



Operating System Capstone

Lecture 1: Course Overview

Tsung Tai Yeh

Monday: 10:10 am– 12:00 pm
Classroom: ED-302

Acknowledgements and Disclaimer

- Slides was developed in the reference with
MIT 6.828 Operating system engineering class, 2018
MIT 6.004 Operating system, 2018
Remzi H. Arpaci-Dusseau etl. , Operating systems: Three easy pieces.
WISC

Outline

- Course overview
- References and text books
- Schedule
- Rating
- Operating system basics

Course overview

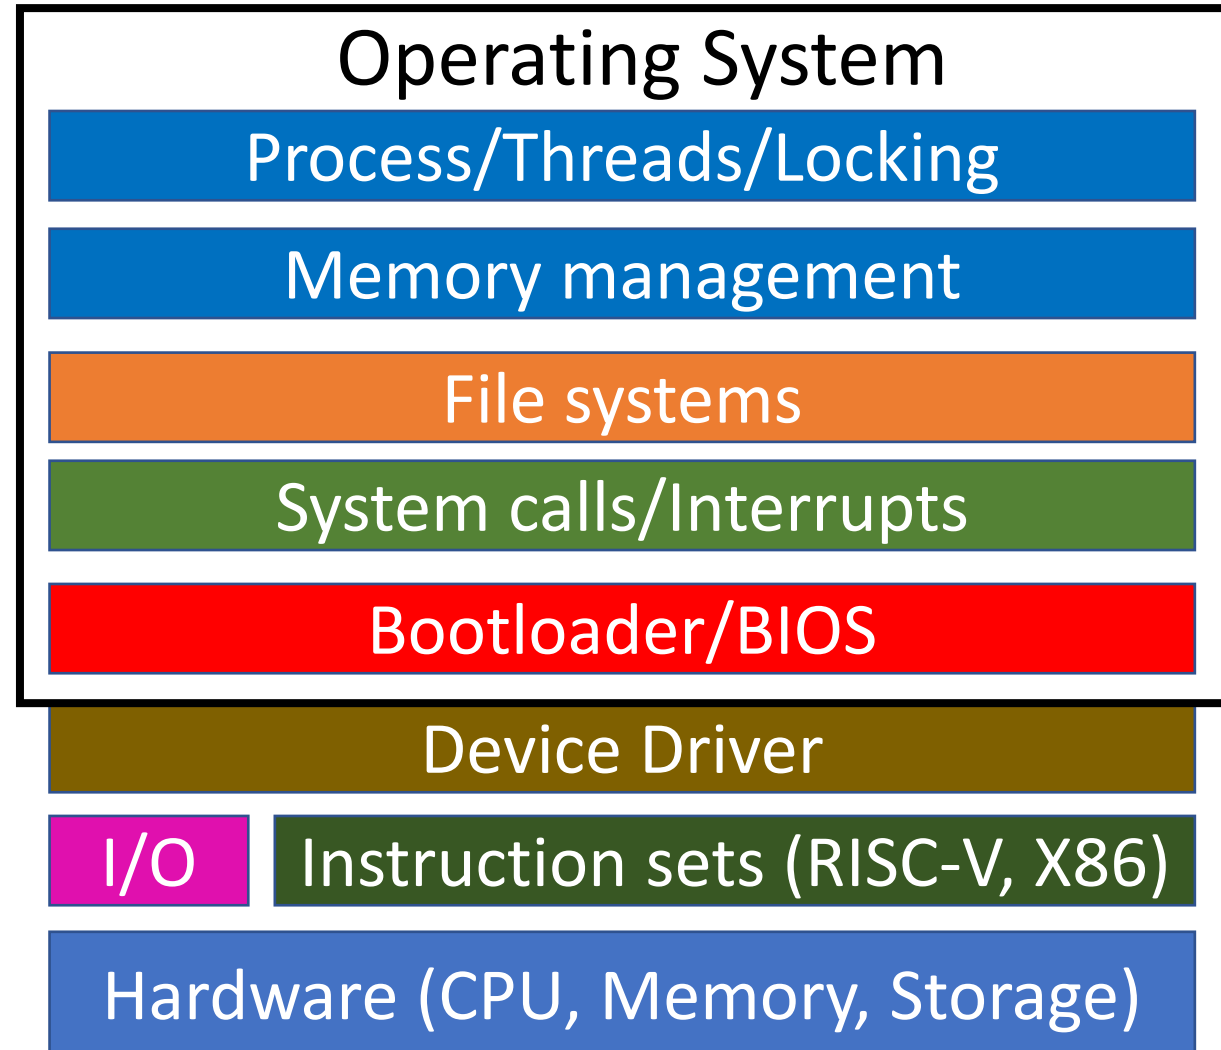
- Instructor: **Tsung Tai Yeh**
- TA team⁺:
- Lecture: M34
- Course web site:
 - <https://shorturl.at/tSX14>
 - <https://nycu-caslab.github.io/OSC2024/>
 - Discord Discussion
 - <https://discord.gg/Gba4rpvNKD>



Course website QR Code

Course overview

- **Operating system design**
 - Hardware + Software
 - Full stack implementation
- **Building a small OS**
- **Lecture + laboratory**
 - Class lecture
 - Reading presentation (10%)
 - 8 labs (95%)



Intended Lecture Outcomes (ILOs)

- What is difference of OSDI class against OS course ?
 - OS concepts + Implementation
 - **Describing** details of the interaction between the computer hardware and OS
 - **Designing** multiple abstractions (system calls, processes, memory management, file systems)
 - **Implementing** bare-metal OS (labs)
 - **Understanding** the OS research on different topics

Lecture

- **Class lecture**

- This lecture also covers each OS topics
- 1 hour lecture – summarize course materials of each topic
- 1 hour reading material presentation
- Lecture materials have shown on the class website
- Students have to preview course materials
- 10:10 am – 12:00 pm on Monday in ED 302

Lab

- Each student will get a Raspberry Pi 3B+ dev. board
- Lab 1-Lab 8 (95%)
 - One lab every two weeks
 - Lab 1 – 5 takes 10% each
 - 6 – 8 takes 15% each
- **Lab Demo**
 - Every student has to demonstrate biweekly your lab work in EC 222
 - TA will check your lab work and ask you questions during your demonstration



Discord Forum QR Code

Reading Presentation

- **Reading Presentation**

- 12 topics, max 5 students are responsible for the presentation of one topics
- Summarize the reading materials
- 1 hour presentation for 1 group
- Each paper presentation takes 10 % of the total score
- Need to prepare 2 – 3 takeaway questions in every topic of the reading material

Takeaway Questions

- What are the purposes of dataflow used by DNN applications?
 - (A) Reduce the data movement across off-chip memory
 - (B) Improve the operational latency
 - (C) Decrease the energy consumption of spatial array accelerator

