

CSE 451: Operating Systems Winter 2007

Module 24 Course Review

Ed Lazowska
lazowska@cs.washington.edu
Allen Center 570

Architectural Support

- Privileged instructions
 - what are they, and who gets to execute them?
 - how does CPU know whether to execute them?
 - why do they need to be privileged?
 - what do they manipulate?
- Events
 - exceptions: what generates them? trap vs. fault?
 - interrupt: what generates them?

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

2

OS Structure

- What are the major components of an OS?
- How are they organized?
 - what is the difference between monolithic, layered, microkernel OS's?
 - advantages and disadvantages?

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

3

Processes

- What is a process? What does it virtualize?
 - differences between program, process, thread?
 - what is contained in process?
 - what does PCB contain?
 - state queues?
 - which states, what transitions are possible?
 - when do transitions happen?
- Process manipulation
 - what does fork() do? how about exec()?
 - how do shells work?

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

4

Threads

- What is a thread?
 - why are they useful?
 - user level vs. kernel level threads?
- How does thread scheduling differ from process scheduling?
 - what operations do threads support?
 - what happens on a thread context switch? what is saved in TCB?
 - preemptive vs. non-preemptive scheduling?

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

5

Synchronization

- Why do we need it?
 - data coordination? execution coordination?
 - what are race conditions? when do they occur?
 - when are resources shared? (variables, heap objects, ...)
- What is mutual exclusion?
 - what is a critical section?
 - what are the requirements of critical sections?
 - mutex, progress, bounded waiting, performance
 - what are mechanisms for programming critical sections?
 - locks, semaphores, monitors, condition variables

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

6

Locks and Semaphores

- What does it mean for acquire/release to be atomic?
- how can locks be implemented?
 - spinlocks? interrupts? OS/thread-scheduler?
 - test-and-set?
 - limitations of locks?
- Semaphores
 - wait vs. signal? difference between semaphore and lock?
 - when do threads block on semaphores? when do they wake?
 - bounded buffers problem
 - producer/consumer
 - readers/writers problem

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

7

Process Scheduling

- Long term vs. short term
- When does scheduling happen?
 - job changes state, interrupts, exceptions, job creation
- Scheduling goals?
 - maximize CPU utilization
 - maximize job throughput
 - minimize (turnaround time | waiting time | response time)
 - batch vs. interactive: what are their goals?
- What is starvation? what causes it?
- FCFS/FIFO, SPT, SRPT, priority, RR, MLFQ...

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

8

Memory Management

- What good is virtual memory?
- Mechanisms for implementing memory management
 - physical vs. virtual addressing
 - partitioning, paging, segmentation
 - page tables, TLB
- Page replacement policies?
- What are overheads related to memory management?

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

9

Virtualizing Memory

- What is difference between a physical and virtual address?
 - fixed vs. variable partitioning
 - base/limit registers..
 - internal vs. external fragmentation
- Paging
 - advantages, disadvantages?
 - what are page tables, PTEs?
 - what are: VPN, PFN, offset? relationship to VA?
 - what's in a PTE? what are modify/reference/valid/prot bits?

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

10

Paging, TLBs

- How to reduce overhead of page table?
 - how do multi-level page tables work?
 - what problem does TLB solve?
 - why do they work?
 - how are they managed?
 - software vs. hardware managed?
- Page faults
 - what is one? how is it used to implement demand paging?
 - what is complete sequence of steps for translating a virtual address to a PA?
 - all the way from TLB access to paging in from disk
- MM tricks
 - shared memory? Mapped files? copy-on-write?

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

11

Page Replacement

- What is page replacement algorithm?
 - what application behavior does it exploit?
 - when is page replacement algorithm invoked?
- Understand:
 - Belady's (optimal), FIFO, LRU, approximations of LRU, LRU clock, working set, page fault frequency
 - what is thrashing? why does it occur and when?

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

12

Disks

- Memory hierarchy and locality
- Physical disk structure
 - platters, surfaces, tracks, sectors, cylinders, arms, heads
- Disk interface
 - how does OS make requests to the disk?
- Disk performance
 - access time = seek + rotation + transfer
- Disk scheduling
 - how does it improve performance?
 - FCFS, SSTF, SCAN, C-SCAN?

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

13

Files and Directories

- What is a file
 - what operations are supported?
 - what characteristics do they have?
 - what are file access methods?
- What is a directory
 - what are they used for?
 - how are they implemented?
 - what is a directory entry?
- How does path name translation work?
- ACLs vs. capabilities
 - matrix
 - advantages and disadvantages of each

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

14

File system data structures

- General strategies?
 - contiguous, linked, indexed?
 - tradeoffs?
- What is an inode?
 - how are they different than directories?
 - how are inodes and directories used to do path resolution, and find files?
- Everything about the Unix File System (UFS)

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

15

FS buffer cache

- What is a buffer cache?
 - why do OS's use them?
- What are differences between caching reads and writes?
 - write-through, write-back, write-behind?
 - read-ahead?

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

16

FFS, JFS, LFS

- What is FFS, how specifically does it improve over original Unix FS?
- How about JFS, what is the key problem that it solves, what are the basic ideas?
- How about LFS, what are the basic ideas, when does it yield an improvement, when does it not?

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

17

RAID

- Basic concepts of RAID
 - stripe files across multiple disks to improve throughput
 - compensate for decreased reliability with parity/ECC
- Sources of improvement as you go from RAID-0 to RAID-5

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

18

Networking

- ISO 7-layer model
- Ethernet protocol
- IP and routing
- TCP principles (sending a long message via postcards)
- Protocol encapsulation

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

19

RPC

- Basic idea – what does it buy you over message passing?
- Subtopics: interface description language, stubs, stub generation, parameter marshaling, binding, runtime/transport, error handling, performance, thread pools
- Transparency: when is distribution transparent, when is it not?

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

20

Distributed systems

- Loosely-coupled, closely-coupled, tightly-coupled
- Grapevine as an example, in detail

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

21

Distributed file systems

- Issues:
 - Basic abstraction, naming, caching, sharing/coherency, replication, performance
- Examples – compare and contrast:
 - NFS
 - AFS
 - Sprite

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

22

Security

- Principals, objects, rights
- Authentication, authorization, auditing
- “Gotchas” with simple password protection
- Trusted third parties in distributed systems
- Spyware
- Confinement

3/4/2007

© 2007 Gribble, Lazowska, Levy, Zahorjan

23