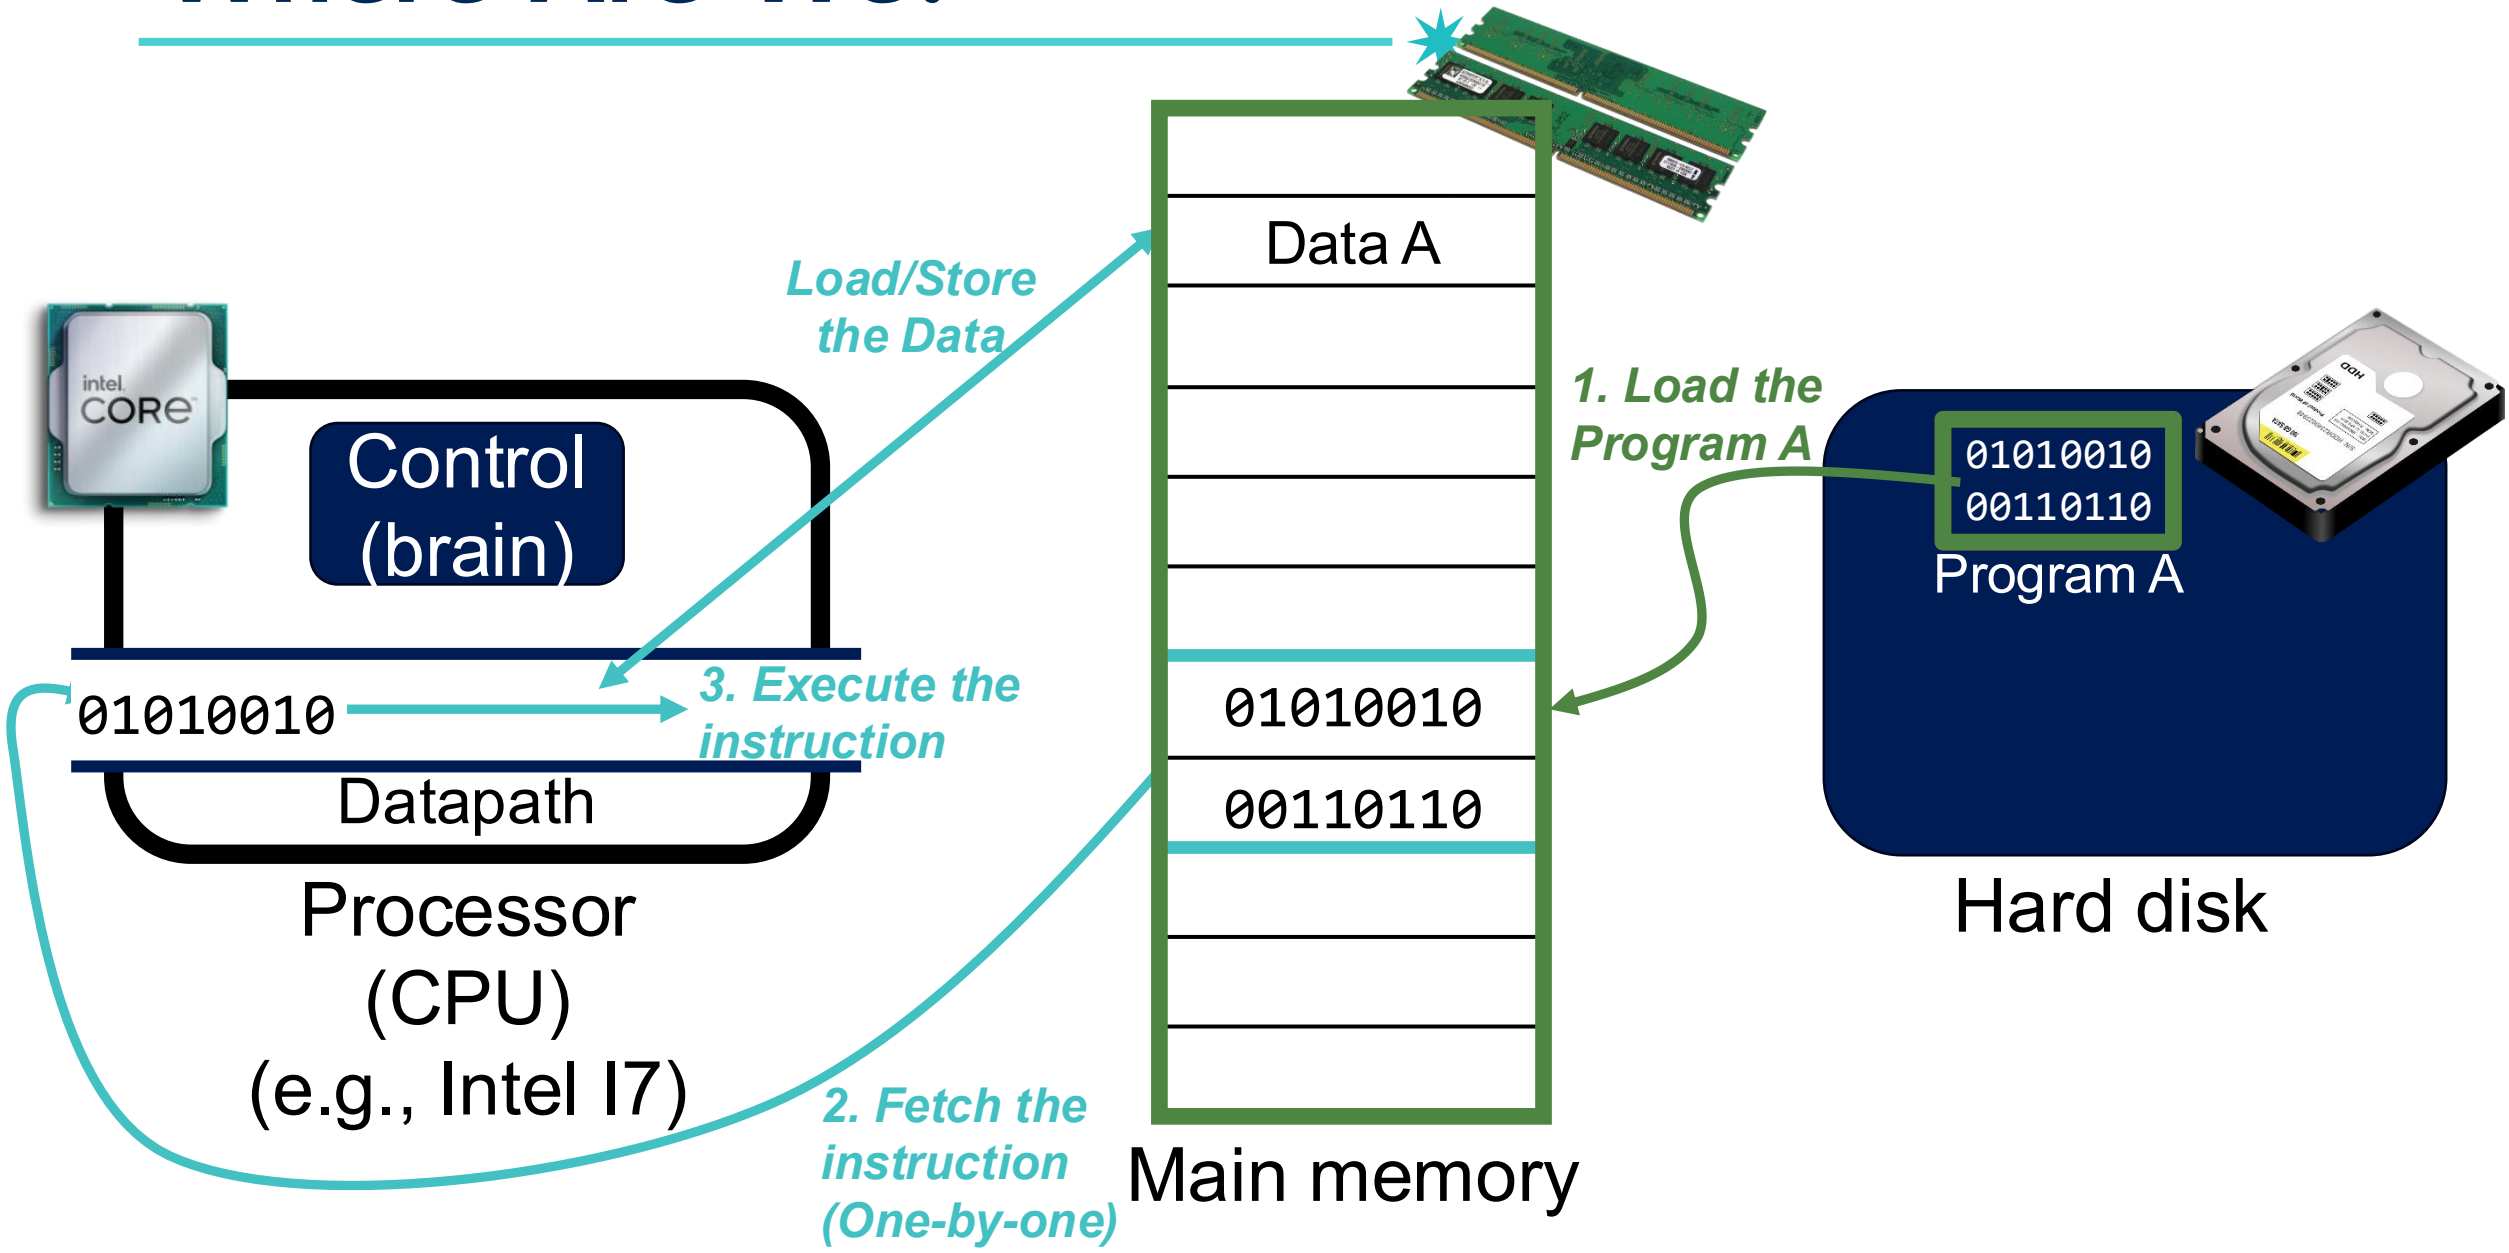


CSE251: System Programming

21. Process (1)

Seongil Wi

Where Are We?



Where Are We?

How data (bit, byte, int, ...) is represented in memory?

Stack, heap, code, data regions

01010010

How is assembly code interpreted, and how does it interact with memory?

Executable and Linkable Format (ELF), Linking Process

1. Load the Program A

01010010
00110110
Program A

Loading process via loader (kernel)

Hard disk

Memory hierarchy (cache, virtual memory)

Main memory

Load/Store



Use the Datapath

(One-by-one)

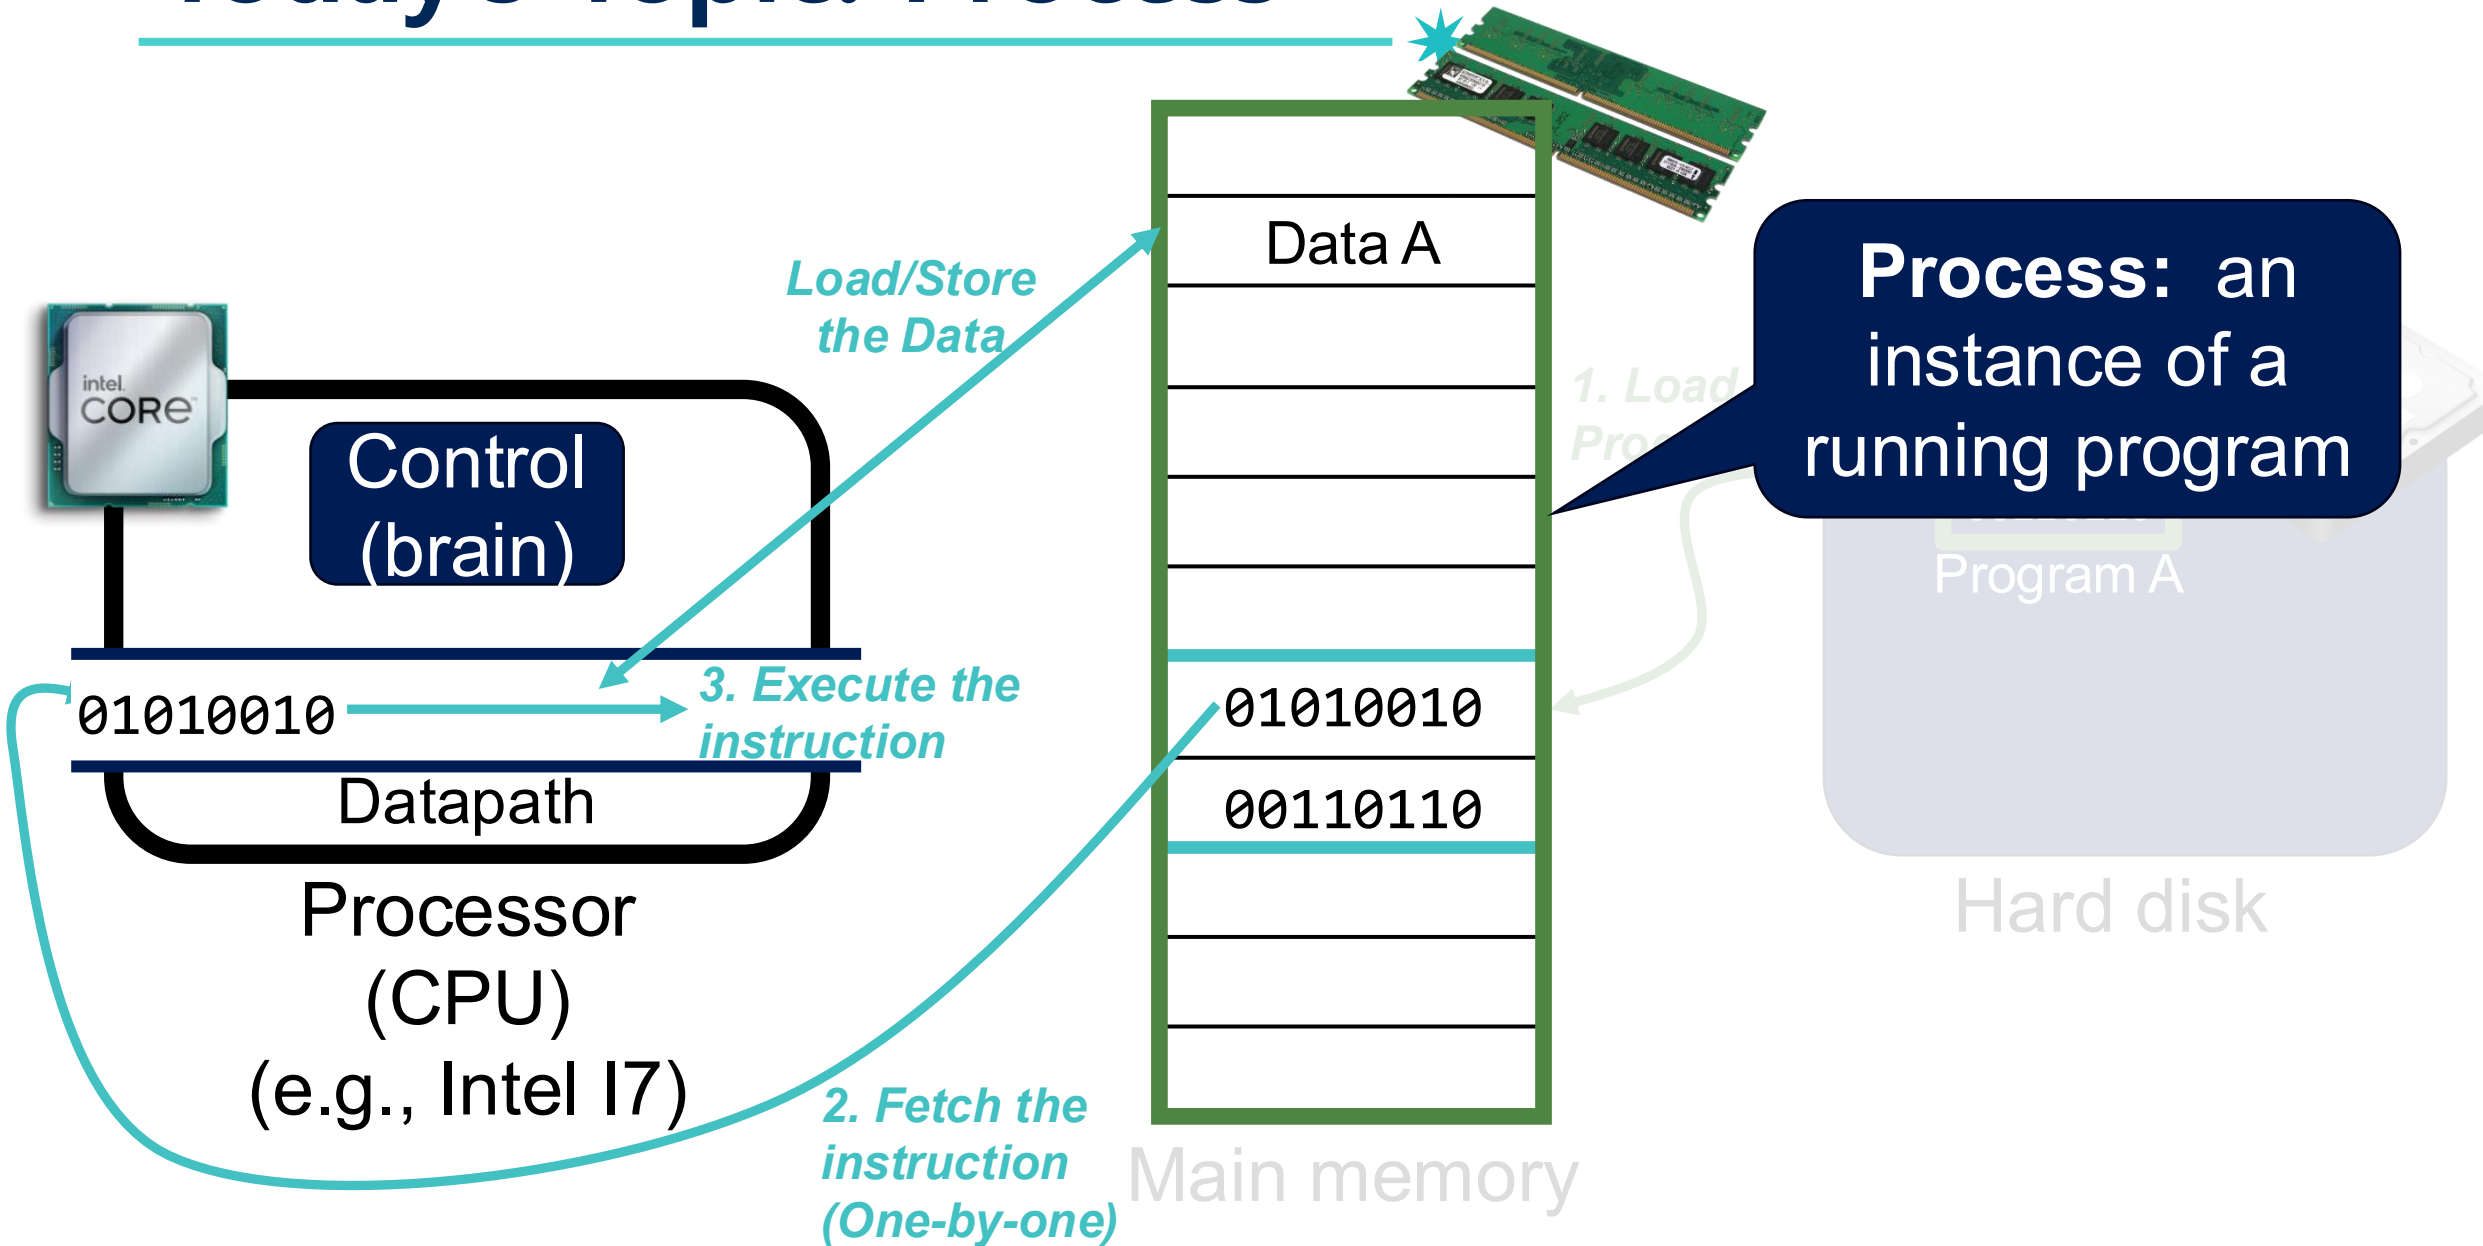
Datapath

Processor

(Brain)



Today's Topic: Process



Processes

Process



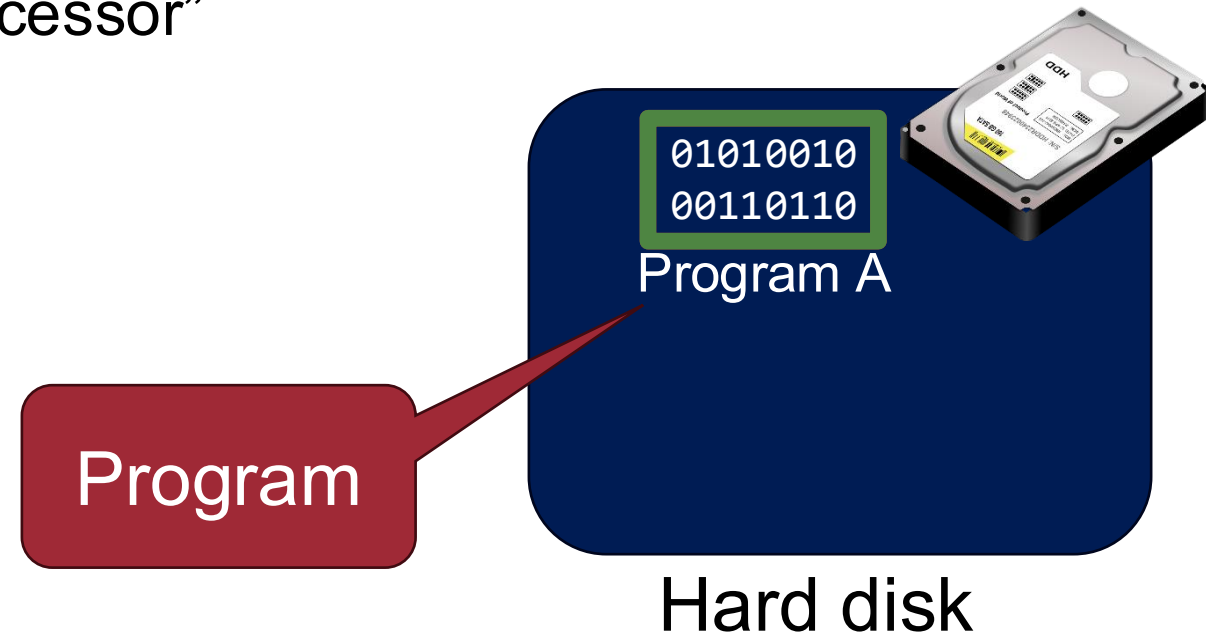
An instance of a program that is being executed (or is ready for execution)

Process



An instance of a program that is being executed (or is ready for execution)

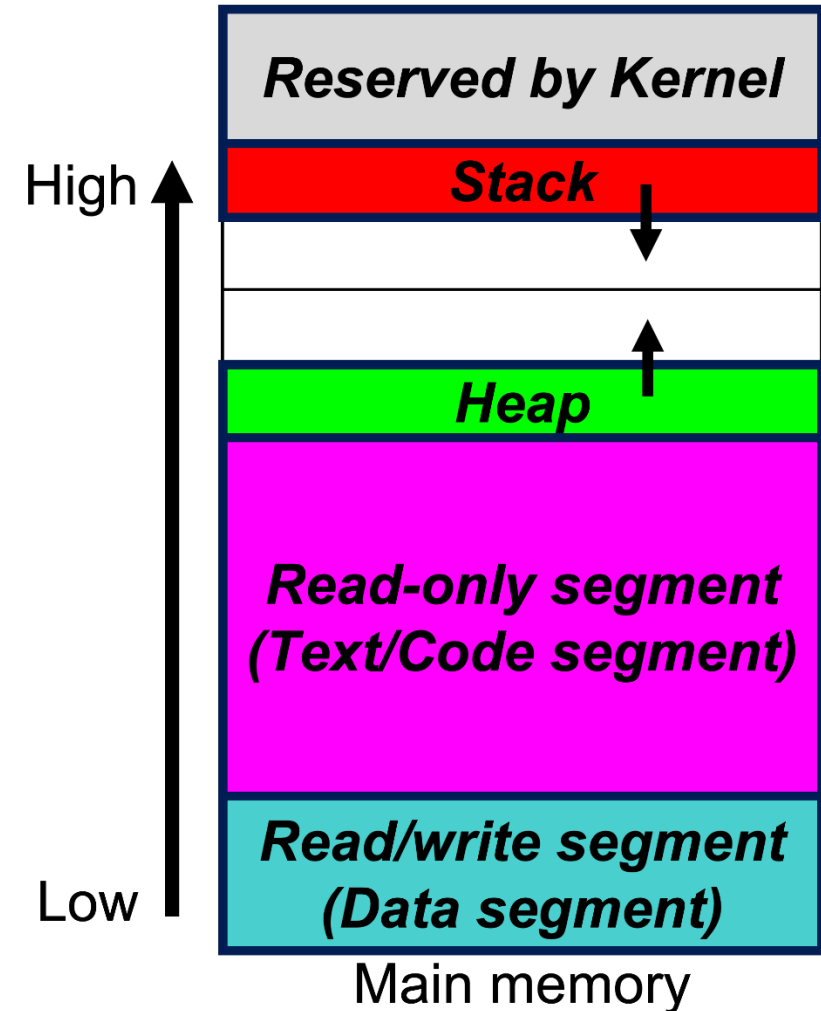
- Not the same as “program” or “processor”



Process

An instance of a program that is being executed (or is ready for execution)

- Not the same as “program” or “processor”
- Consists of:
 - Memory (code, heap, stack, etc.)




