

Operating Systems

Common Concurrency Problems

Nipun Batra

Condition Variables

- **wait (cond_t *cv, mutex_t *lock)**
 - Assumes lock is held when wait() is called
 - Puts caller to sleep + atomically releases lock
 - When awoken, reacquires lock before returning
- **signal (cond_t *cv)**
 - Wake a single waiting thread
 - If there is no waiting thread, just return, do nothing

Condition Variables **wait** Pseudocode

- **wait (cond_t *cv, mutex_t *lock)**
 - Assume lock is held initially
 - pthread_mutex_unlock(lock)
 - block_on_condition(cv)
 - Move from Ready to Waiting/Blocked State
 - Move to Ready State
 - pthread_mutex_lock(lock)

Condition Variables **wait** Pseudocode

- Which of these are atomic?
 - **wait (cond_t *cv, mutex_t *lock)**
 - Assume lock is held initially
 - pthread_mutex_unlock(lock)
 - block_on_condition(cv)
 - Move from Ready to Waiting/Blocked State
 - Move to Ready State
 - pthread_mutex_lock(lock)

Condition Variables **wait** Pseudocode

- Which of these are atomic?
- **wait (cond_t *cv, mutex_t *lock)**
 - Assume lock is held initially
 - pthread_mutex_unlock(lock)
 - block_on_condition(cv)
 - Move from Ready to Waiting/Blocked State
 - Move to Ready State
 - pthread_mutex_lock(lock)

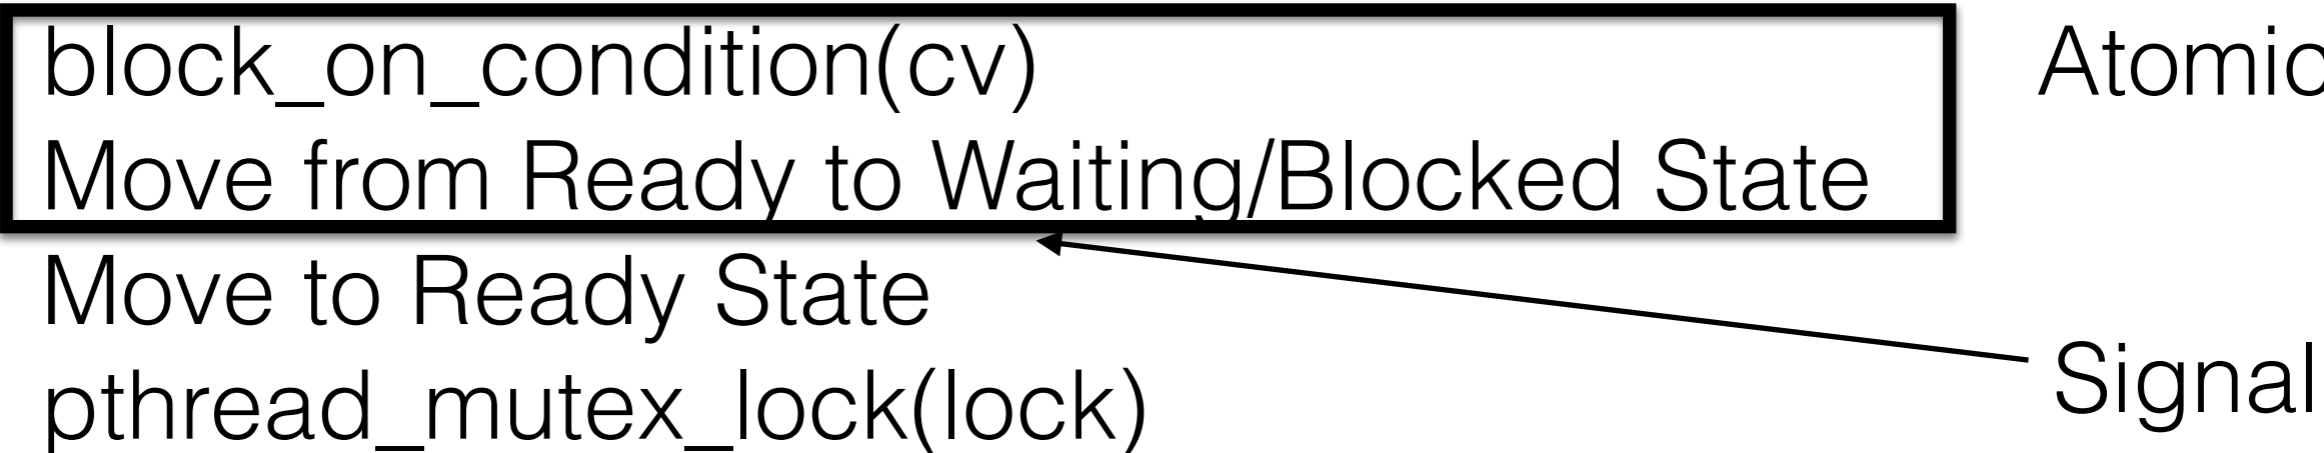
Condition Variables **wait** Pseudocode

- Which of these are atomic?
- **wait (cond_t *cv, mutex_t *lock)**
 - Assume lock is held initially
 - pthread_mutex_unlock(lock)
 - block_on_condition(cv)
 - Move from Ready to Waiting/Blocked State
 - Move to Ready State
 - pthread_mutex_lock(lock)
- Atomic

Condition Variables **wait** Pseudocode

- Which of these are atomic?
- **wait (cond_t *cv, mutex_t *lock)**
 - Assume lock is held initially
 - pthread_mutex_unlock(lock)
 - block_on_condition(cv) Atomic
 - Move from Ready to Waiting/Blocked State
 - Move to Ready State
 - pthread_mutex_lock(lock) Signal

Condition Variables **wait** Pseudocode

- Which of these are atomic?
- **wait (cond_t *cv, mutex_t *lock)**
 - Assume lock is held initially
 - pthread_mutex_unlock(lock)
 - block_on_condition(cv) Atomic
 - Move from Ready to Waiting/Blocked State
 - Move to Ready State
 - pthread_mutex_lock(lock) Signal
- 
- The diagram illustrates the state transitions and atomicity annotations for the wait function pseudocode. A box highlights the steps 'block_on_condition(cv)' and 'Move from Ready to Waiting/Blocked State', with the label 'Atomic' pointing to this box. An arrow points from the label 'Signal' to the step 'Move to Ready State'.

Condition Variables **wait** Pseudocode

- Which of these are atomic?
- **wait (cond_t *cv, mutex_t *lock)**
 - Assume lock is held initially
 - pthread_mutex_unlock(lock)
 - block_on_condition(cv) Atomic
 - Move from Ready to Waiting/Blocked State
 - Move to Ready State
 - pthread_mutex_lock(lock) Signal
-

When does wait return?

Condition Variables **wait** Pseudocode

- Which of these are atomic?
- **wait (cond_t *cv, mutex_t *lock)**
 - Assume lock is held initially
 - pthread_mutex_unlock(lock)
 - block_on_condition(cv) Atomic
 - Move from Ready to Waiting/Blocked State
 - Move to Ready State
 - pthread_mutex_lock(lock)
- Signal
-

When does wait return?

When lock has been acquired.

Condition Variables **wait** Pseudocode

- Which of these are atomic?
- **wait (cond_t *cv, mutex_t *lock)**
 - Assume lock is held initially
 - pthread_mutex_unlock(lock)
 - block_on_condition(cv) Atomic
 - Move from Ready to Waiting/Blocked State
 - Move to Ready State
 - pthread_mutex_lock(lock)
- Signal ←

When does wait return?

When lock has been acquired.

What if some other thread holds the lock — block!

Condition Variables **wait** Pseudocode

- Which of these are atomic?
- **wait (cond_t *cv, mutex_t *lock)**
 - Assume lock is held initially
 - pthread_mutex_unlock(lock)
 - block_on_condition(cv) Atomic
 - Move from Ready to Waiting/Blocked State
 - Move to Ready State ← Signal
 - pthread_mutex_lock(lock)

When does wait return?

When lock has been acquired.

What if some other thread holds the lock — block!

Why do we need to lock again?

