

Lecture Note 5. Concurrency: Semaphore and Deadlock

April 2, 2021
Jongmoo Choi

Dept. of software
Dankook University

<http://embedded.dankook.ac.kr/~choijm>

(This slide is made by Jongmoo Choi. Please let him know when you want to distribute this slide)

Contents

- From Chap 30~32 of the OSTEP
- Chap 30. Condition Variables
- Chap 31. Semaphores
- Chap 32. Common Concurrency Problems
- Chap 33. Event-based Concurrency
- Chap 34. Summary

Chap. 30 Condition Variables

■ Locks

- ✓ Mainly focusing on **mutual exclusion**

■ Condition variables

- ✓ Focusing on **synchronization** (not only mutual exclusion but also ordering)
- ✓ Specifically, used for checking whether a condition is true
 - E.g.: 1) whether a child has completed. 2) whether a buffer is filled

```
1 void *child(void *arg) {
2     printf("child\n");
3     // XXX how to indicate we are done?
4     return NULL;
5 }
6
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    // XXX how to wait for child?
12    printf("parent: end\n");
13    return 0;
14 }
```

Figure 30.1: A Parent Waiting For Its Child

Chap. 30 Condition Variables

- Feasible solution 1: busy waiting with a variable

```
1  volatile int done = 0;
2
3  void *child(void *arg) {
4      printf("child\n");
5      done = 1;
6      return NULL;
7  }
8
9  int main(int argc, char *argv[]) {
10     printf("parent: begin\n");
11     pthread_t c;
12     Pthread_create(&c, NULL, child, NULL); // create child
13     while (done == 0)
14         ; // spin
15     printf("parent: end\n");
16     return 0;
17 }
```

Figure 30.2: Parent Waiting For Child: Spin-based Approach

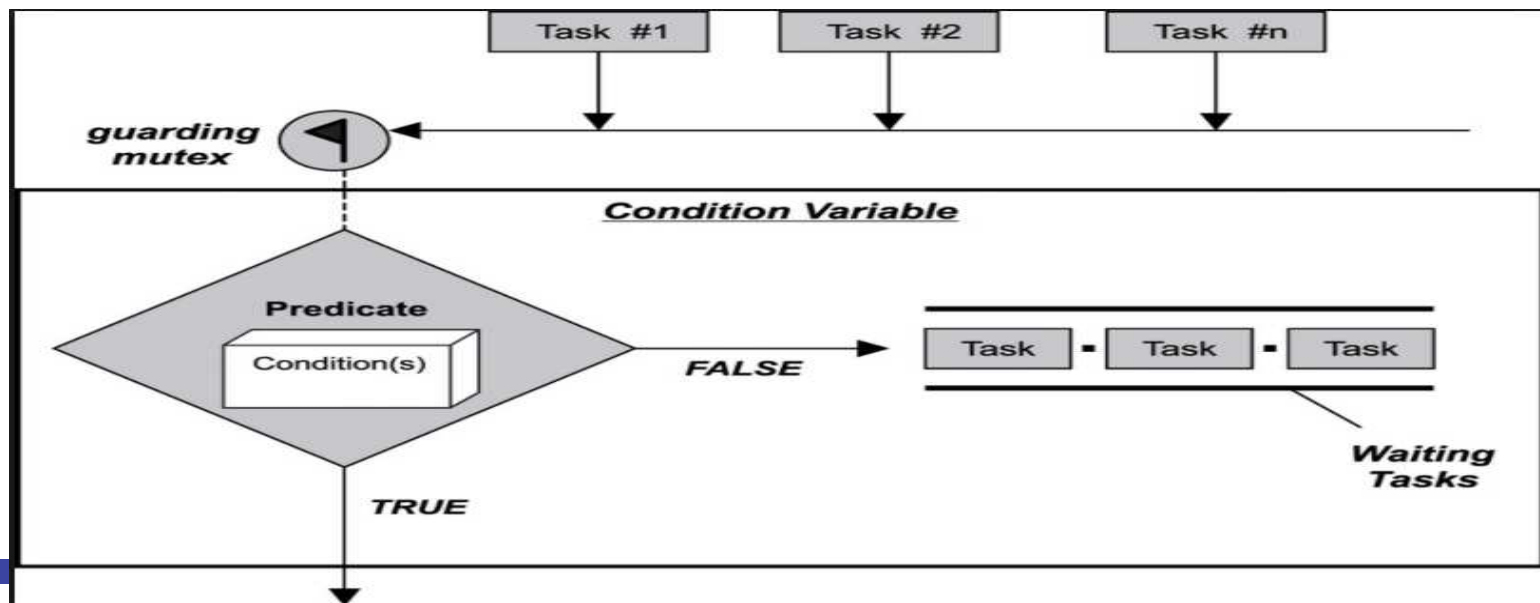
- ✓ Generally work, but inefficient (waste CPU time), sometimes incorrect on multiple children case

30.1 Definition and Routines

■ Feasible solution 2: condition variable

- ✓ An explicit queue that threads can put themselves on when some state of execution (i.e., some condition) is not as desired
- ✓ Some other thread, when it changes state, can then wake one (or more) of those waiting threads and thus allow them to continue.
- ✓ pthread APIs

```
pthread_cond_wait(pthread_cond_t *c, pthread_mutex_t *m);  
pthread_cond_signal(pthread_cond_t *c);
```



30.1 Definition and Routines

■ Feasible solution 2: condition variable

✓ Condition variable example

```
1  int done = 0;
2  pthread_mutex_t m = PTHREAD_MUTEX_INITIALIZER;
3  pthread_cond_t c = PTHREAD_COND_INITIALIZER;
4
5  void thr_exit() {
6      pthread_mutex_lock(&m);
7      done = 1;
8      pthread_cond_signal(&c);
9      pthread_mutex_unlock(&m);
10 }
11
12 void *child(void *arg) {
13     printf("child\n");
14     thr_exit();
15     return NULL;
16 }
17
18 void thr_join() {
19     pthread_mutex_lock(&m);
20     while (done == 0)
21         pthread_cond_wait(&c, &m);
22     pthread_mutex_unlock(&m);
23 }
24
25 int main(int argc, char *argv[]) {
26     printf("parent: begin\n");
27     pthread_t p;
28     pthread_create(&p, NULL, child, NULL);
29     thr_join();
30     printf("parent: end\n");
31     return 0;
32 }
```

