

TCSS 422: OPERATING SYSTEMS

Condition Variables, Producer/Consumer

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

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma



OBJECTIVES

- Assignment 1 – 2/15
- Assignment 2
- Midterm – Postponed until 2/20
- Feedback 2/11
- Practice midterm
- Parallel programming with P-threads cont'd
- Chapter 30 – Condition Variables

February 13, 2019	TCSS422: Operating Systems [Winter 2019] School of Engineering and Technology, University of Washington - Tacoma	L10.2
-------------------	---	-------

FEEDBACK FROM 2/11

- Could you please elaborate more on how pthread wait works?
- What happens to a lock if a thread that currently has the lock is terminated for some reason without releasing it? Does this produce a deadlock as no other thread can access the lock?
- Unclear about concurrent queue, how it works
- The main thing that was unclear was using the head and the tail locks while implementing linked list threads. Even after reviewing the class recording I'm still not sure how it works or why it is necessary to have those 2 locks.
 - Is this the concurrent queue?

February 13, 2019

TCS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.3

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

February 13, 2019

TCS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.4

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.)
```

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.5

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.)
```

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.6

FEEDBACK - 2

- (sloppy counter) You mentioned that threads aren't guaranteed to remain on the same core due to the CPU scheduler. Are there ways to force the scheduler to assign threads to the same core for its entire duration or is this just not possible?
 - Yes, this is called CPU pinning, and there is a C API for this
 - See the Ch. 29 example online of the sloppy counter
- Not clear why we need the sloppy counter
- Not clear when/where (synchronized) counters are used

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.7

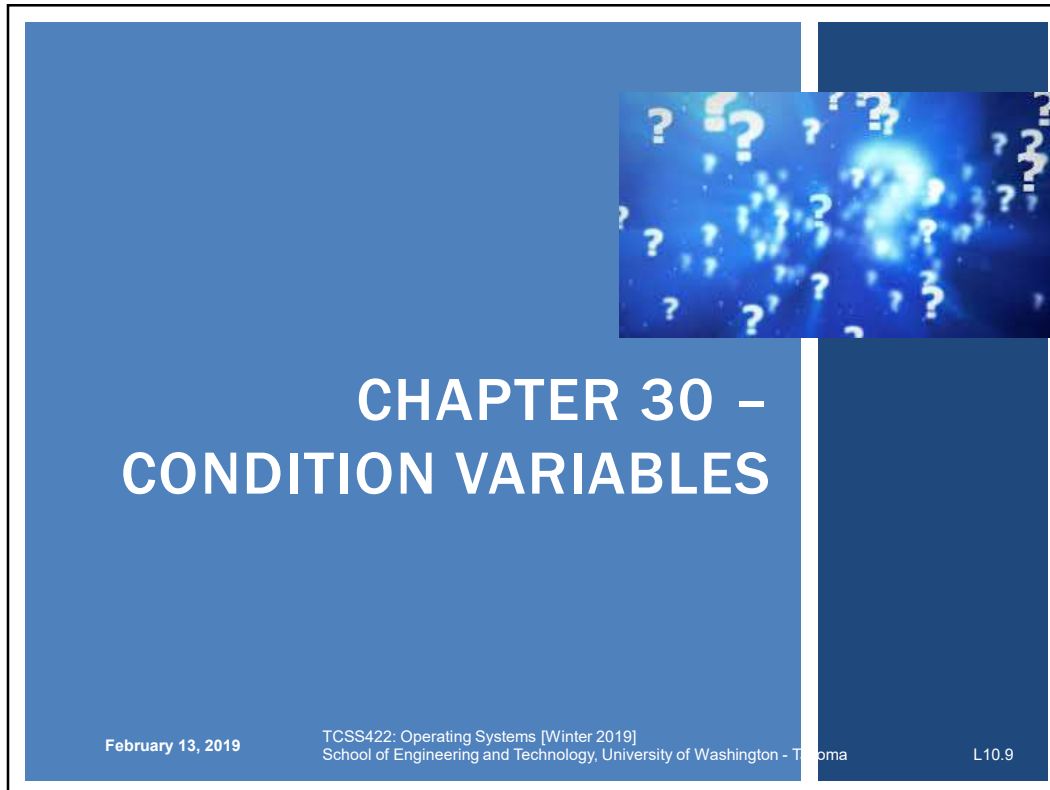
FEEDBACK - 3

- What does it mean to "consume" a matrix?
- Confused about producer/consumer matrix generation, in particular the part of lecture when we were talking about what happens if conditional variables are not used to coordinate exchanges of locks.
 - A clarification on why we need to add locks to the producer/consumer model.
- I'm still confused a bit on how the matrix code works

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.8



CHAPTER 30 – CONDITION VARIABLES

February 13, 2019 TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma L10.9

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 ...

