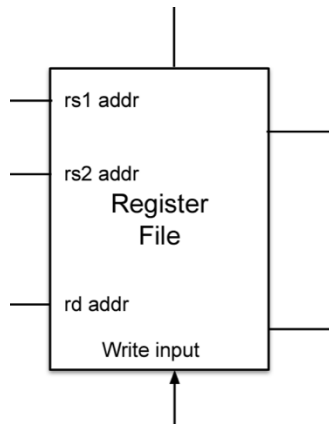
Truth: most things aren't singly cored anymore

- ❖ Does anyone know the # of cores in an Intel i9 CPU?
 - Trick question; *depends on the model.*
 - Intel® Core™ i9-7900X X-series Processor
 - *10 Cores!*
 - *Highest model has 18 cores!*


A Single Core



Fetch and decode instructions.




Register Set and Memory State



Execution Unit (ALU)

Let's focus on just one core for now

- ❖ With one core, we can run one process!
 - let's open up chrome
- ❖ **Question: what if you want to open 2 more applications?**



Let's focus on just one core for now

- ❖ We want to run 3 things on a single core processor.
- ❖ **Things can not run *via parallelism* (simultaneously).**
 - *Why? We only have ONE CORE. Only one ALU to go around.*



running chrome



running spotify



THIS EXAMPLE IS NOT POSSIBLE

Let's focus on just one core for now

- ❖ We want to run 3 things on a single core processor.
- ❖ **Our Solution: Concurrency**
 - we can switch between different processes



running chrome



Let's focus on just one core for now

- ❖ We want to run 3 things on a single core processor.
- ❖ **Our Solution: Concurrency**
 - we can switch between different processes



running spotify

Let's focus on just one core for now

- ❖ We want to run 3 things on a single core processor.
- ❖ **Our Solution: Concurrency**
 - we can switch between different processes



running discord



Let's focus on just one core for now

- ❖ We want to run 3 things on a single core processor.
- ❖ **Our Solution: Concurrency**
 - Make it seem like things are running *simultaneously*



running spotify

Let's focus on just one core for now

- ❖ We want to run 3 things on a single core processor.
- ❖ **Our Solution: Concurrency**
 - Make it seem like things are running *simultaneously*



running chrome



What does concurrency look like?



Chrome Instructions

```

chrome_http_response:
pc → sw a0, -12(s0)
    li a1, 1
    sw a1, -28(s0)
    li a1, 2
    sw a1, -32(s0)
    sh a0, -24(s0)
    lui a0, 2
    addi a0, a0, -112
    sw a0, -36(s0)
    call bind
    lw a1, -36(s0)
    sh a0, -24(s0)
    lui a0, %hi(.L.str)
    addi a0, a0,
    %lo(.L.str)
  
```



Spotify Instructions

```

spotify_gen_wav:
pc → li a1, 10
    mul a0, a0, a1
    fcvt.d.w ft0, a0
    lui a0, %hi(.LCPI2_0)
    fld ft1, %lo(.LCPI2_0)(a0)
    fdiv.d fa0, ft0, ft1
    call sin
    fsd fa0, -24(s0)
    lw a1, -12(s0)
    fld ft0, -24(s0)
    lui a0, %hi(.L.str.7)
    addi a0, a0,
    %lo(.L.str.7)
  
```

running chrome



*pc designates the next inst. to exec here.

What does concurrency look like?



Chrome Instructions

```

chrome_http_response:
  sw a0, -12(s0)
pc → li a1, 1
  sw a1, -28(s0)
  li a1, 2
  sw a1, -32(s0)
  sh a0, -24(s0)
  lui a0, 2
  addi a0, a0, -112
  sw a0, -36(s0)
  call bind
  lw a1, -36(s0)
  sh a0, -24(s0)
  lui a0, %hi(.L.str)
  addi a0, a0,
  %lo(.L.str)

```

running chrome



Spotify Instructions

```

spotify_gen_wav:
pc → li a1, 10
  mul a0, a0, a1
  fcvt.d.w ft0, a0
  lui a0, %hi(.LCPI2_0)
  fld ft1, %lo(.LCPI2_0)(a0)
  fdiv.d fa0, ft0, ft1
  call sin
  fsd fa0, -24(s0)
  lw a1, -12(s0)
  fld ft0, -24(s0)
  lui a0, %hi(.L.str.7)
  addi a0, a0,
  %lo(.L.str.7)

```



*pc designates the next inst. to exec here.

What does concurrency look like?



Chrome Instructions

```
chrome_http_response:
    sw a0, -12(s0)
    li a1, 1
pc → sw a1, -28(s0)
    li a1, 2
    sw a1, -32(s0)
    sh a0, -24(s0)
    lui a0, 2
    addi a0, a0, -112
    sw a0, -36(s0)
    call bind
    lw a1, -36(s0)
    sh a0, -24(s0)
    lui a0, %hi(.L.str)
    addi a0, a0,
    %lo(.L.str)
```



Spotify Instructions

```
spotify_gen_wav:
pc → li a1, 10
    mul a0, a0, a1
    fcvt.d.w ft0, a0
    lui a0, %hi(.LCPI2_0)
    fld ft1, %lo(.LCPI2_0)(a0)
    fdiv.d fa0, ft0, ft1
    call sin
    fsd fa0, -24(s0)
    lw a1, -12(s0)
    fld ft0, -24(s0)
    lui a0, %hi(.L.str.7)
    addi a0, a0,
    %lo(.L.str.7)
```

running chrome



*pc designates the next inst. to exec here.

What does concurrency look like?



Chrome Instructions

```
chrome_http_response:
    sw a0, -12(s0)
    li a1, 1
    sw a1, -28(s0)
    pc → li a1, 2
    sw a1, -32(s0)
    sh a0, -24(s0)
    lui a0, 2
    addi a0, a0, -112
    sw a0, -36(s0)
    call bind
    lw a1, -36(s0)
    sh a0, -24(s0)
    lui a0, %hi(.L.str)
    addi a0, a0,
    %lo(.L.str)
```



Spotify Instructions

```
spotify_gen_wav:
    pc → li a1, 10
    mul a0, a0, a1
    fcvt.d.w ft0, a0
    lui a0, %hi(.LCPI2_0)
    fld ft1, %lo(.LCPI2_0)(a0)
    fdiv.d fa0, ft0, ft1
    call sin
    fsd fa0, -24(s0)
    lw a1, -12(s0)
    fld ft0, -24(s0)
    lui a0, %hi(.L.str.7)
    addi a0, a0,
    %lo(.L.str.7)
```

running chrome



*pc designates the next inst. to exec here.

What does concurrency look like?



Chrome Instructions

```
chrome_http_response:
    sw a0, -12(s0)
    li a1, 1
    sw a1, -28(s0)
    li a1, 2
    sw a1, -32(s0)
    sh a0, -24(s0)
    lui a0, 2
    addi a0, a0, -112
    sw a0, -36(s0)
    call bind
    lw a1, -36(s0)
    sh a0, -24(s0)
    lui a0, %hi(.L.str)
    addi a0, a0,
    %lo(.L.str)
```

pc 



Spotify Instructions

```
spotify_gen_wav:
    li a1, 10
    mul a0, a0, a1
    fcvt.d.w ft0, a0
    lui a0, %hi(.LCPI2_0)
    fld ft1, %lo(.LCPI2_0)(a0)
    fdiv.d fa0, ft0, ft1
    call sin
    fsd fa0, -24(s0)
    lw a1, -12(s0)
    fld ft0, -24(s0)
    lui a0, %hi(.L.str.7)
    addi a0, a0,
    %lo(.L.str.7)
```

pc 

running chrome



*pc designates the next inst. to exec here.

What does concurrency look like?



Chrome Instructions

```
chrome_http_response:
    sw a0, -12(s0)
    li a1, 1
    sw a1, -28(s0)
    li a1, 2
    sw a1, -32(s0)
    sh a0, -24(s0)
    lui a0, 2
    addi a0, a0, -112
    sw a0, -36(s0)
    call bind
    lw a1, -36(s0)
    sh a0, -24(s0)
    lui a0, %hi(.L.str)
    addi a0, a0,
    %lo(.L.str)
```

pc →



Spotify Instructions

```
spotify_gen_wav:
    li a1, 10
    mul a0, a0, a1
    fcvt.d.w ft0, a0
    lui a0, %hi(.LCPI2_0)
    fld ft1, %lo(.LCPI2_0)(a0)
    fdiv.d fa0, ft0, ft1
    call sin
    fsd fa0, -24(s0)
    lw a1, -12(s0)
    fld ft0, -24(s0)
    lui a0, %hi(.L.str.7)
    addi a0, a0,
    %lo(.L.str.7)
```

pc →

running chrome



*pc designates the next inst. to exec here.

What does concurrency look like?



Chrome Instructions

```
chrome_http_response:
  sw a0, -12(s0)
  li a1, 1
  sw a1, -28(s0)
  li a1, 2
  sw a1, -32(s0)
  sh a0, -24(s0)
  lui a0, 2
  addi a0, a0, -112
  sw a0, -36(s0)
  call bind
  lw a1, -36(s0)
  sh a0, -24(s0)
  lui a0, %hi(.L.str)
  addi a0, a0,
  %lo(.L.str)
```

pc →



Spotify Instructions

```
spotify_gen_wav:
  li a1, 10
  mul a0, a0, a1
  fcvt.d.w ft0, a0
  lui a0, %hi(.LCPI2_0)
  fld ft1, %lo(.LCPI2_0)(a0)
  fdiv.d fa0, ft0, ft1
  call sin
  fsd fa0, -24(s0)
  lw a1, -12(s0)
  fld ft0, -24(s0)
  lui a0, %hi(.L.str.7)
  addi a0, a0,
  %lo(.L.str.7)
```

pc →

running chrome



*pc designates the next inst. to exec here.

What does concurrency look like?



Chrome Instructions

```

chrome_http_response:
    sw a0, -12(s0)
    li a1, 1
    sw a1, -28(s0)
    li a1, 2
    sw a1, -32(s0)
    sh a0, -24(s0)
    lui a0, 2
    addi a0, a0, -112
    sw a0, -36(s0)
    call bind
    lw a1, -36(s0)
    sh a0, -24(s0)
    lui a0, %hi(.L.str)
    addi a0, a0,
    %lo(.L.str)
  
```

pc →



Spotify Instructions

```

spotify_gen_wav:
    li a1, 10
    mul a0, a0, a1
    fcvt.d.w ft0, a0
    lui a0, %hi(.LCPI2_0)
    fld ft1, %lo(.LCPI2_0)(a0)
    fdiv.d fa0, ft0, ft1
    call sin
    fsd fa0, -24(s0)
    lw a1, -12(s0)
    fld ft0, -24(s0)
    lui a0, %hi(.L.str.7)
    addi a0, a0,
    %lo(.L.str.7)
  
```

pc →



running spotify

*pc designates the next inst. to exec here.

What does concurrency look like?



Chrome Instructions

```

chrome_http_response:
  sw a0, -12(s0)
  li a1, 1
  sw a1, -28(s0)
  li a1, 2
  sw a1, -32(s0)
  sh a0, -24(s0)
  lui a0, 2
  addi a0, a0, -112
  sw a0, -36(s0)
  call bind
  lw a1, -36(s0)
  sh a0, -24(s0)
  lui a0, %hi(.L.str)
  addi a0, a0,
  %lo(.L.str)
  
```

pc →



Spotify Instructions

```

spotify_gen_wav:
  li a1, 10
  mul a0, a0, a1
  fcvt.d.w ft0, a0
  lui a0, %hi(.LCPI2_0)
  fld ft1, %lo(.LCPI2_0)(a0)
  fdiv.d fa0, ft0, ft1
  call sin
  fsd fa0, -24(s0)
  lw a1, -12(s0)
  fld ft0, -24(s0)
  lui a0, %hi(.L.str.7)
  addi a0, a0,
  %lo(.L.str.7)
  
```

pc →



running spotify

*pc designates the next inst. to exec here.

What does concurrency look like?



Chrome Instructions

```

chrome_http_response:
  sw a0, -12(s0)
  li a1, 1
  sw a1, -28(s0)
  li a1, 2
  sw a1, -32(s0)
  sh a0, -24(s0)
  lui a0, 2
  addi a0, a0, -112
  sw a0, -36(s0)
  call bind
  lw a1, -36(s0)
  sh a0, -24(s0)
  lui a0, %hi(.L.str)
  addi a0, a0,
  %lo(.L.str)
  
```

pc →



Spotify Instructions

```

spotify_gen_wav:
  li a1, 10
  mul a0, a0, a1
  fcvt.d.w ft0, a0
  lui a0, %hi(.LCPI2_0)
  fld ft1, %lo(.LCPI2_0)(a0)
  fdiv.d fa0, ft0, ft1
  call sin
  fsd fa0, -24(s0)
  lw a1, -12(s0)
  fld ft0, -24(s0)
  lui a0, %hi(.L.str.7)
  addi a0, a0,
  %lo(.L.str.7)
  
```

pc →



running spotify

*pc designates the next inst. to exec here.

What does concurrency look like?



Chrome Instructions

```
chrome_http_response:
    sw a0, -12(s0)
    li a1, 1
    sw a1, -28(s0)
    li a1, 2
    sw a1, -32(s0)
    sh a0, -24(s0)
    lui a0, 2
    addi a0, a0, -112
    sw a0, -36(s0)
    call bind
    lw a1, -36(s0)
    sh a0, -24(s0)
    lui a0, %hi(.L.str)
    addi a0, a0,
    %lo(.L.str)
```

pc →



Spotify Instructions

```
spotify_gen_wav:
    li a1, 10
    mul a0, a0, a1
    fcvt.d.w ft0, a0
    lui a0, %hi(.LCPI2_0)
    fld ft1, %lo(.LCPI2_0)(a0)
    fdiv.d fa0, ft0, ft1
    call sin
    fsd fa0, -24(s0)
    lw a1, -12(s0)
    fld ft0, -24(s0)
    lui a0, %hi(.L.str.7)
    addi a0, a0,
    %lo(.L.str.7)
```

pc →



running spotify

*pc designates the next inst. to exec here.

What does concurrency look like?



Chrome Instructions

```
chrome_http_response:
    sw a0, -12(s0)
    li a1, 1
    sw a1, -28(s0)
    li a1, 2
    sw a1, -32(s0)
    sh a0, -24(s0)
    lui a0, 2
    addi a0, a0, -112
    sw a0, -36(s0)
    call bind
    lw a1, -36(s0)
    sh a0, -24(s0)
    lui a0, %hi(.L.str)
    addi a0, a0,
    %lo(.L.str)
```

pc →



Spotify Instructions

```
spotify_gen_wav:
    li a1, 10
    mul a0, a0, a1
    fcvt.d.w ft0, a0
    lui a0, %hi(.LCPI2_0)
    fld ft1, %lo(.LCPI2_0)(a0)
    fdiv.d fa0, ft0, ft1
    call sin
    fsd fa0, -24(s0)
    lw a1, -12(s0)
    fld ft0, -24(s0)
    lui a0, %hi(.L.str.7)
    addi a0, a0,
    %lo(.L.str.7)
```

pc →



running spotify

*pc designates the next inst. to exec here.

What does concurrency look like?



Chrome Instructions

```
chrome_http_response:
  sw a0, -12(s0)
  li a1, 1
  sw a1, 28(s0)
```



Spotify Instructions

```
spotify_gen_wav:
  li a1, 10
  mul a0, a0, a1
  fsub.d ft0, ft0, a0
```

Only one instruction executes at a time on a single-core processor.

To allow multiple processes to run, the operating system takes turns giving each process access to the CORE, one at a time.

```
sw a0, -56(s0)
call bind
lw a1, -36(s0)
sh a0, -24(s0)
lui a0, %hi(.L.str.7)
addi a0, a0, %lo(.L.str.7)
```

```
lw a1, -12(s0)
fld ft0, -24(s0)
lui a0, %hi(.L.str.7)
addi a0, a0, %lo(.L.str.7)
```

**Note: Do not ask how this is done today...
take OS with us next semester... 😊**



Let's bring back this code

```
void answer_emails() {  
    // I'm Jeff Besos  
    // My Inbox has 1,000,000 emails  
    for (auto& email : inbox) {  
        email.send("Sorry, I'm on my yacht...");  
    }  
}  
  
int main(int argc, char* argv[]) {  
    answer_emails();  
    return 0;  
}  
  
auto_responder.c
```

Looping:

keeping track of email
loading email into mem

Send:

creates email
formats data for response
connection to email server
send over network!

Most of these operations are relatively quick!

Except for one...

Connecting to an email server and sending an email can easily take 50+ ms.