Exercise: order using condition variables

Correct Solution

```
1 void *child(void *arg) {
2     printf("child\n");
3     thread_exit()
4     return NULL; }

7 int main(int argc, char *argv[]) {
8     printf("parent: begin\n");
9     pthread_t c;
10    Pthread_create(&c, NULL, child, NULL); // create child
11    thread_join()
12    printf("parent: end\n");
13    return 0; }
```

Exercise: order using condition variables

Correct Solution

```
void thread_exit {  
    mutex_lock(&m)  
    Done = 1  
    cond_signal(&c)  
    mutex_unlock(&m)
```

```
1 void *child(void *arg) {  
2     printf("child\n");  
3     thread_exit()  
4     return NULL; }  
  
7 int main(int argc, char *argv[]) {  
8     printf("parent: begin\n");  
9     pthread_t c;  
10    Pthread_create(&c, NULL, child, NULL); // create child  
11    thread_join()  
12    printf("parent: end\n");  
13    return 0; }
```

Exercise: order using condition variables

Correct Solution

```
void thread_exit {  
    mutex_lock(&m)  
    Done = 1  
    cond_signal(&c)  
    mutex_unlock(&m)
```

```
void thread_join {  
    mutex_lock(&m)           //w  
    while (done==0)         //x  
        cond_wait(&c, &m) //y  
    mutex_unlock(&m) }      //z
```

```
1 void *child(void *arg) {  
2     printf("child\n");  
3     thread_exit()  
4     return NULL; }  
  
7 int main(int argc, char *argv[]) {  
8     printf("parent: begin\n");  
9     pthread_t c;  
10    Pthread_create(&c, NULL, child, NULL); // create child  
11    thread_join()  
12    printf("parent: end\n");  
13    return 0; }
```

Exercise: Build a lock using semaphores

```
1 sem_t m;  
2 sem_init(&m, 0, 1);  
3  
4 sem_wait(&m);  
5 //critical section here  
6 sem_post(&m);
```

Refresher Notes

```
1 int sem_wait(sem_t *s) {  
2   s->value -= 1  
3   wait if s->value < 0  
4 }  
  
1 int sem_post(sem_t *s) {  
2   s->value += 1  
3   wake one waiting thread if any  
4 }
```




Value | Thread 0 State | Thread 1 State T2 State

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call sem_wait()	Running		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call sem_wait()	Running		Ready
-1		Ready	decrement sem	Running		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call sem_wait()	Running		Ready
-1		Ready	decrement sem	Running		Ready
-1		Ready	(sem < 0) → sleep	sleeping		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call sem_wait()	Running		Ready
-1		Ready	decrement sem	Running		Ready
-1		Ready	(sem < 0) → sleep	sleeping		Ready
-1		Running	Switch → T2	sleeping		Running

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call sem_wait()	Running		Ready
-1		Ready	decrement sem	Running		Ready
-1		Ready	(sem < 0) → sleep	sleeping		Ready
-1		Running	Switch → T2	sleeping		Running
-1					Call sem_wait()	Running

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call sem_wait()	Running		Ready
-1		Ready	decrement sem	Running		Ready
-1		Ready	(sem < 0) → sleep	sleeping		Ready
-1		Running	Switch → T2	sleeping		Running
-1					Call sem_wait()	Running
-2	(crit sect: end)	Running		sleeping	Decrement sem	Running

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call sem_wait()	Running		Ready
-1		Ready	decrement sem	Running		Ready
-1		Ready	(sem < 0) → sleep	sleeping		Ready
-1		Running	Switch → T2	sleeping		Running
-1					Call sem_wait()	Running
-2	(crit sect: end)	Running		sleeping	Decrement sem	Running
-2	call sem_post()	Running		sleeping	(sem<0) -> Sleep	Sleeping

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call sem_wait()	Running		Ready
-1		Ready	decrement sem	Running		Ready
-1		Ready	(sem < 0) → sleep	sleeping		Ready
-1		Running	Switch → T2	sleeping		Running
-1					Call sem_wait()	Running
-2	(crit sect: end)	Running		sleeping	Decrement sem	Running
-2	call sem_post()	Running		sleeping	(sem < 0) -> Sleep	Sleeping
			Switch T1			

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call sem_wait()	Running		Ready
-1		Ready	decrement sem	Running		Ready
-1		Ready	(sem < 0) → sleep	sleeping		Ready
-1		Running	Switch → T2	sleeping		Running
-1					Call sem_wait()	Running
-2	(crit sect: end)	Running		sleeping	Decrement sem	Running
-2	call sem_post()	Running		sleeping	(sem < 0) -> Sleep	Sleeping
			Switch T1			
-1	Increment sem					

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call sem_wait()	Running		Ready
-1		Ready	decrement sem	Running		Ready
-1		Ready	(sem < 0) → sleep	sleeping		Ready
-1		Running	Switch → T2	sleeping		Running
-1					Call sem_wait()	Running
-2	(crit sect: end)	Running		sleeping	Decrement sem	Running
-2	call sem_post()	Running		sleeping	(sem < 0) -> Sleep	Sleeping
			Switch T1			
-1	Increment sem					
-1	wake(T1)	Running		Ready		

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call sem_wait()	Running		Ready
-1		Ready	decrement sem	Running		Ready
-1		Ready	(sem < 0) → sleep	sleeping		Ready
-1		Running	Switch → T2	sleeping		Running
-1					Call sem_wait()	Running
-2	(crit sect: end)	Running		sleeping	Decrement sem	Running
-2	call sem_post()	Running		sleeping	(sem < 0) -> Sleep	Sleeping
			Switch T1			
-1	Increment sem					
-1	wake(T1)	Running		Ready		
-1	sem_post() returns	Running		Ready		

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call sem_wait()	Running		Ready
-1		Ready	decrement sem	Running		Ready
-1		Ready	(sem < 0) → sleep	sleeping		Ready
-1		Running	Switch → T2	sleeping		Running
-1					Call sem_wait()	Running
-2	(crit sect: end)	Running		sleeping	Decrement sem	Running
-2	call sem_post()	Running		sleeping	(sem<0) -> Sleep	Sleeping
			Switch T1			
-1	Increment sem					
-1	wake(T1)	Running		Ready		
-1	sem_post() returns	Running		Ready		
-1	Interrupt; Switch → T1	Ready		Running		

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call sem_wait()	Running		Ready
-1		Ready	decrement sem	Running		Ready
-1		Ready	(sem < 0) → sleep	sleeping		Ready
-1		Running	Switch → T2	sleeping		Running
-1					Call sem_wait()	Running
-2	(crit sect: end)	Running		sleeping	Decrement sem	Running
-2	call sem_post()	Running		sleeping	(sem < 0) -> Sleep	Sleeping
			Switch T1			
-1	Increment sem					
-1	wake(T1)	Running		Ready		
-1	sem_post() returns	Running		Ready		
-1	Interrupt; Switch → T1	Ready		Running		

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call sem_wait()	Running		Ready		Ready
0	sem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call sem_wait()	Running		Ready
-1		Ready	decrement sem	Running		Ready
-1		Ready	(sem < 0) → sleep	sleeping		Ready
-1		Running	Switch → T2	sleeping		Running
-1					Call sem_wait()	Running
-2	(crit sect: end)	Running		sleeping	Decrement sem	Running
-2	call sem_post()	Running		sleeping	(sem<0) -> Sleep	Sleeping
			Switch T1			
-1	Increment sem					
-1	wake(T1)	Running		Ready		
-1	sem_post() returns	Running		Ready		
-1	Interrupt; Switch → T1	Ready		Running		

Can T1 run before sem_post returns?

NO!

