Jin-Soo Kim
(*jinsoo.kim@snu.ac.kr*)

Systems Software &
Architecture Lab.

Seoul National University

Fall 2022

Introduction to Operating Systems

Operating System Internals

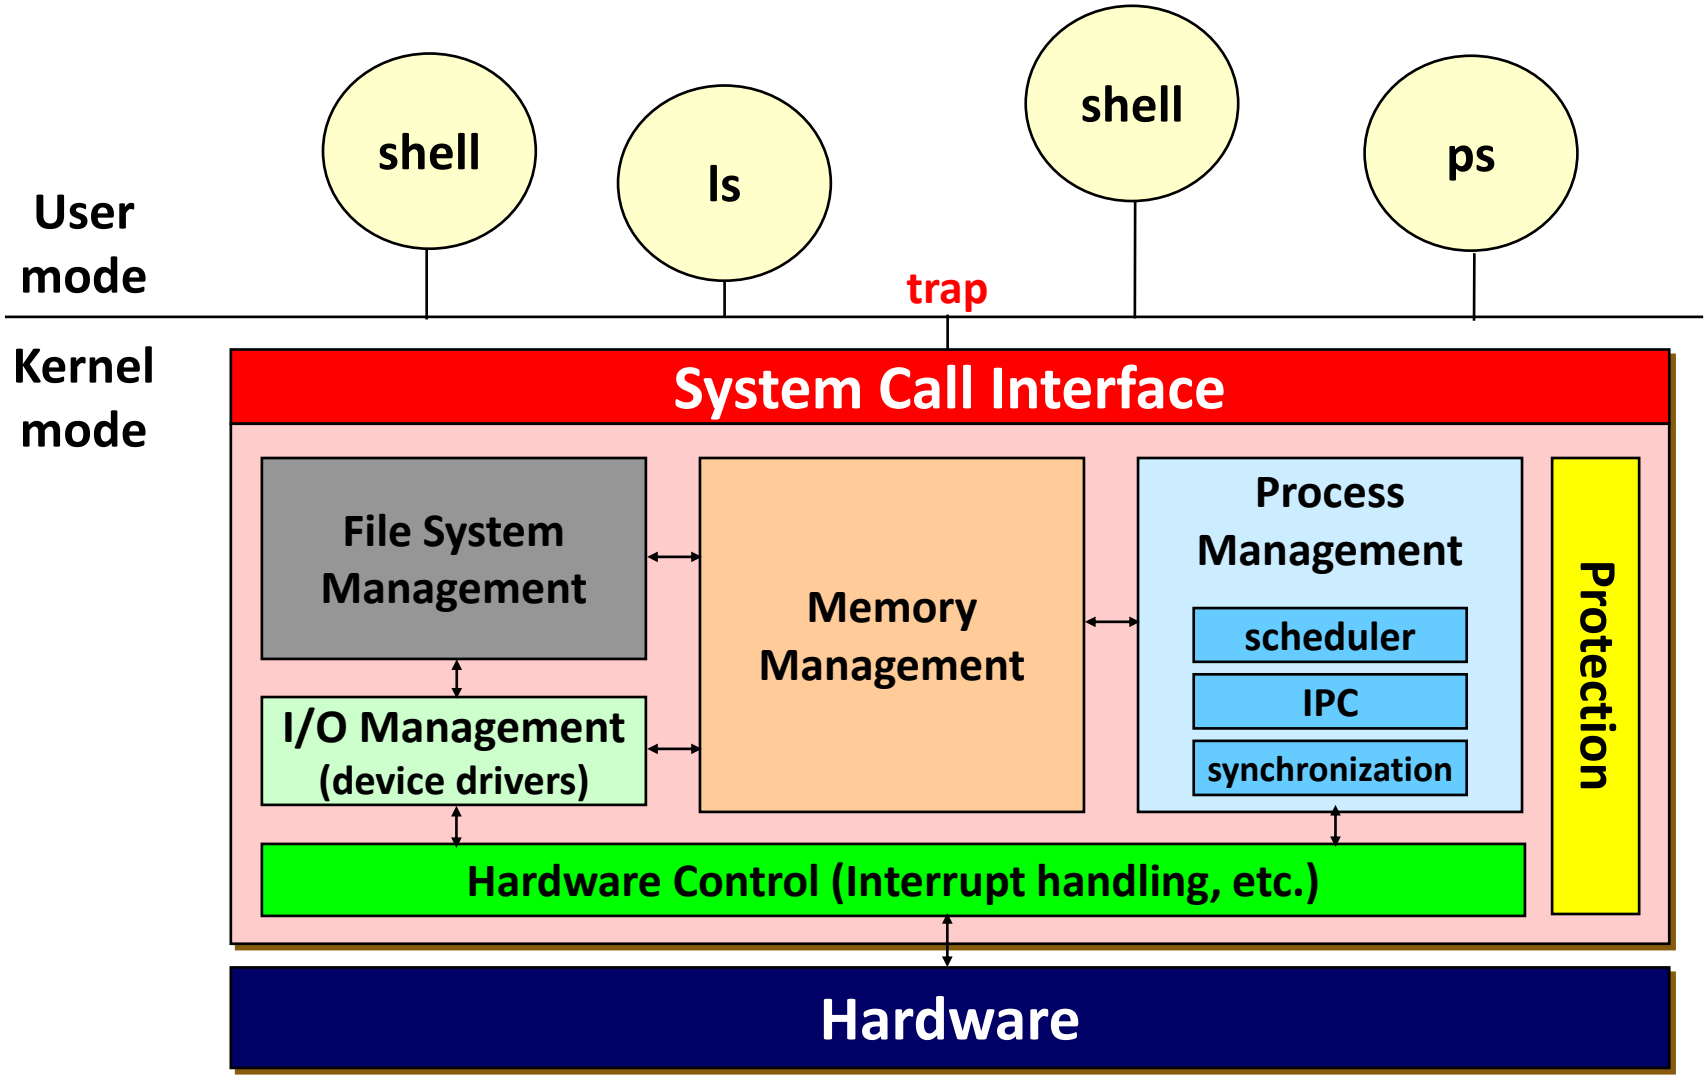

OS: Application View

- OS provides an execution environment for running programs
- OS provides a(an) _____ view of the underlying computer system
  - What are the correct abstractions?
  - How much of hardware should be exposed?
- Typical OS abstractions
  - Processors → Processes, Threads
  - Memory → Address space (virtual memory)
  - Storage → Volumes, Directories, Files
  - I/O Devices → Files (+ ioctls)
  - Networks → Files (sockets, pipes, …)

OS: System View

- OS manages various resources of a computer system

- Sharing
- _____
- Fairness
