

Operating Systems: Processes and Threads

keleher

February 20, 2024

1. Synchronization
2. Bounded counter

- Lock operations: **acquire** and **release**
- `lck ← Lock()` // define a lock
- `lck.acq()` // acquire the lock; **blocking**
 - call only if caller does not hold `lck`
 - returns only when no other thread holds `lck`
- `lck.rel()` // release the lock; **non-blocking**
 - call only if caller holds `lck`
- `lck.rel()` does not give priority to threads blocked in `lck.acq()`

- Condition variable operations: `wait`, `signal` and `signal_all`
- A condition variable is associated with a lock
- `cv ← Condition(lck)` // condition variable associated with `lck`
- `cv.wait()` // wait on `cv`; **blocking**
 - call only if caller already holds `lck`
 - atomically release `lck` and wait on `cv`
when awakened: acquire `lck` and return
- `cv.signal()` // signal `cv`; **non-blocking**
 - call only if caller holds `lck`
 - wake up a thread (if any) waiting on `cv`
- `cv.signal_all()` // wake up all threads waiting on `cv`
- `lck.acq()` does not give priority to threads blocked in `cv.wait()`

```
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 }
```

The diagram consists of two red arrows pointing to lines 21 and 22 of the code. The first arrow points to line 21, `Pthread_cond_wait(&c, &m);`, and is accompanied by the text "atomically release lock and wait()". The second arrow points to line 22, `Pthread_mutex_unlock(&m);`, and is accompanied by the text "re-acquires lock on wake".

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

- Two cases:

1. parent creates child, continues running:

- 1.1 parent acquires lock
- 1.2 checks for child done (no)
- 1.3 go to sleep via `wait()`
- 1.4 child eventually runs, exits
- 1.5 parent wakes

2. child runs immediately

- 2.1 child signals, exits
- 2.2 parent wakes

- What if we didn't have the `done` state variable?

```
1 void thr_exit() {
2     Pthread_mutex_lock (&m);
3     Pthread_cond_signal (&c);
4     Pthread_mutex_unlock (&m);
5 }
6
7 void thr_join() {
8     Pthread_mutex_lock (&m);
9     Pthread_cond_wait (&c, &m);
10    Pthread_mutex_unlock (&m);
11 }
```

Figure 30.4: **Parent Waiting: No State Variable**

- Fine if parent runs first....

- What if we didn't have the associated `lock`?



```
1 void thr_exit() {
2     done = 1;
3     Pthread_cond_signal(&c);
4 }
5
6 void thr_join() {
7     if (done == 0)
8         Pthread_cond_wait(&c);
9 }
```

Figure 30.5: **Parent Waiting: No Lock**

- Not good if parent check, then child runs....
- Referred to as a “TOCTOU” (Time_of_Check to Time_of_Use) vulnerability

- Semaphore: variable with a non-negative integer `count`
- Semaphore operations: `P()` and `V()`
- `sem` \leftarrow Semaphore(`N`) // define semaphore with `count` `N` (≥ 0)
- `sem.P()` // blocking
 - wait until `sem.count` > 0 then decrease `sem.count` by 1; return
 - checking `sem.count` > 0 and decrementing are one atomic step
- `sem.V()` // non-blocking
 - atomically increase `sem.count` by 1; return
- `V()` does not give priority to threads blocked in `P()`

- *await B: S*, where *S* is a code chunk (*no blocking or infinite loop*) and *B* is a boolean condition (*no side effects*):
 - execute *S* only if *B* holds, all in one atomic step
 - if *B* does not hold, wait
- *atomic S*: short for *await True: S*
- Example: Given a linked list *x* with non-blocking functions *add()* and *rmv()*. To allow multiple threads to call these functions simultaneously, simply wrap them as follows:
 - *await True : add()*
 - *await (xnotempty) : rmv()*

- For a multi-threaded program to achieve anything, we have to assume that its threads execute with non-zero speed (but otherwise arbitrarily varying)
- Making this precise is simple for non-blocking statements but not for blocking statements (eg, acquire, wait, P, await)
- A thread at a **non-blocking** statement T eventually gets past T
 - Achieved if every unblocked thread periodically gets cpu cycles
- A thread at a **blocking** statement T eventually gets past T if T is *continuously* unblocked or *repeatedly* (but not continuously) unblocked
 - Achieved in most implementations only in a probabilistic sense, not in a deterministic sense

1. Synchronization
2. Bounded counter

Program P0:

- x, y: global int variables; initially 0
- up(), down() // callable by multiple threads simultaneously
- up() increments x only if $x < 100$, and returns $2*x$
- down() decrements x only if $x > 0$, and returns $2*x$

■ up():

```
int z
```

```
await (x < 100):
```

```
    x ← x+1
```

```
    z ← x
```

```
return 2*z
```

■ down():

```
int z
```

```
await (x > 0):
```

```
    x ← x-1
```

```
    z ← x
```

```
return 2*z
```

Program P1:

- `x, y` // as in P0
- `lck ← Lock()`
- `cvNF ← Condition(lck)` // for guard (`x < 100`)
- `cvNE ← Condition(lck)` // for guard (`x > 0`)

- `up()`:
 - `int z`
 - `lck.acq()`
 - `while (not x < 100):`
 - `cvNF.wait()`
 - `x ← x+1`
 - `z ← x`
 - `cvNE.signal()`
 - `lck.rel()`
 - `return 2*z`
- `down()`:
 - `int z`
 - `lck.acq()`
 - `while (not x > 0):`
 - `cvNE.wait()`
 - `x ← x-1`
 - `z ← x`
 - `cvNF.signal()`
 - `lck.rel()`
 - `return 2*z`

Program P2:

- `x, y` // as in P0
- `lck ← Lock()`
- `cv ← Condition(lck)` // for both guards

- `up():`
 - `int z`
 - `lck.acq()`
 - `while (not x < 100):`
 - `cv.wait()`
 - `x ← x+1`
 - `z ← x`
 - `cv.signal_all()`
 - `lck.rel()`
 - `return 2*z`

- `down():`
 - `int z`
 - `lck.acq()`
 - `while (not x > 0):`
 - `cv.wait()`
 - `x ← x-1`
 - `z ← x`
 - `cv.signal_all()`
 - `lck.rel()`
 - `return 2*z`

Program P3:

- `x, y` // as in P1
 - `mutex ← Semaphore(1)` // for lck
 - `gateNF ← Semaphore(0)` // for cvNF
 - `gateNE ← Semaphore(0)` // for cvNE
-
- `up():`
 - `int z`
 - `mutex.P()`
 - `while (not x < 100)`
 - `mutex.V()`
 - `gateNF.P()`
 - `mutex.P()`
 - `x ← x+1`
 - `z ← x`
 - `gateNE.V()`
 - `mutex.V()`
 - `return z`

- `down():`
 - `int z`
 - `mutex.P()`
 - `while (not x > 0)`
 - `mutex.V()`
 - `gateNE.P()`
 - `mutex.P()`
 - `x ← x-1`
 - `z ← x`
 - `gateNF.V()`
 - `mutex.V()`
 - `return z`