Process

An instance of a program that is being executed (or is ready for execution)

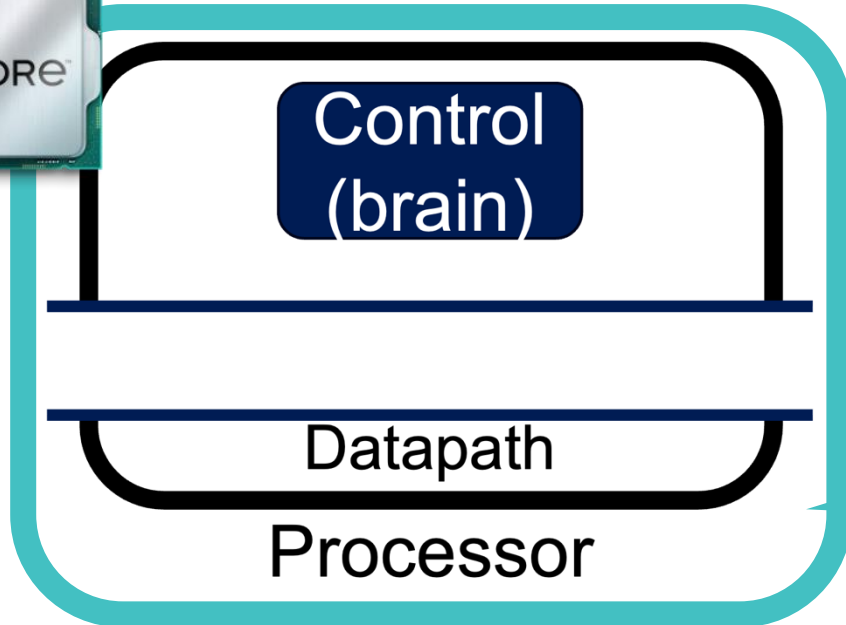
- Not the same as “program” or “processor”

- Consists of:

- Memory (code, heap, stack, etc.)
- Registers used to manage execution (stack pointer, program counter, ...)
- Other resources




Traditional Model: Uniprocessor




Uniprocessor: one processor (core) per chip

Traditional Model: Uniprocessor



Only do one thing: simply reads and executes a sequence of instructions, one at a time

Traditional Model: Uniprocessor




Q. How can **multiple programs run concurrently** on a single processor?




Computer runs many processes simultaneously

- *Applications for one or more users*
 - *Web browsers, email clients, editors, ...*
- *Background tasks*
 - *Monitoring network & I/O devices*

Traditional Model: Uniprocessor



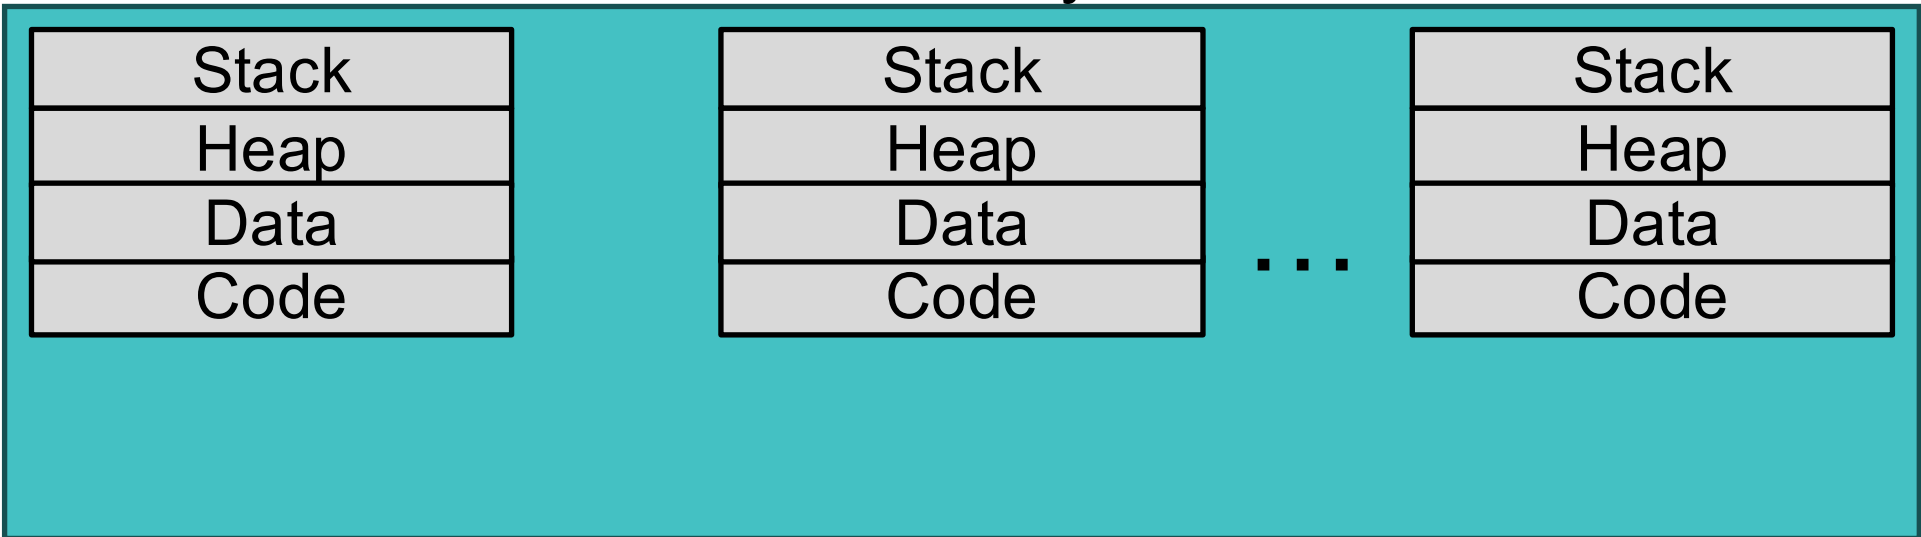


Q. How can **multiple programs run concurrently** on a single processor?

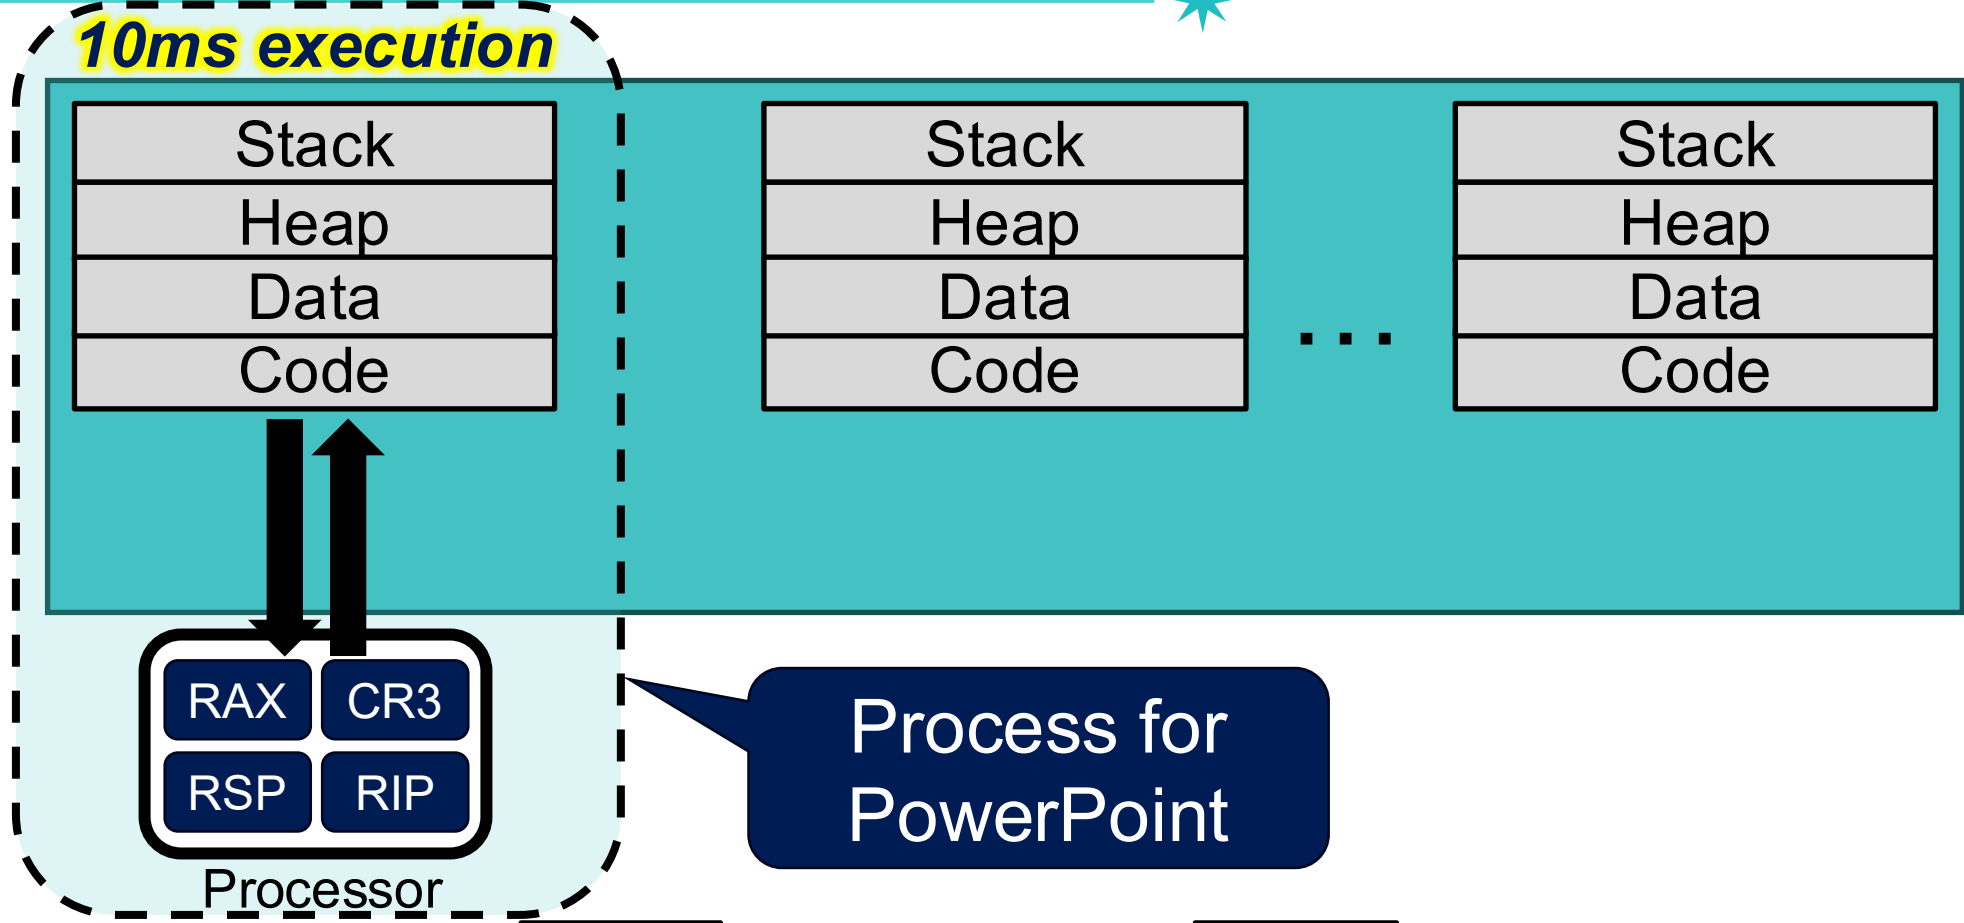


Solution: Multitasking

Memory



Solution: Multitasking



inst₁
inst₂
inst₃
inst₄
...

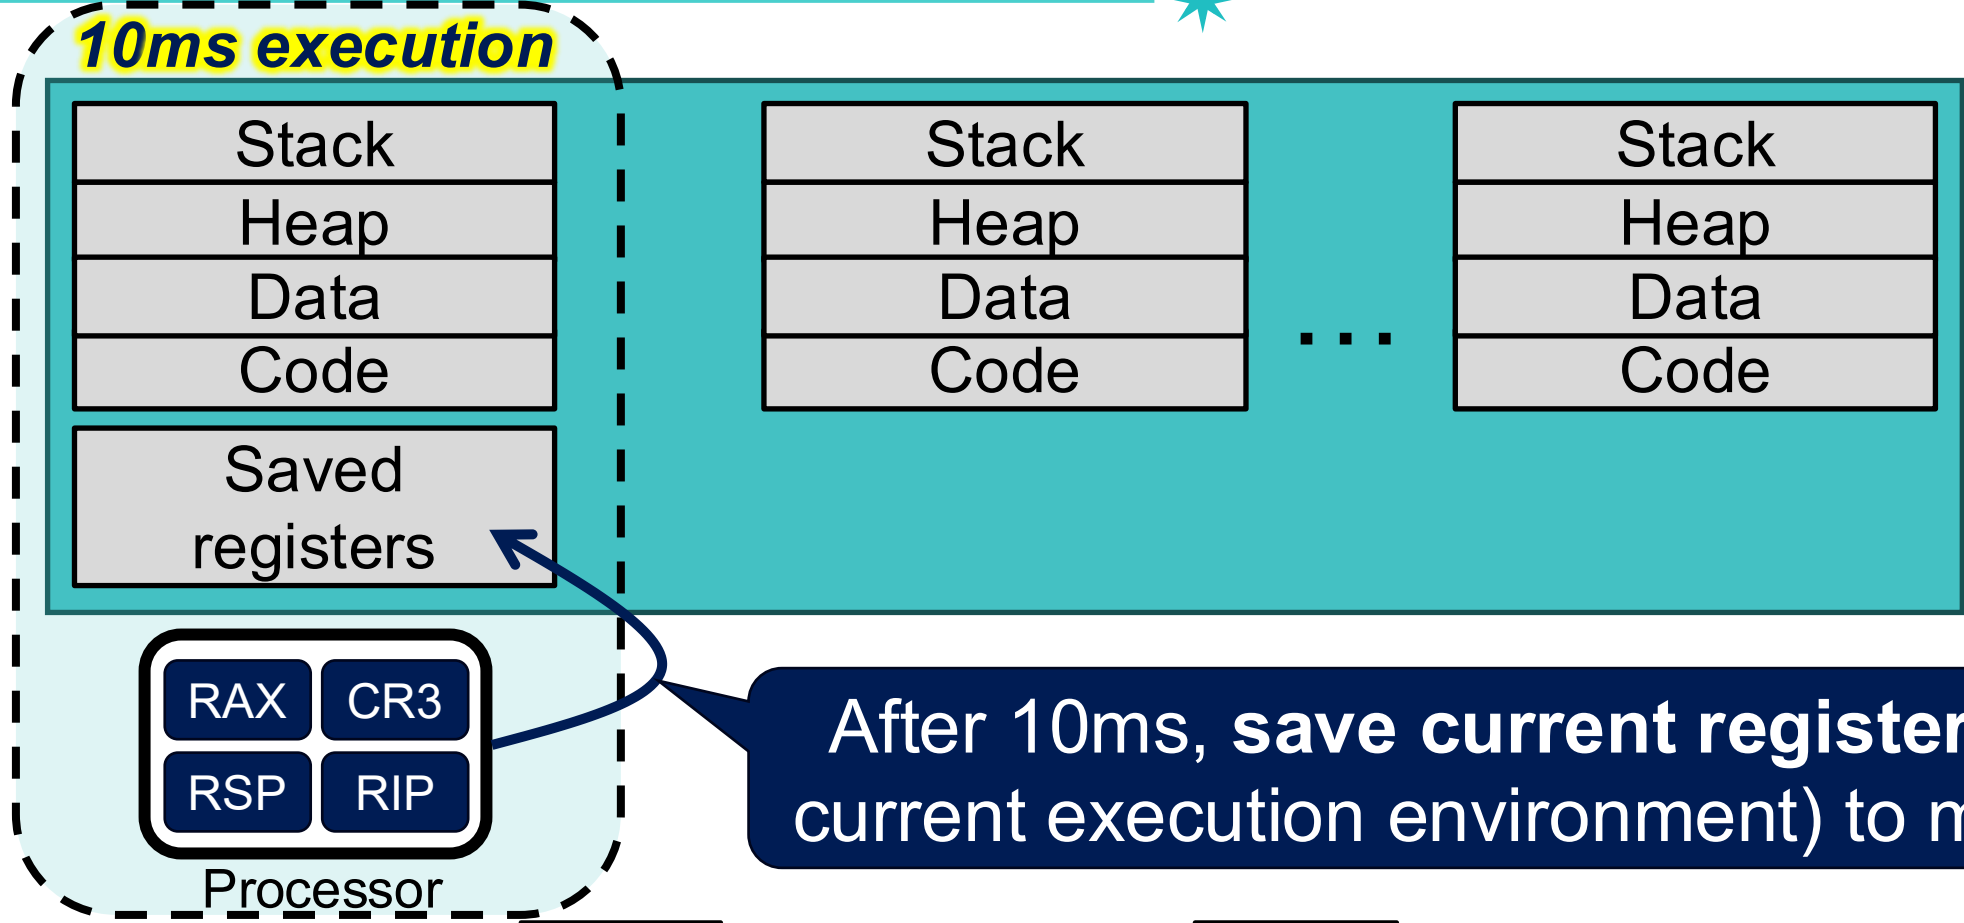


inst₁
inst₂
inst₃
inst₄
...



inst₁
inst₂
inst₃
inst₄
...

Solution: Multitasking



```
inst1  
inst2  
inst3  
inst4  
...
```

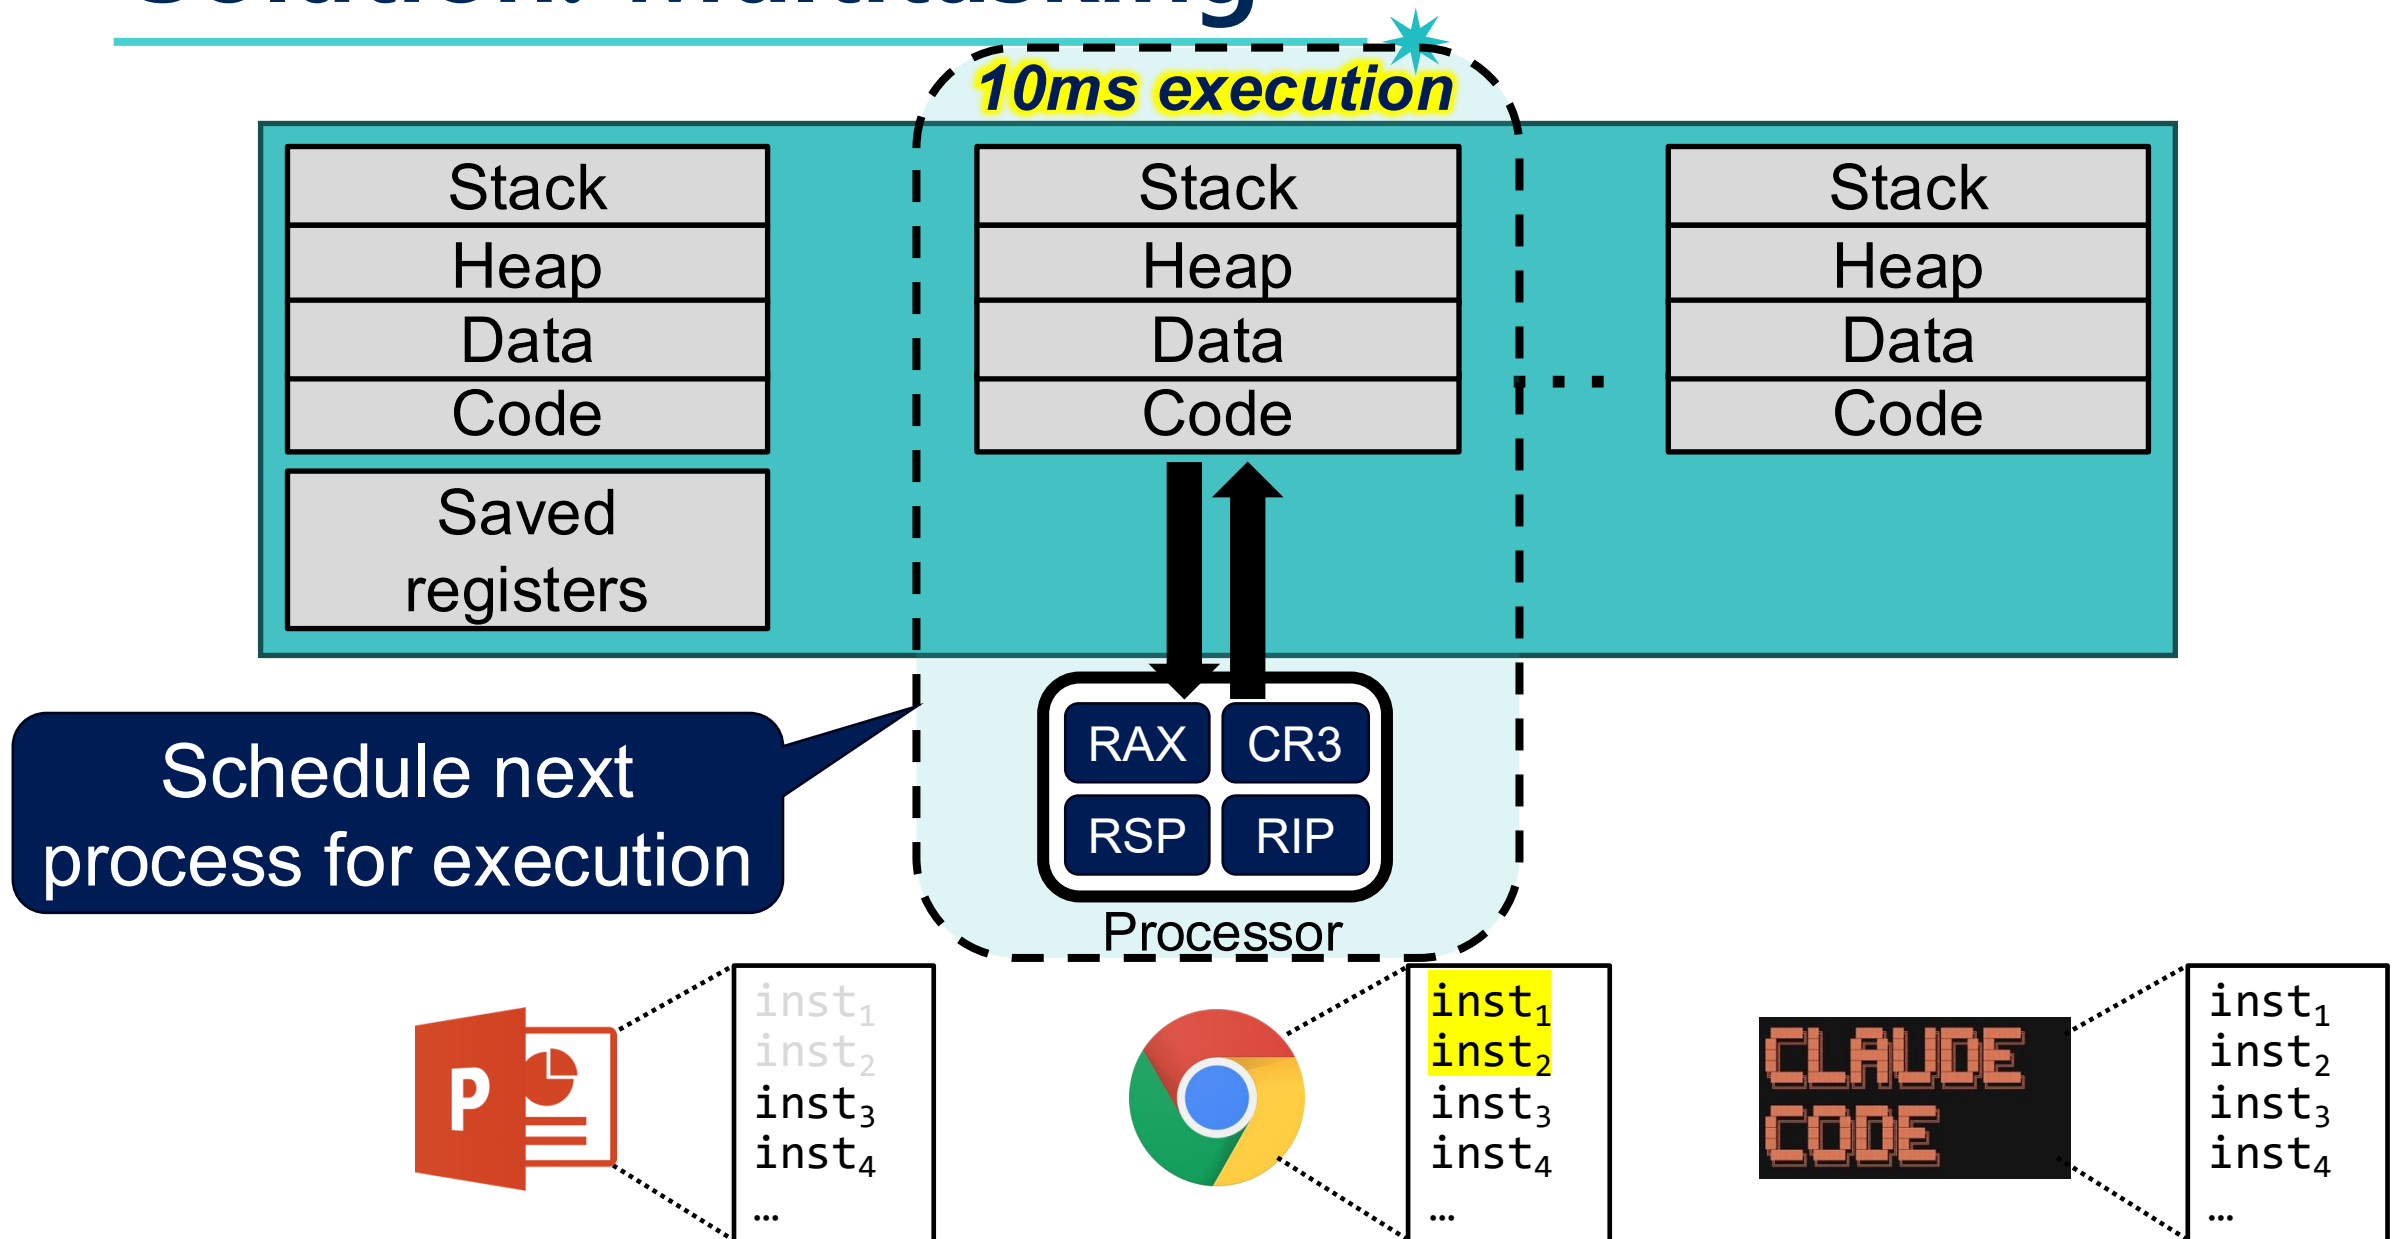


```
inst1  
inst2  
inst3  
inst4  
...
```

