


TCSS 422: OPERATING SYSTEMS

Lock-Based Data Structures, Midterm Review

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

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma



OBJECTIVES – 4/30

- **Questions from 4/28**
- **C Tutorial (Apr 30 11:59p AOE)**
- **Assignment 1 (May 7 11:59p AOE)**
- **Chapter 29: Lock Based Data Structures**
 - Sloppy Counter
 - Concurrent Structures: Linked List, Queue, Hash Table
- **Midterm Review**

April 30, 2020	TCSS422: Operating Systems [Spring 2020] School of Engineering and Technology, University of Washington - Tacoma	L10.2
----------------	---	-------

MATERIAL / PACE

- Please classify your perspective on material covered in today’s class (39 respondents):
 - 1-mostly review, 5-equal new/review, 10-mostly new
 - **Average – 7.30 (↑ from 7.21)**
- Please rate the pace of today’s class:
 - 1-slow, 5-just right, 10-fast
 - **Average – 5.92 (↑ from 5.53)**

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.3

FEEDBACK FROM 4/28

- C tutorial (v.11)
- Question 6 – rephrased as:
 - What is ALWAYS the ASCII integer value for the last character?

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.4

OBJECTIVES – 4/30

- Questions from 4/28
- C Tutorial (Apr 30 11:59p AOE)
- Assignment 1 (May 7 11:59p AOE)
- Chapter 29: Lock Based Data Structures
 - Sloppy Counter
 - Concurrent Structures: Linked List, Queue, Hash Table
- Midterm Review

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.5

OBJECTIVES – 4/30


- Questions from 4/28
- C Tutorial (Apr 30 11:59p AOE)
- Assignment 1 (May 7 11:59p AOE)
- Chapter 29: Lock Based Data Structures
 - Sloppy Counter
 - Concurrent Structures: Linked List, Queue, Hash Table
- Midterm Review

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.6

CHAPTER 29 – LOCK BASED DATA STRUCTURES



April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.7

OBJECTIVES – 4/30

- Questions from 4/28
- C Tutorial (Apr 30 11:59p AOE)
- Assignment 1 (May 7 11:59p AOE)
- **Chapter 29: Lock Based Data Structures**
 - Sloppy Counter
 - Concurrent Structures: Linked List, Queue, Hash Table
- Midterm Review

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.8

SLOPPY COUNTER

- Provides single logical shared counter
 - Implemented with local counters for each ~CPU core
 - 4 CPU cores = 4 local counters & 1 global counter
 - Local counters are synchronized via local locks
 - Global counter is updated periodically
 - Global counter has lock to protect global counter value
 - Sloppiness threshold (S):
 - Update threshold of global counter with local values
 - Small (S): more updates, more overhead
 - Large (S): fewer updates, more performant, less synchronized
- Local counters for each CPU Core:
 - Distribute work evenly, each core has independent local counter

April 30, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.9

SLOPPY COUNTER – MAIN POINTS

- Main idea:
RELAX synchronization requirement for counting
 - Instead of synchronizing global count variable each time:
counter=counter+1
 - Synchronization occurs only every so often:
e.g. every 1000 counts
- Relaxing the synchronization requirement drastically reduces locking API overhead by trading-off split-second accuracy of the counter
- Sloppy counter: trade-off accuracy for speed
 - It's sloppy because it's not so accurate (until the end)

April 30, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.10

SLOPPY COUNTER - 2

- Update threshold (S) = 5
- Synchronized across four CPU cores
- Threads update local CPU counters

Time	L_1	L_2	L_3	L_4	G
0	0	0	0	0	0
1	0	0	1	1	0
2	1	0	2	1	0
3	2	0	3	1	0
4	3	0	3	2	0
5	4	1	3	3	0
6	5 \rightarrow 0	1	3	4	5 (from L_1)
7	0	2	4	5 \rightarrow 0	10 (from L_4)