Schedule

Week	Date	Lecture Topics	Readings	Lab/Project	Materials	Presentation
1	2/19	OS Introduction [slide] [slide]		[slide]	[Note1]	
2	2/26	Assembler,Linker,Loader [slide]			[GOT/PLT] [Note2]	
3	3/4	No Class		Lab 0/1 due		
4	3/11	Boot Loader [slide] [slide]	[Bootling 1-3]		[Note3]	
5	3/18	Process [slide] [slide]	[Bootling 4-6]	Lab 2 due	[Note4]	
6	3/25	Interrupt and exceptions [slide] [slide]	[init 1-5]		[Note5]	
7	4/1	Virtual Memory [slide] [slide]	[init 6-10]	Lab 3 due	[Note6]	
8	4/8	Paging [slide]	[SysCall 1-3]		[Note7]	
9	4/15	Memory allocation [slide] [slide]	[SysCall 4-6]	Lab 4 due	[mimalloc] [Note8]	
10	4/22	Locking [slide] [slide]	[Interrupt 1-5]			
11	4/29	Multi-core Locks [slide] [slide]	[Interrupt 6-10]	Lab 5 due	[RCU]	
12	5/6	File system [slide] [slide]	[MM 1-3]			
13	5/13	Flash Translation Layer [slide] [slide]	[lock1-3]	Lab 6 due		
14	5/20	Device driver [slide] [slide]	[lock4-6]			
15	5/27	Virtual Machine	[FileSystem]	Lab 7 due	[VM1] [VM2] [VM3]	

References and text books

- OS/2 references
 - Andrew S Tanenbaum, and Albert S Woodhull, "Operating Systems Design and Implementation (3rd Edition)"
 - Marshall Kirk McKusick, Keith Bostic, Michael J. Karels, and John S. Quarterman, "The Design and Implementation of the 4.4 BSD Operating System"
- Linux Kernel
 - Robert Love, "Linux Kernel Development (3rd Edition)"
 - Michael Beck, Harald Bohme, Mirko Dziadzka, Ulrich Kunitz, Robert Magnus, and Dirk Verworner, "Linux Kernel Internals (2nd Edition)"
 - Daniel P. Bovet, and Marco Cesati, "Understanding the Linux Kernel, Third Edition"

References and text books

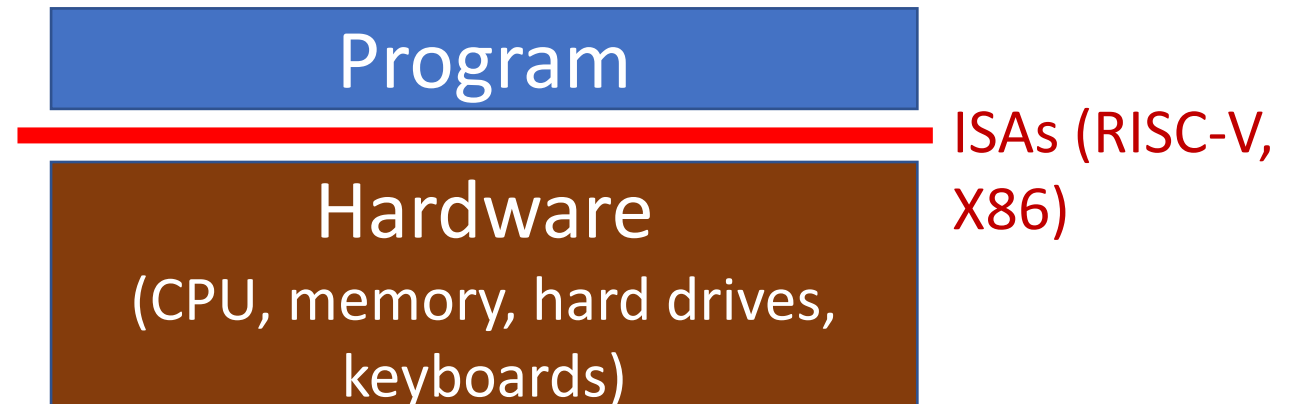
- Network subsystem
 - Klaus Wehrle, Frank Pahlke, Hartmut Ritter, Daniel Muller, and Marc Bechler, "Linux Networking Architecture"
 - Christian Benvenuti, "Understanding Linux Network Internals"
- Device Drivers
 - Sreekrishnan Venkateswaran, "Essential Linux Device Drivers"
 - Jonathan Corbet, Alessandro Rubini, and Greg Kroah-Hartman, "Linux Device Drivers, 3rd Edition"

References and text books

- One text book is not enough
 - Usually you need to refer at least three text books on the same topic and then you realize you have to understand another three topics
- Even you have all text books are not enough
 - Usually you have to “try, try and see”
- Even you understand the codes (or can write the codes) are not enough
 - Usually you have to study the books again and think carefully and deeply

Single-User Machines

- Hardware executes a single program
- The program can access directly all hardware resources in the machine
- The instruction set architecture (ISA) is the interface between software and hardware
- However
 - Most computer systems aren't work like this !

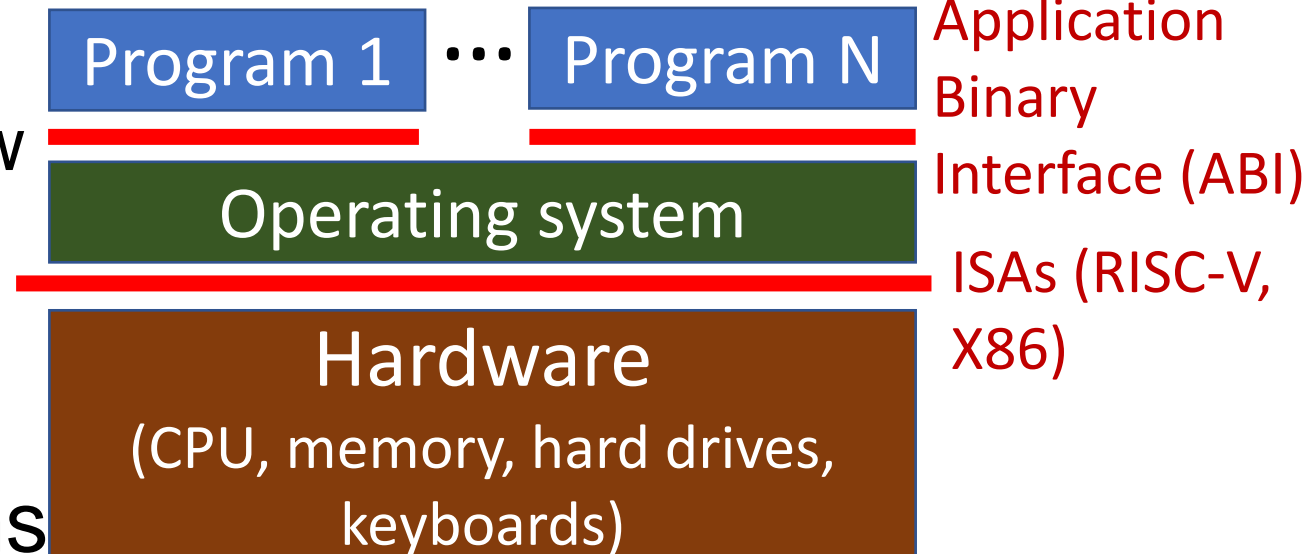


Operating systems

- Multiple executing programs share the machine
- Each program cannot access hardware resource directly