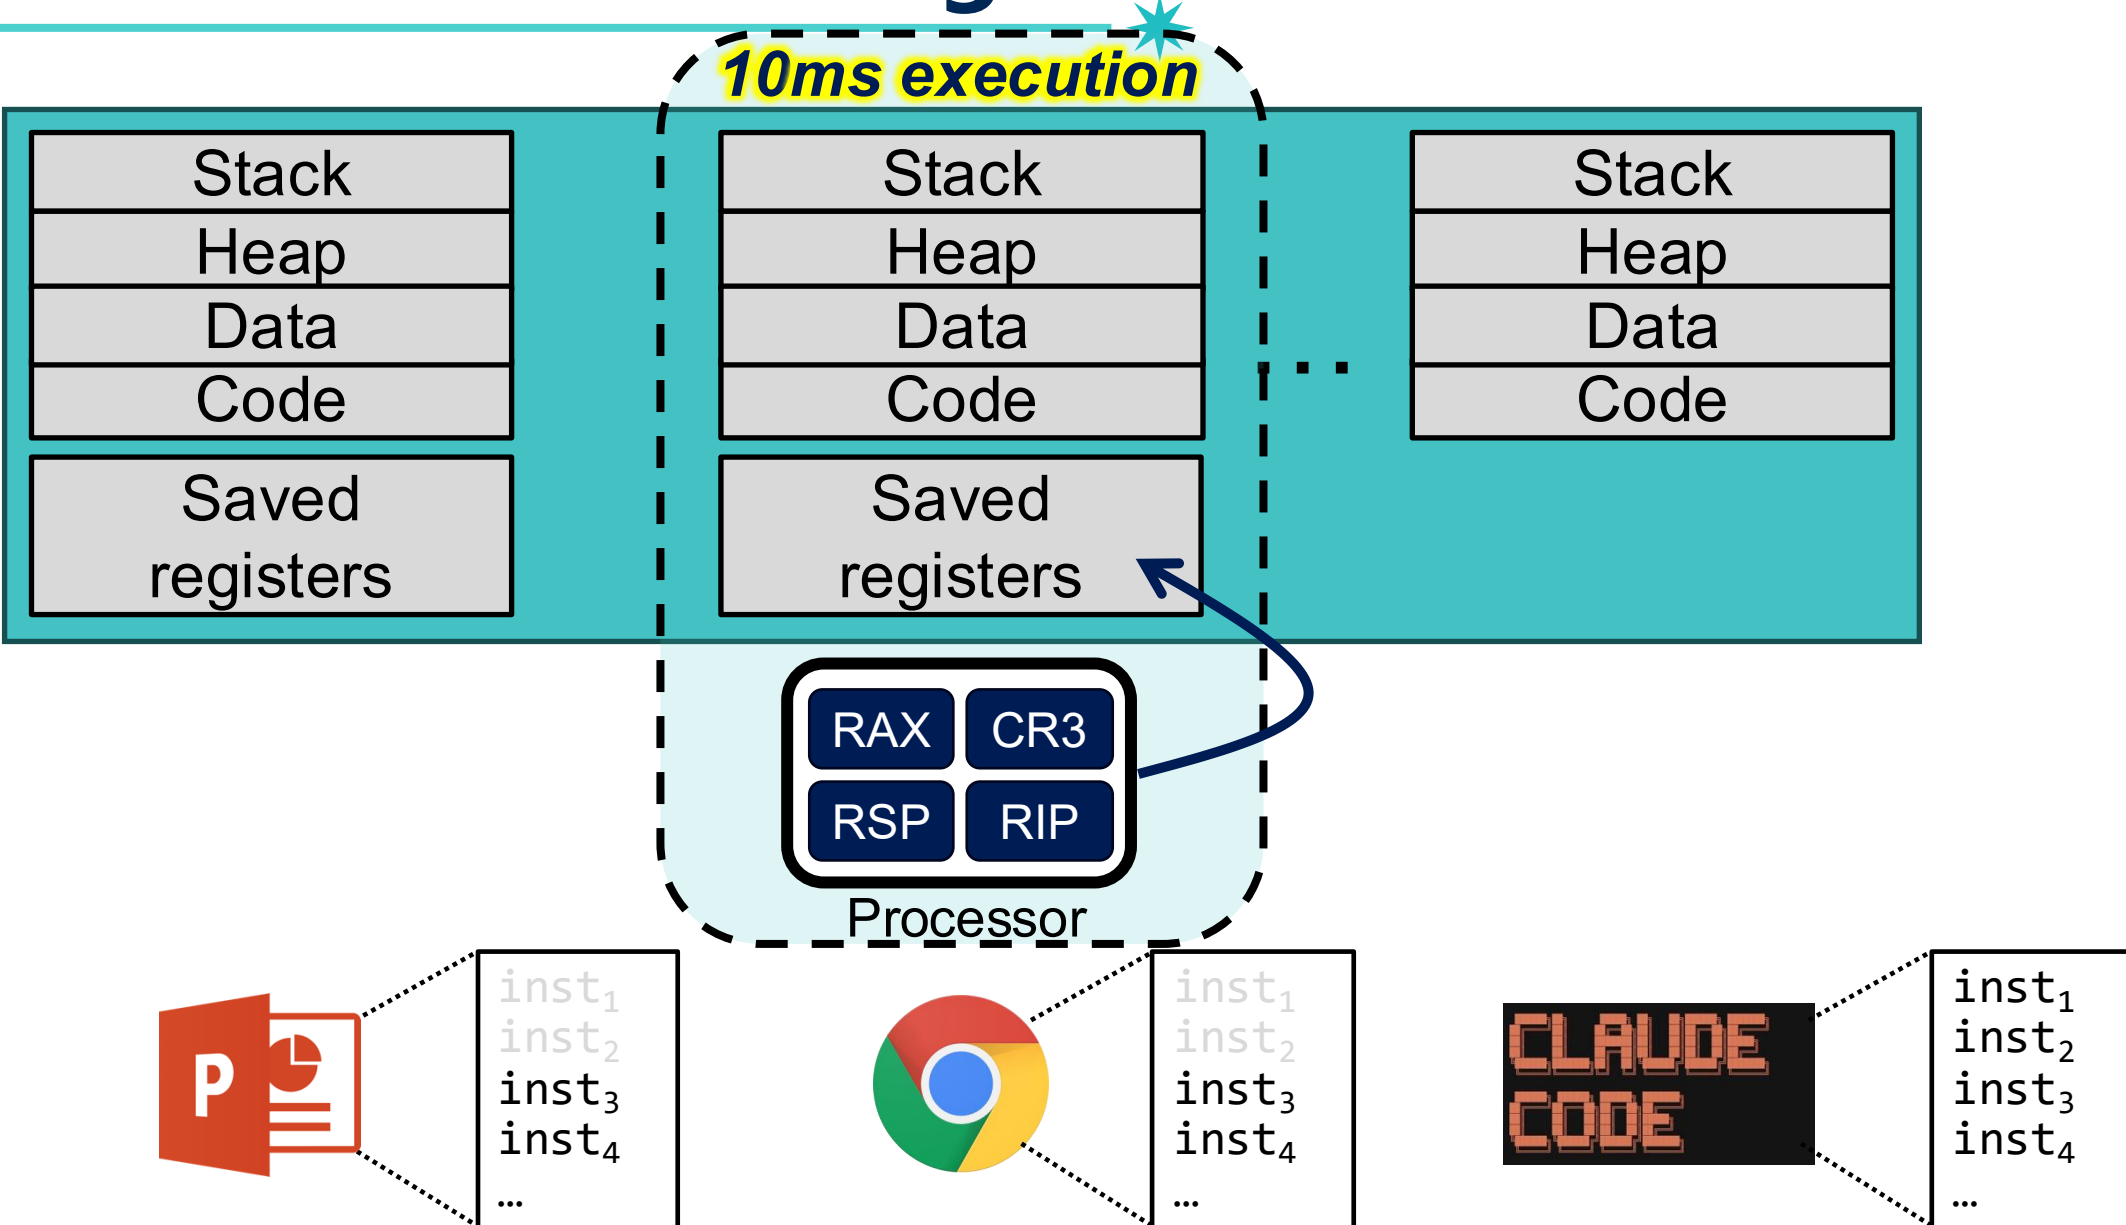


```
inst1  
inst2  
inst3  
inst4  
...
```

