

Locks, Semaphores, and Producer- Consumer Problem

CS 571: Operating Systems (Spring 2020)

Lecture 3

Yue Cheng

Review: Threads

Threads

- **Processes** vs. **threads**
 - Parent and child processes do not share address space
 - Inter-process communication w/ message passing or shared memory
- Threads created by one process share address space, open files, global variables, etc.
- Much cheaper and more flexible inter-thread communication and cooperation

A Simple Example Using pthread

```
1  #include <stdio.h>
2  #include <assert.h>
3  #include <pthread.h>
4
5  void *mythread(void *arg) {
6      printf("%s\n", (char *) arg);
7      return NULL;
8  }
9
10 int
11 main(int argc, char *argv[]) {
12     pthread_t p1, p2;
13     int rc;
14     printf("main: begin\n");
15     rc = pthread_create(&p1, NULL, mythread, "A"); assert(rc == 0);
16     rc = pthread_create(&p2, NULL, mythread, "B"); assert(rc == 0);
17     // join waits for the threads to finish
18     rc = pthread_join(p1, NULL); assert(rc == 0);
19     rc = pthread_join(p2, NULL); assert(rc == 0);
20     printf("main: end\n");
21     return 0;
22 }
```

Thread Trace 1

main

Thread 1

Thread2

starts running
prints "main: begin"
creates Thread 1
creates Thread 2
waits for T1

Thread Trace 1

main

Thread 1

Thread2

starts running
prints "main: begin"
creates Thread 1
creates Thread 2
waits for T1

runs
prints "A"
returns

Thread Trace 1

<u>main</u>	<u>Thread 1</u>	<u>Thread2</u>
starts running		
prints "main: begin"		
creates Thread 1		
creates Thread 2		
waits for T1		
	runs	
	prints "A"	
	returns	
waits for T2		

Thread Trace 1

<u>main</u>	<u>Thread 1</u>	<u>Thread2</u>
starts running prints "main: begin" creates Thread 1 creates Thread 2 waits for T1	runs prints "A" returns	
waits for T2		runs prints "B" returns

Thread Trace 1

<u>main</u>	<u>Thread 1</u>	<u>Thread2</u>
starts running		
prints "main: begin"		
creates Thread 1		
creates Thread 2		
waits for T1		
	runs	
	prints "A"	
	returns	
waits for T2		
		runs
		prints "B"
		returns
prints "main: end"		

Thread Trace 2

main	Thread 1	Thread2
starts running prints "main: begin" creates Thread 1		

Thread Trace 2

main	Thread 1	Thread2
starts running prints "main: begin" creates Thread 1		
	runs prints "A" returns	

Thread Trace 2

main	Thread 1	Thread2
starts running prints "main: begin" creates Thread 1		
	runs prints "A" returns	
creates Thread 2		

Thread Trace 2

main	Thread 1	Thread2
starts running prints "main: begin" creates Thread 1		
	runs prints "A" returns	
creates Thread 2		
		runs prints "B" returns

Thread Trace 2

main	Thread 1	Thread2
starts running prints "main: begin" creates Thread 1		
	runs prints "A" returns	
creates Thread 2		
		runs prints "B" returns
waits for T1 <i>returns immediately; T1 is done</i>		
waits for T2 <i>returns immediately; T2 is done</i>		
prints "main: end"		

Thread Trace 2

main	Thread 1	Thread2
starts running prints "main: begin" creates Thread 1		
	runs prints "A" returns	
creates Thread 2		
		runs prints "B" returns
waits for T1 <i>returns immediately; T1 is done</i>		
waits for T2 <i>returns immediately; T2 is done</i>		
prints "main: end"		

What would a 3rd thread trace look like?

Synchronization

- Race Conditions
- The Critical Section Problem
- Synchronization Hardware and Locks
- Semaphores

Threaded Counting Example

```
1 #include <stdio.h>
2 #include "common.h"
3
4 static volatile int counter = 0;
5
6 //
7 // mythread()
8 //
9 // Simply adds 1 to counter repeatedly, in a loop
10 // No, this is not how you would add 10,000,000 to
11 // a counter, but it shows the problem nicely.
12 //
13 void *mythread(void *arg)
14 {
15     printf("%s: begin\n", (char *) arg);
16     int i;
17     for (i = 0; i < 1e7; i++) {
18         counter = counter + 1;
19     }
20     printf("%s: done\n", (char*) arg);
21     return NULL;
22 }
23
24 //
25 // main()
26 //
27 // Just launches two threads (pthread_create)
28 // and then waits for them (pthread_join)
29 //
30 int main(int argc, char *argv[])
31 {
32     pthread_t p1, p2;
33     printf("main: begin (counter = %d)\n", counter);
34     Pthread_create(&p1, NULL, mythread, "A");
35     Pthread_create(&p2, NULL, mythread, "B");
36
37     // join waits for the threads to finish
38     Pthread_join(p1, NULL);
39     Pthread_join(p2, NULL);
40     printf("main: done with both (counter = %d)\n", counter);
41     return 0;
42 }
```

```
$ git clone https://github.com/tddg/demo-ostep-code
```

```
$ cd demo-ostep-code/threads-intro
```

```
$ make
```

```
$ ./t1 <loop_count>
```

Try it yourself

Back-to-Back Runs

Run 1...

main: begin (counter = 0)

A: begin

B: begin

A: done

B: done

main: done with both (counter = 10706438)

Run 2...

main: begin (counter = 0)

A: begin

B: begin

A: done

B: done

main: done with both (counter = 11852529)

What exactly Happened??

What exactly Happened??

```
% otool -t -v thread_rc
```

[Mac OS X]

```
% objdump -d thread_rc
```

[Linux]

...

```
00000000100000d52    movl 0x2f8e %eax
00000000100000d58    addl $0x1, %eax
00000000100000d5b    movl %eax, 0x2f8e
```

