


TCSS 422: OPERATING SYSTEMS

Limited Direct Execution II, Intro to Schedulers



Wes J. Lloyd
School of Engineering and Technology
University of Washington - Tacoma

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

OBJECTIVES – 4/13

- **Questions from 4/8**
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
 - Cooperative multi-tasking
 - Context switching and preemptive multi-tasking
- Chapter 7: Scheduling Introduction
 - Scheduling metrics
 - Turnaround time, Jain's Fairness Index, Response time
 - FIFO, SJF, STCF, RR schedulers
- Chapter 8: Multi-level Feedback Queue

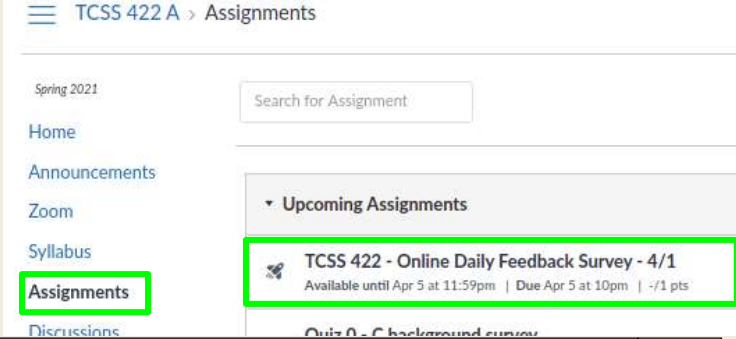
April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.2

ONLINE DAILY FEEDBACK SURVEY

- Daily Feedback Quiz in Canvas – Available After Each Class
- Extra credit available for completing surveys **ON TIME**
- Tuesday surveys: due by ~ Wed @ 11:59p
- Thursday surveys: due ~ Mon @ 11:59p



April 13, 2021	TCSS422: Computer Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.3
----------------	--	------

TCSS 422 - Online Daily Feedback Survey - 4/1

Quiz Instructions

Question 1 0.5 pts

On a scale of 1 to 10, please classify your perspective on material covered in today's class:

1	2	3	4	5	6	7	8	9	10
Mostly Review To Me				Equal New and Review					Mostly New to Me

Question 2 0.5 pts

Please rate the pace of today's class:

1	2	3	4	5	6	7	8	9	10
Slow				Just Right					Fast

April 13, 2021	TCSS422: Computer Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.4
----------------	--	------

MATERIAL / PACE

- Please classify your perspective on material covered in today's class (56 respondents):
 - 1-mostly review, 5-equal new/review, 10-mostly new
 - **Average – 6.91 (no change - previous 6.91)**
- Please rate the pace of today's class:
 - 1-slow, 5-just right, 10-fast
 - **Average – 5.65 (↓ - previous 5.67)**

April 13, 2021

TCSS422: Computer Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.5

FEEDBACK

- **Could you fork another child process from a child process to have a grandparent/grandchild process?**
 - YES
- **Can the exec() function be thought of as opening another thread to a different program?**
 - NO, exec runs another program in the existing process
 - Control is transferred to another executable and is not returned
- **Does exec transfer control back to the main program that called external program?**
 - No, the PID is transferred to the new executable and does not come back. If you want to preserve the original program the idea is to fork and run exec with a child process

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.6

FEEDBACK - 2

- What are the advantages of using lower-level APIs such as `open()` compared to the specialized versions with additional features like `fopen()`? Is this similar to the control tradeoff? Introducing unnecessary overhead and the like?
 - `fopen()` and other functions like it are provided largely out of convenience for developers
 - Specialized wrappers such as `fopen()` abstract additional functionality to make it more easily accessible for programmers
- With the use of standard out and standard error when EXEC with file redirection, I'm still not sure about the steps from L4.30

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.7

EXEC WITH FILE REDIRECTION (OUTPUT)

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <string.h>
#include <fcntl.h>
#include <sys/wait.h>

int
main(int argc, char *argv[]){
    int rc = fork();
    if (rc < 0) { // fork failed; exit
        fprintf(stderr, "fork failed\n");
        exit(1);
    } else if (rc == 0) { // child: redirect standard output to a file
        close(STDOUT_FILENO);
        open("./p4.output", O_CREAT|O_WRONLY|O_TRUNC, S_IRWXU);
        ...
    }
}
```

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.8

FILE MODE BITS