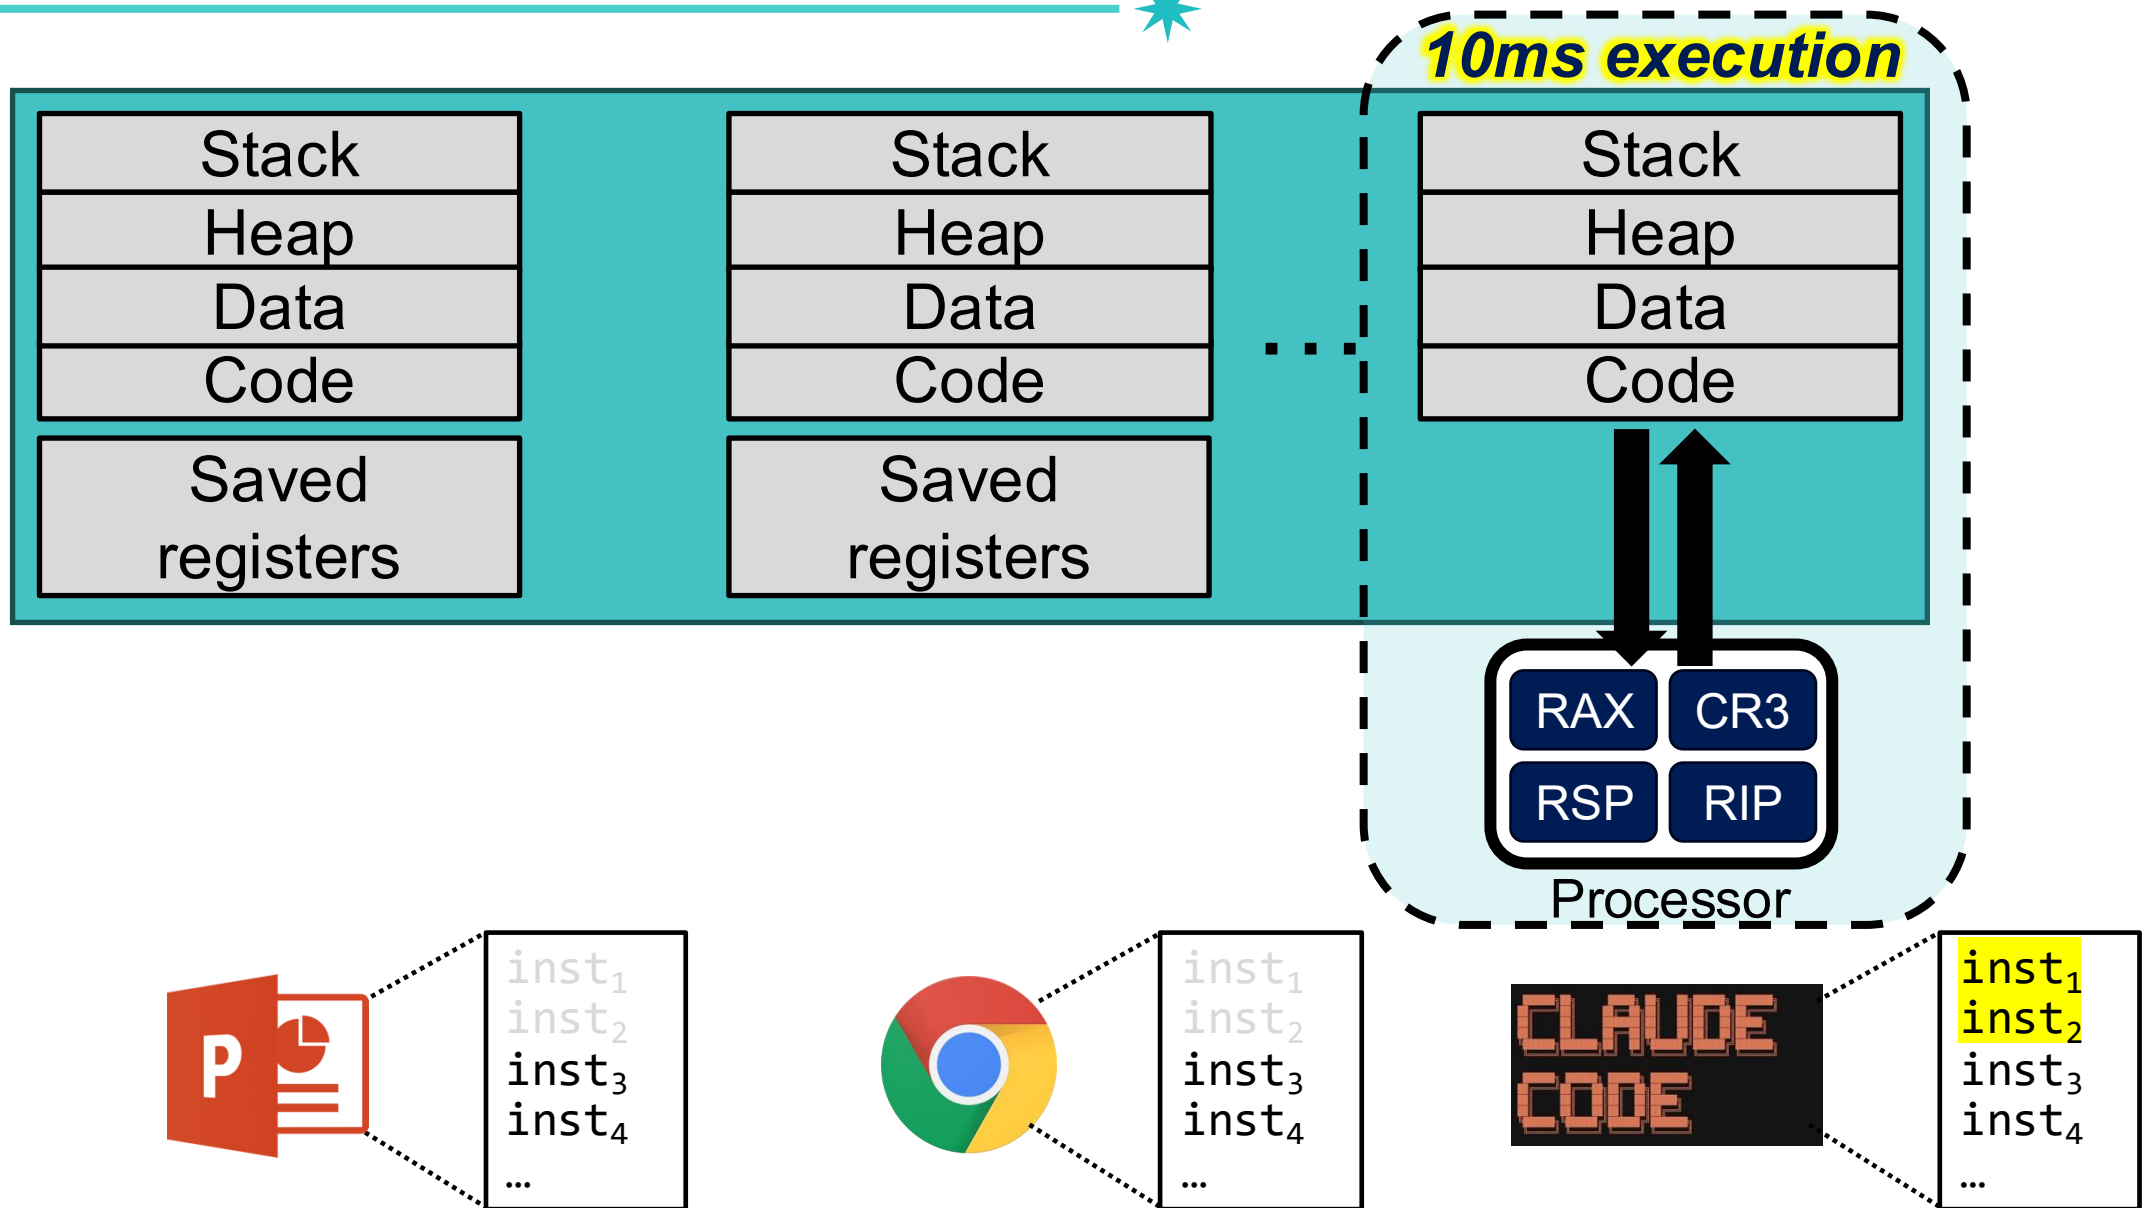
Solution: Multitasking



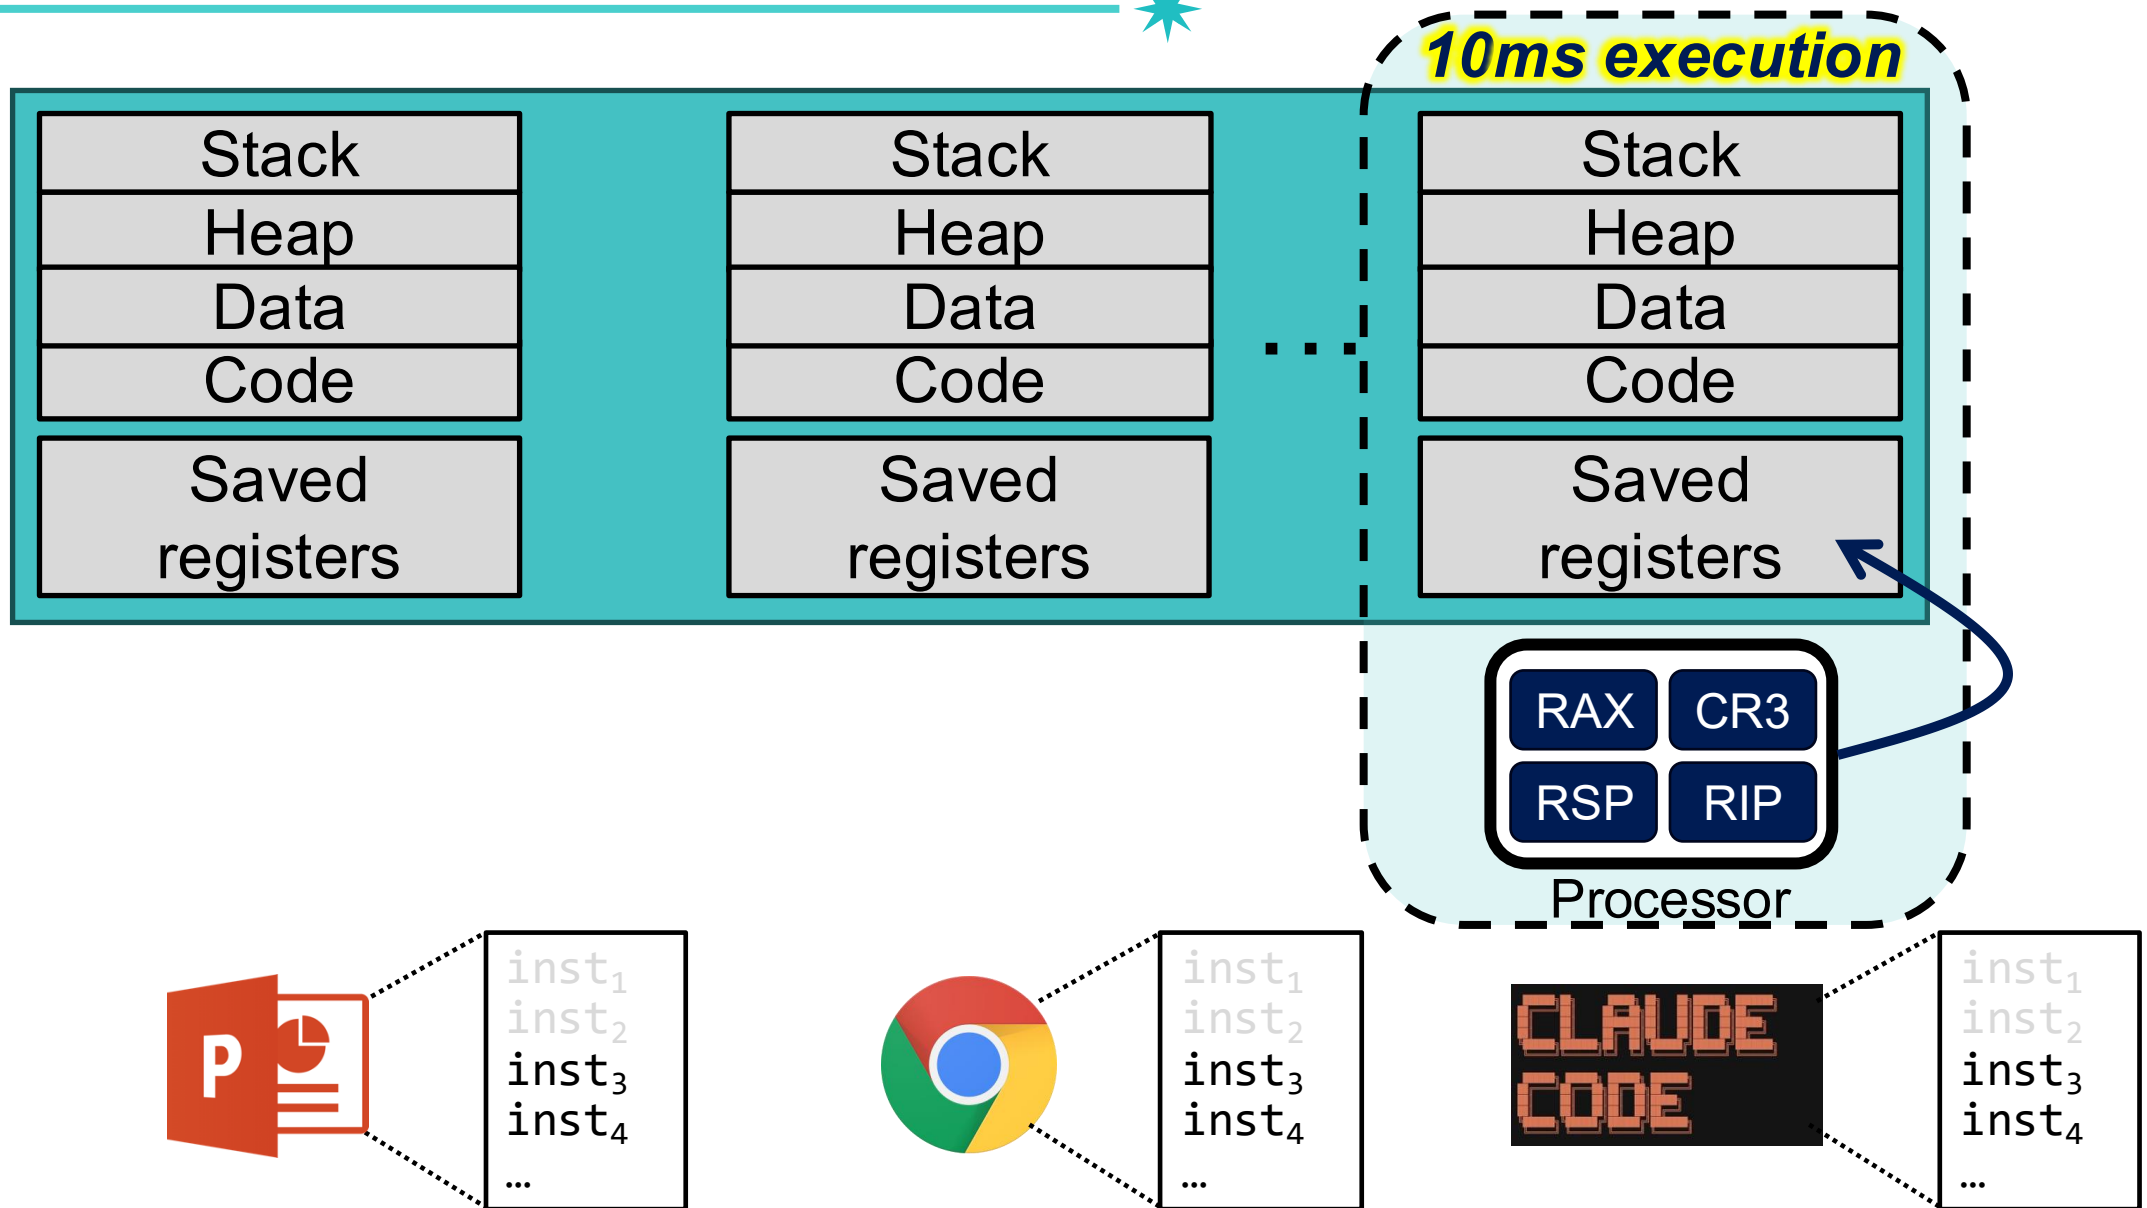
Solution: Multitasking



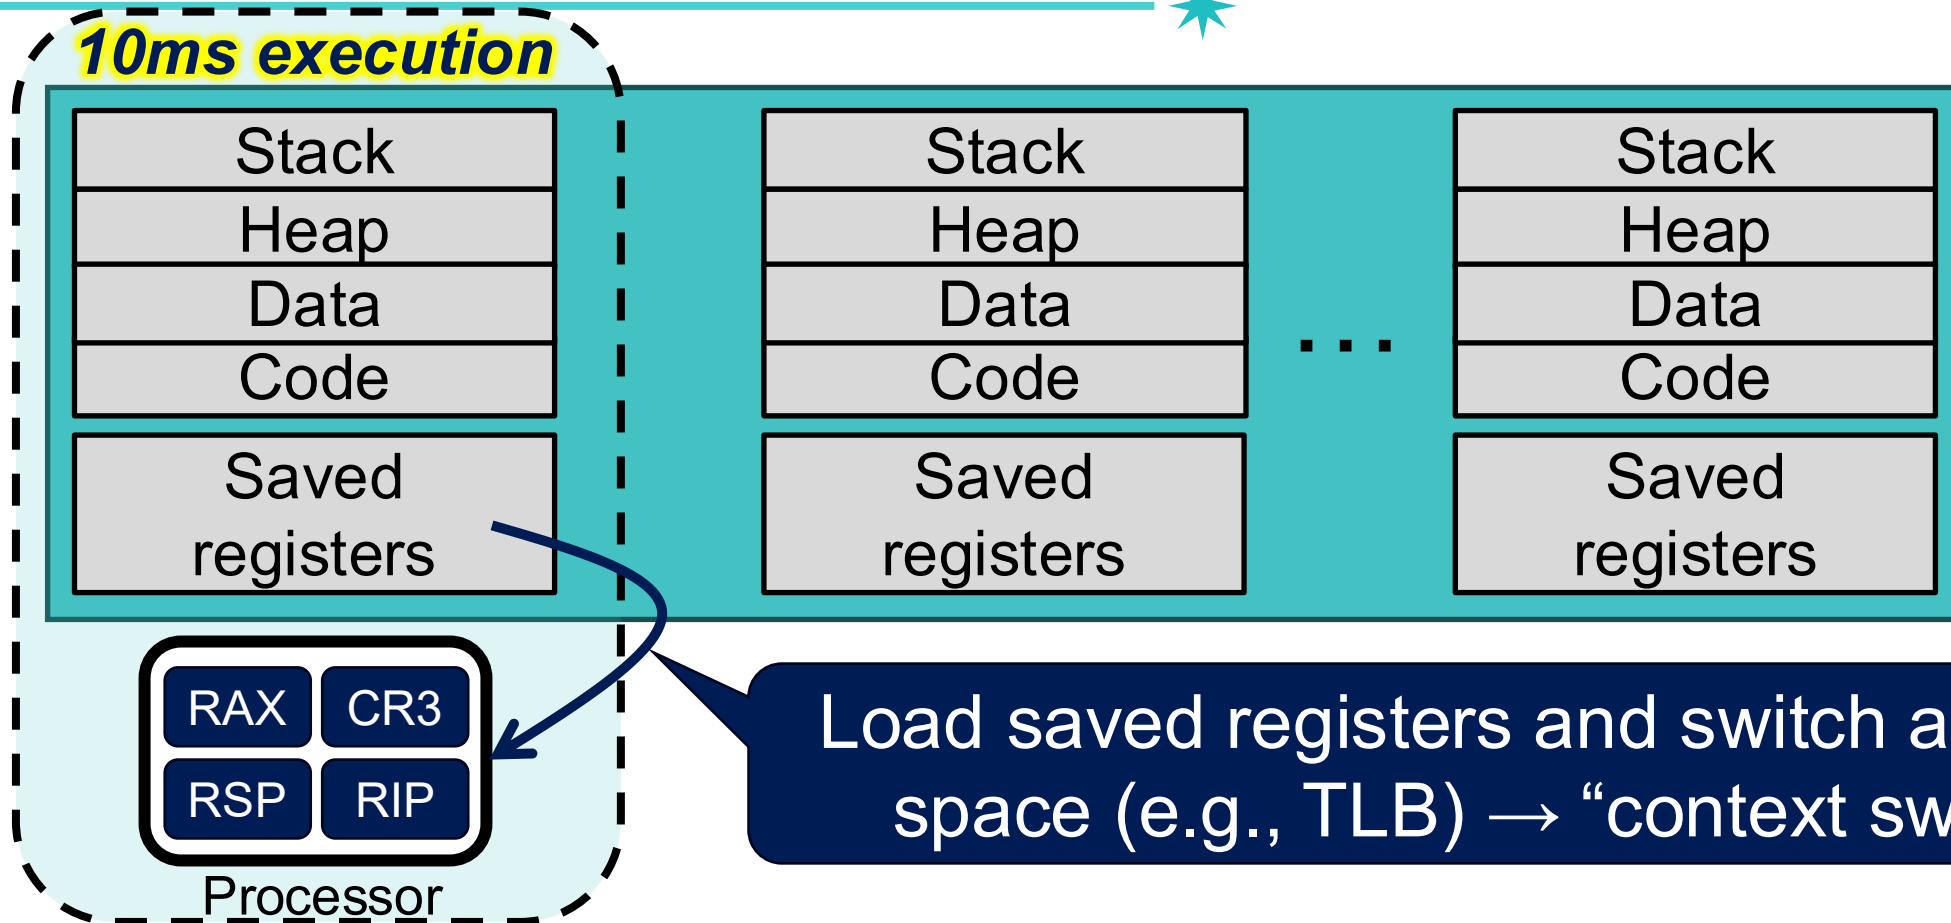
Solution: Multitasking



Solution: Multitasking



Solution: Multitasking



inst₁
inst₂
inst₃
inst₄
...

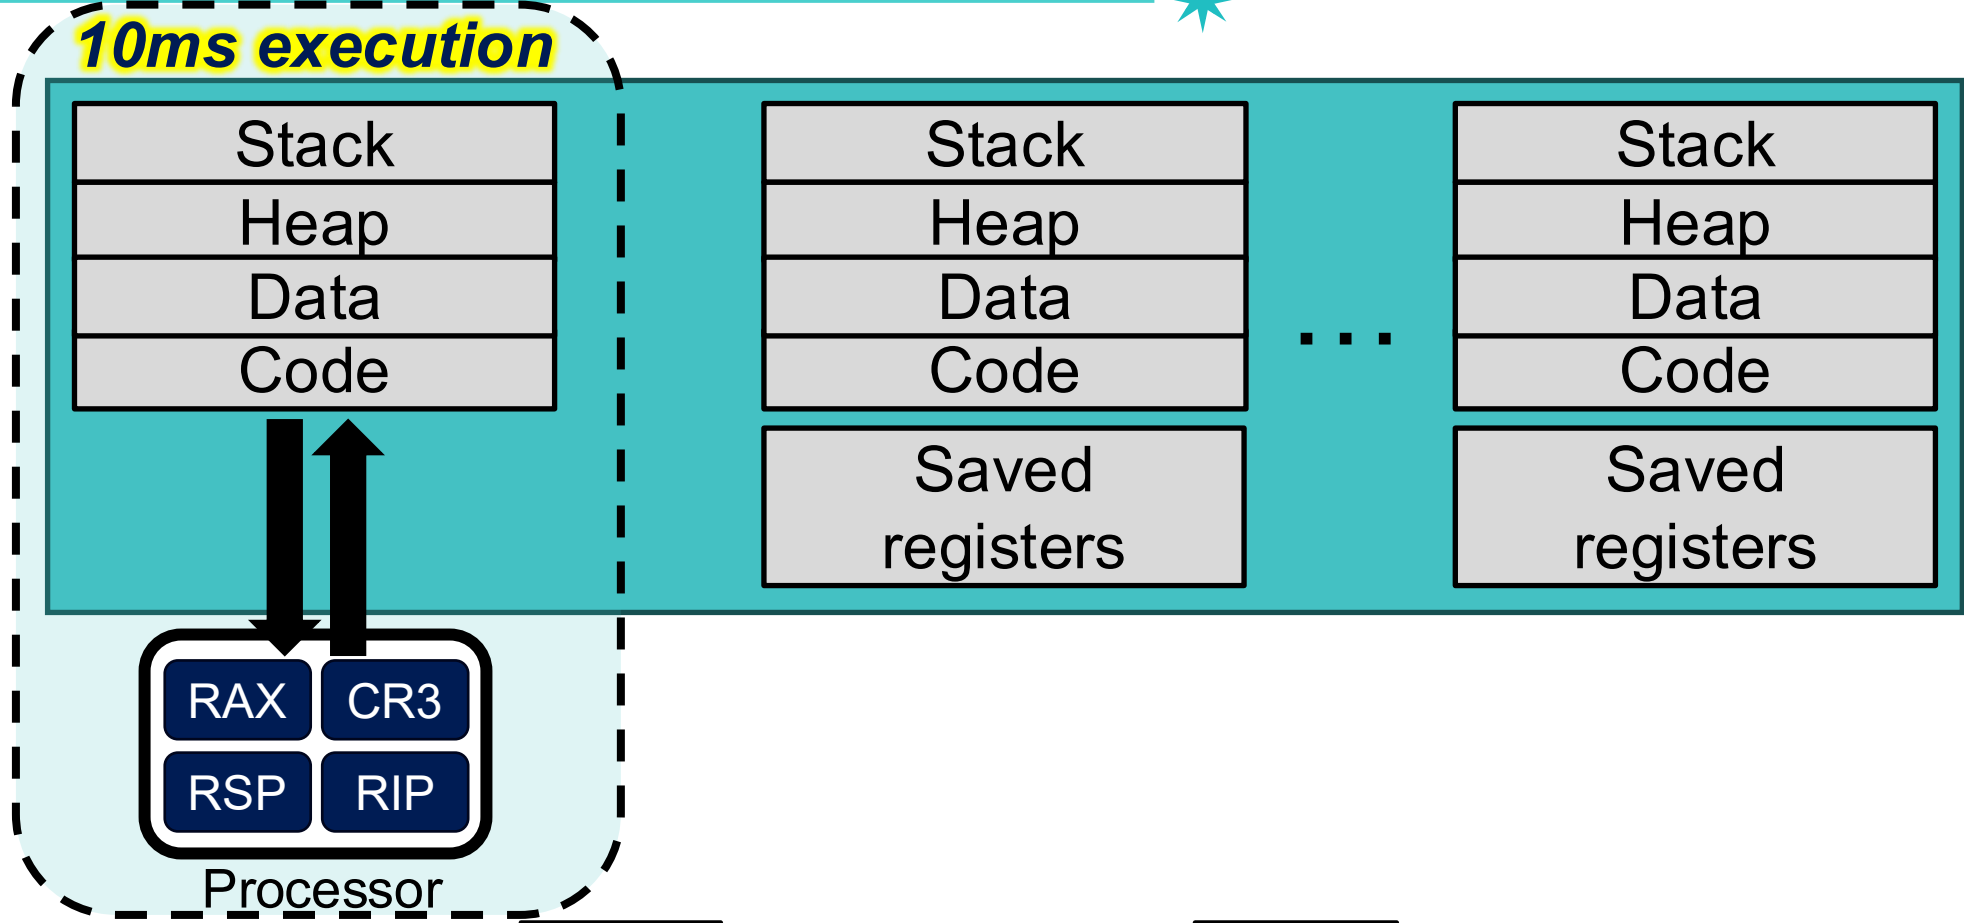


inst₁
inst₂
inst₃
inst₄
...



inst₁
inst₂
inst₃
inst₄
...

Solution: Multitasking



```
inst1  
inst2  
inst3  
inst4  
...
```



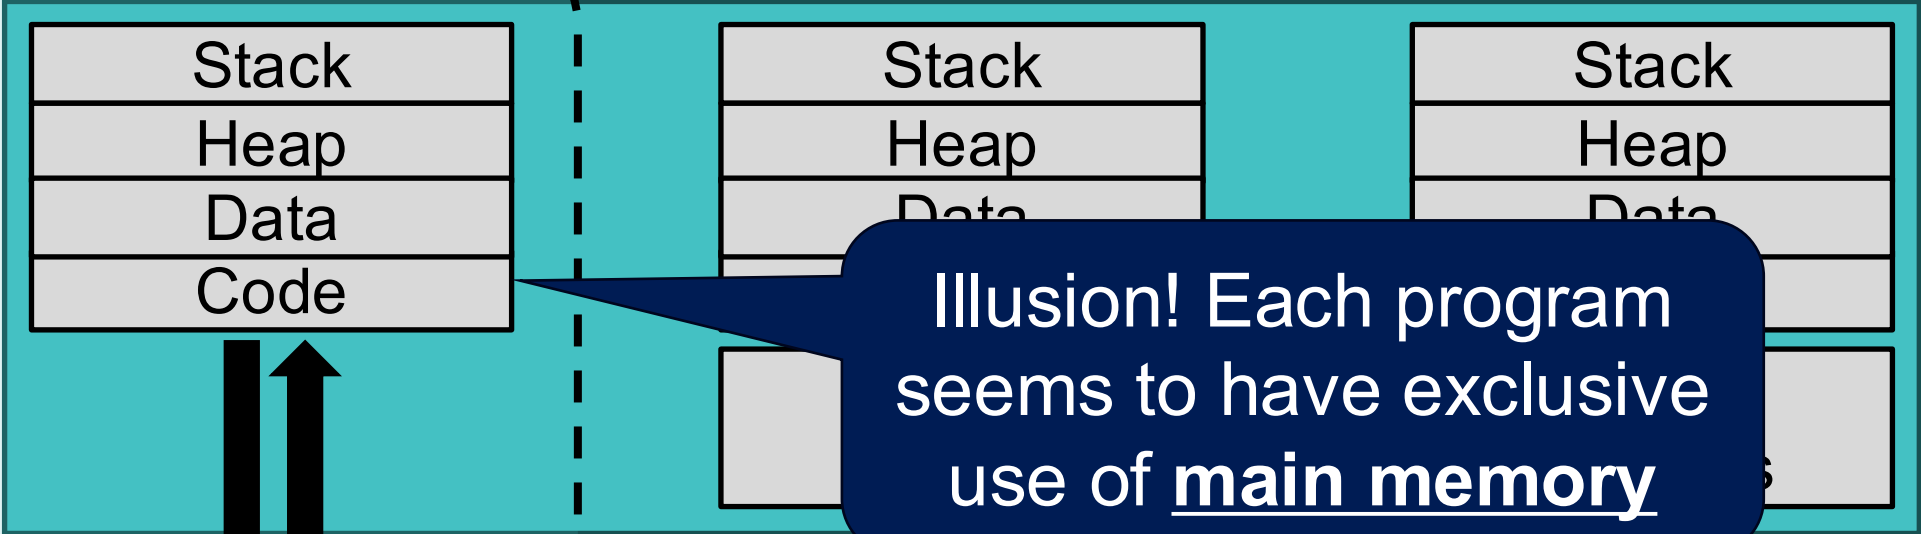
```
inst1  
inst2  
inst3  
inst4  
...
```




```
inst1  
inst2  
inst3  
inst4  
...
```

Solution: Multitasking

10ms execution



Illusion! Each program seems to have exclusive use of main memory



Processor

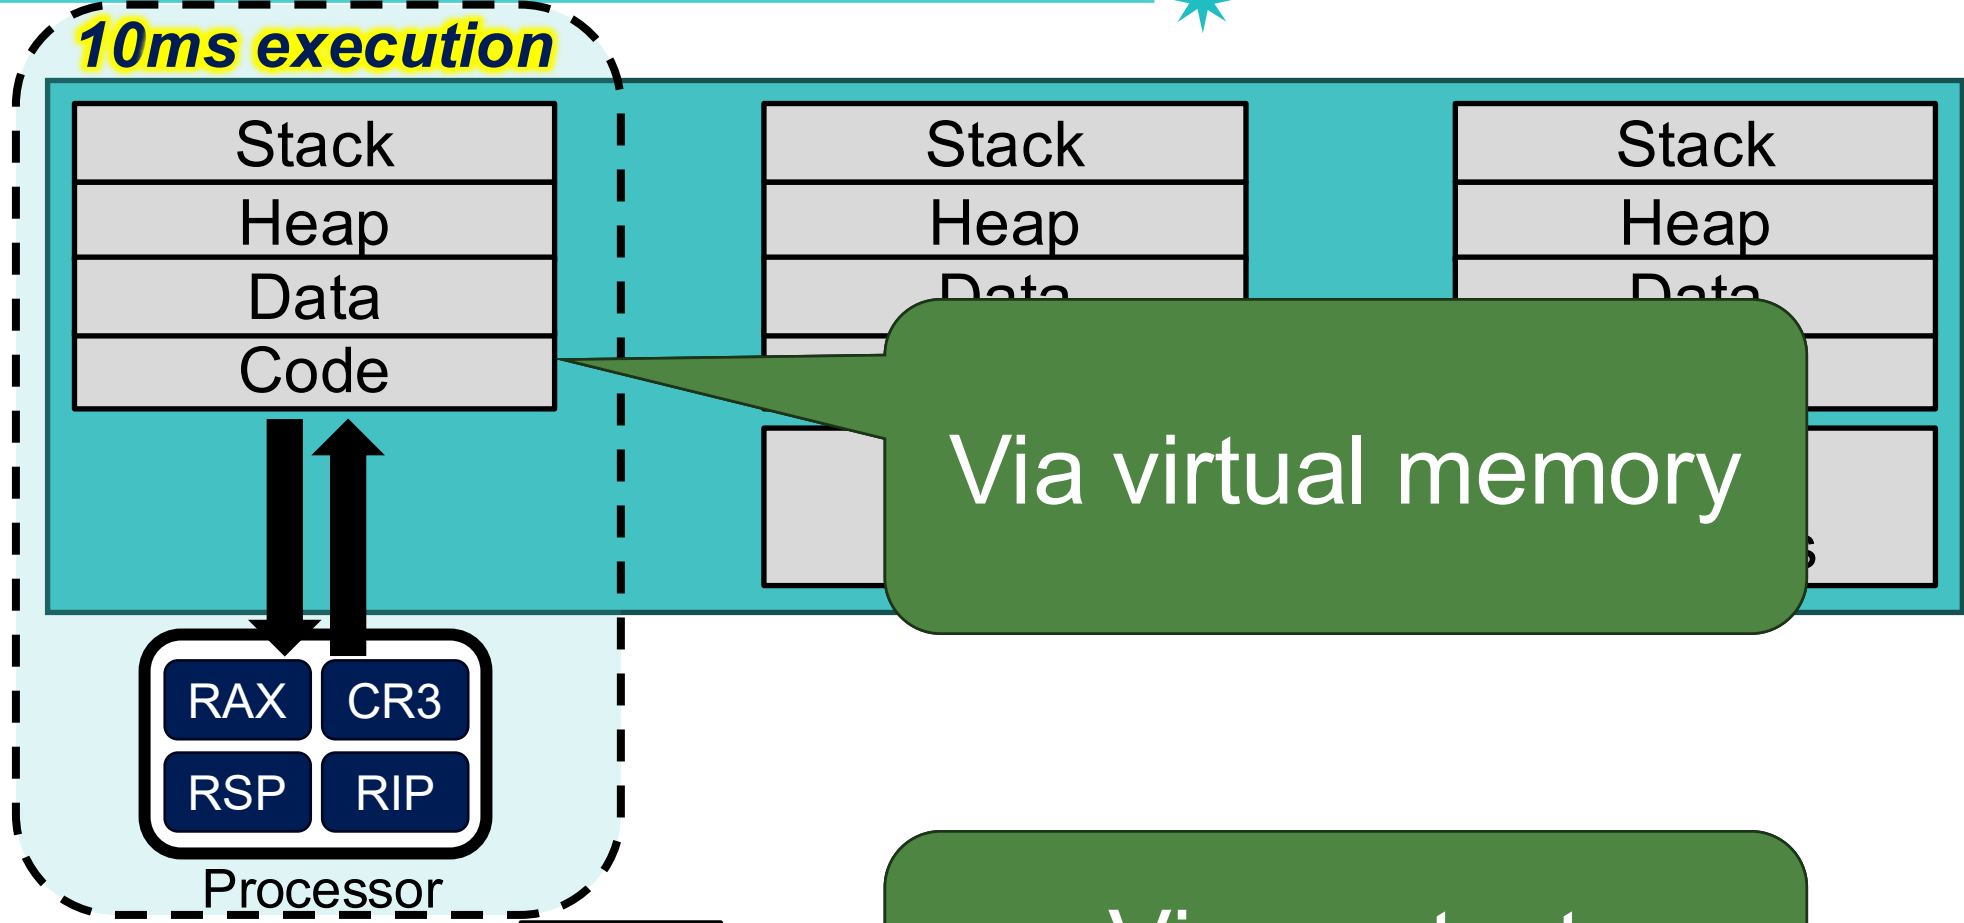


```
inst1  
inst2  
inst3  
inst4  
...
```

Illusion! Each program seems to have exclusive use of the CPU

```
inst1  
inst2  
inst3  
inst4  
...
```

Solution: Multitasking



Processor

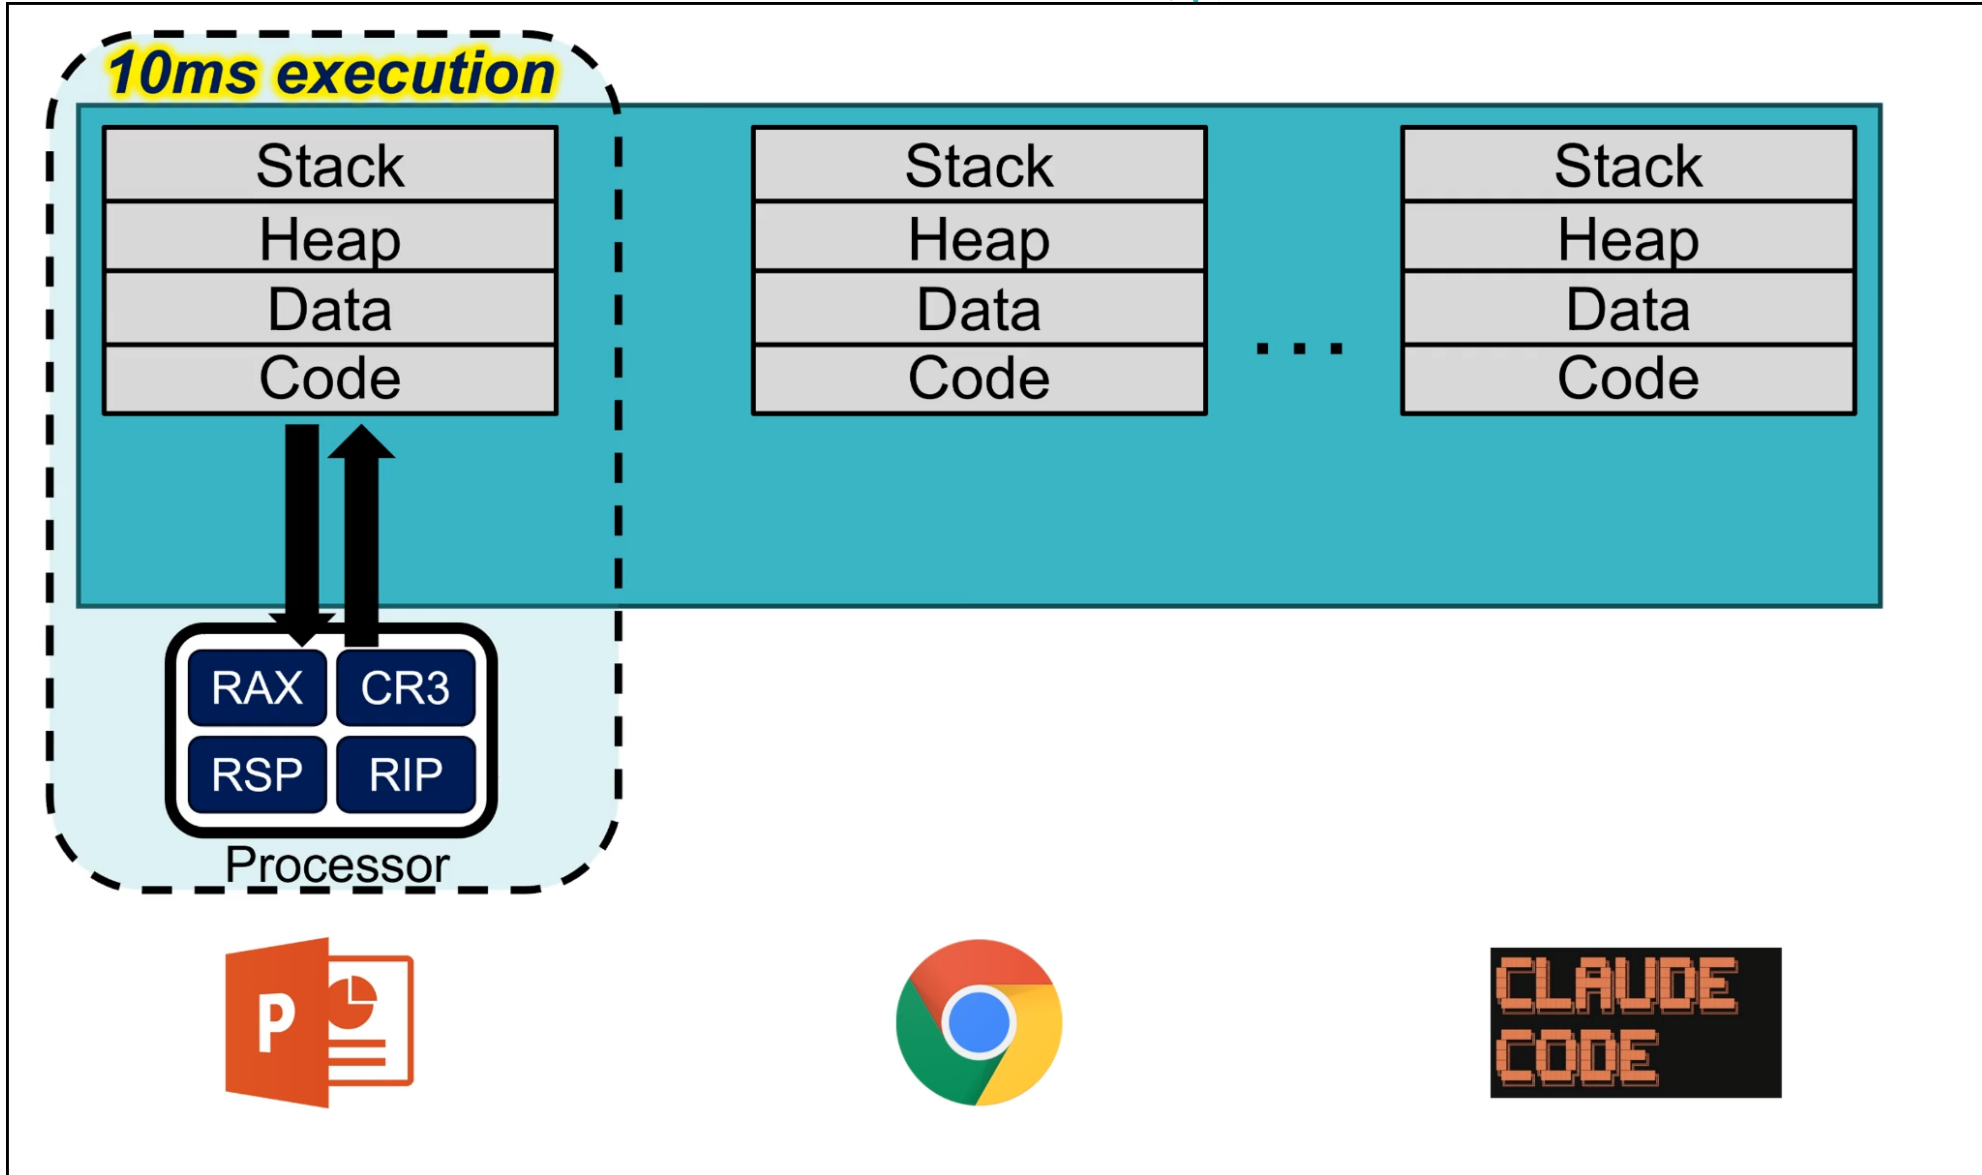


inst₁
inst₂
inst₃
inst₄
...



