


# TCSS 422: OPERATING SYSTEMS

## Locks, Lock-Based Data Structures, Condition Variables

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

April 28, 2020

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



## OBJECTIVES – 4/28

- **Questions from 4/23**
- **C Tutorial (Apr 30 11:59p AOE)**
- **Assignment 1 (May 7 11:59p AOE)**
- **Chapter 28: Locks**
  - Introduction, Lock Granularity
  - Spin Locks, Test and Set, Compare and Swap
- **Chapter 29: Lock Based Data Structures**
  - Sloppy Counter
  - Concurrent Structures: Linked List, Queue, Hash Table
- **Chapter 30: Condition Variables**
  - Producer/Consumer
  - Covering Conditions

April 28, 2020	TCSS422: Operating Systems [Spring 2020] School of Engineering and Technology, University of Washington - Tacoma	L9.2
----------------	---	------

## MATERIAL / PACE

- Please classify your perspective on material covered in today's class (49 respondents):
  - 1-mostly review, 5-equal new/review, 10-mostly new
  - **Average – 7.21 (↓ from 7.32)**
- Please rate the pace of today's class:
  - 1-slow, 5-just right, 10-fast
  - **Average – 5.53 (↓ from 5.63)**

April 28, 2020

TCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L9.3

## FEEDBACK FROM 4/23

- **Is the `pthread_cond_t` data type a queue in itself, OR does the system maintain a queue behind the scenes when an instance of the type is initialized?**
- The system maintains a FIFO queue behind the scenes for each condition variable.
- Any thread that waits on the condition adds itself to the queue
- When a signal is fired on the condition variable, threads are woken up in FIFO order

April 28, 2020

TCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L9.4

## FEEDBACK - 2

- Is one instance of the condition data type required per thread requiring access, or is only one instance required per critical region?
- Only one instance of a condition variable is required for each critical region
- Condition variables are associated with a lock variable
- The lock variable protects the critical section
- The condition variable enforces FIFO access to the critical section

April 28, 2020

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

L9.5

## CONDITION VARIABLE EXAMPLE

- Wait example:

```
pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;
pthread_cond_t cond = PTHREAD_COND_INITIALIZER;

pthread_mutex_lock(&lock);
while (initialized == 0)
    pthread_cond_wait(&cond, &lock);
// Perform work that requires lock
a = a + b;
pthread_mutex_unlock(&lock);
```

- wait puts thread to sleep, and automatically releases lock
- meanwhile: another thread sets the state

```
pthread_mutex_lock(&lock);
initialized = 1;
pthread_cond_signal(&cond);
pthread_mutex_unlock(&lock);
```

State variable set,  
Enables other thread(s)  
to proceed above.

some other code  
in the program

- when awoken, lock is reacquired, state variable passes test, we can then execute  $a=a+b$  !!!

April 28, 2020

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

L9.6

## CONDITION VARIABLES

- Are usually associated with a primitive state variable (int or Boolean)
  - Variable tracks if it is okay to proceed:
  - **\*\* Are preconditions for execution met? \*\***
- Introduced in Chapter 27 (Thread API)
- Covered in depth in Chapter 30

April 28, 2020

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

L9.7

## OBJECTIVES – 4/28

- Questions from 4/23
- **C Tutorial (Apr 30 11:59p AOE)**
- Assignment 1 (May 7 11:59p AOE)
- Chapter 28: Locks
  - Introduction, Lock Granularity
  - Spin Locks, Test and Set, Compare and Swap
- Chapter 29: Lock Based Data Structures
  - Sloppy Counter
  - Concurrent Structures: Linked List, Queue, Hash Table
- Chapter 30: Condition Variables
  - Producer/Consumer
  - Covering Conditions

April 28, 2020

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

L9.8

## OBJECTIVES – 4/28

- Questions from 4/23
- C Tutorial (Apr 30 11:59p AOE)
- Assignment 1 (May 7 11:59p AOE)
- Chapter 28: Locks
  - Introduction, Lock Granularity
  - Spin Locks, Test and Set, Compare and Swap
- Chapter 29: Lock Based Data Structures
  - Sloppy Counter
  - Concurrent Structures: Linked List, Queue, Hash Table
- Chapter 30: Condition Variables
  - Producer/Consumer
  - Covering Conditions