```
S_IRWXU
read, write, execute/search by owner
S_IRUSR
read permission, owner
S_IWUSR
write permission, owner
S_IXUSR
execute/search permission, owner
S_IRWXG
read, write, execute/search by group
S_IRGRP
read permission, group
S_IWGRP
write permission, group
S_IXGRP
execute/search permission, group
S_IRWXO
read, write, execute/search by others
S_IROTH
read permission, others
S_IWOTH
write permission, others
```

April 13, 2021	TCCS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.9
----------------	---	------

EXEC W/ FILE REDIRECTION (OUTPUT) - 2

```
...
// now exec "wc"...
char *myargs[3];
myargs[0] = strdup("wc");           // program: "wc" (word count)
myargs[1] = strdup("p4.c");         // argument: file to count
myargs[2] = NULL;                   // marks end of array
execvp(myargs[0], myargs);         // runs word count
} else {                             // parent goes down this path (main)
    int wc = wait(NULL);
}
return 0;
}
```

```
prompt> ./p4
prompt> cat p4.output
32 109 846 p4.c
prompt>
```

April 13, 2021	TCCS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.10
----------------	---	-------

FEEDBACK - 3

- **Do system calls and traps execute with limited direct execution?**
 - These are privileged operations that are executed in the kernel, with direct execution.
- **In layman operating systems like Windows and iOS, what processes are already trusted? How difficult is it to make other processes trusted for the purposes of LDE?'**
 - User processes by default are not trusted
 - They run with Limited Direct Execution
 - Only operating system kernel code is trusted
 - In Linux this can be the kernel itself or kernel modules

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.11

FEEDBACK - 4

- **What is kernel? Is it the same thing as OS?**
 - Yes, the kernel is the binary executable that runs with the computer boots
 - In Linux this is the `"/boot/vmlinuz"` executable
 - Identify file properties of your kernel with the command:
`sudo file /boot/vmlinuz-$(uname -r)`
- **Can you please explain the differences between the kernel, hardware, and the program?**
 - The Linux kernel is the binary file `/boot/vmlinuz...` that is largely written in C & Assembly that implements the Linux OS
 - The hardware is the laptop, desktop, or virtual machine
 - The program is a user program that you write such as the examples we've reviewed in class (i.e. `fork.c`)

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.12

FEEDBACK - 5

- **What are trap tables and trap handlers?**
 - **TRAP TABLE:**
The x86 processor uses a table known as the interrupt descriptor table (IDT) to determine how to transfer control when a trap occurs. The x86 allows up to 256 different interrupt or exception entry points into the kernel, each with a different interrupt vector.
 - **TRAP HANDLERS:**
Trap handlers are OS kernel functions that are pointed to by the trap table. These are “event handlers” that respond to various traps.
- **What is the difference between system APIs and system calls? Are Fork(), wait(), exec() both system APIs and system calls?**
 - A system call is the action of actually calling the system API
 - A system API is the function as defined in the OS

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.13

OBJECTIVES – 4/13

- Questions from 4/8
- **Assignment 0**
- **C Tutorial - Pointers, Strings, Exec in C**
- **Chapter 6: Limited Direct Execution**
 - Cooperative multi-tasking
 - Context switching and preemptive multi-tasking
- **Chapter 7: Scheduling Introduction**
