


## TCCS 422: OPERATING SYSTEMS

### Condition Variables

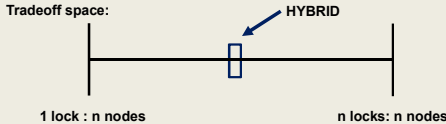
Wes J. Lloyd  
Institute of Technology  
University of Washington - Tacoma



## FEEDBACK – 4/25

- Quiz: How to create a thread safe data struct
  - Quizzes like these are great, I learn a lot!
- How hybrid approach is implemented?
- Hybrid approach: presumably for linked-list locking

Tradeoff space:



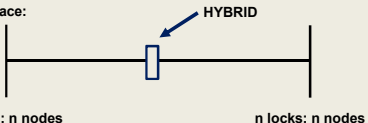
1 lock : n nodes                      n locks: n nodes

April 27, 2017	TCCS422: Operating Systems [Spring 2017] Institute of Technology, University of Washington - Tacoma	L9.2
----------------	--	------

## FEEDBACK - 2

- Consider the tradeoff space: Concurrent Linked List
- Ratio of locks : Nodes

Tradeoff space:



1 lock : n nodes                      n locks: n nodes

- Which design is best for fast list traversal?
- Which design is best for optimal concurrency?
  - Many threads working within the structure at same time
- If we add locks:
 

How does list traversal change?  
 How does concurrency change?

April 27, 2017	TCCS422: Operating Systems [Spring 2017] Institute of Technology, University of Washington - Tacoma	L9.3
----------------	--	------

## FEEDBACK - 3

- Quiz 3 verifications
- Program 2
  - Posted
- Midterm: Thursday May 4 – Primary Coverage:
  - CPU Scheduling (Virtualizing the CPU)
  - Chapters 4, 6, 7, 8, 9
  - Concurrency
  - Chapters 26, 27, 28, 29, 30, 32\*
    - \* - deadlocks: common causes, how to avoid

April 27, 2017	TCCS422: Operating Systems [Spring 2017] Institute of Technology, University of Washington - Tacoma	L9.4
----------------	--	------

## OBJECTIVES

- Sloppy Counter, demo
- Condition variables
- Consumer/Producer
- Covering condition

April 27, 2017	TCCS422: Operating Systems [Spring 2017] Institute of Technology, University of Washington - Tacoma	L9.5
----------------	--	------

## 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 25, 2017	TCCS422: Operating Systems [Spring 2017] Institute of Technology, University of Washington - Tacoma	L8.6
----------------	--	------

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> )

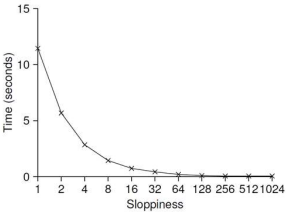
April 25, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L8.7

THRESHOLD VALUE S

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



April 25, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L8.8

SLOPPY COUNTER - EXAMPLE

- Example implementation (sloppybasic.c)
- Also with CPU affinity (sloppy.c)

April 25, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L8.9

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 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.10

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 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.11

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 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.12

## CONDITION VARIABLES - QUESTIONS

- Why would we want to put waiting threads on a queue... why not use a stack?
- Using condition variables eliminates busy waiting by putting a thread "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 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.13

## MATRIX GENERATOR

Matrix generation example

Chapter 30  
signal.c

April 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.14

## 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 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.15

## 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 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.16

## PRODUCER / CONSUMER



April 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.17

## 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 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.18

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 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.19

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 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.20

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 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.21

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
```

April 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.22

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 ?

April 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.23

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		Ready	0	
	c2	Running	Ready	Ready		Ready	0	
	c3	Sleep	Ready	Ready		Ready	0	Nothing to get
		Sleep	Ready	Ready	p1	Running	0	
		Sleep	Ready	Ready	p2	Running	0	
		Ready	Ready	Ready	p5	Running	1	Buffer now full $T_{c1}$ awoken
		Ready	Ready	Ready	p6	Running	1	
		Ready	Ready	Ready	p1	Running	1	
		Ready	Ready	Ready	p2	Running	1	
		Ready	Ready	Ready	p3	Sleep	1	Buffer full; sleep $T_{c2}$ sneaks in ...
	c1	Running	Sleep	Sleep		Sleep	1	
	c2	Running	Sleep	Sleep		Sleep	1	
	c4	Running	Ready	Ready		Ready	0	... and grabs data $T_p$ awoken
	c5	Running	Ready	Ready		Ready	0	
	c6	Running	Ready	Ready		Ready	0	Oh oh! No data

April 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.24

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 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.25

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

Legend

c1/p1- lock  
c2/p2- check var  
c3/p3- wait  
c4- put()  
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 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.26

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()  
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 ...

April 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.27

TWO CONDITIONS

- Use two condition variables: empty & full
  - One condition handles the producer
  - the other the consumer

```
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 == 1)  
9              pthread_cond_wait(&empty, &mutex);  
10         put(i);  
11         pthread_cond_signal(&full);  
12         pthread_mutex_unlock(&mutex);  
13     }  
14  
15 }
```

April 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.28

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 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.29

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)  
9              pthread_cond_wait(&empty, &mutex);  
10         put(i);  
11         pthread_cond_signal(&full);  
12         pthread_mutex_unlock(&mutex);  
13     }  
14  
15 }  
16  
17 void *consumer(void *arg) {  
18     int i;  
19     for (i = 0; i < loops; i++) {  
20         pthread_mutex_lock(&mutex);  
21         while (count == 0)  
22             pthread_cond_wait(&full, &mutex);  
23         int tmp = get(i);  
24     }  
25 }
```

April 27, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.30

FINAL P/C - 3

```
(Cont.)
23 pthread_cond_signal(&cond_v); // 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 27, 2017

TCSS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.31

COVERING CONDITIONS

- A condition that covers **all** cases (conditions):
- Excellent use case for pthread\_cond\_broadcast

- Consider memory allocation:
  - What if 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

April 27, 2017

TCSS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.32

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 }
```

April 27, 2017

TCSS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.33

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 27, 2017

TCSS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.34

QUESTIONS



April 27, 2017

TCSS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L9.35