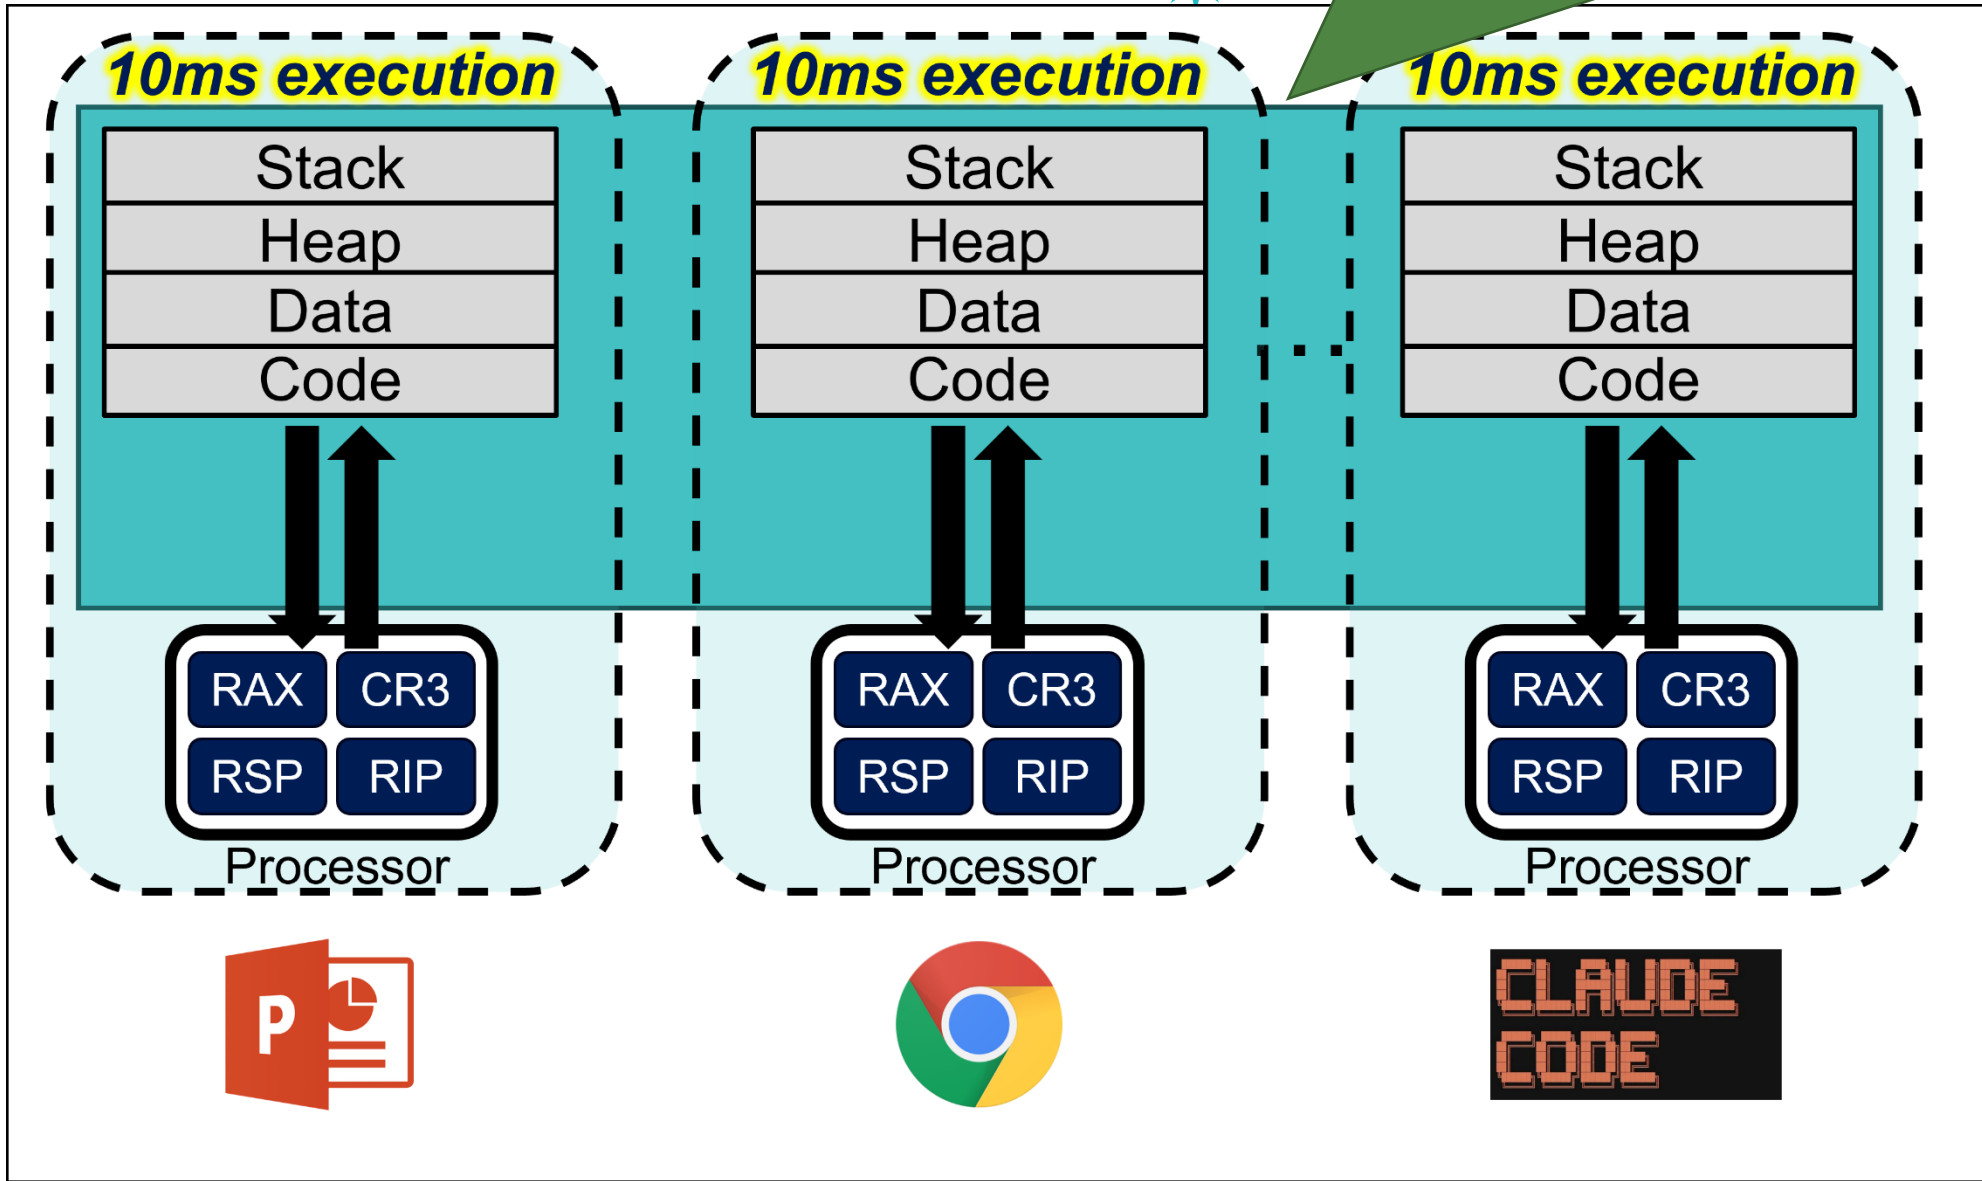
inst₁
inst₂
inst₃
inst₄
...

Illusion Demo Video



Illusion Logical View

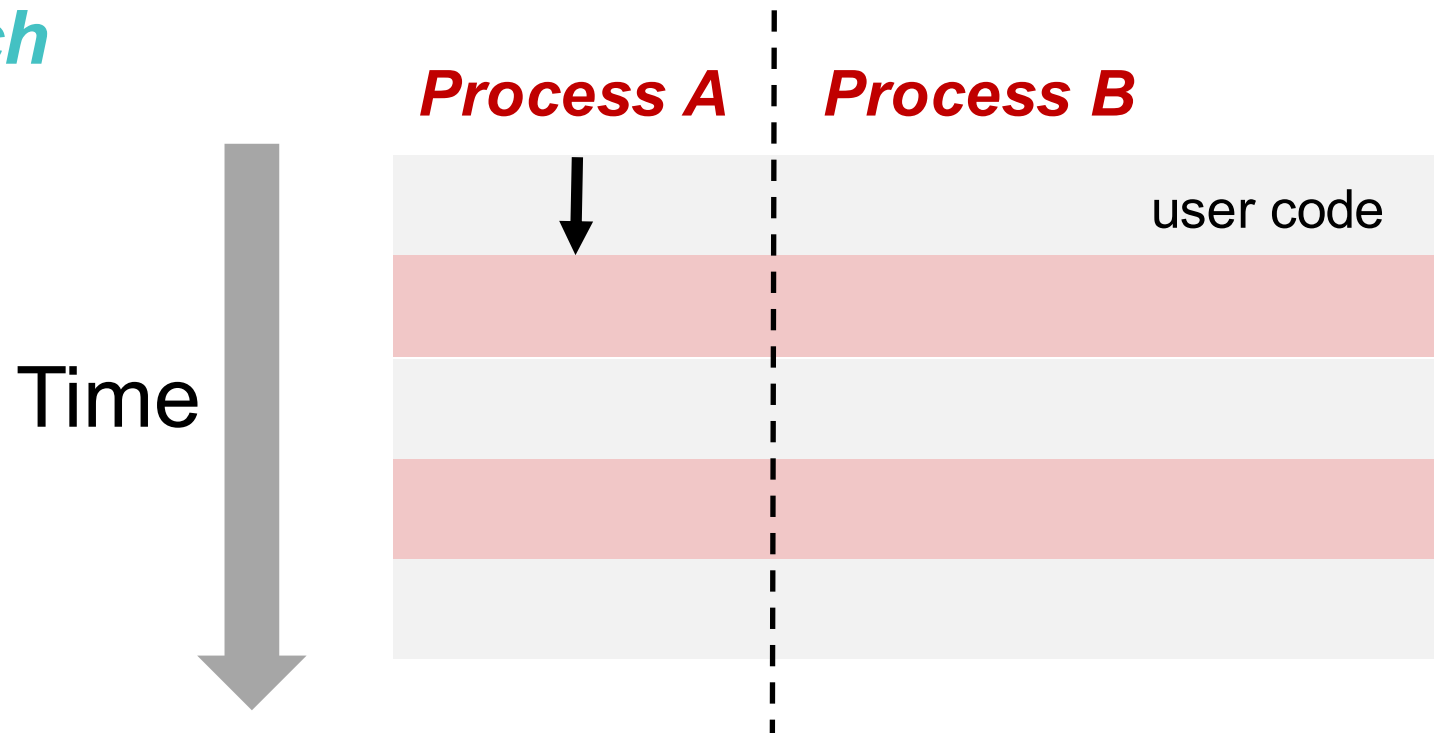
Through context switching, the OS makes a single CPU appear as multiple logical CPUs, one for each process



Context Switching



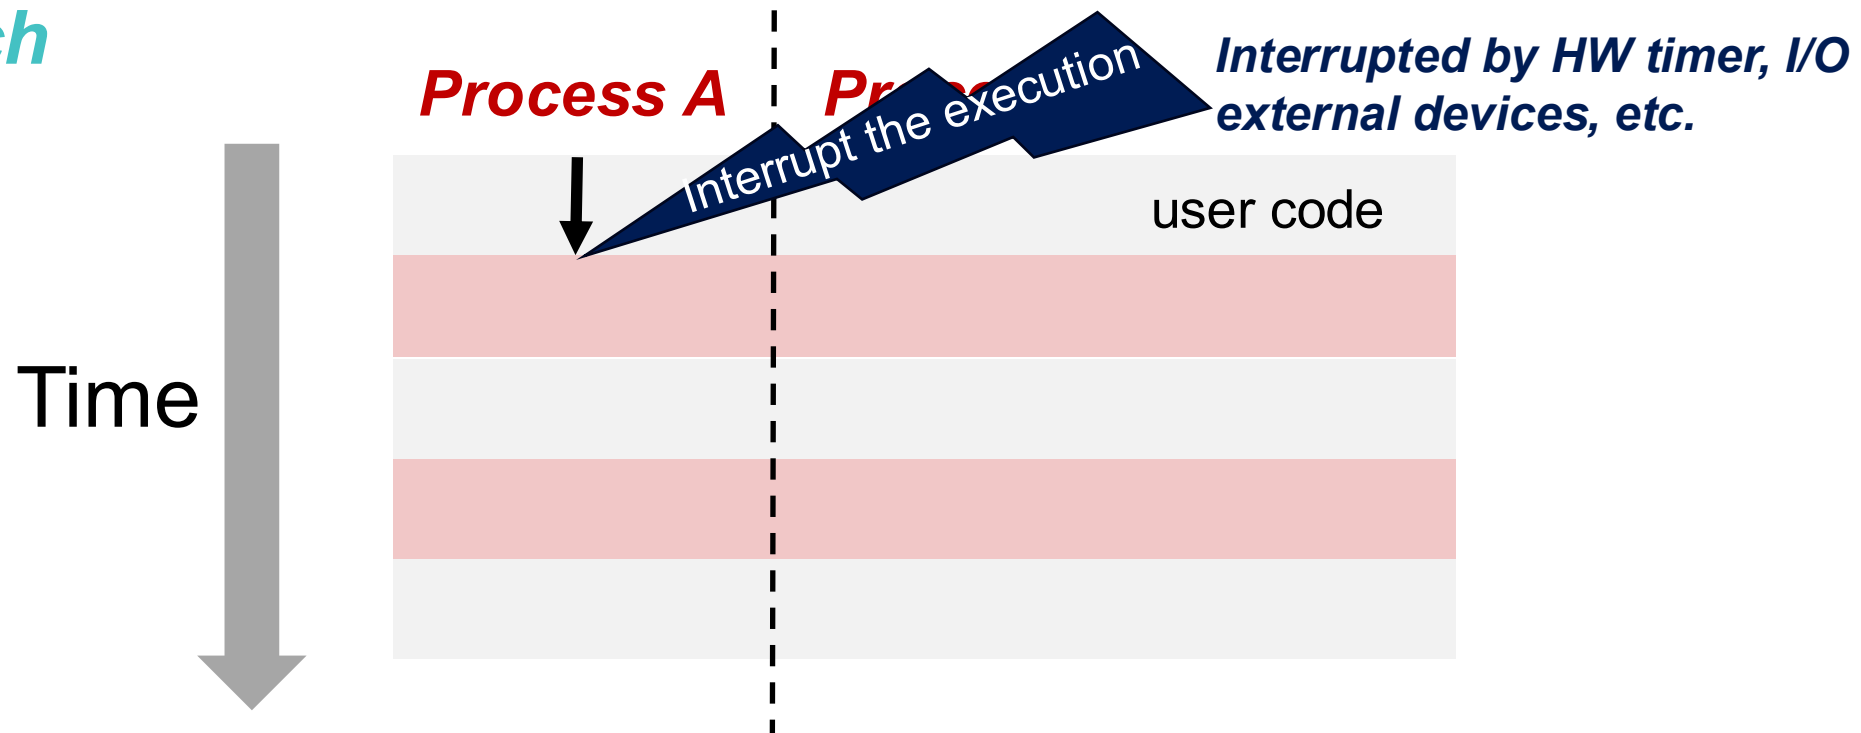
- Processes are managed by a shared chunk of memory-resident OS code called the **kernel**
 - Important: the kernel is not a separate process, but rather runs as part of some existing process
- Control flow passes from one process to another via a **context switch**



Context Switching



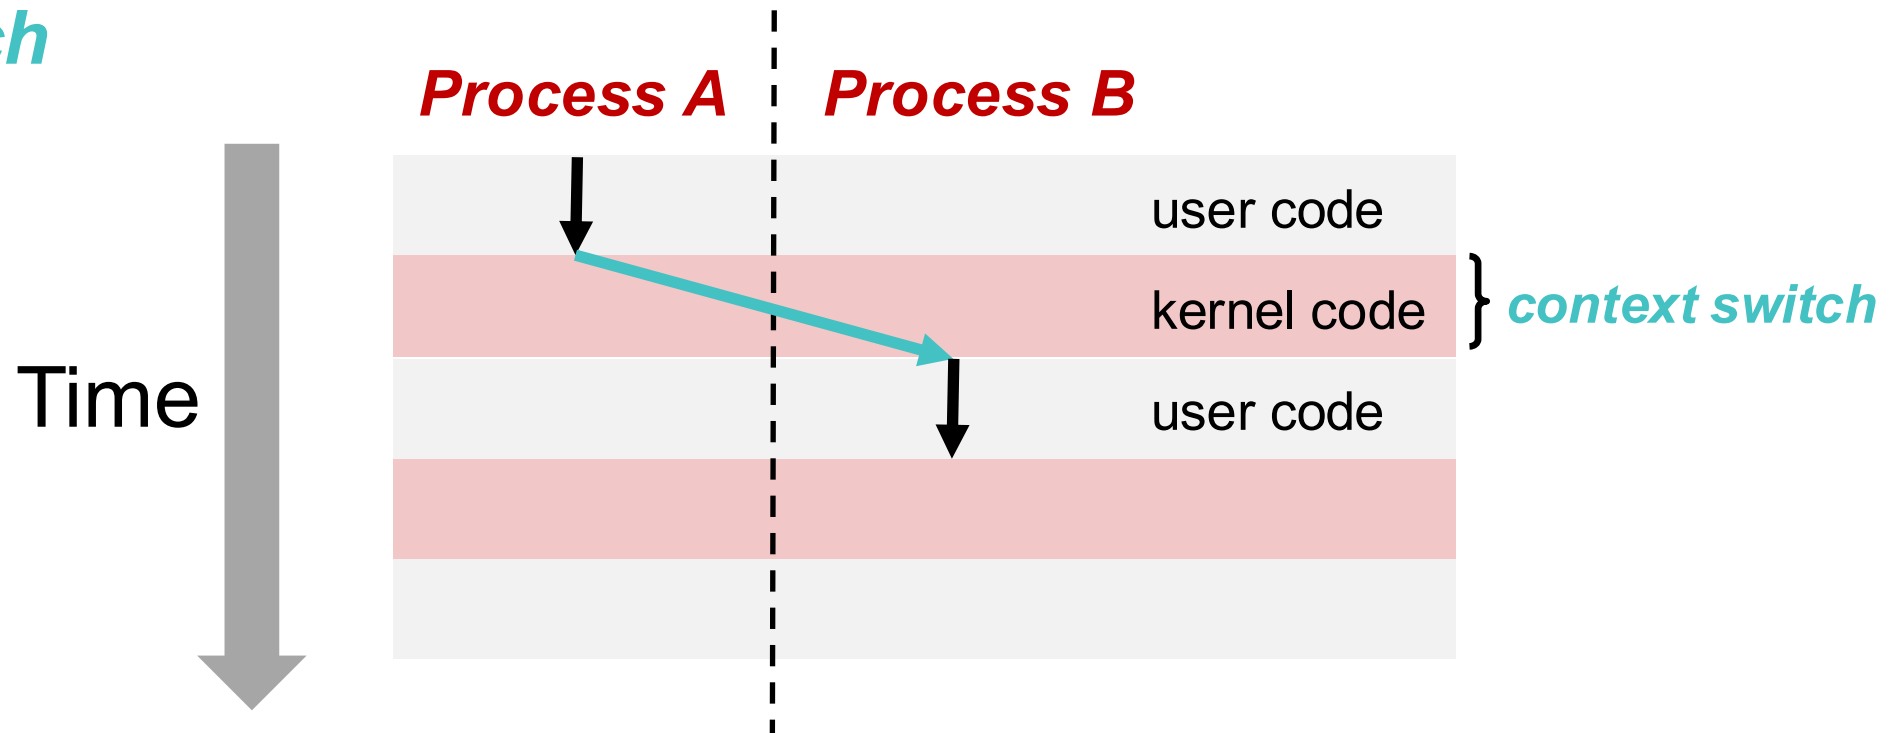
- Processes are managed by a shared chunk of memory-resident OS code called the **kernel**
 - Important: the kernel is not a separate process, but rather runs as part of some existing process
- Control flow passes from one process to another via a **context switch**



