

FEEDBACK 1/14

- fork() will a modification to the data of a forked process trigger Copy-and-Write to duplicate all process data?, or only that which was modified...?
 - Only modified data saves memory and time
 - How can this be tested?
- Example: exec.c revisited example using wait()
 - Recompiled exec.c
 - Unable to reproduce bug (parent existing before child)
 - Old version was compiled with previous Linux kernel, suspect changes may have caused different behavior
 - Check out: waitpid() and capturing return code from wait()

January 16, 2019

TCSS422: Operating Systems [Winter 2019]

School of Engineering and Technology, University of Washington - Tacoma

FEEDBACK - 2

- Why exec, instead of just calling another function like in Java?
 - Fork(), exec(), and wait() allow fine grained control over processes
 - Goal: remotely invoke another executable program, from an existing C program
 - Can fork() a new processes to delegate to run an external program
 - Can redirect input, output, stderr of processes
- Exec.. is still a little complex
 - 6 versions: execl(), execlp(), execle(), execv(), execvp(), execvpe()
 - "p" duplicates the PATH of the shell
 - "e" allows the environment (variables) to be passed in as an array of pointers to NULL-terminated strings, list is NULL terminated
 - "v" provide cmd & args as array (vector)
 - "I" provide cmd & args as (<u>list)</u> of individual params sent to the function (see execl.c example)

January 16, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L4.3

exec()

- Supports running an external program
- 6 types: execl(), execlp(), execle(), execv(), execvp(), execvpe()
- execl(), execlp(), execle(): const char *arg (example: execl.c)

Provide cmd and args as individual params to the function Each arg is a pointer to a null-terminated string **ODD**: pass a variable number of args: (arg0, arg1, .. argn)

Execv(), execvp(), execvpe() (example: exec.c) Provide cmd and args as an Array of pointers to strings

Strings are null-terminated First argument is name of command being executed Fixed number of args passed in

January 16, 2019

TCSS422: Operating Systems [Winter 2019]

School of Engineering and Technology, University of Washington - Tacoma

14.4

FEEDBACK - 3

- What is the (too much / too little) control tradeoff for Operating System designs?
 - Who (what) has too much/little control?
 - What are consequences of too much control?
 - ... consequences of too little control?
- What is direct vs. limited direct execution with respect to operating systems?



- Direct code runs directly on HW, OS can not intervene, processes can "run away"...
 - Limited direct code runs on HW, but can not directly access protected resources (memory, devices), OS can intervene and *preempt* a process

January 16, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L4.5

FEEDBACK - 4

- How do we identify system processes?
- All kernel processes (threads) are children of process (pid==2)
- "ps" command to display processes owned by the kernel:

```
ps --ppid 2 -p 2 -o
uname,pid,ppid,pcpu,pmem,vsz,rssize,start,time,cmd |
more
```

"ps" command to display non-kernel processes:

```
ps --ppid 2 -p 2 -o
uname, pid, ppid, pcpu, pmem, vsz, rssize, start, time, cmd -
-deselect | more
```

What is significant regarding the output?

TCSS422: Operating Systems [Winter 2019] January 16, 2019

School of Engineering and Technology, University of Washington - Tacoma

FEEDBACK - 5

- Hard time following along. Overwhelming amount of information. Can't seem to find guideline/direction
 - AFTER CLASS :
 - (1) Review Slides
 - (2) Read chapters (in this case Ch. 5 & 6)
 - (3) While reviewing/reading, **make list** of confusing topics
 - (4) Formulate questions about these topics
 - THEN
 - (1) Visit professor after class, during office hours, or make an appointment --or--
 - (2) Email list of questions to professor

January 16, 2019

TCSS422: Operating Systems [Winter 2019] School of Engineering and Technology, University of Washington - Tacoma

L4.7

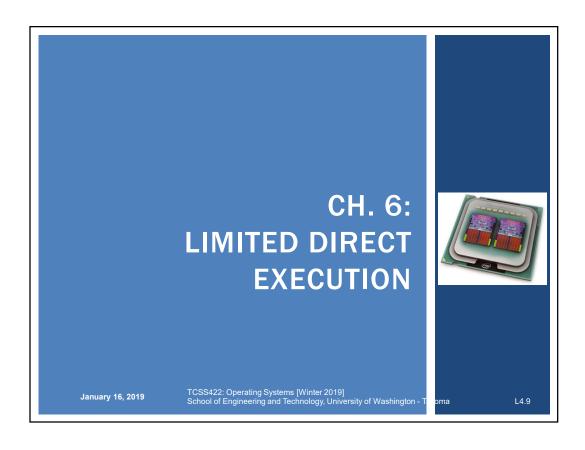
OBJECTIVES

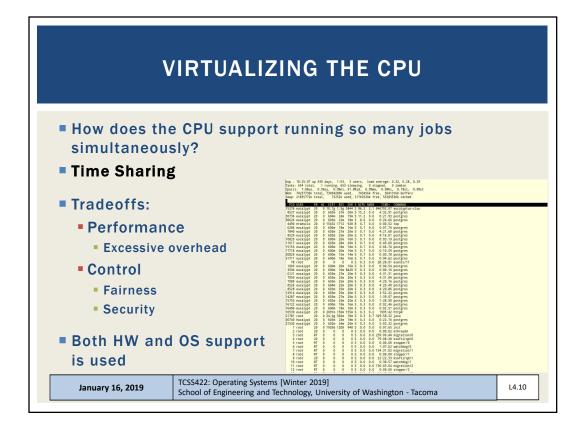
- Active Reading Quiz 1
- Assignment 0 / Linux Tutorial
- C Tutorial
- Chapter 6 Limited Direct Execution
- CPU Scheduling:
- Chapter 7 Scheduling Introduction
- Chapter 8 Multi-level Feedback Queue (MLFQ)