...

```
counter = counter + 1;
```

Concurrent Access to the Same Memory Address

OS

Thread 1

Thread 2

Value



Enter into critical section

`movl 0x2f8e, %eax`

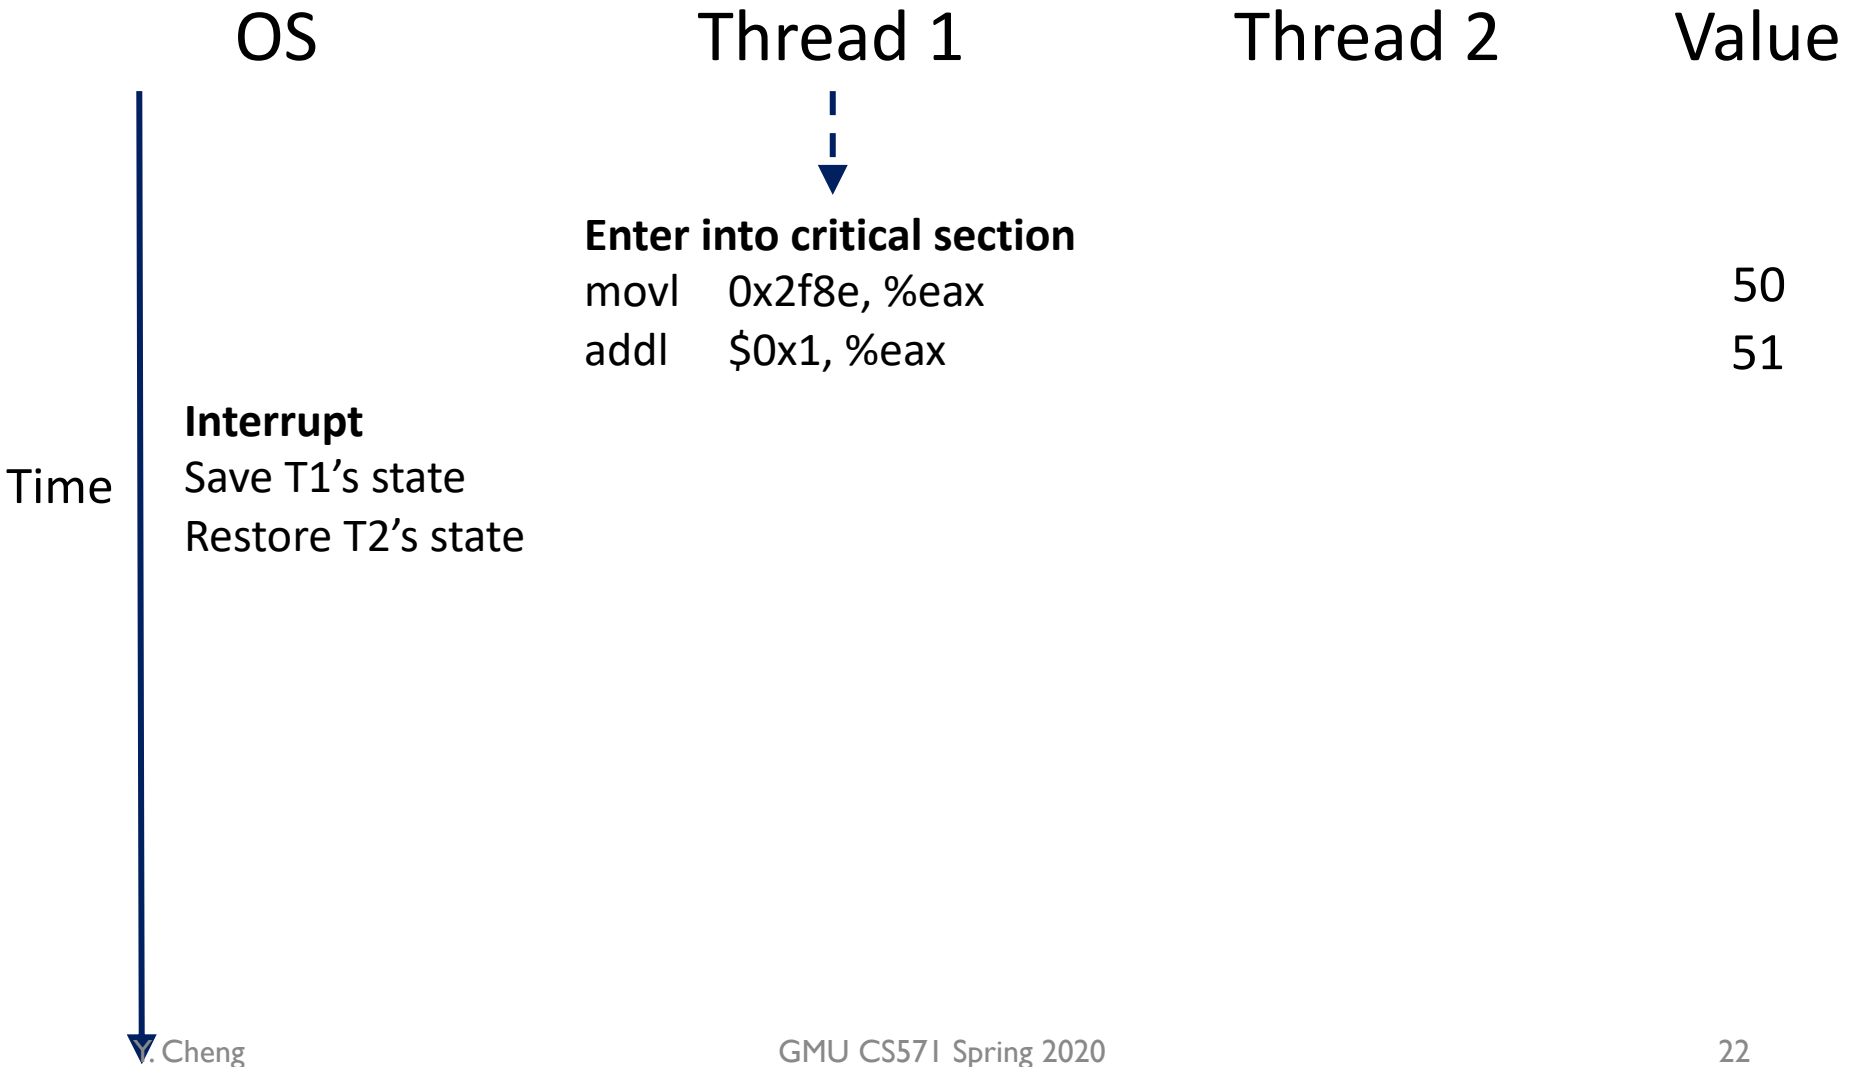
`addl $0x1, %eax`

50