- An operating system (OS)
 - Control these programs how they share hardware

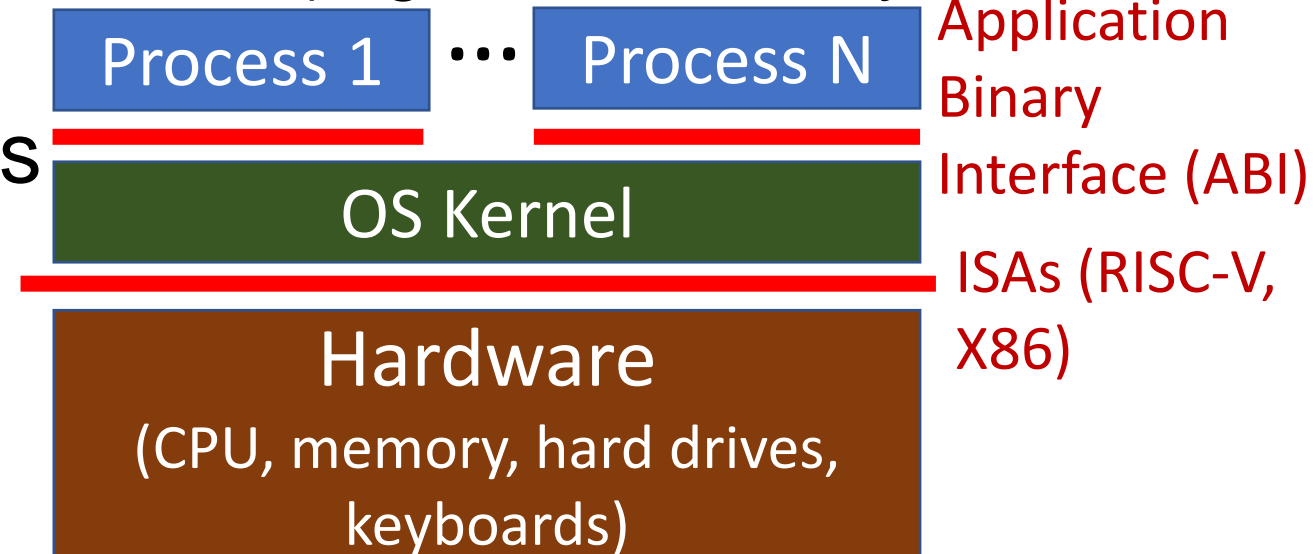
- The application binary interface (ABI) is the interface between programs and the OS



Process vs. Program

- A program is a collection of instructions
- A **process** is an instance of a program that is being executed
 - Include program code + other state (registers, memory, and other resources)

- The **OS kernel** is a process with special privileges

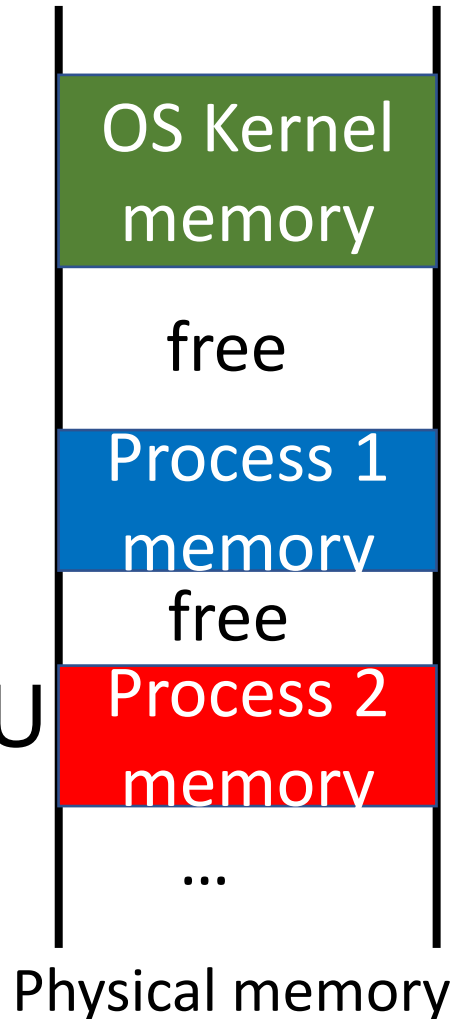


Goals of operating systems

- An operating system is to support several activities at once
 - Many running program as **processes**
- **Protection and privacy**
 - Process multiplexing
 - Processes cannot access each other's data (isolation)
- **Abstraction**
 - OS hides details of underlying hardware
 - Hardware resource manager

Operating systems: The big picture

- The OS kernel provides a **private address space** to each process
 - Each process is allocated space in physical memory by the OS
 - A process is not allowed to access the memory of other processes
- The OS kernel **schedules processes** into the CPU
 - Each process is given a fraction of CPU time
 - A process cannot use more CPU time than allowed
 - Context switch