April 28, 2020

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

L9.9

## MIDTERM


- Tuesday May 5<sup>th</sup>
- 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 28, 2020

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

L9.10

# CHAPTER 28 – LOCKS



April 28, 2020

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

L9.11

## OBJECTIVES – 4/28

- Questions from 4/23
- C Tutorial (Apr 30 11:59p AOE)
- Assignment 1 (May 7 11:59p AOE)
- **Chapter 28: Locks**
  - Introduction, Lock Granularity
  - Spin Locks, Test and Set, Compare and Swap
- **Chapter 29: Lock Based Data Structures**
  - Sloppy Counter
  - Concurrent Structures: Linked List, Queue, Hash Table
- **Chapter 30: Condition Variables**
  - Producer/Consumer
  - Covering Conditions

April 28, 2020

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

L9.12

# LOCKS



- Ensure critical section(s) are executed atomically-as a *unit*
  - Only one thread is allowed to execute a critical section at any given time
  - Ensures the code snippets are “mutually exclusive”

- Protect a global counter:

```
balance = balance + 1;
```

- A “critical section”:

```
1 lock_t mutex; // some globally-allocated lock 'mutex'
2 ...
3 lock(&mutex);
4 balance = balance + 1;
5 unlock(&mutex);
```

April 28, 2020

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

L9.13

## LOCKS - 2

- Lock variables are called “MUTEX”
  - Short for mutual exclusion (that’s what they guarantee)
- Lock variables store the state of the lock
- States
  - **Locked** (acquired or held)
  - **Unlocked** (available or free)
- Only 1 thread can hold a lock

April 28, 2020

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

L9.14

## LOCKS - 3

- `pthread_mutex_lock(&lock)`
  - Try to acquire lock
  - If lock is free, calling thread will acquire the lock
  - Thread with lock enters critical section
    - Thread “owns” the lock
- No other thread can acquire the lock before the owner releases it.

April 28, 2020

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

L9.15

## LOCKS - 4

- Program can have many mutex (lock) variables to “serialize” many critical sections
- Locks are also used to protect data structures
  - Prevent multiple threads from changing the same data simultaneously
  - Programmer can make sections of code “granular”
    - Fine grained – means just one grain of sand at a time through an hour glass
  - Similar to relational database transactions
    - DB transactions prevent multiple users from modifying a table, row, field

April 28, 2020

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


L9.16



## FINE GRAINED?

### ■ *Is this code a good example of “fine grained parallelism”?*

```
pthread_mutex_lock(&lock);
a = b++;
b = a * c;
*d = a + b + c;
FILE * fp = fopen ("file.txt", "r");
fscanf(fp, "%s %s %s %d", str1, str2, str3, &e);
ListNode *node = mylist->head;
int i=0
while (node) {
    node->title = str1;
    node->subheading = str2;
    node->desc = str3;
    node->end = *e;
    node = node->next;
    i++
}
e = e - i;
pthread_mutex_unlock(&lock);
```