April 30, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.11

THRESHOLD VALUE S

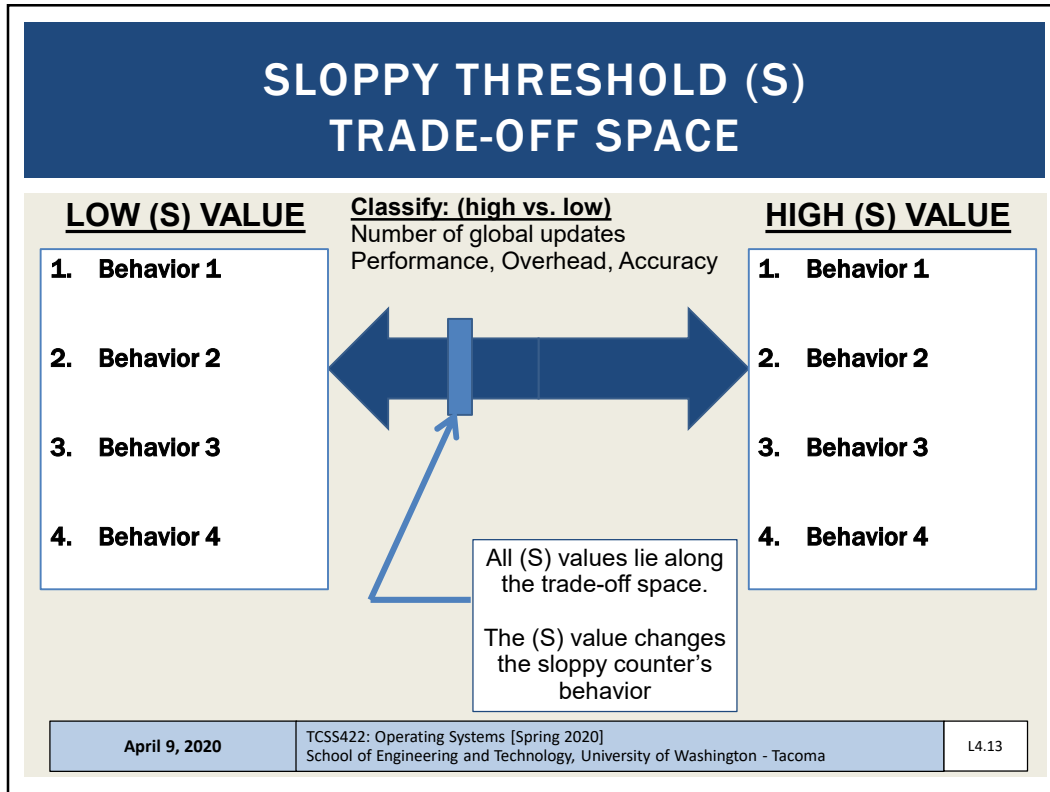
- Consider 4 threads increment a counter 1000000 times each
- Low $S \rightarrow$ What is the consequence?
- High $S \rightarrow$ What is the consequence?

Sloppiness	Time (seconds)
1	12
2	6
4	3
8	1.5
16	0.8
32	0.4
64	0.2
128	0.1
256	0.05
512	0.02
1024	0.01

April 30, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.12



SLOPPY COUNTER - EXAMPLE

- Example implementation
- Chapter 29
 - `sloppybasic.c`
- Also a version with CPU affinity:
 - `sloppy.c`
 - Local counters pinned to specific CPUs
 - Uses `sched_setaffinity` API call for CPU pinning

April 30, 2020	TCSS422: Operating Systems [Spring 2020] School of Engineering and Technology, University of Washington - Tacoma	L10.14
----------------	---	--------

CONCURRENT LINKED LIST - 1

- Simplification - only basic list operations shown
- Structs and initialization:

```
1 // basic node structure
2 typedef struct __node_t {
3     int key;
4     struct __node_t *next;
5 } node_t;
6
7 // basic list structure (one used per list)
8 typedef struct __list_t {
9     node_t *head;
10    pthread_mutex_t lock;
11 } list_t;
12
13 void List_Init(list_t *L) {
14     L->head = NULL;
15     pthread_mutex_init(&L->lock, NULL);
16 }
17
18 (Cont.)
```

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.15

CONCURRENT LINKED LIST - 2

- Insert – adds item to list
- Everything is critical!
 - There are two unlocks

```
(Cont.)
18 int List_Insert(list_t *L, int key) {
19     pthread_mutex_lock(&L->lock);
20     node_t *new = malloc(sizeof(node_t));
21     if (new == NULL) {
22         perror("malloc");
23         pthread_mutex_unlock(&L->lock);
24         return -1; // fail
25     }
26     new->key = key;
27     new->next = L->head;
28     L->head = new;
29     pthread_mutex_unlock(&L->lock);
30     return 0; // success
31 }
32
33 (Cont.)
```