Semaphores Implementation

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API	Our implementation

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

```
1 #include <semaphore.h>
2 sem_t s;
```

Our implementation

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

```
1 #include <semaphore.h>
2 sem_t s;
```

Our implementation

```
1 typedef struct __Zem_t {
2     int value;
3     pthread_cond_t cond;
4     pthread_mutex_t lock;
5 } Zem_t;
6
```

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

API	
-----	--

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API	Our implementation

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

```
1 int sem_init(sem_t *s,  
int init_val) {  
2   s->value=init_val;  
3 }
```

Our implementation

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

```
1 int sem_init(sem_t *s,  
int init_val) {  
2     s->value=init_val;  
3 }
```

Our implementation

```
1 // only one thread can call this  
2 void Zem_init(Zem_t *s, int value) {  
3     s->value = value;  
4     Cond_init(&s->cond);  
5     Mutex_init(&s->lock);  
6 }
```

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

API	
-----	--

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API	Our implementation

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

Our implementation

```
1 int sem_post(sem_t *s) {  
2     s->value += 1  
3     wake one waiting thread if  
any  
4 }
```

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

Our implementation

```
1 int sem_post(sem_t *s) {  
2     s->value += 1  
3     wake one waiting thread if  
any  
4 }
```

Atomic operation

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

```
1 int sem_post(sem_t *s) {  
2     s->value += 1  
3     wake one waiting thread if  
any  
4 }
```

Our implementation

```
void Zem_post(Zem_t *s) {  
23 Mutex_lock(&s->lock);  
24 s->value++;  
25 Cond_signal(&s->cond);  
26 Mutex_unlock(&s->lock);  
27 }
```

Atomic operation

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

API	
-----	--

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API	Our implementation

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

Our implementation

```
1 int sem_wait(sem_t *s) {  
2   s->value -= 1  
3   wait if s->value < 0  
4 }
```

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

```
1 int sem_wait(sem_t *s) {  
2   s->value -= 1  
3   wait if s->value < 0  
4 }
```

Our implementation

Atomic operation

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

```
1 int sem_wait(sem_t *s) {  
2   s->value -= 1  
3   wait if s->value < 0  
4 }
```

Atomic operation

Our implementation

```
void Zem_wait(Zem_t *s) {  
1   Mutex_lock(&s->lock);  
2   while (s->value <= 0)  
3     Cond_wait(&s->cond, &s->lock);  
4   s->value--;  
5   Mutex_unlock(&s->lock);  
6 }
```

Semaphores Implementation

- Build semaphores using **locks** and **condition variables**
- Any critical section should require **locking**
- **Signal** and **Wait** on condition
- Don't maintain the invariant that the value of the semaphore, when negative, reflects the number of waiting threads

API

```
1 int sem_wait(sem_t *s) {  
2   s->value -= 1  
3   wait if s->value < 0  
4 }
```

Atomic operation

Our implementation

```
void Zsem_wait(Zsem_t *s) {  
1   Mutex_lock(&s->lock);  
2   while (s->value <= 0) Should it be ==0?  
3   Cond_wait(&s->cond, &s->lock);  
4   s->value--;  
5   Mutex_unlock(&s->lock);  
6 }
```



Value	Thread 0	State	Thread 1	State	T2	State
-------	----------	-------	----------	-------	----	-------

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call zem_wait()	Running		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call zem_wait()	Running		Ready
0		Ready	(zem ≤ 0) → sleep	sleeping		Ready

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call zem_wait()	Running		Ready
0		Ready	(zem ≤ 0) → sleep	sleeping		Ready
0		Running	Switch → T2	sleeping		Running

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call zem_wait()	Running		Ready
0		Ready	(zem ≤ 0) → sleep	sleeping		Ready
0		Running	Switch → T2	sleeping		Running
0					Call zem_wait()	Running

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call zem_wait()	Running		Ready
0		Ready	(zem ≤ 0) → sleep	sleeping		Ready
0		Running	Switch → T2	sleeping		Running
0					Call zem_wait()	Running
0					(zem < 0) -> Sleep	Sleeping

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call zem_wait()	Running		Ready
0		Ready	(zem ≤ 0) → sleep	sleeping		Ready
0		Running	Switch → T2	sleeping		Running
0					Call zem_wait()	Running
0					(zem < 0) -> Sleep	Sleeping
0	(crit sect: end)	Running		sleeping	Switch → T1	Sleeping

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call zem_wait()	Running		Ready
0		Ready	(zem ≤ 0) → sleep	sleeping		Ready
0		Running	Switch → T2	sleeping		Running
0					Call zem_wait()	Running
0					(zem < 0) -> Sleep	Sleeping
0	(crit sect: end)	Running		sleeping	Switch → T1	Sleeping
0	call sem_post()	Running		sleeping		

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call zem_wait()	Running		Ready
0		Ready	(zem ≤ 0) → sleep	sleeping		Ready
0		Running	Switch → T2	sleeping		Running
0					Call zem_wait()	Running
0					(zem < 0) -> Sleep	Sleeping
0	(crit sect: end)	Running		sleeping	Switch → T1	Sleeping
0	call sem_post()	Running		sleeping		
	Acquires lock	Running				

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call zem_wait()	Running		Ready
0		Ready	(zem ≤ 0) → sleep	sleeping		Ready
0		Running	Switch → T2	sleeping		Running
0					Call zem_wait()	Running
0					(zem < 0) -> Sleep	Sleeping
0	(crit sect: end)	Running		sleeping	Switch → T1	Sleeping
0	call sem_post()	Running		sleeping		
	Acquires lock	Running				
1	Increments zem	Running				

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call zem_wait()	Running		Ready
0		Ready	(zem ≤ 0) → sleep	sleeping		Ready
0		Running	Switch → T2	sleeping		Running
0					Call zem_wait()	Running
0					(zem < 0) -> Sleep	Sleeping
0	(crit sect: end)	Running		sleeping	Switch → T1	Sleeping
0	call sem_post()	Running		sleeping		
	Acquires lock	Running				
1	Increments zem	Running				
1	Wakes up thread T1	Running		Ready		

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call zem_wait()	Running		Ready
0		Ready	(zem ≤ 0) → sleep	sleeping		Ready
0		Running	Switch → T2	sleeping		Running
0					Call zem_wait()	Running
0					(zem < 0) -> Sleep	Sleeping
0	(crit sect: end)	Running		sleeping	Switch → T1	Sleeping
0	call sem_post()	Running		sleeping		
	Acquires lock	Running				
1	Increments zem	Running				
1	Wakes up thread T1	Running		Ready		
			Condition wait will return once it gets lock	Ready		

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call zem_wait()	Running		Ready
0		Ready	(zem ≤ 0) → sleep	sleeping		Ready
0		Running	Switch → T2	sleeping		Running
0					Call zem_wait()	Running
0					(zem < 0) -> Sleep	Sleeping
0	(crit sect: end)	Running		sleeping	Switch → T1	Sleeping
0	call sem_post()	Running		sleeping		
	Acquires lock	Running				
1	Increments zem	Running				
1	Wakes up thread T1	Running		Ready		
			Condition wait will return once it gets lock	Ready		
1	zem_post() returns	Running		Ready		

Value	Thread 0	State	Thread 1	State	T2	State
1		Running		Ready		Ready
1	call zem_wait()	Running		Ready		Ready
0	zem_wait() retruns	Running		Ready		Ready
0	(crit set: begin)	Running		Ready		Ready
0	Interrupt; Switch → T1	Ready		Running		Ready
0		Ready	call zem_wait()	Running		Ready
0		Ready	(zem ≤ 0) → sleep	sleeping		Ready
0		Running	Switch → T2	sleeping		Running
0					Call zem_wait()	Running
0					(zem < 0) -> Sleep	Sleeping
0	(crit sect: end)	Running		sleeping	Switch → T1	Sleeping
0	call sem_post()	Running		sleeping		
	Acquires lock	Running				
1	Increments zem	Running				
1	Wakes up thread T1	Running		Ready		
			Condition wait will return once it gets lock	Ready		
1	zem_post() returns	Running		Ready		
0			Zem = zem - 1	Running		