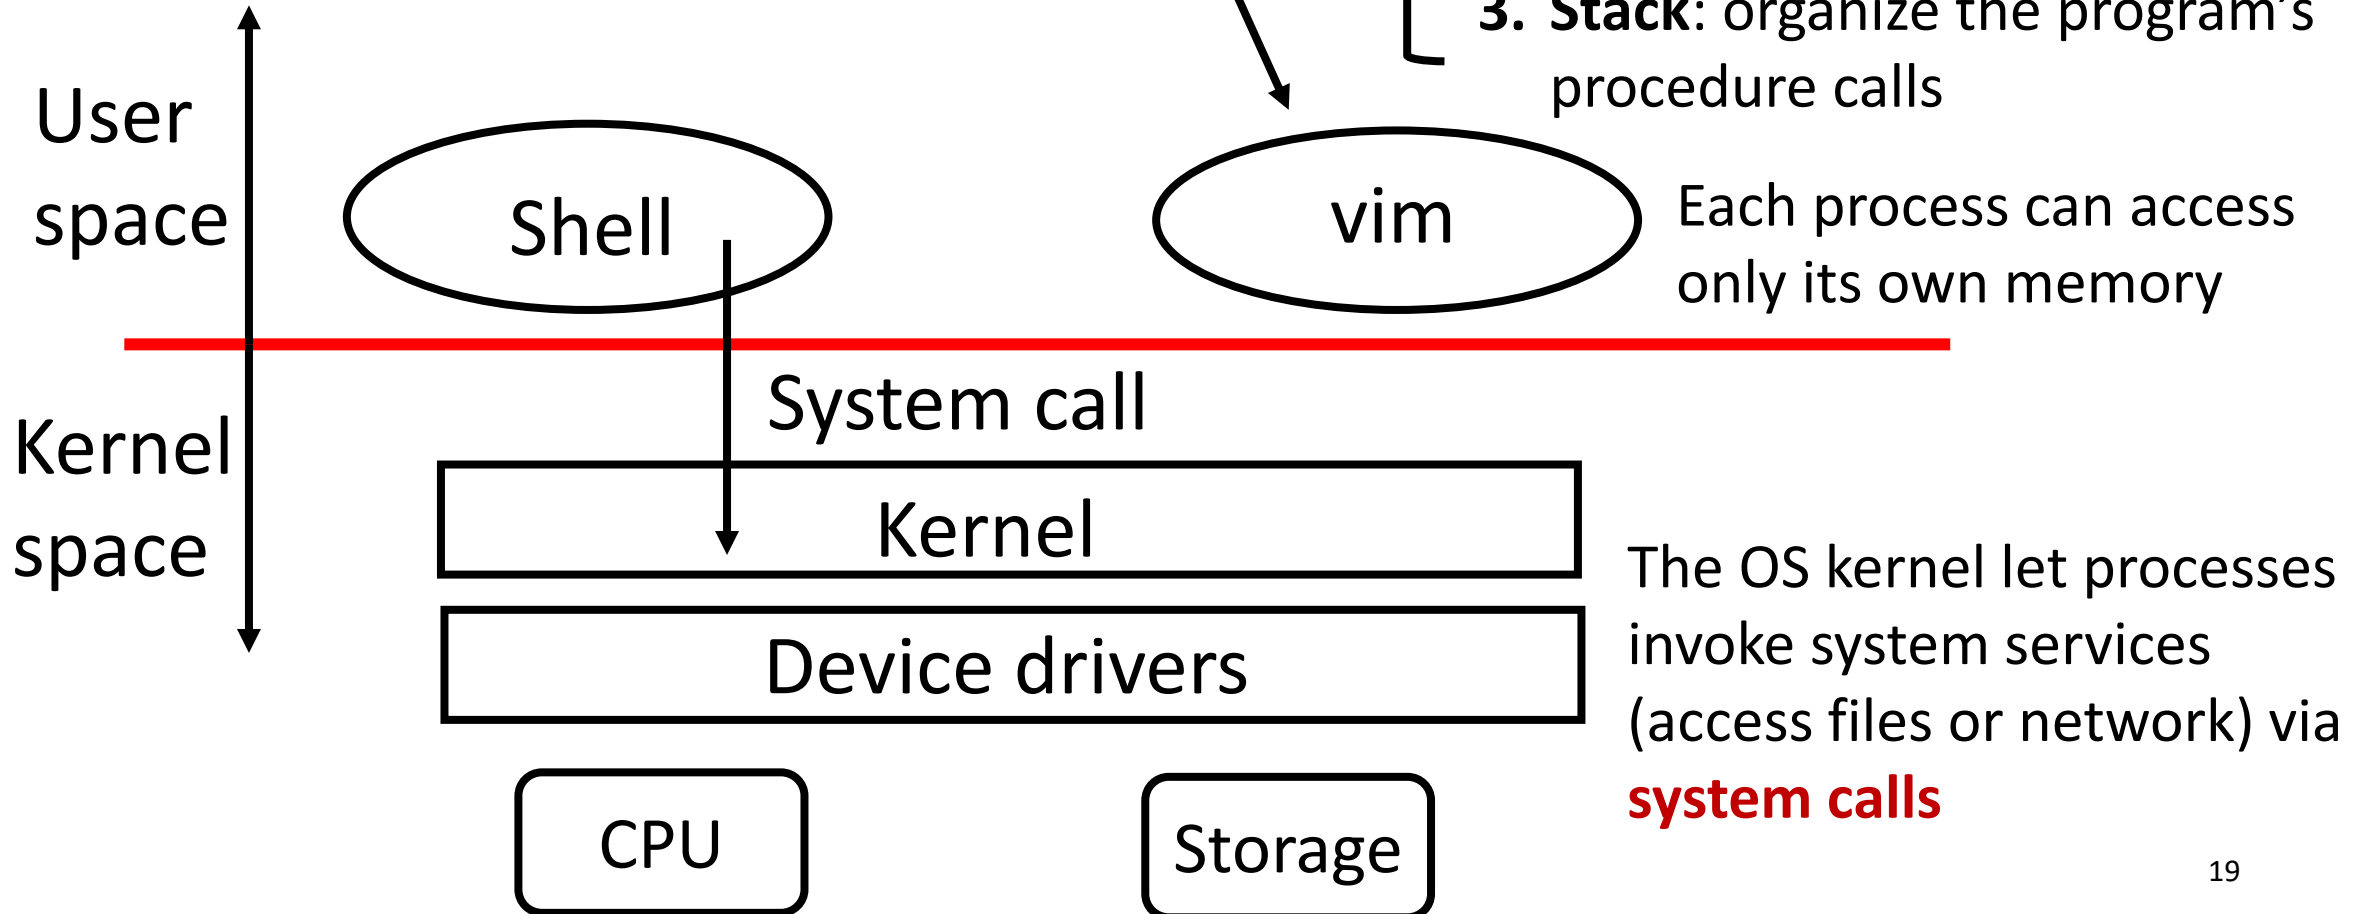


Each running program is called **process**

How does OS work?

Process

1. **Instructions:** implement the program's computation
2. **Data:** the variables on which the computation acts
3. **Stack:** organize the program's procedure calls



Implementing an OS

- The OS works as a **virtual machine (VM)** to each process
 - Each process believes it runs on its own machine
- Virtual machines can be implemented entirely in software, but at a performance cost
 - For instance, python programs are 10 – 100x slower than native Linux programs because python interpreter overheads
- We want to support operating systems with minimal overheads
 - Need hardware support for virtual machine

User and kernel mode

- Two modes of execution: **user** and **kernel (supervisor)**
 - OS kernel runs in supervisor mode
 - All other processes run in user mode
- **In the kernel mode**
 - Privilege instructions and register are available
 - Interrupts and exceptions to safely transition from user to supervisor mode
- **Virtual memory**
 - Provide private address spaces and abstract the storage resources of the machine

What services does an OS kernel provides?