Figure 30.3: Parent Waiting For Child: Use A Condition Variable

✓ Note: 1) wait(): unlock/lock **implicitly**, 2) while instead of if in join()

30.2 Producer/Consumer (Bounded Buffer) Problem

- The famous Producer/Consumer problem (also known as bounded buffer problem)
 - ✓ Scenario
 - Producers generate data items and place them in a buffer
 - Consumers grab the items from the buffer and consume them
 - e.g. DB server, streaming server, pipe, cache, ...
 - ✓ Issue
 - Mutual exclusion
 - Empty case: no data
 - Full case: no available buffer



30.2 Producer/Consumer (Bounded Buffer) Problem

- Basic structure: without considering sharing
 - ✓ Shared buffer: put(), get() interfaces
 - Assumption: space for only one item (single buffer) → relax later
 - ✓ Producer/Consumer: producer(), consumer()

```
1 int buffer;
2 int count = 0; // initially, empty
3
4 void put(int value) {
5     assert(count == 0);
6     count = 1;
7     buffer = value;
8 }
9
10 int get() {
11     assert(count == 1);
12     count = 0;
13     return buffer;
14 }
```




Figure 30.6: The Put And Get Routines (v1)

```
1 void *producer(void *arg) {
2     int i;
3     int loops = (int) arg;
4     for (i = 0; i < loops; i++) {
5         put(i);
6     }
7 }
8
9 void *consumer(void *arg) {
10     while (1) {
11         int tmp = get();
12         printf("%d\n", tmp);
13     }
14 }
```

Figure 30.7: Producer/Consumer Threads (v1)

30.2 Producer/Consumer (Bounded Buffer) Problem

■ Solution 1: Now consider sharing

- ✓ Mutual exclusion: mutex
- ✓ Ordering: condition variable

```
1 int loops; // must initialize somewhere...
2 cond_t cond;
3 mutex_t mutex;
4
5 void *producer(void *arg) {
6     int i;
7     for (i = 0; i < loops; i++) {
8         pthread_mutex_lock(&mutex); // p1
9         if (count == 1) // p2
10            pthread_cond_wait(&cond, &mutex); // p3
11        put(i); // p4
12        pthread_cond_signal(&cond); // p5
13        pthread_mutex_unlock(&mutex); // p6
14    }
15 }
16
17 void *consumer(void *arg) {
18     int i;
19     for (i = 0; i < loops; i++) {
20        pthread_mutex_lock(&mutex); // c1
21        if (count == 0) // c2
22            pthread_cond_wait(&cond, &mutex); // c3
23        int tmp = get(); // c4
24        pthread_cond_signal(&cond); // c5
25        pthread_mutex_unlock(&mutex); // c6
26        printf("%d\n", tmp);
27    }
28 }
```

count



buffer

Figure 30.8: Producer/Consumer: Single CV And If Statement

☛ Is it correct?

30.2 Producer/Consumer (Bounded Buffer) Problem

- Solution 1 (cont')
 - ✓ Wake up C1, but run C2

```

16 void *consumer(void *arg) {
17     int i;
18     for (i = 0; i < loops; i++) {
19         Pthread_mutex_lock(&mutex); // c1
20         if (count == 0) // c2
21             Pthread_cond_wait(&cond, &mutex); // c3
22         int tmp = get(); // c4
23         Pthread_cond_signal(&cond); // c5
24         Pthread_mutex_unlock(&mutex); // c6
25         printf("%d\n", tmp);
26     }
27 }
28
4 void *producer(void *arg) {
5     int i;
6     for (i = 0; i < loops; i++) {
7         Pthread_mutex_lock(&mutex); // p1
8         if (count == 1) // p2
9             Pthread_cond_wait(&cond, &mutex); // p3
10        put(i); // p4
11        Pthread_cond_signal(&cond); // p5
12        Pthread_mutex_unlock(&mutex); // p6
13    }
14 }

```

T_{c1}	State	T_{c2}	State	T_p	State	Count	Comment
c1	Running		Ready		Ready	0	
c2	Running		Ready		Ready	0	
c3	Sleep		Ready		Ready	0	Nothing to get
	Sleep		Ready	p1	Running	0	
	Sleep		Ready	p2	Running	0	
	Sleep		Ready	p4	Running	1	Buffer now full
	Ready		Ready	p5	Running	1	T_{c1} awoken
	Ready		Ready	p6	Running	1	
	Ready		Ready	p1	Running	1	
	Ready		Ready	p2	Running	1	
	Ready		Ready	p3	Sleep	1	Buffer full; sleep
	Ready	c1	Running		Sleep	1	T_{c2} sneaks in ...
	Ready	c2	Running		Sleep	1	
	Ready	c4	Running		Sleep	0	... and grabs data
	Ready	c5	Running		Ready	0	T_p awoken
	Ready	c6	Running		Ready	0	
c4	Running		Ready		Ready	0	Oh oh! No data

Figure 30.9: Thread Trace: Broken Solution (v1)

30.2 Producer/Consumer (Bounded Buffer) Problem

■ Solution 2

- ✓ while instead of if

```
1  int loops;
2  cond_t cond;
3  mutex_t mutex;
4
5  void *producer(void *arg) {
6      int i;
7      for (i = 0; i < loops; i++) {
8          Pthread_mutex_lock(&mutex);           // p1
9          while (count == 1)                   // p2
10             Pthread_cond_wait(&cond, &mutex); // p3
11          put(i);                               // p4
12          Pthread_cond_signal(&cond);          // p5
13          Pthread_mutex_unlock(&mutex);        // p6
14      }
15  }
16
17 void *consumer(void *arg) {
18     int i;
19     for (i = 0; i < loops; i++) {
20         Pthread_mutex_lock(&mutex);           // c1
21         while (count == 0)                   // c2
22             Pthread_cond_wait(&cond, &mutex); // c3
23         int tmp = get();                      // c4
24         Pthread_cond_signal(&cond);          // c5
25         Pthread_mutex_unlock(&mutex);        // c6
26         printf("%d\n", tmp);
27     }
28 }
```

Figure 30.10: Producer/Consumer: Single CV And While

➤ Now, is it correct?

