



## TCSS 422: OPERATING SYSTEMS

### Condition Variables

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



## 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)

October 21, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L9.4

## OBJECTIVES

- Condition variables
- Consumer/Producer
- Covering condition

October 21, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L9.2

## 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

October 21, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L9.5

## 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 ...

October 21, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L9.3

## EXAMPLE: MATRIX GENERATOR

- Thread A continuously generates 2-D matrices
- Thread B computes the average value of 2-D matrices
- Thread B has nothing to do before Thread A generates a matrix
- To simplify memory management, Thread A and Thread B share a pointer to the most recently created matrix
- Thread A can't generate a new array with the shared pointer before Thread B completes calculating an average

October 21, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L9.6

## MATRIX GENERATOR - 2

- Solution ?
  - Thread B could employ a spin lock to continuously check a state variable
- Issues
  - Continuously making a comparison to check a state variable will occupy a CPU core
  - And without hardware support ensuring atomicity in critical sections can not be guaranteed

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.7

## MATRIX GENERATOR - 5

- If the child thread is not created fast enough, or if there is a OS context before the child obtains the lock to generate the array, the program could deadlock...
- The program may still execute (no deadlock), but shared data could be changed out of sequence leading to errors
- Try adding printf statement to observe whether the child thread or parent (int main) is deadlocking...

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.10

## MATRIX GENERATOR - 3

Matrix generation example

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.8

## SUBTLE RACE CONDITION

```

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

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.11

## MATRIX GENERATOR - 4


- Consider the importance of the **state** variable
- What will the code do if we remove it?
- Consider
  - rows = 1000
  - cols = 1000
- Consider
  - rows = 10
  - cols = 10

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.9

## PRODUCER / CONSUMER



October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.12

## 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
  - In our example the 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

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.13

## 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 }
```

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.16

## 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 send output from grep to wc as it is produced
  - File stream

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.14

## 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
```

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.17

## PUT/GET ROUTINES

- Buffer is a one element shared data structure
- 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 }
```

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.15

## 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 ?

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.18

EXECUTION TRACE							
Two threads							
$T_{c1}$	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
	Sleep		Ready	p5	Running	1	$T_{c1}$ awoken
	Ready		Ready	p6	Running	1	
	Ready		Ready	p1	Running	1	
	Ready		Ready	p2	Running	1	
	Ready		Ready	p3	Running	1	Buffer full: sleep
	Ready		Ready	p3	Sleep	1	$T_{c2}$ sneaks in ...
	Ready		Ready	c1	Running	1	
	Ready		Ready	c2	Running	1	
	Ready		Ready	c4	Running	0	... and grabs data
	Ready		Ready	c5	Running	0	$T_p$ awoken
	Ready		Ready	c6	Running	0	
	Ready		Ready		Ready	0	Oh oh! No data

EXECUTION TRACE - 3							
$T_{c2}$ runs, no data to consume							
$T_{c1}$	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		Ready		Sleep	0	
	Sleep		Ready		Sleep	0	Everyone asleep ...

PRODUCER/CONSUMER SYNCHRONIZATION							
<ul style="list-style-type: none"> <li>When producer threads awake, they do not check if there is any data in the buffer...</li> <li>Need while, not if</li> <li>What if <math>T_p</math> puts a value, wakes <math>T_{c1}</math> whom consumes the value</li> <li>Then <math>T_p</math> has a value to put, but <math>T_{c1}</math>'s signal on <math>\&amp;cond</math> wakes <math>T_{c2}</math></li> <li>There is nothing for <math>T_{c2}</math> consume, so <math>T_{c2}</math> sleeps</li> <li><math>T_{c1}</math>, <math>T_{c2}</math>, and <math>T_p</math> all sleep forever</li> <li><math>T_{c1}</math> needs to wake <math>T_p</math> to <math>T_{c2}</math></li> </ul>							

TWO CONDITIONS							
<ul style="list-style-type: none"> <li>Add a second condition</li> <li>One condition handles the producer</li> <li>the other the consumer</li> </ul> <pre> 1  while (t_empty, fill); 2  mutex_t_mutex; 3 4  void *producer(void *arg) { 5      int i; 6      for (i = 0; i &lt; loops; i++) { 7          pthread_mutex_lock(&amp;mutex); 8          while (count == 1) 9              pthread_cond_wait(&amp;empty, &amp;mutex); 10         put(i); 11         pthread_cond_signal(&amp;fill); 12         pthread_mutex_unlock(&amp;mutex); 13     } 14 } 15 </pre>							

EXECUTION TRACE - 2							
$T_{c1}$	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	p1	Running	0	Nothing to get
	Sleep		Sleep	p2	Running	0	
	Sleep		Sleep	p4	Running	1	Buffer now full
	Sleep		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)
	Ready		Sleep		Sleep	1	Recheck condition
	Ready		Sleep		Sleep	0	$T_{c1}$ grabs data
	Ready		Ready		Sleep	0	Oops! Woke $T_{c2}$

FINAL PRODUCER/CONSUMER							
<ul style="list-style-type: none"> <li>Change buffer from int, to int buffer[MAX]</li> <li>Add indexing variables</li> <li>Type: two variables named fill, need separate namespaces</li> </ul> <pre> 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 } </pre>							

## FINAL P/C - 2

```

1 //cond_t empty, fill?
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         while (count == MAX) // p2
9             pthread_cond_wait(&empty, &mutex); // p3
10        put(i); // p4
11        pthread_cond_signal(&fill); // 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
20         while (count == 0) // c2
21             pthread_cond_wait(&fill, &mutex); // c3
22         int tmp = get(i); // c4
23     }
24 }

```

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.25

## 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); // Check available memory
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 }

```

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.28

## 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

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.26

## COVER CONDITIONS - 3

- Broad awakens all blocked threads requesting memory
- Will evaluate each request: (bytesLeft < size)
  - Reject: requests which cannot be fulfilled
    - Insufficient memory
  - Process all requests which **can** be fulfilled
    - with newly available memory
- Overhead
  - Many threads may be awoken which take no action

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.29

## COVERING 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 with low memory resources
- PREVENT: Out of memory - - queue requests until memory is free

October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.27

## QUESTIONS



October 21, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L9.30