CS3841 Operating Systems

Dr. Walter Schilling

Fall, 2012

You may use 1 8.5 x 11 inch sheet of paper with notes and other supporting material for the exam.
The exam is scheduled for Monday, November 12, 2012 from 11:00 - 13:00.

1. Week #1
   (a) Lecture #1 Working in C
      i. Draw the C flow of C compilation from source code to object code.
      ii. Explain the purpose for the preprocessor, compiler, and linker within the C compilation model
      iii. Using the gcc compiler, generate the output for the preprocessor stage of compilation
      iv. Explain the concept of a dependency.
      v. Create a GNU Make file which automatically generates dependencies, creates preprocessed source code, and links a given C application.
   (b) Lecture #2 Introduction to Operating Systems
      i. Compare and Contrast the User View and System View of an operating system.
      ii. Explain the difference between user mode and kernel mode within an operating system.
      iii. Draw the storage structure hierarchy for a computer system.
      iv. Explain the difference between a trap and an interrupt.
      v. Explain, in the context of an operating system, multiprogramming.
      vi. Explain, in the context of an operating system, time sharing.
      vii. Understand and use the ls, man, cd, rm, cp, cat, more, less, tar, sort, kill, and ps commands.
   (c) Lecture #3 Operating Systems Structures
      i. List and characterize operating systems services (User interface, program execution, IO, file system manipulation, communications, error detection, resource allocation, accounting, protection and security)
      ii. Compare and contrast the command interpreter and graphical user interface approaches to interface with the computer.
      iii. Compare and contrast approaches to command interpreter implementation
      iv. List various UNIX shells
      v. Explain how a system call is made
      vi. Explain the concept of a system call
      vii. Explain the usage of the malloc and free operations within the C programming language.
      viii. Construct simple C programs which use malloc and free to solve problems.
      ix. Implement Screen and File I/O in C, showing how the system calls are invoked
2. Week #2
   (a) Lecture #1 Operating Systems Design and Virtual Machines
      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 relationship between a layered architecture and a virtual machine.
      viii. List the benefits of using a virtual machine.

(b) Lecture #2 Processes

i. Explain the flow of control when an operating system boots
ii. Define the term process
iii. Draw a graphical representation of a process in memory
iv. Explain the concept of process state
v. Draw a state transition diagram for process states
vi. List the contents of a process control block
vii. Explain what the process scheduler is responsible for doing within the operating system.
viii. Explain the concept of process dispatching
ix. Obtain information about the executing processes under Windows and Linux

(c) Lecture #3 Process Operations

i. Explain how a CPU Context switch occurs
ii. List reasons why a context switch would occur
iii. Explain why context switching can be bad
iv. Compare and contrast IO Bound and CPU Bound processes
v. Explain the purpose for the UNIX fork, wait, and exec commands.
vi. Construct programs using the fork, wait, and exec unix commands
vii. Explain how a process is terminated.
viii. Execute a UNIX command in the background using the shell
ix. Use the UNIX command shell to terminate a process

3. Week #3

(a) Lecture #1

i. Explain why it is important to allow processes to execute in parallel.
ii. List two methods for interprocess communication
iii. Explain the difference between indirect and direct communication in terms of message passing.
iv. Explain how UML sequence diagrams can be used to represent interprocess communications.
v. List the advantages and disadvantages of using shared memory for interprocess communication.
vi. List the advantages and disadvantages of using pipes for interprocess communication.
vii. Construct a rudimentary program using pipes.

(b) Lecture #2

i. Define a socket
ii. Explain the difference between “big-endian” and “little-endian”.

(c) Lecture #3

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 concept of the join call relative to a thread
vii. Implement multi-threaded software using POSIX threads in C.

4. Week #4

(a) Lecture #1

i. Explain the interaction between threads and fork?
ii. Explain the risks of improper termination of threads

(b) Lecture #2

i. Explain the CPU and IO Burst cycle used for scheduling
ii. Explain the relationship between an IO bound program and CPU bound program in terms of CPU bursts
iii. List the five reasons why the scheduler may be invoked

iv. 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?
v. Explain the purpose for the dispatcher and scheduler within the operating system.
vi. Define CPU utilization, Throughput, Turnaround time, Waiting time, Response time in terms of their impact on scheduling.
vii. Explain the operation of a FIFO scheduler
viii. Explain the convoy effect of FCFS Scheduling

(c) Lecture #3

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. Using priority scheduling, draw a schedule for a set of jobs
v. Explain round robin scheduling
vi. Explain the relationship between quantum length and performance.
vii. For all scheduling algorithms
A. Draw GANTT Chart showing processing sequence
B. Calculate the average waiting time
viii. Justify the design decisions for the Linux kernel based upon scheduling theory

5. Week #5

(a) Lecture #1

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

(b) Lecture #2

i. Perform basic synchronization using PThreads mutexs.
ii. Explain the concept of a semaphore
iii. Compare and contrast a counting semaphore with a binary semaphore
iv. Explain the concept of a spinlock.
v. Define a deadlock within a process synchronization system.
vi. Explain the concept of a priority inversion.