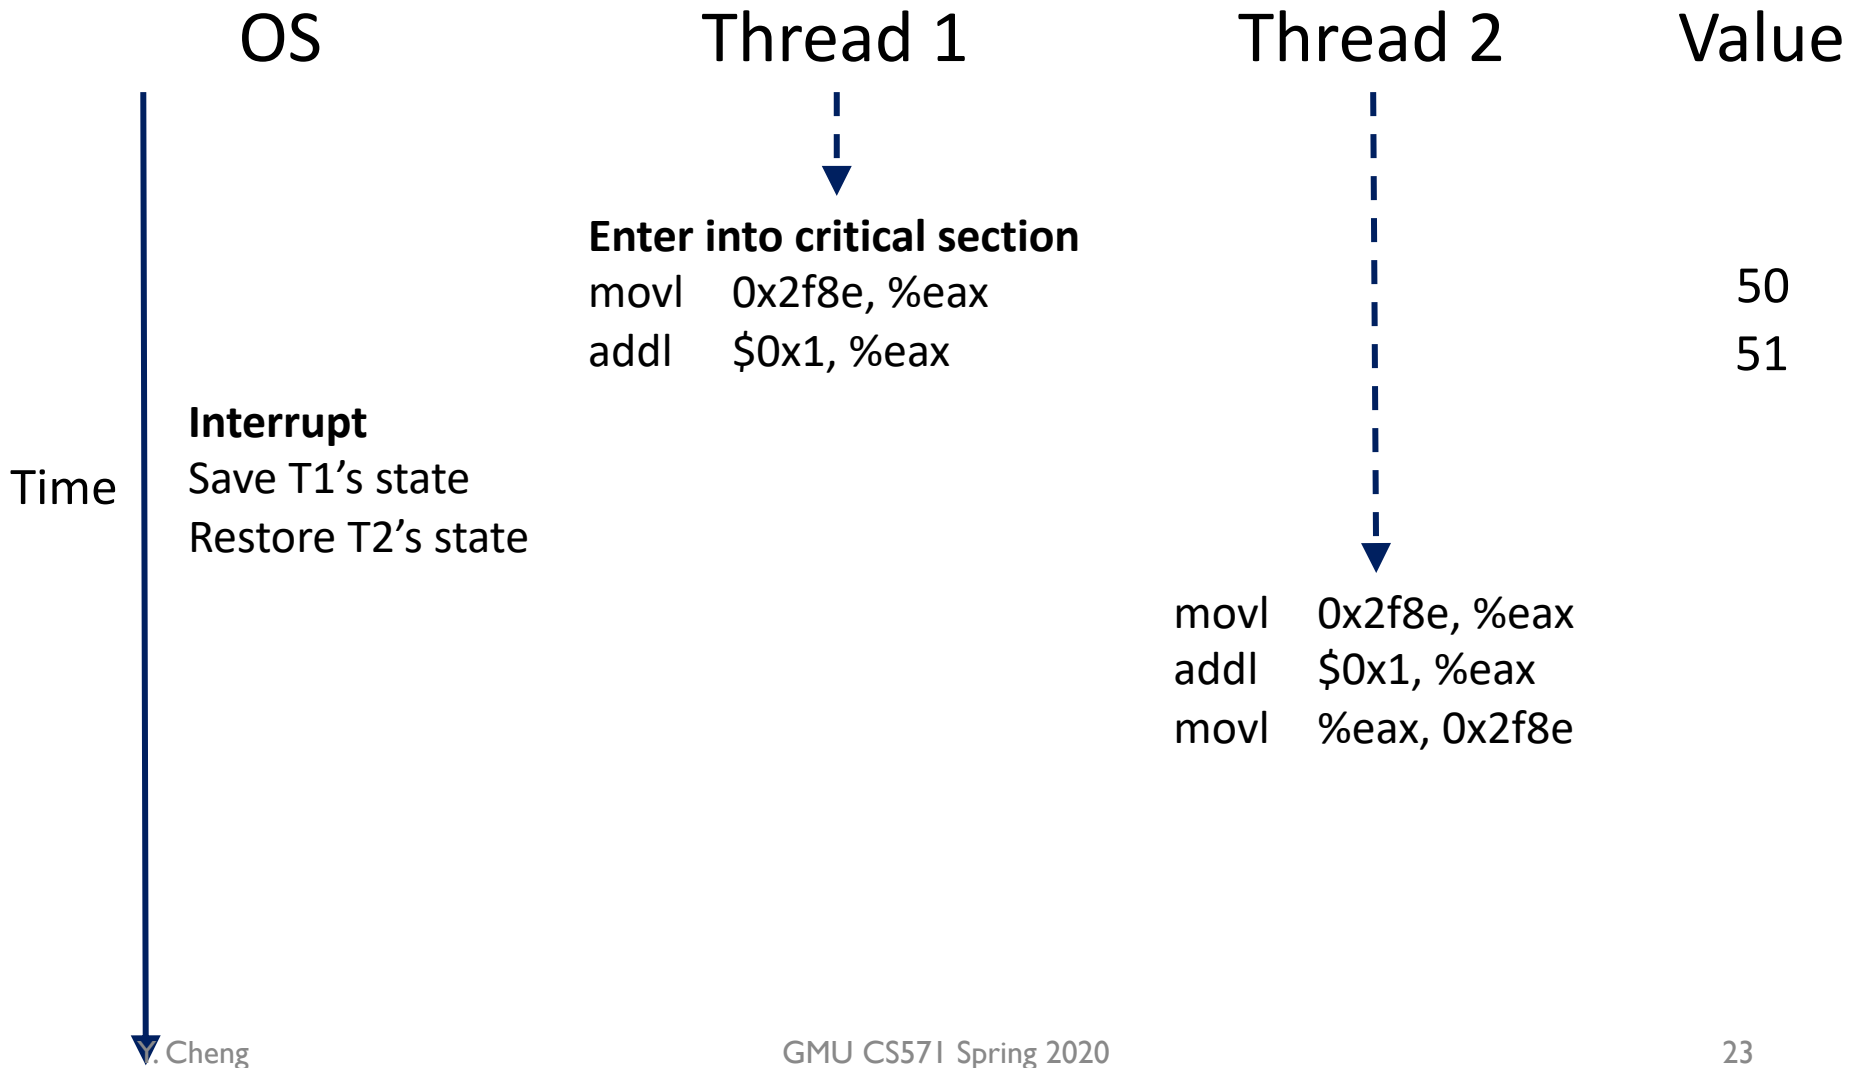
51

Time

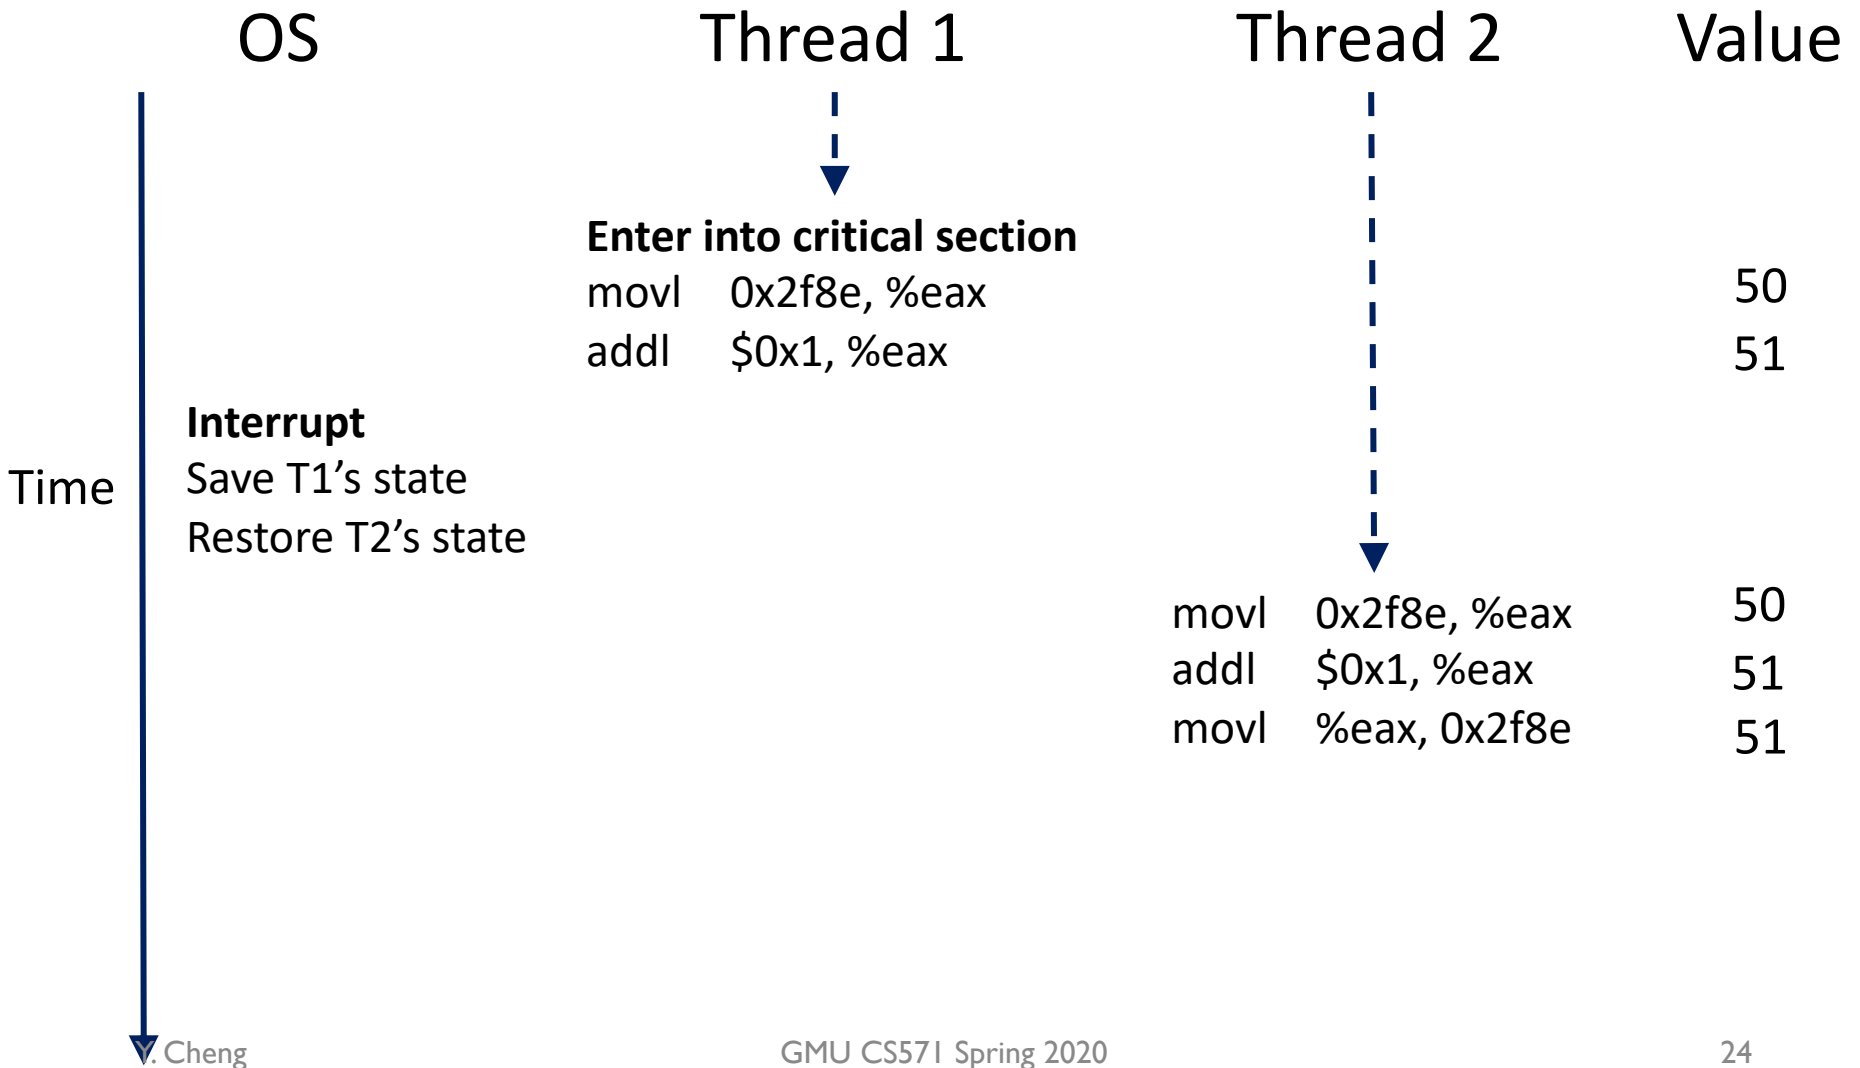
Concurrent Access to the Same Memory Address



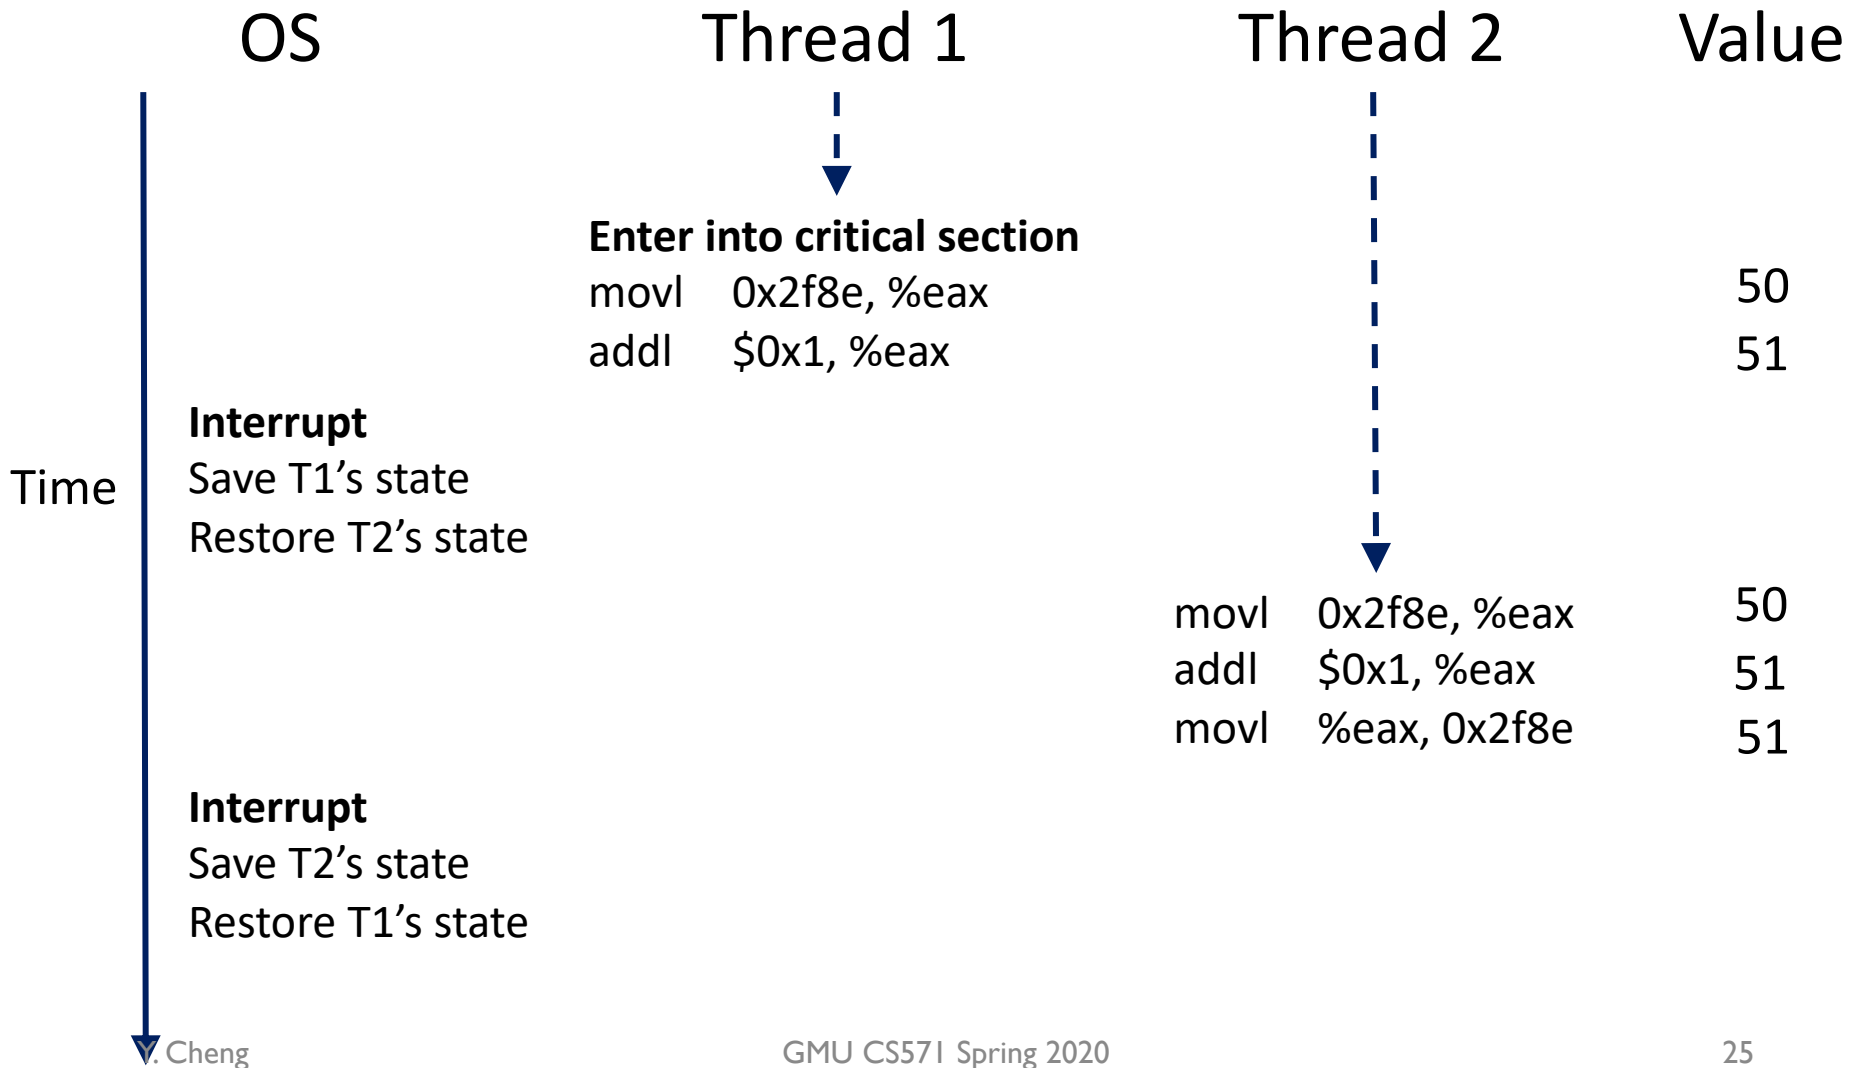
Concurrent Access to the Same Memory Address