 - Scheduling metrics
 - Turnaround time, Jain’s Fairness Index, Response time
 - FIFO, SJF, STCF, RR schedulers
- **Chapter 8: Multi-level Feedback Queue**

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.14

OBJECTIVES – 4/13

- Questions from 4/8
- Assignment 0
- **C Tutorial - Pointers, Strings, Exec in C**
- Chapter 6: Limited Direct Execution
 - Cooperative multi-tasking
 - Context switching and preemptive multi-tasking
- Chapter 7: Scheduling Introduction
 - Scheduling metrics
 - Turnaround time, Jain's Fairness Index, Response time
 - FIFO, SJF, STCF, RR schedulers
- Chapter 8: Multi-level Feedback Queue

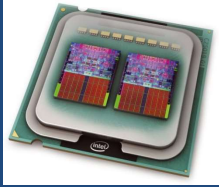
April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.15
----------------	---	-------

OBJECTIVES – 4/13

- Questions from 4/8
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
 - **Cooperative multi-tasking**
 - Context switching and preemptive multi-tasking
- Chapter 7: Scheduling Introduction
 - Scheduling metrics
 - Turnaround time, Jain's Fairness Index, Response time
 - FIFO, SJF, STCF, RR schedulers
- Chapter 8: Multi-level Feedback Queue

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.16
----------------	---	-------

CH. 6: LIMITED DIRECT EXECUTION



April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L3.17

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
 - 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?

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.18
----------------	---	-------

MULTITASKING

- How/when should the OS regain control of the CPU to switch between processes?
- Cooperative multitasking (mostly pre-32 bit)
 - < Y
 - Op
 - When performing I/O
 - Illegal operations
- (POLLEV)
What problems could you for see with this approach?

A process gets stuck in an infinite loop.
→ **Reboot the machine**

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.19
----------------	---	-------

What problems exist for regaining the control of the CPU with cooperative multitasking OSes?

W

of the CPU with cooperative multitasking OSes?

Start the presentation to see live content. Still no live content? Install the app or get help at PollEv.com/app

Total Results

QUESTION: MULTITASKING

- What problems exist for regaining the control of the CPU with cooperative multitasking OSes?

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.21

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?

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.22

MULTITASKING - 2

- Preemptive multitasking (32 & 64 bit OSes)
- >= Mac OSX, Windows 95+

- Timer
 - Rais
 - 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?

A **timer interrupt** gives OS the ability to run again on a CPU.

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.23
----------------	---	-------

W

For an OS that uses a system timer to force arbitrary context switches to share the CPU, what is a good value (in seconds) for the timer interrupt?

April 13, 2021 Start the presentation TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma Total Res: L5.24

QUESTION: TIME SLICE

- For an OS that uses a system timer to force arbitrary context switches to share the CPU, what is a good value (in seconds) for the timer interrupt?

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.25

OBJECTIVES – 4/13

- Questions from 4/8
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
 - Cooperative multi-tasking
 - **Context switching and preemptive multi-tasking**
- Chapter 7: Scheduling Introduction