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.16

CONCURRENT LINKED LIST - 3

- Lookup – checks list for existence of item with key
- Once again everything is critical
 - Note - there are also two unlocks

```
(Cont.)
32
32  int List_Lookup(list_t *L, int key) {
33      pthread_mutex_lock(&L->lock);
34      node_t *curr = L->head;
35      while (curr) {
36          if (curr->key == key) {
37              pthread_mutex_unlock(&L->lock);
38              return 0; // success
39          }
40          curr = curr->next;
41      }
42      pthread_mutex_unlock(&L->lock);
43      return -1; // failure
44  }
```

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.17

CONCURRENT LINKED LIST

- First Implementation:
 - Lock **everything** inside Insert() and Lookup()
 - If malloc() fails lock must be released
 - Research has shown “**exception-based control flow**” to be error prone
 - 40% of Linux OS bugs occur in rarely taken code paths
 - Unlocking in an exception handler is considered a poor coding practice
 - There is nothing specifically wrong with this example however
- Second Implementation ...

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.18

CCL – SECOND IMPLEMENTATION

■ Init and Insert

```
1  void List_Init(list_t *L) {
2      L->head = NULL;
3      pthread_mutex_init(&L->lock, NULL);
4  }
5
6  void List_Insert(list_t *L, int key) {
7      // synchronization not needed
8      node_t *new = malloc(sizeof(node_t));
9      if (new == NULL) {
10         perror("malloc");
11         return;
12     }
13     new->key = key;
14
15     // just lock critical section
16     pthread_mutex_lock(&L->lock);
17     new->next = L->head;
18     L->head = new;
19     pthread_mutex_unlock(&L->lock);
20 }
21
```

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.19

CCL – SECOND IMPLEMENTATION - 2

■ Lookup

```
(Cont.)
22  int List_Lookup(list_t *L, int key) {
23      int rv = -1;
24      pthread_mutex_lock(&L->lock);
25      node_t *curr = L->head;
26      while (curr) {
27          if (curr->key == key) {
28              rv = 0;
29              break;
30          }
31          curr = curr->next;
32      }
33      pthread_mutex_unlock(&L->lock);
34      return rv; // now both success and failure
35  }
```

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.20

CONCURRENT LINKED LIST PERFORMANCE

- Using a single lock for entire list is not very performant
- Users must “wait” in line for a single lock to access/modify any item
- Hand-over-hand-locking (lock coupling)
 - Introduce a lock for each node of a list
 - Traversal involves handing over previous node’s lock, acquiring the next node’s lock...
 - Improves lock granularity
 - Degrades traversal performance
- Consider hybrid approach
 - Fewer locks, but more than 1
 - Best lock-to-node distribution?



April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.21

MICHAEL AND SCOTT CONCURRENT QUEUES

- Improvement beyond a single master lock for a queue (FIFO)
- Two locks:
 - One for the **head** of the queue
 - One for the **tail**
- Synchronize enqueue and dequeue operations
- Add a dummy node
 - Allocated in the queue initialization routine
 - Supports separation of head and tail operations
- Items can be added and removed by separate threads at the same time

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.22

CONCURRENT QUEUE

■ Remove from queue

```
1  typedef struct __node_t {
2      int value;
3      struct __node_t *next;
4  } node_t;
5
6  typedef struct __queue_t {
7      node_t *head;
8      node_t *tail;
9      pthread_mutex_t headLock;
10     pthread_mutex_t tailLock;
11 } queue_t;
12
13 void Queue_Init(queue_t *q) {
14     node_t *tmp = malloc(sizeof(node_t));
15     tmp->next = NULL;
16     q->head = q->tail = tmp;
17     pthread_mutex_init(&q->headLock, NULL);
18     pthread_mutex_init(&q->tailLock, NULL);
19 }
20 (Cont.)
```