Time Analysis

```

void answer_emails() {
    // I'm Jeff Besos
    // My Inbox has 1,000,000 emails
    for (auto& email : inbox) {
        email.send("Sorry, I'm on my yacht...");
    }
}

int main(int argc, char* argv[]) {
    answer_emails();
    return 0;
}
auto_responder.c

```

Looping:

keeping track of email
loading email into mem

Send:

creates email
formats data for response
connection to email server
send over network!

Total Elapsed Time

10 ms

Cal Current Email
Loading Email
invoking send()



150 ms

Time it takes to send the email over the network

This means *calling the function*, not *doing the sending over the network*

pollev.com/cis2400

What's so wrong with this? Discuss!

```
void answer_emails() {  
    // My Inbox has 1,000,000 emails  
    for (auto& email : inbox) {  
        email.send("Sorry, I'm out of town...");  
    }  
}  
  
int main(int argc, char* argv[]) {  
    answer_emails();  
    return 0;  
}  
  
auto_responder.c
```

10 ms

Cal Current Email
Loading Email
invoking send()

150 ms

Time it takes to send the email
over the network

Put your short responses into poll everywhere 😊

pollev.com/cis2400

What do we spend most of our time doing?

```
void answer_emails() {  
  // My Inbox has 1,000,000 emails  
  for (auto& email : inbox) {  
    email.send("Sorry, I'm out of town...");  
  }  
}  
  
int main(int argc, char* argv[]) {  
  answer_emails();  
  return 0;  
}                                     auto_responder.c
```

10 ms

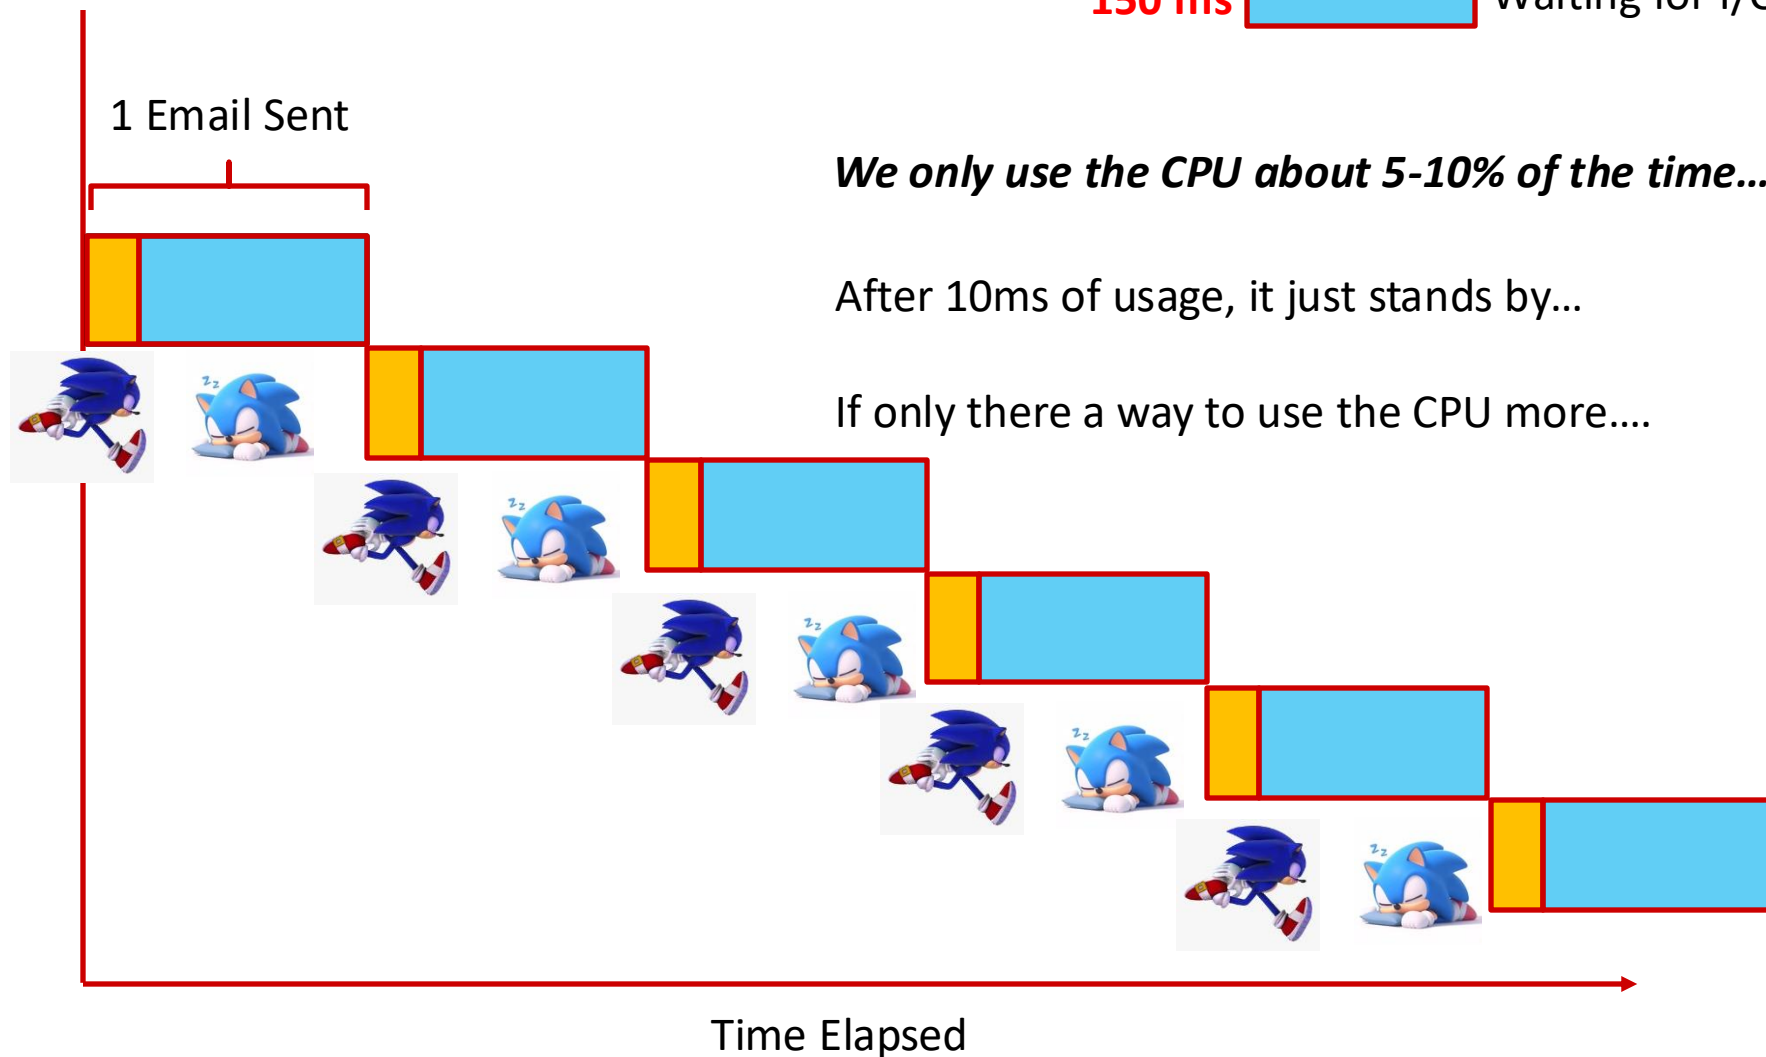

Cal Current Email
Loading Email
invoking send()

150 ms

Time it takes to send the email
over the network

- ❖ We spend more time in the I/O operations
 - Establishing Connection with Email Server
 - Sending Email Over Network

Let's Visualize the Bottleneck



*yes it's not to scale

CPU Utilization


- ❖ When a process waits for I/O, the CPU remains idle, wasting valuable processing time.
- ❖ Our goal is to **maximize CPU utilization** and ensure it stays actively engaged in useful work.
- ❖ What if we could send the next email while waiting for the network I/O?



Before Threads, There Were Processes

- ❖ A "program" in *execution* is called a *process*.

Memory Layout (State)




Instructions Executed

```
pc   ➔ send:  
    addi sp, sp, -16  
    sw ra, 12(sp)  
    sw s0, 8(sp)  
    addi s0, sp, 16  
    // loop omitted  
    call send_smtp_data  
    // exit loop  
    lw ra, 12(sp)  
    lw s0, 8(sp)  
    addi sp, sp, 16  
    jalr x0, ra, 0  
  
main:  
    //  
    //calls send
```

Sharing the CPU: It's like a microwave

- ❖ When everyone wants to make cup of ramen at 2AM in your dorm, you probably *share* a microwave
 - (unless you all have a microwave in each of your rooms then this is an example of parallelism via multiple CPU Cores, don't ask...)


Queue to use Microwave




Sharing the CPU: It's like a microwave

- ❖ When everyone wants to make cup of ramen at 2AM in your dorm, you probably *share* a microwave
 - (unless you all have a microwave in each of your rooms then this is an example of parallelism via multiple CPU Cores, don't ask...)

Queue to use Microwave



Waiting For Microwave:



Sharing the CPU: It's like a microwave

- ❖ When everyone wants to make cup of ramen at 2AM in your dorm, you probably *share* a microwave
 - (unless you all have a microwave in each of your rooms then this is an example of parallelism via multiple CPU Cores, don't ask...)


Queue to use Microwave



Waiting For Microwave:



Done:




Sharing the CPU: It's like a microwave

- ❖ When everyone wants to make cup of ramen at 2AM in your dorm, you probably *share* a microwave
 - (unless you all have a microwave in each of your rooms then this is an example of parallelism via multiple CPU Cores, don't ask...)


Queue to use Microwave



Waiting For Microwave:



Done:



Sharing the CPU: It's like a microwave


- ❖ When everyone wants to make cup of ramen at 2AM in your dorm, you probably *share* a microwave
 - (unless you all have a microwave in each of your rooms then this is an example of parallelism via multiple CPU Cores, don't ask...)

Did you notice how the microwave was always being used?



We achieved 99% Microwave Utilization


Done:



Sharing the Microwave Without Threads

- ❖ When everyone wants to make cup of ramen at 2AM in your dorm, you probably *share* a microwave
 - (unless you all have a microwave in each of your rooms then this is an example of parallelism via multiple CPU Cores, don't ask...)


Queue to use Microwave




Sharing the Microwave Without Threads

- ❖ When everyone wants to make cup of ramen at 2AM in your dorm, you probably *share* a microwave
 - (unless you all have a microwave in each of your rooms then this is an example of parallelism via multiple CPU Cores, don't ask...)

Queue to use Microwave




Waiting For Microwave:




Sharing the Microwave Without Threads

- ❖ When everyone wants to make cup of ramen at 2AM in your dorm, you probably *share* a microwave
 - (unless you all have a microwave in each of your rooms then this is an example of parallelism via multiple CPU Cores, don't ask...)


Queue to use Microwave



Waiting For Microwave:




Wait!!!! You can't use it yet -- it has to cool down!




Sharing the Microwave Without Threads

- ❖ When everyone wants to make cup of ramen at 2AM in your dorm, you probably *share* a microwave
 - (unless you all have a microwave in each of your rooms then this is an example of parallelism via multiple CPU Cores, don't ask...)

Queue to use Microwave



Waiting For Microwave:



Wait!!!! You can't use it yet -- it has to cool down!

So even if the microwave isn't in use
the person takes **ownership** of the microwave and hogs it...


Sharing the Microwave Without Threads

- ❖ As you can see, this is incredibly inefficient...

Queue to use Microwave



Waiting For Microwave:




Done:



you can use it
now Travis...

Process vs Threads?

One Process: `./cook_ramen_together`



The Core

Process vs Threads?

One Process: `./cook_ramen_together`

```
cook_ramen_microwave();  
cook_ramen_microwave();  
cook_ramen_microwave();
```



The Core

Process vs Threads?

One Process: `./cook_ramen_together`

```
for(int i = 0; i < 3; i++){  
    cook_ramen_microwave();  
}
```



The Core

Process vs Threads?

One Process: `./cook_ramen_together`

```
for(int i = 0; i < 3; i++){  
    cook_ramen_microwave();  
}
```

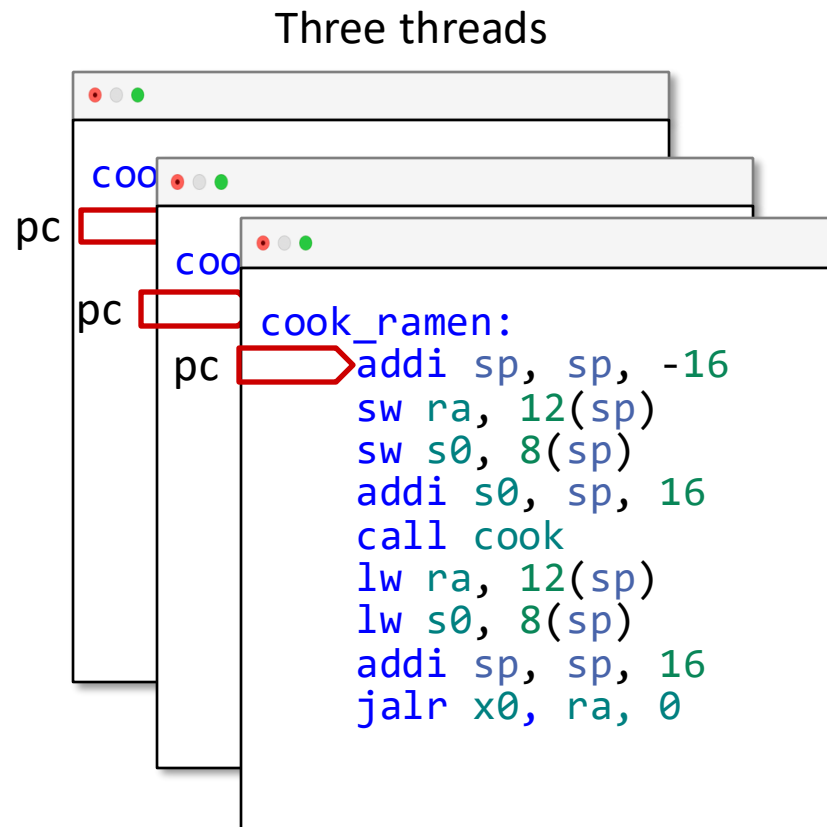
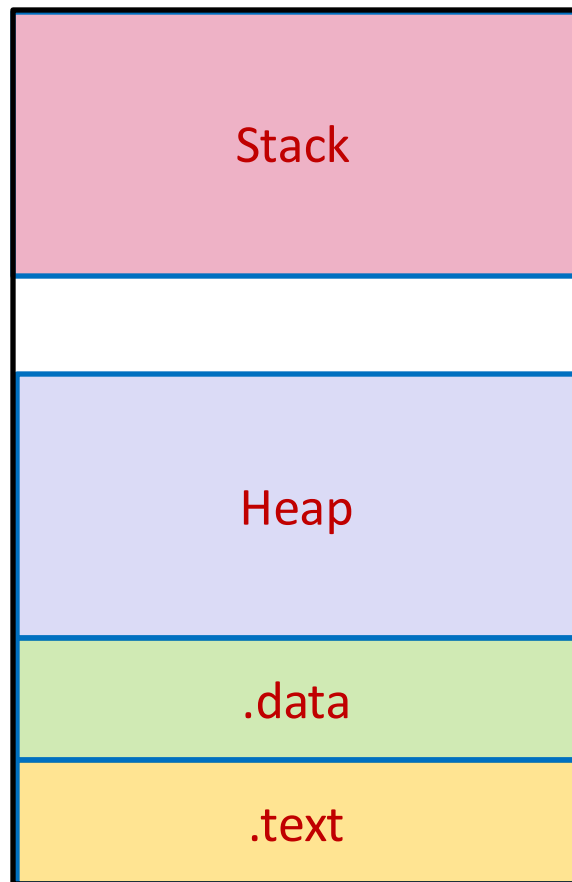


The Core

We all want to cook the ramen –
We just each need to run *our own* `cook_ramen_mico()` function.


Process vs Threads?