- Processes
- Memory allocation
- File systems
- Security
- Others: users, networking, terminals, etc..

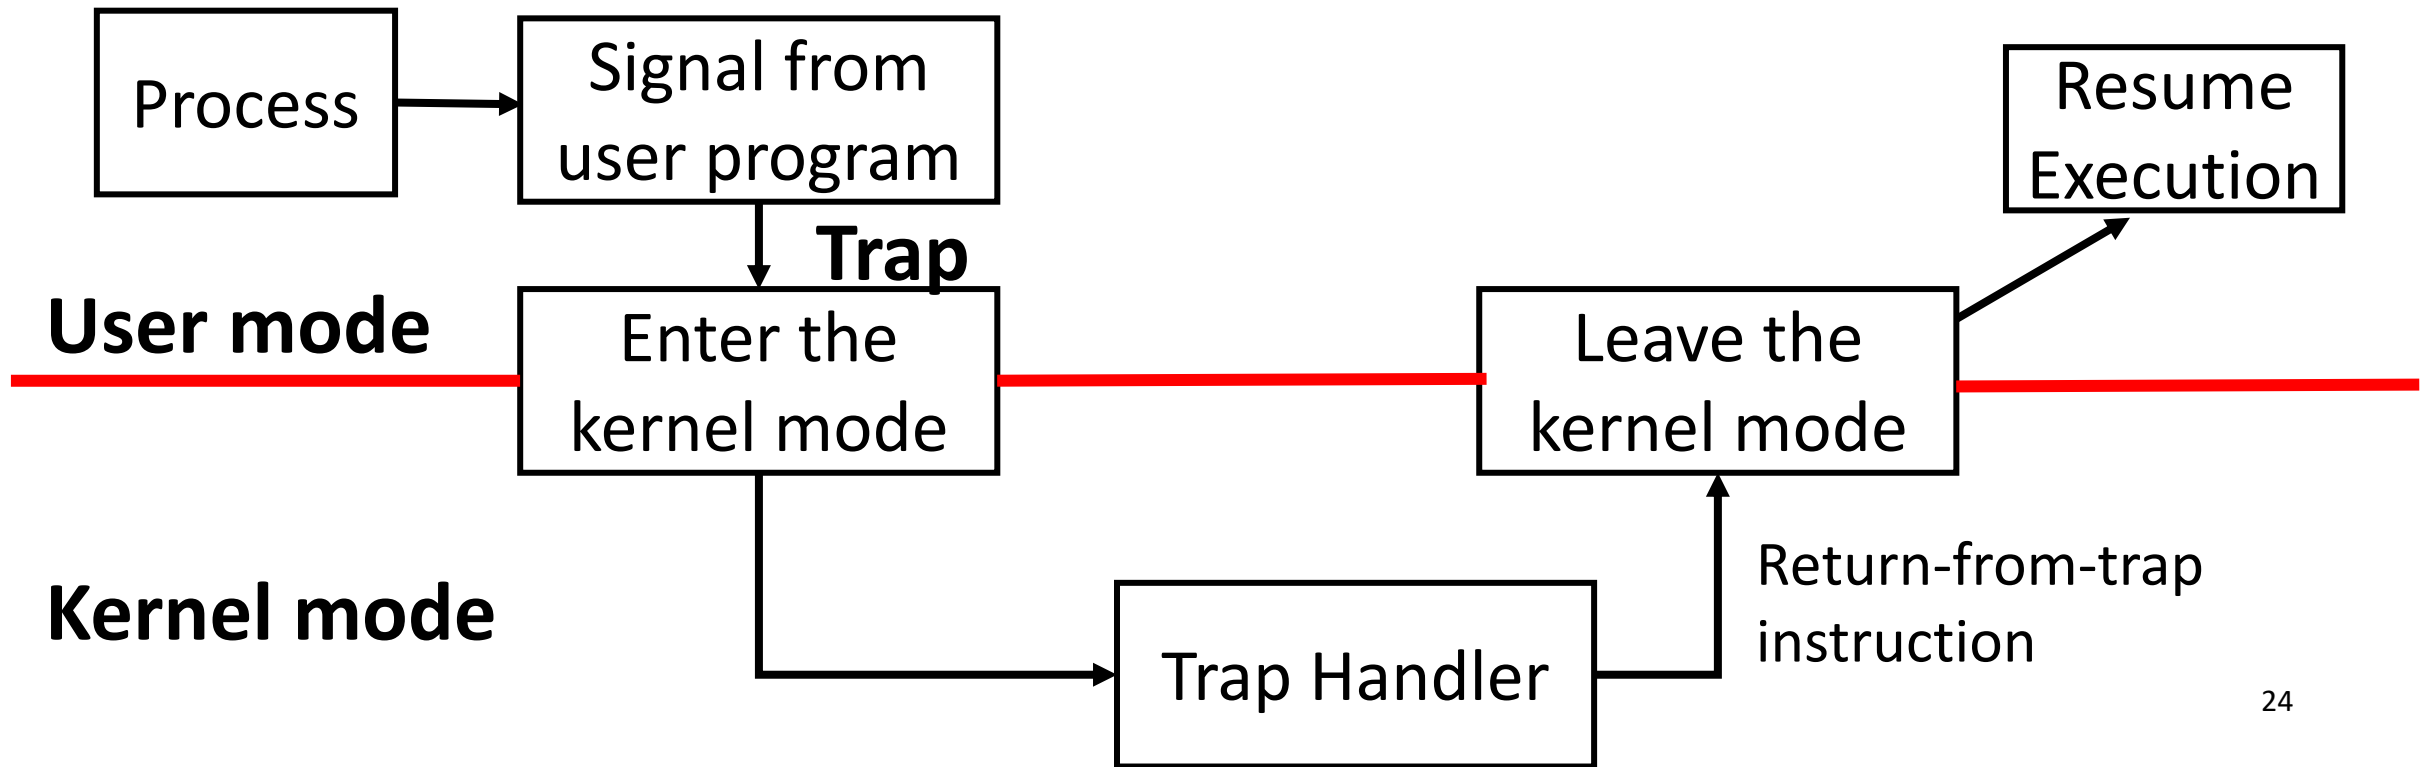
Process and thread

- Each process has a thread of execution
 - The state of a thread (local variables, function call return address) is stored on the thread's stacks
 - Each process has two stacks: a user stack and a kernel stack

Process	Thread
Process is any in-execution program	Thread is the segment of a process
Process is isolated	Thread share memory
Process has its own process control block (PCB) and address space	Thread has parent's PCB, its own TCB, stack, and address space
Process takes more time for creation	Thread takes less time for creation

