

# TCSS 422: OPERATING SYSTEMS

## Limited Direct Execution, Introduction to OS/CPU Scheduling



Wes J. Lloyd

School of Engineering and Technology  
University of Washington - Tacoma

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington

Tacoma

1

## TEXT BOOK COUPON

- 15% off textbook code: **AAC72SAVE15**
- <https://www.lulu.com/shop/andrea-arpaci-dusseau-and-remzi-arpaci-dusseau/operating-systems-three-easy-pieces-hardcover-version-110/hardcover/product-15gjeeky.html?q=three+easy+pieces+operating+systems&page=1&pageSize=4>
- With coupon textbook is only \$33.79 + tax & shipping

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.2

2

## TCSS 422 – OFFICE HRS – WINTER 2026

- **Office Hours plan for Winter:**
- **Tuesday 2:30 - 3:30 pm Instructor Wes, Zoom**
- **Tue/Thur 6:00 - 7:00 pm Instructor Wes, CP 229/Zoom**
- **Tue 6:00 – 7:00 pm GTA Robert, Zoom/MDS 302**
- **Wed 1:00 – 2:00 pm GTA Robert, Zoom/MDS 302**
- **Instructor is available after class at 6pm in CP 229 each day**

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.3

3

## OBJECTIVES – 1/22

- **Questions from 1/20**
- **Assignment 0**
- **C Tutorial - Pointers, Strings, Exec in C**
- **Chapter 6: Limited Direct Execution**
- **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**
  - **MLFQ Scheduler**
  - **Job Starvation**
  - **Gaming the Scheduler**
  - **Examples**

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.4

4

## 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

TCSS 422 A > Assignments

Spring 2021

Home

Announcements

Zoom

Syllabus

Assignments

Discussions

January 22, 2026

TCSS422: Computer Operating Systems [Spring 2025]  
School of Engineering and Technology, University of Washington - Tacoma

L5.5

5

### 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				

January 22, 2026

TCSS422: Computer Operating Systems [Spring 2025]  
School of Engineering and Technology, University of Washington - Tacoma

L5.6

6

## MATERIAL / PACE

- Please classify your perspective on material covered in today's class (32 of 46 respondents – 69.5%) :
- **\*\*\*SURVEY ISSUE\*\*\***: Canvas Survey Initially Unavailable, **NOW POSTED AND AVAILABLE** until Mon Jan 26
- 1-mostly review, 5-equal new/review, 10-mostly new
- **Average – 7.03 (↑ - previous 6.54)**
  
- Please rate the pace of today's class:
- 1-slow, 5-just right, 10-fast
- **Average – 5.08 (↑ - previous 4.73)**

January 22, 2026

TCSS422: Computer Operating Systems [Spring 2025]  
School of Engineering and Technology, University of Washington - Tacoma

L5.7

7

## FEEDBACK FROM 1/20

- **Where would you use a fork() ?**
- The fork() API is used to create a new process in Linux
- **\*What is strange\*** –how new processes are created in Linux
- Instead of creating a brand-new squeaky-clean process, with an empty heap, stack, & data segments, fork() clones the existing process
- A complete copy of the parent process is made – BUT...
  - The copy is made using **COW: copy on write**
  - Memory is only cloned when data changes
- COW allows fork() to occur in a huge program, while actually doing minimal memory cloning

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.8

8

## FEEDBACK - 2

- Does OS always fork the parent process? Or is there any case to fork the child process?

- The fork() API can be called inside any process, to clone the existing process, to create a new one

- Inside a main program, you can have nested forks:

```

parent-fork()
  /   \
parent  child-1-fork()
  /   \
child-1  child-2-fork()
  /   \
child2  child3

```

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.9

9

## FEEDBACK - 3

- What does the fork bomb do if it doesn't call fork() literally?