 - Scheduling metrics
 - Turnaround time, Jain's Fairness Index, Response time
 - FIFO, SJF, STCF, RR schedulers
- Chapter 8: Multi-level Feedback Queue

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.26

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.

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.27

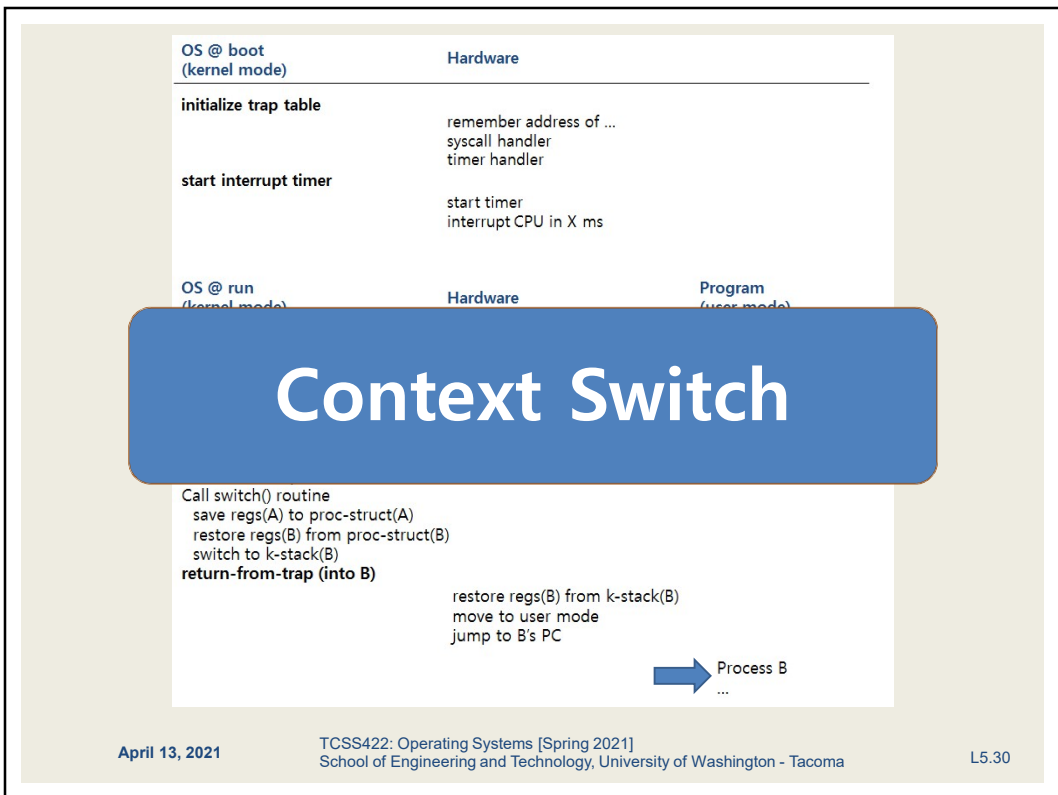
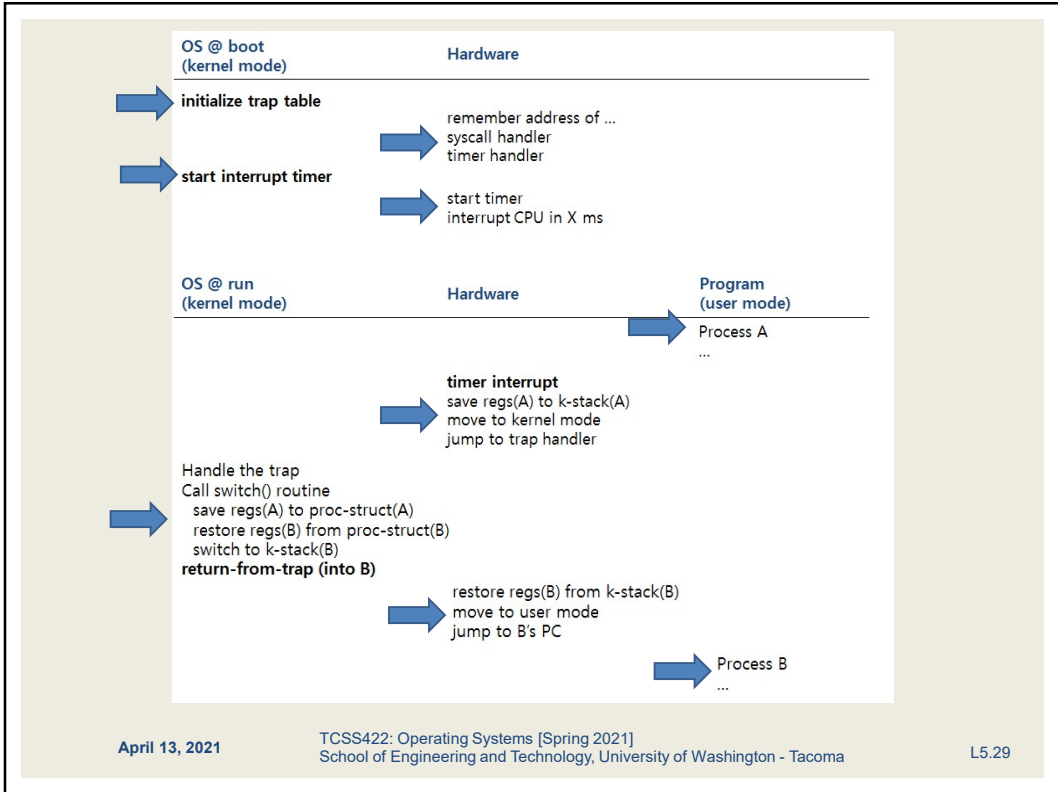
CONTEXT SWITCH - 2

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

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.28



INTERRUPTED INTERRUPTS

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

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.31

PREEMPTIVE KERNEL

- Use “locks” as markers of regions of non-preemptibility (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

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.32

**WE WILL RETURN AT
5:10PM**




April 13, 2021 TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma L5.33

OBJECTIVES – 4/13

- Questions from 4/8
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
 - Cooperative multi-tasking
 - Context switching and preemptive multi-tasking
- **Chapter 7: Scheduling Introduction**
 - Scheduling metrics
 - Turnaround time, Jain's Fairness Index, Response time
 - FIFO, SJF, STCF, RR schedulers
- Chapter 8: Multi-level Feedback Queue

April 13, 2021 TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma L5.34

CHAPTER 7- SCHEDULING: INTRODUCTION



April 13, 2021 TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma L5.35

SCHEDULING METRICS

- **Metrics:** A standard measure to quantify to what degree a system possesses some property. Metrics provide repeatable 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?

April 13, 2021 TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma L5.36

SCHEDULING METRICS - 2

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

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

- 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

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.37

OBJECTIVES - 4/13

- Questions from 4/8
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
 - Cooperative multi-tasking
 - Context switching and preemptive multi-tasking
- Chapter 7: Scheduling Introduction
 - Scheduling metrics
 - Turnaround time, Jain's Fairness Index, Response time
 - **FIFO** SJF, STCF, RR schedulers
- Chapter 8: Multi-level Feedback Queue

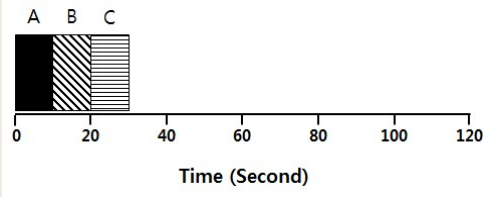
April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.38

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 = $\frac{10 + 20 + 30}{3} = 20 \text{ sec}$

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.39
----------------	---	-------

OBJECTIVES – 4/13

- Questions from 4/8
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
 - Cooperative multi-tasking
 - Context switching and preemptive multi-tasking
- Chapter 7: Scheduling Introduction
 - Scheduling metrics
 - Turnaround time, Jain's Fairness Index, Response time
 - FIFO, **SJF**, STCF, RR schedulers
- Chapter 8: Multi-level Feedback Queue

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.40
----------------	---	-------

SJF: SHORTEST JOB FIRST

- Given that we know execution times in advance:
 - Run in order of duration, shortest to longest
 - Non preemptive scheduler
 - This is not realistic
 - Arrival: A B C, duration a=100 sec, b/c=10sec

$$\text{Average turnaround time} = \frac{10 + 20 + 120}{3} = 50 \text{ sec}$$