30.2 Producer/Consumer (Bounded Buffer) Problem

■ Solution 2 (cont')

- ✓ Signal to P, but wake up C2

```

16 void *consumer(void *arg) {
17     int i;
18     for (i = 0; i < loops; i++) {
19         Pthread_mutex_lock(&mutex); // c1
20         while (count == 0) // c2
21             Pthread_cond_wait(&cond, &mutex); // c3
22         int tmp = get(); // c4
23         Pthread_cond_signal(&cond); // c5
24         Pthread_mutex_unlock(&mutex); // c6
25         printf("%d\n", tmp);
26     }
4 void *producer(void *arg) {
5     int i;
6     for (i = 0; i < loops; i++) {
7         Pthread_mutex_lock(&mutex); // p1
8         while (count == 1) // p2
9             Pthread_cond_wait(&cond, &mutex); // p3
10        put(i); // p4
11        Pthread_cond_signal(&cond); // p5
12        Pthread_mutex_unlock(&mutex); // p6
13    }

```

T_{c1}	State	T_{c2}	State	T_p	State	Count	Comment
c1	Running		Ready		Ready	0	
c2	Running		Ready		Ready	0	
c3	Sleep		Ready		Ready	0	Nothing to get
	Sleep	c1	Running		Ready	0	
	Sleep	c2	Running		Ready	0	
	Sleep	c3	Sleep		Ready	0	Nothing to get
	Sleep		Sleep	p1	Running	0	
	Sleep		Sleep	p2	Running	0	
	Sleep		Sleep	p4	Running	1	Buffer now full
	Ready		Sleep	p5	Running	1	T_{c1} awoken
	Ready		Sleep	p6	Running	1	
	Ready		Sleep	p1	Running	1	
	Ready		Sleep	p2	Running	1	
	Ready		Sleep	p3	Sleep	1	Must sleep (full)
c2	Running		Sleep		Sleep	1	Recheck condition
c4	Running		Sleep		Sleep	0	T_{c1} grabs data
c5	Running		Ready		Sleep	0	Oops! Woke T_{c2}
c6	Running		Ready		Sleep	0	
c1	Running		Ready		Sleep	0	
c2	Running		Ready		Sleep	0	
c3	Sleep		Ready		Sleep	0	Nothing to get
	Sleep	c2	Running		Sleep	0	
	Sleep	c3	Sleep		Sleep	0	Everyone asleep...

Figure 30.11: Thread Trace: Broken Solution (v2)

30.2 Producer/Consumer (Bounded Buffer) Problem

■ Solution 3 (final)

✓ Two condition variables

- Indicate explicitly which thread I want to send my signal.

```
1  cond_t  empty, fill;
2  mutex_t mutex;
3
4  void *producer(void *arg) {
5      int i;
6      for (i = 0; i < loops; i++) {
7          Pthread_mutex_lock(&mutex);
8          while (count == 1)
9              Pthread_cond_wait(&empty, &mutex);
10         put(i);
11         Pthread_cond_signal(&fill);
12         Pthread_mutex_unlock(&mutex);
13     }
14 }
15
16 void *consumer(void *arg) {
17     int i;
18     for (i = 0; i < loops; i++) {
19         Pthread_mutex_lock(&mutex);
20         while (count == 0)
21             Pthread_cond_wait(&fill, &mutex);
22         int tmp = get();
23         Pthread_cond_signal(&empty);
24         Pthread_mutex_unlock(&mutex);
25         printf("%d\n", tmp);
26     }
27 }
```

Figure 30.12: Producer/Consumer: Two CVs And While

30.2 Producer/Consumer (Bounded Buffer) Problem

Multiple buffers cases: final solution

```
1  int buffer[MAX];
2  int fill_ptr = 0;
3  int use_ptr = 0;
4  int count = 0;
5
6  void put(int value) {
7      buffer[fill_ptr] = value;
8      fill_ptr = (fill_ptr + 1) % MAX;
9      count++;
10 }
11
12 int get() {
13     int tmp = buffer[use_ptr];
14     use_ptr = (use_ptr + 1) % MAX;
15     count--;
16     return tmp;
17 }
```




Figure 30.13: The Correct Put And Get Routines

```
1  cond_t empty, fill;
2  mutex_t mutex;
3
4  void *producer(void *arg) {
5      int i;
6      for (i = 0; i < loops; i++) {
7          Pthread_mutex_lock(&mutex);           // p1
8          while (count == MAX)                 // p2
9              Pthread_cond_wait(&empty, &mutex); // p3
10         put(i);                               // p4
11         Pthread_cond_signal(&fill);          // p5
12         Pthread_mutex_unlock(&mutex);        // p6
13     }
14 }
15
16 void *consumer(void *arg) {
17     int i;
18     for (i = 0; i < loops; i++) {
19         Pthread_mutex_lock(&mutex);           // c1
20         while (count == 0)                   // c2
21             Pthread_cond_wait(&fill, &mutex); // c3
22         int tmp = get();                     // c4
23         Pthread_cond_signal(&empty);         // c5
24         Pthread_mutex_unlock(&mutex);        // c6
25         printf("%d\n", tmp);
26     }
27 }
```

Figure 30.14: The Correct Producer/Consumer Synchronization



Quiz for 6th-Week 2nd-Lesson

■ Quiz

- ✓ 1. Explain the three issues that we need to consider for the producer/consumer problem.
- ✓ 2. Discuss whether the below program is correct or not? If incorrect, describe why?
- ✓ Due: until 6 PM Friday of this week (9th, April)

