

# CS3844 Operating Systems

Dr. Walter Schilling

Winter, 2013-2014

For the final exam, you may bring one 8.5 x 11 inch sheet of paper with notes.

The final exam will consist of approximately 50 to 60 percent material from after the midterm and 40 to 50 percent from before the midterm.

Course and Section	Exam Time and location
CS 3844 001	Friday 8:00 AM - 10:00 AM S206
CS 3844 002	Thursday 2:00 PM - 4:00 PM R301

## 1. Week #1

### (a) Lecture #1 (An Introduction to Operating Systems)

- i. Draw an abstract view of a computer system.
- ii. Explain the difference between the user view and the system view of an operating system.
- iii. Explain the difference between the kernel, systems programs, and applications programs.
- iv. Draw a representation of a modern computer system.

### (b) Lecture #2 (System Calls)

- i. List the two modes of operation for a microprocessor
- ii. Define the concept of a trap
- iii. Explain how a system call is made
- iv. Describe various methods for handling parameters as they are passed to System calls.
- v. List some examples of system calls
- vi. Understand how a system call may manifest itself in x86 assembly language code.
- vii. Draw the C flow of C compilation from source code to object code.
- viii. Explain the purpose for the preprocessor, compiler, and linker within the C compilation model
- ix. Using the gcc compiler, generate the output for the preprocessor stage of compilation
- x. Explain the concept of a dependency.
- xi. Create a GNU Make file which automatically generates dependencies, creates preprocessed source code, and links a given C application.

### (c) Lecture #3 (Operating Systems Structures)

- i. Compare and contrast simple structured operating systems, layered operating systems, microkernels, and module based operating systems.
- ii. List the limitations of the MS-DOS operating system.
- iii. Draw a picture for a layered operating system.
- iv. List the advantages of a layered operating system.
- v. List the problems of designing a layered operating system.
- vi. Explain the fundamental purpose for the microkernel within a microkernel based operating system.
- vii. Explain the concept of a module based operating system
- viii. List the advantages and disadvantages of each approach to operating systems design.
- ix. Explain how to change the executing shell on a Linux machine.
- x. Explain the relationship between a shell and a process in Linux.

## 2. Week #2

### (a) Lecture #1 (Processes)

- i. Define the term process

- ii. Define the terms job, user programs, and tasks.
- iii. Explain the contents of the text section, data section, heap, and stack of a program
- iv. Draw a graphical representation of a process in memory
  - v. Explain the concept of process state
- vi. Draw a state transition diagram for process states
- vii. List the contents of a process control block
- viii. Explain what the process scheduler is responsible for doing within the operating system.
- ix. Explain the concept of process dispatching
  - x. List the processes executing on a Linux workstation.
- xi. Explain how the kill program can be used to terminate an executing program.

(b) Lecture #2 (Process Context Switching)

- i. Explain how a CPU Context switch occurs
- ii. Explain how the hardware may impact the time necessary for a context switch (i.e. Sun Ultra Sparc)
- iii. List reasons why a context switch would occur
- iv. Explain why context switching can be bad
  - v. Compare and contrast IO Bound and CPU Bound processes
- vi. Explain how the OS performs a context switch
- vii. Construct code which switches the stack as part of an OS context switch
- viii. Understand how process state is stored within a process control block
  - ix. Explain how to terminate a process

(c) Lecture #3 (No class - Rockwell Collins Trip)

3. Week #3

(a) Lecture #1 (Process Operations)

- i. Explain the concept of the parent - child relationship between processes
- ii. Explain the purpose for the UNIX fork, wait, and exec commands.
- iii. Explain what happens when the exit method is called.
- iv. Explain the concept of cascading termination
  - v. Explain the purpose for the UNIX fork, wait, and exec commands.
- vi. Construct programs using the fork, wait, and exec unix commands

(b) Lecture #2 (Interprocess Communications)

- i. Explain why processes should be allowed to operate in parallel
- ii. List two methods for sharing information between processes
- iii. Draw graphical representation for processes communicating using shared memory and message passing systems.
- iv. Explain the flow necessary to use a shared memory partition
  - v. List and define the 4 main synchronization types that can be employed with message passing systems
- vi. List the advantages of each mechanism for interprocess communication
- vii. Construct a rudimentary program which uses pipes to communicate between processes.

