


TCSS 422: OPERATING SYSTEMS

Condition Variables, Concurrency Problems



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

November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

FEEDBACK FROM 11/5

- How does `pthread_join()` join thread values?

NAME

`pthread_join` - join with a terminated thread

SYNOPSIS

```
#include <pthread.h>
```

```
int pthread_join(pthread_t thread, void **retval);
```

2nd parameter provides a `void **` pointer

Can return pointer to any user defined struct

November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.2

OBJECTIVES

- Quiz 3 – Synchronized Array
- Multi-threaded Programming
 - Chapter 30 – Condition Variables
 - Chapter 32 – Concurrency Problems
- Memory Virtualization
 - Chapters 13, 14, 15, 16....

November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.3

FORK() COPY ON WRITE

- Processes and Threads share the code segment.
- From: <https://en.wikipedia.org/wiki/Copy-on-write>
- When `fork()` is called, a copy of all parent process pages is created, and loaded into a separate memory location by the OS for the child process.
- But this is not needed in certain cases.
- If a child executes an "exec" call or exits very soon after the `fork()`, there is no need to copy the parent process' pages.
- As an optimization, Linux uses a technique called **copy-on-write (COW)**.

November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.4

COPY ON WRITE - 2

- When the fork() occurs, parent process pages are ***NOT*** copied for the child process.
- Pages are shared between the parent and child.
- When a process (parent or child) modifies a memory page, a separate copy of the page is made for that process (parent or child) which performed the modification.
- This process uses the newly copied page rather than the shared one in future references.
- The other process (the one which did not modify the shared page) continues to use the original copy of the page (which is now no longer shared).
- This technique is called copy-on-write since the page is copied only when some process modifies to it.
- Binary C files are unmodified, with COW they are shared

November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.5



CHAPTER 30 – CONDITION VARIABLES

November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.6

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 ...

November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.7

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)



November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.8

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

November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.9

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?**

November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.10

<h1>MATRIX GENERATOR</h1>		
<p>Matrix generation example</p> <p>Chapter 30 signal.c</p>		
November 7, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L12.11

<h1>MATRIX GENERATOR</h1>		
<ul style="list-style-type: none">▪ 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"		
November 7, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L12.12

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 }
```

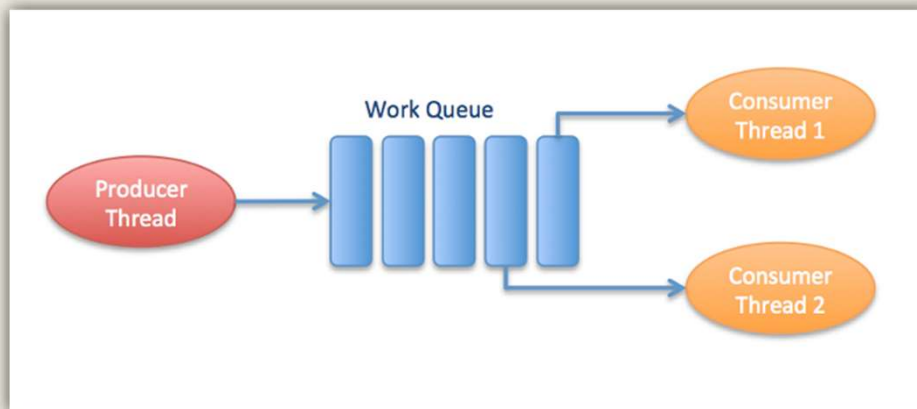
- Parent thread calls `thr_join()` and executes the comparison
- The context switches to the child
- The child runs `thr_exit()` and signals the parent, but the parent is not waiting yet.
- The signal is lost
- The parent deadlocks

November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.13

PRODUCER / CONSUMER



November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.14

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

November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.15

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

November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.16

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 }
```

November 7, 2018

TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.17

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 }
```

November 7, 2018

TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.18

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
```

November 7, 2018

TCSS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma

L12.19

QUESTIONS

