

CSE 153
**Design of Operating
Systems**

Fall 2018

Midterm Review

Midterm

- in class on Thursday
- Covers material through synchronization and deadlock
- Based upon lecture material and modules of the book indicated on the class schedule
 - ◆ Closed book. No additional sheets of notes

Overview

- Architectural support for Oses
- Processes
- Threads
- Synchronization

Arch Support for OSes

- Types of architecture support
 - ◆ Manipulating privileged machine state
 - ◆ Generating and handling events
 - ◆ Instruction for locking

Privileged Instructions

- What are privileged instructions?
 - ◆ Who gets to execute them?
 - ◆ How does the CPU know whether they can be executed?
 - ◆ Difference between user and kernel mode
- Why do they need to be privileged?
- What do they manipulate?
 - ◆ Protected control registers
 - ◆ Memory management
 - ◆ I/O devices

Events

- Events
 - ◆ Synchronous: faults (exceptions), system calls
 - ◆ Asynchronous: interrupts
- What are faults, and how are they handled?
- What are system calls, and how are they handled?
- What are interrupts, and how are they handled?
 - ◆ How do I/O devices use interrupts?
- What is the difference between exceptions and interrupts?

Processes

- What is a process?
- What resource does it virtualize?
- What is the difference between a process and a program?
- What is contained in a process?

Process Data Structures

- Process Control Blocks (PCBs)
 - ◆ What information does it contain?
 - ◆ How is it used in a context switch?
- State queues
 - ◆ What are process states?
 - ◆ What is the process state graph?
 - ◆ When does a process change state?
 - ◆ How does the OS use queues to keep track of processes?

Process Manipulation

- What does CreateProcess on Windows do?
- What does fork() on Unix do?
 - ◆ What does it mean for it to “return twice”?
- What does exec() on Unix do?
 - ◆ How is it different from fork?
- How are fork and exec used to implement shells?

Threads

- What is a thread?
 - ◆ What is the difference between a thread and a process?
 - ◆ How are they related?
- Why are threads useful?
- What is the difference between user-level and kernel-level threads?
 - ◆ What are the advantages/disadvantages of one over another?

Thread Implementation

- How are threads managed by the run-time system?
 - ◆ Thread control blocks, thread queues
 - ◆ How is this different from process management?
- What operations do threads support?
 - ◆ Fork, yield, sleep, etc.
 - ◆ What does thread yield do?
- What is a context switch?
- What is the difference between non-preemptive scheduling and preemptive thread scheduling?
 - ◆ Voluntary and involuntary context switches

Synchronization

- Why do we need synchronization?
 - ◆ Coordinate access to shared data structures
 - ◆ Coordinate thread/process execution
- What can happen to shared data structures if synchronization is not used?
 - ◆ Race condition
 - ◆ Corruption
 - ◆ Bank account example
- When are resources shared?
 - ◆ Global variables, static objects
 - ◆ Heap objects

Mutual Exclusion

- What is mutual exclusion?
- What is a critical section?
 - ◆ What guarantees do critical sections provide?
 - ◆ What are the requirements of critical sections?
 - » Mutual exclusion (safety)
 - » Progress (liveness)
 - » Bounded waiting (no starvation: liveness)
 - » Performance
- How does mutual exclusion relate to critical sections?
- What are the mechanisms for building critical sections?
 - ◆ Locks, semaphores

Locks

- What does Acquire do?
- What does Release do?
- What does it mean for Acquire/Release to be atomic?
- How can locks be implemented?
 - ◆ Spinlocks
 - ◆ Disable/enable interrupts
- How does test-and-set work?
 - ◆ What kind of lock does it implement?
- What are the limitations of using spinlocks, interrupts?
 - ◆ Inefficient, interrupts turned off too long

Semaphores

- What is a semaphore?
 - ◆ What does Wait/P/Decrement do?
 - ◆ What does Signal/V/Increment do?
 - ◆ How does a semaphore differ from a lock?
 - ◆ What is the difference between a binary semaphore and a counting semaphore?
- When do threads block on semaphores?
- When are they woken up again?
- Using semaphores to solve synchronization problems
 - ◆ Readers/Writers problem
 - ◆ Bounded Buffers problem

Deadlock

- Deadlock happens when processes are waiting on each other and cannot make progress
- What are the conditions for deadlock?
 - ◆ Mutual exclusion
 - ◆ Hold and wait
 - ◆ No preemption
 - ◆ Circular wait
- How to visualize, represent abstractly?
 - ◆ Resource allocation graph (RAG)
 - ◆ Waits for graph (WFG)