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.23

CONCURRENT QUEUE - 2

■ Add to queue

```
(Cont.)
21 void Queue_Enqueue(queue_t *q, int value) {
22     node_t *tmp = malloc(sizeof(node_t));
23     assert(tmp != NULL);
24
25     tmp->value = value;
26     tmp->next = NULL;
27
28     pthread_mutex_lock(&q->tailLock);
29     q->tail->next = tmp;
30     q->tail = tmp;
31     pthread_mutex_unlock(&q->tailLock);
32 }
(Cont.)
```

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.24

CONCURRENT HASH TABLE

- Consider a simple hash table
 - Fixed (static) size
 - Hash maps to a bucket
 - Bucket is implemented using a concurrent linked list
 - One lock per hash (bucket)
 - Hash bucket is a linked lists


April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.25

INSERT PERFORMANCE – CONCURRENT HASH TABLE

- Four threads – 10,000 to 50,000 inserts
 - iMac with four-core Intel 2.7 GHz CPU



**The simple concurrent hash table scales
magnificently.**

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.26

CONCURRENT HASH TABLE

```
1  #define BUCKETS (101)
2
3  typedef struct _hash_t {
4      list_t lists[BUCKETS];
5  } hash_t;
6
7  void Hash_Init(hash_t *H) {
8      int i;
9      for (i = 0; i < BUCKETS; i++) {
10         List_Init(&H->lists[i]);
11     }
12 }
13
14 int Hash_Insert(hash_t *H, int key) {
15     int bucket = key % BUCKETS;
16     return List_Insert(&H->lists[bucket], key);
17 }
18
19 int Hash_Lookup(hash_t *H, int key) {
20     int bucket = key % BUCKETS;
21     return List_Lookup(&H->lists[bucket], key);
22 }
```

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.27

Which is a major advantage of using concurrent data structures in your programs?

Locks are encapsulated within data structure code ensuring thread safety.

Lock granularity tradeoff already optimized inside data structure

Multiple threads can more easily share data

All of the above

None of the above

Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app


LOCK-FREE DATA STRUCTURES

- Lock-free data structures in Java
 - `Java.util.concurrent.atomic` package
 - Classes:
 - `AtomicBoolean`
 - `AtomicInteger`
 - `AtomicIntegerArray`
 - `AtomicIntegerFieldUpdater`
 - `AtomicLong`
 - `AtomicLongArray`
 - `AtomicLongFieldUpdater`
 - `AtomicReference`
 - See: <https://docs.oracle.com/en/java/javase/11/docs/api/java.base/java/util/concurrent/atomic/package-summary.html>

April 30, 2020	TCSS422: Operating Systems [Spring 2020] School of Engineering and Technology, University of Washington - Tacoma	L10.29
----------------	---	--------

TCSS 422 WILL RETURN AT ~2:30PM

April 30, 2020	TCSS422: Operating Systems [Spring 2020] School of Engineering and Technology, University of Washington - Tacoma	L10.30
----------------	---	--------





MIDTERM REVIEW

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.3
1

OBJECTIVES – 4/30

- Questions from 4/28
- C Tutorial (Apr 30 11:59p AOE)
- Assignment 1 (May 7 11:59p AOE)
- Chapter 29: Lock Based Data Structures
 - Sloppy Counter
 - Concurrent Structures: Linked List, Queue, Hash Table
- Midterm Review

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.32

MIDTERM

- Tuesday May 5th
- ONLINE via Canvas (for 3 hrs 12:30 – 3:30p)
- Additional hour provided in case of internet issues, etc.
- Open book, note, internet
- Individual work

- Preparation:
- Practice quiz: CPU scheduling (*to be posted*)
 - Auto grading w/ multiple attempts allowed as study aid
- Practice THURSDAY – first hour of lecture
 - Series of problems presented with some time to solve
 - Will then work through solutions
- Second hour – new material not on midterm

April 30, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.33

FIFO EXAMPLE

- Operation of CPU schedulers can be visualized with timing graphs.
- The graph below depicts a FIFO scheduler where three jobs arrive in the sequence A, B, C, where job A runs for 10 time slices, job B for 5 time slices, and job C for 10 time slices.