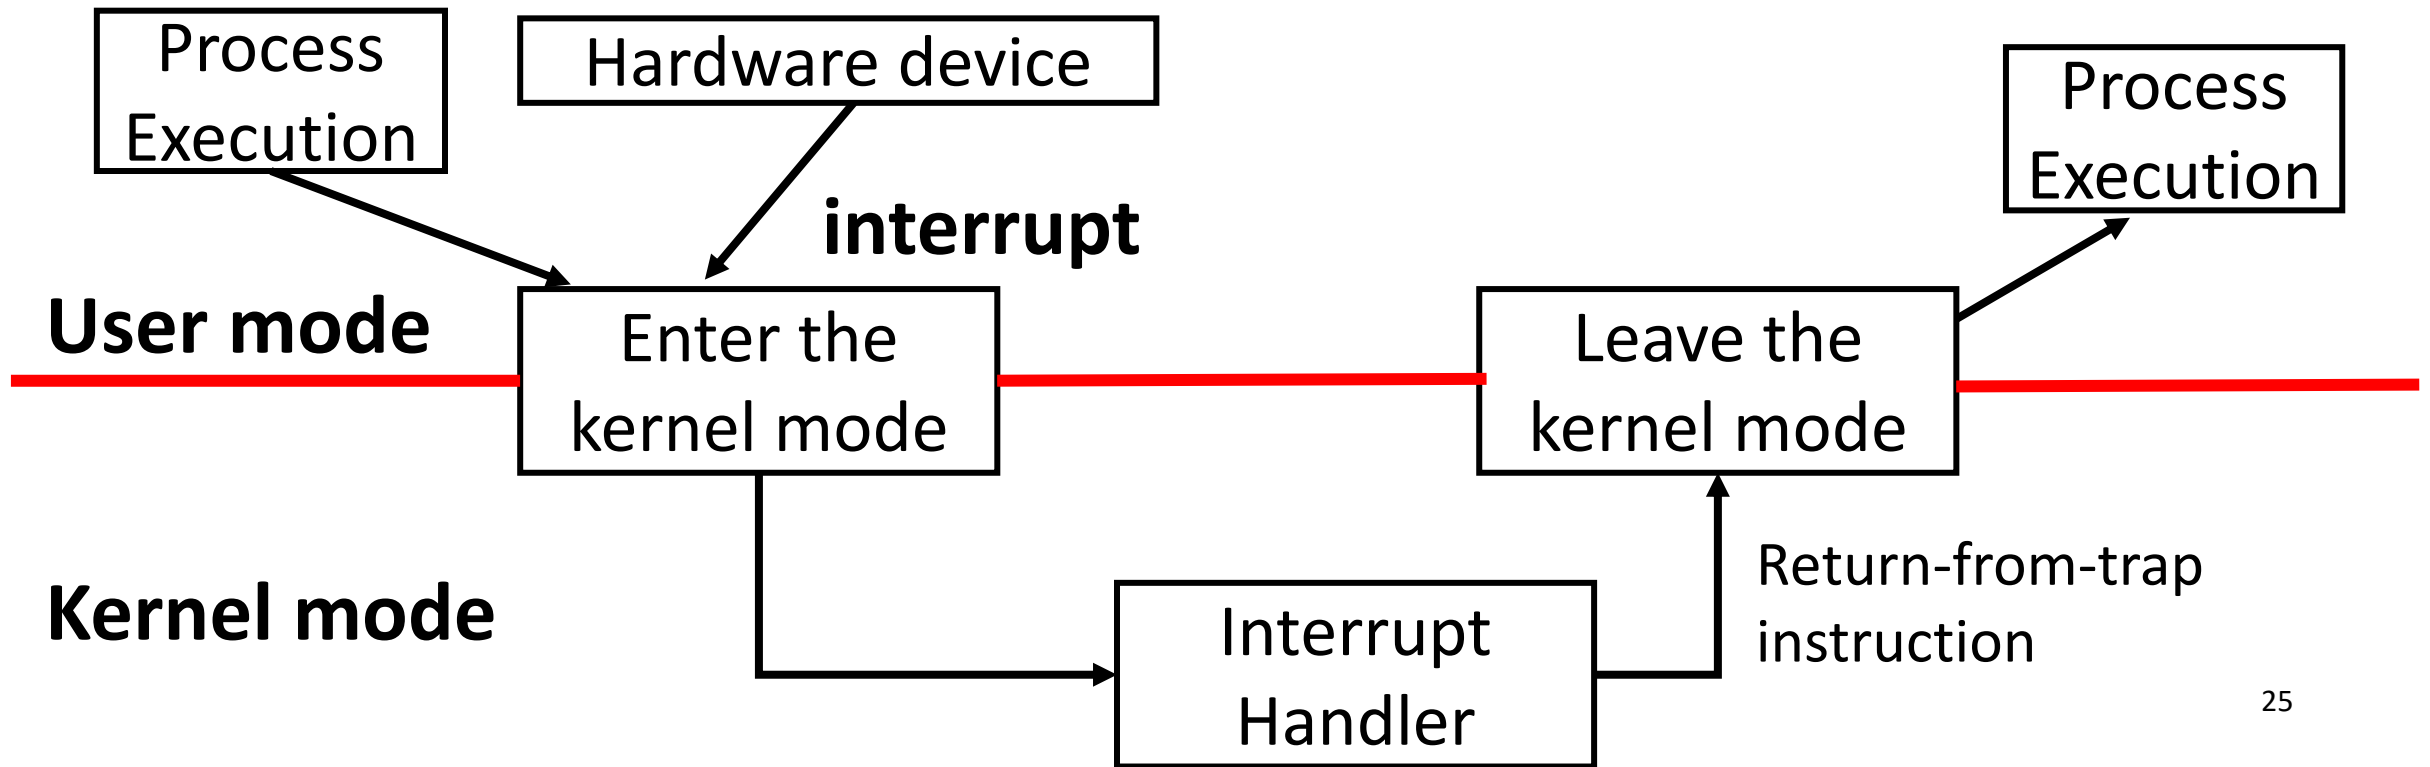
System call

- Using a **trap** that is a synchronous interrupt triggered by an exception in a user process to execute functionality.
1. During a trap, the execution of a process is set as high priority compared to user code
 2. When the OS detects a trap, it pauses the user process
 3. The OS resumes the execution when the system call is completed



What is the interrupt in the OS ?

- An interrupt is a hardware or software signal and notifies the processor that a critical process needs urgent execution
- Using to interrupt present working process



What is the interrupt in the OS ?


- **interrupt service routine (ISR)**
 - A specific bus control line handles interrupts in I/O devices
- A CPU contains a specific interrupt pin known as INT pin for the interrupt
 - The INT pin connects hardware devices such as keyboards, NIC cards
 - OS can invoke the keyboard interrupt handler routine to do interrupt
 - **Multiple hardware devices share a single INT pin** using an interrupt controller
 - To determine which device produced the interrupt, the processor contacts the interrupt controller.