Concurrent Access to the Same Memory Address



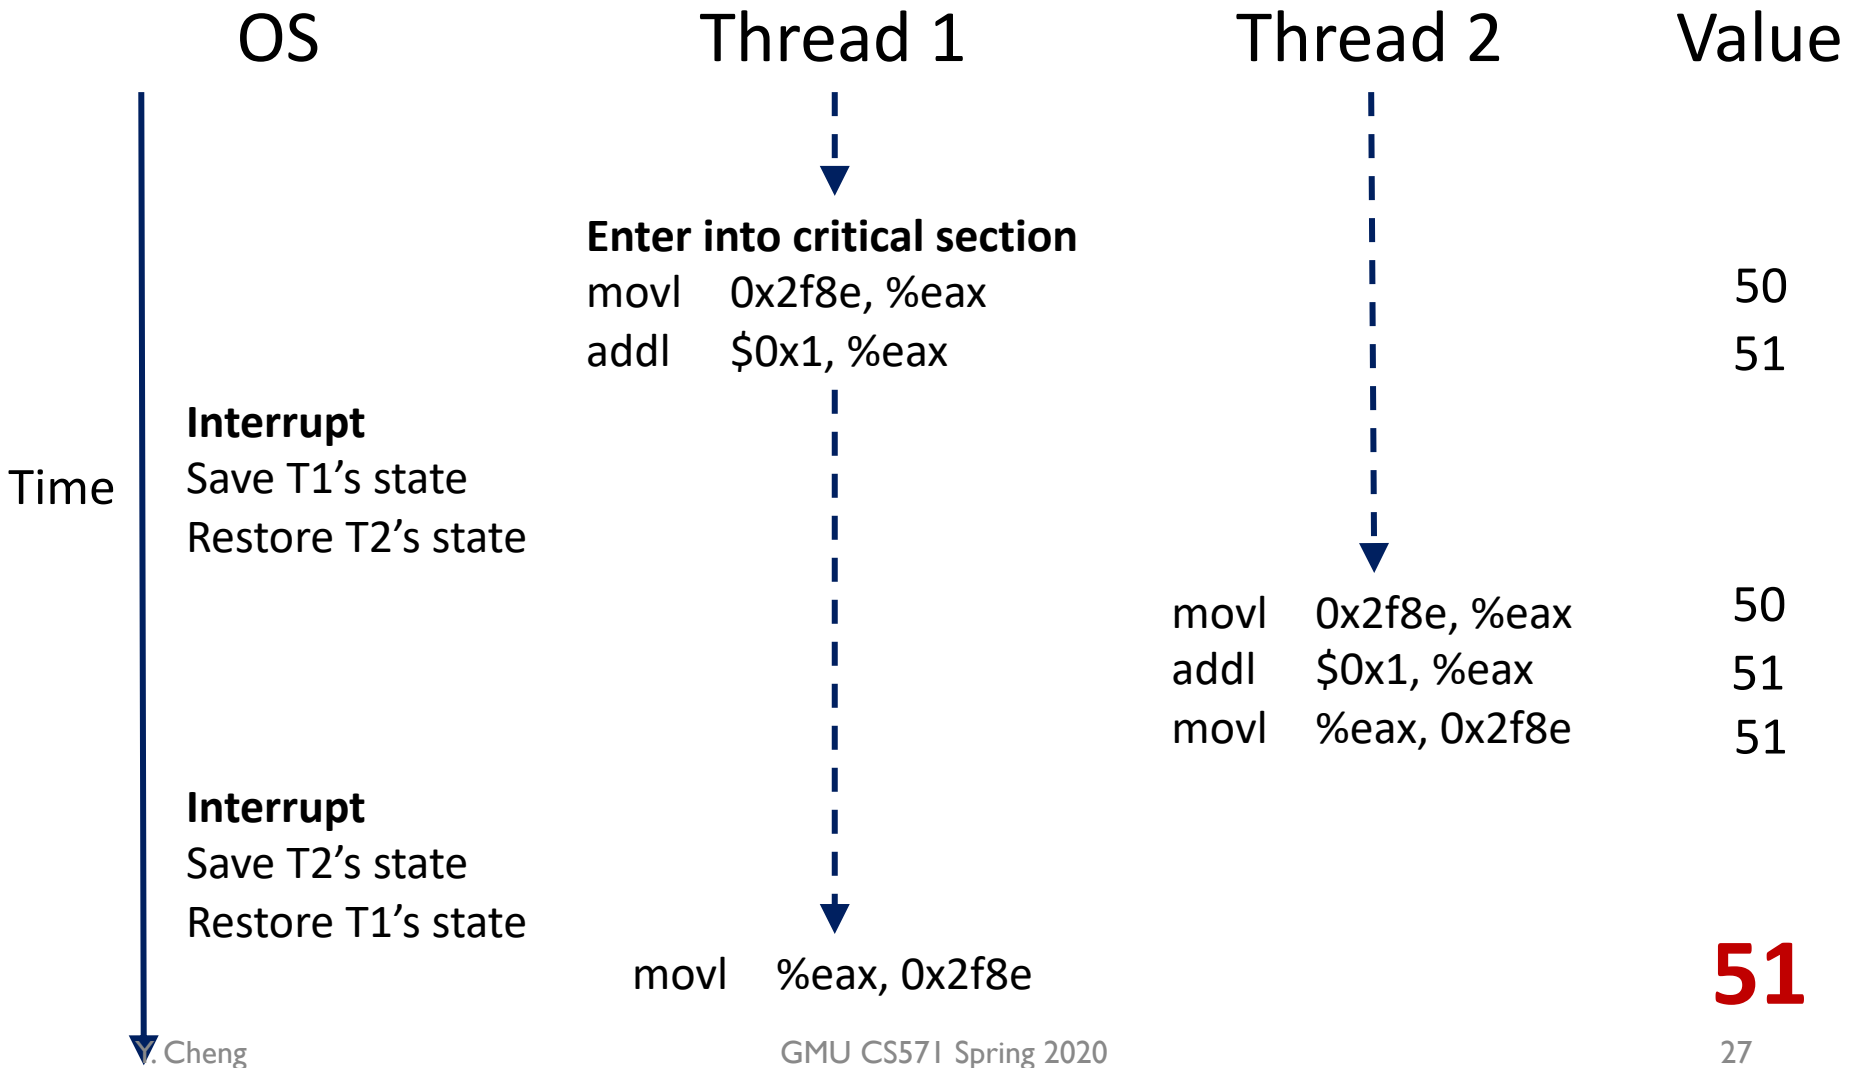
Concurrent Access to the Same Memory Address