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.41
----------------	---	-------

SJF: WITH RANDOM ARRIVAL

- If jobs arrive at any time: duration a=100s, b/c=10s
- A @ t=0sec, B @ t=10sec, C @ t=10sec

$$\text{Average turnaround time} = \frac{100 + (110 - 10) + (120 - 10)}{3} = 103.33 \text{ sec}$$

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.42
----------------	---	-------

OBJECTIVES – 4/13

- Questions from 4/8
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
 - Cooperative multi-tasking
 - Context switching and preemptive multi-tasking
- Chapter 7: Scheduling Introduction
 - Scheduling metrics
 - Turnaround time, Jain's Fairness Index, Response time
 - FIFO, SJF, **STCF**, RR schedulers
- Chapter 8: Multi-level Feedback Queue

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.43
----------------	---	-------

SCTF: SHORTEST TIME TO COMPLETION FIRST

- Consider: duration $a=100\text{sec}$, $b/c=10\text{sec}$
 - $A_{\text{len}}=100$ $A_{\text{arrival}}=0$
 - $B_{\text{len}}=10$, $B_{\text{arrival}}=10$, $C_{\text{len}}=10$, $C_{\text{arrival}}=10$

Time (Second)

Average turnaround time = $\frac{(120 - 0) + (20 - 10) + (30 - 10)}{3} = 50 \text{ sec}$

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.44
----------------	---	-------

OBJECTIVES – 4/13

- Questions from 4/8
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
 - Cooperative multi-tasking
 - Context switching and preemptive multi-tasking
- Chapter 7: Scheduling Introduction
 - Scheduling metrics
 - Turnaround time, Jain's Fairness Index, **Response time**
 - FIFO, SJF, STCF, RR schedulers
- Chapter 8: Multi-level Feedback Queue

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.45

SCHEDULING METRICS - 3

- Scheduling Metric #3: **Response Time**
- Time from when job arrives until it starts execution

$$T_{response} = T_{firstrun} - T_{arrival}$$

- STCF, SJF, FIFO
 - can perform poorly with respect to response time

What scheduling algorithm(s) can help minimize response time?

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.46

OBJECTIVES – 4/13

- Questions from 4/8
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
 - Cooperative multi-tasking
 - Context switching and preemptive multi-tasking
- Chapter 7: Scheduling Introduction
 - Scheduling metrics
 - Turnaround time, Jain's Fairness Index, Response time
 - FIFO, SJF, STCF, **RR** schedulers
- Chapter 8: Multi-level Feedback Queue

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.47
----------------	---	-------

RR: ROUND ROBIN

- Run each job awhile, then switch to another distributing the CPU evenly (fairly)
- Scheduling Quantum is called a time slice
- Time a multiple time period.

Process	Burst Time
P1	12

RR is fair, but performs poorly on metrics such as turnaround time

Round Robin scheduling algorithm Gantt chart

Scheduling Quantum = 5 seconds

P1	P2	P3	P4	P5	P1	P2	P4	P1	
0	5	10	14	19	24	29	32	37	39

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.48
----------------	---	-------

RR EXAMPLE

- ABC arrive at time=0, each run for 5 seconds

SJF (Bad for Response Time)

OVERHEAD not considered

$$T_{average\ response} = \frac{0 + 5 + 10}{3} = 5sec$$

RR with a time-slice of 1sec (Good for Response Time)

$$T_{average\ response} = \frac{0 + 1 + 2}{3} = 1sec$$

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.49
----------------	---	-------

ROUND ROBIN: TRADEOFFS

Short Time Slice

Fast Response Time

High overhead from context switching

Long Time Slice

Slow Response Time

Low overhead from context switching

