

Lecture 15: Semaphores

Mythili Vutukuru

IIT Bombay

What is a semaphore?

- Synchronization primitive like condition variables
- Semaphore is a variable with an underlying counter
- Two functions on a semaphore variable
 - Up/post increments the counter
 - Down/wait decrements the counter and blocks the calling thread if the resulting value is negative
- A semaphore with init value 1 acts as a simple lock (binary semaphore = mutex)

```
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);
```

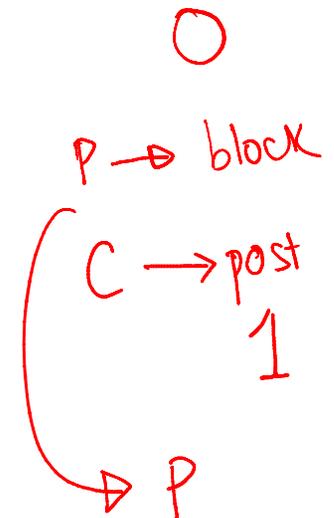
Handwritten annotations: A red circle around 'X' in line 2, a red arrow pointing from the circle to 'sem_wait' in line 4, and a red arrow pointing from 'sem_wait' to 'critical section here' in line 5. To the right of the code, the handwritten expression '1 - 0 - 1' is written in red.

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

Semaphores for ordering

- Can be used to set order of execution between threads like CV
- Example: parent waiting for child (init = 0)

```
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 }
```



Example: Producer/Consumer (1)

- Need two semaphores for signaling
 - One to track empty slots, and make producer wait if no more empty slots
 - One to track full slots, and make consumer wait if no more full slots
- One semaphore to act as mutex for buffer

```
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 }
```

Example: Producer/Consumer (2)

```
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);
10         put(i);
11         sem_post(&mutex);
12         sem_post(&full);
13     }
14 }
15
16 void *consumer(void *arg) {
17     int i;
18     for (i = 0; i < loops; i++) {
19         sem_wait(&full);
20         sem_wait(&mutex);
21         int tmp = get();
22         sem_post(&mutex);
23         sem_post(&empty);
24         printf("%d\n", tmp);
25     }
26 }
```

Incorrect solution with deadlock

- What if lock is acquired before signaling?
- Waiting thread sleeps with mutex and the signaling thread can never wake it up

```
5 void *producer(void *arg) {
6     int i;
7     for (i = 0; i < loops; i++) {
8         sem_wait(&mutex);
9         sem_wait(&empty);
10        put(i);
11        sem_post(&full);
12        sem_post(&mutex);
13    }
14 }
15
16 void *consumer(void *arg) {
17     int i;
18     for (i = 0; i < loops; i++) {
19         sem_wait(&mutex);
20         sem_wait(&full);
21         int tmp = get();
22         sem_post(&empty);
23         sem_post(&mutex);
24         printf("%d\n", tmp);
25     }
26 }
```