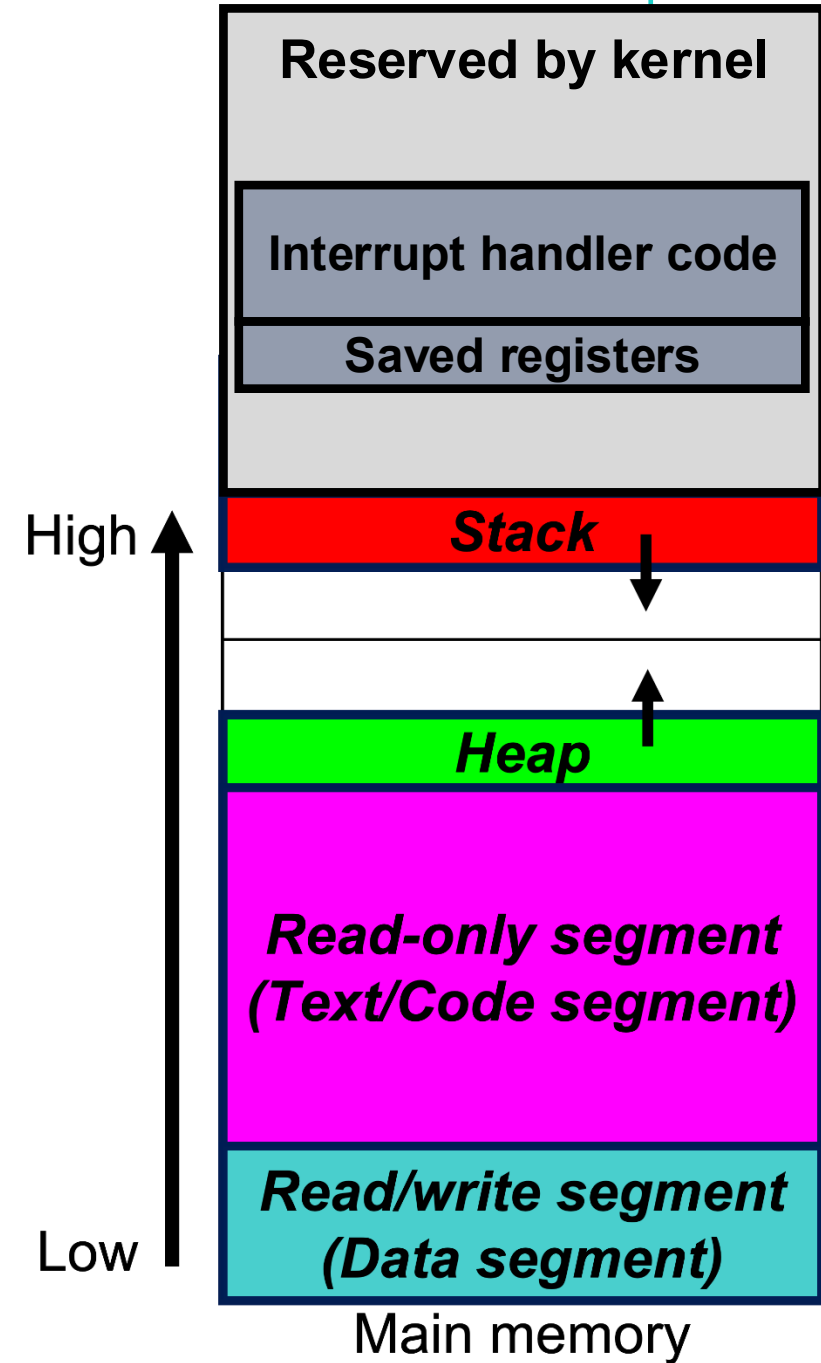
Context Switching



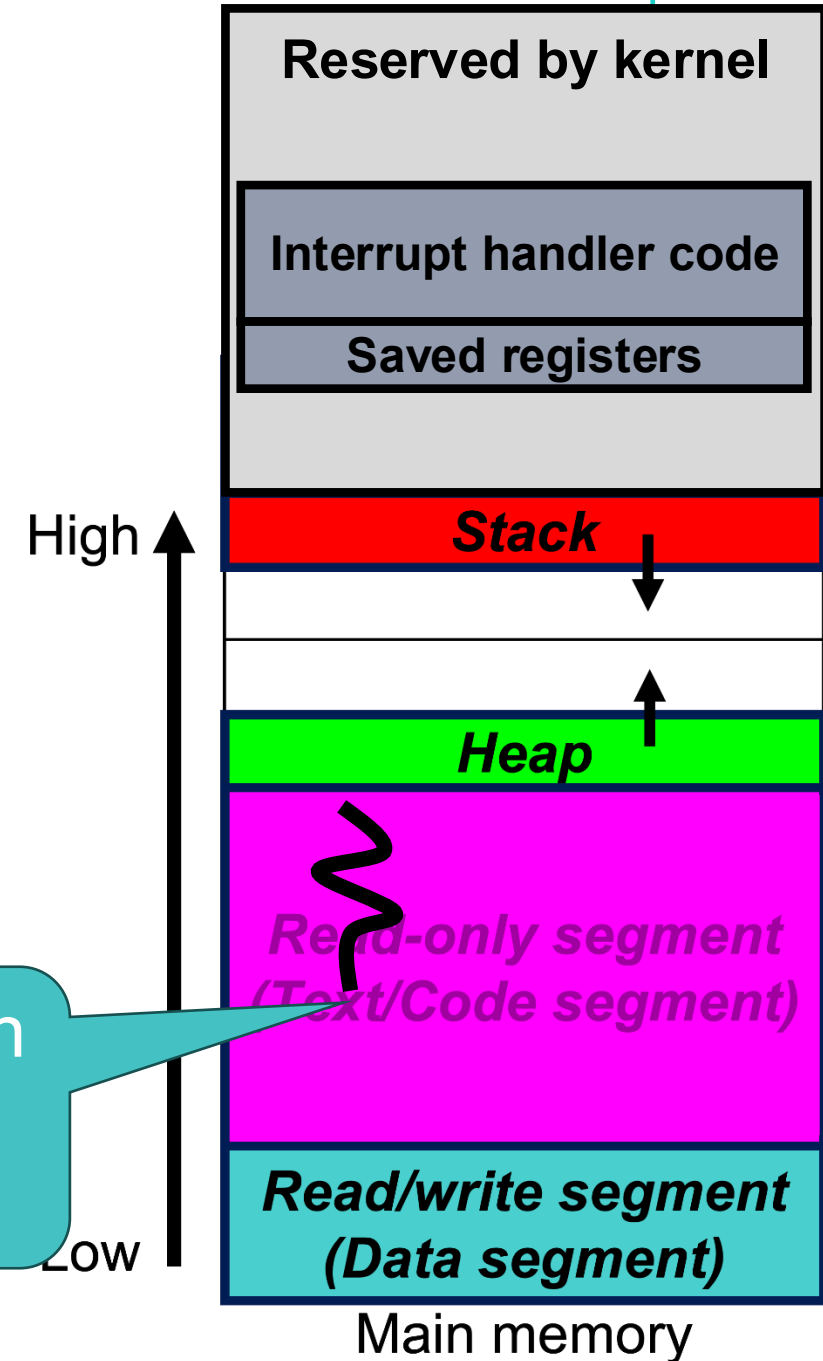
- Processes are managed by a shared chunk of memory-resident OS code called the **kernel**
 - Important: the kernel is not a separate process, but rather runs as part of some existing process
- Control flow passes from one process to another via a **context switch**



Context Switching

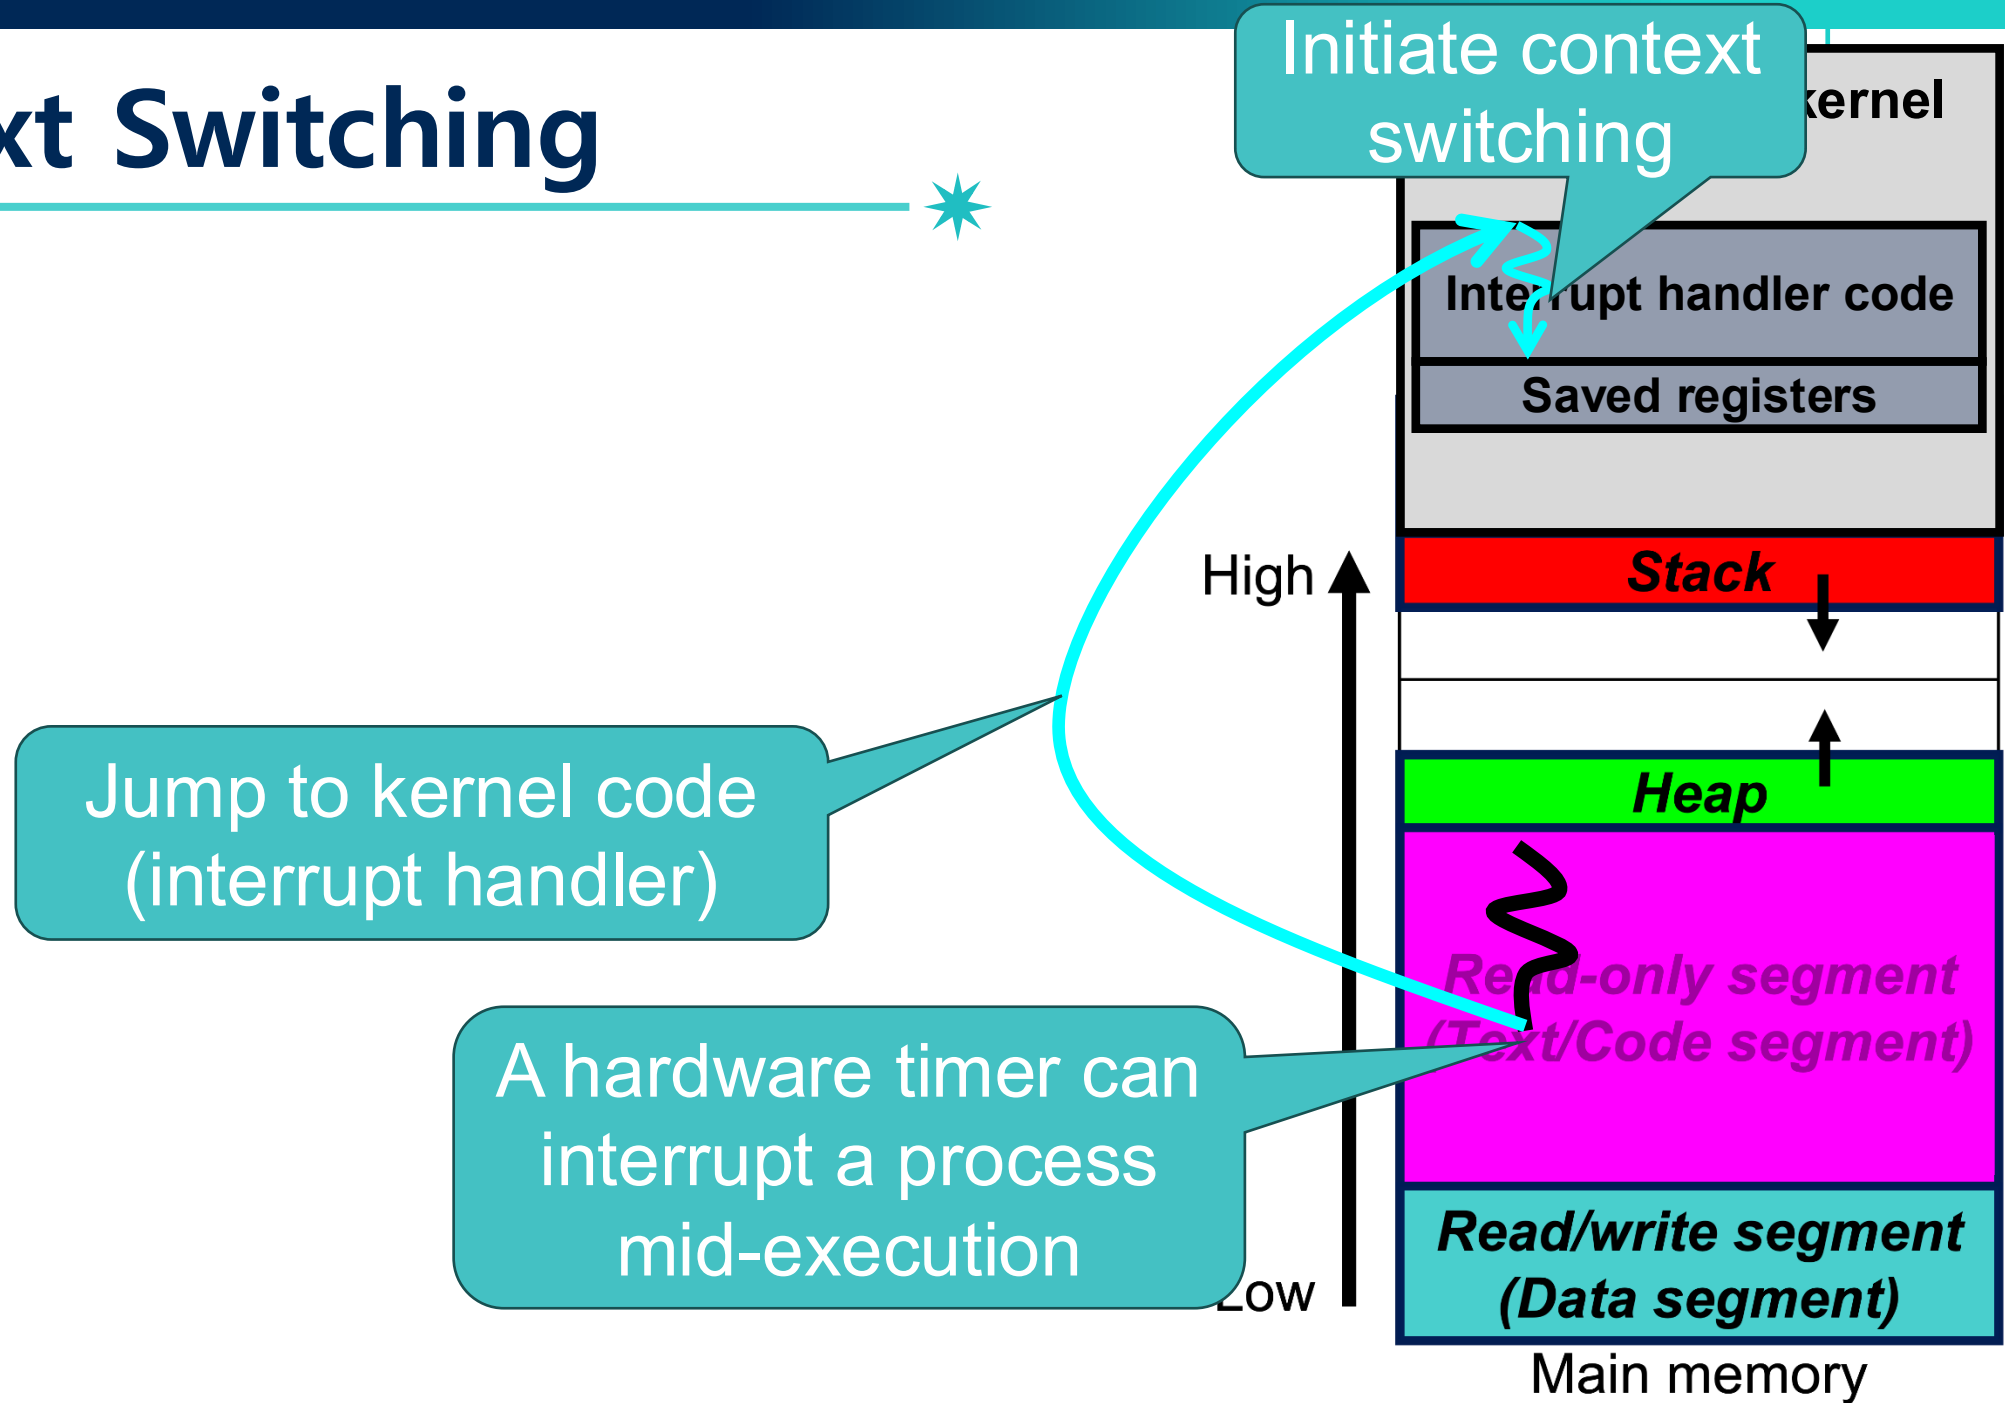


Context Switching



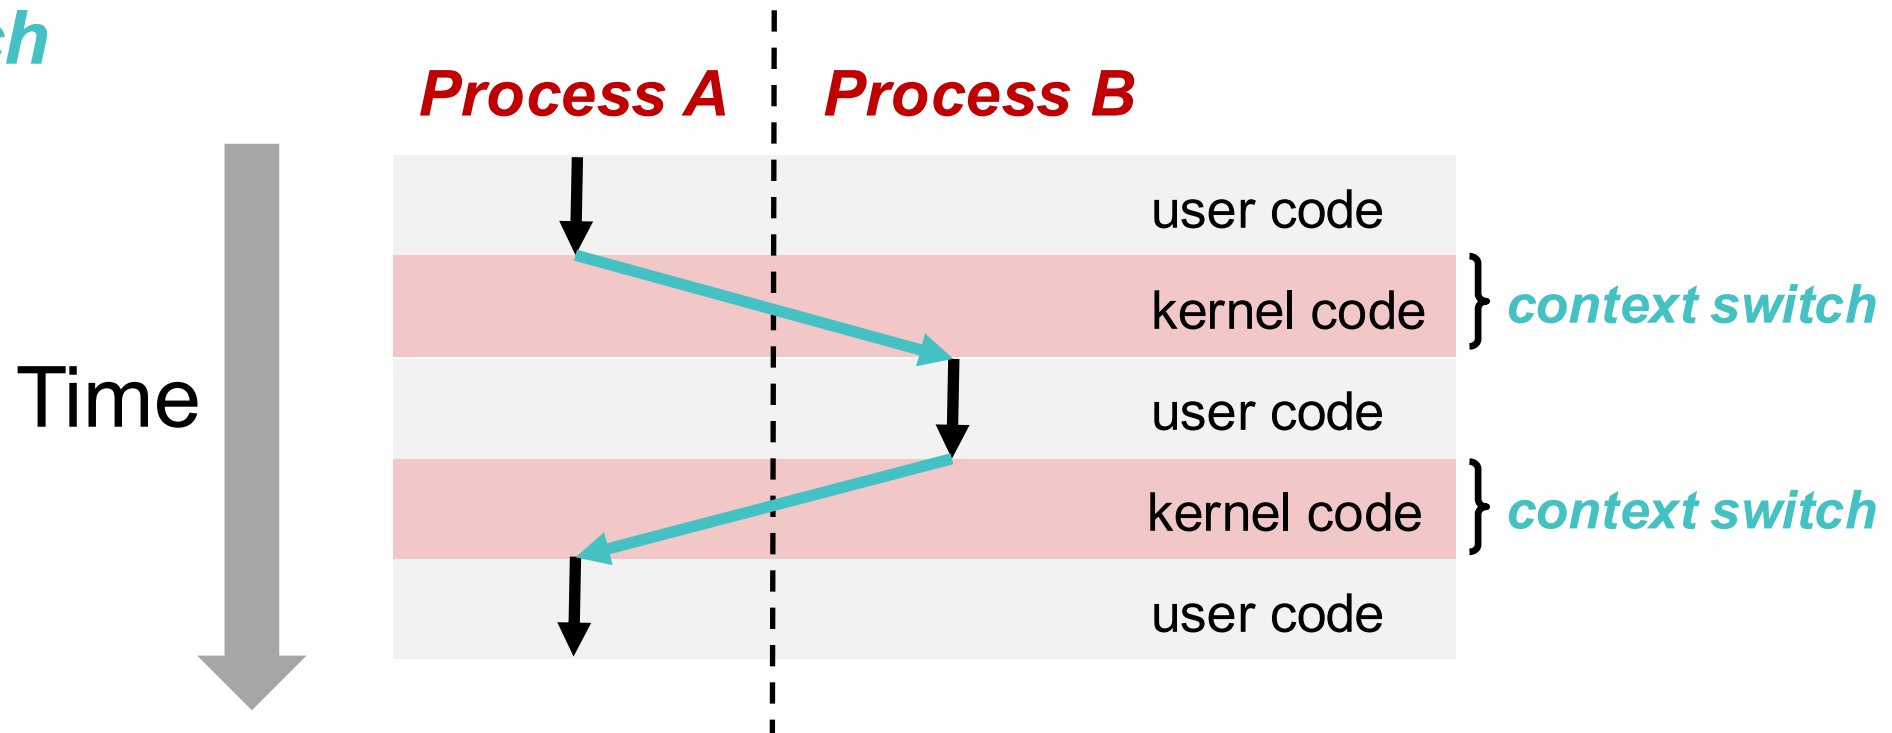
A hardware timer can interrupt a process mid-execution

Context Switching



Context Switching

- Processes are managed by a shared chunk of memory-resident OS code called the **kernel**
 - Important: the kernel is not a separate process, but rather runs as part of some existing process
- Control flow passes from one process to another via a **context switch**



Multitasking



- A **single processor** executes multiple processes **concurrently**
 - Process executions interleaved through context switching
 - Address spaces managed by virtual memory system
 - Register values for non-executing processes saved in memory

Summary: Process




An instance of a program that is being executed (or is ready for execution)


- Not the same as “program” or “processor”

- Provides each program with two key abstractions:
 - ***Private address space***
 - Each program seems to have exclusive use of main memory
 - Provided by kernel mechanism called **virtual memory**
 - ***Logical control flow***
 - Each program seems to have exclusive use of the CPU
 - Provided by kernel mechanism called **context switching**

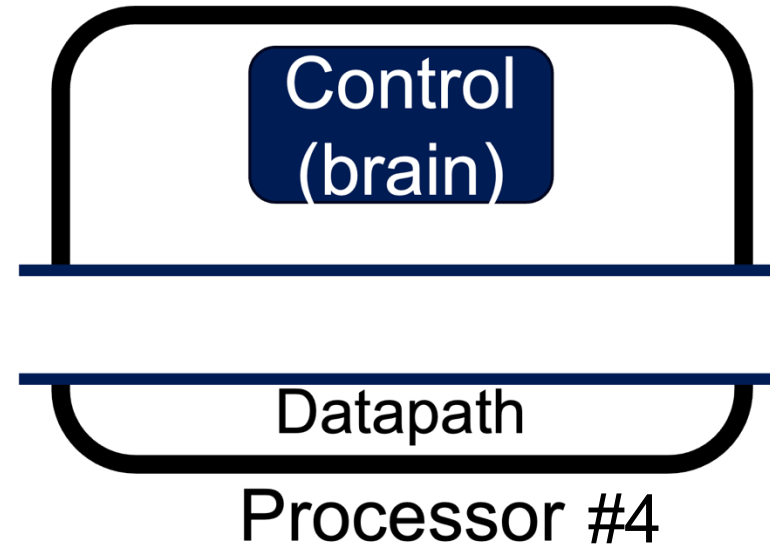
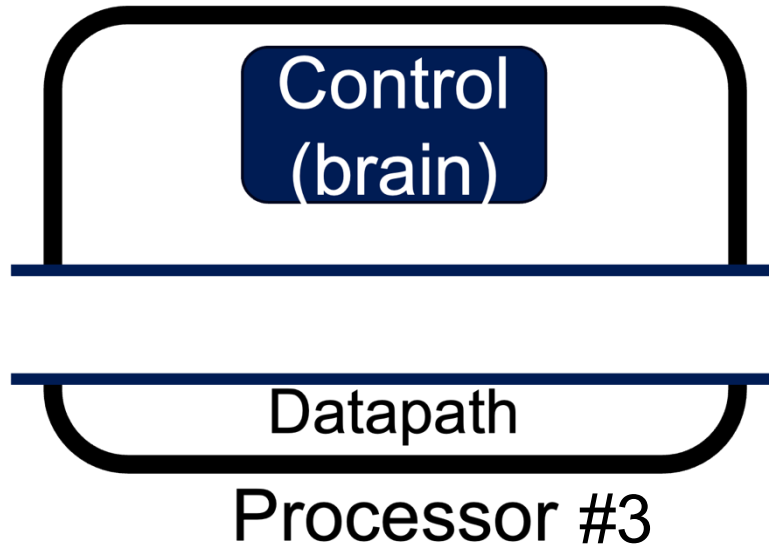
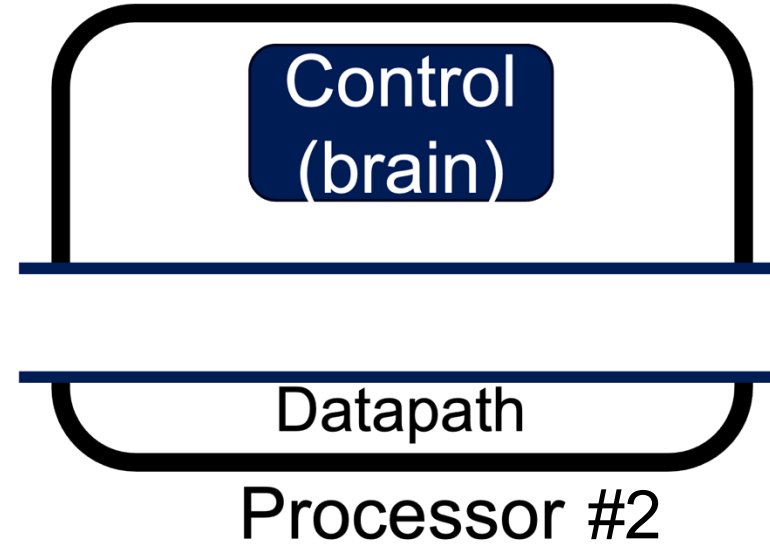
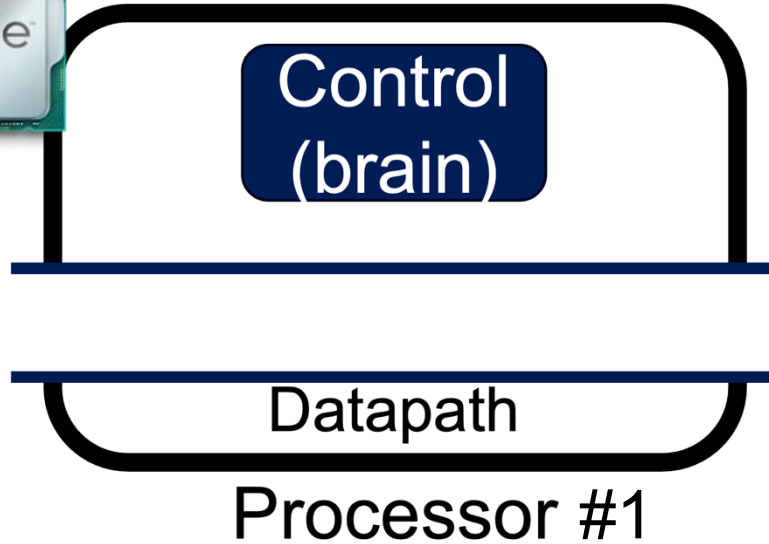
Nowadays: Uniprocessor to Multiprocessors