April 28, 2020

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

L9.17

## FINE GRAINED PARALLELISM

```
pthread_mutex_lock(&lock_a);
pthread_mutex_lock(&lock_b);
a = b++;
pthread_mutex_unlock(&lock_b);
pthread_mutex_unlock(&lock_a);

pthread_mutex_lock(&lock_b);
b = a * c;
pthread_mutex_unlock(&lock_b);

pthread_mutex_lock(&lock_d);
*d = a + b + c;
pthread_mutex_unlock(&lock_d);

FILE * fp = fopen ("file.txt", "r");
pthread_mutex_lock(&lock_e);
fscanf(fp, "%s %s %s %d", str1, str2, str3, &e);
pthread_mutex_unlock(&lock_e);

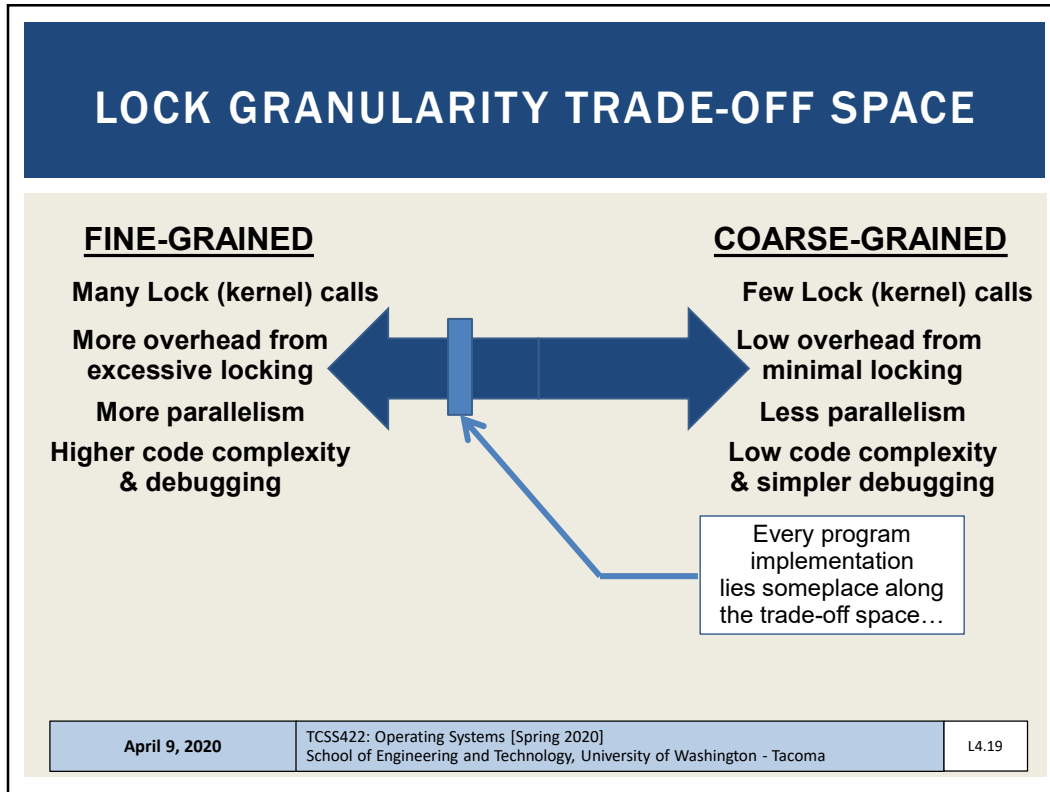
ListNode *node = mylist->head;
int i=0 . . .
```



April 28, 2020

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


L9.18



## EVALUATING LOCK IMPLEMENTATIONS

- **Correctness**
  - Does the lock work?
  - Are critical sections mutually exclusive? (*atomic-as a unit?*)
  
- **Fairness**
  - Do all threads that compete for a lock have a fair chance of acquiring it?
  
- **Overhead**

What makes a good lock?



April 28, 2020	TCSS422: Operating Systems [Spring 2020] School of Engineering and Technology, University of Washington - Tacoma	L9.20
----------------	---	-------

## BUILDING LOCKS

- Locks require hardware support
  - To minimize overhead, ensure fairness and correctness
  - Special “atomic-as a unit” instructions to support lock implementation
  - Atomic-as a unit exchange instruction
    - XCHG
  - Compare and exchange instruction
    - CMPXCHG
    - CMPXCHG8B
    - CMPXCHG16B

April 28, 2020

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

L9.21

## HISTORICAL IMPLEMENTATION

- To implement mutual exclusion
  - Disable interrupts upon entering critical sections

```
1  void lock() {  
2      DisableInterrupts();  
3  }  
4  void unlock() {  
5      EnableInterrupts();  
6  }
```

- Any thread could disable system-wide interrupt
  - What if lock is never released?
- On a multiprocessor processor each CPU has its own interrupts
  - Do we disable interrupts for all cores simultaneously?
- While interrupts are disabled, they could be lost
  - If not queued...

April 28, 2020

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

L9.22

## SPIN LOCK IMPLEMENTATION

- Operate without atomic-as a *unit* assembly instructions
- “Do-it-yourself” Locks
- Is this lock implementation: (1)Correct? (2)Fair? (3)Performant?



```

1  typedef struct __lock_t { int flag; } lock_t;
2
3  void init(lock_t *mutex) {
4      // 0 → lock is available, 1 → held
5      mutex->flag = 0;
6  }
7
8  void lock(lock_t *mutex) {
9      while (mutex->flag == 1) // TEST the flag
10         ; // spin-wait (do nothing)
11     mutex->flag = 1; // now SET it !
12 }
13
14 void unlock(lock_t *mutex) {
15     mutex->flag = 0;
16 }
    
```