- **Time slice impact:**
 - Turnaround time (for earlier example):
 $ts(1,2,3,4,5) = 14, 14, 13, 14, 10$
 - Fairness: round robin is always fair, $J=1$

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.50
----------------	---	-------

SCHEDULING WITH I/O

- **STCF scheduler**
 - A: CPU=50ms, I/O=40ms, 10ms intervals
 - B: CPU=50ms, I/O=0ms
 - Consider A as 10ms subjobs (CPU, then I/O)
- **Without considering I/O:**

CPU utilization = 100/140 = 71%

Time (msec)

Poor Use of Resources

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.51
----------------	---	-------

SCHEDULING WITH I/O - 2

- **When a job initiates an I/O request**
 - A is blocked, waits for I/O to compute, frees CPU
 - STCF scheduler assigns B to CPU
- **When I/O completes → raise interrupt**
 - Unblock A, STCF goes back to executing A: (10ms sub-job)

Cpu utilization = 100/100 = 100%

Time (msec)

Overlap Allows Better Use of Resources

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.52
----------------	---	-------

W Which scheduler, thus far, best address fairness and average response time of jobs?

Respond at PollEv.com/wesleylloyd641

Text **WESLEYLLOYD641** to **22333** once to join, then **1, 2, 3, 4, 5...**

First In - First Out (FIFO)	1
Shortest Job First (SJF)	2
Shortest Time to Completion First (STCF)	3
Round Robin	4
None of the Above	5
All of the Above	6

Start the presentation to see live content. Still no live content? Install the app or get help at PollEv.com/app Total Results

QUESTION: SCHEDULING FAIRNESS


- Which scheduler, this far, best addresses fairness and average response time of jobs?
- First In – First Out (FIFO)
- Shortest Job First (SJF)
- Shortest Time to Completion First (STCF)
- Round Robin (RR)
- None of the Above
- All of the Above

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.54
----------------	---	-------

SCHEDULING METRICS

- Consider Three jobs (A, B, C) that require:
 $time_A=400ms$, $time_B=100ms$, and $time_C=200ms$
- All jobs arrive at $time=0$ in the sequence of A B C.
- Draw a scheduling graph to help compute the **average response time (ART)** and **average turnaround time (ATT)** scheduling metrics for the FIFO scheduler.

Example:



April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.55

When poll is active, respond at [PollEv.com/wesleylloyd641](https://poll-ev.com/wesleylloyd641)

Text **WESLEYLLOYD641** to **22333** once to join

What is the Average Response Time of the FIFO scheduler?

Start the presentation to see live content. Still no live content? Install the app or get help at [PollEv.com/app](https://poll-ev.com/app)

When poll is active, respond at PollEv.com/wesleylloyd641
Text **WESLEYLLOYD641** to **22333** once to join

What is the Average Turnaround Time of the FIFO scheduler?

Start the presentation to see live content. Still no live content? Install the app or get help at PollEv.com/app

SCHEDULING METRICS

- Consider Three jobs (A, B, C) that require:
 $time_A=400ms$, $time_B=100ms$, and $time_C=200ms$
- All jobs arrive at $time=0$ in the sequence of A B C.
- Draw a scheduling graph to help compute the **average response time (ART)** and **average turnaround time (ATT)** scheduling metrics for the SJF scheduler.

Example:

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.58
----------------	---	-------

When poll is active, respond at PollEv.com/wesleylloyd641
Text **WESLEYLLOYD641** to **22333** once to join

What is the Average Response Time of the Shortest Job First Scheduler?