- fork() is a C API

- In bash, the fork-bomb creates a function called ":"

- The fork bomb is

```
:(){ :|:&};:
```

- If we rename ":" to "funca", what we are really doing is defining a bash function call "funca":

```
funca(){ funca | funca & };
```

- Each call to funca makes 2 more funca calls

- Each funca call is run as a separate process

- Every command in bash is run as a separate process

- After the funca definition, funca is called (with ":")

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.10

10

## FEEDBACK - 4

- **How do you call a custom program with the exec APIs ?**
- The same as you would call any Linux command
- It must be in the system path
- If your executable is in the present working directory, then “.” must be in the path
- Check the path with:  
`echo $PATH`
- If the executable is not in the path, a fully qualified pathname is required to run the executable
- The ‘which’ command reports what command will be run
- Myargs should always include the command and arguments

January 22, 2026

TCSS422: Operating Systems [Winter 2026]

School of Engineering and Technology, University of Washington - Tacoma

L5.11

11

## OBJECTIVES - 1/22

- Questions from 1/20
- **Assignment 0**
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
- 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
  - MLFQ Scheduler
  - Job Starvation
  - Gaming the Scheduler
  - Examples

January 22, 2026

TCSS422: Operating Systems [Winter 2026]

School of Engineering and Technology, University of Washington - Tacoma

L5.12

12

## OBJECTIVES – 1/22

- Questions from 1/20
- Assignment 0
- **C Tutorial - Pointers, Strings, Exec in C**
- Chapter 6: Limited Direct Execution
- 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
  - MLFQ Scheduler
  - Job Starvation
  - Gaming the Scheduler
  - Examples

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.13

13

## FINISH CHAPTER 5

- Switch to Lecture 4 Slides
- Slides L4.51 to L4.60 (thru system calls and traps)

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.14

14

## OBJECTIVES – 1/22

- Questions from 1/20
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- **Chapter 6: Limited Direct Execution**
- 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
  - MLFQ Scheduler
  - Job Starvation
  - Gaming the Scheduler
  - Examples

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.15

15

## CH. 6: LIMITED DIRECT EXECUTION



January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.16

16

## CHAPTER 6

- **Chapter 6: Limited Direct Execution**
  - Direct execution
  - Limited direct execution
  - CPU modes
  - System calls and traps
  - **Cooperative multi-tasking**
  - Context switching and preemptive multi-tasking

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.17

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 see with this approach?

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.18


18

## MULTITASKING

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

January 22, 2026 TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma L5.19

19



When poll is active respond at PollEv.com /weslloyd Send weslloyd and your message to 22333

QR code

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

Too many processes calling yield at once? 

Program has too much control over OS, too much trust to user 

SEE MORE 

the OS is at the mercy of whomever wrote a

Current responses

20

## QUESTION: MULTITASKING

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

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.21

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)
  - (POLLIEV) What is a good interval for the timer interrupt?

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.22

22

## MULTITASKING - 2


- **Preemptive multitasking** (32 & 64 bit OSes)
- >= Mac OSX, Windows 95+
- **Timer**
  - Raises a interrupt
  - Intel

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

  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 22, 2026	TCSS422: Operating Systems [Winter 2026] School of Engineering and Technology, University of Washington - Tacoma	L5.23
------------------	---	-------

23



When poll is active respond at [PollEv.com/weslloyd](http://PollEv.com/weslloyd) Send weslloyd and your message to 22333

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?

Join by Web [PollEv.com/weslloyd](http://PollEv.com/weslloyd) Loading... Join by QR code Scan with your camera app

Join by Text  
Send **weslloyd** and your message to 22333

**Current responses**

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?