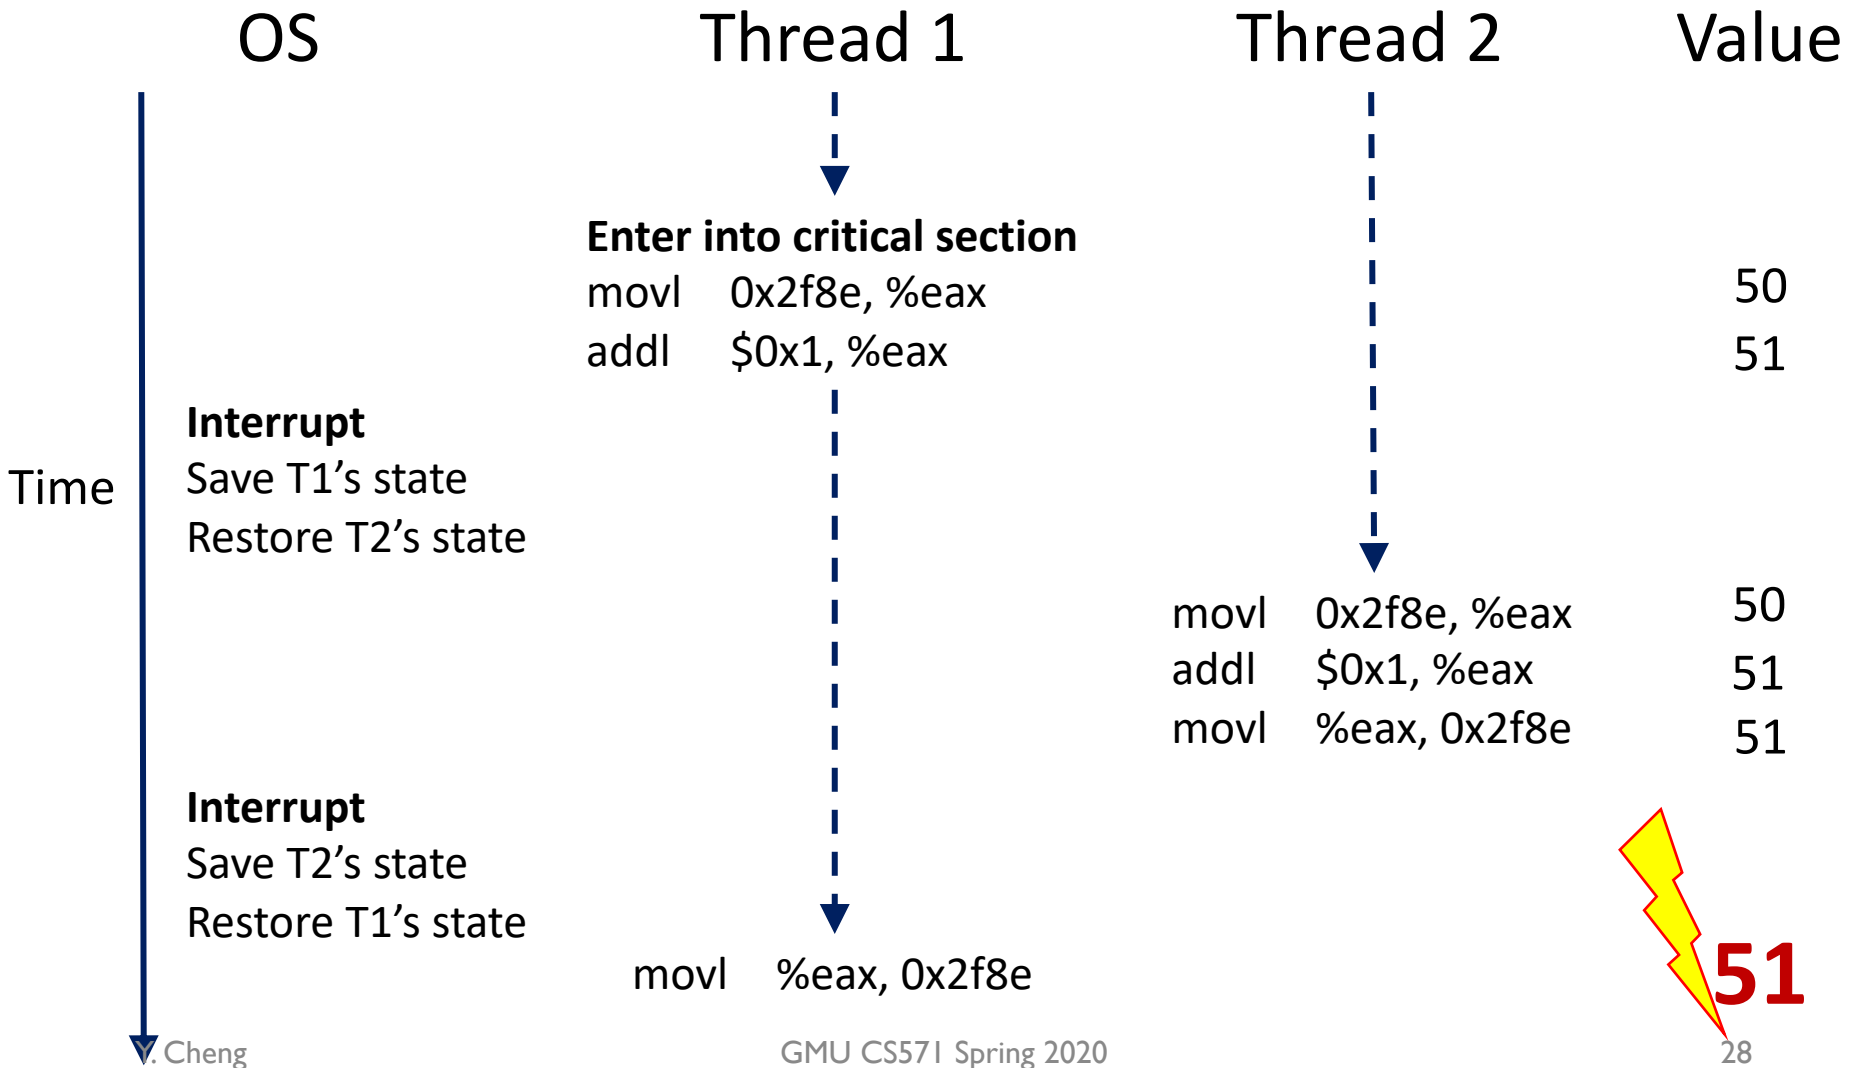
Concurrent Access to the Same Memory Address



Concurrent Access to the Same Memory Address



Concurrent Access to the Same Memory Address



Takeaway

- Observe: In a **time-shared** system, **the exact instruction execution order** cannot be predicted
 - Deterministic vs. **Non-deterministic**
- Any possible orders could happen, which result in different output across runs