One Process: `./cook_ramen_together`




Multi-Threaded Process

Thread One Instructions

```
pc  send:  
addi sp, sp, -16  
sw ra, 12(sp)  
sw s0, 8(sp)  
addi s0, sp, 16  
// loop omitted  
call send_smtp_data  
// exit loop  
lw ra, 12(sp)  
lw s0, 8(sp)  
addi sp, sp, 16  
jalr x0, ra, 0
```

Thread Two Instructions

```
pc  send:  
addi sp, sp, -16  
sw ra, 12(sp)  
sw s0, 8(sp)  
addi s0, sp, 16  
// loop omitted  
call send_smtp_data  
// exit loop  
lw ra, 12(sp)  
lw s0, 8(sp)  
addi sp, sp, 16  
jalr x0, ra, 0
```

Wait, **two program counters in the same process**? Yup! (Don't worry about how this is possible)

Wait, **two copies of the same instructions**? No! (They share this region...)

*important to know that these threads are running in the same *process*

Multi-Threaded Program

Memory Layout (State)



Thread Two
Instructions

```
send:
  addi sp, sp, -16
  sw ra, 12(sp)
  sw s0, 8(sp)
  addi s0, sp, 16
  // loop omitted
  call send_smtp_data
  // exit loop
```

pc

Thread 2
Instructions

```
send:
  addi sp, sp, -16
  sw ra, 12(sp)
  sw s0, 8(sp)
  addi s0, sp, 16
  // loop omitted
  call send_smtp_data
  // exit loop
```

pc

It's just the same instructions in the text segment!
They share this region of memory...

Multi-Threaded Program

Memory Layout (State)



Thread Two
Instructions

```

pc → send:
    addi sp, sp, -16
    sw ra, 12(sp)
    sw s0, 8(sp)
    addi s0, sp, 16
    // loop omitted
    call send_smtp_data
    // exit loop
  
```

Thread 2
Instructions

```

nd:
    addi sp, sp, -16
    sw ra, 12(sp)
    sw s0, 8(sp)
    addi s0, sp, 16
    // loop omitted
    call send_smtp_data
    // exit loop
  
```

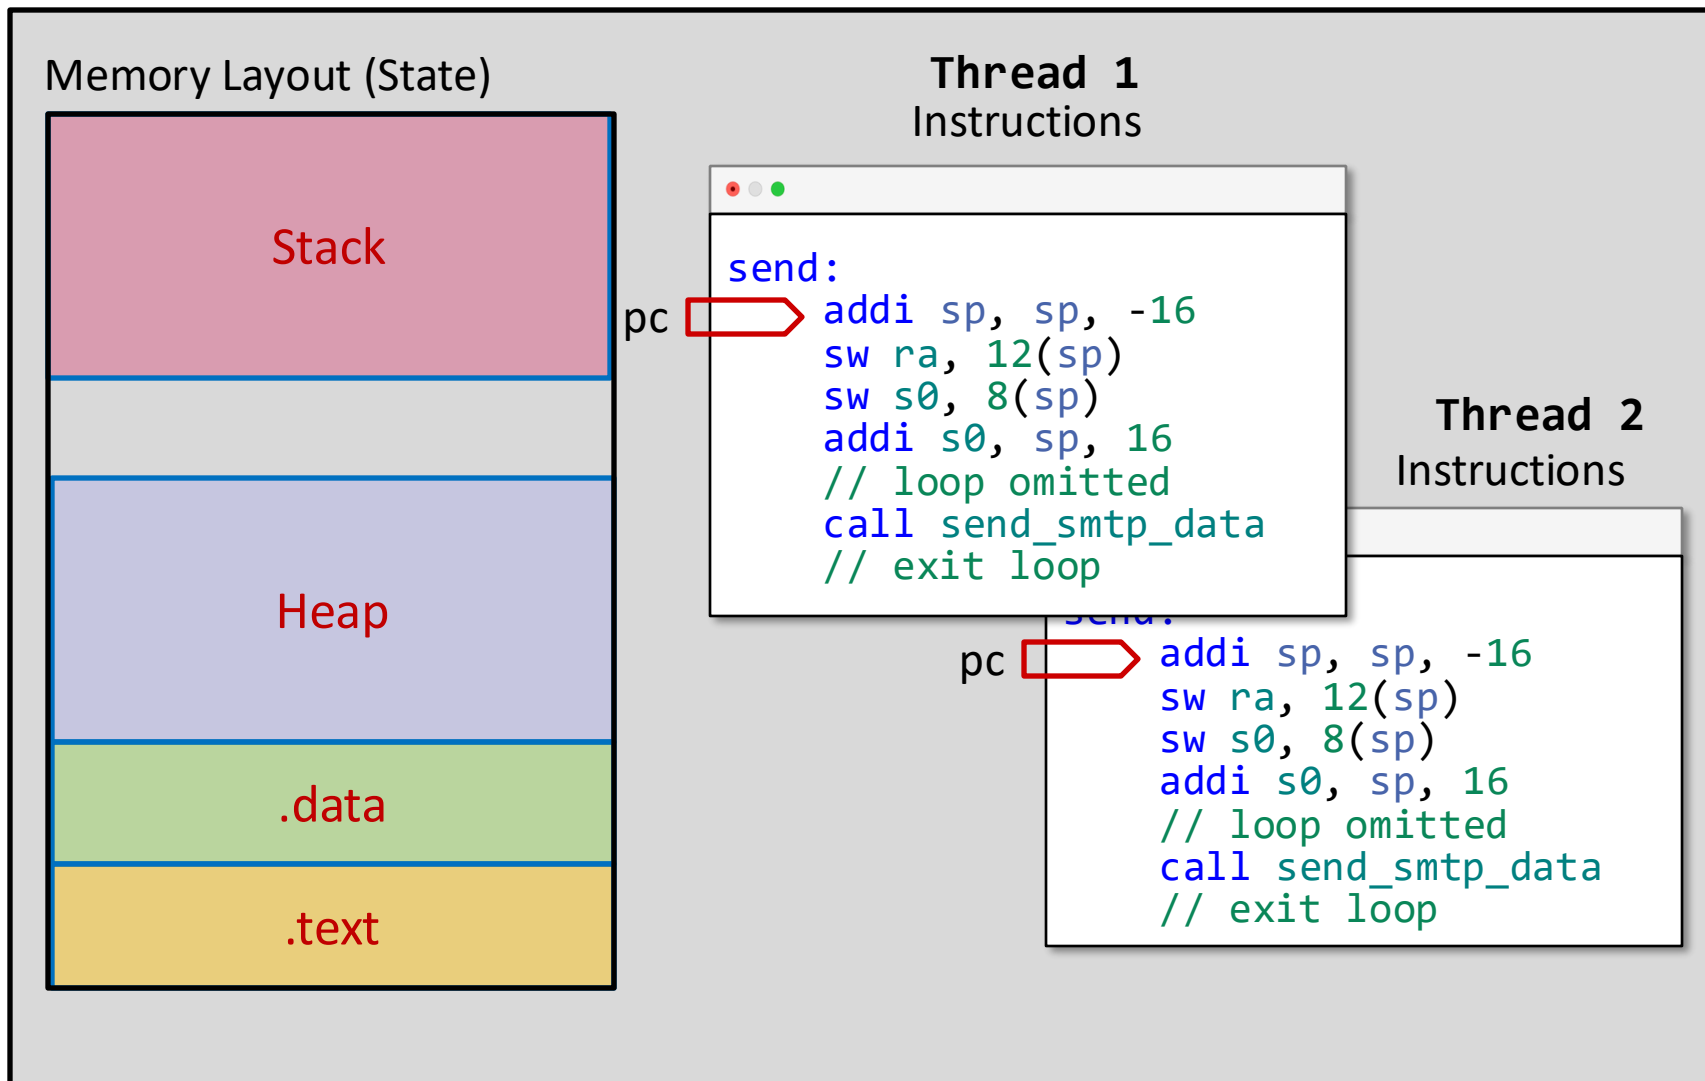
These two threads also share the processor!



These two threads share the entire memory space!


The stack, the heap, data (global vars), and the text!

One Process with Two Threads




Concurrency: Processor Sharing

Thread One Instructions

```
pc  send:  
addi sp, sp, -16  
sw ra, 12(sp)  
sw s0, 8(sp)  
addi s0, sp, 16  
// loop omitted  
call send_smtp_data  
// exit loop  
lw ra, 12(sp)  
lw s0, 8(sp)  
addi sp, sp, 16  
jalr x0, ra, 0
```

Thread Two Instructions

```
pc  send:  
addi sp, sp, -16  
sw ra, 12(sp)  
sw s0, 8(sp)  
addi s0, sp, 16  
// loop omitted  
call send_smtp_data  
// exit loop  
lw ra, 12(sp)  
lw s0, 8(sp)  
addi sp, sp, 16  
jalr x0, ra, 0
```

This thread is waiting for CPU...



Concurrency: Processor Sharing

Thread One Instructions

```
send:  
addi sp, sp, -16  
sw ra, 12(sp)  
sw s0, 8(sp)  
addi s0, sp, 16  
// loop omitted  
call send_smtp_data  
// exit loop  
lw ra, 12(sp)  
lw s0, 8(sp)  
addi sp, sp, 16  
jalr x0, ra, 0
```

Thread Two Instructions

```
send:  
addi sp, sp, -16  
sw ra, 12(sp)  
sw s0, 8(sp)  
addi s0, sp, 16  
// loop omitted  
call send_smtp_data  
// exit loop  
lw ra, 12(sp)  
lw s0, 8(sp)  
addi sp, sp, 16  
jalr x0, ra, 0
```

This thread is waiting for CPU...



Concurrency: Processor Sharing

Thread One Instructions

```
send:  
  addi sp, sp, -16  
  sw ra, 12(sp)  
  sw s0, 8(sp)  
  addi s0, sp, 16  
  // loop omitted  
  call send_smtp_data  
  // exit loop  
  lw ra, 12(sp)  
  lw s0, 8(sp)  
  addi sp, sp, 16  
  jalr x0, ra, 0
```

pc



Thread Two Instructions

```
send:  
  addi sp, sp, -16  
  sw ra, 12(sp)  
  sw s0, 8(sp)  
  addi s0, sp, 16  
  // loop omitted  
  call send_smtp_data  
  // exit loop  
  lw ra, 12(sp)  
  lw s0, 8(sp)  
  addi sp, sp, 16  
  jalr x0, ra, 0
```

pc



This thread is waiting for CPU...



Concurrency: Processor Sharing

Thread One Instructions

```
send:  
  addi sp, sp, -16  
  sw ra, 12(sp)  
  sw s0, 8(sp)  
  addi s0, sp, 16  
  // loop omitted  
  call send_smtp_data  
  // exit loop  
  lw ra, 12(sp)  
  lw s0, 8(sp)  
  addi sp, sp, 16  
  jalr x0, ra, 0
```

pc 

Thread Two Instructions

```
send:  
  addi sp, sp, -16  
  sw ra, 12(sp)  
  sw s0, 8(sp)  
  addi s0, sp, 16  
  // loop omitted  
  call send_smtp_data  
  // exit loop  
  lw ra, 12(sp)  
  lw s0, 8(sp)  
  addi sp, sp, 16  
  jalr x0, ra, 0
```

pc 

This thread is waiting for CPU...



Concurrency: Processor Sharing

Thread One Instructions

```
send:  
  addi sp, sp, -16  
  sw ra, 12(sp)  
  sw s0, 8(sp)  
  addi s0, sp, 16  
  // loop omitted  
pc → call send_smtp_data  
  // exit loop  
  lw ra, 12(sp)  
  lw s0, 8(sp)  
  addi sp, sp, 16  
  jalr x0, ra, 0
```

Thread Two Instructions

```
send:  
pc → addi sp, sp, -16  
  sw ra, 12(sp)  
  sw s0, 8(sp)  
  addi s0, sp, 16  
  // loop omitted  
  call send_smtp_data  
  // exit loop  
  lw ra, 12(sp)  
  lw s0, 8(sp)  
  addi sp, sp, 16  
  jalr x0, ra, 0
```

This thread is waiting for CPU...



Concurrency: Processor Sharing

Thread One Instructions

```
send:
  addi sp, sp, -16
  sw ra, 12(sp)
  sw s0, 8(sp)
  addi s0, sp, 16
  // loop omitted
  call send_smtp_data
  // exit loop
  lw ra, 12(sp)
  lw s0, 8(sp)
  addi sp, sp, 16
  jalr x0, ra, 0
```

Thread Two Instructions

```
send:
  addi sp, sp, -16
  sw ra, 12(sp)
  sw s0, 8(sp)
  addi s0, sp, 16
  // loop omitted
  call send_smtp_data
  // exit loop
  lw ra, 12(sp)
  lw s0, 8(sp)
  addi sp, sp, 16
  jalr x0, ra, 0
```

This thread is waiting for I/O.

Let's switch to the other thread...



Concurrency: Processor Sharing

Thread One Instructions

```
send:
  addi sp, sp, -16
  sw ra, 12(sp)
  sw s0, 8(sp)
  addi s0, sp, 16
  // loop omitted
  call send_smtp_data
  // exit loop
  lw ra, 12(sp)
  lw s0, 8(sp)
  addi sp, sp, 16
  jalr x0, ra, 0
```

pc



Thread Two Instructions

```
send:
  addi sp, sp, -16
  sw ra, 12(sp)
  sw s0, 8(sp)
  addi s0, sp, 16
  // loop omitted
  call send_smtp_data
  // exit loop
  lw ra, 12(sp)
  lw s0, 8(sp)
  addi sp, sp, 16
  jalr x0, ra, 0
```

pc



This thread is waiting for I/O.
Let's switch to the other thread...



Concurrency: Processor Sharing

Thread One Instructions

```

send:
  addi sp, sp, -16
  sw ra, 12(sp)
  sw s0, 8(sp)
  addi s0, sp, 16
  // loop omitted
  call send_smtp_data
  // exit loop
  lw ra, 12(sp)
  lw s0, 8(sp)
  addi sp, sp, 16
  jalr x0, ra, 0
  
```

pc



Thread Two Instructions

```

send:
  addi sp, sp, -16
  sw ra, 12(sp)
  sw s0, 8(sp)
  addi s0, sp, 16
  // loop omitted
  call send_smtp_data
  // exit loop
  lw ra, 12(sp)
  lw s0, 8(sp)
  addi sp, sp, 16
  jalr x0, ra, 0
  
```

pc



This thread is waiting for I/O.
Let's switch to the other thread...



Concurrency: Processor Sharing

Thread One Instructions

```
send:
  addi sp, sp, -16
  sw ra, 12(sp)
  sw s0, 8(sp)
  addi s0, sp, 16
  // loop omitted
  call send_smtp_data
  // exit loop
  lw ra, 12(sp)
  lw s0, 8(sp)
  addi sp, sp, 16
  jalr x0, ra, 0
```

pc



Thread Two Instructions

```
send:
  addi sp, sp, -16
  sw ra, 12(sp)
  sw s0, 8(sp)
  addi s0, sp, 16
  // loop omitted
  call send_smtp_data
  // exit loop
  lw ra, 12(sp)
  lw s0, 8(sp)
  addi sp, sp, 16
  jalr x0, ra, 0
```

pc



This thread is waiting for I/O.
Let's switch to the other thread...




Concurrency: Processor Sharing

Thread One Instructions

```

send:
  addi sp, sp, -16
  sw ra, 12(sp)
  sw s0, 8(sp)
  addi s0, sp, 16
  // loop omitted
  call send_smtp_data
  // exit loop
  lw ra, 12(sp)
  lw s0, 8(sp)
  addi sp, sp, 16
  jalr x0, ra, 0
  
```


pc 

This thread is *waiting for CPU now..*

Thread Two Instructions

```

send:
  addi sp, sp, -16
  sw ra, 12(sp)
  sw s0, 8(sp)
  addi s0, sp, 16
  // loop omitted
  call send_smtp_data
  // exit loop
  lw ra, 12(sp)
  lw s0, 8(sp)
  addi sp, sp, 16
  jalr x0, ra, 0
  
```

pc 

This thread is waiting for I/O.