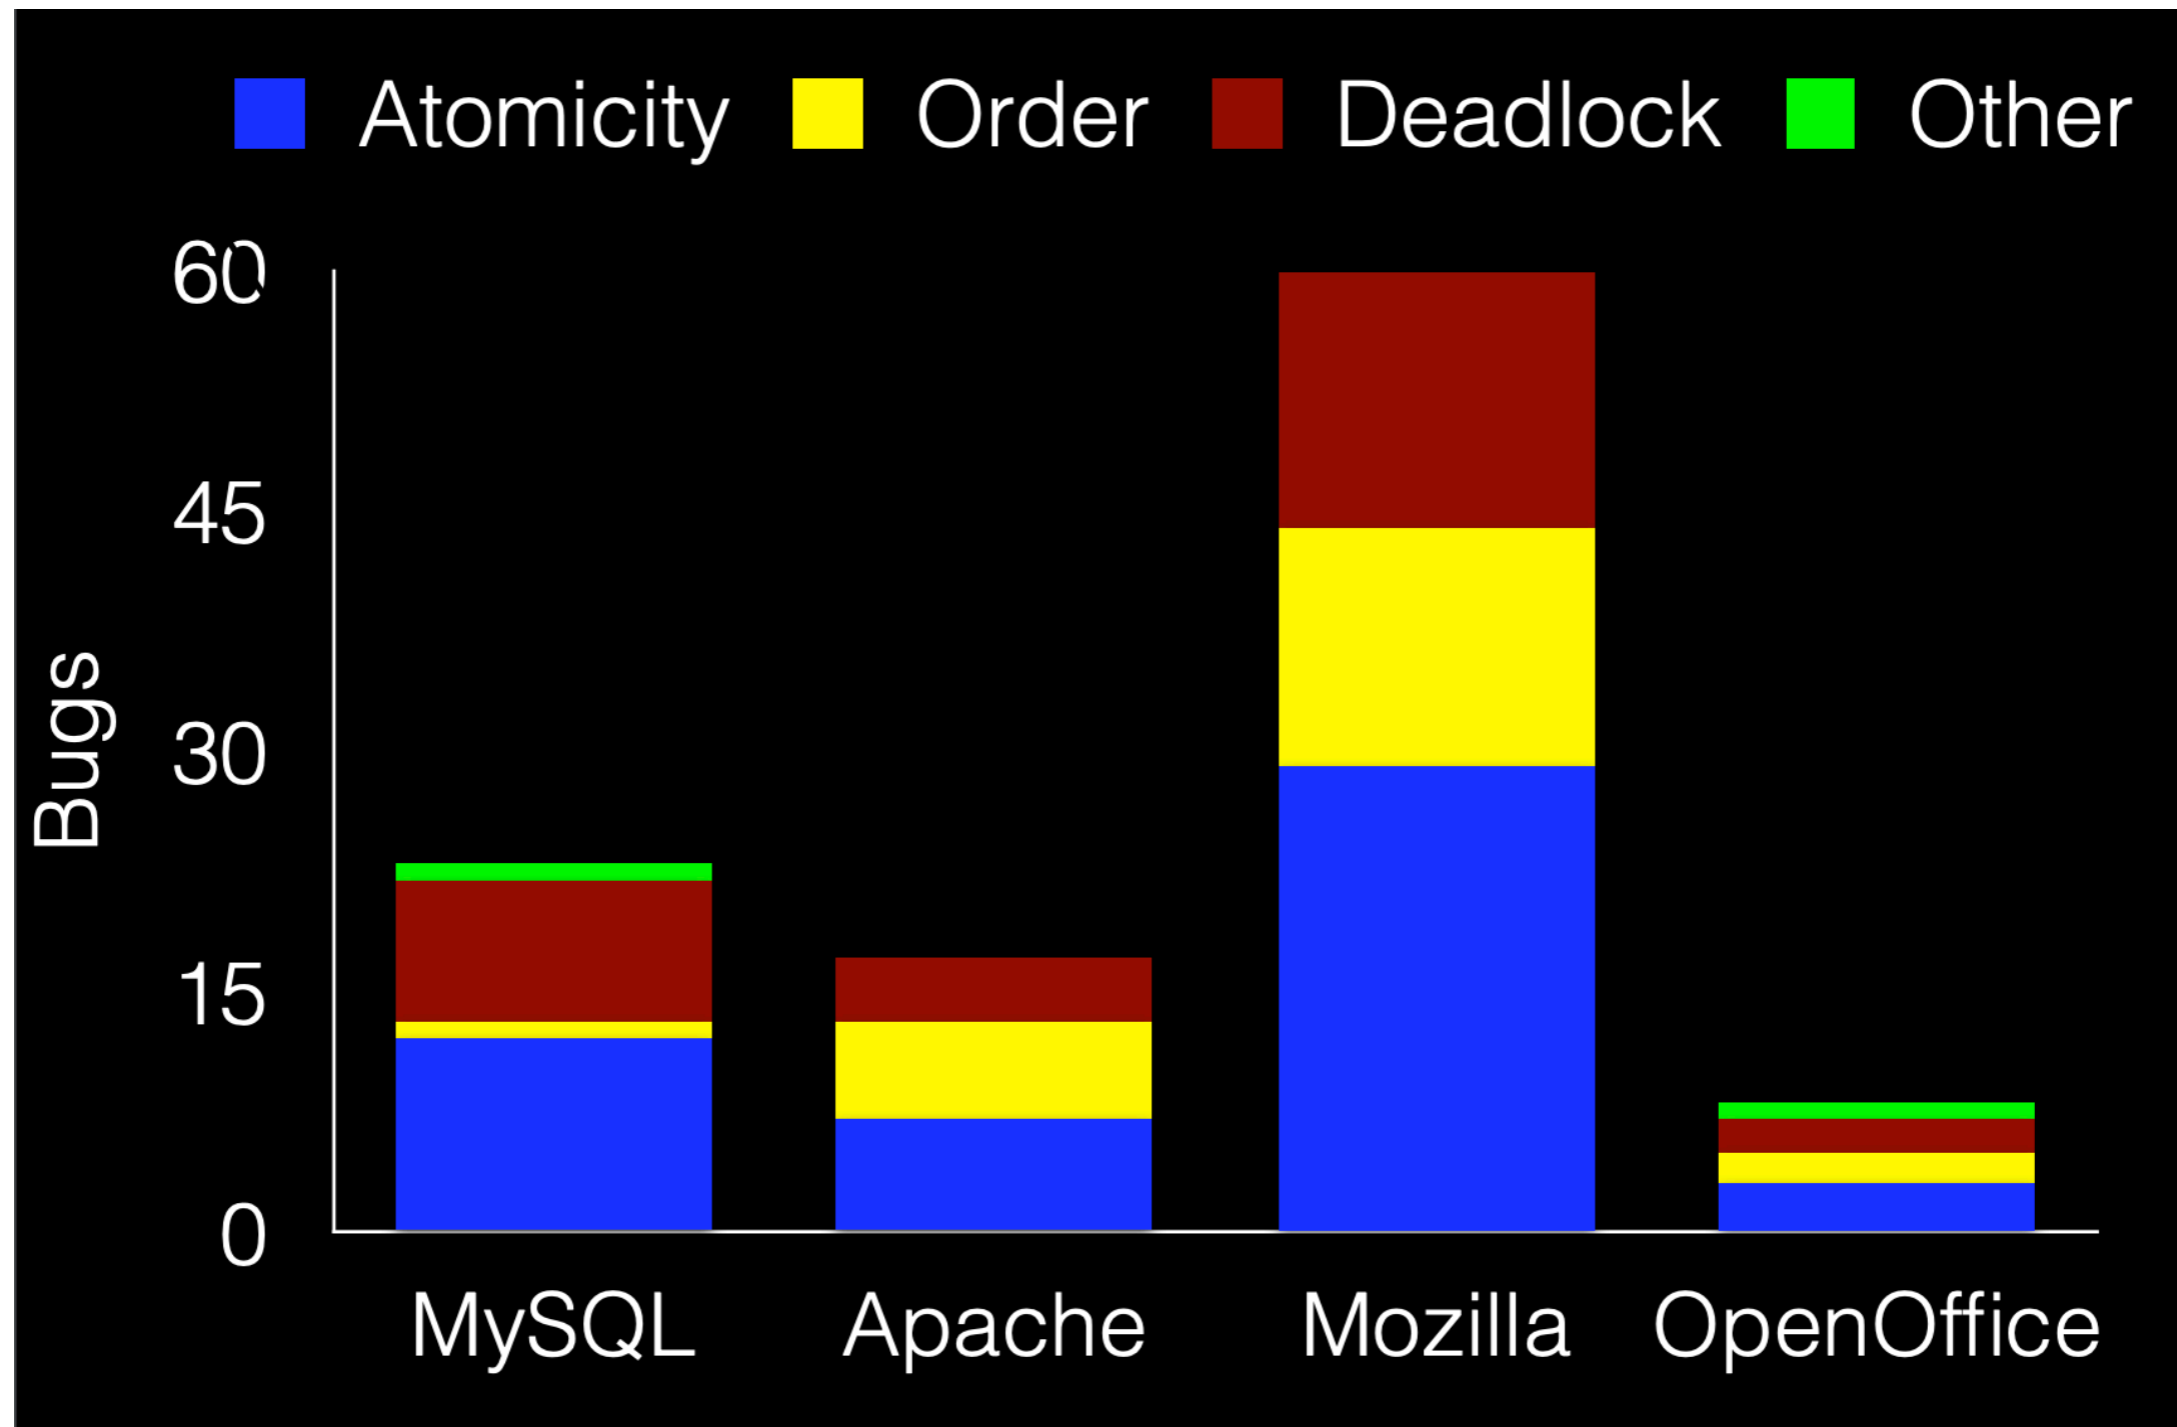
Another Implementation ...

```
typedef struct {
    int value;
    struct process *list;
} semaphore;

wait(semaphore *S) {
    S->value--;
    if (S->value < 0) {
        add this process to S->list;
        block();
    }
}

signal(semaphore *S) {
    S->value++;
    if (S->value <= 0) {
        remove a process P from S->list;
        wakeup(P);
    }
}
```

Concurrency Bugs



Types of bugs in 4 major projects from 500K bug reports

Concurrency Bugs — Atomicity

Concurrency Bugs — Atomicity

MySQL bug ...