January 16, 2019

TCSS422: Operating Systems [Winter 2019]

School of Engineering and Technology, University of Washington - Tacoma





LIMITED DIRECT EXECUTION

- OS implements LDE to support time/resource sharing
- Limited direct execution means "only limited" processes can execute DIRECTLY on the CPU in trusted mode
- TRUSTED means the process is trusted, and it can do anything... (e.g. it is a system / kernel level process)
- Enabled by protected (safe) control transfer
- CPU supported context switch
- Provides data isolation

January 16, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L4.11

CPU MODES

- Utilize CPU Privilege Rings (Intel x86)
 - rings 0 (kernel), 1 (VM kernel), 2 (unused), 3 (user)

access <no access

User mode:

Application is running, but w/o direct I/O access

Kernel mode:

OS kernel is running performing restricted operations

January 16, 2019

TCSS422: Operating Systems [Winter 2019]

School of Engineering and Technology, University of Washington - Tacoma

CPU MODES

- <u>User mode: ring 3 untrusted</u>
 - Some instructions and registers are disabled by the CPU
 - Exception registers
 - HALT instruction
 - MMU instructions
 - OS memory access
 - I/O device access
- Kernel mode: ring 0 trusted
 - All instructions and registers enabled

January 16, 2019

TCSS422: Operating Systems [Winter 2019] School of Engineering and Technology, University of Washington - Tacoma

L4.13

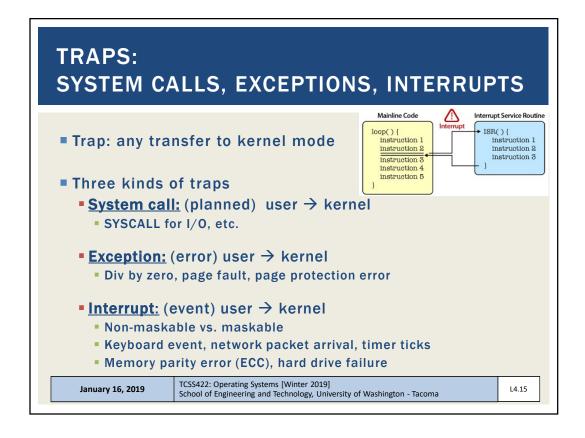
SYSTEM CALLS

- Implement restricted "OS" operations
- Kernel exposes key functions through an API:
 - Device I/O (e.g. file I/O)
 - Task swapping: context switching between processes
 - Memory management/allocation: malloc()
 - Creating/destroying processes

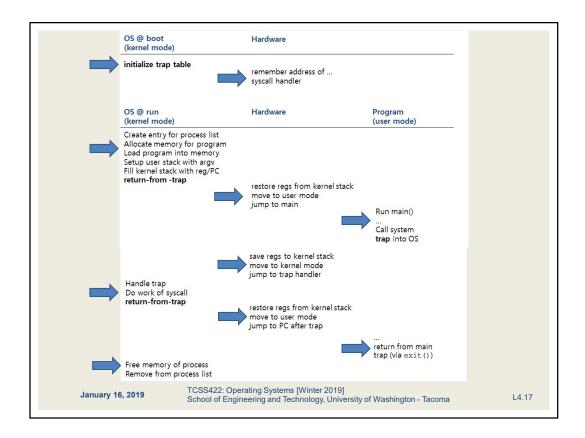
January 16, 2019

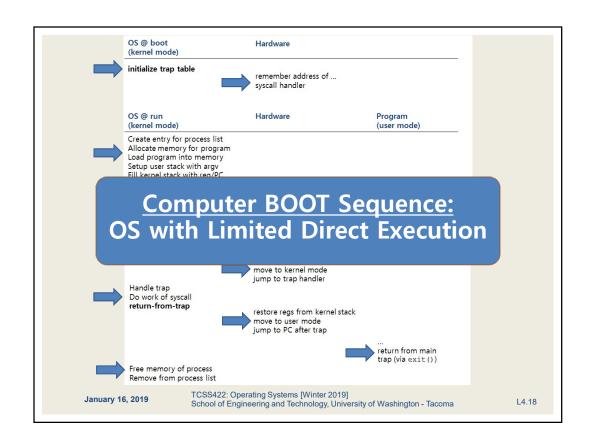
TCSS422: Operating Systems [Winter 2019]

School of Engineering and Technology, University of Washington - Tacoma



EXCEPTION TYPES					
Exception type	Synchronous vs. asynchronous	User request vs.	User maskable vs. nonmaskable	Within vs. between Instructions	Resume vs. terminate
I/O device request	Asynchronous	Coerced	Nonmaskable	Between	Resume
Invoke operating system	Synchronous	User request	Nonmaskable	Between	Resume
Tracing instruction execution	Synchronous	User request	User maskable	Between	Resume
Breakpoint	Synchronous	User request	User maskable	Between	Resume
Integer arithmetic overflow	Synchronous	Coerced	User maskable	Within	Resume
Floating-point arithmetic overflow or underflow	Synchronous	Coerced	User maskable	Within	Resume
Page fault	Synchronous	Coerced	Nonmaskable	Within	Resume
Misaligned memory accesses	Synchronous	Coerced	User maskable	Within	Resume
Memory protection violation	Synchronous	Coerced	Nonmaskable	Within	Resume
Using undefined instruction	Synchronous	Coerced	Nonmaskable	Within	Terminate
Hardware malfunction	Asynchronous	Coerced	Nonmaskable	Within	Terminate
Power fallure	Asynchronous	Coerced	Nonmaskable	Within	Terminate





MULTITASKING

- How/when should the OS regain control of the CPU to switch between processes?
- Cooperative multitasking (mostly pre 32-bit)
 - < Windows 95, Mac OSX</p>
 - Opportunistic: running programs must give up control
 - User programs must call a special yield system call
 - When performing I/O
 - Illegal operations
 - (POLLEV)

What problems could you for see with this approach?