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.25

25

## 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?
  - Typical time slice for process execution is **10 to 100 milliseconds**
  - Typical context switch overhead is (*switch between processes*) **0.01 milliseconds**
    - 0.1% of the time slice (1/1000<sup>th</sup>)

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.26

26

## CHAPTER 6

- **Chapter 6: Limited Direct Execution**

- Direct execution
- Limited direct execution
- CPU modes
- System calls and traps
- Cooperative multi-tasking
- **Context switching and preemptive multi-tasking**

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.27

27

## 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 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.28

28

## 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 22, 2026

TCSS422: Operating Systems [Winter 2026]


School of Engineering and Technology, University of Washington - Tacoma

L5.29

29



30



31

## CLASS BREAK - QUESTIONS

- When dealing with OS control trade-offs, what does an OS have to do so that there is not too much overhead?**
  - OS functions must use fast, lightweight data structures for the process context, and fast scheduling algorithms
  - Context switches must be engineered to occur rapidly
  - Memory virtualization approach (how the system translates between virtual and physical addresses) needs to be super fast
  - CPUs include a special cache to help address translation called the Translation Lookaside Buffer (TLB). Essentially once an address has been translated, the cached translation is used until the address falls out of the cache
- Should we memorize Linux terminal commands?**
  - Programming/scripting involves much trial-and-error. Knowing more commands without looking them up enables more trials in less time

January 22, 2026      TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma      L5.32

32

## EXTRA CREDIT SEMINAR – FRIDAY

- **Research Talk:** Securing the Blindspots: From Formal Verification to Generative Verification
- Dr. Shakthi Yaras Weerasinghe, Postdoctoral Research Associate at the University of Arizona
- Friday Jan 23, 12:20pm - 1:20pm Milgard 110
- **Abstract:** This seminar explores the critical challenge of "authorization blindspots"—latent security vulnerabilities created when decentralized microservice policies diverge. Attendees will follow a research trajectory designed to detect, and validate these risks through a novel, multi-layered framework. We will first explore Automated Formal Verification, demonstrating how microservice source code can be encoded into SMT constraint models to deterministically identify end-to-end inconsistencies without manual specification. The discussion will then move to Generative Validation, illustrating how formal static analysis can be hybridized with Large Language Models. Participants will see how policy-enriched Intermediate Representations (IR) are used to guide LLMs in generating executable, downstream-aware tests that empirically validate security drifts with high semantic validity.

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.33

33

WE WILL RETURN AT  
~4:55PM



January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.34

34

## 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

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.35

35

## 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

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.36

36

## OBJECTIVES – 1/22

- Questions from 1/20
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
- **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
  - MLFQ Scheduler
  - Job Starvation
  - Gaming the Scheduler
  - Examples

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.37

37

## CHAPTER 7- SCHEDULING: INTRODUCTION



January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.38

38

## 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?

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.39

39

## 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

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.40

40

With n=3 and  $x_1=.2$ ,  $x_2=.7$ ,  $x_3=.1$ 

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

$$\begin{aligned} & \frac{1}{3 \cdot (.2)^2 + (.7)^2 + (.1)^2} \\ & 3 \cdot (.04 + .49 + .01) = \frac{1}{1.62} = .617 \\ & 3 \cdot (.54) \end{aligned}$$

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.41

41

With n=3 and  $x_1=.2$ ,  $x_2=.7$ ,  $x_3=.1$ 

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

$$\begin{aligned} & \frac{(.2 + .7 + .1)^2}{3 \cdot (.2)^2 + (.7)^2 + (.1)^2} = \frac{1}{1.62} = .617 \\ & 3 \cdot (.04 + .49 + .01) \\ & 3 \cdot (.54) \\ & 1.62 \end{aligned}$$

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.42

42

With n=3 and  $x_1=.33$ ,  $x_2=.33$ ,  $x_3=.33$ 

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

$$\begin{aligned} & \overline{3 \cdot \left( .33^2 + .33^2 + .33^2 \right)} \\ & \overline{3 \cdot ( .1089 + .1089 + .1089 )} \\ & \overline{3 \cdot ( .3267 )} = \overline{.9801} = 1 \end{aligned}$$

