


TCSS 422: OPERATING SYSTEMS

Three Easy Pieces:
CVs, Concurrency Problems,
Intro to Memory Virtualization



Wes J. Lloyd
Institute of Technology
University of Washington - Tacoma

May 7, 2018

TCSS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

OBJECTIVES

- Midterm Review
- Assignment 2 – Matrix Task Processor
- Condition Variables – Ch. 30
- Concurrency Problems – Ch. 32
- Memory Virtualization Is next...**
- Address Spaces – Ch. 13
- Memory API – Ch. 14
- Address Translation – Ch. 15
- Segmentation – Ch. 15

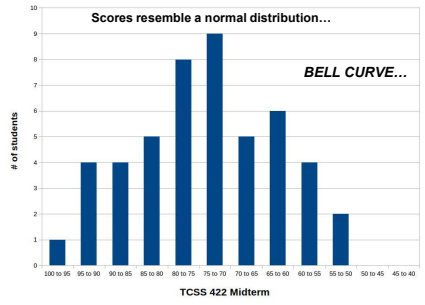
May 7, 2018

TCSS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.2

SPRING 2018 MIDTERM

Scores resemble a normal distribution...



BELL CURVE...

TCSS 422 Midterm

May 7, 2018

TCSS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.3

MIDTERM - 2


- Average: 73.44%
- Mode: 81
- Median: 73.5
- Lower quartile \leq 64.75%
- Second quartile \leq 73.5%
- Third quartile \leq 81%
- Upper quartile \leq 97%
- Curve: +12
- Be concerned if you're in the lower quartile
 - More effort is needed

November 7, 2017

TCSS558: Applied Distributed Computing [Fall 2017]
Institute of Technology, University of Washington - Tacoma

L11.4

CHAPTER 30 –
CONDITION VARIABLES




May 7, 2018

TCSS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.5

CONDITION VARIABLES



- Support a signaling mechanism to alert threads when preconditions have been satisfied
- Eliminate busy waiting
- Alert one or more threads to “consume” a result, or respond to state changes in the application
- Threads are placed on an **explicit queue** (FIFO) to wait for signals
- Signal**: wakes one thread
broadcast wakes all (ordering by the OS)

May 7, 2018

TCSS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.6

Slides by Wes J. Lloyd

L10.1

CONDITION VARIABLES - 3

Condition variable

```
pthread_cond_t c;
```

- Requires initialization

Condition API calls

```
pthread_cond_wait(pthread_cond_t *c, pthread_mutex_t *m); // wait()
pthread_cond_signal(pthread_cond_t *c); // signal()
```

wait() accepts a mutex parameter

- Releases lock, puts thread to sleep

signal()

- Wakes up thread, awakening thread acquires lock

May 7, 2018

TCS5422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.7

MATRIX GENERATOR

Matrix generation example

Chapter 30
signal.c

May 7, 2018

TCS5422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.8

SUBTLE RACE CONDITION: WITHOUT A WHILE

```
1 void thr_exit() {
2     done = 1;
3     pthread_cond_signal(&c);
4 }
5
6 void thr_join() {
7     if (done == 0)
8         pthread_cond_wait(&c);
9 }
```


- Parent thread calls thr_join() and executes the comparison
- The context switches to the child
- The child runs thr_exit() and signals the parent, but the parent is not waiting yet.
- The signal is lost**
- The parent deadlocks

May 7, 2018

TCS5422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.9

PRODUCER / CONSUMER



May 7, 2018

TCS5422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.10

PRODUCER / CONSUMER

Producer

- Produces items – consider the child matrix maker
- Places them in a buffer
 - Example: the buffer is only 1 element (single array pointer)

Consumer

- Grabs data out of the buffer
- Our example: parent thread receives dynamically generated matrices and performs an operation on them
 - Example: calculates average value of every element (integer)

Multithreaded web server example

- Http requests placed into work queue; threads process

May 7, 2018

TCS5422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.11

PRODUCER / CONSUMER - 2

Producer / Consumer is also known as **Bounded Buffer**

Bounded buffer

- Similar to piping output from one Linux process to another
- grep pthread signal.c | wc -l
- Synchronized access:
 - sends output from grep → wc as it is produced
- File stream