January 16, 2019

TCSS422: Operating Systems [Winter 2019] School of Engineering and Technology, University of Washington - Tacoma

L4.19

MULTITASKING

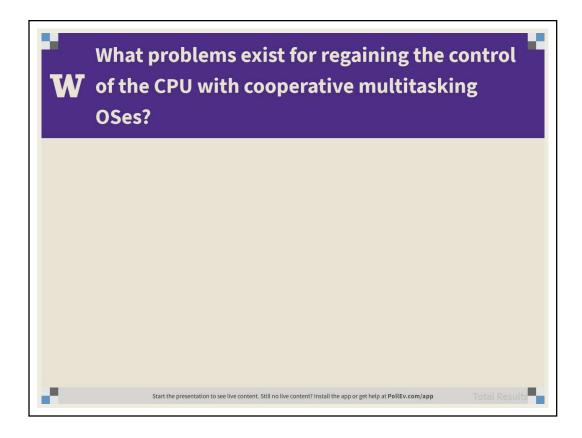
- How/when should the OS regain control of the CPU to switch between processes?
- Cooperative multitacking (mostly pro 32 hit)
 - A process gets stuck in an infinite loop.
 - → Reboot the machine
 - Wnen periorining i/ o
 - Illegal operations
 - (POLLEV)

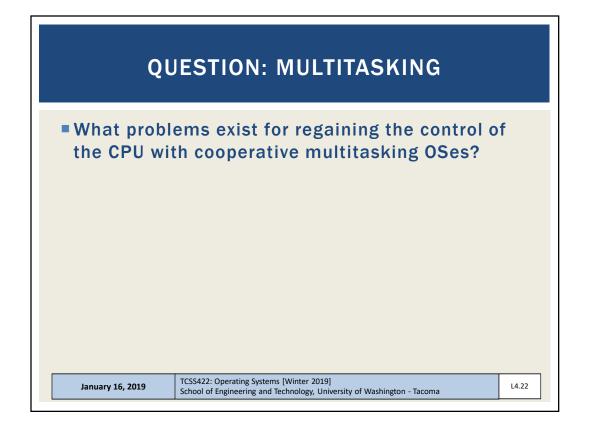
What problems could you for see with this approach?

January 16, 2019

TCSS422: Operating Systems [Winter 2019]

School of Engineering and Technology, University of Washington - Tacoma





MULTITASKING - 2

- Preemptive multitasking (32 & 64 bit OSes)
- >= Mac OSX, Windows 95+
- **■** Timer interrupt
 - Raised at some regular interval (in ms)
 - Interrupt handling
 - 1. Current program is halted
 - 2. Program states are saved
 - 3. OS Interrupt handler is run (kernel mode)
- (PollEV) What is a good interval for the timer interrupt?

January 16, 2019

TCSS422: Operating Systems [Winter 2019] School of Engineering and Technology, University of Washington - Tacoma

L4.23

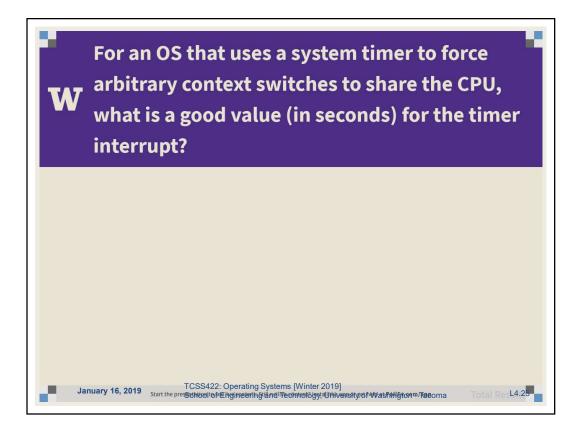
MULTITASKING - 2

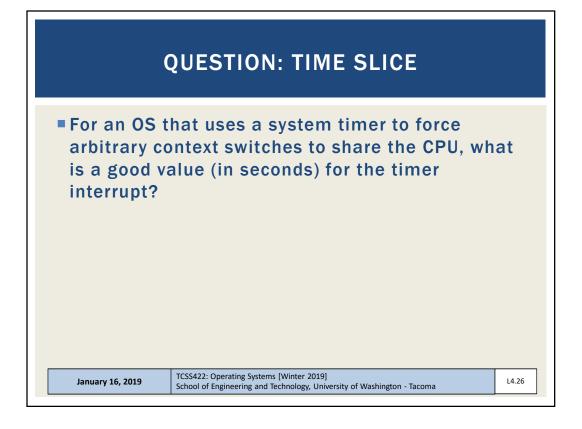
- Preemptive multitasking (32 & 64 bit OSes)
- >= Mac OSX, Windows 95+
- Timer
 - Rais A timer interrupt gives OS the ability to run again on a CPU.
 - Inter
 - 1. Current program is halted
 - 2. Program states are saved
 - 3. OS Interrupt handler is run (kernel mode)
- (PollEV) What is a good interval for the timer interrupt?

January 16, 2019

TCSS422: Operating Systems [Winter 2019]

School of Engineering and Technology, University of Washington - Tacoma





CONTEXT SWITCH

- Preemptive multitasking initiates "trap" into the OS code to determine:
- Whether to continue running the current process, or switch to a different one.
- If the decision is made to switch, the OS performs a context switch swapping out the current process for a new one.