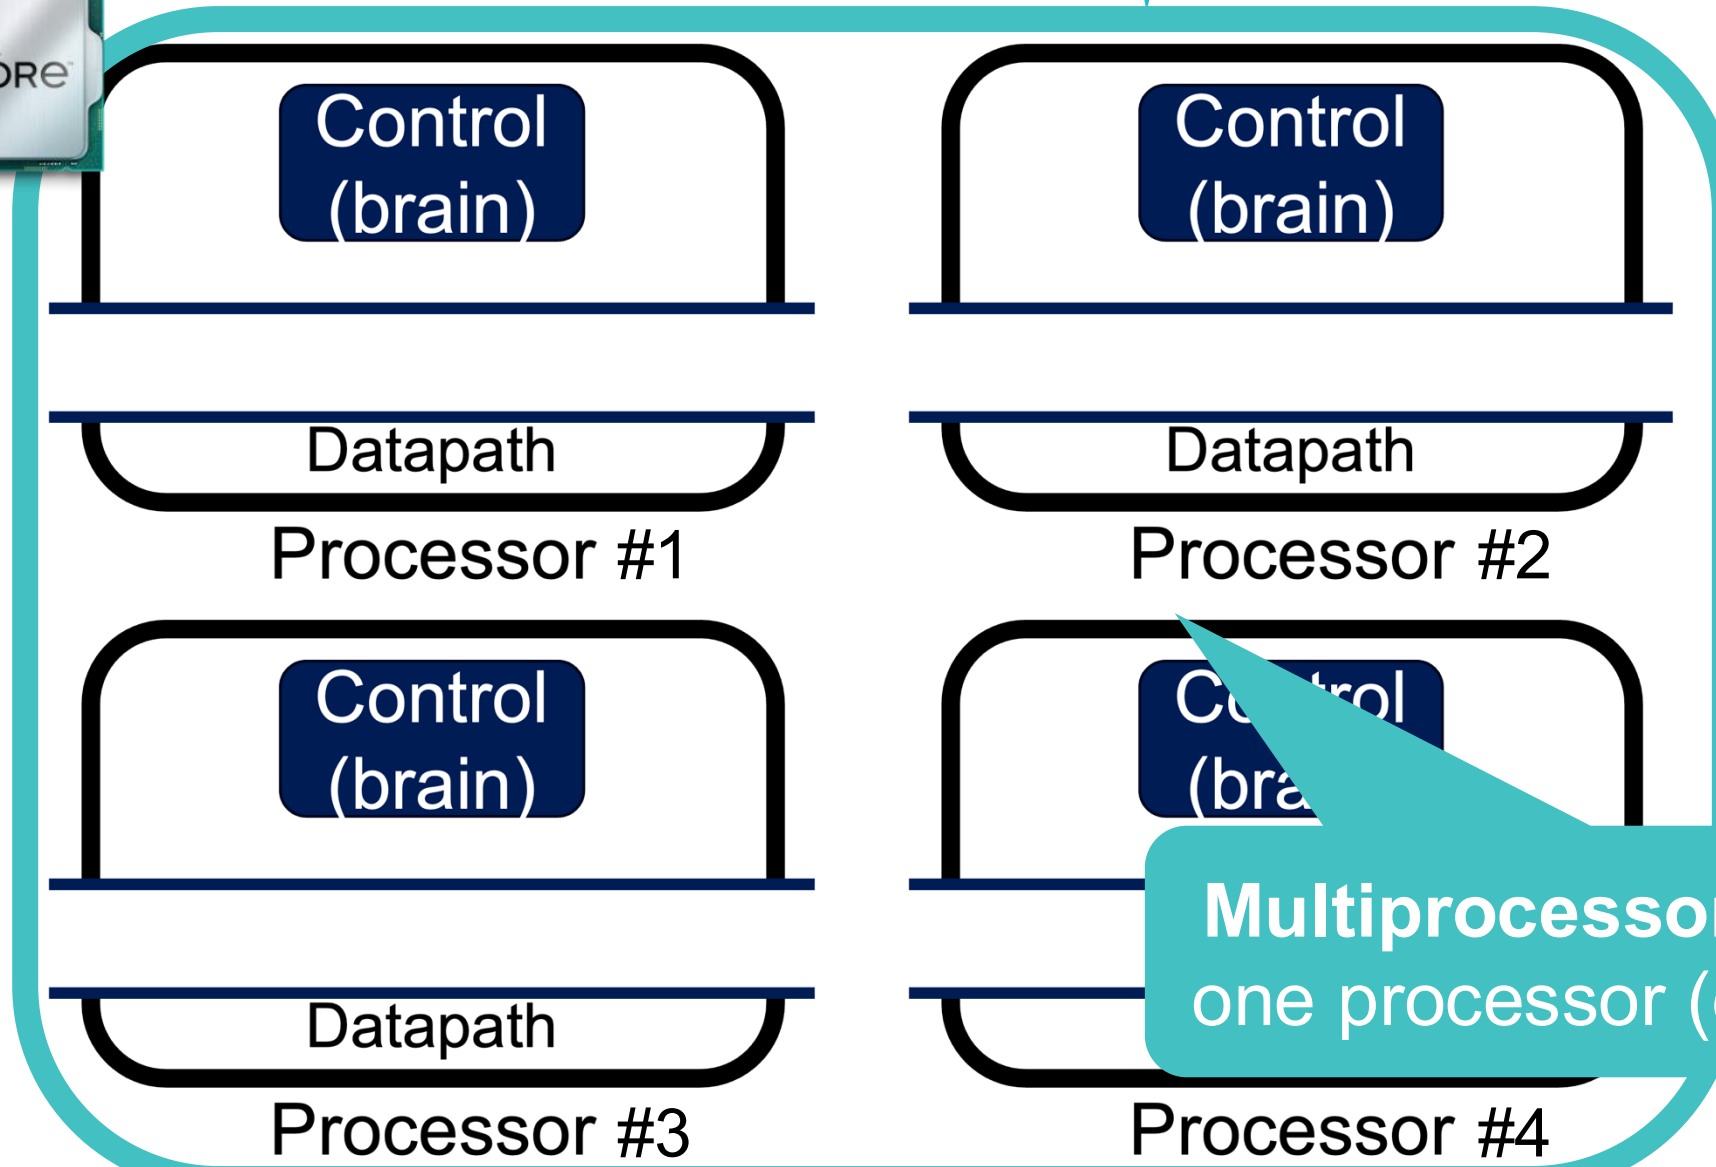
Nowadays: Uniprocessor to Multiprocessors



Nowadays: Uniprocessor to Multiprocessors

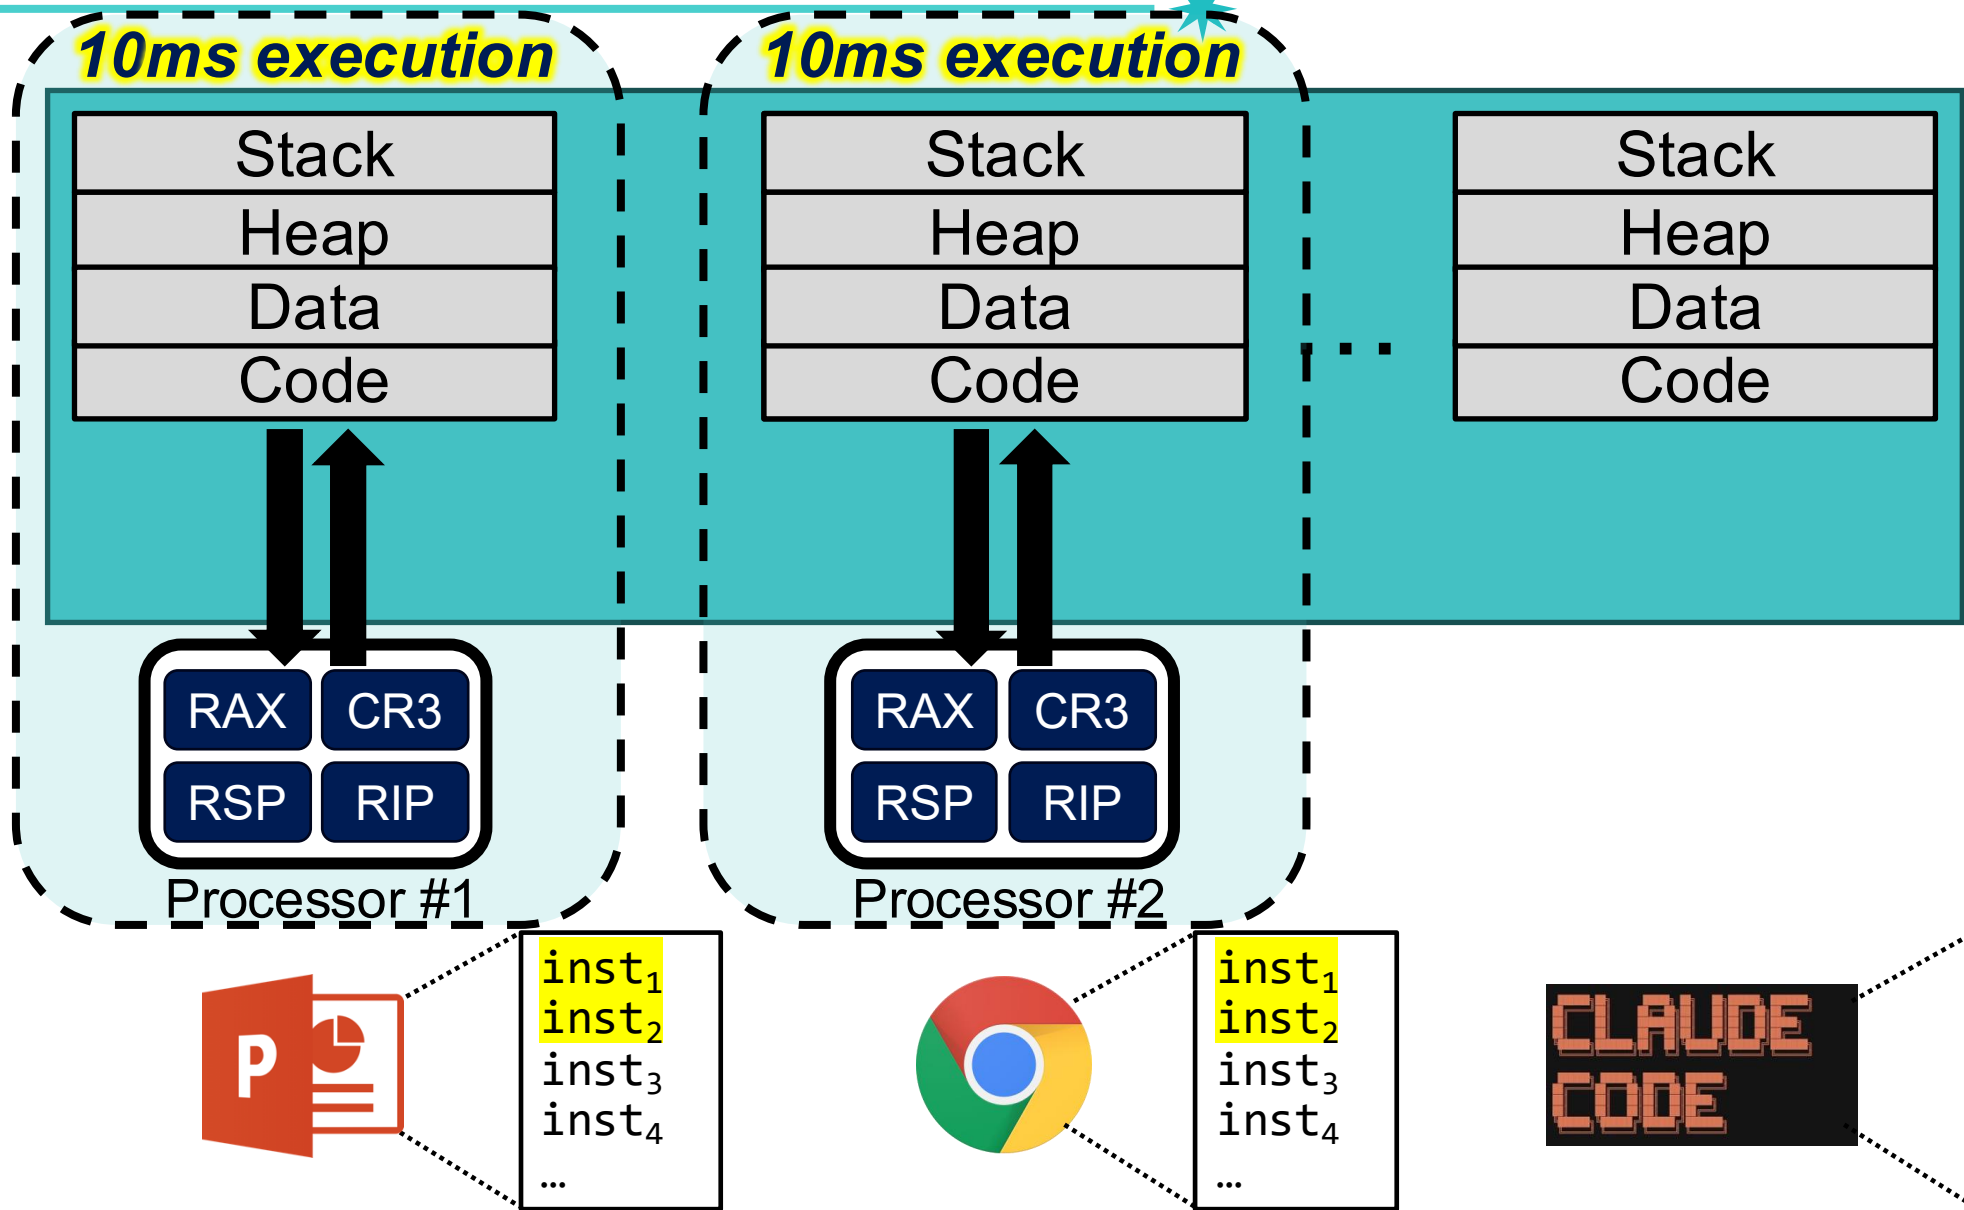


Nowadays: Uniprocessor to Multiprocessors

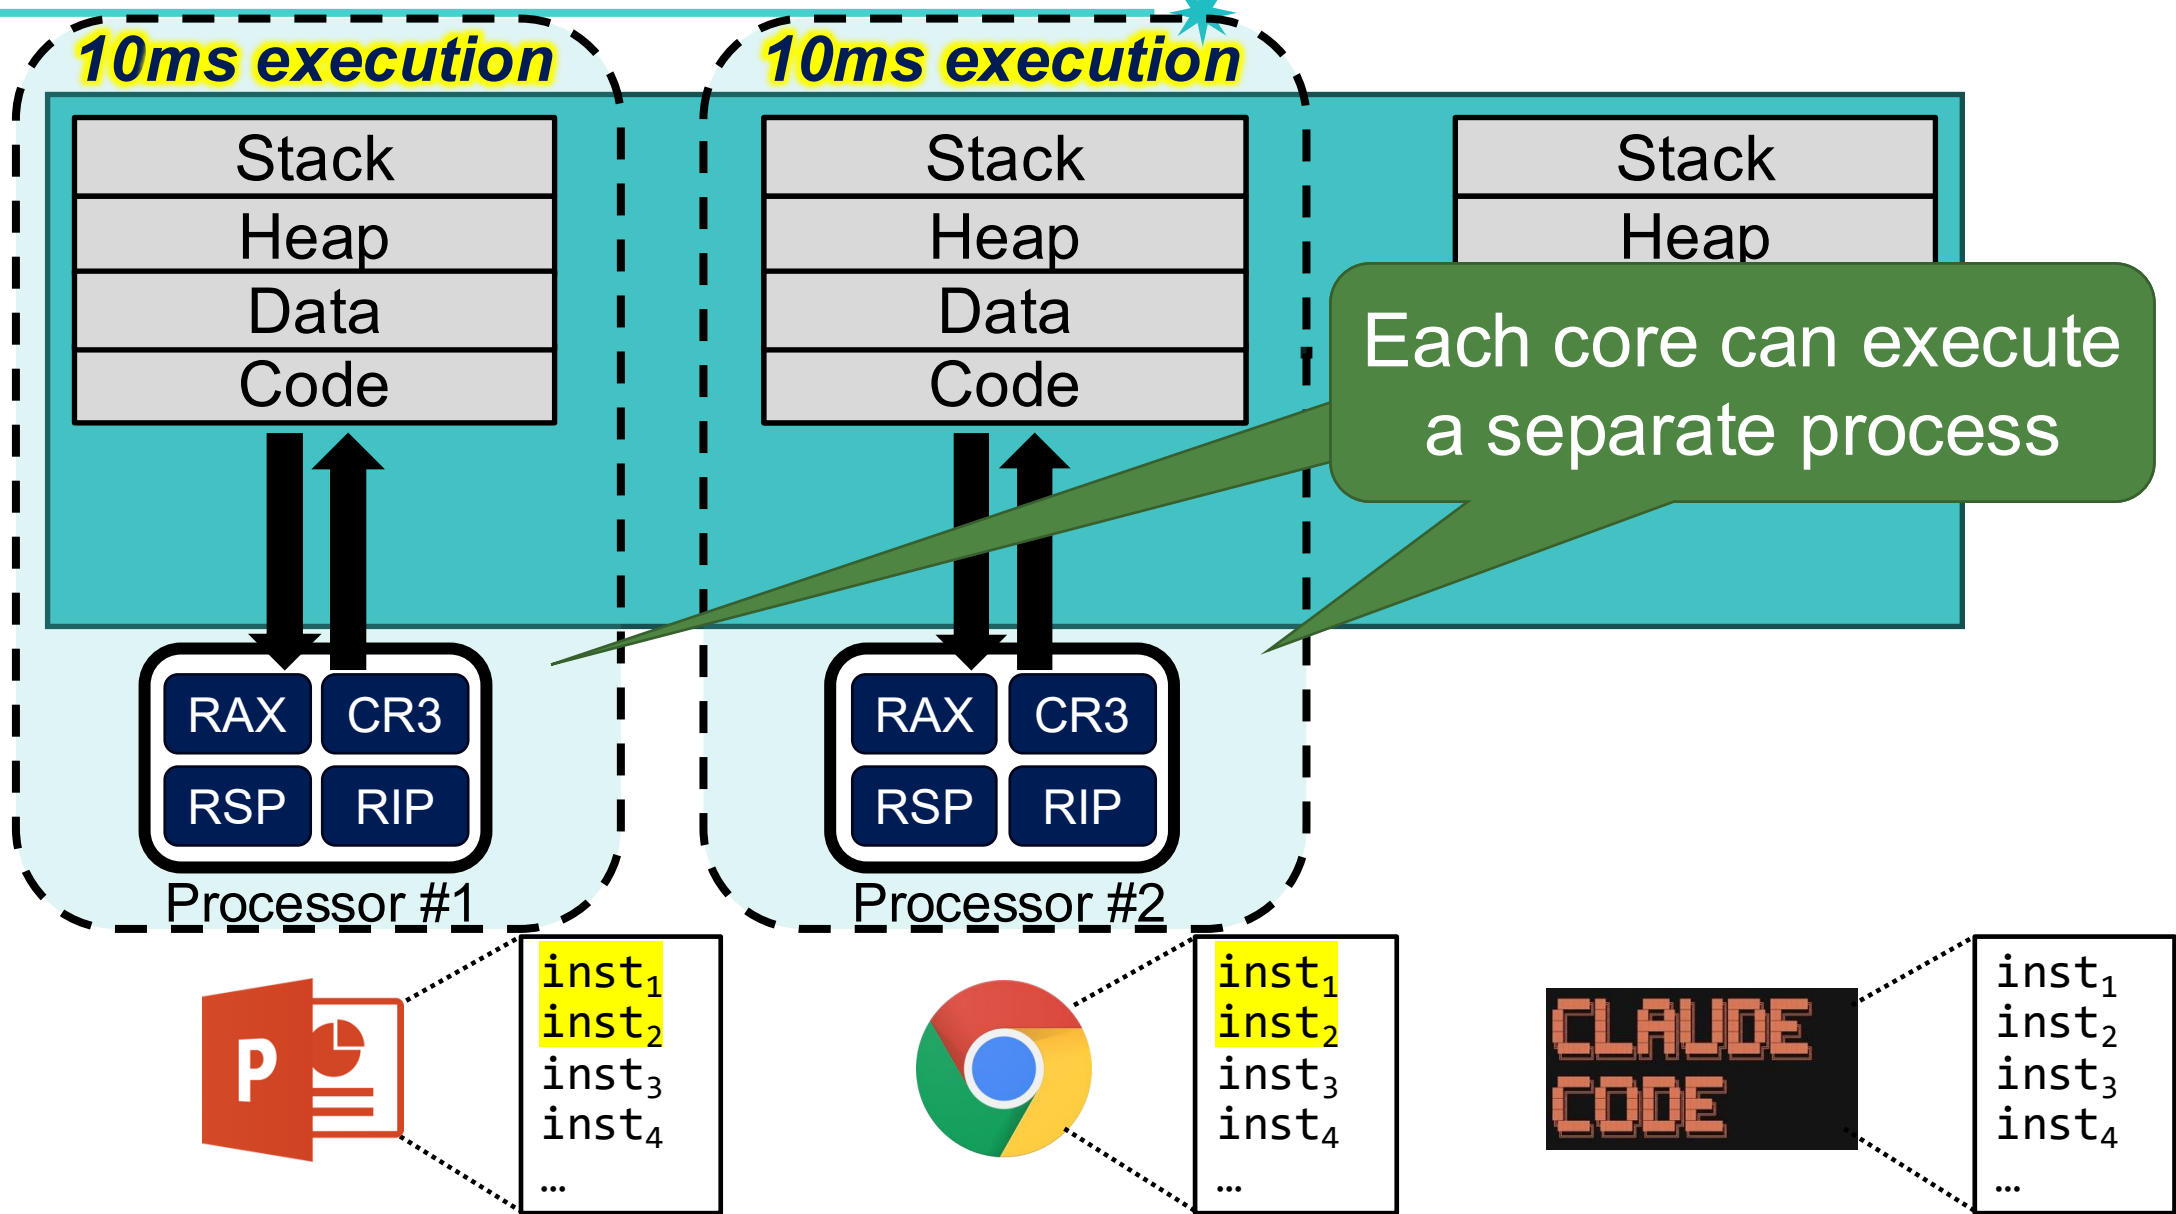


Multiprocessors: more than one processor (core) per chip

Multiprocessors (Multicore Processors)



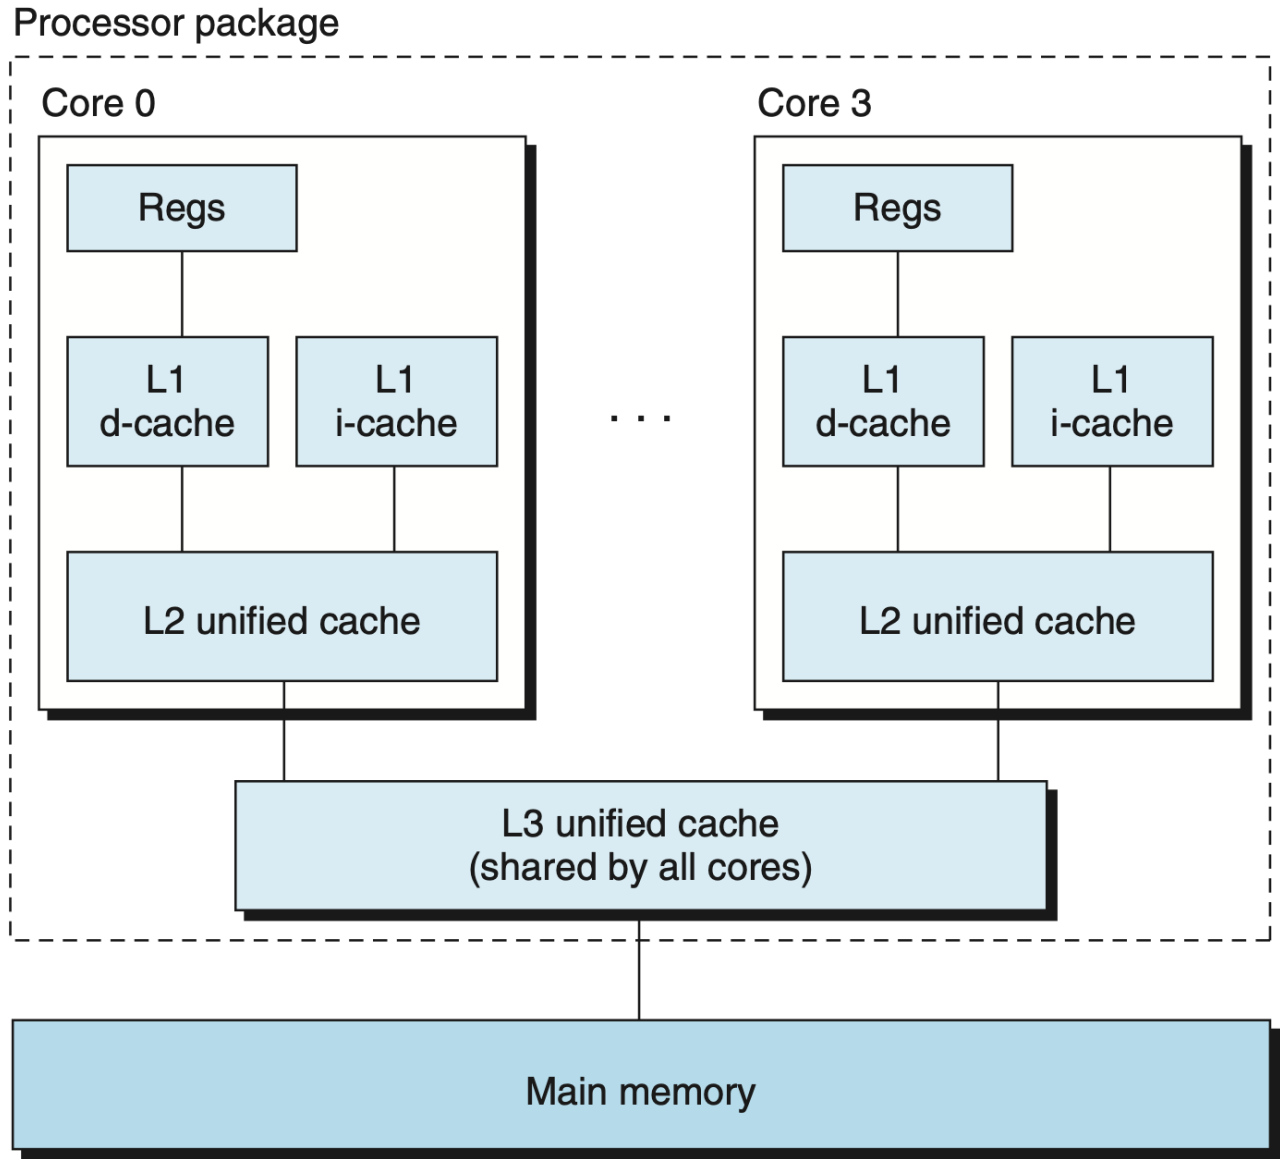
Multiprocessors (Multicore Processors)