- Efficiency
- …

Resources

- **CPU**
- **Memory**
- **I/O devices**
- **Queues**
- **Energy**
- **…**

OS: Implementation View

- OS is highly-concurrent, _____ software

- Two kinds of events
  - System calls
  - Interrupts

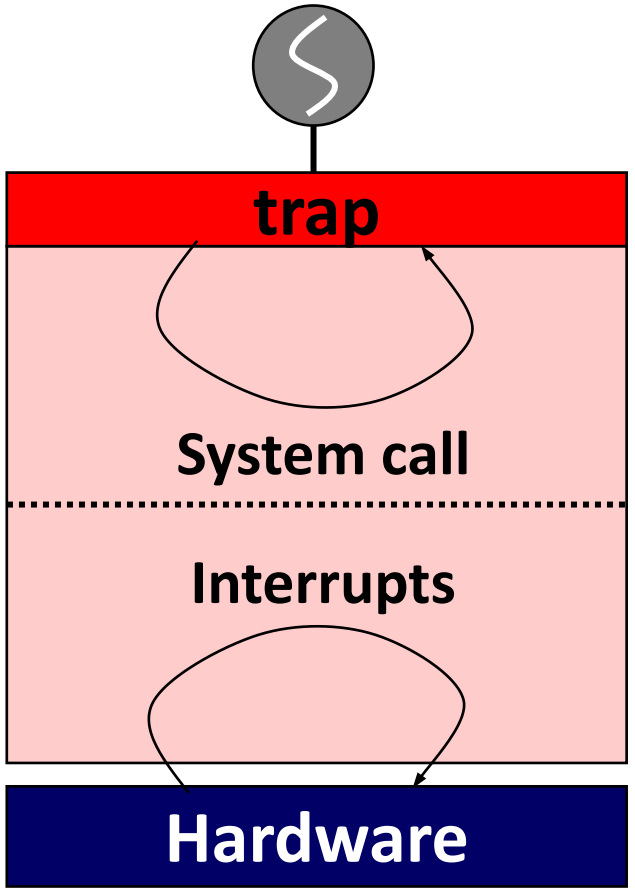

Unix Features

- Process control
  - `fork()`, `exec()`, `wait()`, `exit()`
  - Pipes for inter-process communication (IPC)
- Hierarchical file systems
  - Special files: uniform I/O, naming, and protection
  - Removable file systems via mount/umount
  - i-node
- Signals
- Shells
  - Standard I/O and I/O redirection, filters
  - Shell scripts

Architectural Support for OS (1)

- CPU modes of operation: kernel vs. user
  - 4 levels in x86: Ring 0 > 1 > 2 > 3
  - 3 levels in RISC-V: Machine > Supervisor > User