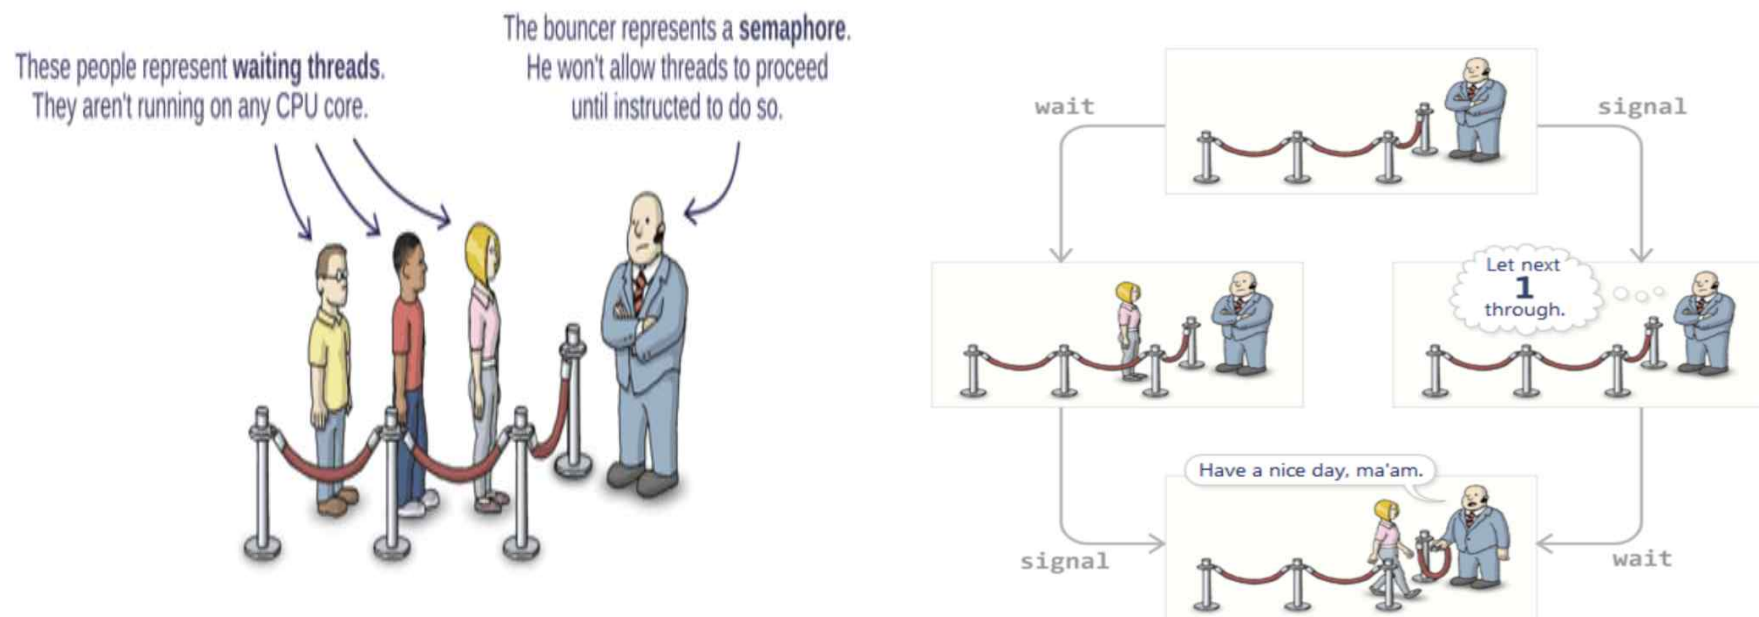
```
1 int loops; // must initialize somewhere...
2 cond_t cond;
3 mutex_t mutex;
4
5 void *producer(void *arg) {
6     int i;
7     for (i = 0; i < loops; i++) {
8         Pthread_mutex_lock(&mutex);           // p1
9         if (count == 1)                       // p2
10            Pthread_cond_wait(&cond, &mutex); // p3
11         put(i);                               // p4
12         Pthread_cond_signal(&cond);          // p5
13         Pthread_mutex_unlock(&mutex);        // p6
14     }
15 }
16
17 void *consumer(void *arg) {
18     int i;
19     for (i = 0; i < loops; i++) {
20         Pthread_mutex_lock(&mutex);           // c1
21         if (count == 0)                       // c2
22            Pthread_cond_wait(&cond, &mutex); // c3
23         int tmp = get();                       // c4
24         Pthread_cond_signal(&cond);          // c5
25         Pthread_mutex_unlock(&mutex);        // c6
26         printf("%d\n", tmp);
27     }
28 }
```

Figure 30.8: Producer/Consumer: Single CV And If Statement

Chap 31. Semaphores

■ Semaphore

- ✓ Well-known structure for concurrency control
 - Can be used as both a lock and a condition variable
 - Binary semaphore, Counting semaphore
 - Can be employed by various concurrency problems including 1) producer/consumer, 2) reader/writer and 3) dining philosophers
- ✓ Invented by the famous Edsger Dijkstra



(Source: <http://preshing.com/20150316/semaphores-are-surprisingly-versatile/>)

31.1 Semaphores: A Definition

■ Semaphore definition

- ✓ An object with an integer value manipulated by three routines
 - `sem_init(semaphore, p_shared, initial_value)`
 - `sem_wait()`: also called as `P()`, `down()` ...
 - Decrease the value of the semaphore (S). Then, either return right away (when $S \geq 0$) or cause the caller to suspend execution waiting for a subsequent post (when $S < 0$)
 - `sem_post()`: also called as `V()`, `up()`, `sem_signal()` ...
 - Increment the value of the semaphore and then, if there is a thread waiting to be woken, wakes **one** of them up
 - Others: `sem_trywait()`, `sem_timewait()`, `sem_destroy()`

```
1  #include <semaphore.h>
2  sem_t s;
3  sem_init(&s, 0, 1);
```

Figure 31.1: Initializing A Semaphore

```
1  int sem_wait(sem_t *s) {
2      decrement the value of semaphore s by one
3      wait if value of semaphore s is negative
4  }
5
6  int sem_post(sem_t *s) {
7      increment the value of semaphore s by one
8      if there are one or more threads waiting, wake one
9  }
```

Figure 31.2: Semaphore: Definitions Of Wait And Post

31.2 Binary Semaphores (Locks)

■ Using a semaphore as a lock

```

1  sem_t m;
2  sem_init(&m, 0, X); // initialize semaphore to X; what should X be?
3
4  sem_wait(&m);
5  // critical section here
6  sem_post(&m);

```

Figure 31.3: A Binary Semaphore (That Is, A Lock)

✓ Running example

- Can support the mutual exclusion
- Note that the value of the semaphore, **when negative, is equal to the number of waiting threads**