January 16, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L4.27

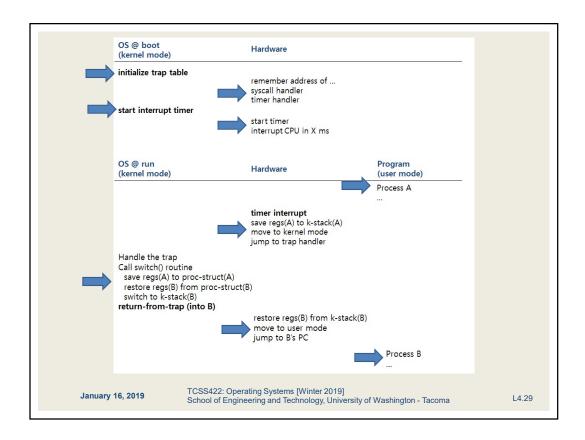
CONTEXT SWITCH - 2

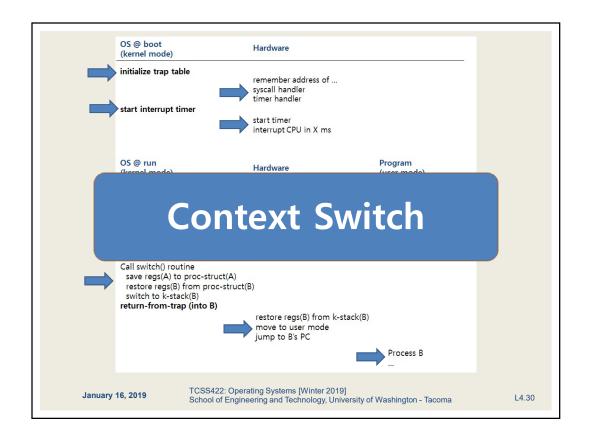
- 1. Save register values of the current process to its kernel stack
 - General purpose registers
 - PC: program counter (instruction pointer)
 - kernel stack pointer
- 2. Restore soon-to-be-executing process from its kernel stack
- 3. Switch to the kernel stack for the soon-to-be-executing process

January 16, 2019

TCSS422: Operating Systems [Winter 2019]

School of Engineering and Technology, University of Washington - Tacoma





INTERRUPTED INTERRUPTS

- What happens if during an interrupt (trap to kernel mode), another interrupt occurs?
- Linux
 - < 2.6 kernel: non-preemptive kernel</p>
 - >= 2.6 kernel: preemptive kernel

January 16, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L4.31

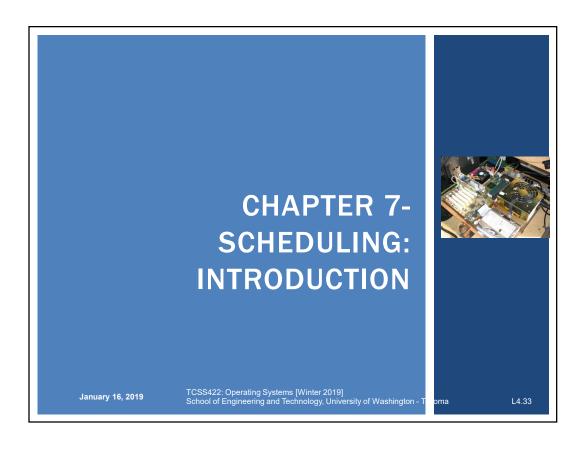
PREEMPTIVE KERNEL

- Use "locks" as markers of regions of nonpreemptibility (non-maskable interrupt)
- Preemption counter (preempt count)
 - begins at zero
 - increments for each lock acquired (not safe to preempt)
 - decrements when locks are released
- Interrupt can be interrupted when preempt count=0
 - It is safe to preempt (maskable interrupt)
 - the interrupt is more important

January 16, 2019

TCSS422: Operating Systems [Winter 2019]

School of Engineering and Technology, University of Washington - Tacoma



SCHEDULING METRICS

- Metrics: A standard measure to quantify to what degree a system possesses some property. Metrics provide <u>repeatable</u> techniques to quantify and compare systems.
- Measurements are the numbers derived from the application of metrics
- Scheduling Metric #1: Turnaround time
- The time at which the job completes minus the time at which the job arrived in the system

 $T_{turnaround} = T_{completion} - T_{arrival}$

How is turnaround time different than execution time?

January 16, 2019

TCSS422: Operating Systems [Winter 2019]

School of Engineering and Technology, University of Washington - Tacoma

SCHEDULING METRICS - 2

- Scheduling Metric #2: Fairness
 - Jain's fairness index
 - Quantifies if jobs receive a fair share of system resources

$$\mathcal{J}(x_1,x_2,\ldots,x_n) = rac{(\sum_{i=1}^n x_i)^2}{n \cdot \sum_{i=1}^n x_i^2}$$

- n processes
- x_i is time share of each process
- worst case = 1/n
- best case = 1
- Consider n=3, worst case = .333, best case=1
- With n=3 and x_1 =.2, x_2 =.7, x_3 =.1, fairness=.62
- With n=3 and x_1 =.33, x_2 =.33, x_3 =.33, fairness=1

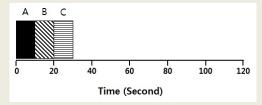
January 16, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L4.35

SCHEDULERS

- FIFO: first in, first out
 - Very simple, easy to implement
- Consider
 - 3 x 10sec jobs, arrival: A B C, duration 10 sec each