April 28, 2020

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

L9.23

## DIY: CORRECT?

- Correctness requires luck... (e.g. *DIY lock is incorrect*)

Thread1	Thread2
call lock() while (flag == 1) interrupt: switch to Thread 2  flag = 1; // set flag to 1 (too!)	call lock() while (flag == 1) flag = 1; interrupt: switch to Thread 1

- Here both threads have “acquired” the lock simultaneously

April 28, 2020

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

L9.24

## DIY: PERFORMANT?

```
void lock(lock_t *mutex)
{
    while (mutex->flag == 1);    // while lock is unavailable, wait...
    mutex->flag = 1;
}
```

- What is wrong with while(<cond>); ?
- Spin-waiting wastes time actively waiting for another thread
- while (1); will “peg” a CPU core at 100%
  - Continuously loops, and evaluates mutex->flag value...
  - Generates heat...

April 28, 2020

TCS5422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L9.25

## TEST-AND-SET INSTRUCTION

- Hardware support required for working locks
- Book presents pseudo code of C implementation
  - TEST-and-SET adds a simple check to the basic spin lock
  - Assumption is this “C code” runs atomically on CPU:

```
1  int TestAndSet(int *ptr, int new) {
2      int old = *ptr;    // fetch old value at ptr
3      *ptr = new;        // store 'new' into ptr
4      return old;        // return the old value
5  }
```

- lock() method checks that TestAndSet doesn't return 1
- Comparison is in the caller
- Can implement the C version (non-atomic) and have some success on a single-core VM

April 28, 2020

TCS5422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L9.26

## DIY: TEST-AND-SET - 2

- C version: requires preemptive scheduler on single core system
- Lock is never released without a context switch
- single-core VM: occasionally will deadlock, doesn't miscount

```
1  typedef struct __lock_t {
2      int flag;
3  } lock_t;
4
5  void init(lock_t *lock) {
6      // 0 indicates that lock is available,
7      // 1 that it is held
8      lock->flag = 0;
9  }
10
11 void lock(lock_t *lock) {
12     while (TestAndSet(&lock->flag, 1) == 1)
13         ; // spin-wait
14 }
15
16 void unlock(lock_t *lock) {
17     lock->flag = 0;
18 }
```

April 28, 2020

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

L9.27

## SPIN LOCK EVALUATION

- **Correctness:**
  - Spin locks with atomic Test-and-Set:  
Critical sections won't be executed simultaneously by (2) threads
- **Fairness:**
  - No fairness guarantee. Once a thread has a lock, nothing forces it to relinquish it...
- **Performance:**
  - Spin locks perform "busy waiting"
  - Spin locks are best for short periods of waiting (< 1 time quantum)
  - Performance is slow when multiple threads share a CPU
    - Especially if "spinning" for long periods

April 28, 2020

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

L9.28

## COMPARE AND SWAP

- Checks that the lock variable has the expected value **FIRST**, before changing its value
  - If so, make assignment
  - Return value at location
- Adds a comparison to TestAndSet
  - Textbook presents C pseudo code
  - Assumption is that the compare-and-swap method runs atomically
- Useful for wait-free synchronization
  - Supports implementation of shared data structures which can be updated atomically (*as a unit*) using the HW support CompareAndSwap instruction
  - Shared data structure updates become “wait-free”
  - Upcoming in Chapter 32

April 28, 2020

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

L9.29

## COMPARE AND SWAP

- Compare and Swap

```
1  int CompareAndSwap(int *ptr, int expected, int new) {
2      int actual = *ptr;
3      if (actual == expected)
4          *ptr = new;
5      return actual;
6  }
```

- Spin lock

```
1
2
3      ; // spin
4  }
```

**C implementation 1-core VM:  
Count is correct, no deadlock**

- X86 provides “**cmpxchg1**” compare-and-exchange instruction
  - **cmpxchg8b**
  - **cmpxchg16b**

April 28, 2020

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

L9.30

## When implementing locks in a high-level language (e.g. C), what is missing that prevents implementation of CORRECT locks?