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	Interrupt; Switch→T1	Ready		Running
0		Ready	call sem_wait()	Running
-1		Ready	decrement sem	Running
-1		Ready	(sem<0)→sleep	Sleeping
-1		Running	Switch→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	Interrupt; Switch→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

Figure 31.5: Thread Trace: Two Threads Using A Semaphore

31.3 Semaphores for Ordering

- Using a semaphore as a conditional variable
 - ✓ Initial semaphore value: 0 (note: it is initialized as 1 for mutex)

```
1  sem_t s;
2
3  void *
4  child(void *arg) {
5      printf("child\n");
6      sem_post(&s); // signal here: child is done
7      return NULL;
8  }
9
10 int
11 main(int argc, char *argv[]) {
12     sem_init(&s, 0, X); // what should X be?
13     printf("parent: begin\n");
14     pthread_t c;
15     Pthread_create(&c, NULL, child, NULL);
16     sem_wait(&s); // wait here for child
17     printf("parent: end\n");
18     return 0;
19 }
```

Figure 31.6: A Parent Waiting For Its Child

- Compare semaphore (this page) with condition variable (page 6) ➔ No "Done" variable

31.4 Producer/Consumer (Bounded Buffer) Problem

■ Using a semaphore for the producer/consumer problem

- ✓ mutex: **binary semaphore**, full/empty: **counting semaphore**

```
1  int buffer[MAX];
2  int fill = 0;
3  int use = 0;
4
5  void put(int value) {
6      buffer[fill] = value;    // Line F1
7      fill = (fill + 1) % MAX; // Line F2
8  }
9
10 int get() {
11     int tmp = buffer[use];    // Line G1
12     use = (use + 1) % MAX;    // Line G2
13     return tmp;
14 }
```

Figure 31.9: 7

```
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);
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 }
```


Figure 31.12: Adding Mutual Exclusion (Correctly)

- Summary of two versions (semaphore in page 20 vs condition variable in page 14)
 - 1) No count variable (owing to counting semaphore)
 - 2) ordering → mutex vs mutex → ordering (See page 40)

31.5 Reader-Writer Locks

■ Producer/Consumer vs. Reader/Writer

- ✓ Producer/Consumer: need mutual exclusion (e.g. list insert/delete)
- ✓ Reader/Writer: need mutual exclusion, but allow multiple readers (e.g. tree lookup and insert)
 - Specific comparison
 - A Producer or Consumer in Critical Section → next Producer or Consumer must wait
 - A writer in Critical Section → 1) next writer or 2) next reader must wait
 - A reader in Critical Section → 1) next writer must wait, 2) but next reader can enter (**better performance**)
 - Issue (related to starvation)
 - Readers in Critical Section + a writer is waiting → a reader arrives : wait or allowed (depending on either writer preference or reader preference)



31.5 Reader-Writer Locks

■ Implementation for reader/writer

- ✓ lock: for mutual exclusion on readers
- ✓ writelock: to allow a write or multiple readers
 - The below implementation prefer readers (writers can starve)

```
1  typedef struct _rwlock_t {
2      sem_t lock;           // binary semaphore (basic lock)
3      sem_t writelock;     // used to allow ONE writer or MANY readers
4      int  readers;        // count of readers reading in critical section
5  } rwlock_t;
6
7  void rwlock_init(rwlock_t *rw) {
8      rw->readers = 0;
9      sem_init(&rw->lock, 0, 1);
10     sem_init(&rw->writelock, 0, 1);
11 }
12
13 void rwlock_acquire_readlock(rwlock_t *rw) {
14     sem_wait(&rw->lock);
15     rw->readers++;
16     if (rw->readers == 1)
17         sem_wait(&rw->writelock); // first reader acquires writelock
18     sem_post(&rw->lock);
19 }
20
21 void rwlock_release_readlock(rwlock_t *rw) {
22     sem_wait(&rw->lock);
23     rw->readers--;
24     if (rw->readers == 0)
25         sem_post(&rw->writelock); // last reader releases writelock
26     sem_post(&rw->lock);
27 }
28
29 void rwlock_acquire_writelock(rwlock_t *rw) {
30     sem_wait(&rw->writelock);
31 }
32
33 void rwlock_release_writelock(rwlock_t *rw) {
34     sem_post(&rw->writelock);
35 }
```

Figure 31.13: A Simple Reader-Writer Lock

31.6 The Dining Philosophers

■ Problem definition

- ✓ There are five “philosophers” sitting around a table.
- ✓ Between each pair of philosophers is a single fork (thus, five total)
- ✓ The philosophers each have times for thinking or for eating
- ✓ In order to eat, a philosopher needs two forks, both the one on their left and the one on their right → shared resource → concurrency

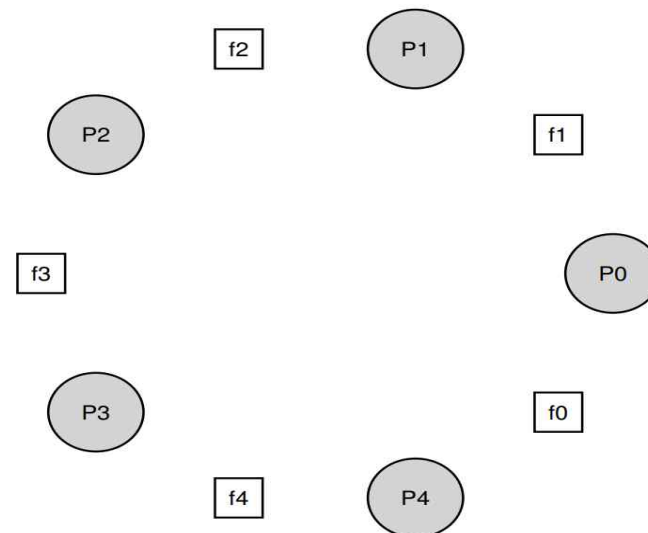



Figure 31.14: The Dining Philosophers

31.6 The Dining Philosophers

■ Solution

- ✓ Basic loop for each philosopher
- ✓ Now question is how to implement `getforks()` and `putforks()`
 - Using five semaphores: `sem_t forks[5]`
 - Obtain semaphore before acquire a fork
- ✓ Cause **Deadlock**
 - All philosophers obtain their left fork, while waiting their right one
 - How to avoid this issue?