Let's switch to **thread 1** and see if the I/O is done



Concurrency: Processor Sharing

Thread One Instructions

```

send:
  addi sp, sp, -16
  sw ra, 12(sp)
  sw s0, 8(sp)
  addi s0, sp, 16
  // loop omitted
  call send_smtp_data
  // exit loop
  lw ra, 12(sp)
  lw s0, 8(sp)
  addi sp, sp, 16
  jalr x0, ra, 0
  
```

Thread Two Instructions

```

send:
  addi sp, sp, -16
  sw ra, 12(sp)
  sw s0, 8(sp)
  addi s0, sp, 16
  // loop omitted
  call send_smtp_data
  // exit loop
  lw ra, 12(sp)
  lw s0, 8(sp)
  addi sp, sp, 16
  jalr x0, ra, 0
  
```

pc

pc

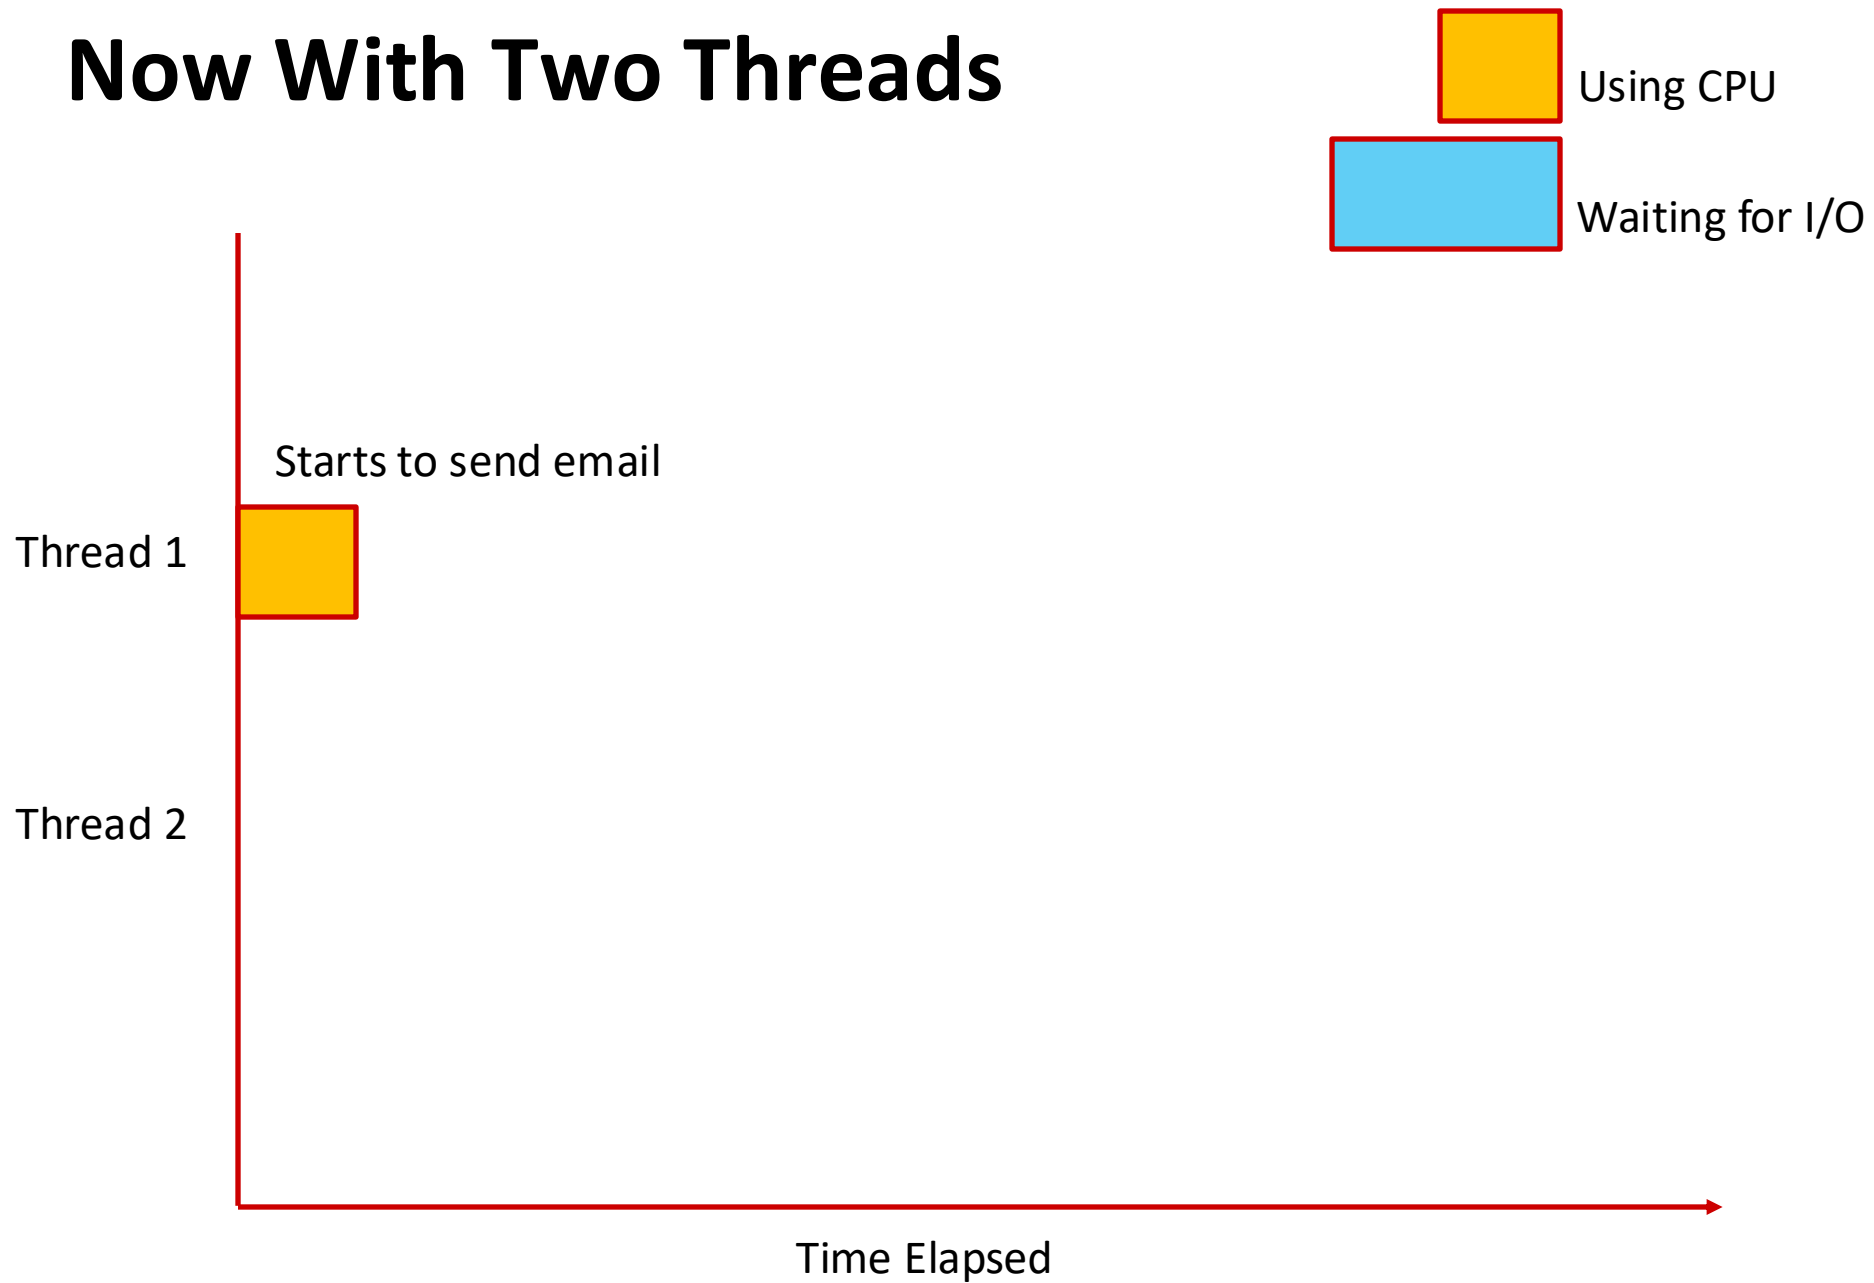
Now this thread continues!
It does the set up to send
another email!



This thread is waiting for
I/O.