Difference between the trap and interrupt?

Trap	Interrupt
A signal emitted by a user program	A signal emitted by a hardware device
Synchronous process	Asynchronous process
Can occur only from software device	Can occur from a hardware or a software
Only generated by a user program ISA	Generated by an OS and user program ISA
Traps are subset of interrupts	Interrupts are superset of traps
Execute a specific functionality in the OS and gives the control to the trap handler	Force the CPU to trigger a specific interrupt handler routine

Exceptions

- Exception: Event that needs to be processed by the OS kernel. The event is usually unexpected or rare



Exception handler
(in OS kernel)

Causes for exceptions

- **Exceptions**

- Synchronous events generated by the process itself
- E.g. illegal instructions, divide-by-0, illegal memory address

- **Interrupts**

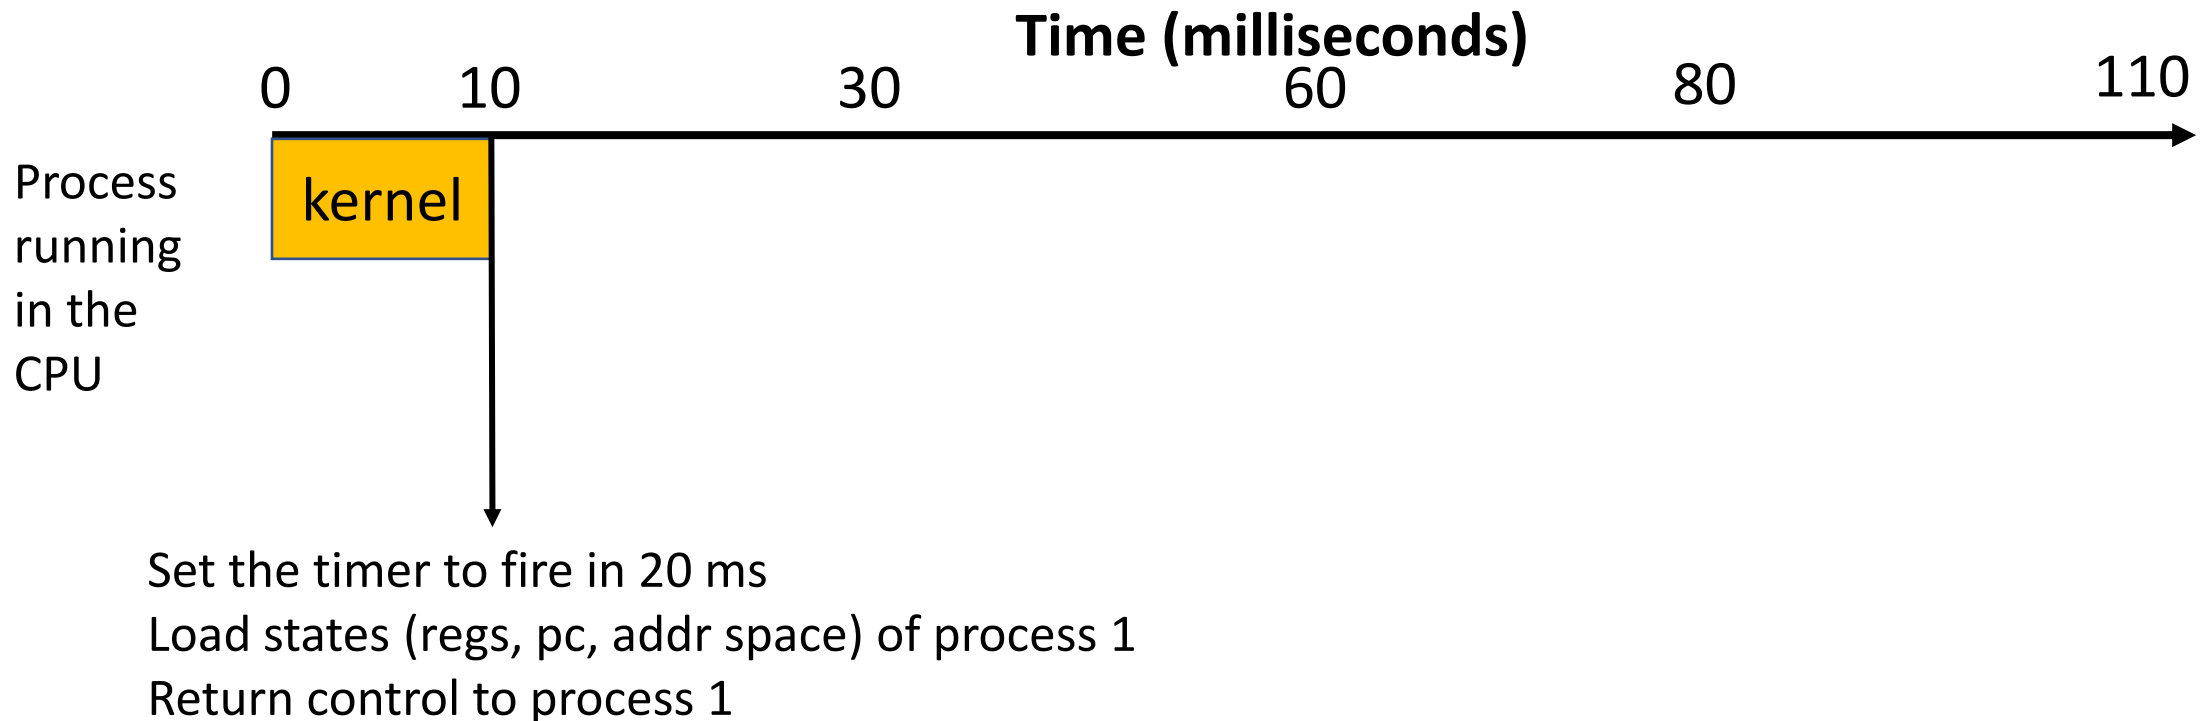
- Asynchronous events generated by I/O devices
- E.g. timer expired, keystroke, packet received, disk transfer complete

Handling exceptions

- When an exception happens, the processor
 - Stop the current process at instruction I_i , completing all the instructions up to I_{i-1} (precise exceptions)
 - Saves the PC of instruction I_i and the reason for the exception in special (privileged) register
 - Enable supervisor mode, disable interrupts, and transfers control to a pre-specified exception handler PC
- After the OS kernel handles the exception, it returns control to the process at instruction I_i
 - Exception is transparent to the process
- If the exception is due to an illegal operation by the program that cannot be fixed, the OS aborts the process