■ New Solutions

- ✓ 1) break ordering, 2) set limit, 3) employ transaction (e.g. the Monitor), 4) more resource, 5) teach philosophers (idea from a student)




Figure 31.14: The Dining Philosophers

```
while (1) {  
    think ();  
    getforks ();  
    eat ();  
    putforks ();  
}
```

(Basic loop)




Figure 31.14: The Dining Philosophers

```
1 void get_forks(int p) {  
2     sem_wait (&forks[left (p)]);  
3     sem_wait (&forks[right (p)]);  
4 }  
5  
6 void put_forks(int p) {  
7     sem_post (&forks[left (p)]);  
8     sem_post (&forks[right (p)]);  
9 }
```

Figure 31.15: The `get_forks ()` And `put_forks ()` Routines

```
1 void get_forks(int p) {  
2     if (p == 4) {  
3         sem_wait (&forks[right (p)]);  
4         sem_wait (&forks[left (p)]);  
5     } else {  
6         sem_wait (&forks[left (p)]);  
7         sem_wait (&forks[right (p)]);  
8     }  
9 }
```

Figure 31.16: Breaking The Dependency In `get_forks ()`



Quiz for 7th-Week 1st-Lesson

■ Quiz

- ✓ 1. Explain the meaning of semaphore value in Figure 31.5. In other words, what does it imply when the value is 1 (or 0 or -1 or even -2)?
- ✓ 2. Discuss the differences between the producer/consumer and reader/writer problem (at least 2 differences).
- ✓ Due: until 6 PM Friday of this week (16th, April)

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	Interrupt; Switch→T1	Ready		Running
0		Ready	call sem_wait()	Running
-1		Ready	decrement sem	Running
-1		Ready	(sem<0)→sleep	Sleeping
-1		Running	Switch→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	Interrupt; Switch→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

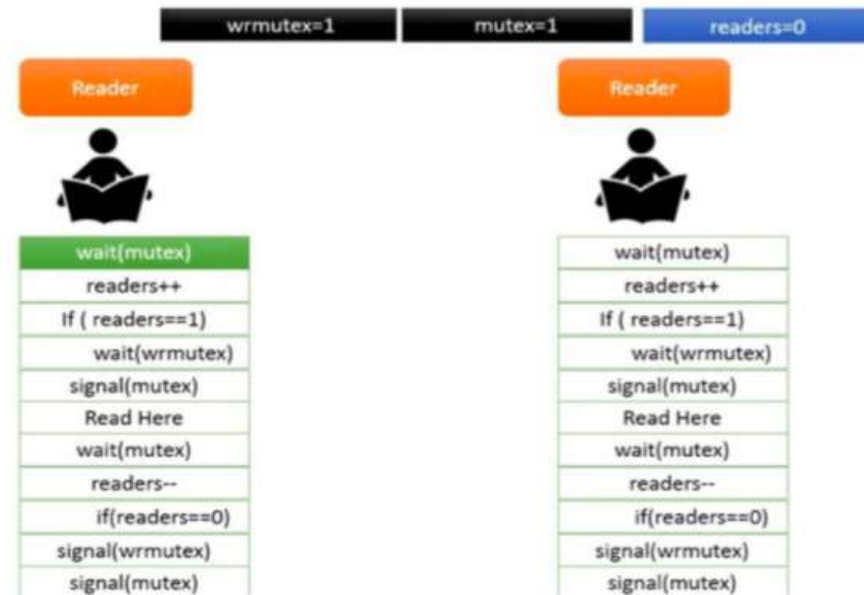


Figure 31.5: Thread Trace: Two Threads Using A Semaphore

(Source: www.chegg.com/)

Chap 32. Common Concurrency Problems

- Concurrency
 - ✓ Pros: can enhance throughput via processing in parallel
 - ✓ Cons: may cause several **troublesome** concurrency bugs (a.k.a. **timing bugs**)
- 32.1 What Types of Concurrency Bugs Exist?

Application	What it does	Non-Deadlock	Deadlock
MySQL	Database Server	14	9
Apache	Web Server	13	4
Mozilla	Web Browser	41	16
OpenOffice	Office Suite	6	2
Total		74	31

Figure 32.1: **Bugs In Modern Applications**

- ✓ Total bugs: 105
 - Deadlock bugs: 31
 - Non-deadlock bugs : 74
- ✓ Differ among applications

32.2 Non-Deadlock Bugs

- Two major types of non-deadlock bugs
 - ✓ Atomicity-Violation Bugs (From MySQL sources)

```
1 Thread 1::
2 if (thd->proc_info) {
3     ...
4     fputs(thd->proc_info, ...);
5     ...
6 }
7
8 Thread 2::
9 thd->proc_info = NULL;
```

- ✓ Order-Violation Bugs

```
1 Thread 1::
2 void init() {
3     ...
4     mThread = PR_CreateThread(mMain, ...);
5     ...
6 }
7
8 Thread 2::
9 void mMain(...) {
10     ...
11     mState = mThread->State;
12     ...
13 }
```

32.2 Non-Deadlock Bugs

■ Solution to Atomicity-Violation Bugs

```
1  pthread_mutex_t proc_info_lock = PTHREAD_MUTEX_INITIALIZER;
2
3  Thread 1::
4  pthread_mutex_lock(&proc_info_lock);
5  if (thd->proc_info) {
6      ...
7      fputs(thd->proc_info, ...);
8      ...
9  }
10 pthread_mutex_unlock(&proc_info_lock);
11
12 Thread 2::
13 pthread_mutex_lock(&proc_info_lock);
14 thd->proc_info = NULL;
15 pthread_mutex_unlock(&proc_info_lock);
```

32.2 Non-Deadlock Bugs

■ Solution to Order-Violation Bugs

```
1  pthread_mutex_t mtLock = PTHREAD_MUTEX_INITIALIZER;
2  pthread_cond_t  mtCond = PTHREAD_COND_INITIALIZER;
3  int mtInit
4      = 0;
5
6  Thread 1::
7  void init() {
8      ...
9      mThread = PR_CreateThread(mMain, ...);
10
11     // signal that the thread has been created...
12     pthread_mutex_lock(&mtLock);
13     mtInit = 1;
14     pthread_cond_signal(&mtCond);
15     pthread_mutex_unlock(&mtLock);
16     ...
17 }
18
19 Thread 2::
20 void mMain(...) {
21     ...
22     // wait for the thread to be initialized...
23     pthread_mutex_lock(&mtLock);
24     while (mtInit == 0)
25         pthread_cond_wait(&mtCond, &mtLock);
26     pthread_mutex_unlock(&mtLock);
27
28     mState = mThread->State;
29     ...
30 }
```