*you might notice this looks familiar

Now With Two Threads



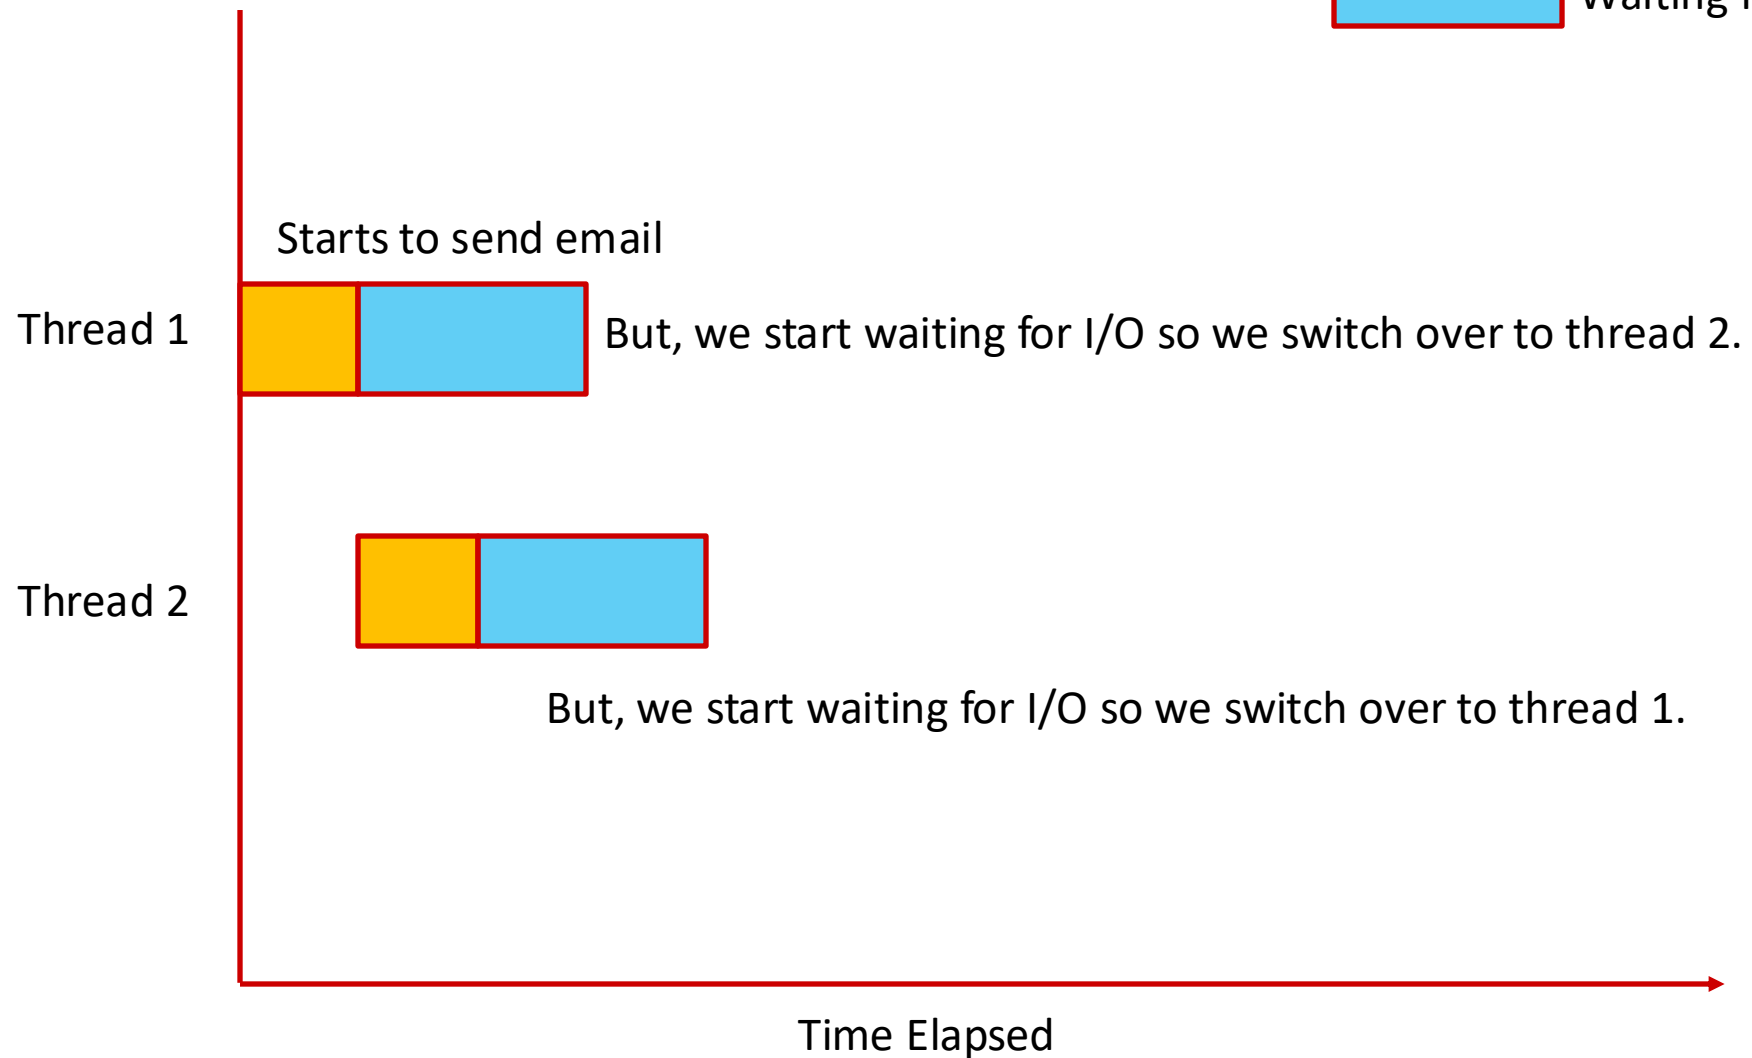
Now With Two Threads



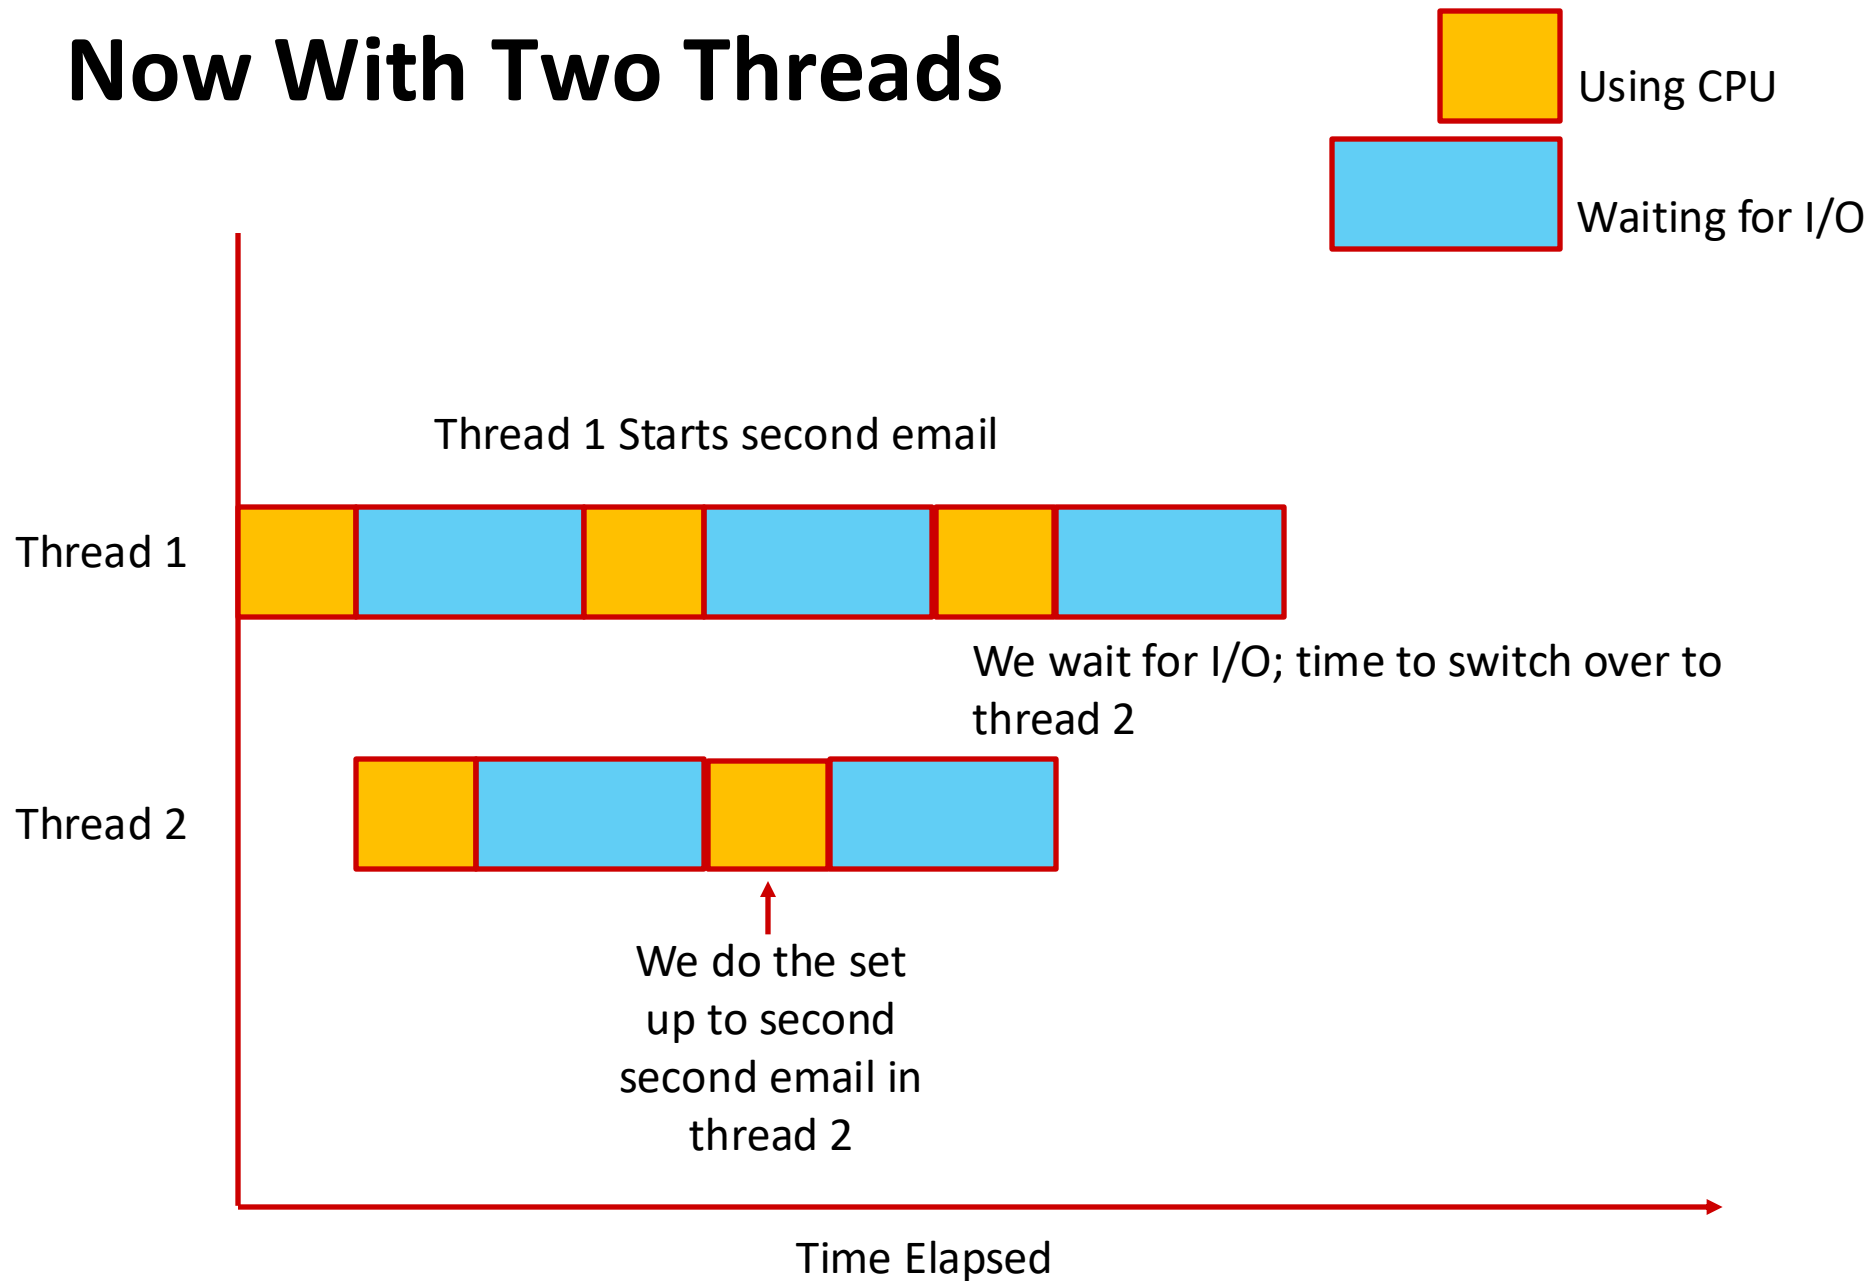
Using CPU



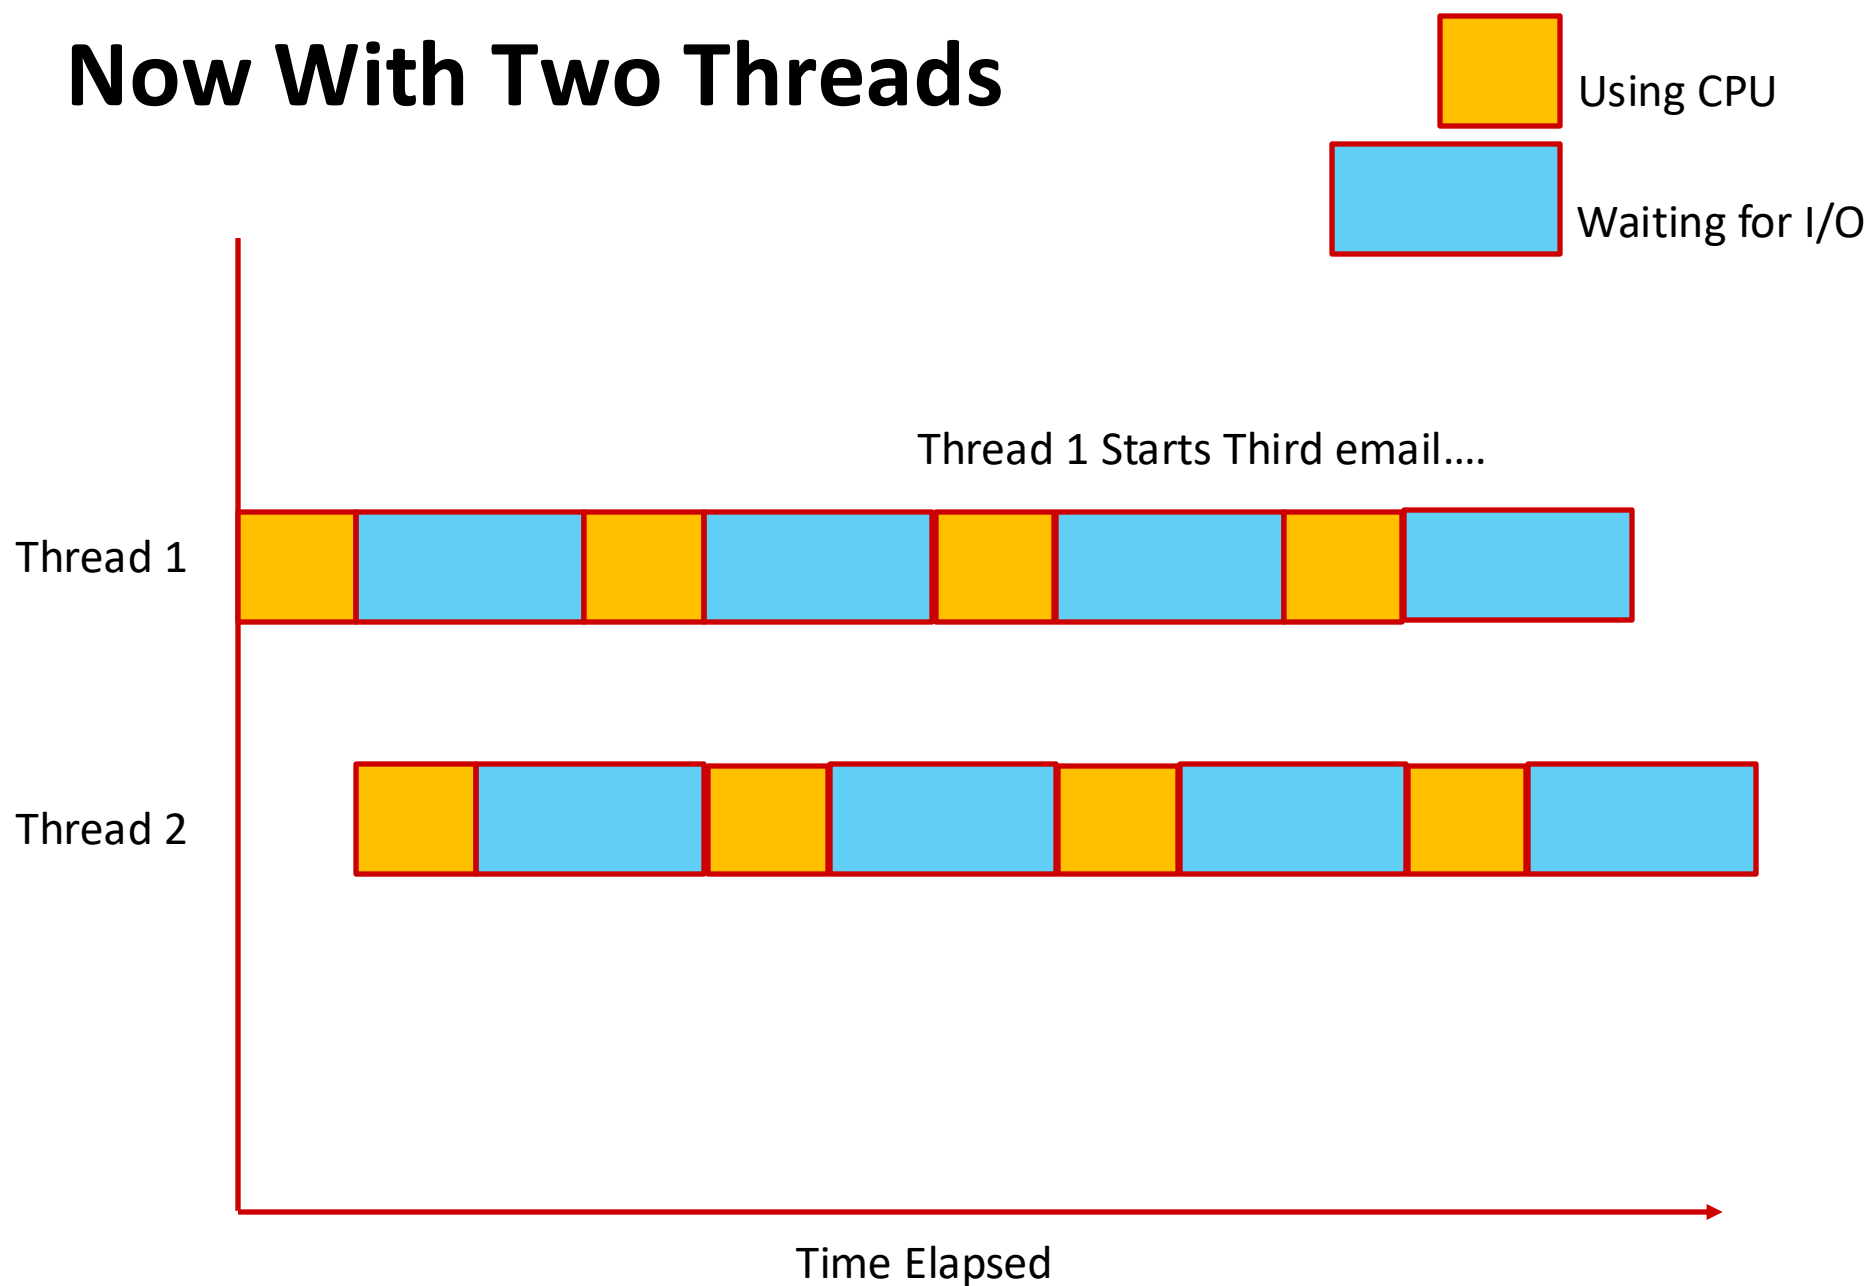
Waiting for I/O



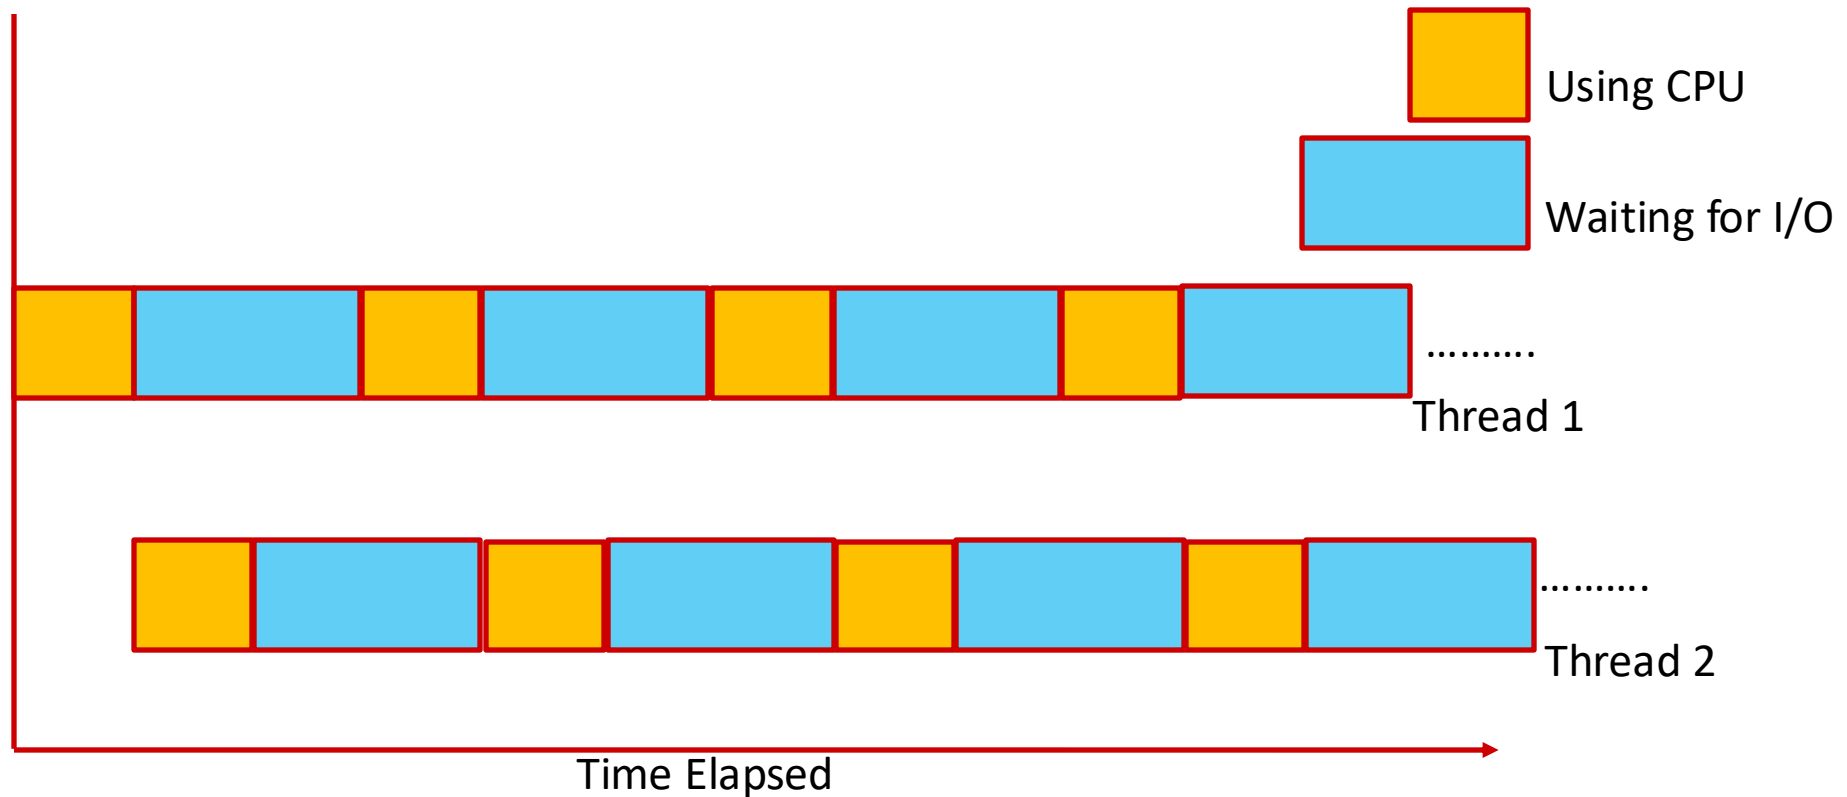
Now With Two Threads



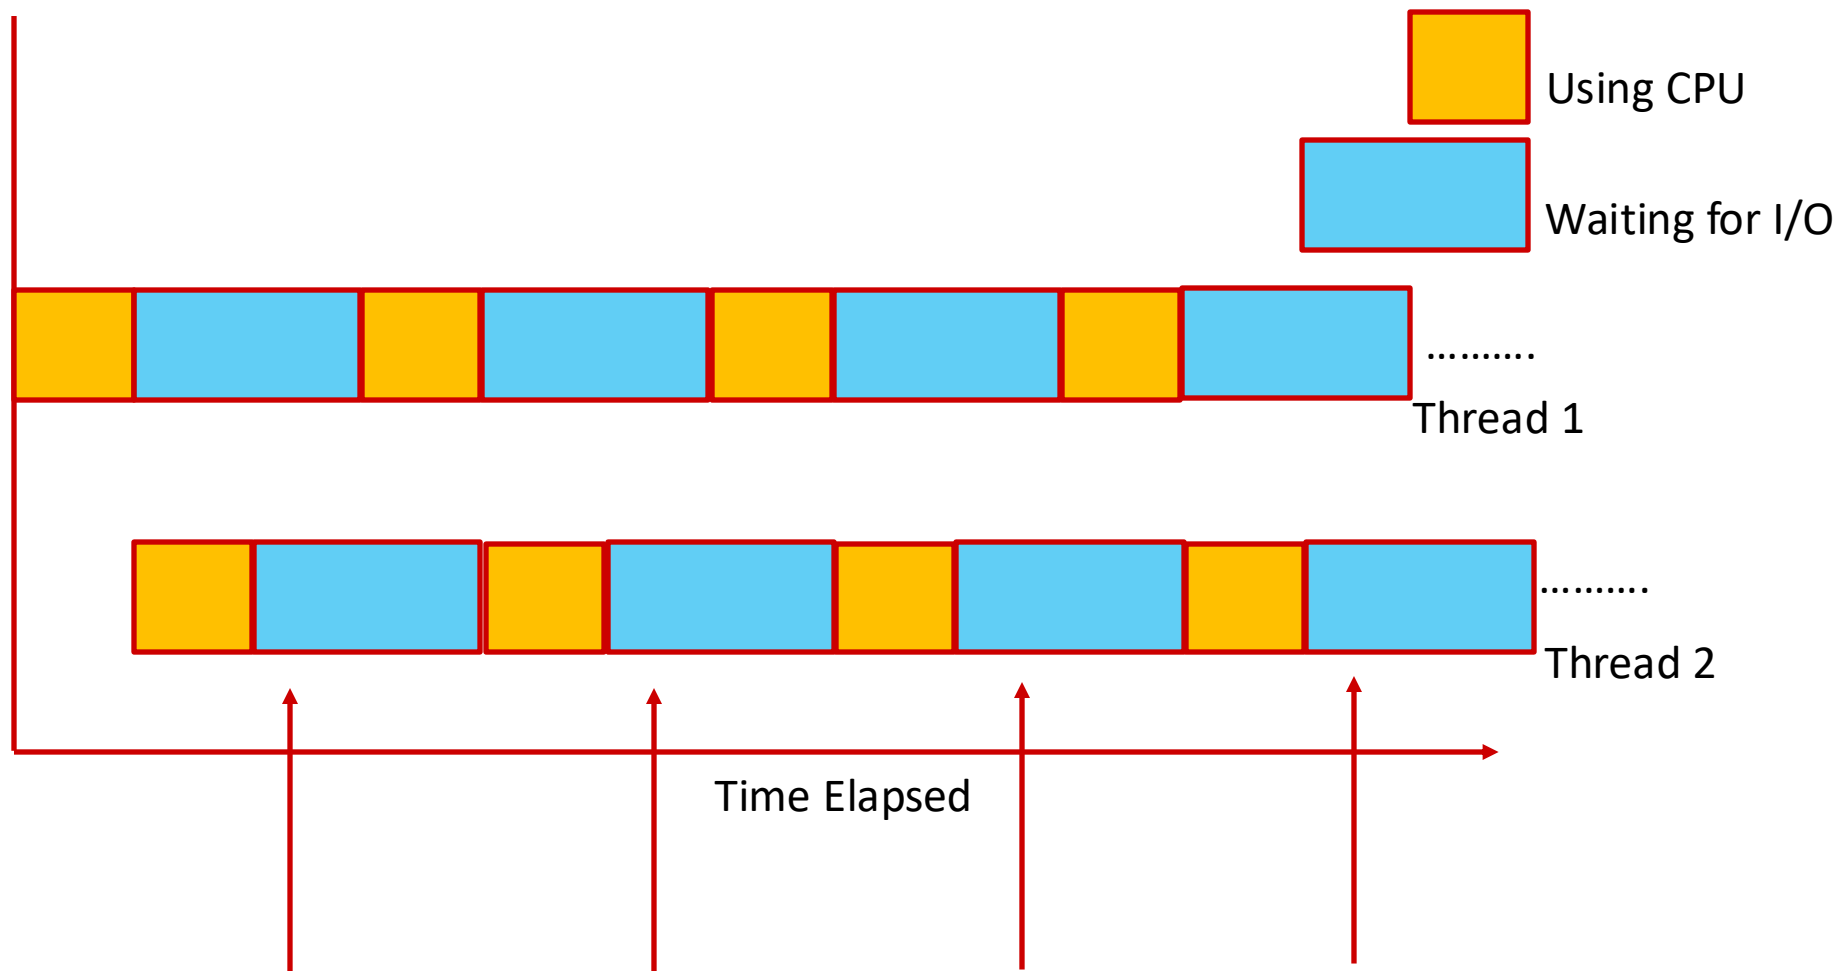
Now With Two Threads



Is the CPU always utilized? Discuss!

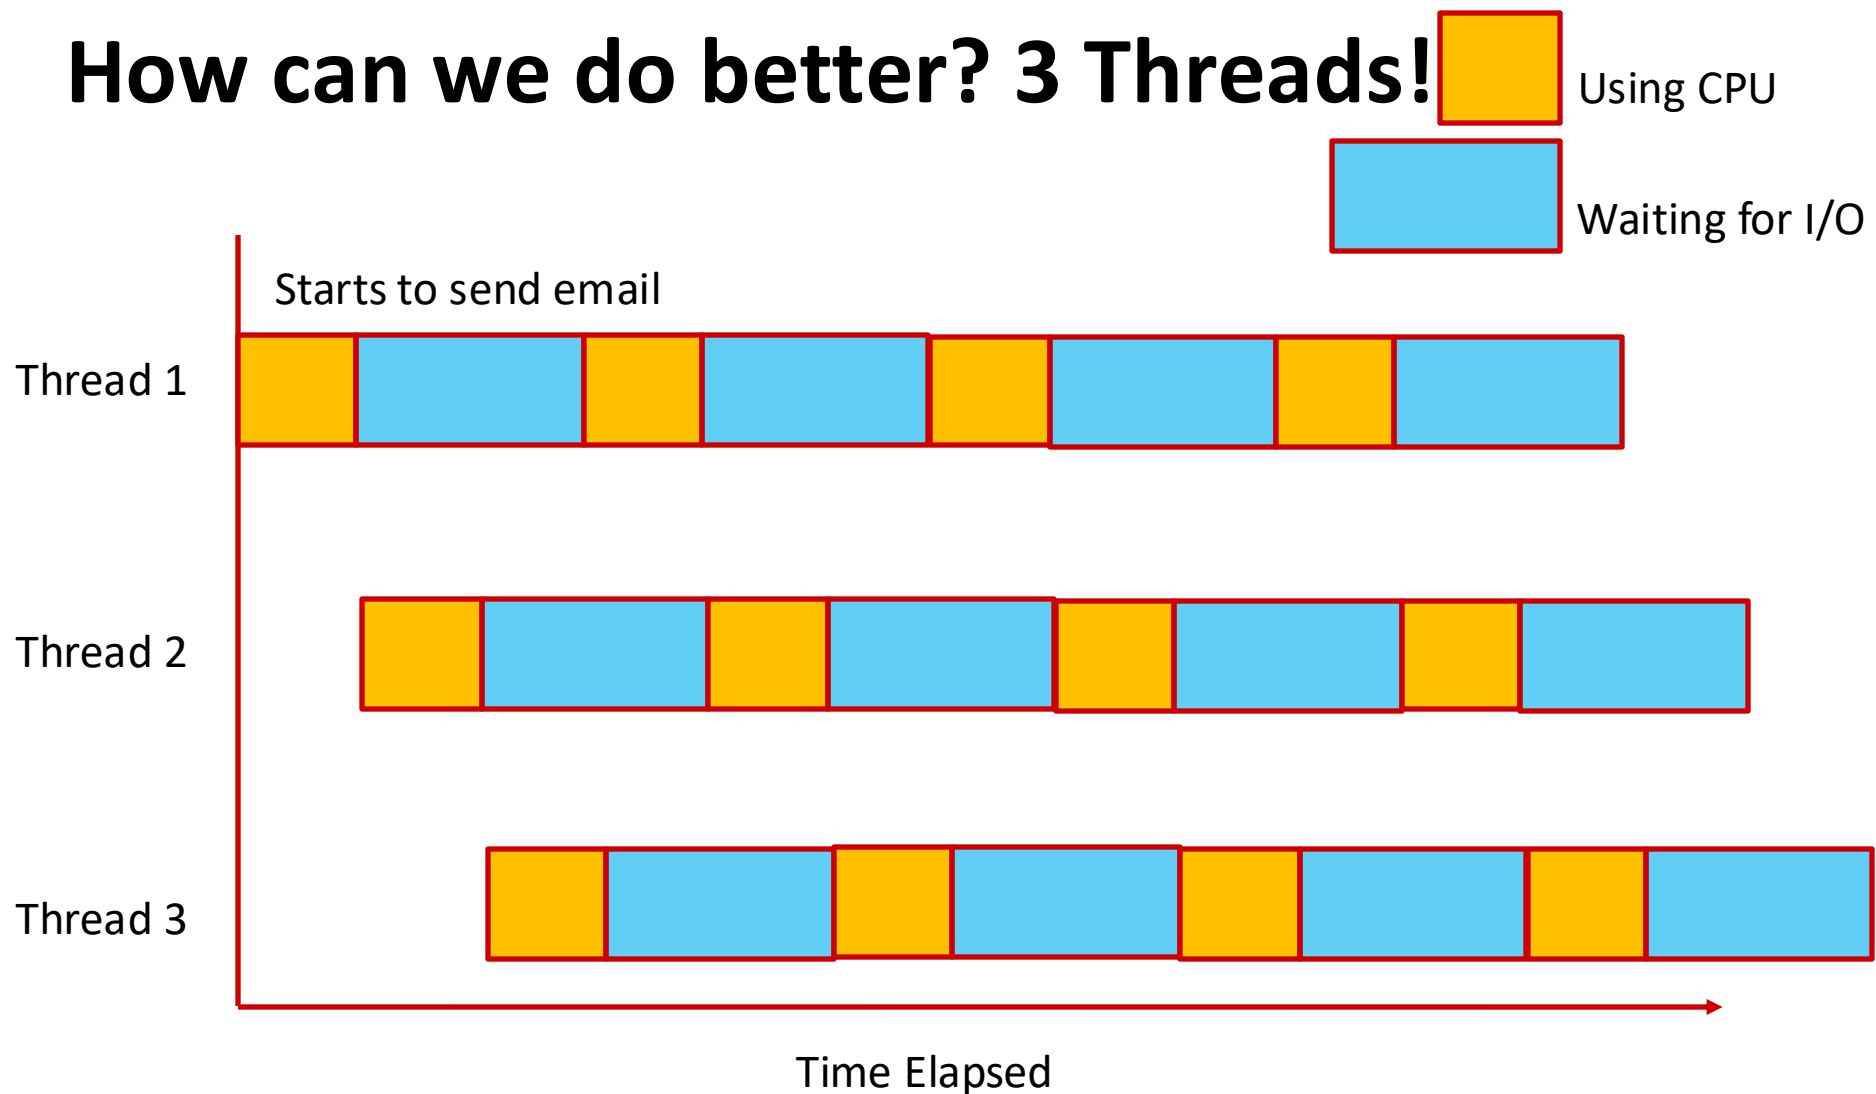


Is the CPU always utilized? No! :/