(c) Lecture #3

i. Exam Review and Catchup

6. Week #6

(a) Lecture #1

i. Midterm Exam

(b) Lecture #2

i. Explain the dining philosophers problem and how it results in a potential deadlock.
ii. List the conditions necessary for a deadlock to occur.
iii. Construct a resource allocation graph from a given problem description.
iv. Analyze a resource allocation graph to determine if a deadlock is present within the system.
v. List three methods for handling deadlocks.

(c) Lecture #3

i. Continuation of other material.

7. Week #7

(a) Lecture #1

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.
iv. Explain what occurs when memory is swapped.
v. Define fragmentation.
vi. In the context of paging, explain frames and pages.
vii. Given a logical address and a page table, calculate the physical address for a piece of memory.
viii. Explain the overhead with using paged memory

(b) Lecture #2

i. Explain how a shared library may be loaded with virtual memory.
ii. Define demand paging.
iii. Explain the purpose for the valid-invalid bit within a virtual memory page table.
iv. Draw a diagram showing the steps to handle a page fault in a virtual memory system.
v. Explain the interaction between copy on write and invocations of the fork command.

(c) Lecture #3

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. Compare and contrast FIFO, OPT, and LRU page replacement algorithms, nothing performance differences and implementation differences.

8. Week #8

(a) Lecture #1

i. Explain thrashing and its causes.
ii. Explain the relationship between memory-mapped files and virtual memory.
iii. Lecture #2
A. List the attributes of a file.
B. List the operations on a file.
C. List the information associated with an open file.
D. Draw the flow for a file write.
E. Draw the flow representing a file read.
F. Compare and contrast direct access and sequential access to files
G. Define the terms partition, volume, and directory.
H. Explain the difference between absolute and relative path names
I. Compare and contrast single level directories, two level directories, tree structured directories, and acyclic-graph directories. What are the advantages of each system? What are the disadvantages of each system?
J. Using UNIX, create a link between one file and another file.
iv. Lecture #3
A. Explain access control.
B. In terms of a UNIX file system, define the terms owner, group, and universe.
C. Using UNIX commands, control the access to a given file and list the access rights to a given file.
D. List commonly used File Systems.
E. Explain the purpose for the mount table.
F. Describe the contents of the UNIX fstab file
G. Using shell commands, change the access for UNIX files.
H. Using shell commands, link a file in UNIX.
I. Explain the difference between absolute and relative path names

(b) Week #9

i. Lecture #1
A. Explain how contiguous allocation of files works within a file system.
B. Critique the effectiveness of contiguous allocation.
C. Explain the concept of linked allocation.

D. Describe the purpose for the FAT.
E. Explain the operation of indexed allocation.
F. Describe the purpose for the UNIX inode.

ii. Lecture #2
A. Define the terms track, sector, cylinder, and platter in terms of a magnetic disk
B. Explain the concept of memory mapped I/O
C. Compare and contrast polling and interrupts for device management.
D. Explain the concept of direct memory access
E. Justify the usage of direct memory access from a performance standpoint
F. Explain the application level interface for devices
G. Compare and contrast blocking and non-blocking I/O
H. Justify the usage of kernel mode for I/O implementation

iii. Lecture #3
A. Continuation...

(c) Week #10

i. Lecture #1
A. Continuation...

ii. Lecture #2
A. Continuation...

iii. Lecture #3
A. Continuation...

1 Lab Outcomes

1. Lab 1: Getting used to Linux

   (a) Demonstrate an ability to use a Linux shell.
   (b) Use the man command to obtain documentation about Linux commands.
   (c) Explain how to list the contents of a directory in multiple forms.
   (d) Navigate the Linux file system by changing directories.
   (e) Manage the creation and deletion of new files and directories from within the command shell.
   (f) Capture the output of a Linux program executing to a file.
   (g) Manage the creation and extraction of zip files and tarballs using the command shell.
   (h) Construct a makefile which will automatically generate the project as well as allow for the clean building of source code.

2. Lab 2: Memory Management and Data Structures in C

   (a) Use malloc and free to manage the allocation and deallocation of dynamic memory.
   (b) Implement a doubly linked list in C.
   (c) Understand the purpose for the void pointer in C.
   (d) Apply appropriate casts to correctly use a void pointer.
   (e) Implement and use C struct to solve a software problem.
   (f) Use test cases to verify the correct operation of a constructed source code module.

3. Lab 3: Counting Words

   (a) Practice C development in a UNIX environment.
   (b) Construct software in C which uses File input and output routines.
   (c) Manage dynamic memory and heap allocation using C methods.
   (d) Use previously developed libraries as a part of a software development.
   (e) Practice the usage of UNIX piping to chain UNIX programs.

4. Lab 4: Counting Words: Part 2
    (a) Use command line parameters to pass data between programs.
    (b) Manage dynamic memory and heap allocation using C methods.
    (c) Manage the spawning of additional processes from within a UNIX program using the fork method.
    (d) Perform interprocess communication between forked processes using pipes.
    (e) Analyze the performance of a program in the UNIX operating environment.
5. Lab 6-8:
    (a) Construct code which uses POSIX threads.
    (b) Construct code which uses sockets to communicate.
    (c) Construct software which protects against race conditions using semaphores and other protection mechanisms.
6. Lab 9-10:
    (a) Understand the operation of dynamic memory management through the implementation of a dynamic memory manager.
    (b) Practice C development in a UNIX environment.
    (c) Construct software which uses C structures and pointer references.