Shared state variable

Condition variables

ATOMIC instructions

Fairness

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](https://pollev.com/app)

## TWO MORE “LOCK BUILDING” CPU INSTRUCTIONS

- Cooperative instructions used together to support synchronization on RISC systems
- No support on x86 processors
  - Supported by RISC: Alpha, PowerPC, ARM
- Load-linked (LL)
  - Loads value into register
  - Same as typical load
  - Used as a mechanism to track competition
- Store-conditional (SC)
  - Performs “mutually exclusive” store
  - Allows only one thread to store value

April 28, 2020

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

L9.32



## LL/SC LOCK

```
1  int LoadLinked(int *ptr) {
2      return *ptr;
3  }
4
5  int StoreConditional(int *ptr, int value) {
6      if (no one has updated *ptr since the LoadLinked to this address) {
7          *ptr = value;
8          return 1; // success!
9      } else {
10         return 0; // failed to update
11     }
12 }
```

- LL instruction loads pointer value (ptr)
- SC only stores if the load link pointer has not changed
- Requires HW support
  - C code is psuedo code

April 28, 2020

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

L9.33

## LL/SC LOCK - 2

```
1  void lock(lock_t *lock) {
2      while (1) {
3          while (LoadLinked(&lock->flag) == 1)
4              ; // spin until it's zero
5          if (StoreConditional(&lock->flag, 1) == 1)
6              return; // if set-it-to-1 was a success: all done
7                      // otherwise: try it all over again
8      }
9  }
10
11 void unlock(lock_t *lock) {
12     lock->flag = 0;
13 }
```

- Two instruction lock

April 28, 2020


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

L9.34

# TCSS 422 WILL RETURN AT ~2:40PM

April 28, 2020

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




L9.35

# CHAPTER 29 – LOCK BASED DATA STRUCTURES

April 28, 2020

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



L9.36

## OBJECTIVES – 4/28

- Questions from 4/23
- C Tutorial (Apr 30 11:59p AOE)
- Assignment 1 (May 7 11:59p AOE)
- Chapter 28: Locks
  - Introduction, Lock Granularity
  - Spin Locks, Test and Set, Compare and Swap
- Chapter 29: Lock Based Data Structures
  - Sloppy Counter
  - Concurrent Structures: Linked List, Queue, Hash Table
- Chapter 30: Condition Variables
  - Producer/Consumer
  - Covering Conditions

April 28, 2020

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

L9.37

## LOCK-BASED CONCURRENT DATA STRUCTURES

- Adding locks to data structures make them **thread safe**.
- Considerations:
  - Correctness
  - Performance
  - Lock granularity

April 28, 2020

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

L9.38

## COUNTER STRUCTURE W/O LOCK

### ■ Synchronization weary --- not thread safe

```
1  typedef struct __counter_t {
2      int value;
3  } counter_t;
4
5  void init(counter_t *c) {
6      c->value = 0;
7  }
8
9  void increment(counter_t *c) {
10     c->value++;
11 }
12
13 void decrement(counter_t *c) {
14     c->value--;
15 }
16
17 int get(counter_t *c) {
18     return c->value;
19 }
```

April 28, 2020

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

L9.39

## CONCURRENT COUNTER

```
1  typedef struct __counter_t {
2      int value;
3      pthread_lock_t lock;
4  } counter_t;
5
6  void init(counter_t *c) {
7      c->value = 0;
8      Pthread_mutex_init(&c->lock, NULL);
9  }
10
11 void increment(counter_t *c) {
12     Pthread_mutex_lock(&c->lock);
13     c->value++;
14     Pthread_mutex_unlock(&c->lock);
15 }
16
```

- Add lock to the counter
- Require lock to change data

April 28, 2020

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

L9.40

## CONCURRENT COUNTER - 2

- Decrease counter
- Get value

```
(Cont.)
17  void decrement(counter_t *c) {
18      pthread_mutex_lock(&c->lock);
19      c->value--;
20      pthread_mutex_unlock(&c->lock);
21  }
22
23  int get(counter_t *c) {
24      pthread_mutex_lock(&c->lock);
25      int rc = c->value;
26      pthread_mutex_unlock(&c->lock);
27      return rc;
28  }
```


April 28, 2020

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

L9.41

## CONCURRENT COUNTERS - PERFORMANCE

- iMac: four core Intel 2.7 GHz i5 CPU
- Each thread increments counter 1,000,000 times



Traditional vs. sloppy counter  
Sloppy Threshold (S) = 1024

**Synchronized counter scales poorly.**

April 28, 2020

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

L9.42

## PERFECT SCALING

- Achieve (N) performance gain with (N) additional resources
- Throughput:
  - Transactions per second (tps)
- 1 core
  - N = 100 tps
- 10 cores (x10)
  - N = 1000 tps (x10)

April 28, 2020

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

L9.43

## SLOPPY COUNTER

- Provides single logical shared counter
  - Implemented using 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
- Why this implementation?  
Why do we want counters local to each CPU Core?