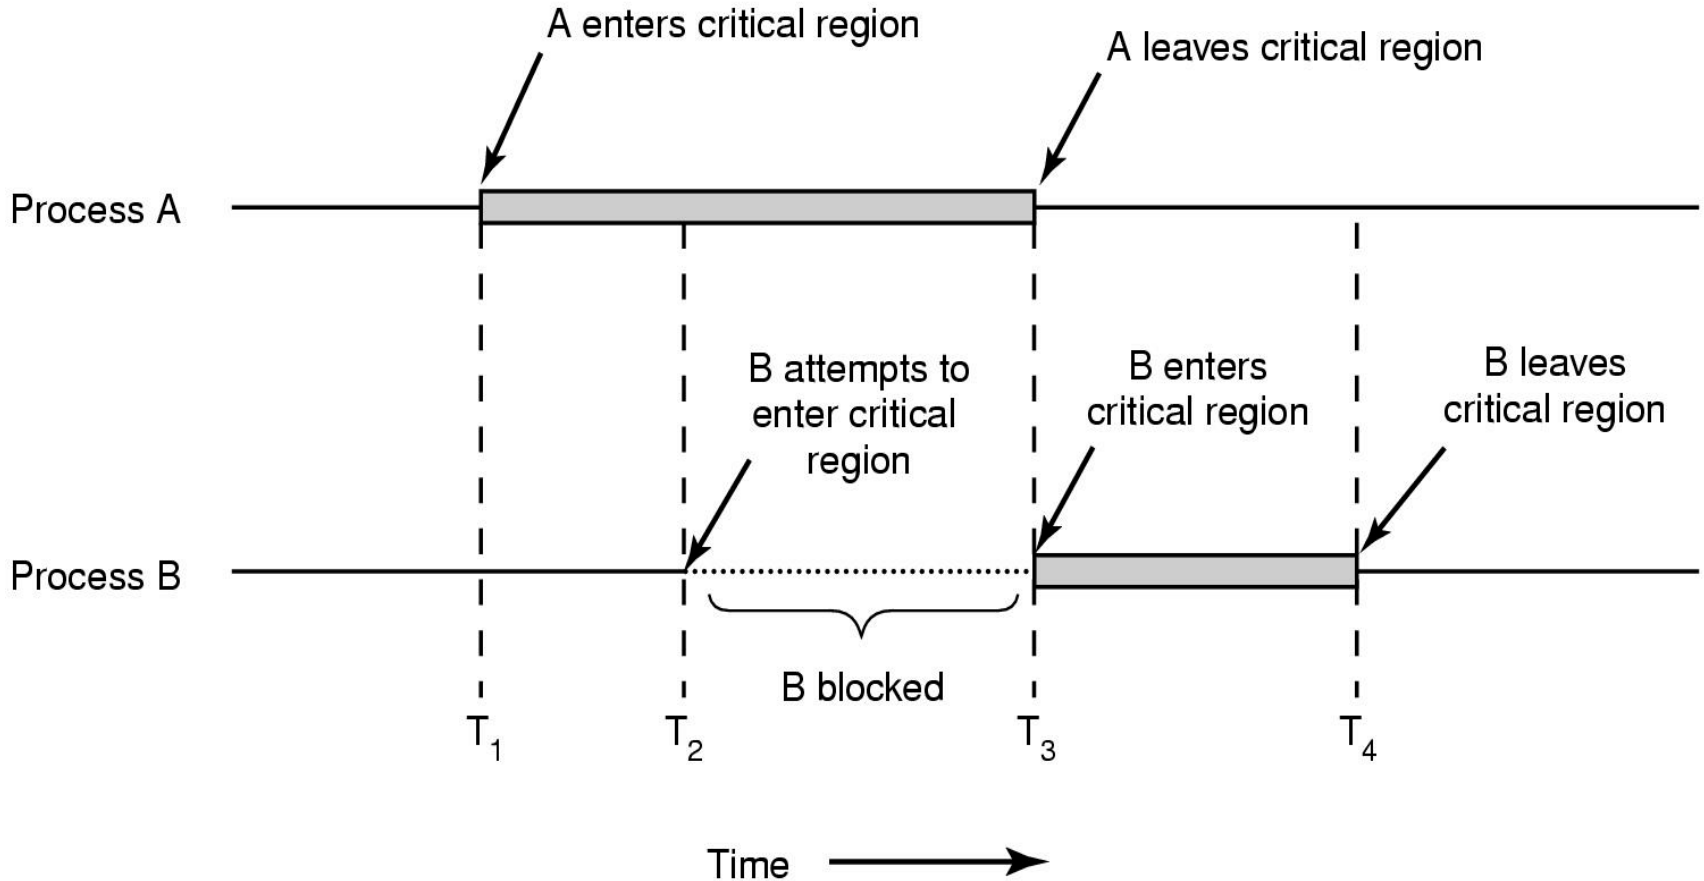
Race Conditions

- Situations like this, where multiple processes are writing or reading some shared data and the final **result depends on who runs precisely when**, are called **race conditions**
 - A serious problem for any concurrent system using shared variables
- Programmers must make sure that some **high-level** code sections are executed **atomically**
 - Atomic operation: It completes in its **entirety without worrying about interruption by any other potentially conflict-causing process**

The Critical-Section Problem

- N processes/threads all competing to access the shared data
- Each process/thread has a code segment, called **critical section (critical region)**, in which the shared data is accessed
- **Problem** – ensure that when one process is executing in its critical section, no other process is allowed to execute in **that** critical section
- The execution of the critical sections by the processes must be **mutually exclusive** in time

Mutual Exclusion



Solving Critical-Section Problem

Any solution to the problem must satisfy **four conditions!**

Mutual Exclusion:

No two processes may be simultaneously inside the same critical section

Bounded Waiting:

No process should have to wait forever to enter a critical section

Progress:

No process executing a code segment unrelated to a given critical section can block another process trying to enter the same critical section

Arbitrary Speed:

No assumption can be made about the relative speed of different processes (though all processes have a non-zero speed)

Using Lock to Protect Shared Data



- Suppose that two threads A and B have access to a shared variable “balance”

Thread A:

```
balance = balance + 1
```

Thread B:

```
balance = balance + 1
```

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



Locks

- A lock is a **variable**
- Two states
 - Available or free
 - Locked or held
- **lock()**: tries to acquire the lock
- **unlock()**: releases the lock that has been acquired by caller

Building a Lock

- Needs help from hardware + OS
- A number of hardware primitives to support a lock
- Goals of a lock
 - Basic task: Mutual exclusion
 - Fairness
 - Performance

First Attempt: A Simple Flag

- How about just using loads/stores instructions?

```
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 }
```

First Attempt: A Simple Flag

- How about just using loads/stores instructions?

```
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 }
```

→ A spin lock

First Attempt: A Simple Flag

- How about just using loads/stores instructions?

```
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 }
```

→ A spin lock

What's the problem?

First Attempt: A Simple Flag

Flag is 0 initially

Thread 1

Thread 2

call `lock()`

while (`flag == 1`)

interrupt: switch to Thread 2

First Attempt: A Simple Flag

Flag is 0 initially

Thread 1

call `lock()`

while (`flag == 1`)

interrupt: switch to Thread 2

Thread 2

Checking that Flag is 0, again...

call `lock()`

while (`flag == 1`)

First Attempt: A Simple Flag

Flag is set to 1 by T2

Thread 1

call `lock()`

while (`flag == 1`)

interrupt: switch to Thread 2

Thread 2

call `lock()`

while (`flag == 1`)

`flag = 1;`

interrupt: switch to Thread 1

First Attempt: A Simple Flag

Flag is set to 1 again! Two threads both in Critical Section

Thread 1

```
call lock ()
```

```
while (flag == 1)
```

```
interrupt: switch to Thread 2
```

```
flag = 1; // set flag to 1 (too!)
```

Thread 2

```
call lock ()
```

```
while (flag == 1)
```

```
flag = 1;
```

```
interrupt: switch to Thread 1
```

First Attempt: A Simple Flag

Flag is set to 1 again! Two threads both in Critical Section

Thread 1

```
call lock ()  
while (flag == 1)  
interrupt: switch to Thread 2
```

```
flag = 1; // set flag to 1 (too!)
```

Thread 2

```
call lock ()  
while (flag == 1)  
flag = 1;  
interrupt: switch to Thread 1
```

Reason:

Lock operation is not atomic!

And therefore, no mutual exclusion!

Getting Help from the Hardware

- One solution supported by hardware may be to use interrupt capability

```
do {  
    lock()  
    critical section;  
    unlock()  
    remainder section;  
} while (1);
```

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

Getting Help from the Hardware

- One solution supported by hardware may be to use interrupt capability

```
do {  
    lock()  
    critical section;  
    unlock()  
    remainder section;  
} while (1);
```

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

Are we done??

Synchronization Hardware

- Many machines provide special **hardware instructions** to help achieve mutual exclusion
- The **TestAndSet (TAS)** instruction tests and modifies the content of a memory word **atomically**
- TAS returns old value pointed to by **old_ptr** and updates said value to **new**

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

Operations performed atomically!

Mutual Exclusion with TAS

- Initially, lock's flag set to 0

```
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, 1 that it is held
7      lock->flag = 0;
8  }
9
10 void lock(lock_t *lock) {
11     while (TestAndSet(&lock->flag, 1) == 1)
12         ; // spin-wait (do nothing)
13 }
14
15 void unlock(lock_t *lock) {
16     lock->flag = 0;
17 }
```

→ A correct spin lock

Busy Waiting and Spin Locks

- This approach is based on **busy waiting**
 - If the critical section is being used, waiting processes loop continuously at the entry point
- A binary “**lock**” variable that uses busy waiting is called a **spin lock**
 - Processes that find the lock unavailable “spin” at the entry
- It actually works (**mutual exclusion**)

Busy Waiting and Spin Locks

- This approach is based on **busy waiting**
 - If the critical section is being used, waiting processes loop continuously at the entry point
- A binary “**lock**” variable that uses busy waiting is called a **spin lock**
 - Processes that find the lock unavailable “spin” at the entry
- It actually works (**mutual exclusion**)
- **Disadvantages?**
 - **Fairness?**
 - **Performance?**

Busy Waiting and Spin Locks

- This approach is based on **busy waiting**
 - If the critical section is being used, waiting processes loop continuously at the entry point
- A binary “**lock**” variable that uses busy waiting is called a **spin lock**
 - Processes that find the lock unavailable “spin” at the entry
- It actually works (**mutual exclusion**)
- **Disadvantages?**
 - **Fairness?** (A: No. Heavy contention may cause starvation)
 - **Performance?** (A: Busy waiting wastes CPU cycles)

A Simple Approach: Just Yield (Win)!

- When you are going to spin, just **give up** the CPU to another process/thread

```
1  void init() {
2      flag = 0;
3  }
4
5  void lock() {
6      while (TestAndSet(&flag, 1) == 1)
7          yield(); // give up the CPU
8  }
9
10 void unlock() {
11     flag = 0;
12 }
```



Lock Worksheet

Semaphores

- Introduced by E. W. Dijkstra
- Motivation: Avoid busy waiting by **blocking** a process execution until some condition is satisfied
- Two operations are defined on a semaphore variable s :
 - `sem_wait(s)` (also called $P(s)$ or $\text{down}(s)$)
 - `sem_post(s)` (also called $V(s)$ or $\text{up}(s)$)

Semaphore Operations

- Conceptually, a semaphore has an integer value. This value is greater than or equal to 0
- `sem_wait(s):`
`s.value-- ; /* Executed atomically */`
`/* wait/block if s.value < 0 (or negative) */`
- A process/thread executing the wait operation on a semaphore with value < 0 being **blocked** until the semaphore's value becomes greater than 0
 - **No busy waiting**
- `sem_post(s):`
`s.value++; /* Executed atomically */`
`/* if one or more process/thread waiting, wake one */`

Semaphore Operations (cont.)

- If multiple processes/threads are blocked on the same semaphore 's', only one of them will be awakened when another process performs post(s) operation
- Who will have higher priority?

Semaphore Operations (cont.)