Concurrency Bugs — Atomicity

MySQL bug ...

1 Thread1:

```
2  if(thd->proc_info){
3    ...
4    fputs(thd->proc_info , ...);
5  ...
6  }
```

Concurrency Bugs — Atomicity

MySQL bug ...

1 Thread1::

```
2  if(thd->proc_info){
3      ...
4      fputs(thd->proc_info , ...);
5  ...
6  }
```

8 Thread2::

```
9  thd->proc_info = NULL;
```

Concurrency Bugs — Atomicity

MySQL bug ...

1 Thread1::

```
2  if(thd->proc_info){
3      ...
4      fputs(thd->proc_info , ...);
5  ...
6  }
```

8 Thread2::

```
9  thd->proc_info = NULL;
```

- Is this problematic?

Concurrency Bugs — Atomicity

MySQL bug ...

1 Thread1::

```
2  if(thd->proc_info){
3      ...
4      fputs(thd->proc_info , ...);
5  ...
6  }
```

8 Thread2::

```
9  thd->proc_info = NULL;
```

- Is this problematic?
 - Yes, else we wouldn't be discussing ...

Concurrency Bugs — Atomicity

MySQL bug ...

1 Thread1::

```
2  if(thd->proc_info){
3    ...
4    fputs(thd->proc_info , ...);
5  ...
6  }
```

8 Thread2::

```
9  thd->proc_info = NULL;
```

- Is this problematic?
 - Yes, else we wouldn't be discussing ...
 - How?

Concurrency Bugs — Atomicity

```
1 pthread_mutex_t lock =  
PTHREAD_MUTEX_INITIALIZER;  
2
```

3 Thread1::

```
4 pthread_mutex_lock(&lock);  
5 if(thd->proc_info){  
6   ...  
7   fputs(thd->proc_info , ...);  
8   ...  
9 }  
10 pthread_mutex_unlock(&lock);
```

1 Thread2::

```
2 pthread_mutex_lock(&lock);  
3 thd->proc_info = NULL;  
4 pthread_mutex_unlock(&lock);
```

Concurrency Bugs — Atomicity

Simple Solution

```
1 pthread_mutex_t lock =  
PTHREAD_MUTEX_INITIALIZER;  
2
```

3 Thread1::

```
4 pthread_mutex_lock(&lock);  
5 if(thd->proc_info){  
6   ...  
7   fputs(thd->proc_info , ...);  
8   ...  
9 }  
10 pthread_mutex_unlock(&lock);
```

1 Thread2::

```
2 pthread_mutex_lock(&lock);  
3 thd->proc_info = NULL;  
4 pthread_mutex_unlock(&lock);
```

Concurrency Bugs — Order Violation

1 Thread1::

```
2 void init(){  
3   mThread =  
PR_CreateThread(mMain, ...);  
4 }  
5
```

6 Thread2::

```
7 void mMain(...){  
8   mState = mThread->State  
9 }
```

Concurrency Bugs — Order Violation

Mozilla bug ...

1 Thread1::

```
2 void init(){  
3   mThread =  
PR_CreateThread(mMain, ...);  
4 }  
5
```

6 Thread2::

```
7 void mMain(...){  
8   mState = mThread->State  
9 }
```

Concurrency Bugs — Order Violation

Mozilla bug ...

1 Thread1::

```
2 void init(){
3   mThread =
4   PR_CreateThread(mMain, ...);
5 }
```

6 Thread2::

```
7 void mMain(...){
8   mState = mThread->State
9 }
```

- Is this problematic?

Concurrency Bugs — Order Violation

Mozilla bug ...

1 Thread1::

```
2 void init(){
3   mThread =
4   PR_CreateThread(mMain, ...);
5 }
```

6 Thread2::

```
7 void mMain(...){
8   mState = mThread->State
9 }
```

- Is this problematic?
 - Yes, else we wouldn't be discussing ...

Concurrency Bugs — Order Violation

Mozilla bug ...

1 Thread1::

```
2 void init(){
3   mThread =
PR_CreateThread(mMain, ...);
4 }
5
```

6 Thread2::

```
7 void mMain(...){
8   mState = mThread->State
9 }
```

- Is this problematic?
 - Yes, else we wouldn't be discussing ...
 - How?

Concurrency Bugs — Order Violation

Concurrency Bugs — Order Violation

```
1 pthread_mutex_t mtLock = PTHREAD_MUTEX_INITIALIZER;  
2 pthread_cond_t mtCond = PTHREAD_COND_INITIALIZER;  
3 int mtInit = 0;
```

Concurrency Bugs — Order Violation

```
1 pthread_mutex_t mtLock = PTHREAD_MUTEX_INITIALIZER;
2 pthread_cond_t mtCond = PTHREAD_COND_INITIALIZER;
3 int mtInit = 0;
```

1 Thread 1::

```
2 void init(){
3   ...
4   mThread = PR_CreateThread(mMain,...);
5
6   // signal that the thread has been created.
7   pthread_mutex_lock(&mtLock);
8   mtInit = 1;
9   pthread_cond_signal(&mtCond);
10  pthread_mutex_unlock(&mtLock);
11  ...
12 }
```

Concurrency Bugs — Order Violation

```
1 pthread_mutex_t mtLock = PTHREAD_MUTEX_INITIALIZER;
2 pthread_cond_t mtCond = PTHREAD_COND_INITIALIZER;
3 int mtInit = 0;
```

1 Thread 1::

```
2 void init(){
3   ...
4   mThread = PR_CreateThread(mMain,...);
5
6   // signal that the thread has been created.
7   pthread_mutex_lock(&mtLock);
8   mtInit = 1;
9   pthread_cond_signal(&mtCond);
10  pthread_mutex_unlock(&mtLock);
11  ...
12 }
```

20 Thread2::

```
21 void mMain(...){
    // wait for the thread to be initialized
    ...
22  pthread_mutex_lock(&mtLock);
23  while(mtInit == 0)
24    pthread_cond_wait(&mtCond,
&mtLock);
25
    pthread_mutex_unlock(&mtLock);
26  mState = mThread->State;
}
```

Concurrency Bugs — Deadlock

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

Thread 2

Lock(L2);

Lock(L1);

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

Thread 2

Lock(L2);

Lock(L1);

- Thread T1 gets Lock L1

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

Thread 2

Lock(L2);

Lock(L1);

- Thread T1 gets Lock L1
- Thread T1 gets Lock L2

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

Thread 2

Lock(L2);

Lock(L1);

- Thread T1 gets Lock L1
- Thread T1 gets Lock L2
- Thread T1 completes critical section

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

Thread 2

Lock(L2);

Lock(L1);

- Thread T1 gets Lock L1
- Thread T1 gets Lock L2
- Thread T1 completes critical section
- Context Switch

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

Thread 2

Lock(L2);

Lock(L1);

- Thread T1 gets Lock L1
- Thread T1 gets Lock L2
- Thread T1 completes critical section
- Context Switch
- Thread T2 gets Lock L2 and Lock L1

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

Thread 2

Lock(L2);

Lock(L1);

- Thread T1 gets Lock L1
- Thread T1 gets Lock L2
- Thread T1 completes critical section
- Context Switch
- Thread T2 gets Lock L2 and Lock L1
- Works :)

Concurrency Bugs — Deadlock

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

Thread 2

Lock(L2);

Lock(L1);

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

Thread 2

Lock(L2);

Lock(L1);