May 7, 2018

TCS5422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.12

PUT/GET ROUTINES

- Buffer is a one element shared data structure (int)
- Producer “puts” data
- Consumer “gets” data
- Shared data structure requires synchronization

```
1  int buffer;  
2  int count = 0; // initially, empty  
3  
4  void put(int value) {  
5      assert(count == 0);  
6      count = 1;  
7      buffer = value;  
8  }  
9  
10 int get() {  
11     assert(count == 1);  
12     count = 0;  
13     return buffer;  
14 }
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.13

PRODUCER / CONSUMER - 3

- Producer adds data
- Consumer removes data (busy waiting)
- Will this code work (w/ln locks) with 2-threads?

1. Producer

2. Consumer

```
1  void *producer(void *arg) {  
2      int i;  
3      int loops = (int) arg;  
4      for (i = 0; i < loops; i++) {  
5          put(i);  
6      }  
7  }  
8  
9  void *consumer(void *arg) {  
10     int i;  
11     while (1) {  
12         int tmp = get();  
13         printf("%d\n", tmp);  
14     }  
15 }
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.14

PRODUCER / CONSUMER - 3

- The shared data structure needs synchronization!

```
1  cond_t cond;  
2  mutex_t mutex;  
3  
4  void *producer(void *arg) {  
5      int i;  
6      for (i = 0; i < loops; i++) {  
7          pthread_mutex_lock(&mutex); // p1  
8          if (count == 1) // p2  
9              pthread_cond_wait(&cond, &mutex); // p3  
10         put(i); // p4  
11         pthread_cond_signal(&cond); // p5  
12         pthread_mutex_unlock(&mutex); // p6  
13     }  
14 }  
15  
16 void *consumer(void *arg) {  
17     int i;  
18     for (i = 0; i < loops; i++) {  
19         pthread_mutex_lock(&mutex); // c1
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.15

PRODUCER/CONSUMER - 4

```
20     if (count == 0) // c2  
21         pthread_cond_wait(&cond, &mutex); // c3  
22     int tmp = get(); // c4  
23     pthread_cond_signal(&cond); // c5  
24     pthread_mutex_unlock(&mutex); // c6  
25     printf("%d\n", tmp);  
26 }  
27 }
```

- This code as-is works with just:
 (1) Producer
 (1) Consumer
- If we scale to (2+) consumer's it fails
 ▪ How can it be fixed ?

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.16

EXECUTION TRACE:
NO WHILE, 1 PRODUCER, 2 CONSUMERS

- Two threads

Legend
c1/p1- lock
c2/p2- check var
c3/p3- wait
c4- put()
p4- get()
c5/p5- signal
c6/p6- unlock

	T _{c1}	State	T _{c2}	State	T _p	State	Count	Comment
	c1	Running		Ready		Ready	0	
	c2	Running		Ready		Ready	0	
	c3	Sleep		Ready		Ready	0	Nothing to get
		Sleep		Ready	p1	Running	0	
		Sleep		Ready	p2	Running	0	
		Sleep		Ready	p4	Running	1	Buffer now full T _{c1} awoken
		Ready		Ready	p5	Running	1	
		Ready		Ready	p6	Running	1	
		Ready		Ready	p1	Running	1	
		Ready		Ready	p2	Running	1	
		Ready		Ready	p3	Sleep	1	Buffer full; sleep
	c1	Running		Sleep		Sleep	1	T _{c2} sneaks in ...
	c2	Running		Sleep		Sleep	1	
	c4	Running		Sleep		Sleep	0	... and grabs data
	c5	Running		Ready		Ready	0	T _p awoken
	c6	Running		Ready		Ready	0	
	c4	Running		Ready		Ready	0	Oh oh! No data

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.17

PRODUCER/CONSUMER
SYNCHRONIZATION

- When producer threads awake, they do not check if there is any data in the buffer...
 - Need while, not if
- What if T_p puts a value, wakes T_{c1} whom consumes the value
- Then T_p has a value to put, but T_{c1}'s signal on &cond wakes T_{c2}
- There is nothing for T_{c2} consume, so T_{c2} sleeps
- T_{c1}, T_{c2}, and T_p all sleep forever
- T_{c1} needs to wake T_p to T_{c2}

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.18

EXECUTION TRACE:
WHILE, 1 CONDITION, 1 PRODUCER, 2 CONSUMERS

Legend

c1/p1- lock
c2/p2- check var
c3/p3- wait
c4- put()
p4- get()
c5/p5- signal
c6/p6- unlock

T _{c1}	State	T _{c2}	State	T _p	State	Count	Comment
c1	Running		Ready		Ready	0	
c2	Running		Ready		Ready	0	
c3	Ready		Ready		Ready	0	Nothing to get
	Sleep	c1	Running		Ready	0	
	Sleep	c2	Running		Ready	0	
	Sleep	c3	Sleep		Ready	0	Nothing to get
	Sleep		Sleep	p1	Running	0	
	Sleep		Sleep	p2	Running	0	
	Sleep		Sleep	p4	Running	1	Buffer now full
	Ready		Sleep	p5	Running	1	T _{c1} awoken
	Ready		Sleep	p6	Running	1	
	Ready		Sleep	p1	Running	1	
	Ready		Sleep	p2	Running	1	
	Ready		Sleep	p3	Sleep	1	Must sleep (full)
c2	Running		Sleep		Sleep	1	Recheck condition
c4	Running		Sleep		Sleep	0	T _{c1} grabs data
c5	Running		Ready		Sleep	0	Oops! Woke T _{c2}

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.19

EXECUTION TRACE – 2
WHILE, 1 CONDITION, 1 PRODUCER, 2 CONSUMERS

T_{c2} runs, no data to consume

Legend

c1/p1- lock
c2/p2- check var
c3/p3- wait
c4- put()
p4- get()
c5/p5- signal
c6/p6- unlock

T _{c1}	State	T _{c2}	State	T _p	State	Count	Comment
...	(cont.)
c6	Running		Ready		Sleep	0	
c1	Running		Ready		Sleep	0	
c2	Running		Ready		Sleep	0	
c3	Sleep		Ready		Sleep	0	Nothing to get
	Sleep	c2	Running		Sleep	0	
	Sleep	c3	Sleep		Sleep	0	Everyone asleep ...

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.20

TWO CONDITIONS

Use two condition variables: empty & full

One condition handles the producer

the other the consumer