FIFO | AAAAAAAAAABBBBBCCCCCCCCC

0 10 15 25

April 30, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.34

Q1- SHORTEST JOB FIRST (SJF) SCHEDULER

■ Draw a scheduling graph for the SJF scheduler without preemption for the following jobs. Draw vertical lines for key events and be sure to label the X-axis times as in the example.

Job	Arrival Time	Job Length
A	T=0	25
B	T=5	10
C	T=10	15

SJF

|

|

|

|

|

0

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.35

Q1 – SJF - 2

What is the response time (RT) and turnaround time (TT) for jobs A, B, and C?

RT Job A: _____

TT Job A: _____

RT Job B: _____

TT Job B: _____

RT Job C: _____

TT Job C: _____

What is the average response time for all jobs? _____

What is the average turnaround time for all jobs? _____

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.36

L10.37

L10.38

Q3 - OPERATING SYSTEM APIs

1. Provide a definition for what is a blocking API call
2. Provide a definition for a non-blocking API call
3. Provide an example of a blocking API call.
Consider APIs used to manage processes and/or threads.
4. Provide an example of a non-blocking API call.
Consider APIs used to manage processes and/or threads.

April 30, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.39

Q4 – OPERATING SYSTEM APIs - II

1. When implementing memory synchronization for a multi-threaded program list one **advantage** of combining the use of a condition variable with a lock variable via the Linux C thread API calls: `pthread_mutex_lock()` and `pthread_cond_wait()`
2. When implementing memory synchronization for a multi-threaded program using locks, list one **disadvantage** of using blocking thread API calls such as the Linux C thread API calls for: `pthread_mutex_lock()` and `pthread_cond_wait()`
3. List (2) factors that cause Linux blocking API calls to introduce **overhead** into programs:

April 30, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.40

Q5 – PERFECT MULTITASKING OPERATING SYSTEM

In a perfect-multi-tasking operating system, every process of the same priority will always receive exactly $1/n^{\text{th}}$ of the available CPU time. Important CPU improvements for multi-tasking include: (1) fast context switching to enable jobs to be swapped in-and-out of the CPU very quickly, and (2) the use of a timer interrupt to preempt running jobs without the user voluntarily yielding the CPU. These innovations have enabled major improvements towards achieving a coveted “Perfect Multi-Tasking System”.

List and describe two challenges that remain complicating the full realization of a Perfect Multi-Tasking Operating System. In other words, what makes it very difficult for all jobs (for example, 10 jobs) of the same priority to receive **EXACTLY** the same runtime on the CPU? Your description must explain why the challenge is a problem for achieving perfect multi-tasking.

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.41

Q6 – ROUND-ROBIN SCHEDULER

Show a scheduling graph for a Round-Robin (RR) scheduler with job preemption where newly arriving jobs will immediately run. Assume a time slice of 3 timer units. Draw vertical lines for key events and be sure to label the X-axis times as in the example.

Job	Arrival Time	Job Length
A	T=0	25
B	T=5	10
C	T=10	15



April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.42

Q6 – RR SCHEDULER - 2

Using the graph, from time t=10 until all jobs complete at t=50, evaluate Jain’s Fairness Index:

Jain’s fairness index is expressed as:

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

Where n is the number of jobs, and x_i is the time share of each process Jain’s fairness index=1 for best case fairness, and 1/n for worst case fairness.

For the time window from t=10 to t=50, what percentage of the CPU time is allocated to each of the jobs A, B, and C?

Job A: _____ Job B: _____ Job C: _____

With these values, calculate Jain’s fairness index from t=10 to t=50.

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.43

Q7 – SLOPPY COUNTER

Below is a tradeoff space graph similar to those we’ve shown in class. Based on the sloppy counter threshold (S), add numbers on the **left** or **right** side of the graph for each of the following tradeoffs:

1. High number of Global Updates

2. High Performance

3. High Overhead

4. High Accuracy

5. Low number of Global Updates

6. Low Performance

7. Low Overhead

8. Low Accuracy

Low sloppy threshold (S)

High sloppy threshold (S)

|-----|

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.44

MULTI-LEVEL FEEDBACK QUEUE

- Review the bonus lecture for examples of Multi-level-feedback-queue problems (MLFQ)
- <https://tinyurl.com/y8ucda5z>

April 30, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L10.45

QUESTIONS



**WILL RETURN IN A FEW
MINUTES**