April 28, 2020

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

L9.44

## SLOPPY COUNTER – MAIN POINTS

- Idea of Sloppy Counter is to RELAX the 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 28, 2020

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

L9.45

## SLOPPY COUNTER - 2

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

Time	L <sub>1</sub>	L <sub>2</sub>	L <sub>3</sub>	L <sub>4</sub>	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 → 0	1	3	4	5 (from L <sub>1</sub> )
7	0	2	4	5 → 0	10 (from L <sub>4</sub> )

April 28, 2020

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

L9.46

THRESHOLD VALUE S

- Consider 4 threads increment a counter 1000000 times each
- Low S → What is the consequence?
- High S → 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 28, 2020

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

L9.47

SLOPPY COUNTER - EXAMPLE

- Example implementation
- Also with CPU affinity

April 28, 2020

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

L9.48



## 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 28, 2020

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

L9.49

## 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 28, 2020

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

L9.50

## 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 28, 2020

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

L9.51

## 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 28, 2020

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

L9.52

## 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 28, 2020

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

L9.53

## 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 28, 2020

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

L9.54

## 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 28, 2020

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

L9.55

## 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 28, 2020

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

L9.56

## 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 28, 2020

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

L9.57

## 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 28, 2020

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

L9.58

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 28, 2020

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

L9.59

INSERT PERFORMANCE –  
CONCURRENT HASH TABLE

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

The graph plots Time (seconds) on the y-axis (0 to 15) against Inserts (Thousands) on the x-axis (0 to 40). The Simple Concurrent List (open circles) shows a steep upward curve, while the Concurrent Hash Table (crosses) shows a nearly flat line.

Inserts (Thousands)	Simple Concurrent List (seconds)	Concurrent Hash Table (seconds)
10	~1.0	~0.2
20	~2.5	~0.2
30	~4.5	~0.2
40	~7.5	~0.2
50	~11.5	~0.2

The simple concurrent hash table scales magnificently.

April 28, 2020

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

L9.60

## 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 28, 2020

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

L9.61

### 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](https://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 28, 2020

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

L9.63

## CHAPTER 30 – CONDITION VARIABLES



April 28, 2020

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

L9.64



## OBJECTIVES – 4/28

- Questions from 4/23
- C Tutorial (Apr 30 11:59p AOE)
- Assignment 1 (May 7 11:59p AOE)
- Chapter 28: Locks
  - Introduction, Lock Granularity
  - Spin Locks, Test and Set, Compare and Swap
- Chapter 29: Lock Based Data Structures
  - Sloppy Counter
  - Concurrent Structures: Linked List, Queue, Hash Table
- Chapter 30: Condition Variables
  - Producer/Consumer
  - Covering Conditions

April 28, 2020

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

L9.65

## CONDITION VARIABLES

- There are many cases where a thread wants to wait for another thread before proceeding with execution
- Consider when a precondition must be fulfilled before it is meaningful to proceed ...

April 28, 2020

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

L9.66

## CONDITION VARIABLES - 2



- 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)

April 28, 2020

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

L9.67

## 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

April 28, 2020

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

L9.68

## CONDITION VARIABLES - QUESTIONS

- Why would we want to put waiting threads on a queue... why not use a stack?
  - Queue (FIFO), Stack (LIFO)
  - Using condition variables eliminates busy waiting by putting threads to “sleep” and yielding the CPU.
- Why do we want to not busily wait for the lock to become available?
- A program has 10-threads, where 9 threads are waiting. The working thread finishes and broadcasts that the lock is available. What happens next?