```
1 void *producer(void *arg) {  
2     mutex_t mutex;  
3  
4     void *producer(void *arg) {  
5         int i;  
6         for (i = 0; i < loops; i++) {  
7             pthread_mutex_lock(&mutex);  
8             while (count == 1)  
9                 pthread_cond_wait(&empty, &mutex);  
10            put(i);  
11            pthread_cond_signal(&full);  
12            pthread_mutex_unlock(&mutex);  
13        }  
14    }  
15 }
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.21

FINAL PRODUCER/CONSUMER

Change buffer from int, to int buffer[MAX]

Add indexing variables

```
1 int buffer[MAX];  
2 int fill = 0;  
3 int use = 0;  
4 int count = 0;  
5  
6 void put(int value) {  
7     buffer[fill] = value;  
8     fill = (fill + 1) % MAX;  
9     count++;  
10 }  
11  
12 int get() {  
13     int tmp = buffer[use];  
14     use = (use + 1) % MAX;  
15     count--;  
16     return tmp;  
17 }
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.22

FINAL P/C - 2

```
1 void *producer(void *arg) {  
2     mutex_t mutex;  
3  
4     void *producer(void *arg) {  
5         int i;  
6         for (i = 0; i < loops; i++) {  
7             pthread_mutex_lock(&mutex);  
8             while (count == MAX)  
9                 pthread_cond_wait(&empty, &mutex);  
10            put(i);  
11            pthread_cond_signal(&full);  
12            pthread_mutex_unlock(&mutex);  
13        }  
14    }  
15  
16 void *consumer(void *arg) {  
17     int i;  
18     for (i = 0; i < loops; i++) {  
19         pthread_mutex_lock(&mutex);  
20         while (count == 0)  
21             pthread_cond_wait(&full, &mutex);  
22         int tmp = get();  
23     }  
24 }
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.23

FINAL P/C - 3

Producer: only sleeps when buffer is full

Consumer: only sleeps if buffers are empty

```
(Cont.)  
23 pthread_cond_signal(&empty);  
24 pthread_mutex_unlock(&mutex);  
25 printf("%d\n", tmp);  
26 }  
27 }
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.24

COVERING CONDITIONS

- A condition that covers **all** cases (conditions):
- Excellent use case for `pthread_cond_broadcast`
- Consider memory allocation:
 - When a program deals with huge memory allocation/deallocation on the heap
 - Access to the heap must be managed when memory is scarce

PREVENT: Out of memory:
- queue requests until memory is free

- Which thread should be woken up?

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.25

COVERING CONDITIONS - 2