• `pthread_cond_wait()`: unlock and lock mutex implicitly before and after sleep (see page 6)

32.3 Deadlock Bugs

■ Deadlock

- ✓ A situation where two or more threads wait for events that never occur

```
Thread 1:  
pthread_mutex_lock(L1);  
pthread_mutex_lock(L2);
```

```
Thread 2:  
pthread_mutex_lock(L2);  
pthread_mutex_lock(L1);
```

- E.g.) When a thread (say Thread 1) is holding a lock (L1) and waiting for another one (L2); **unfortunately**, the thread (Thread 2) that holds lock L2 is waiting for L1 to be released.




Figure 32.2: The Deadlock Dependency Graph

32.3 Deadlock Bugs

- **Deadlock: 4 Conditions**
 - ✓ Mutual exclusion
 - ✓ Hold-and-Wait
 - ✓ No preemption for resource
 - ✓ Circular wait

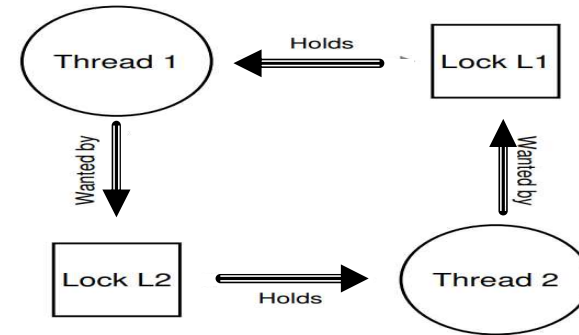




Figure 32.2: The Deadlock Dependency Graph



(a) Deadlock Possible



(b) Deadlock

32.3 Deadlock Bugs

- How to handle Deadlock: three strategies
 - ✓ 1. Deadlock Prevention
 - ✓ 2. Deadlock Avoidance via Scheduling
 - ✓ 3. Deadlock Detection and Recovery



Approach	Resource Allocation Policy	Different Schemes	Major Advantages	Major Disadvantages
Prevention	Conservative; undercommits resources	Requesting all resources at once	<ul style="list-style-type: none"> • Works well for processes that perform a single burst of activity • No preemption necessary 	<ul style="list-style-type: none"> • Inefficient • Delays process initiation • Future resource requirements must be known by processes
		Preemption	<ul style="list-style-type: none"> • Convenient when applied to resources whose state can be saved and restored easily 	<ul style="list-style-type: none"> • Preempts more often than necessary
		Resource ordering	<ul style="list-style-type: none"> • Feasible to enforce via compile-time checks • Needs no run-time computation since problem is solved in system design 	<ul style="list-style-type: none"> • Disallows incremental resource requests
Avoidance	Midway between that of detection and prevention	Manipulate to find at least one safe path	<ul style="list-style-type: none"> • No preemption necessary 	<ul style="list-style-type: none"> • Future resource requirements must be known by OS • Processes can be blocked for long periods
Detection	Very liberal; requested resources are granted where possible	Invoke periodically to test for deadlock	<ul style="list-style-type: none"> • Never delays process initiation • Facilitates online handling 	<ul style="list-style-type: none"> • Inherent preemption losses

(Source: "Operating systems: Internals and Design Principle" by W. Stalling)

32.3 Deadlock Bugs

■ Deadlock prevention

✓ This strategy seeks to prevent one of the 4 Deadlock conditions

✓ 1. Hold-and-wait

- Acquire all locks at once, atomically

✓ 2. No Preemption

- Release lock if it can not hold another lock
- Concern: 1) may cause [Livelock](#), 2) sometimes require undo
 - Two threads could both be repeatedly attempting this sequence and repeatedly failing to acquire both locks → add random delay

✓ 3. Circular Wait


- A total ordering on lock acquisition
- E.g.) The comment at the top of the source code in Linux: “*i_mutex*” before *i_mmap_mutex*”

```
1 pthread_mutex_lock(prevention); // begin lock acquisition
2 pthread_mutex_lock(L1);
3 pthread_mutex_lock(L2);
4 ...
5 pthread_mutex_unlock(prevention); // end
```

(Acquire all locks atomically)

```
1 top:
2 pthread_mutex_lock(L1);
3 if (pthread_mutex_trylock(L2) != 0) {
4     pthread_mutex_unlock(L1);
5     goto top;
6 }
```

(Release lock if it can not hold another lock)



32.3 Deadlock Bugs

■ Deadlock prevention (cont')

✓ 4. Mutual Exclusion:

- “lock free” approach: no lock but support mutual exclusion
 - Using powerful hardware instructions, we can build data structures in a manner that does not require explicit locking
- Atomic integer operation with compare-and-swap (chapter 28.9 in LN 4)

```
void increment(counter_t *c) {  
    Pthread_mutex_lock(&c->lock);  
    c->value++;  
    Pthread_mutex_unlock(&c->lock);  
}
```

Using Lock

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

Lock free

- List management (39 page in LN4)

```
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 }
```

Using Lock

Lock free

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

```
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) == 0);  
8 }
```

☛ Lock free: applicable only some specific cases vs Lock: general


32.3 Deadlock Bugs

■ Deadlock Avoidance via Scheduling

- ✓ Instead of prevention, try to avoid by scheduling threads in a way as to guarantee no deadlock can occur.
 - E.g.) two CPUs, four threads, T1 wants to use L1 and L2, T2 also wants both, T3 wants L1 only, T4 wants nothing



- E.g. 2) more contention (negative for load balancing)



- No deadlock, but under-utilization → [A conservative approach](#)

32.3 Deadlock Bugs

■ Deadlock Avoidance via Scheduling (cont')

✓ Famous algorithm: Banker's algorithm

- E.g.) Multiple processes with single resource case (also applicable to multiple resources case)