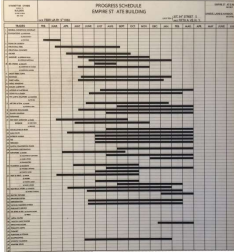
Start the presentation to see live content. Still no live content? Install the app or get help at PollEv.com/app

When poll is active, respond at PollEv.com/wesleylloyd641
Text **WESLEYLLOYD641** to **22333** once to join

What is the Average Turnaround Time of the Shortest Job First Scheduler?

Start the presentation to see live content. Still no live content? Install the app or get help at PollEv.com/app

**CHAPTER 8 –
MULTI-LEVEL FEEDBACK
QUEUE (MLFQ) SCHEDULER**



April 13, 2021 TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma L5.61

OBJECTIVES

- **Chapter 8: Multi-level Feedback Queue**
 - MLFQ Scheduler
 - Job Starvation
 - Gaming the Scheduler
 - Examples

April 13, 2021 TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma L5.62

MULTI-LEVEL FEEDBACK QUEUE

- Objectives:
 - Improve turnaround time:
Run shorter jobs first
 - Minimize response time:
Important for interactive jobs (UI)
- Achieve without a priori knowledge of job length

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.63
----------------	---	-------

MLFQ - 2

Round-Robin
within a Queue

- Multiple job queues
- Adjust job priority based on observed behavior
- Interactive Jobs
 - Frequent I/O → keep priority high
 - Interactive jobs require fast response time (GUI/UI)
- Batch Jobs
 - Require long periods of CPU utilization
 - Keep priority low

[High Priority] Q8 → (A) → (B)

Q7

Q6

Q5

Q4 → (C)

Q3

Q2

[Low Priority] Q1 → (D)

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.64
----------------	---	-------

MLFQ: DETERMINING JOB PRIORITY

- New arriving jobs are placed into highest priority queue
- If a job uses its entire time slice, priority is reduced (↓)
 - Jobs appears CPU-bound (“batch” job), not interactive (GUI/UI)
- If a job relinquishes the CPU for I/O priority stays the same


MLFQ approximates SJF

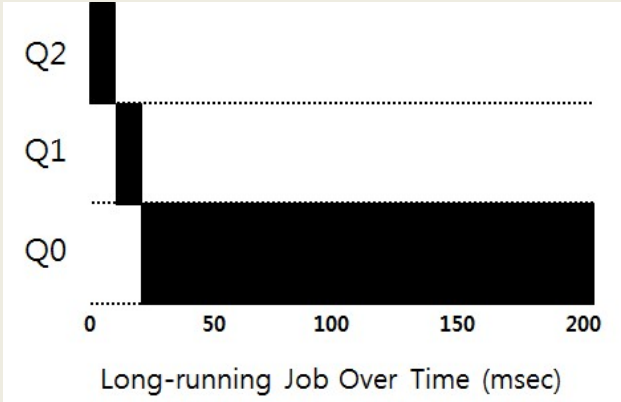
April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.65
----------------	---	-------

MLFQ: LONG RUNNING JOB

- Three-queue scheduler, time slice=10ms

Priority





April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.66
----------------	---	-------

MLFQ: BATCH AND INTERACTIVE JOBS

- $A_{arrival_time} = 0ms, A_{run_time} = 200ms,$
- $B_{run_time} = 20ms, B_{arrival_time} = 100ms$

Priority

Scheduling multiple jobs (ms)

A:

B:

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.67
----------------	---	-------

MLFQ: BATCH AND INTERACTIVE - 2

- Continuous interactive job (B) with long running batch job (A)
 - Low response time is good for B
 - A continues to make progress

The MLFQ approach keeps interactive job(s) at the highest priority

A Mixed I/O-intensive and CPU-intensive Workload (msec)

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.68
----------------	---	-------

MLFQ: ISSUES

- **Starvation**

[High Priority] Q8 → (A) → (B) → (C) → (D) → (E) → (F)

Q7

Q6

Q5

Q4

Q3

Q2

[Low Priority] Q1 → (G) → (H) *CPU bound batch job(s)*

April 13, 2021 TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma L5.69

MLFQ: ISSUES - 2

- **Gaming the scheduler**
 - Issue I/O operation at 99% completion of the time slice
 - Keeps job priority fixed – never lowered
- **Job behavioral change**
 - CPU/batch process becomes an interactive process

[High Priority] Q8 → (A) → (B) → (C) → (D) → (E) → (F)

Q7

Q6

Q5

Q4

Q3


Q2

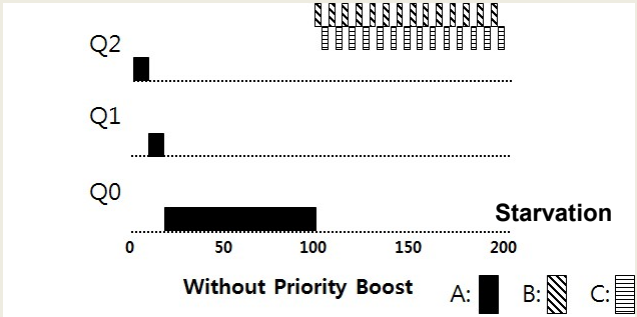
[Low Priority] Q1 → (G) → (H) *CPU bound batch job(s)*

Priority becomes stuck →