- Thread T1 gets Lock L1

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

- Thread T1 gets Lock L1
- Context Switch

Thread 2

Lock(L2);

Lock(L1);

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

- Thread T1 gets Lock L1
- Context Switch
- Thread T2 gets Lock L2

Thread 2

Lock(L2);

Lock(L1);

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

- Thread T1 gets Lock L1
- Context Switch
- Thread T2 gets Lock L2
- Context Switch

Thread 2

Lock(L2);

Lock(L1);

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

Thread 2

Lock(L2);

Lock(L1);

- Thread T1 gets Lock L1
- Context Switch
- Thread T2 gets Lock L2
- Context Switch
- Thread T1 waits since it doesn't have Lock 2

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

Thread 2

Lock(L2);

Lock(L1);

- Thread T1 gets Lock L1
- Context Switch
- Thread T2 gets Lock L2
- Context Switch
- Thread T1 waits since it doesn't have Lock 2
- Context Switch

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

Lock(L2);

Thread 2

Lock(L2);

Lock(L1);

- Thread T1 gets Lock L1
- Context Switch
- Thread T2 gets Lock L2
- Context Switch
- Thread T1 waits since it doesn't have Lock 2
- Context Switch
- Thread t2 waits since it doesn't have Lock 1

Concurrency Bugs — Deadlock

Thread 1

Lock(L1);

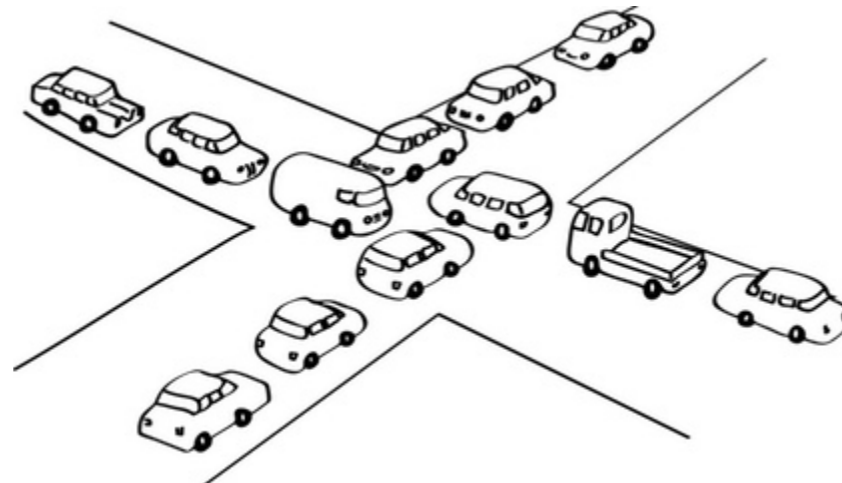
Lock(L2);

- Thread T1 gets Lock L1
- Context Switch
- Thread T2 gets Lock L2
- Context Switch
- Thread T1 waits since it doesn't have Lock 2
- Context Switch
- Thread t2 waits since it doesn't have Lock 1

Thread 2

Lock(L2);

Lock(L1);



Concurrency Bugs — Deadlock

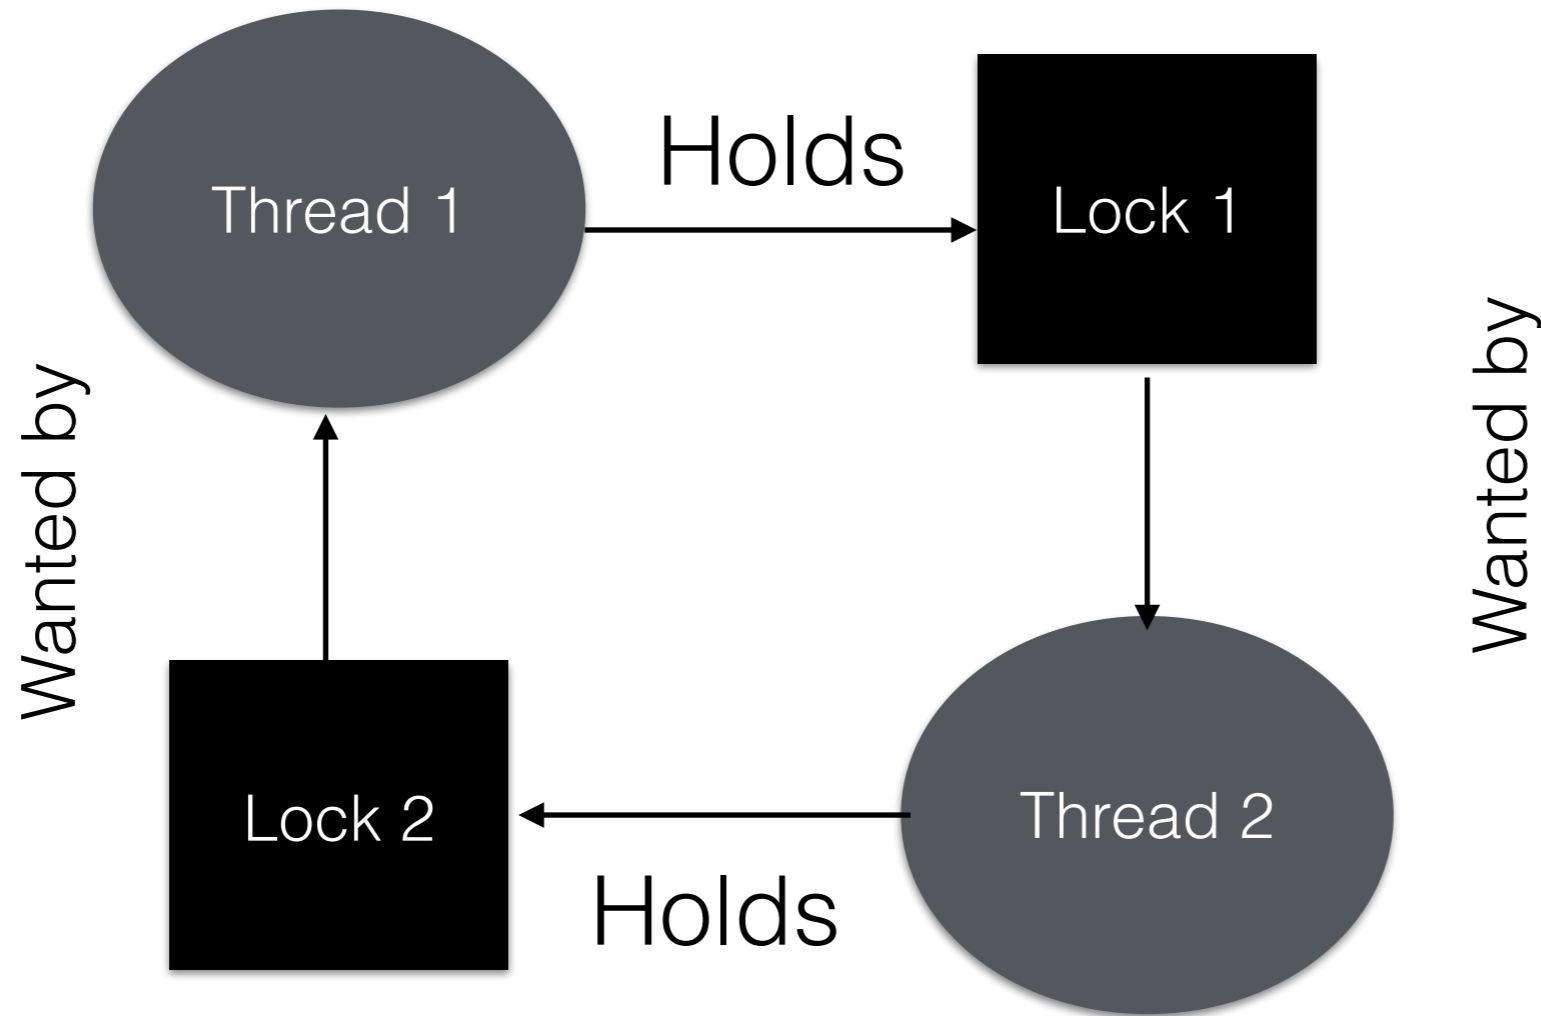
Dependency Graphs

Thread 1

Lock(L1);
Lock(L2);

Thread 2

Lock(L2);
Lock(L1);