4. Week #4

(a) Lecture #1 (Interprocess Communications (Sockets and RPC))

- i. Define a socket
- ii. Define the acronym RPC
- iii. Explain how an RPC executes, specifically in regards to stubs and the concept of marshaling.
- iv. Explain the difference between “big-endian” and “little-endian”.
  - v. Construct a simple application in Linux using RPC.
- vi. Communicate between two programs using an RPC call.

(b) Lecture #2 (Threads)

- i. Explain the concept of a thread
- ii. Draw a representation of a single threaded process and a multi-threaded process.

- iii. Compare and Contrast the advantages and disadvantages of threads versus processes
- iv. Explain how multi-threaded program can be useful in a multi-core environment.
- v. Explain the difference between kernel threads and user threads
- vi. Explain the difference between many to one, one to one, and many to many models of thread behavior
- vii. Using C, construct a simple multithreaded POSIX compliant application.

(c) Lecture #3 (Thread Problems)

- i. Explain the interaction between threads and fork?
- ii. Explain the difference between asynchronous and deferred cancelation.
- iii. Explain the risks of improper termination of threads
- iv. Define the concept of a UNIX Signal.
- v. Explain the concept of a thread pool
- vi. Explain the concept of a UNIX signal.
- vii. Explain how a UNIX signal can be handled.
- viii. Construct rudimentary programs which handle a UNIX signal.

5. Week #5

(a) Lecture #1 (Process Synchronization)

- i. Define race condition.
- ii. Define critical section.
- iii. Explain the design ramifications of a preemptive kernel versus a non-preemptive kernel in terms of critical sections.
- iv. Define mutual exclusion
- v. Define an atomic operation
- vi. Explain the problem of deadlocks and starvation
- vii. Explain the problem of priority inversion and justify why priority inversion solves this problem
- viii. List 4 mechanisms that can be used to prevent deadlocks.
- ix. Explain the difference between a mutex and a semaphore.
- x. Explain the concept of a counting semaphore.
- xi. Perform basic synchronization using PThreads mutexes.
- xii. Use semaphores to perform basic synchronization operations.

(b) Lecture #2 (Process Synchronization - Implementation)

- i. Be able to construct a basic synchronization primitive in C source code.
- ii. Explain the purpose for each entry within a semaphore structure.

(c) Lecture #3 (The Dining Philosophers)

- i. Construct a resource allocation graph.
- ii. List the conditions necessary for a deadlock to occur.
- iii. Explain the dining philosophers problem and how it results in a potential deadlock.
- iv. Construct a resource allocation graph from a given problem description.
- v. Analyze a resource allocation graph to determine if a deadlock is present within the system.
- vi. Analyze source code to determine if a deadlock is present based on execution traces.

6. Week #6

(a) Lecture #1 (Scheduling)

- i. Explain the CPU and IO Burst cycle used for scheduling
- ii. Recognize the distribution of CPU activities on a system
- iii. Explain the relationship between an IO bound program and CPU bound program in terms of CPU bursts
- iv. List the four reasons why the scheduler may be invoked
- v. Compare and Contrast Pre-emptive and non-preemptive scheduling. What are the advantages of one system versus the other, and how is the operating system different based on the two approaches?
- vi. Explain the purpose for the dispatcher and scheduler within the operating system.
- vii. Define dispatch latency

viii. Define CPU utilization, Throughput, Turnaround time, Waiting time, Response time in terms of their impact on scheduling.

(b) Lecture #2 (FIFO Scheduling)

- i. Explain the operation of a FIFO scheduler
- ii. Calculate the average waiting time for a given set of processes using FCFS scheduling
- iii. Construct a GANTT chart for a given set of processes
- iv. Calculate the throughput for a given system.
- v. Explain the convoy effect of FCFS scheduling
- vi. List the advantages and disadvantages of FCFS scheduling.

(c) Lecture #3 (Midterm Exam)

- i. Obtain a successful exam score.