February 13, 2019	TCSS422: Operating Systems [Winter 2019] School of Engineering and Technology, University of Washington - Tacoma	L10.10
-------------------	---	--------

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)

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.11

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

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.12

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?

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.13

MATRIX GENERATOR

Matrix generation example

Chapter 30
signal.c

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.14

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"

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.15

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

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.16

PRODUCER / CONSUMER



February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.17

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

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.18

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

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.19

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 }
```

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.20

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 }
```

February 13, 2019

TCCS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.21

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);
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

February 13, 2019

TCCS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.22

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 }
```

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 ?

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.23

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	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 _{c1} 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 _{c2} sneaks in ...
		Ready	c2	Running		Sleep	1	
		Ready	c4	Running		Sleep	0	... and grabs data
		Ready	c5	Running		Ready	0	T _p awoken
		Ready	c6	Running		Ready	0	
	c4	Running		Ready		Ready	0	Oh oh! No data

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.24

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}

February 13, 2019

TCCS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.25

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}

February 13, 2019

TCCS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.26

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()
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 ...

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.27

TWO CONDITIONS

Use two condition variables: empty & full

One condition handles the producer

the other the consumer

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

→

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);

}

}

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.28

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 }
```

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.29

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);           // p1
8          while (count == MAX)                  // p2
9              pthread_cond_wait(&empty, &mutex); // p3
10         put(i);                                // p4
11         pthread_cond_signal(&full);           // 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(&full, &mutex); // c3
22         int tmp = get();                      // c4
```

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.30

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**

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.31

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?**

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.32

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 }

```

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.33


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

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.34



COMPUTER BOOT SEQUENCE:
OS WITH DIRECT EXECUTION

■ What if programs could directly control the CPU / system?

OS	Program
1. Create entry for process list 2. Allocate memory for program 3. Load program into memory 4. Set up stack with <code>argc / argv</code> 5. Clear registers 6. Execute <code>call main()</code>	7. Run <code>main()</code> 8. Execute <code>return from main()</code>
9. Free memory of process 10. Remove from process list	

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.36

COMPUTER BOOT SEQUENCE:
OS WITH DIRECT EXECUTION

■ What if programs could directly control the CPU / system?

OS	Program
1. Create entry for process list 2. Allocate memory for	
Without <i>limits</i> on running programs, the OS wouldn't be in control of anything and would "just be a library"	
5. Clear registers 6. Execute call <code>main()</code>	7. Run <code>main()</code> 8. Execute return from <code>main()</code>
9. Free memory of process 10. Remove from process list	

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.37

DIRECT EXECUTION - 2

■ With direct execution:

How does the OS stop a program from running, and switch to another to support **time sharing**?

How do programs share disks and perform I/O if they are given direct control? Do they know about each other?

With direct execution, how can dynamic memory structures such as linked lists grow over time?

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.38

CONTROL TRADEOFF

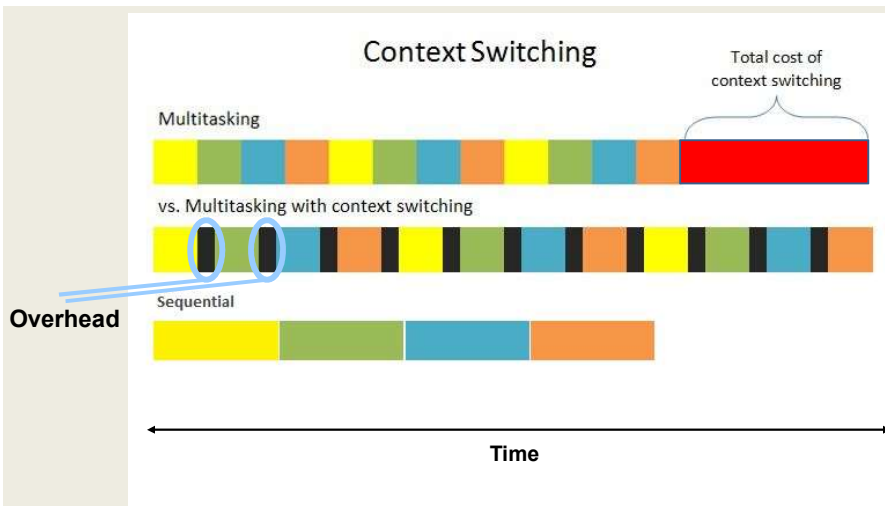
- **Too little control:**
 - No security
 - No time sharing
- **Too much control:**
 - Too much OS overhead
 - Poor performance for compute & I/O
 - Complex APIs (system calls), difficult to use

February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.39

CONTEXT SWITCHING OVERHEAD



February 13, 2019

TCSS422: Operating Systems [Winter 2019]
School of Engineering and Technology, University of Washington - Tacoma

L10.40