Why Deadlocks Occur

Why Deadlocks Occur

- Encapsulation

Why Deadlocks Occur

- Encapsulation
 - Example: Java vector addAll method

Why Deadlocks Occur

- Encapsulation
 - Example: Java vector addAll method
 - $V1 = [1, 4, 5]$

Why Deadlocks Occur

- Encapsulation
 - Example: Java vector addAll method
 - $V1 = [1, 4, 5]$
 - $V2 = [6, 7, 8]$

Why Deadlocks Occur

- Encapsulation
 - Example: Java vector addAll method
 - $V1 = [1, 4, 5]$
 - $V2 = [6, 7, 8]$
 - $V1.addAll(V2) \rightarrow [1, 4, 5, 6, 7, 8]$

Why Deadlocks Occur

- Encapsulation
 - Example: Java vector addAll method
 - $V1 = [1, 4, 5]$
 - $V2 = [6, 7, 8]$
 - $V1.addAll(V2) \longrightarrow [1, 4, 5, 6, 7, 8]$
 - If addAll is multithreaded & assuming it locks V1, V2 in that order

Why Deadlocks Occur

- Encapsulation
 - Example: Java vector addAll method
 - $V1 = [1, 4, 5]$
 - $V2 = [6, 7, 8]$
 - $V1.addAll(V2) \longrightarrow [1, 4, 5, 6, 7, 8]$
 - If addAll is multithreaded & assuming it locks V1, V2 in that order
 - If some thread all calls $V2.add(V1)$, we have a deadlock!

Why Deadlocks Occur

- Encapsulation
 - Example: Java vector addAll method
 - $V1 = [1, 4, 5]$
 - $V2 = [6, 7, 8]$
 - $V1.addAll(V2) \rightarrow [1, 4, 5, 6, 7, 8]$
 - If addAll is multithreaded & assuming it locks V1, V2 in that order
 - If some thread all calls $V2.add(V1)$, we have a deadlock!
- Complex dependencies

Why Deadlocks Occur

- Encapsulation
 - Example: Java vector addAll method
 - $V1 = [1, 4, 5]$
 - $V2 = [6, 7, 8]$
 - $V1.addAll(V2) \rightarrow [1, 4, 5, 6, 7, 8]$
 - If addAll is multithreaded & assuming it locks V1, V2 in that order
 - If some thread all calls $V2.add(V1)$, we have a deadlock!
- Complex dependencies
 - Virtual memory system calls filesystem

Why Deadlocks Occur

- Encapsulation
 - Example: Java vector addAll method
 - $V1 = [1, 4, 5]$
 - $V2 = [6, 7, 8]$
 - $V1.addAll(V2) \rightarrow [1, 4, 5, 6, 7, 8]$
 - If addAll is multithreaded & assuming it locks V1, V2 in that order
 - If some thread all calls $V2.add(V1)$, we have a deadlock!
- Complex dependencies
 - Virtual memory system calls filesystem
 - Filesystem calls virtual memory

Concurrency Bugs — Deadlock

Dependency Graphs

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
-----------	-------------

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
Hold-and-wait	Threads hold resources allocated to them while waiting for additional resources

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
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.

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
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 thread holds one or more resources that are being requested by the next thread in the chain

Concurrency Bugs — Deadlock

Dependency Graphs

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
-----------	-------------

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
Hold-and-wait	Threads hold resources allocated to them while waiting for additional resources

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
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.

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
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 thread holds one or more resources that are being requested by the next thread in the chain

Preventing Circular Wait

Thread 1

Lock(L1);

Lock(L2);

Deadlock Version

Thread 2

Lock(L2);

Lock(L1);

Preventing Circular Wait

- Provide a total ordering of lock acquisition

Thread 1

Lock(L1);

Lock(L2);

Deadlock Version

Thread 2

Lock(L2);

Lock(L1);

Preventing Circular Wait

- Provide a total ordering of lock acquisition

Thread 1

Lock(L1);

Lock(L2);

Deadlock Version

Thread 2

Lock(L2);

Lock(L1);

Preventing Circular Wait

- Provide a total ordering of lock acquisition

Thread 1

Lock(L1);
Lock(L2);

Deadlock Version

Thread 2

Lock(L2);
Lock(L1);

Thread 1

Lock(L1);
Lock(L2);

Non-deadlock Version

Thread 2

Lock(L1);
Lock(L2);

Preventing Circular Wait

Preventing Circular Wait

- Provide a total ordering of lock acquisition

Preventing Circular Wait

- Provide a total ordering of lock acquisition

Preventing Circular Wait

- Provide a total ordering of lock acquisition
- Define a function `do_something` which works correct even if two threads call it as:

Preventing Circular Wait

- Provide a total ordering of lock acquisition
- Define a function `do_something` which works correct even if two threads call it as:
 - T1 — `do_something (L1, L2)` and

Preventing Circular Wait

- Provide a total ordering of lock acquisition
- Define a function `do_something` which works correct even if two threads call it as:
 - T1 — `do_something (L1, L2)` and
 - T2 — `do_something (L2, L1)`

Preventing Circular Wait

- Provide a total ordering of lock acquisition
- Define a function `do_something` which works correct even if two threads call it as:
 - T1 — `do_something (L1, L2)` and
 - T2 — `do_something (L2, L1)`

`do something(mutex t *m1, mutex t *m2)`

Preventing Circular Wait

- Provide a total ordering of lock acquisition
- Define a function `do_something` which works correct even if two threads call it as:
 - T1 — `do_something (L1, L2)` and
 - T2 — `do_something (L2, L1)`

```
do something(mutex t *m1, mutex t *m2)
```

```
if (m1 > m2)
```

```
{ // grab locks in high-to-low address order
```

```
pthread_mutex_lock(m1);
```

```
pthread_mutex_lock(m2); }
```

```
else {
```

```
pthread_mutex_lock(m2);
```

```
pthread_mutex_lock(m1); }
```

```
// Code assumes that m1 != m2 (it is not the same lock)
```

Concurrency Bugs — Deadlock

Dependency Graphs

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
-----------	-------------

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
Hold-and-wait	Threads hold resources allocated to them while waiting for additional resources

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
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.

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
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 thread holds one or more resources that are being requested by the next thread in the chain

Preventing Hold and Wait

Thread 1

Lock(L1);

Lock(L2);

Deadlock Version

Thread 2

Lock(L2);

Lock(L1);

Preventing Hold and Wait

- Acquire **all** locks **at once atomically**

Thread 1

Lock(L1);

Lock(L2);

Deadlock Version

Thread 2

Lock(L2);

Lock(L1);

Preventing Hold and Wait

- Acquire **all** locks **at once atomically**

Thread 1

Lock(L1);

Lock(L2);

Deadlock Version

Thread 2

Lock(L2);

Lock(L1);

Preventing Hold and Wait

- Acquire **all** locks **at once atomically**

Thread 1

Lock(L1);
Lock(L2);

Deadlock Version

Thread 2

Lock(L2);
Lock(L1);

Thread 1

Lock(ALL)

Lock(L1);
Lock(L2);

.....

Unlock(ALL);

Non-deadlock Version

Thread 2

Lock(ALL)

Lock(L1);
Lock(L2);

...

Unlock(ALL)

Preventing Hold and Wait

- Acquire **all** locks **at once atomically**

Thread 1

Lock(ALL)

Lock(L1);

Lock(L2);

.....

Unlock(ALL);

Non-deadlock Version

Thread 2

Lock(ALL)

Lock(L1);

Lock(L2);

...

Unlock(ALL)

Preventing Hold and Wait

- Acquire **all** locks **at once atomically**
- Cons

Thread 1

Lock(ALL)

Lock(L1);

Lock(L2);

.....

Unlock(ALL);

Non-deadlock Version

Thread 2

Lock(ALL)

Lock(L1);

Lock(L2);

...

Unlock(ALL)

Preventing Hold and Wait

- Acquire **all** locks **at once atomically**
- Cons
 - Requires us to know which all locks will be required ahead of time

Thread 1

Lock(ALL)

Lock(L1);

Lock(L2);

.....

Unlock(ALL);

Non-deadlock Version

Thread 2

Lock(ALL)

Lock(L1);

Lock(L2);

...

Unlock(ALL)

Preventing Hold and Wait

- Acquire **all** locks **at once atomically**
- Cons
 - Requires us to know which all locks will be required ahead of time
 - Reduction of concurrency

Thread 1

Lock(ALL)

Lock(L1);

Lock(L2);

.....