- Protected or _____ instructions
  - Direct I/O access (e.g., `in`/`out` instructions in x86)
  - Accessing system registers
  - Memory state management
  - …

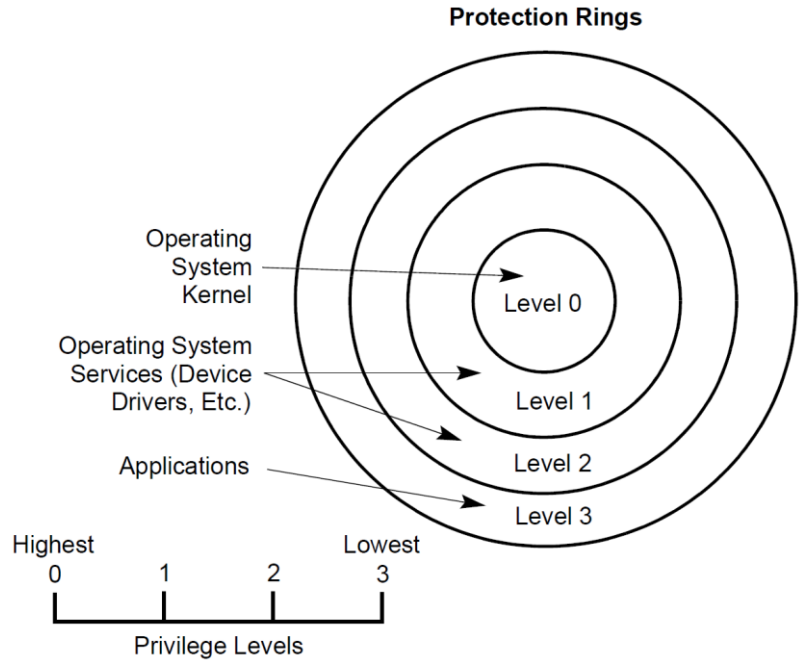

Architectural Support for OS (2)

- Interrupts
  - Generated by hardware devices
  - External interrupts vs. IPIs
  - Asynchronous

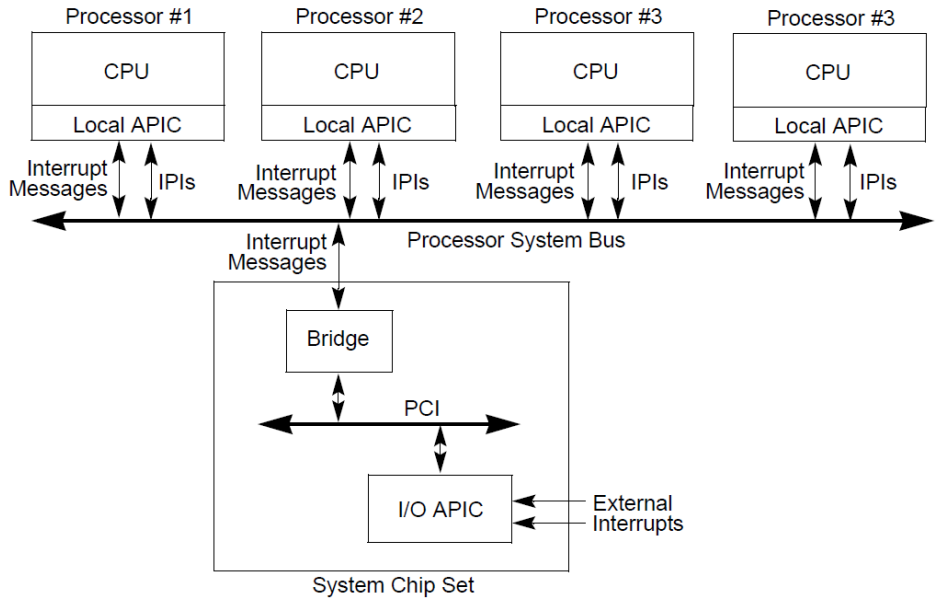

- Exceptions
  - Generated by software executing instructions
    - Faults (unintentional, but possibly recoverable): page faults, protection faults, …
    - Traps (intentional): `syscall` instruction in x86_64 or `ecall` instruction in RISC-V
    - Aborts (unintentional and unrecoverable): parity error, machine error, …
  - Synchronous
  - Exception handling is logically same as interrupt handling

Architectural Support for OS (3)

- Memory protection
  - Segmentation
  - Paging
- Timer
- DMA (Direct Memory Access)
- Atomic instructions
  - Atomic inc/dec
  - Test-and-Set
  - Compare-and-Swap
  - LL (Load Locked) & SC (Store Conditional)
  - …