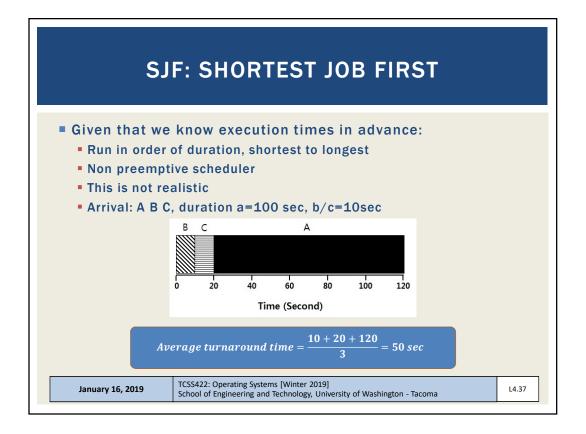


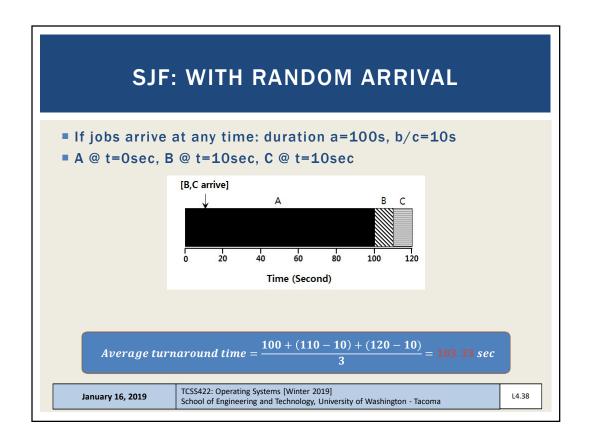
Average turnaround time =

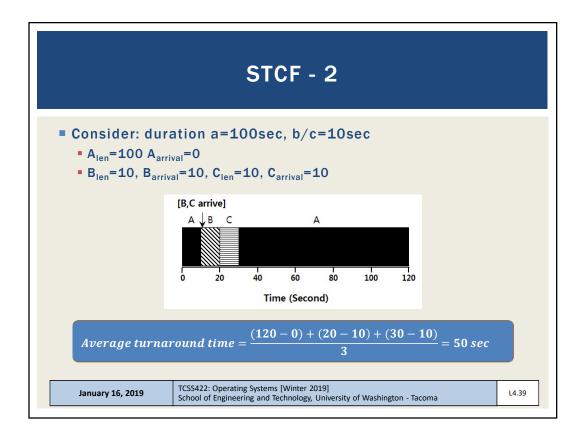
January 16, 2019

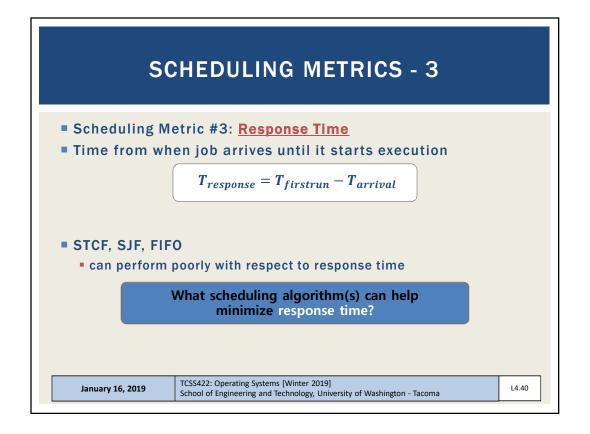
TCSS422: Operating Systems [Winter 2019]

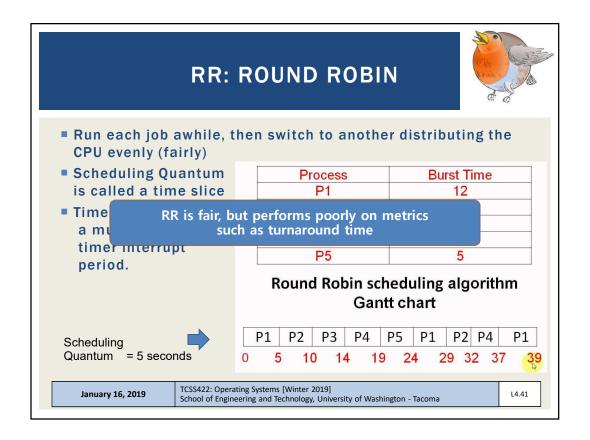
School of Engineering and Technology, University of Washington - Tacoma