Here, there is only waiting for I/O until there's more work to do.

How can we do better? 3 Threads!



Here, the CPU or core is consistently active, with no idle time spent waiting for additional tasks to process.

How can we do better? 3 Threads!



Here, the CPU or core is **consistently active**, with **no idle time** spent waiting for additional tasks to process.

As soon as **one thread starts waiting for I/O** there's always another waiting for the CPU.


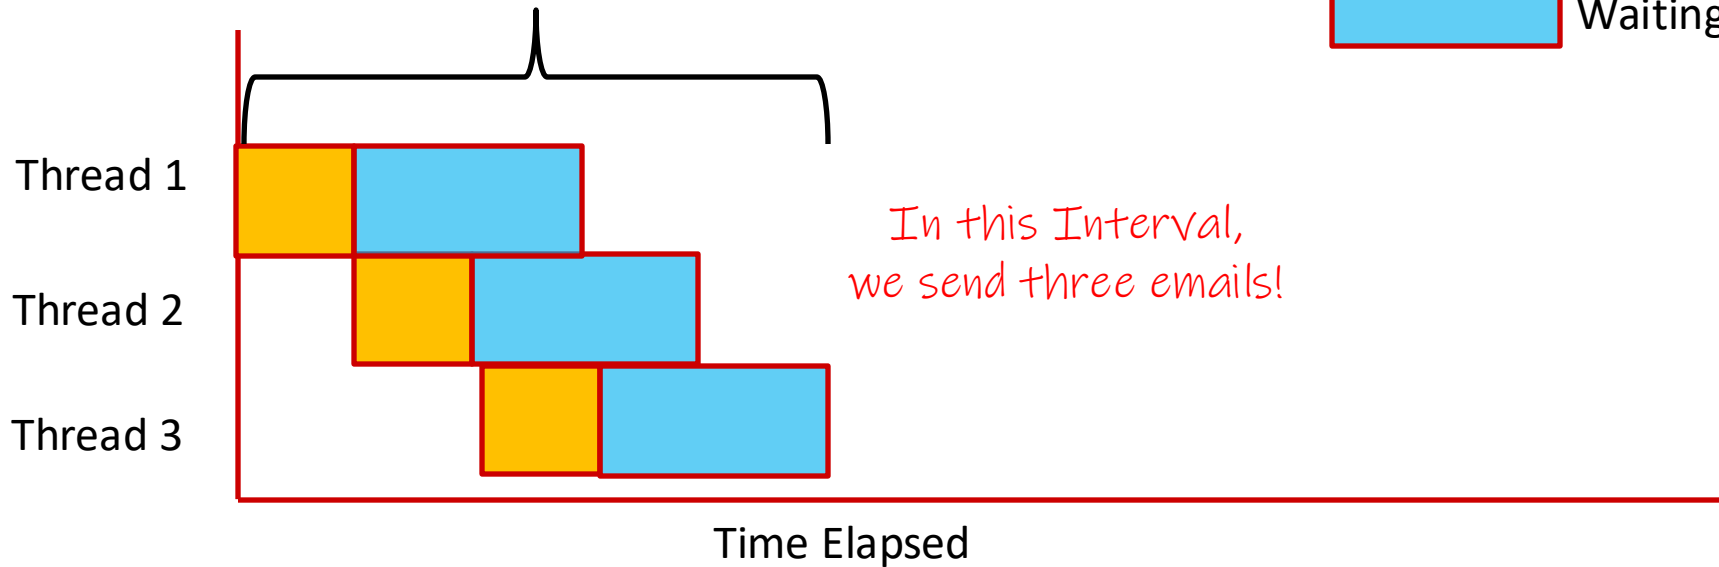
How is this better?