Unlock(ALL);

Non-deadlock Version

Thread 2

Lock(ALL)

Lock(L1);

Lock(L2);

...

Unlock(ALL)

Concurrency Bugs — Deadlock

Dependency Graphs

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
-----------	-------------

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
Hold-and-wait	Threads hold resources allocated to them while waiting for additional resources

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
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.

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
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 thread holds one or more resources that are being requested by the next thread in the chain

Pre-emption

Thread 1

Lock(L1);
Lock(L2);

Deadlock Version

Thread 2

Lock(L2);
Lock(L1);

Non-deadlock Version

```
1 top:  
2 lock(L1);  
3 if( tryLock(L2) == -1 ){  
4 unlock(L1);  
5 goto top;  
6 }
```

Pre-emption

Thread 1

Lock(L1);
Lock(L2);

Deadlock Version

Non-deadlock Version

```
1 top:  
2 lock(L1);  
3 if( tryLock(L2) == -1 ){  
4 unlock(L1);  
5 goto top;  
6 }
```

Thread 2

Lock(L2);
Lock(L1);

Pre-emption

- Livelock: Both threads running this sequence repeatedly

Thread 1

Lock(L1);
Lock(L2);

Deadlock Version

Non-deadlock Version

```
1 top:  
2 lock(L1);  
3 if( tryLock(L2) == -1 ){  
4 unlock(L1);  
5 goto top;  
6 }
```

Thread 2

Lock(L2);
Lock(L1);

Pre-emption

- Livelock: Both threads running this sequence repeatedly
- How to solve?

Thread 1

Lock(L1);
Lock(L2);

Deadlock Version

Non-deadlock Version

```
1 top:  
2 lock(L1);  
3 if( tryLock(L2) == -1 ){  
4 unlock(L1);  
5 goto top;  
6 }
```

Thread 2

Lock(L2);
Lock(L1);

Pre-emption

- Livelock: Both threads running this sequence repeatedly
- How to solve?
 - Add random delay

Thread 1

Lock(L1);
Lock(L2);

Deadlock Version

Non-deadlock Version

```
1 top:  
2 lock(L1);  
3 if( tryLock(L2) == -1 ){  
4 unlock(L1);  
5 goto top;  
6 }
```

Thread 2

Lock(L2);
Lock(L1);

Concurrency Bugs — Deadlock

Dependency Graphs

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
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 thread holds one or more resources that are being requested by the next thread in the chain

Prevention — Mutual exclusion

Prevention — Mutual exclusion

- Use atomic instructions!

Prevention — Mutual exclusion

- Use atomic instructions!

```
1 int CompareAndSwap(int *address, int expected, int new){  
2   if(*address == expected){  
3     *address = new;  
4     return 1; // success  
5   }  
6   return 0;  
7 }
```

Prevention — Mutual exclusion

- Use atomic instructions!

```
1 int CompareAndSwap(int *address, int expected, int new){  
2   if(*address == expected){  
3     *address = new;  
4     return 1; // success  
5   }  
6   return 0;  
7 }
```

- Use above code to implement atomic increment: $x = x+k$

Prevention — Mutual exclusion

- Use atomic instructions!

```
1 int CompareAndSwap(int *address, int expected, int new){
2   if(*address == expected){
3     *address = new;
4     return 1; // success
5   }
6   return 0;
7 }
```

- Use above code to implement atomic increment: $x = x+k$

```
1 void AtomicIncrement(int *value, int amount){
...Fill Here
5 }
```

Prevention — Mutual exclusion

- Use atomic instructions!

```
1 int CompareAndSwap(int *address, int expected, int new){
2   if(*address == expected){
3     *address = new;
4     return 1; // success
5   }
6   return 0;
7 }
```

- Use above code to implement atomic increment: $x = x+k$

```
1 void AtomicIncrement(int *value, int amount){
2   do{
3     int old = *value;
4   }while( CompareAndSwap(value, old, old+amount)!=0);
5 }
```

Prevention — Mutual exclusion

- List insertion using atomic instructions

Prevention — Mutual exclusion

- List insertion using atomic instructions

```
1 void insert(int value){
2  node_t * n = malloc(sizeof(node_t));
3  assert( n != NULL );
4  n->value = value ;
5  n->next = head;
6  head = n;
7 }
```

Prevention — Mutual exclusion

- List insertion using atomic instructions

```
1 void insert(int value){
2  node_t * n = malloc(sizeof(node_t));
3  assert( n != NULL );
4  n->value = value ;
5  n->next = head;
6  head = n;
7 }
```

- Where is the race condition?

Prevention — Mutual exclusion

- Mutex based solution

Prevention — Mutual exclusion

- Mutex based solution

```
1 void insert(int value){
2  node_t * n = malloc(sizeof(node_t));
3  assert( n != NULL );
4  n->value = value ;
5  lock(listlock); // begin critical section
6  n->next = head;
7  head = n;
8  unlock(listlock); //end critical section
9 }
```

Prevention — Mutual exclusion

- Atomic instruction based?

HINT

```
1 int CompareAndSwap(int *address, int expected, int new){  
2   if(*address == expected){  
3     *address = new;  
4     return 1; // success  
5   }  
6   return 0;  
7 }
```

Prevention — Mutual exclusion

- Atomic instruction based?

HINT

```
1 int CompareAndSwap(int *address, int expected, int new){
2   if(*address == expected){
3     *address = new;
4     return 1; // success
5   }
6   return 0;
7 }
```

```
1 void insert(int value) {
2   node_t *n = malloc(sizeof(node_t));
3   assert(n != NULL);
4   n->value = value;
5   do {
6     n->next = head;
7   } while (CompareAndSwap(&head, n->next, n));
8 }
```

Deadlock Avoidance via Scheduling

Only one of T1 & T2 will
run at a given time ...

Deadlock Avoidance via Scheduling

- Global knowledge about which threads might be acquired is needed ahead of time

Only one of T1 & T2 will run at a given time ...

Deadlock Avoidance via Scheduling

- Global knowledge about which threads might be acquired is needed ahead of time
- Assume 2 processors and 4 threads and following lock requirements. How will you schedule to avoid deadlocks?

Only one of T1 & T2 will run at a given time ...

Deadlock Avoidance via Scheduling

- Global knowledge about which threads might be acquired is needed ahead of time
- Assume 2 processors and 4 threads and following lock requirements. How will you schedule to avoid deadlocks?
 - Hint: you can think serial!

Only one of T1 & T2 will run at a given time ...

Deadlock Avoidance via Scheduling

- Global knowledge about which threads might be acquired is needed ahead of time
- Assume 2 processors and 4 threads and following lock requirements. How will you schedule to avoid deadlocks?
 - Hint: you can think serial!

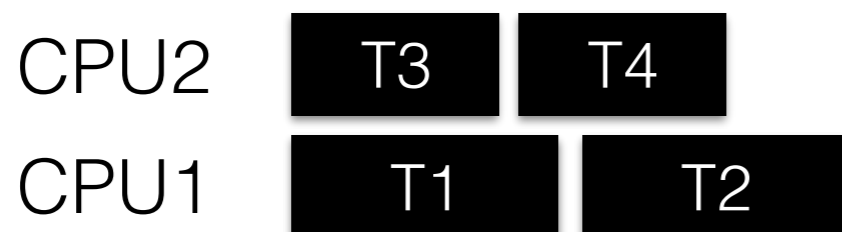
	T1	T2	T3	T4
L1	yes	yes	no	no
L2	yes	yes	yes	no

Only one of T1 & T2 will run at a given time ...

Deadlock Avoidance via Scheduling

- Global knowledge about which threads might be acquired is needed ahead of time
- Assume 2 processors and 4 threads and following lock requirements. How will you schedule to avoid deadlocks?
 - Hint: you can think serial!

	T1	T2	T3	T4
L1	yes	yes	no	no
L2	yes	yes	yes	no



Only one of T1 & T2 will run at a given time ...

Deadlock Avoidance via Scheduling

Deadlock Avoidance via Scheduling

	T1	T2	T3	T4
L1	yes	yes	yes	no
L2	yes	yes	yes	no

Deadlock Avoidance via Scheduling

	T1	T2	T3	T4
L1	yes	yes	yes	no
L2	yes	yes	yes	no