Multiprocessors (Multicore Processors)

- Multiple CPUs on single chip
- Share main memory (and some caches)
- Each can execute a separate process
 - Scheduling of processors onto cores done by kernel
- This is called “Parallelism”

Recall: Intel Core i7 Cache Hierarchy



Multiprocessing Example

- Running “top” on your environment!

```
top - 10:57:12 up 42 days, 17:53, 1 user, load average: 2.01, 2.02, 2.00
Tasks: 702 total, 3 running, 697 sleeping, 2 stopped, 0 zombie
%Cpu(s): 0.4 us, 5.8 sy, 0.0 ni, 93.7 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
MiB Mem : 515851.0 total, 1663.2 free, 4095.0 used, 510092.8 buff/cache
MiB Swap: 30518.0 total, 30512.0 free, 6.0 used. 508397.0 avail Mem
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
4056347	compu	20	0	44756	27036	3560	R	100.0	0.0	11764:27	lackey-amd64-li
4061301	compu	20	0	44756	27036	3560	R	99.7	0.0	11759:11	lackey-amd64-li
1550125	su3604	20	0	13812	4676	3496	R	0.7	0.0	0:00.07	top
1365471	su3604	20	0	14120	5856	4364	S	0.3	0.0	0:00.65	sshd
1542286	root	20	0	0	0	0	I	0.3	0.0	0:03.33	kworker/u64:0-edac-poller
1550032	root	20	0	0	0	0	I	0.3	0.0	0:00.01	kworker/24:1-events
1	root	20	0	173920	18356	8528	S	0.0	0.0	1:01.53	systemd
2	root	20	0	0	0	0	S	0.0	0.0	0:06.84	kthreadd
3	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	rcu_gp
4	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	rcu_par_gp
6	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	kworker/0:0H-events_highpri
10	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	mm_percpu_wq
11	root	20	0	0	0	0	S	0.0	0.0	0:00.00	rcu_tasks_rude_

702 "tasks", 3 of which are active

Pressing Example

Each process is identified by Process ID (PID), user account, and command line

```
top - 10:57:12 up 42 days, 17:53, 1 user, load average: 2.01, 2.02, 2.00
```

```
Tasks: 702 total, 3 running, 697 sleeping, 2 stopped, 0 zombie
```

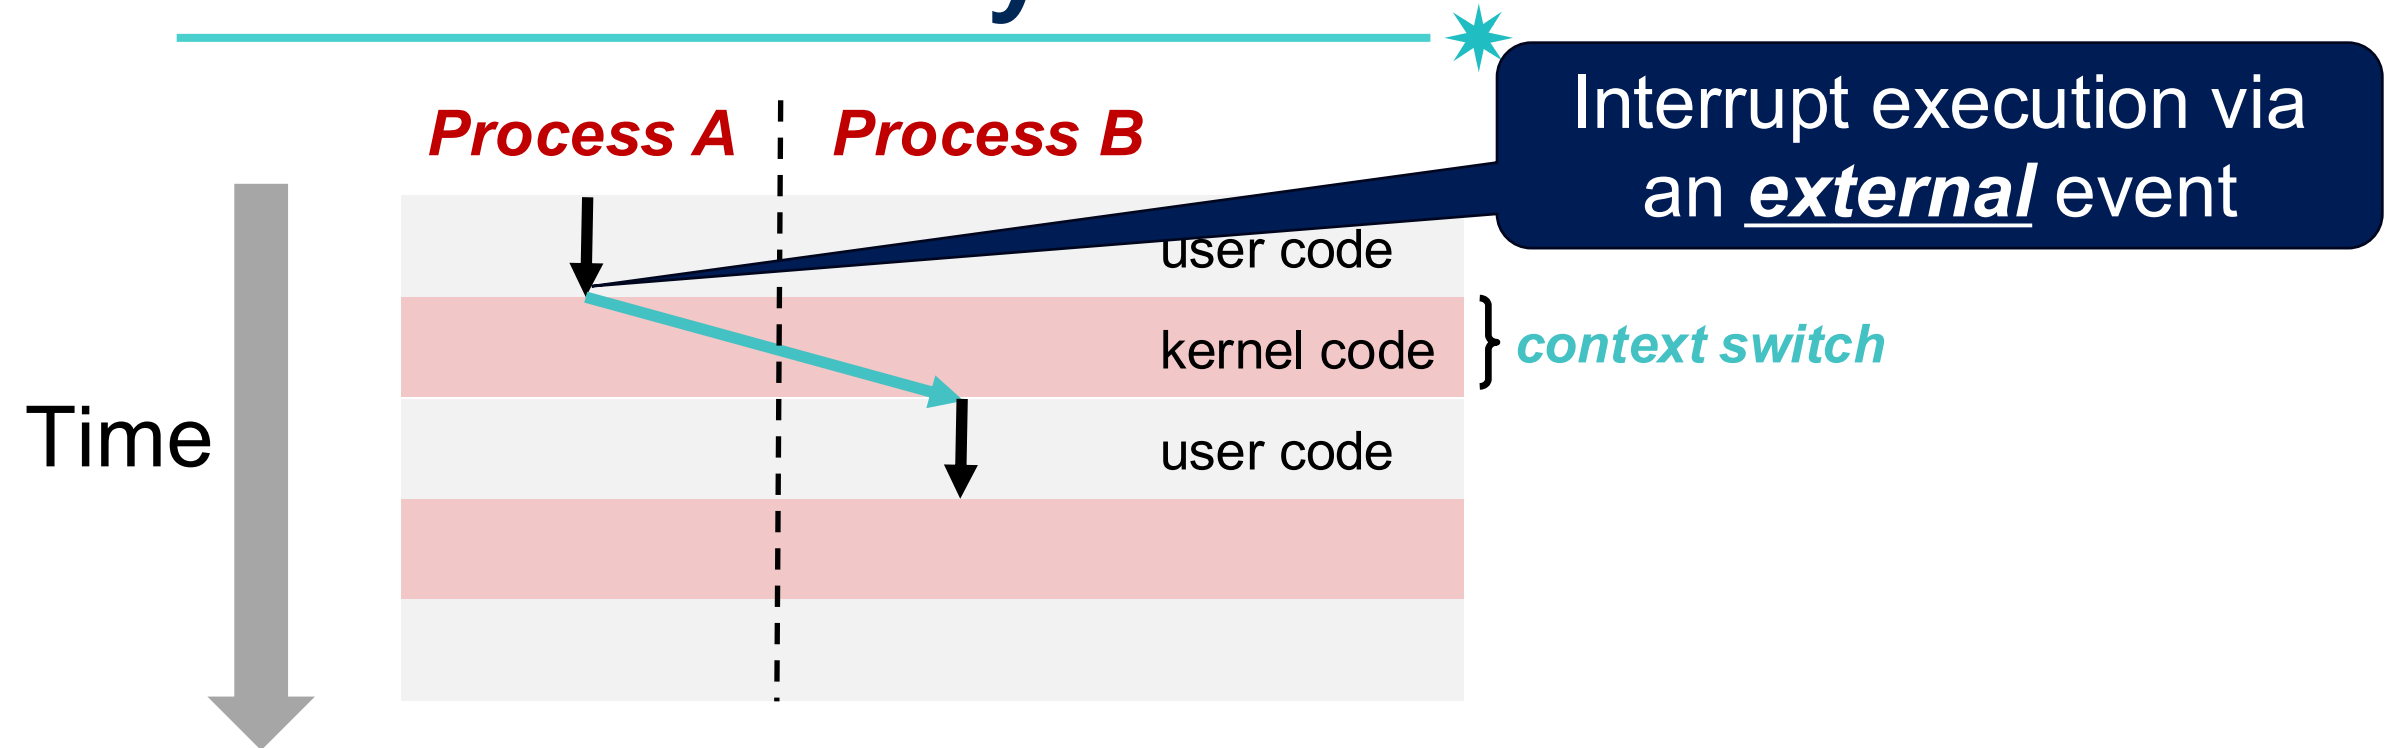
```
%Cpu(s): 0.4 us, 5.8 sy, 0.0 ni, 93.7 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
```

```
MiB Mem : 515851.0 total, 1663.2 free, 4095.0 used, 510092.8 buff/cache
```

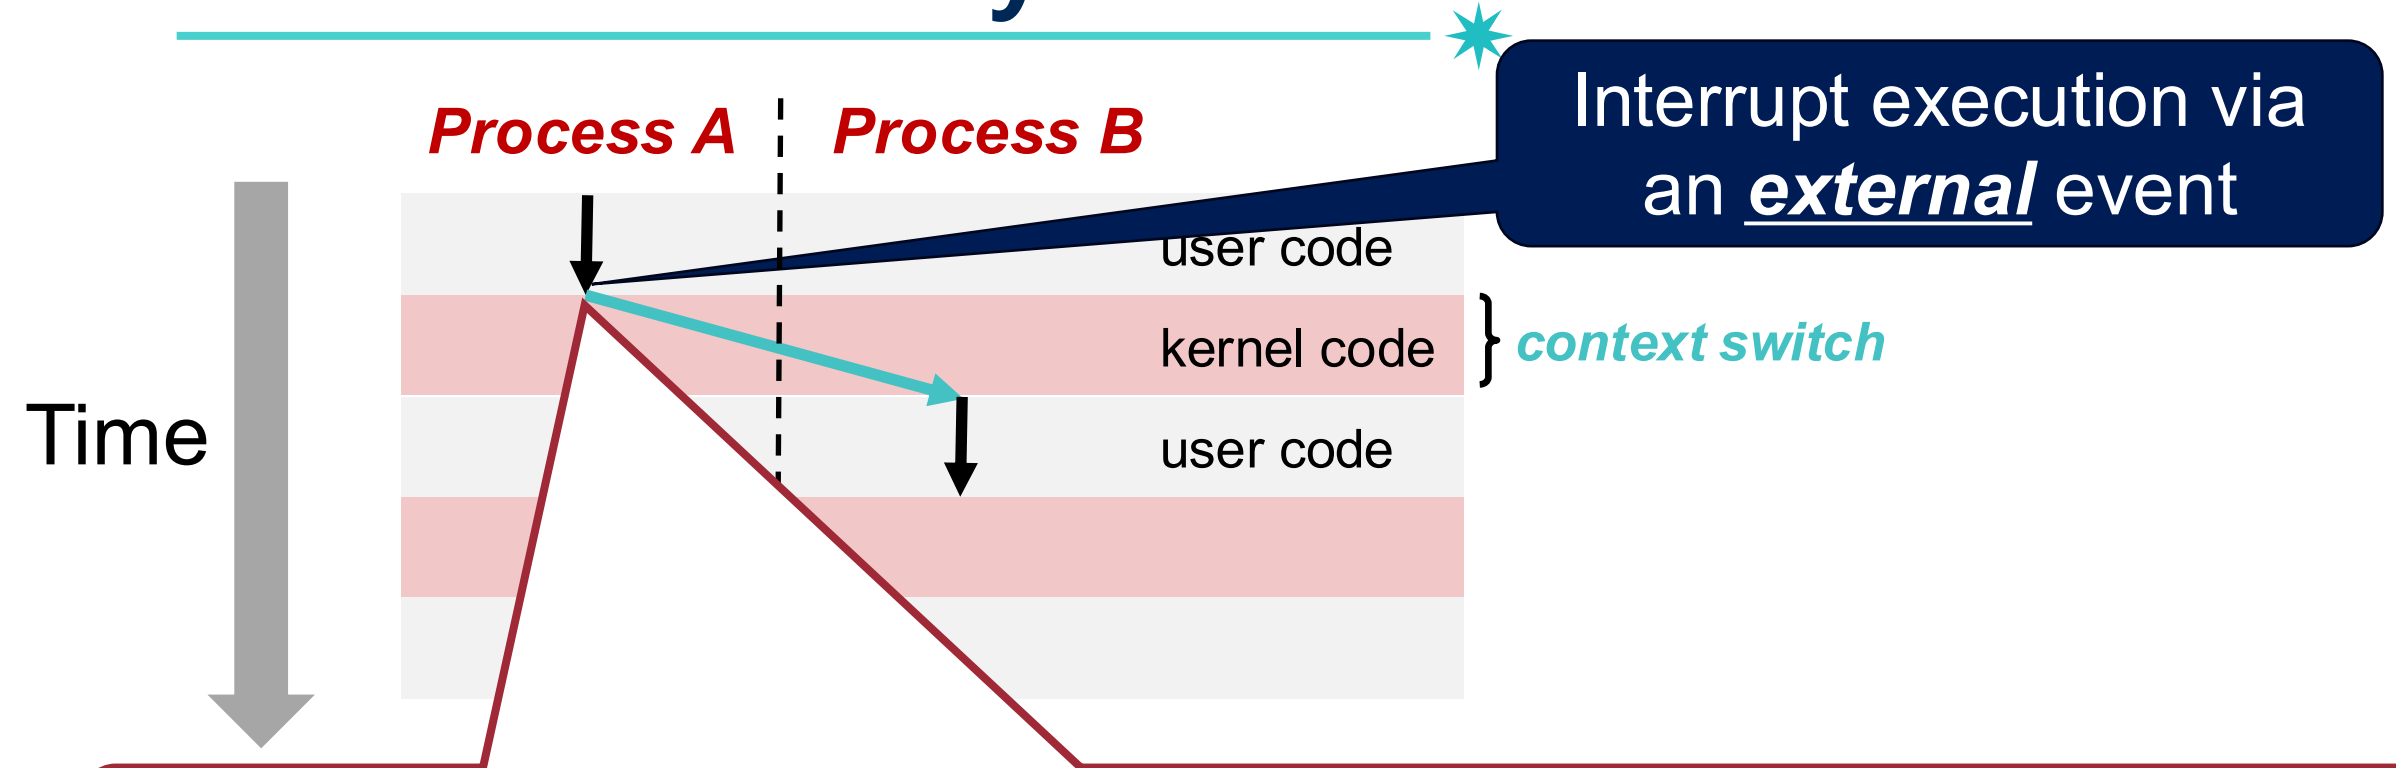
```
MiB Swap: 30518.0 total, 30512.0 free, 6.0 used, 508397.0 avail Mem
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
4056347	compu	20	0	44756	27036	3560	R	100.0	0.0	11764:27	lackey-amd64-li
4061301	compu	20	0	44756	27036	3560	R	99.7	0.0	11759:11	lackey-amd64-li
1550125	su3604	20	0	13812	4676	3496	R	0.7	0.0	0:00.07	top
1365471	su3604	20	0	14120	5856	4364	S	0.3	0.0	0:00.65	sshd
1542286	root	20	0	0	0	0	I	0.3	0.0	0:03.33	kworker/u64:0-edac-poller
1550032	root	20	0	0	0	0	I	0.3	0.0	0:00.01	kworker/24:1-events
1	root	20	0	173920	18356	8528	S	0.0	0.0	1:01.53	systemd
2	root	20	0	0	0	0	S	0.0	0.0	0:06.84	kthreadd
3	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	rcu_gp
4	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	rcu_par_gp
6	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	kworker/0:0H-events_highpri
10	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	mm_percpu_wq
11	root	20	0	0	0	0	S	0.0	0.0	0:00.00	rcu_tasks_rude_

Motivation of System Calls




Motivation of System Calls



Sometimes, the program itself needs help from the kernel!
 (e.g., print something to the screen, read/write a file, create a new process, ...)



Q. How can a user program intentionally request a privileged operation?



System Calls!

System Calls

System Calls



- A request to the kernel to perform a **privileged operation**

System Calls



- A request to the kernel to perform a **privileged operation**
- Whenever a program wants to cause an **effect outside its own process**, it must ask the kernel for help
- It can be called directly from your program or indirectly through a library (e.g., `printf()`)

For example,

- Read/write files
- Get current time
- Allocate RAM (`sbrk`)
- Create new processes

Invoke a System Call



- We can invoke a system call by executing a special instruction
 - `int 0x80` on x86 Linux
 - `syscall` on x86-64 and MIPS Linux
 - `svc` on ARM Linux
 - `ecall` on RISC-V Linux

Linux Syscall Convention

Instruction	System Call Num.	Arguments	Return
syscall	rax	rdi, rsi, rdx, r10, r8, r9	rax

```
$ cat /usr/include/asm/unistd_64.h
```

```
#ifndef _ASM_UNISTD_64_H
#define _ASM_UNISTD_64_H
```

```
#define __NR_read 0
#define __NR_write 1
#define __NR_open 2
#define __NR_close 3
#define __NR_stat 4
#define __NR_fstat 5
```

```
...
```

For example, the `printf()` function contains the following assembly code:

```
printf:
    ...
    mov    rax, 1
    syscall
    ...
```

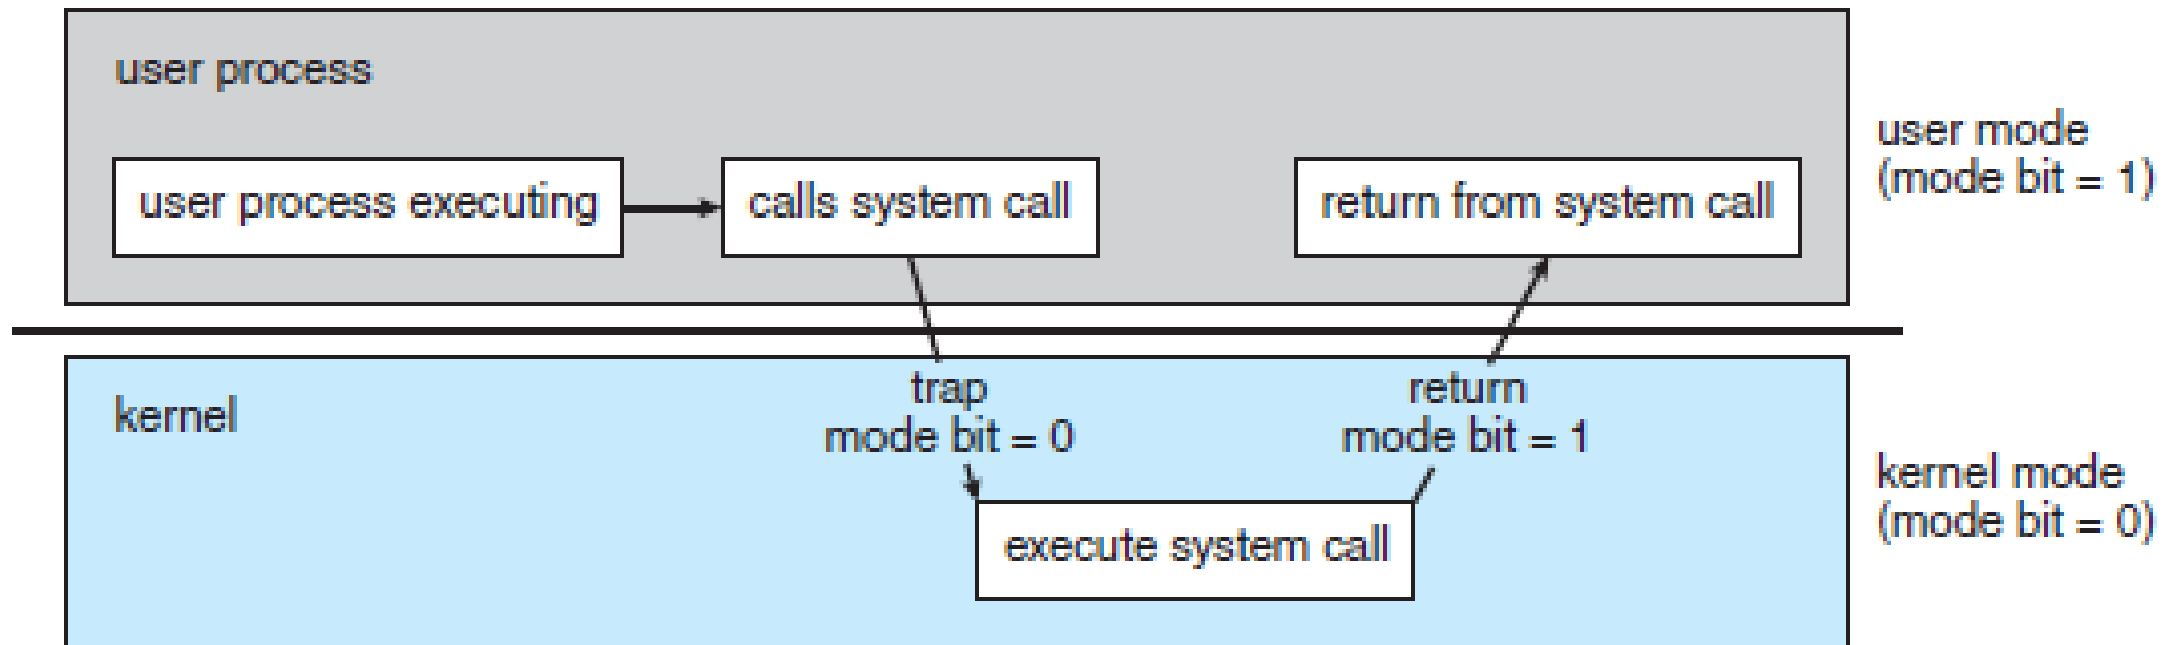
All the System Calls



accept	fchdir	getsockname	mbind	pidfd_getfd	rt_sigaction	sethostname	timer_gettime
accept4	fchmod	getsockopt	membarrier	pidfd_open	rt_sigpending	setitimer	timer_settime
acct	fchmodat	gettid	memfd_create	pidfd_send_signal	rt_sigprocmask	setns	timerfd_create
add_key	fchown	gettimeofday	memfd_secret	pipe2	rt_sigqueueinfo	setpgid	timerfd_gettime
adjtimex	fchownat	getuid	migrate_pages	pivot_root	rt_sigreturn	setpriority	timerfd_settime
bind	fdatasync	getxattr	mincore	pkey_alloc	rt_sigsuspend	setregid	times
bpf	fgetxattr	init_module	mknodat	pkey_free	rt_sigtimedwait	setresgid	tkill
brk	finit_module	inotify_add_watch	mlock	pkey_mprotect	rt_tsigqueueinfo	setresuid	umask
capget	flistxattr	inotify_init1	mlock2	ppoll	sched_get_priority_max	setreuid	umount2
capset	flock	inotify_rm_watch	mlockall	prctl	sched_get_priority_min	setrlimit	uname
chdir	fremovexattr	io_cancel	mount	pread64	sched_getaffinity	setsid	unlinkat
chroot	fsetxattr	io_destroy	mount_setattr	preadv	sched_getattr	setsockopt	unshare
clock_adjtime	fsmount	io_getevents	move_mount	preadv2	sched_getparam	settimeofday	userfaultfd
clock_getres	fsopen	io_pgetevents	move_pages	prlimit64	sched_getscheduler	setuid	utimensat
clock_gettime	fspick	io_setup	mprotect	process_madvise	sched_rr_get_interval	setxattr	vhangup
clock_nanosleep	fsync	io_submit	mq_getsetattr	process_mrelease	sched_setaffinity	shmat	vmsplice
clock_settime	futex	io_uring_enter	mq_notify	process_vm_readv	sched_setattr	shmctl	wait4
clone	futex_waitv	io_uring_register	mq_open	process_vm_writev	sched_setparam	shmdt	waitid
clone3	get_mempolicy	io_uring_setup	mq_timedreceive	pselect6	sched_setscheduler	shmget	write
close	get_robust_list	ioctl	mq_timedsend	ptrace	sched_yield	shutdown	writev
close_range	getcpu	ioprio_get	mq_unlink	pwrite64	seccomp	sigaltstack	
connect	getcwd	ioprio_set	mremap	pwritev	semctl	sigaltd4	
copy_file_range	getdents64	kcmp	msgctl	quotactl	semget	socket	
delete_module	getegid	kexec_file_load	msgget	quotactl_fd	semop	socketpair	
dup	geteuid	kexec_load	msgrcv	read	semtimedop	splice	
dup3	getgid	keyctl	msgsnd	readahead	sendmmsg	statx	
epoll_create1	getgroups	kill	msync	readlinkat	sendmsg	swapoff	
epoll_ctl	getitimer	landlock_add_rule	munlock	readv	sendto	swapon	
epoll_pwait	getpeername	landlock_create_rule_set	munlockall	reboot	set_mempolicy	symlinkat	
epoll_pwait2	getpgid	landlock_restrict_self	munmap	recvfrom	set_mempolicy_home_node	sync	
eventfd2	getpid	lgetxattr	name_to_handle_at	recvmsg	set_robust_list	sync_file_range	
execve	getppid	linkat	nanosleep	recvmmsg	set_tid_address	sync_file_range2	
execveat	getpriority	listen	nfsservctl	remap_file_pages	setdomainname	syncfs	
exit	getrandom	listxattr	open_by_handle_at	removexattr	setfsgid	sysinfo	
exit_group	getresgid	llistxattr	open_tree	renameat	setfsuid	syslog	
faccessat	getresuid	lookup_dcookie	openat	renameat2	setgid	tee	
faccessat2	getrlimit	lremovexattr	openat2	request_key	setgroups	tgkill	
fallocate	getrusage	lsetxattr	perf_event_open	restart_syscall		timer_create	
fanotify_init	getsid	madvise	personality	rseq		timer_delete	
fanotify_mark						timer_getoverrun	

System Call

- The programmatic way in which a computer program requests a service from the operating system



Question?