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.43

43

With n=3 and  $x_1=.33$ ,  $x_2=.33$ ,  $x_3=.33$ 

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

$$\begin{aligned} & \overline{.33 + .33 + .33} \rightarrow 1 \\ & \overline{3 \cdot \left( .33^2 + (.33)^2 + (.33)^2 \right)} \rightarrow .9801 \\ & \overline{3 \cdot ( .1089 + .1089 + .1089 )} \\ & \overline{3 \cdot ( .3267 )} \\ & \overline{3 \cdot ( ( .333 )^2 + \dots \dots \dots )} \\ & \overline{3 \cdot ( .9801 )} \overline{\overline{.998}} \rightarrow 1.002 \end{aligned}$$

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.44

44

## OBJECTIVES – 1/22

- Questions from 1/20
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
- 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
  - MLFQ Scheduler
  - Job Starvation
  - Gaming the Scheduler
  - Examples

January 22, 2026


TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.45

45

## SCHEDULERS

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



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

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.46

46

## OBJECTIVES – 1/22

- Questions from 1/20
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
- 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
  - MLFQ Scheduler
  - Job Starvation
  - Gaming the Scheduler
  - Examples

January 22, 2026

TCSS422: Operating Systems [Winter 2026]


School of Engineering and Technology, University of Washington - Tacoma

L5.47

47

## 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



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

January 22, 2026

TCSS422: Operating Systems [Winter 2026]


School of Engineering and Technology, University of Washington - Tacoma

L5.48

48

## 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



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

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.49

49

## OBJECTIVES – 1/22

- Questions from 1/20
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
- 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
  - MLFQ Scheduler
  - Job Starvation
  - Gaming the Scheduler
  - Examples

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma


L5.50

50

## STCF: SHORTEST TIME TO COMPLETION FIRST



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



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

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.51

51

## OBJECTIVES – 1/22

- Questions from 1/20
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
- 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
  - MLFQ Scheduler
  - Job Starvation
  - Gaming the Scheduler
  - Examples

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.52

52

## 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?

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.53

53

## OBJECTIVES – 1/22

- Questions from 1/20
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
- 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
  - MLFQ Scheduler
  - Job Starvation
  - Gaming the Scheduler
  - Examples

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.54

54

## RR: ROUND ROBIN





- Run each job awhile, then switch to another distributing the CPU evenly (fairly)
- Scheduling Quantum is called a time slice
- Time quantum is a maximum time a process can run in a period.

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

Process	Burst Time
P1	12

**Round Robin scheduling algorithm Gantt chart**

Scheduling Quantum = 5 seconds




January 22, 2026 | TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma | L5.55

55


## RR EXAMPLE

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

**SJF (Bad for Response Time)**



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



**OVERHEAD not considered**

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

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

January 22, 2026 | TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma | L5.56

56

## 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,  $J=1$

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
 School of Engineering and Technology, University of Washington - Tacoma

L5.57


57

## 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:



January 22, 2026


TCSS422: Operating Systems [Winter 2026]  
 School of Engineering and Technology, University of Washington - Tacoma

L5.58

58

## 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)



January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.59

59

When poll is active respond at [PollEv.com/weslloyd](https://PollEv.com/weslloyd) Send weslloyd to 22333

**W** Which scheduler, thus far, best address fairness and average response time of jobs? 0

First In - First Out (FIFO)

Shortest Job First (SJF)

Shortest Time to Completion First (STCF)

Round Robin

SEE MORE

Current responses

60

## 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

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.61

61

## 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:



January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.62

62

When poll is active respond at [PollEv.com/weslloyd](https://PollEv.com/weslloyd) Send weslloyd and your message to 22333

**W** What is the Average Response Time of the FIFO scheduler? 0

Join by Web Loading...

[PollEv.com/weslloyd](https://PollEv.com/weslloyd)

Join by Text Scan with your camera app

Send **weslloyd** and your message to **22333**

Join by QR code 

Current responses

63

When poll is active respond at [PollEv.com/weslloyd](https://PollEv.com/weslloyd) Send weslloyd and your message to 22333

**W** What is the Average Turnaround Time of the FIFO scheduler? 0

Join by Web Loading...