- If multiple processes/threads are blocked on the same semaphore 's', only one of them will be awakened when another process performs post(s) operation
- Who will have higher priority?
 - A: FIFO, or whatever queuing strategy

Attacking Critical Section Problem with Semaphores

- Declare and define a semaphore:

```
sem_t s;  
sem_init(&s, 0, 1); /* initially s = 1 */
```

- Routine of Thread 0 & 1:

```
do {  
    sem_wait(s);  
    critical section  
  
    sem_post(s);  
    remainder section  
} while (1);
```

Binary semaphore,
which is a lock



Attacking Critical Section Problem with Semaphores

- **Single** thread using a **binary** semaphore

Value of Semaphore	Thread 0	Thread 1
1		

Attacking Critical Section Problem with Semaphores

- **Single** thread using a **binary** semaphore

Value of Semaphore	Thread 0	Thread 1
1		
1	call <code>sem_wait()</code>	
0	<code>sem_wait()</code> returns	

Attacking Critical Section Problem with Semaphores

- **Single** thread using a **binary** semaphore

Value of Semaphore	Thread 0	Thread 1
1		
1	call <code>sem_wait()</code>	
0	<code>sem_wait()</code> returns	
0	(crit sect)	
0	call <code>sem_post()</code>	

Attacking Critical Section Problem with Semaphores

- **Single** thread using a **binary** semaphore

Value of Semaphore	Thread 0	Thread 1
1		
1	call <code>sem_wait()</code>	
0	<code>sem_wait()</code> returns	
0	(crit sect)	
0	call <code>sem_post()</code>	
1	<code>sem_post()</code> returns	

Attacking Critical Section Problem with Semaphores

- **Two** threads using a **binary** semaphore

Value	Thread 0	State	Thread 1	State
1		Running		Ready

Attacking Critical Section Problem with Semaphores

- Two threads using a binary semaphore

Value	Thread 0	State	Thread 1	State
1		Running		Ready
1	call <code>sem_wait()</code>	Running		Ready
0	<code>sem_wait()</code> returns	Running		Ready
0	(crit sect: begin)	Running		Ready

Attacking Critical Section Problem with Semaphores

- **Two** threads using a **binary** semaphore

Value	Thread 0	State	Thread 1	State
1		Running		Ready
1	<code>call sem_wait()</code>	Running		Ready
0	<code>sem_wait()</code> returns	Running		Ready
0	<code>(crit sect: begin)</code>	Running		Ready
0	<i>Interrupt; Switch → T1</i>	Ready		Running

Attacking Critical Section Problem with Semaphores

- **Two** threads using a **binary** semaphore

Value	Thread 0	State	Thread 1	State
1		Running		Ready
1	call <code>sem_wait()</code>	Running		Ready
0	<code>sem_wait()</code> returns	Running		Ready
0	(crit sect: begin)	Running		Ready
0	<i>Interrupt; Switch</i> → T1	Ready		Running
0		Ready	call <code>sem_wait()</code>	Running
-1		Ready	decrement sem	Running
-1		Ready	(sem < 0) → sleep	Sleeping

Attacking Critical Section Problem with Semaphores

- **Two** threads using a **binary** semaphore

Value	Thread 0	State	Thread 1	State
1		Running		Ready
1	call <code>sem_wait()</code>	Running		Ready
0	<code>sem_wait()</code> returns	Running		Ready
0	(crit sect: begin)	Running		Ready
0	<i>Interrupt; Switch</i> → T1	Ready		Running
0		Ready	call <code>sem_wait()</code>	Running
-1		Ready	decrement sem	Running
-1		Ready	(sem < 0) → sleep	Sleeping
-1		Running	<i>Switch</i> → T0	Sleeping

Attacking Critical Section Problem with Semaphores

- **Two** threads using a **binary** semaphore

Value	Thread 0	State	Thread 1	State
1		Running		Ready
1	call sem_wait()	Running		Ready
0	sem_wait() returns	Running		Ready
0	(crit sect: begin)	Running		Ready
0	<i>Interrupt; Switch</i> →T1	Ready		Running
0		Ready	call sem_wait()	Running
-1		Ready	decrement sem	Running
-1		Ready	(sem<0)→sleep	Sleeping
-1		Running	<i>Switch</i> →T0	Sleeping
-1	(crit sect: end)	Running		Sleeping
-1	call sem_post()	Running		Sleeping
0	increment sem	Running		Sleeping
0	wake(T1)	Running		Ready
0	sem_post() returns	Running		Ready

Attacking Critical Section Problem with Semaphores

- **Two** threads using a **binary** semaphore

Value	Thread 0	State	Thread 1	State
1		Running		Ready
1	call sem_wait()	Running		Ready
0	sem_wait() returns	Running		Ready
0	(crit sect: begin)	Running		Ready
0	<i>Interrupt; Switch</i> →T1	Ready		Running
0		Ready	call sem_wait()	Running
-1		Ready	decrement sem	Running
-1		Ready	(sem<0)→sleep	Sleeping
-1		Running	<i>Switch</i> →T0	Sleeping
-1	(crit sect: end)	Running		Sleeping
-1	call sem_post()	Running		Sleeping
0	increment sem	Running		Sleeping
0	wake(T1)	Running		Ready
0	sem_post() returns	Running		Ready
0	<i>Interrupt; Switch</i> →T1	Ready		Running

Attacking Critical Section Problem with Semaphores

- **Two** threads using a **binary** semaphore

Value	Thread 0	State	Thread 1	State
1		Running		Ready
1	call sem_wait()	Running		Ready
0	sem_wait() returns	Running		Ready
0	(crit sect: begin)	Running		Ready
0	<i>Interrupt; Switch</i> →T1	Ready		Running
0		Ready	call sem_wait()	Running
-1		Ready	decrement sem	Running
-1		Ready	(sem<0)→sleep	Sleeping
-1		Running	<i>Switch</i> →T0	Sleeping
-1	(crit sect: end)	Running		Sleeping
-1	call sem_post()	Running		Sleeping
0	increment sem	Running		Sleeping
0	wake(T1)	Running		Ready
0	sem_post() returns	Running		Ready
0	<i>Interrupt; Switch</i> →T1	Ready		Running
0		Ready	sem_wait() returns	Running
0		Ready	(crit sect)	Running
0		Ready	call sem_post()	Running
1		Ready	sem_post() returns	Running

Classical Problems of Synchronization


- Producer-Consumer Problem
 - Semaphore version
 - Condition Variable
 - A CV-based version

Today

- Readers-Writers Problem
- Dining-Philosophers Problem

Producer-Consumer Problem


- The **bounded-buffer** producer-consumer problem assumes that there is a buffer of size N
- The producer process puts items to the buffer area
- The consumer process consumes items from the buffer
- The producer and the consumer execute **concurrently**



Example: Unix Pipes


- A pipe may have many writers and readers
- Internally, there is a **finite-sized buffer**
- Writers **add data** to the buffer
- Readers **remove data** from the buffer

Example: Unix Pipes




Example: Unix Pipes

Write




Example: Unix Pipes




Example: Unix Pipes


Write




Example: Unix Pipes



Example: Unix Pipes




Example: Unix Pipes




Example: Unix Pipes

Write




Example: Unix Pipes




Example: Unix Pipes

Read




Example: Unix Pipes

Read



Example: Unix Pipes


Read



Note: reader must **wait**


Example: Unix Pipes

Write



Example: Unix Pipes

Write



Example: Unix Pipes

Write



Note: writer must **wait**

Example: Unix Pipes

- Implementation
 - Reads/writes to buffer require **locking**
 - When buffers are **full**, writers (producers) **must wait**
 - When buffers are **empty**, readers (consumers) **must wait**

Linux Pipe Commands

```
% ps aux | less
```



Pipe

```
% cat file | grep <str>
```



Pipe

Producer-Consumer Model: Parameters

- Shared data:
`sem_t full, empty;`

- Initially:

```
full = 0          /* The number of full buffers */
```

```
empty = MAX      /* The number of empty buffers */
```

First Attempt: MAX = 1

```
1 sem_t empty;
2 sem_t full;
3
4 void *producer(void *arg) {
5     int i;
6     for (i = 0; i < loops; i++) {
7         sem_wait(&empty);           // line P1
8         put(i);                     // line P2
9         sem_post(&full);           // line P3
10    }
11 }
12
13 void *consumer(void *arg) {
14     int i, tmp = 0;
15     while (tmp != -1) {
16         sem_wait(&full);           // line C1
17         tmp = get();               // line C2
18         sem_post(&empty);         // line C3
19         printf("%d\n", tmp);
20     }
21 }
22
23 int main(int argc, char *argv[]) {
24     // ...
25     sem_init(&empty, 0, MAX); // MAX buffers are empty to begin with...
26     sem_init(&full, 0, 0);    // ... and 0 are full
27     // ...
28 }
```

```
1 int buffer[MAX];
2 int fill = 0;
3 int use = 0;
4
5 void put(int value) {
6     buffer[fill] = value;
7     fill = (fill + 1) % MAX;
8 }
9
10 int get() {
11     int tmp = buffer[use];
12     use = (use + 1) % MAX;
13     return tmp;
14 }
```

Put and Get routines

First Attempt: MAX = 10?

```
1 sem_t empty;
2 sem_t full;
3
4 void *producer(void *arg) {
5     int i;
6     for (i = 0; i < loops; i++) {
7         sem_wait(&empty);           // line P1
8         put(i);                     // line P2
9         sem_post(&full);           // line P3
10    }
11 }
12
13 void *consumer(void *arg) {
14     int i, tmp = 0;
15     while (tmp != -1) {
16         sem_wait(&full);           // line C1
17         tmp = get();               // line C2
18         sem_post(&empty);         // line C3
19         printf("%d\n", tmp);
20     }
21 }
22
23 int main(int argc, char *argv[]) {
24     // ...
25     sem_init(&empty, 0, MAX); // MAX buffers are empty to begin with...
26     sem_init(&full, 0, 0);    // ... and 0 are full
27     // ...
28 }
```

```
1 int buffer[MAX];
2 int fill = 0;
3 int use = 0;
4
5 void put(int value) {
6     buffer[fill] = value;
7     fill = (fill + 1) % MAX;
8 }
9
10 int get() {
11     int tmp = buffer[use];
12     use = (use + 1) % MAX;
13     return tmp;
14 }
```


Put and Get routines

First Attempt: MAX = 10?


fill = 0

empty = 10

Producer 0: **Running**

```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
         sem_wait(&empty);
        put(i);
        sem_post(&full);
    }
}
```