April 28, 2020

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

L9.69

## MATRIX GENERATOR

Matrix generation example

Chapter 30  
signal.c

April 28, 2020

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

L9.70

## MATRIX GENERATOR

- The main thread, and worker thread (generates matrices) share a single matrix pointer.
- What would happen if we don't use a condition variable to coordinate exchange of the lock?
- Let's try "nosignal.c"

April 28, 2020

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

L9.71

## 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

April 28, 2020

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

L9.72

## PRODUCER / CONSUMER



April 28, 2020

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

L9.73

## 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

April 28, 2020

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

L9.74

## 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

April 28, 2020

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

L9.75

## 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 }
```

April 28, 2020

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

L9.76

## PRODUCER / CONSUMER - 3

- Producer adds data
- Consumer removes data (busy waiting)
- Will this code work (spin 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 }
```

April 28, 2020

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

L9.77

## 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);
8          if (count == 1)
9              pthread_cond_wait(&cond, &mutex);
10         put(i);
11         pthread_cond_signal(&cond);
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     }
21 }
```

Producer

// p1  
// p2  
// p3  
// p4  
// p5  
// p6

// c1

April 28, 2020

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

L9.78

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 }
```

Consumer

■ 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 ?

April 28, 2020

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

L9.79

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<sub>c1</sub>

T<sub>c2</sub>

T<sub>p</sub>

T <sub>c1</sub>	State	T <sub>c2</sub>	State	T <sub>p</sub>	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
	Ready		Ready	p5	Running	1	T <sub>c1</sub> awoken
	Ready		Ready	p6	Running	1	
	Ready		Ready	p1	Running	1	
	Ready		Ready	p2	Running	1	
	Ready		Ready	p3	Sleep	1	Buffer full; sleep
	Ready	c1	Running		Sleep	1	T <sub>c2</sub> sneaks in ...
	Ready	c2	Running		Sleep	1	
	Ready	c4	Running		Sleep	0	... and grabs data
	Ready	c5	Running		Ready	0	T <sub>p</sub> awoken
	Ready	c6	Running		Ready	0	
c4	Running		Ready		Ready	0	Oh oh! No data

April 28, 2020

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

L9.80



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  $T_{c2}$

April 28, 2020

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

L9.81

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	Sleep		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 $T_{c2}$

April 28, 2020

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

L9.82

EXECUTION TRACE – 2

WHILE, 1 CONDITION, 1 PRODUCER, 2 CONSUMERS

T<sub>c2</sub> 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 <sub>c1</sub>	State	T <sub>c2</sub>	State	T <sub>p</sub>	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 ...

April 28, 2020

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

L9.83

TWO CONDITIONS

Use two condition variables: empty & full

One condition handles the producer

the other the consumer

1 cond t empty, full;

mutex\_t mutex;

void \*producer(void \*arg) {

int i;

for (i = 0; i < loops; i++) {

pthread\_mutex\_lock(&mutex);

while (count == 1)

pthread\_cond\_wait(&empty, &mutex);

put(i);

pthread\_cond\_signal( &full);

pthread\_mutex\_unlock(&mutex);

}

}

April 28, 2020

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

L9.84

## 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 }
```

April 28, 2020

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

L9.85

## FINAL P/C - 2

```
1  cond_t empty, full
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) // p1
9              pthread_cond_wait(&empty, &mutex); // p2
10         put(i); // p3
11         pthread_cond_signal(&full); // p4
12         pthread_mutex_unlock(&mutex); // p5
13     }
14 }
15
16 void *consumer(void *arg) {
17     int i;
18     for (i = 0; i < loops; i++) {
19         pthread_mutex_lock(&mutex); // c1
20         while (count == 0) // c2
21             pthread_cond_wait(&full, &mutex); // c3
22         int tmp = get(); // c4
```

April 28, 2020

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

L9.86

## FINAL P/C - 3

```
(Cont.)
23      pthread_cond_signal(&empty);           // c5
24      pthread_mutex_unlock(&mutex);          // c6
25      printf("%d\n", tmp);
26      }
27      }
```

- **Producer: only sleeps when buffer is full**
- **Consumer: only sleeps if buffers are empty**

April 28, 2020

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

L9.87

## 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?**

April 28, 2020

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

L9.88

## 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); // Broadcast
23     pthread_mutex_unlock(&m);
24 }
```

Check available memory

Broadcast

April 28, 2020

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

L9.89

## 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

April 28, 2020

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

L9.90

# QUESTIONS



WILL RETURN IN A FEW  
MINUTES