```
1 // how many bytes of the heap are free?
2 int bytesLeft = MAX_HEAP_SIZE;
3
4 // need lock and condition too
5 cond_t c;
6 mutex_t m;
7
8 void *
9 allocate(int size) {
10     pthread_mutex_lock(&m);
11     while (bytesLeft < size)
12         pthread_cond_wait(&c, &m);
13     void *ptr = ...; // get mem from heap
14     bytesLeft -= size;
15     pthread_mutex_unlock(&m);
16     return ptr;
17 }
18
19 void free(void *ptr, int size) {
20     pthread_mutex_lock(&m);
21     bytesLeft += size;
22     pthread_cond_signal(&c);
23     pthread_mutex_unlock(&m);
24 }
```

Check available memory

Broadcast

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.26


COVER CONDITIONS - 3

- Broadcast awakens all blocked threads requesting memory
- Each thread evaluates if there's enough memory: (`bytesLeft < size`)
 - Reject: requests that cannot be fulfilled- go back to sleep
 - Insufficient memory
 - Run: requests which **can** be fulfilled
 - with newly available memory!
- **Overhead**
 - Many threads may be awoken which can't execute

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.27



CHAPTER 32 –
CONCURRENCY
PROBLEMS

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.28

OBJECTIVES

- Chapter 32:
 - Non-deadlock concurrency bugs
 - Deadlock causes
 - Deadlock prevention

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.29

CONCURRENCY BUGS IN
OPEN SOURCE SOFTWARE

- “Learning from Mistakes – A Comprehensive Study on Real World Concurrency Bug Characteristics”
 - Shan Lu et al.
 - Architectural Support For Programming Languages and Operating Systems (ASPLOS 2008), Seattle WA

Application	What it does	Non-Deadlock	Deadlock
MySQL	Database Server	14	9
Apache	Web Server	13	4
Mozilla	Web Browser	41	16
Open Office	Office Suite	6	2
Total		74	31

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.30

NON-DEADLOCK BUGS

- Majority of concurrency bugs
- Most common:
 - Atomicity violation: forget to use locks
 - Order violation: failure to initialize lock/condition before use

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.31

ATOMICITY VIOLATION - MYSQL

- Two threads access the `proc_info` field in struct `thd`
- `NULL` is 0 in C
- Serialized access to shared memory among separate threads is not enforced (e.g. non-atomic)
- Simple example:

Programmer intended
variable to be accessed
atomically...

```
1 Thread1::
2 if(thd->proc_info){
3     ...
4     fputs(thd->proc_info, ...);
5     ...
6 }
7
8 Thread2::
9 thd->proc_info = NULL;
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.32

ATOMICITY VIOLATION - SOLUTION

- Add locks for all uses of: `thd->proc_info`

```
1 pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;
2
3 Thread1::
4 pthread_mutex_lock(&lock);
5 if(thd->proc_info){
6     ...
7     fputs(thd->proc_info, ...);
8     ...
9 }
10 pthread_mutex_unlock(&lock);
11
12 Thread2::
13 pthread_mutex_lock(&lock);
14 thd->proc_info = NULL;
15 pthread_mutex_unlock(&lock);
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.33

ORDER VIOLATION BUGS

- Desired order between memory accesses is flipped
- E.g. something is checked before it is set
- Example:

```
1 Thread1::
2 void init(){
3     mThread = PR_CreateThread(mMain, ...);
4 }
5
6 Thread2::
7 void mMain(...){
8     mState = mThread->State
9 }
```

- What if `mThread` is not initialized?

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.34

ORDER VIOLATION - SOLUTION

- Use condition variable to enforce order

```
1 pthread_mutex_t mtLock = PTHREAD_MUTEX_INITIALIZER;
2 pthread_cond_t mtCond = PTHREAD_COND_INITIALIZER;
3 int mInit = 0;
4
5 Thread 1::
6 void init(){
7     ...
8     mThread = PR_CreateThread(mMain, ...);
9
10    // signal that the thread has been created.
11    pthread_mutex_lock(&mtLock);
12    mInit = 1;
13    pthread_cond_signal(&mtCond);
14    pthread_mutex_unlock(&mtLock);
15    ...
16 }
17
18 Thread2::
19 void mMain(...){
20    ...
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.35