Producer 1: Runnable

```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
         sem_wait(&empty);
        put(i);
        sem_post(&full);
    }
}
```


First Attempt: MAX = 10?

fill = 0

empty = 9


Producer 0: **Running**

```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&empty);
        put(i);
        sem_post(&full);
    }
}
```




Producer 1: Runnable

```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&empty);
        put(i);
        sem_post(&full);
    }
}
```



```
void put(int value) {
    buffer[fill] = value;
    fill = (fill + 1) % MAX;
}
```




First Attempt: MAX = 10?

fill = 0

empty = 9


Producer 0: **Running**

```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&empty);
        put(i);
        sem_post(&full);
    }
}
```




Producer 1: **Runnable**

```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&empty);
        put(i);
        sem_post(&full);
    }
}
```



```
void put(int value) {
    buffer[fill] = value;
    Interrupted ...
    fill = (fill + 1) % MAX;
}
```




First Attempt: MAX = 10?

fill = 0

empty = 9


Producer 0: **Sleeping**

```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&empty);
        put(i);
        sem_post(&full);
    }
}
```




Producer 1: Runnable

```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&empty);
        put(i);
        sem_post(&full);
    }
}
```



```
void put(int value) {
    buffer[fill] = value;
    Interrupted ...
    fill = (fill + 1) % MAX;
}
```




First Attempt: MAX = 10?

fill = 0

empty = 9


Producer 0: Runnable

```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&empty);
        put(i);
        sem_post(&full);
    }
}
```



Producer 1: **Running**

```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&empty);
        put(i);
        sem_post(&full);
    }
}
```



```
void put(int value) {
    buffer[fill] = value;
    Interrupted ...
    fill = (fill + 1) % MAX;
}
```

First Attempt: MAX = 10?


fill = 0

Overwrite!

empty = 8

Producer 0: Runnable


```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&empty);
        put(i);
        sem_post(&full);
    }
}
```




```
void put(int value) {
    buffer[fill] = value;
    Interrupted ...
    fill = (fill + 1) % MAX;
}
```

Producer 1: **Running**

```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&empty);
        put(i);
        sem_post(&full);
    }
}
```



```
void put(int value) {
     buffer[fill] = value;
    fill = (fill + 1) % MAX;
}
```

One More Parameter: A `mutex` lock

- Shared data:

```
sem_t full, empty;
```

- Initially:

```
full = 0;      /* The number of full buffers */  
empty = MAX;  /* The number of empty buffers */  
mutex = 1;    /* Semaphore controlling the access  
              to the buffer pool */
```

Add “Mutual Exclusion”

```
1  sem_t empty;
2  sem_t full;
3  sem_t mutex;
4
5  void *producer(void *arg) {
6      int i;
7      for (i = 0; i < loops; i++) {
8          sem_wait(&mutex);           // line p0 (NEW LINE)
9          sem_wait(&empty);           // line p1
10         put(i);                     // line p2
11         sem_post(&full);             // line p3
12         sem_post(&mutex);           // line p4 (NEW LINE)
13     }
14 }
15
16 void *consumer(void *arg) {
17     int i;
18     for (i = 0; i < loops; i++) {
19         sem_wait(&mutex);           // line c0 (NEW LINE)
20         sem_wait(&full);             // line c1
21         int tmp = get();             // line c2
22         sem_post(&empty);           // line c3
23         sem_post(&mutex);           // line c4 (NEW LINE)
24         printf("%d\n", tmp);
25     }
26 }
27
28 int main(int argc, char *argv[]) {
29     // ...
30     sem_init(&empty, 0, MAX); // MAX buffers are empty to begin with...
31     sem_init(&full, 0, 0);    // ... and 0 are full
32     sem_init(&mutex, 0, 1);   // mutex=1 because it is a lock (NEW LINE)
33     // ...
34 }
```

Add “Mutual Exclusion”

```
1  sem_t empty;
2  sem_t full;
3  sem_t mutex;
4
5  void *producer(void *arg) {
6      int i;
7      for (i = 0; i < loops; i++) {
8          sem_wait(&mutex);           // line p0 (NEW LINE)
9          sem_wait(&empty);           // line p1
10         put(i);                       // line p2
11         sem_post(&full);             // line p3
12         sem_post(&mutex);           // line p4 (NEW LINE)
13     }
14 }
15
16 void *consumer(void *arg) {
17     int i;
18     for (i = 0; i < loops; i++) {
19         sem_wait(&mutex);           // line c0 (NEW LINE)
20         sem_wait(&full);             // line c1
21         int tmp = get();             // line c2
22         sem_post(&empty);           // line c3
23         sem_post(&mutex);           // line c4 (NEW LINE)
24         printf("%d\n", tmp);
25     }
26 }
27
28 int main(int argc, char *argv[]) {
29     // ...
30     sem_init(&empty, 0, MAX); // MAX buffers are empty to begin with...
31     sem_init(&full, 0, 0);    // ... and 0 are full
32     sem_init(&mutex, 0, 1);   // mutex=1 because it is a lock (NEW LINE)
33     // ...
34 }
```


**What if consumer
gets to run first??**

Adding “Mutual Exclusion”

mutex = 1
full = 0
empty = 10


Producer 0: Runnable

```
void *producer(void *arg) {  
    int i;  
    for (i = 0; i < loops; i++) {  
        sem_wait(&mutex);  
        sem_wait(&empty);  
        put(i);  
        sem_post(&full);  
        sem_post(&mutex);  
    }  
}
```



Consumer 0: **Running**

```
void *consumer(void *arg) {  
    int i;  
    for (i = 0; i < loops; i++) {  
        sem_wait(&mutex);  
        sem_wait(&full);  
        int tmp = get();  
        sem_post(&empty);  
        sem_post(&mutex);  
        printf("%d\n", tmp);  
    }  
}
```



Adding “Mutual Exclusion”


mutex = 0

full = 0

empty = 10


Producer 0: Runnable

```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&mutex);
        sem_wait(&empty);
        put(i);
        sem_post(&full);
        sem_post(&mutex);
    }
}
```



Consumer 0: **Running**

```
void *consumer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&mutex);
        sem_wait(&full);
        int tmp = get();
        sem_post(&empty);
        sem_post(&mutex);
        printf("%d\n", tmp);
    }
}
```



Consumer 0 is waiting for full to be greater than or equal to 0

Adding “Mutual Exclusion”


mutex = -1

full = -1

empty = 10


Producer 0: **Running**

```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&mutex);
        sem_wait(&empty);
        put(i);
        sem_post(&full);
        sem_post(&mutex);
    }
}
```



Consumer 0: Runnable

```
void *consumer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&mutex);
        sem_wait(&full);
        int tmp = get();
        sem_post(&empty);
        sem_post(&mutex);
        printf("%d\n", tmp);
    }
}
```



Consumer 0 is **waiting** for full to be greater than or equal to 0

Adding “Mutual Exclusion”

Deadlock!!


mutex = -1

full = -1

empty = 10


Producer 0: **Running**

```
void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&mutex);
        sem_wait(&empty);
        put(i);
        sem_post(&full);
        sem_post(&mutex);
    }
}
```



Consumer 0: **Runnable**

```
void *consumer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        sem_wait(&mutex);
        sem_wait(&full);
        int tmp = get();
        sem_post(&empty);
        sem_post(&mutex);
        printf("%d\n", tmp);
    }
}
```




Producer 0 **gets stuck** at acquiring `mutex` which has been locked by Consumer 0!

Consumer 0 is **waiting** for `full` to be greater than or equal to 0

Deadlocks

- A set of threads are said to be in a *deadlock* state when *every* thread in the set is waiting for an event that can be caused *only* by another thread in the set



A typical deadlock dependency graph

Conditions for Deadlock

- **Mutual exclusion**
 - Threads claim exclusive control of resources that require e.g., a thread grabs a lock
- **Hold-and-wait**
 - Threads hold resources allocated to them while waiting for additional resources
- **No preemption**
 - Resources cannot be forcibly removed from threads that are holding them
- **Circular wait**
 - There exists a circular chain of threads such that each holds one or more resources that are being requests by next thread in chain

Correct Mutual Exclusion

```
1 sem_t empty;
2 sem_t full;
3 sem_t mutex;
4
5 void *producer(void *arg) {
6     int i;
7     for (i = 0; i < loops; i++) {
8         sem_wait(&empty);           // line p1
9         sem_wait(&mutex);           // line p1.5 (MOVED MUTEX HERE...)
10        put(i);                       // line p2
11        sem_post(&mutex);            // line p2.5 (... AND HERE)
12        sem_post(&full);             // line p3
13    }
14 }
15
16 void *consumer(void *arg) {
17     int i;
18     for (i = 0; i < loops; i++) {
19         sem_wait(&full);             // line c1
20         sem_wait(&mutex);           // line c1.5 (MOVED MUTEX HERE...)
21         int tmp = get();             // line c2
22         sem_post(&mutex);            // line c2.5 (... AND HERE)
23         sem_post(&empty);           // line c3
24         printf("%d\n", tmp);
25     }
26 }
27
28 int main(int argc, char *argv[]) {
29     // ...
30     sem_init(&empty, 0, MAX); // MAX buffers are empty to begin with...
31     sem_init(&full, 0, 0);     // ... and 0 are full
32     sem_init(&mutex, 0, 1);    // mutex=1 because it is a lock
33     // ...
34 }
```

Mutex wraps just around critical section!

Mutex wraps just around critical section!

Producer-Consumer Solution

- Make sure that
 1. The producer and the consumer do not access the buffer area and related variables at the same time
 2. No item is made available to the consumer if all the buffer slots are empty
 3. No slot in the buffer is made available to the producer if all the buffer slots are full

Semaphore Worksheet