April 13, 2021 TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma L5.70

RESPONDING TO BEHAVIOR CHANGE





Without Priority Boost A: B: C:


- **Priority Boost**
 - **Reset all jobs to topmost queue after some time interval S**

April 13, 2021

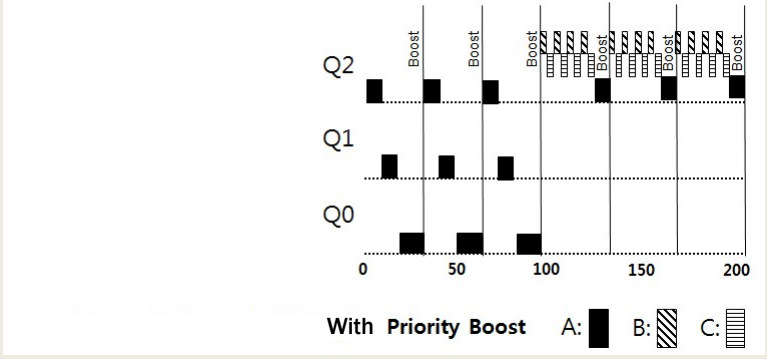
TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.71

RESPONDING TO BEHAVIOR CHANGE - 2



- **With priority boost**
 - **Prevents starvation**



With Priority Boost A: B: C:

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.72

KEY TO UNDERSTANDING MLFQ – PB

- Without priority boost:
- **Rule 1:** If Priority(A) > Priority(B), A runs (B doesn't).
- **Rule 2:** If Priority(A) = Priority(B), A & B run in RR.
- **KEY:** If time quantum of a higher queue is filled, then we don't run any jobs in lower priority queues!!!

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.73
----------------	---	-------

STARVATION EXAMPLE

- **Consider 3 queues:**
- Q2 - HIGH PRIORITY - Time Quantum 10ms
- Q1 - MEDIUM PRIORITY - Time Quantum 20 ms
- Q0 - LOW PRIORITY - Time Quantum 40 ms
- Job A: 200ms no I/O
- Job B: 5ms then I/O
- Job C: 5ms then I/O
- Q2 fills up, starves Q1 & Q0
- A makes no progress

Without Priority Boost

A: B: C:

April 13, 2021	TCSS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.74
----------------	---	-------

PREVENTING GAMING

- Improved time accounting:
 - Track total job execution time in the queue
 - Each job receives a fixed time allotment
 - When allotment is exhausted, job priority is lowered

Without(Left) and With(Right) Gaming Tolerance

April 13, 2021	TCCS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.75
----------------	---	-------

MLFQ: TUNING

- Consider the tradeoffs:
 - How many queues?
 - What is a good time slice?
 - How often should we “Boost” priority of jobs?
 - What about different time slices to different queues?

Example) 10ms for the highest queue, 20ms for the middle, 40ms for the lowest

April 13, 2021	TCCS422: Operating Systems [Spring 2021] School of Engineering and Technology, University of Washington - Tacoma	L5.76
----------------	---	-------

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

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.77

MLFQ RULE SUMMARY

- The refined set of MLFQ rules:
- **Rule 1:** If $\text{Priority}(A) > \text{Priority}(B)$, A runs (B doesn't).
- **Rule 2:** If $\text{Priority}(A) = \text{Priority}(B)$, 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.

April 13, 2021

TCSS422: Operating Systems [Spring 2021]
School of Engineering and Technology, University of Washington - Tacoma

L5.78

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 lose points.



EXAMPLE

- **Question:**
- Given a system with a quantum length of 10 ms in its highest queue, how often would you have to boost jobs back to the highest priority level to guarantee that a single long-running (and potentially starving) job gets at least 5% of the CPU?
- Some combination of n short jobs runs for a total of 10 ms per cycle without relinquishing the CPU
 - E.g. 2 jobs = 5 ms ea; 3 jobs = 3.33 ms ea, 10 jobs = 1 ms ea
 - n jobs always uses full time quantum (10 ms)
 - Batch jobs starts, runs for full quantum of 10ms
 - All other jobs run and context switch totaling the quantum per cycle
 - If 10ms is 5% of the CPU, when must the priority boost be ???
 - **ANSWER** → *Priority boost should occur every 200ms*

April 13, 2021

TCCS422: Operating Systems [Spring 2021]
 School of Engineering and Technology, University of Washington - Tacoma

L5.80