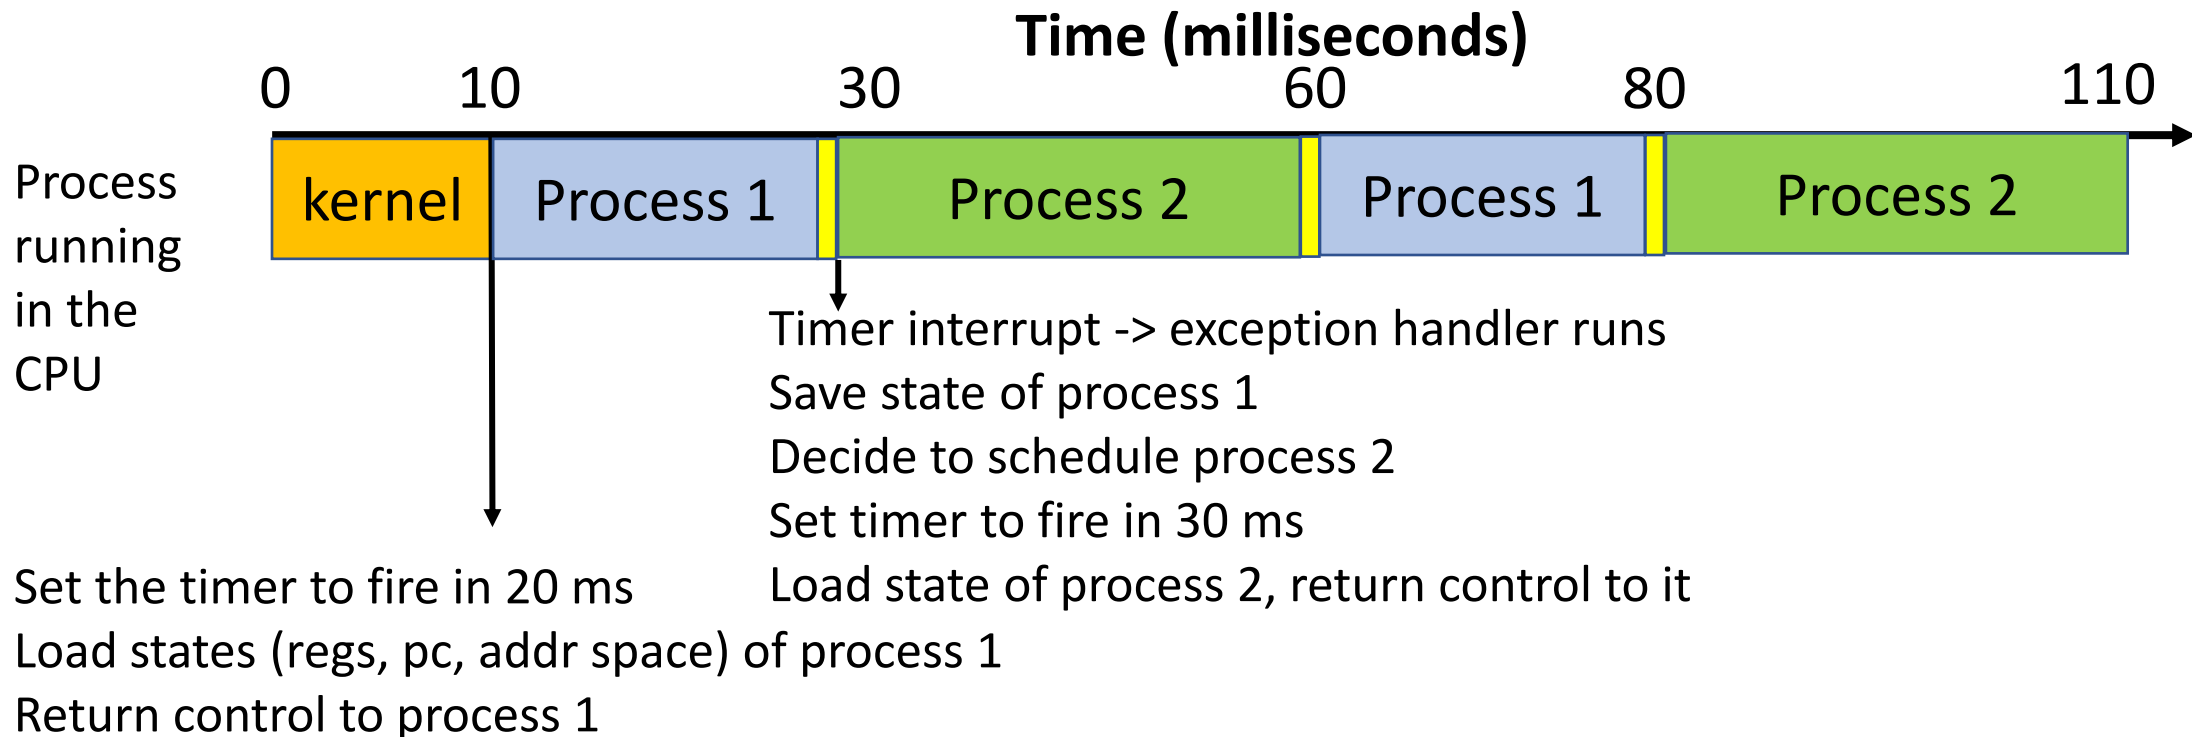
Case study 1: CPU scheduling

- The OS kernel schedules processes into the CPU
 - Each process is given a fraction of CPU time
 - Enabled by timer interrupts
 - Kernel sets timer, which raises an interrupt after a specified time



Case study 1: CPU scheduling

- The OS kernel schedules processes into the CPU
 - Each process is given a fraction of CPU time
 - Enabled by timer interrupts
 - Kernel sets timer, which raises an interrupt after a specified time



Kernel organization

- **Monolithic kernel**

- Entire operating system resides in the kernel space
- The implementations of all system calls run in kernel mode
- E.g. Unix, Linux


- **Good**

- Easy for subsystems to cooperate
- One cache shared by file system and virtual memory

- **Bad**

- Interactions are complex
- Mistake is fatal because an error in kernel model will result in the kernel to fail
- No isolation within kernel

Kernel organization



- **Microkernel**

- Move most OS functionality to user-space
- Kernel can be small, mostly IPC
- The hope:
 - Simple kernel can be fast and reliable

- **Microkernel wins:**

- Fast IPC
- separate services force kernel developers to think about modularity

- **Microkernel losses:**

- kernel can't be tiny: needs to know about processes and memory
- it's hard to split the kernel into lots of service processes!

Kernel organization

- **Exokernel (1995)**

- Philosophy: eliminate all abstractions, let app do what it wants
- An exokernel would not provide address space, pipes, etc.
- libOSes implements abstractions, each app can have its own custom libOS
- Why? Kernel may be faster due to streamlining, simplicity
- Apps may be faster because they can customize libOS