ORDER VIOLATION - SOLUTION 2

```
21 // wait for the thread to be initialized ...
22 pthread_mutex_lock(&mtLock);
23 while(mInit == 0)
24     pthread_cond_wait(&mtCond, &mtLock);
25 pthread_mutex_unlock(&mtLock);
26
27 mState = mThread->State;
28 ...
29 }
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.36

NON-DEADLOCK BUGS - 1

- 97% of Non-Deadlock Bugs were
 - Atomicity
 - Order violations
- Consider what is involved in “spotting” these bugs in code
- Desire for automated tool support (IDE)

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.37

NON-DEADLOCK BUGS - 2

- Atomicity
 - How can we tell if a given variable is shared?
 - Can search the code for uses
 - How do we know if all instances of its use are shared?
 - Can some non-synchronized (non-atomic) uses be legal?
 - Before threads are created, after threads exit
 - Must verify the scope
- Order violation
 - Must consider all variable accesses
 - Must known desired order

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.38

DEADLOCK BUGS

- Presence of a cycle in code
- Thread 1 acquires lock L1, waits for lock L2
- Thread 2 acquires lock L2, waits for lock L1

Thread 1:
lock (L1);
lock (L2);

Thread 2:
lock (L2);
lock (L1);

- Both threads can block, unless one manages to acquire both locks

```
graph LR
    T1((Thread 1)) -- Holds --> L1[Lock L1]
    L1 -- "Wanted by" --> T2((Thread 2))
    T2 -- Holds --> L2[Lock L2]
    L2 -- "Wanted by" --> T1
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.39

REASONS FOR DEADLOCKS

- Complex code
 - Must avoid circular dependencies – can be hard to find...
- Encapsulation hides potential locking conflicts
 - Easy-to-use APIs embed locks inside
 - Programmer doesn't know they are there
 - Consider the Java Vector class:

```
1 Vector v1,v2;
2 v1.AddAll(v2);
```

- Vector is thread safe (synchronized) by design
- If there is a v2.AddAll(v1); call at nearly the same time deadlock could result

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.40

CONDITIONS FOR DEADLOCK

- Four conditions are required for dead lock to occur

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
Hold-and-wait	Threads hold resources allocated to them while waiting for additional resources
No preemption	Resources cannot be forcibly removed from threads that are holding them.
Circular wait	There exists a circular chain of threads such that each thread holds one more resources that are being requested by the next thread in the chain

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.41

PREVENTION – MUTUAL EXCLUSION

- Build wait-free data structures
 - Eliminate locks altogether
 - Build structures using CompareAndSwap atomic CPU (HW) instruction
- C pseudo code for CompareAndSwap
- Hardware executes this code atomically

```
1 int CompareAndSwap(int *address, int expected, int new){
2     if(*address == expected){
3         *address = new;
4         return 1; // success
5     }
6     return 0;
7 }
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.42

PREVENTION – MUTUAL EXCLUSION - 2

■ Recall atomic increment

```
1 void AtomicIncrement(int *value, int amount){
2     do{
3         int old = *value;
4     }while( CompareAndSwap(value, old, old+amount)==0);
5 }
```

- Compare and Swap tries over and over until successful
- CompareAndSwap is guaranteed to be atomic
- When it runs it is **ALWAYS** atomic (at HW level)

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.43

MUTUAL EXCLUSION: LIST INSERTION

■ Consider list insertion

```
1 void insert(int value){
2     node_t * n = malloc(sizeof(node_t));
3     assert( n != NULL );
4     n->value = value ;
5     n->next = head;
6     head = n;
7 }
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.44

MUTUAL EXCLUSION – LIST INSERTION - 2

■ Lock based implementation

```
1 void insert(int value){
2     node_t * n = malloc(sizeof(node_t));
3     assert( n != NULL );
4     n->value = value ;
5     lock(listlock); // begin critical section
6     n->next = head;
7     head = n;
8     unlock(listlock); //end critical section
9 }
```

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.45

MUTUAL EXCLUSION – LIST INSERTION - 3

■ Wait free (no lock) implementation

```
1 void insert(int value) {
2     node_t *n = malloc(sizeof(node_t));
3     assert( n != NULL );
4     n->value = value;
5     do {
6         n->next = head;
7     } while (!CompareAndSwap(&head, n->next, n));
8 }
```

- Assign &head to n (new node ptr)
- Only when head = n->next

May 7, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L10.46

QUESTIONS