Using CPU



Waiting for I/O




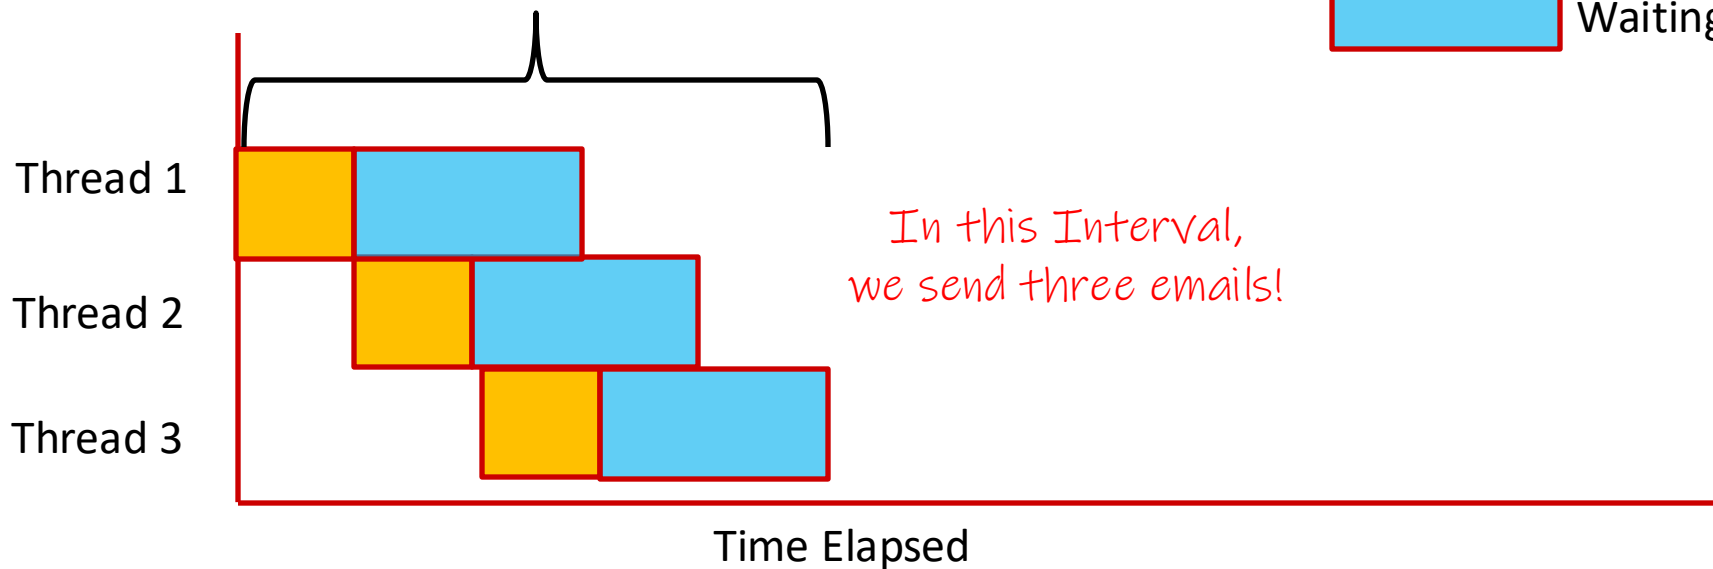
How is this better?





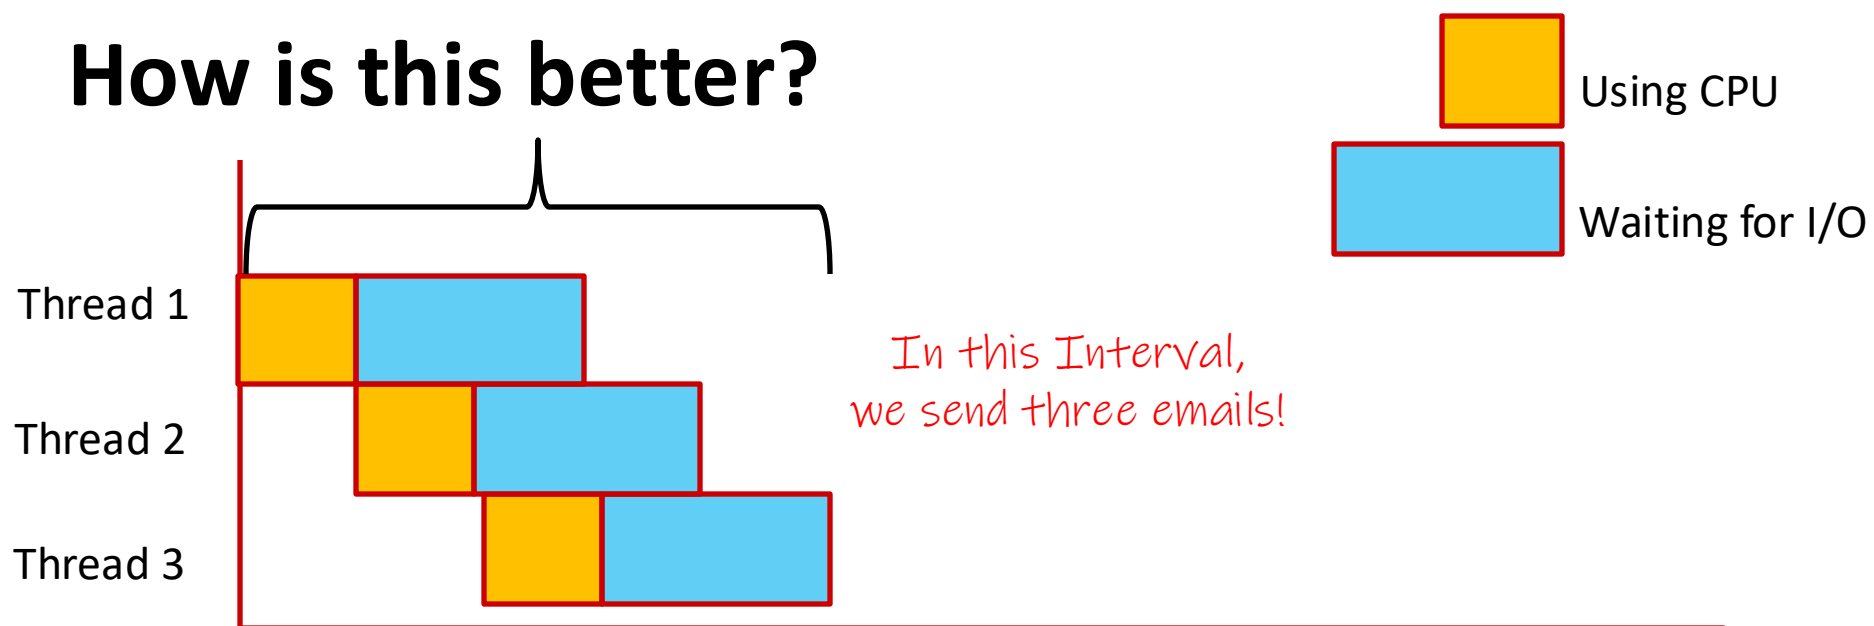
Using CPU



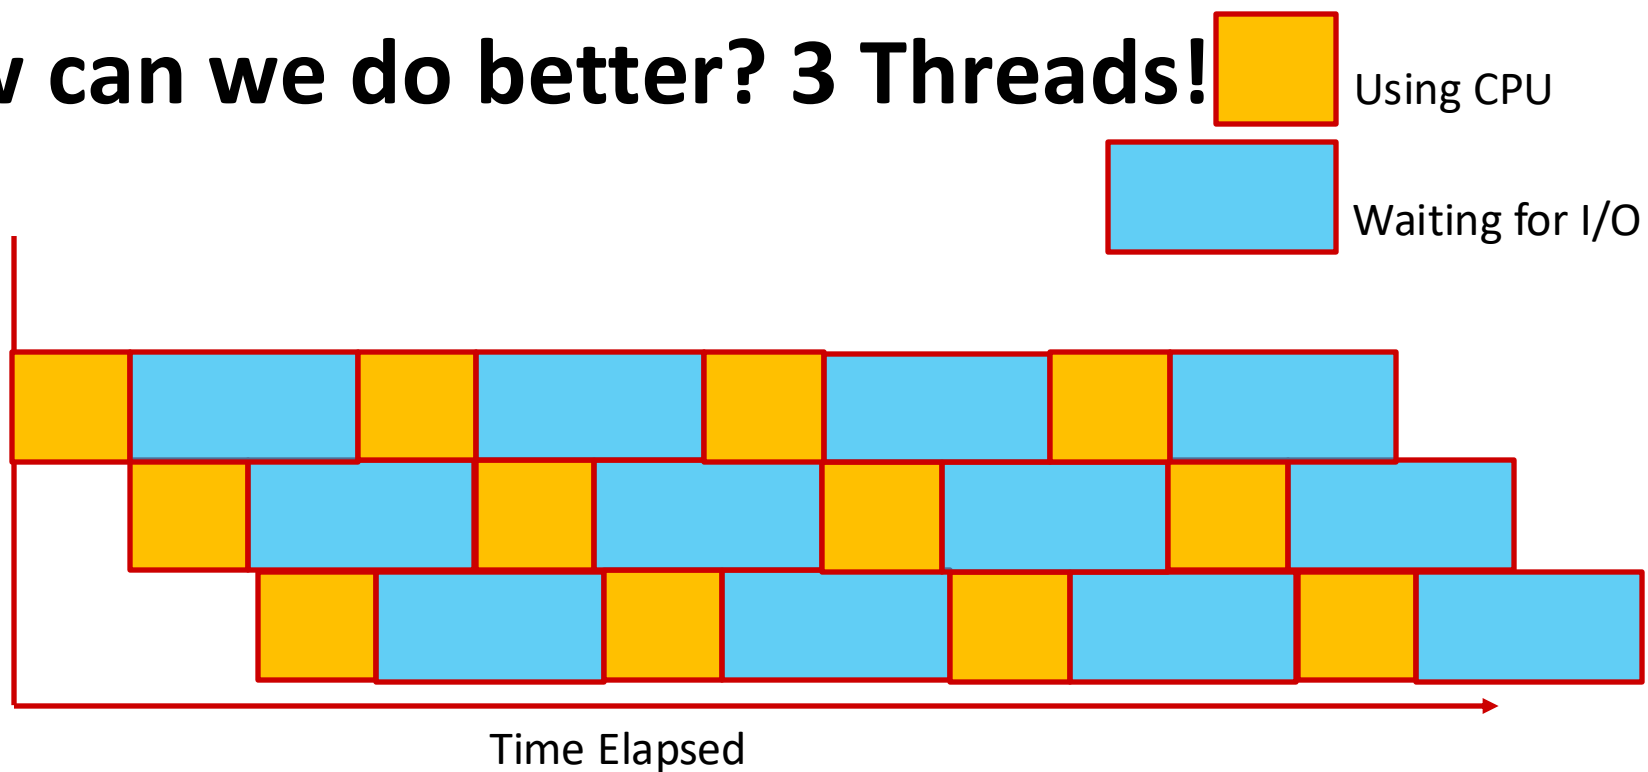
Waiting for I/O



How is this better?



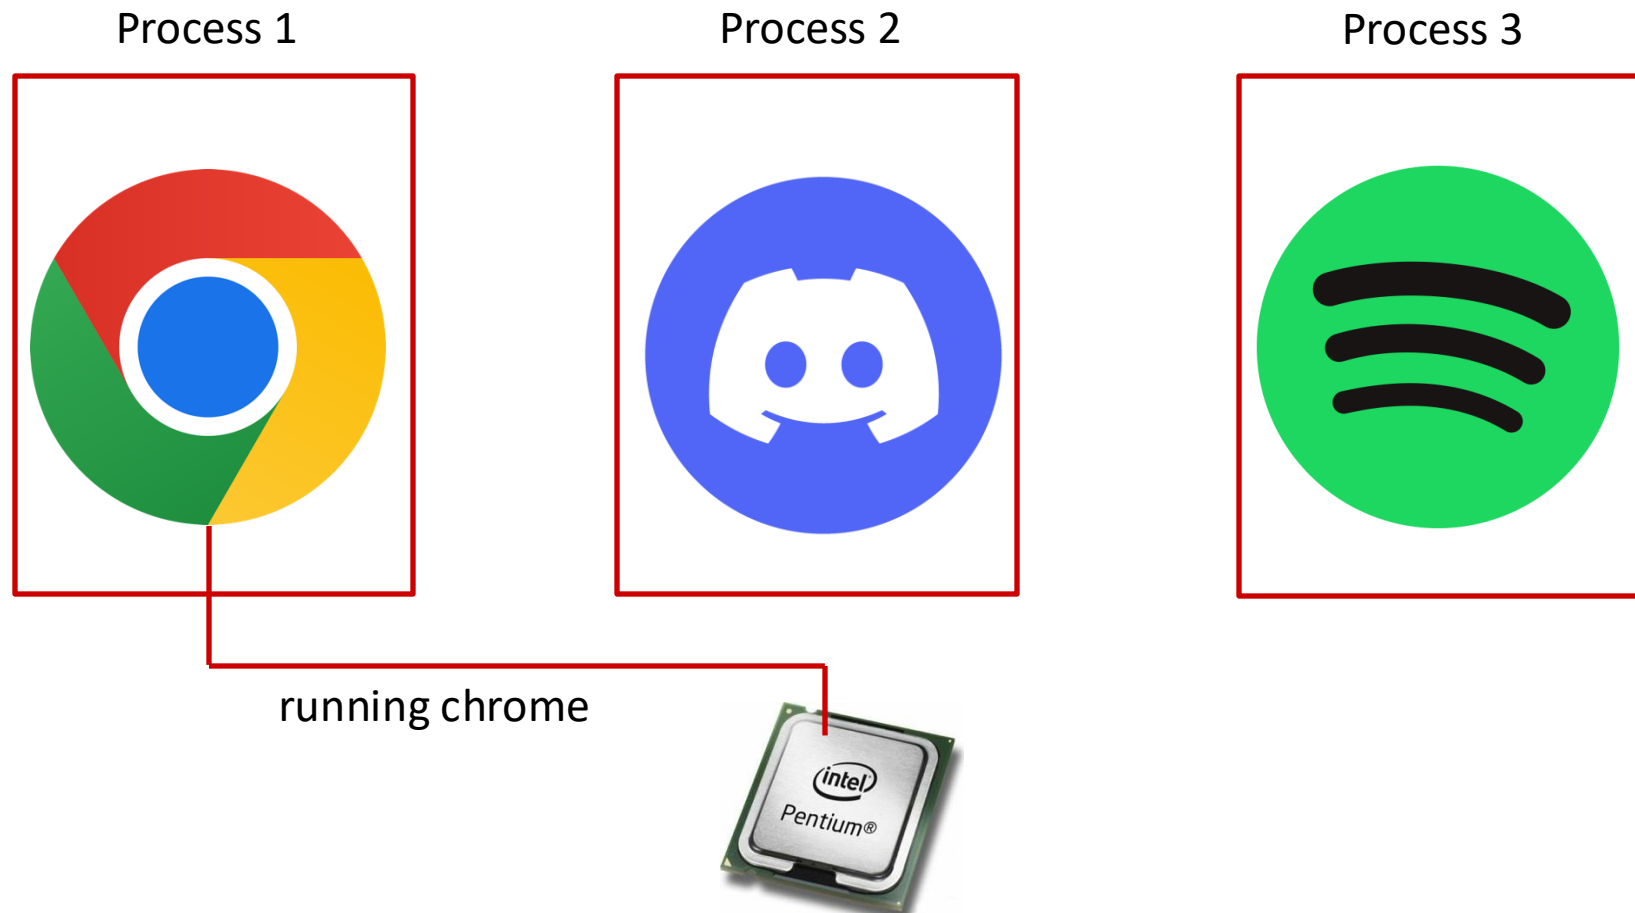
How can we do better? 3 Threads!



In general, using more threads allows multiple tasks to be handled simultaneously, which means more work gets done in the same amount of time.

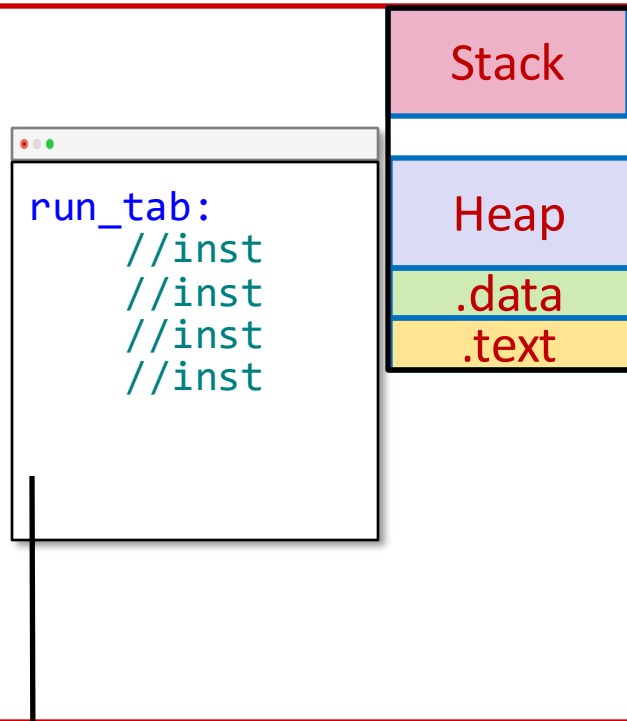
Processes vs. Threads: What's the Difference?