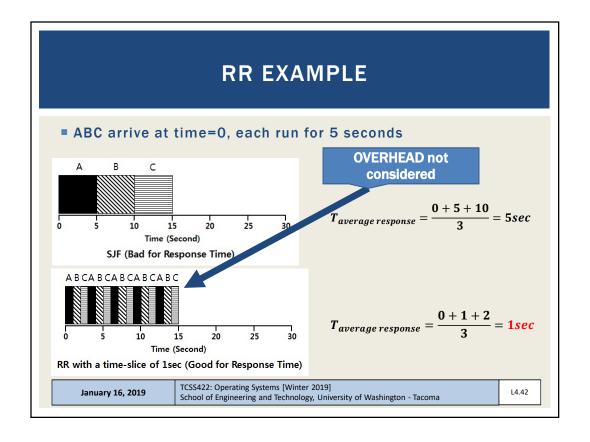


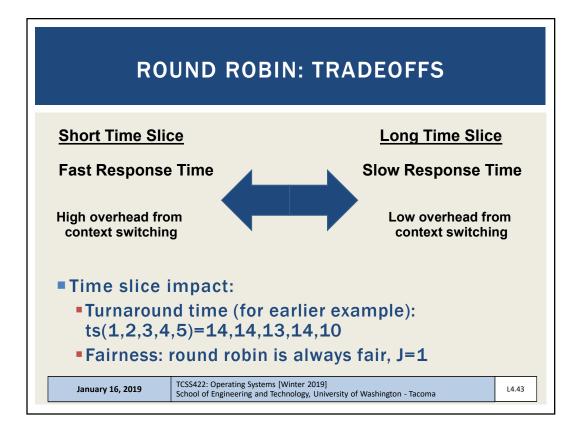


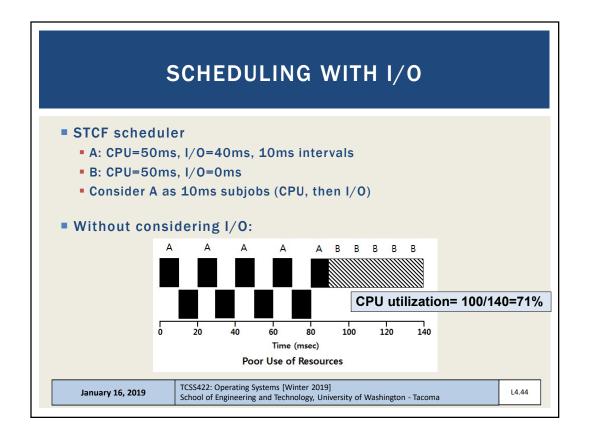


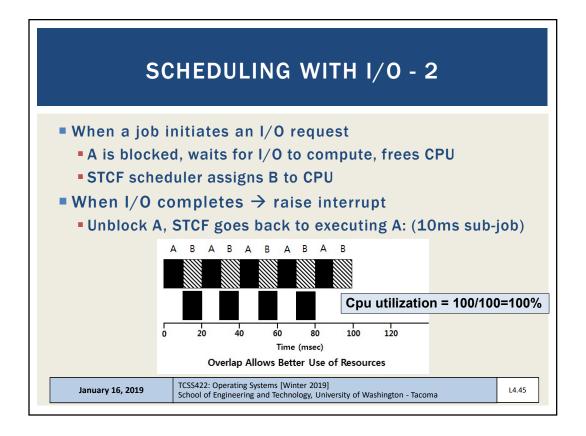


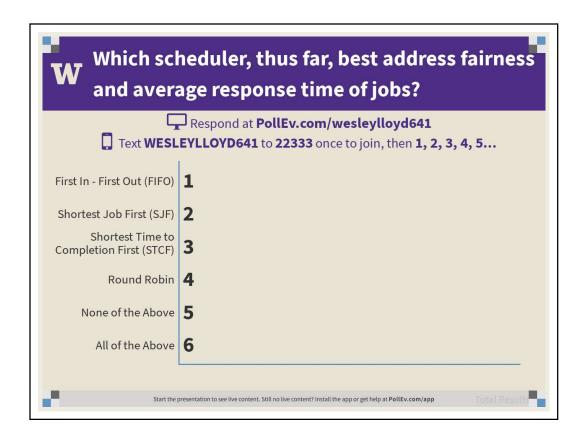


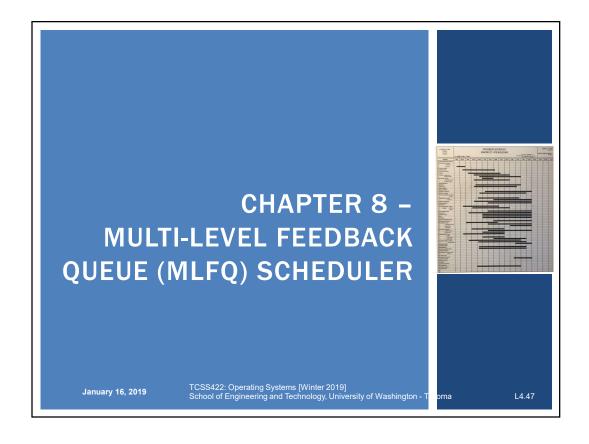


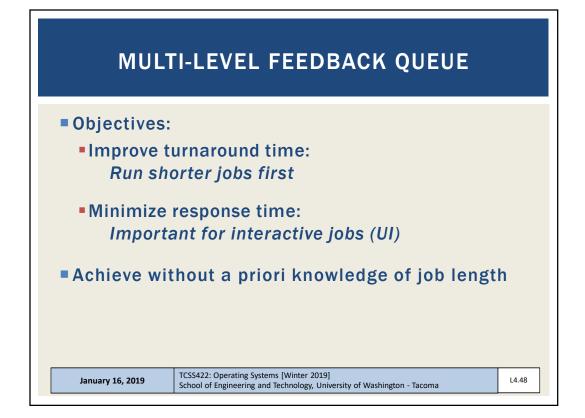


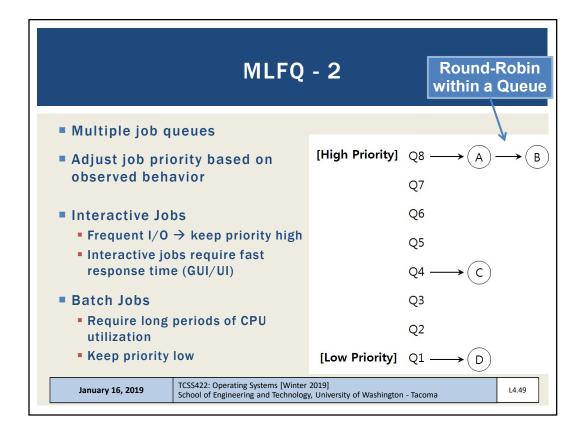




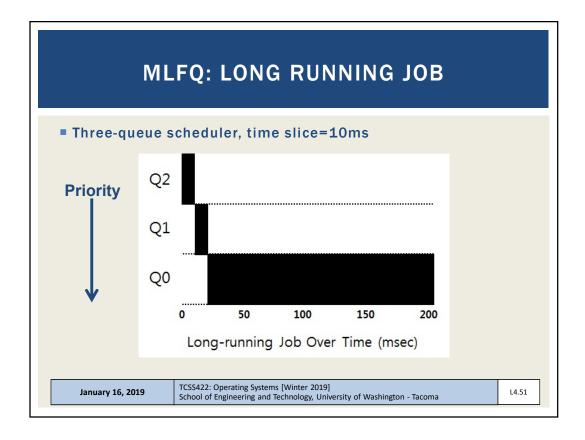


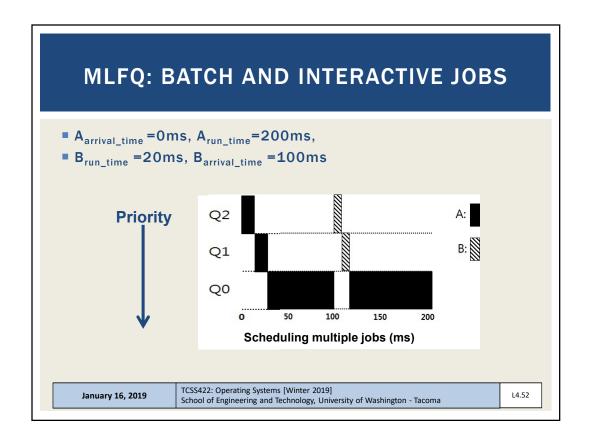


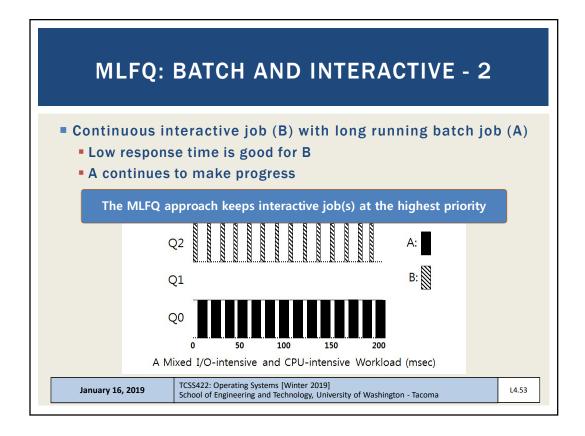


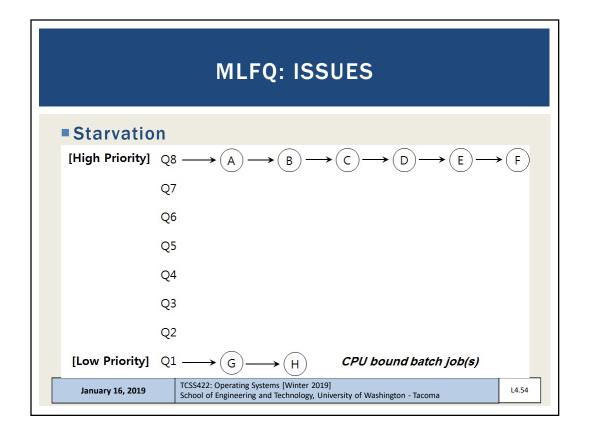


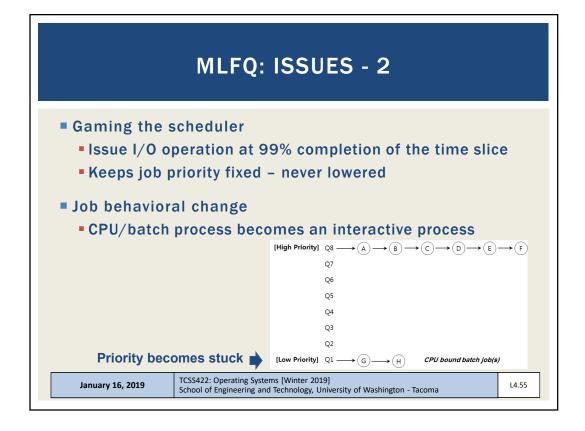


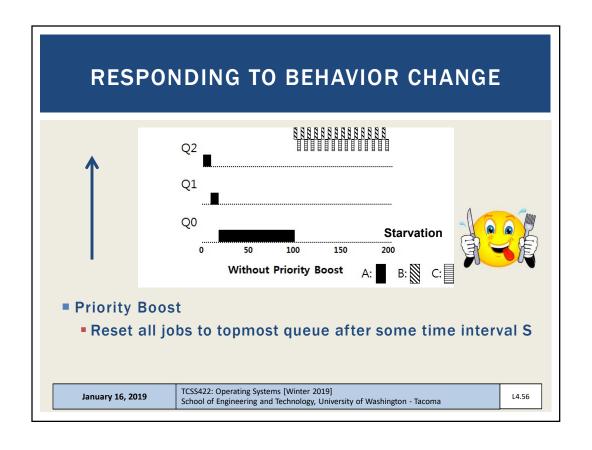


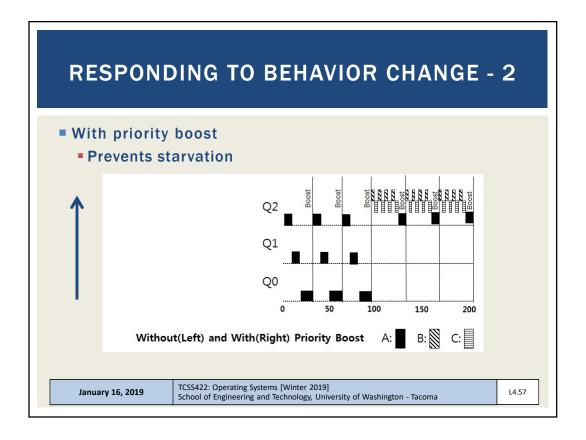


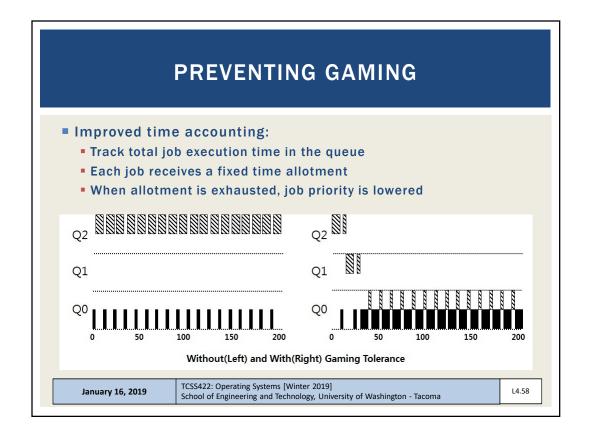


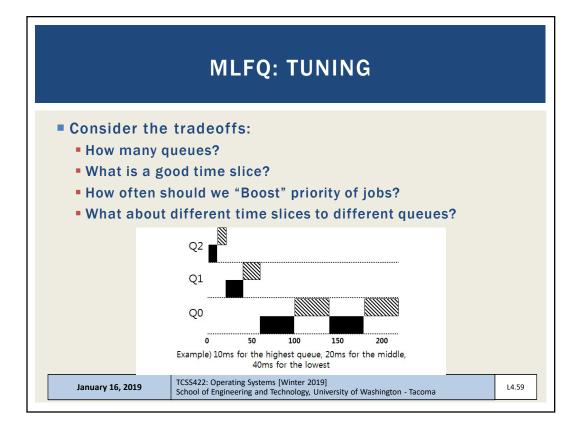












PRACTICAL EXAMPLE ■ Oracle Solaris MLFQ implementation ■ 60 Queues → w/ slowly increasing time slice (high to low priority) ■ Provides sys admins with set of editable table(s) ■ Supports adjusting time slices, boost intervals, priority changes, etc. ■ Advice ■ Provide OS with hints about the process ■ Nice command → Linux | January 16, 2019 | TCSS422: Operating Systems (Winter 2019) | School of Engineering and Technology, University of Washington - Tacoma | L4.60 | L4.60 |

MLFQ RULE SUMMARY

- The refined set of MLFQ rules:
- Rule 1: If Priority(A) > Priority(B), A runs (B doesn't).
- Rule 2: If Priority(A) = Priority(B), A & B run in RR.
- Rule 3: When a job enters the system, it is placed at the highest priority.
- Rule 4: Once a job uses up its time allotment at a given level (regardless of how many times it has given up the CPU), its priority is reduced(i.e., it moves down on queue).
- Rule 5: After some time period S, move all the jobs in the system to the topmost queue.

January 16, 2019 TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L4.61

Jackson deploys a 3-level MLFQ scheduler. The time slice is 1 for high priority jobs, 2 for medium priority, and 4 for low priority. This MLFQ scheduler performs a Priority Boost every 6 timer units. When the priority boost fires, the current job is preempted, and the next scheduled job is run in round-robin order.

 Job
 Arrival Time
 Job Length

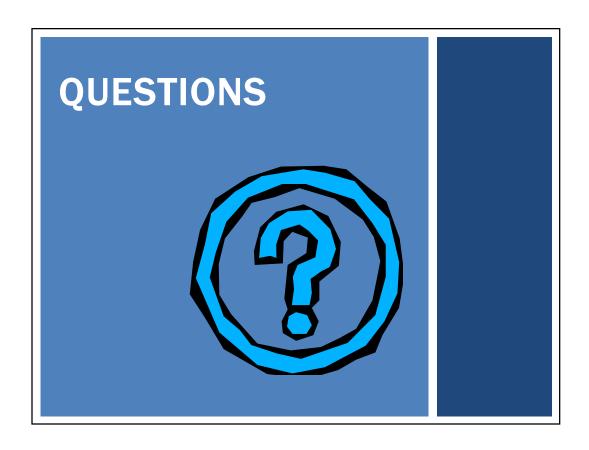
 A
 T=0
 4

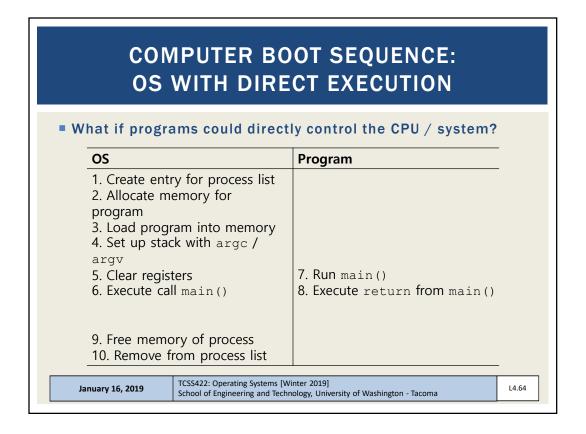
 B
 T=0
 16

 C
 T=0
 8

(11 points) Show a scheduling graph for the MLFQ scheduler for the jobs above. Draw vertical lines for key events and be sure to label the X-axis times as in the example. Please draw clearly. An unreadable graph will loose points.

HIGH |
MED |
LOW |





COMPUTER BOOT SEQUENCE: OS WITH DIRECT EXECUTION

What if programs could directly control the CPU / system?

OS	Program			
 Create entry for process list Allocate memory for 				
Without <i>limits</i> on running programs, the OS wouldn't be in control of anything and would " <u>just be a library</u> "				
5. Clear registers 6. Execute call main()	7. Run main() 8. Execute return from main()			
9. Free memory of process10. Remove from process list				

January 16, 2019

TCSS422: Operating Systems [Winter 2019] School of Engineering and Technology, University of Washington - Tacoma

DIRECT EXECUTION - 2

■ With direct execution:

How does the OS stop a program from running, and switch to another to support time sharing?

How do programs share disks and perform I/O if they are given direct control? Do they know about each other?

With direct execution, how can dynamic memory structures such as linked lists grow over time?

January 16, 2019

TCSS422: Operating Systems [Winter 2019]

School of Engineering and Technology, University of Washington - Tacoma

L4.66

CONTROL TRADEOFF Too little control: No security No time sharing Too much control: Too much OS overhead Poor performance for compute & I/O Complex APIs (system calls), difficult to use