7. Week #7

(a) Lecture #1 (SJF Scheduling)

- i. Explain the algorithm for SJF Scheduling
- ii. Explain why exponential averaging can be used to estimate the shortest job burst.
- iii. Calculate the exponential average based on a series of CPU bursts and an initial estimate.
- iv. List the advantages and disadvantages of SJF scheduling

(b) Lecture #2 (Priority Scheduling)

- i. Explain priority scheduling.
- ii. Using priority scheduling, draw a schedule for a set of jobs
- iii. Define starvation in terms of processor scheduling
- iv. Demonstrate how processor aging can solve the process of starvation
- v. Explain how the Linux scheduler is implemented.
- vi. Explain how the Linux scheduler obtains  $O(1)$  scheduling.

(c) Lecture #3 (Memory Management)

- i. Explain how the base and limit registers are used to trigger a trap.
- ii. List three methods of address binding and explain the difference.
- iii. Explain the difference between a logical address and a physical address.

8. Week #8

(a) Lecture #1 (Swapping)

- i. Explain what occurs when memory is swapped.
- ii. Define backing store
- iii. Define fragmentation.
- iv. Explain the usage for relocation and limit registers
- v. In the context of paging, explain frames and pages.
- vi. Calculate how long it will take to perform a memory swap.

(b) Lecture #2 (Paging)

- i. Define the terms frame and page
- ii. Explain the concept of a page table
- iii. Explain how to determine the number of bits in a page number and a page offset
- iv. Explain the operation of a translation look-aside buffer (TLB)
- v. Calculate the term hit ratio
- vi. Calculate the effective memory access time
- vii. Explain the purpose for the valid-invalid bit in a page table

(c) Lecture #3 (Virtual Memory)

- i. Explain the relationship between Virtual memory and physical memory.
- ii. Explain the concept of a virtual address space.
- iii. Draw a diagram showing how virtual memory may allow two processes to share a library.
- iv. Explain how processes are swapped in and out of virtual memory.

- v. Explain the operation of the page table.
- vi. List the steps necessary to handle a page fault.
- vii. Explain the concept of copy on write.

## 9. Week #9

### (a) Lecture #1 (Page Replacement Algorithms)

- i. Explain the concept of page replacement.
- ii. Define victim frame.
- iii. Explain the purpose for the dirty bit within a virtual memory system.
- iv. Explain thrashing.
- v. Compare and contrast FIFO, OPT, and LRU page replacement algorithms, noting performance differences and implementation differences.
- vi. Explain Belady's anomaly.

### (b) Lecture #2 (Memory Mapped IO)

- i. Explain the concept of memory mapped IO.
- ii. Explain how processes can share the same version of a file by mapping to the same location in memory.
- iii. Explain the concept of memory mapped IO.
- iv. Explain how processes can share the same version of a file by mapping to the same location in memory.
- v. Explain the purpose for the mmap command to map a file into memory.
- vi. Construct a simple application which uses mmap.

### (c) Lecture #3 (File System Structure)

- i. List common file types within file systems
- ii. Explain the concept of a partition
- iii. Compare and contrast single level directories, two level directories, tree structured directories, and acyclic-graph directories
- iv. Explain the difference between absolute and relative path names
  - v. Explain how the file system is organized within UNIX
  - vi. Explain the concept of mounting a file system
  - vii. Describe the contents of the UNIX fstab file
  - viii. Using shell commands, change the access for UNIX files.
  - ix. Using shell commands, link a file in UNIX.

## 10. Week #10

### (a) Lecture #1 (File System Implementation)

- i. Explain the concepts of a layered file system
- ii. List the contents of a file control block.
- iii. Explain the concept of a partition and mounting of a partition.
- iv. Explain the relationship between continuous allocation, linked allocation, and indexed allocation within a file system

### (b) Lecture #2 (File Protection)

- i. Explain the goals of protection.
- ii. Explain the difference between users and groups.
- iii. In Unix, be capable of changing the group access for a file.
- iv. In Linux, explain how access is allowed to a file.
  - v. Explain how access may be conveyed from a group or user.
  - vi. Explain in Unix how a file may be converted to be executable.

### (c) Lecture #3 (Final Exam Review)

- i. Prepare to pass the test with an exemplary grade.