- ❖ Each process has its own memory space—makes sense, right?
- ❖ ***Why should Spotify have access to Chrome's memory?***



Processes vs. Threads: What's the Difference?

Process 1



- ❖ We want to run three tabs.
- ❖ It makes sense to run each tab separately!



Processes vs. Threads: What's the Difference?

Process 1



Stack

Heap

.data

.text

```
run_tab:
//inst
//inst
//inst
//inst
```

Process 2



Stack

Heap

.data

.text

```
run_tab:
//inst
//inst
//inst
//inst
```

Process 2



Stack

Heap

.data

.text

```
run_tab:
//inst
//inst
//inst
//inst
```

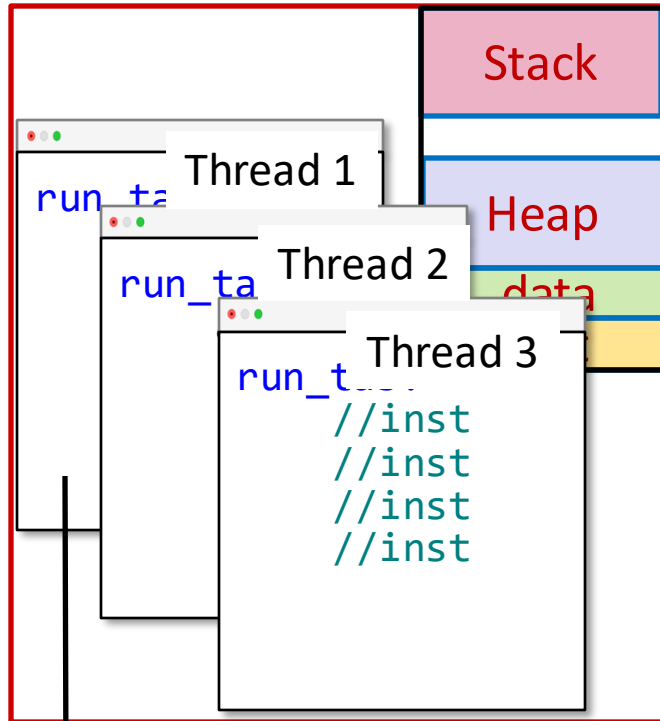
running chrome



Bad: We are allocating too many resources for just three tabs.

Processes vs. Threads: What's the Difference?

Process 1



Why not something like this?

All tabs share the same memory since they're running the same application—this makes sense.

However, each tab operates independently, maintaining its own execution context.

This is One Process with Three Threads



Processes vs. Threads: What's the Difference?

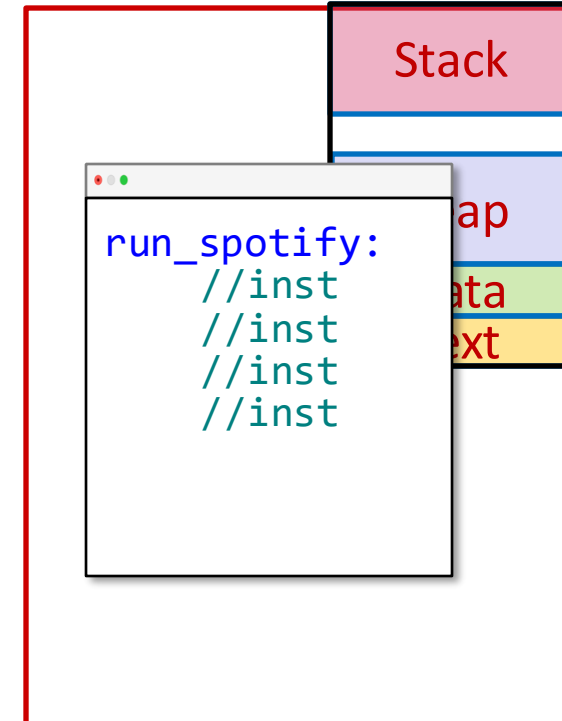
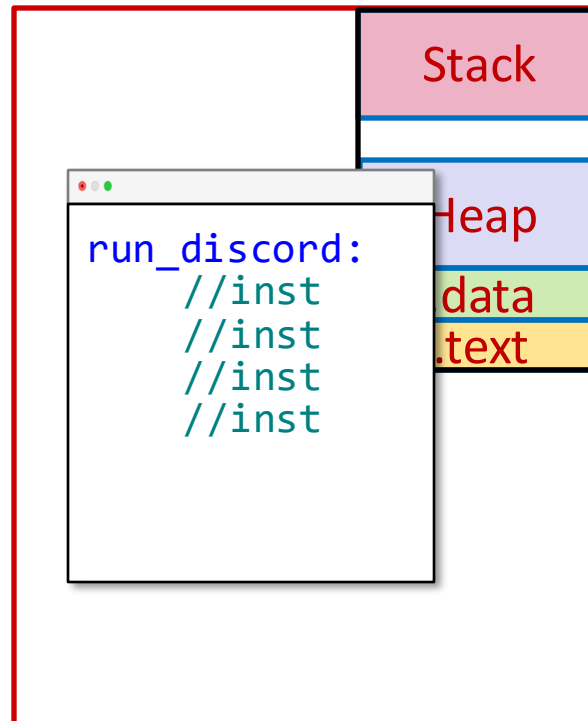
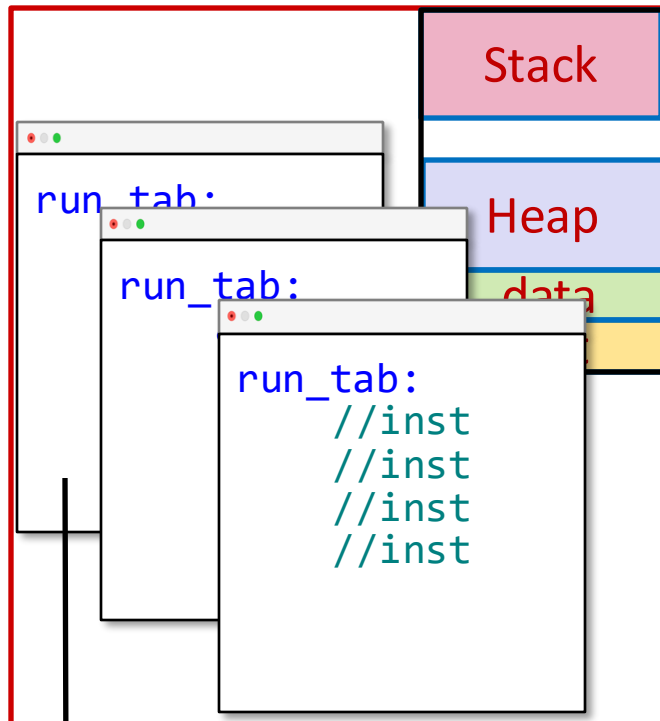
Process 1



Process 2



Process 3



running chrome



Processes vs. Threads: What's the Difference?

Process 1



Process 2



Process 3



Stack

Stack

Stack

run_tab:

Heap

Heap

Heap

run_tab:

data

data

data

run_tab:

```
//inst
//inst
//inst
//inst
```

run_discord:

```
//inst
//inst
//inst
//inst
```

run_spotify:

```
//inst
//inst
//inst
//inst
```

running chrome



Processes vs. Threads: What's the Difference?

Process 1



Process 2



Process 3



Stack

Stack

Stack

run_tab:

Heap

Heap

Heap

run_tab:

data

data

data

run_tab:

.text

.text

```
//inst
//inst
//inst
//inst
```

```
run_discord:
  //inst
  //inst
  //inst
  //inst
```

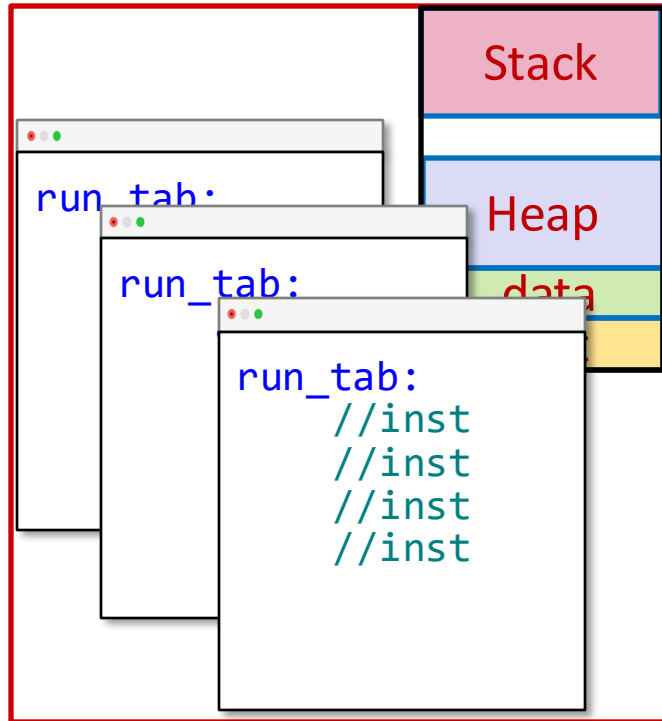
```
run_spotify:
  //inst
  //inst
  //inst
  //inst
```

running chrome

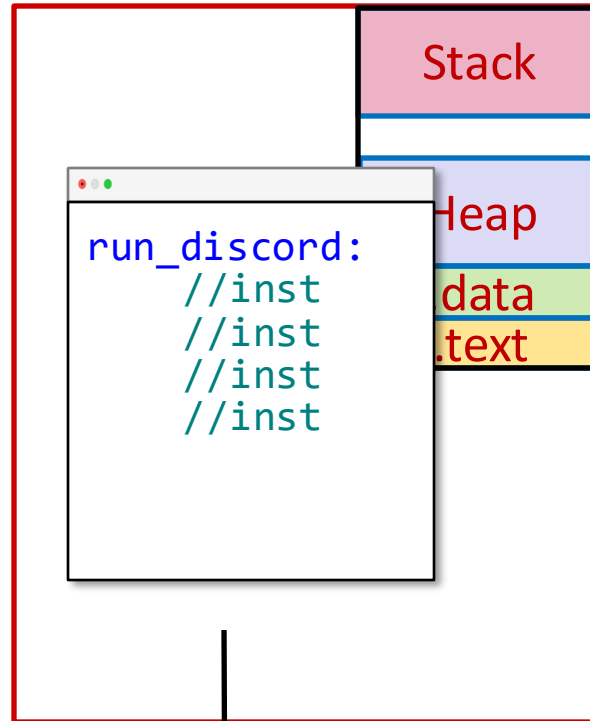


Processes vs. Threads: What's the Difference?

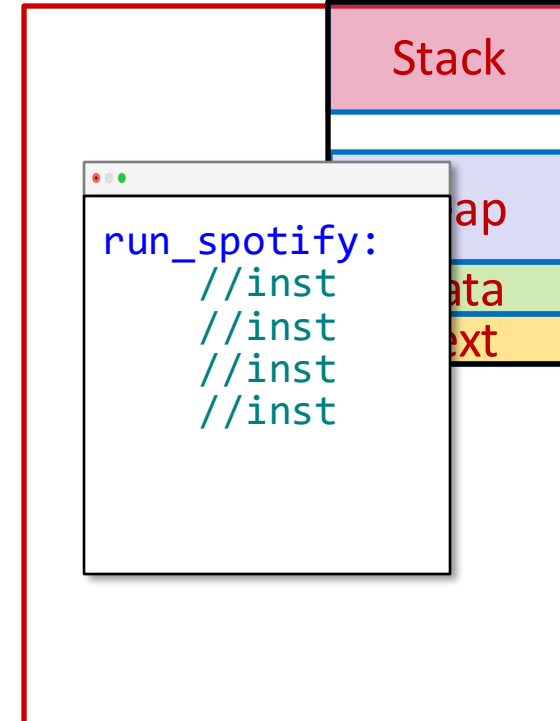
Process 1



Process 2



Process 3

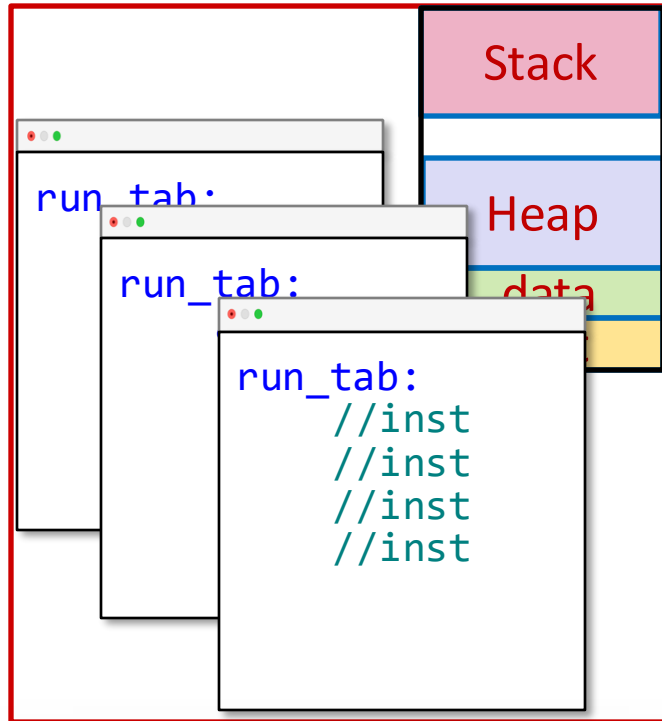


running discord

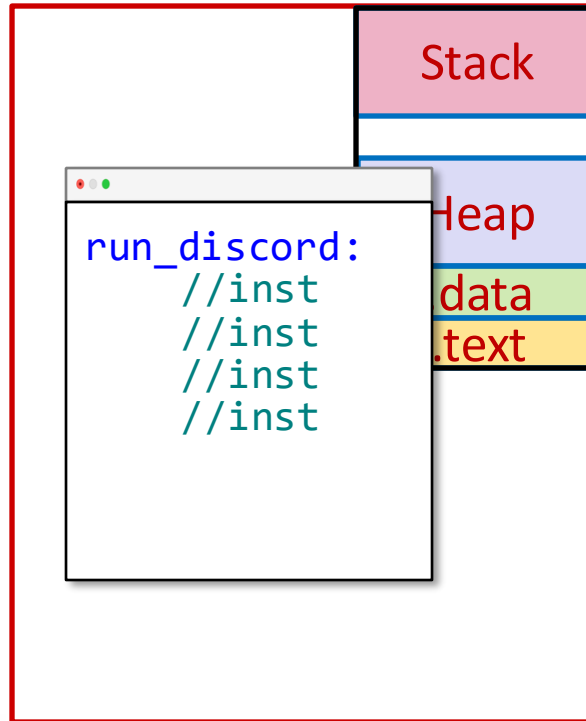


Processes vs. Threads: What's the Difference?

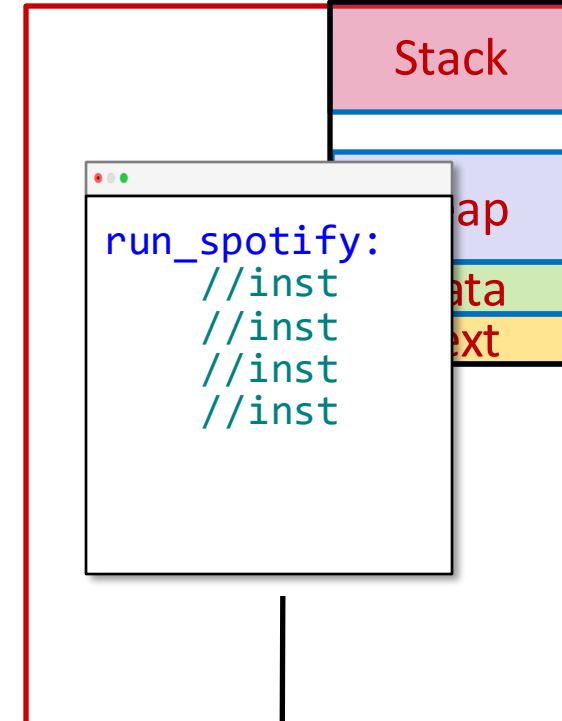
Process 1



Process 2



Process 3




Now we can switch what runs on the CPU within the same process.



running spotify

Example: Visualizing the Mandelbrot Set



Compute Intensity

We need to compute for each pixel, if it belongs in the set in addition.

Currently: ~960000 values

Let's compare how threads help us here:

Non-Threaded vs Fully Threaded Implementations

And that's it! 😊