[PollEv.com/weslloyd](https://PollEv.com/weslloyd)

Join by Text Scan with your camera app

Send **weslloyd** and your message to **22333**

Join by QR code 

64

## 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:



January 22, 2026


TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.65

65

A screenshot of a PollEv poll interface. The poll question is: "What is the Average Response Time of the Shortest Job First Scheduler?". The poll is currently active, indicated by a "0" in a blue box. Below the question, there are two main join options: "Join by Web" and "Join by Text". The "Join by Web" option is linked to the URL "PollEv.com/weslloyd". The "Join by Text" option instructs users to "Send weslloyd and your message to 22333". To the right of these options is a QR code with the text "Join by QR code" and "Scan with your camera app". The poll interface includes standard navigation buttons for activities, moderation, visual settings, and editing.

66



Join by Web PollEv.com/weslloyd Join by Text Send weslloyd and your message to 22333

**W** What is the Average Turnaround Time of the Shortest Job First Scheduler? 0


Join by Web Loading... Join by QR code  
Scan with your camera app

PollEv.com/weslloyd

Join by Text  
Send **weslloyd** and your message to 22333

Current responses

67



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

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Thomas

L5.68

68

## OBJECTIVES – 1/22

- Questions from 1/20
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
- 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
  - **MLFQ Scheduler**
  - Job Starvation
  - Gaming the Scheduler
  - Examples

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.69

69

## 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

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.70

70

## 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)

January 22, 2026 | TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma | L5.71

71

## MLFQ: DETERMINING JOB PRIORITY

- New arriving jobs are placed into highest priority queue
- If a job uses its entire time slice, priority is reduced (↓)
  - Job 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

January 22, 2026 | TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma | L5.72

72

## MLFQ: LONG RUNNING JOB

- Three-queue scheduler, time slice=10ms



January 22, 2026


TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.73

73

## MLFQ: BATCH AND INTERACTIVE JOBS

- $A_{arrival\_time} = 0\text{ms}$ ,  $A_{run\_time} = 200\text{ms}$ ,
- $B_{run\_time} = 20\text{ms}$ ,  $B_{arrival\_time} = 100\text{ms}$



January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma


L5.74

74

## 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



January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.75

75

## OBJECTIVES – 1/22

- Questions from 1/20
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
- 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
  - MLFQ Scheduler
  - Job Starvation**
  - Gaming the Scheduler
  - Examples

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.76

76

## 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)*

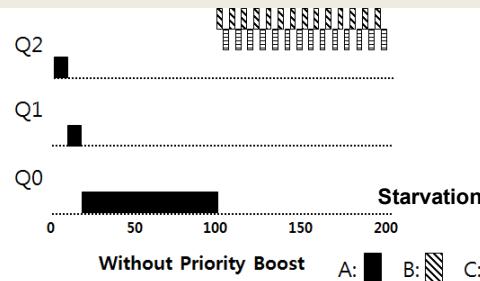
January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.77

77

## RESPONDING TO BEHAVIOR CHANGE



### ■ Priority Boost

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

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma


L5.78

78

## RESPONDING TO BEHAVIOR CHANGE - 2



- With priority boost
  - Prevents starvation



January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.79

79

## 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!!!

January 22, 2026


TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.80

80

## 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



January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.81

81

## OBJECTIVES – 1/22

- Questions from 1/20
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
- 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
  - MLFQ Scheduler
  - Job Starvation
  - **Gaming the Scheduler**
- Examples

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.82

82

## 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



January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma


L5.83

83

## 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



January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma


L5.84

84

## 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

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.85

85

## OBJECTIVES – 1/22

- Questions from 1/20
- Assignment 0
- C Tutorial - Pointers, Strings, Exec in C
- Chapter 6: Limited Direct Execution
- 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
  - MLFQ Scheduler
  - Job Starvation
  - Gaming the Scheduler
  - Examples

January 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.86

86

## 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 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.87

87

## 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 22, 2026

TCSS422: Operating Systems [Winter 2026]  
School of Engineering and Technology, University of Washington - Tacoma

L5.88

88

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.




89

## 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

90



91