	Has	Max
A	0	5
B	0	6
C	0	3
D	0	7

Initial State: Free = 10


	Has	Max
A	2	5
B	0	6
C	1	3
D	5	7

State 1: Free = 2

	Has	Max
A	2	5
B	1	6
C	1	3
D	5	7

State 2: Free = 1

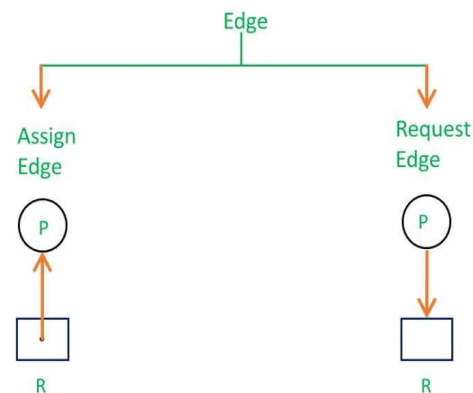
- Safe and unsafe state
 - Try to stay in safe state while allocating resources




32.3 Deadlock Bugs

■ Deadlock Detection and Recovery


- ✓ Allow deadlocks to occasionally occur, and then take a detection and recovery action
 - E.g.) If an OS froze once a year, you would just reboot it (but failure is a norm in a Cloud/Bigdata platform)
 - Many DB systems employ active deadlock detection approach
- ✓ How to detect?
 - Periodically, build **resource allocation graph**, checking in for cycles
- ✓ How to recovery?
 - Select a victim (youngest or least locks)



Meaning of Node and Edge in Resource allocation graph



Resource allocation graph Example without Deadlock



Resource allocation graph Example with Deadlock

(Source: <https://www.slideshare.net/AbhinawRai/deadlock-51330115>)

32.4 Summary

- Concurrency method
 - ✓ Lock, Condition variable, Semaphore, ...
- Well-known concurrency problems
 - ✓ The Producer/Consumer problem
 - ✓ The Reader/Writer problem
 - ✓ The Dining philosopher problem
- Concurrency bugs
 - ✓ Non-Deadlock bugs
 - ✓ Deadlock bugs
- Deadlock approach
 - ✓ Prevention
 - ✓ Avoidance
 - ✓ Detection and Recovery

TIP: DON'T ALWAYS DO IT PERFECTLY (TOM WEST'S LAW)


Tom West, famous as the subject of the classic computer-industry book *Soul of a New Machine* [K81], says famously: "Not everything worth doing is worth doing well", which is a terrific engineering maxim. If a bad thing happens rarely, certainly one should not spend a great deal of effort to prevent it, particularly if the cost of the bad thing occurring is small. If, on the other hand, you are building a space shuttle, and the cost of something going wrong is the space shuttle blowing up, well, perhaps you should ignore this piece of advice.



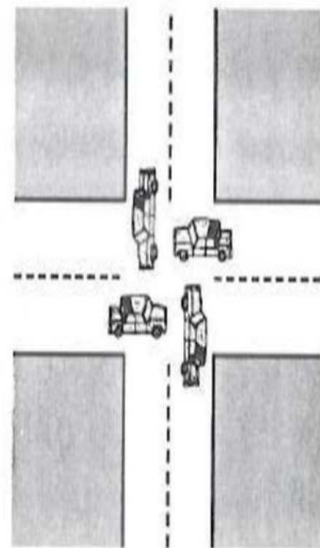
Quiz for 7th-Week 2nd-Lesson

■ Quiz


- ✓ 1. Explain 4 conditions for deadlock using the below left figure.
- ✓ 2. Is there a deadlock in the below right figure of this page?
- ✓ Due: until 6 PM Friday of this week (16th, April)



(a) Deadlock Possible



(b) Deadlock



(Source: www.chegg.com/)

Appendix 1

■ 31.4 Producer/Consumer (Bounded Buffer) Problem

✓ Second attempt: Adding 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 }
```

✚ Is it correct?

Figure 31.11: Adding Mutual Exclusion (Incorrectly)

Appendix 1

■ 31.7 How to Implement Semaphores

✓ Using mutex and condition variable

```
1  typedef struct __Zem_t {
2      int value;
3      pthread_cond_t cond;
4      pthread_mutex_t lock;
5  } Zem_t;
6
7  // only one thread can call this
8  void Zem_init(Zem_t *s, int value) {
9      s->value = value;
10     Cond_init(&s->cond);
11     Mutex_init(&s->lock);
12 }
13
14 void Zem_wait(Zem_t *s) {
15     Mutex_lock(&s->lock);
16     while (s->value <= 0)
17         Cond_wait(&s->cond, &s->lock);
18     s->value--;
19     Mutex_unlock(&s->lock);
20 }
21
22 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 }
```

Figure 31.16: Implementing Zemaphores With Locks And CVs

Appendix 2

■ 30.3 pthread_cond_broadcast: Covering Conditions

- ✓ Memory allocation library for multi-thread env.
- ✓ Issue: which one to wake up?
 - E.g.) no free space, T1 asks 100B, T2 asks 10B, Both sleep → T3 free 50B → T2 wakeup: okay, T1 wakeup: sleep again, but T2 also sleeps
- ✓ pthread_cond_broadcast() instead of pthread_cond_signal()

```
1 // how many bytes of the heap are free?
2 int bytesLeft = MAX_HEAP_SIZE;
3
4 // need lock and condition too
5 cond_t c;
6 mutex_t m;
7
8 void *
9 allocate(int size) {
10     pthread_mutex_lock(&m);
11     while (bytesLeft < size)
12         pthread_cond_wait(&c, &m);
13     void *ptr = ...; // get mem from heap
14     bytesLeft -= size;
15     pthread_mutex_unlock(&m);
16     return ptr;
17 }
18
19 void free(void *ptr, int size) {
20     pthread_mutex_lock(&m);
21     bytesLeft += size;
22     pthread_cond_signal(&c); // whom to signal??
23     pthread_mutex_unlock(&m);
24 }
```

Figure 30.13: Covering Conditions: An Example

- ✦ Please read carefully the program in Figure 30.13, Figure 30.14 and Figure 31.12. It will be great helpful when you do the Lab. 2 ^^