


TCSS 422: OPERATING SYSTEMS

Semaphores

Wes J. Lloyd
Institute of Technology
University of Washington - Tacoma



OBJECTIVES

- Semaphores - API
- Uses
- Reader/Writer Locks
- Dining Philosophers

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.2

ANOTHER APPROACH TO CONCURRENCY

- We've looked at **Locks** (ch. 28) and **Conditions** (ch. 30) to provide atomicity in critical sections for concurrency
- Now we'll look at "**semaphores**"
- Provide same functionality
- With different "packaging"

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.3

THE SEMAPHORE

- Semaphores (struct in Linux):
- Contains:
 - Lock
 - Integer: (essentially a counter)
 - List: (thread wait list)

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.4

SEMAPHORE API

■ sem_init():

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

■ Initializes new semaphore:

- First param- address of a semaphore
- Second param: 0- single process, 1- multiprocess
"1" can be used with fork() to synchronize processes
- Third param: initial value of counter

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.5

SEMAPHORE API - 2

■ sem_wait():

- Decrements the value of the semaphore counter, and returns
- Adds thread to wait queue if counter <= 0 until and blocks it

```
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 }
```

- The negative value corresponds to the number of queued, waiting threads

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.6

SEMAPHORE API - 3

- sem_post():
 - Increases the semaphore counter by 1.
 - Awakens a thread on the wait queue (if any)
 - (when counter < 0)

```
1 int sem_post(sem_t *s) {
2     increment the value of semaphore s by one
3     if there are one or more threads waiting, wake one
4 }
```

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.7

SEMAPHORE AS A LOCK

- What should the value of X be below?
 - Consider two threads entering this code, one immediately after the other
 - What should the first thread do?
 - The second thread do?

```
sem_t m;
sem_init(&m, 0, X); // initialize semaphore to X

sem_wait(&m);       // similar to lock
// critical section goes here
sem_post(&m);       // similar to unlock
```

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.8

TWO THREADS AND A SEMAPHORE

- Semaphore as a lock:

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

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.9

SEMAPHORE AS A CONDITION VARIABLE

- Semaphores can be thought of as “mutants”
 - They can be used as locks, or condition variables
- Consider an example

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.10

SEMAPHORE AS A CONDITION VARIABLE - 2

- What should be the value of X ?

```
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 }
```

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.11

ORDERING OF EXECUTION – 1 OF 2

- Parent calls sem_wait() before child calls sem_post()

Value	Parent	State	Child	State
0	Create (Child)	Running	(Child exists; is runnable)	Ready
0	call sem_wait()	Running		Ready
-1	decrement sem	Running		Ready
-1	(sem < 0) → sleep	sleeping		Ready
-1	Switch → Child	sleeping	child runs	Running
-1		sleeping	call sem_post()	Running
0		sleeping	increment sem	Running
0		Ready	wake (Parent)	Running
0		Ready	sem_post() returns	Running
0		Ready	Interrupt; Switch → Parent	Ready
0	sem_wait() returns	Running		Ready

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.12

ORDERING OF EXECUTION - 2 OF 2

- Child runs, calls `sem_post()` before parent calls `sem_wait()`

Value	Parent	State	Child	State
0	Create (Child)	Running	(Child exists; is runnable)	Ready
0	Interrupt; switch-Child	Ready	child runs	Running
0		Ready	call <code>sem_post()</code>	Running
1		Ready	increment sem	Running
1		Ready	wake (nobody)	Running
1		Ready	<code>sem_post()</code> returns	Running
1	parent runs	Running	Interrupt; Switch-Parent	Ready
1	call <code>sem_wait()</code>	Running		Ready
0	decrement sem	Running		Ready
0	(sem<0) → awake	Running		Ready
0	<code>sem_wait()</code> returns	Running		Ready

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.13

PRODUCER/CONSUMER W/ SEMAPHORES

- Producer: `put()`
- Consumer: `get()`
- With `MAX=1`, 1 consumer thread, 1 producer thread:

```

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 void get() {
11     int tmp = buffer[use]; // line g1
12     use = (use + 1) % MAX; // line g2
13     return tmp;
14 }

```

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.14

PRODUCER/CONSUMER W/ SEMAPHORES - 2

```

1  sem_t empty;
2  sem_t full;
3
4  void *producer(void *arg) {
5      int i;
6      for (i = 0; i < loops; i++) {
7          sem_wait(&empty); // line P1
8          put(i); // line P2
9          sem_post(&full); // line P3
10     }
11 }
12
13 void *consumer(void *arg) {
14     int i, tmp = 0;
15     while (tmp != -1) {
16         sem_wait(&full); // line C1
17         tmp = get(i); // line C2
18         sem_post(&empty); // line C3
19         printf("%d\n", tmp);
20     }
21 }
22 ...

```

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.15

PRODUCER/CONSUMER W/ SEMAPHORES - 3

```

21 int main(int argc, char *argv[]) {
22     // ...
23     sem_init(&empty, 0, MAX); // MAX buffers are empty to begin with...
24     sem_init(&full, 0, 0); // ... and 0 are full
25     // ...
26 }

```

- This code is sufficient for any size buffer with 1 producer, 1 consumer
- Try it out
- But what happens if we add multiple producers and consumers?
- Try it out
- Must consider critical sections

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.16

MULTI THREAD P/C SEMAPHORES W/ MUTUAL EXCLUSION

- Which part of the code is the critical section?

```

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 (Cont.)

```

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.17

MULTI THREAD P/C SEMAPHORES W/ MUTUAL EXCLUSION - 2

```

(Cont.)
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(i); // line c2
22         sem_post(&empty); // line c3
23         sem_post(&mutex); // line c4 (NEW LINE)
24         printf("%d\n", tmp);
25     }
26 }

```

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.18

EXECUTION FLOW

- With one producer, one consumer
 - Consumer acquires mutex (the lock)
 - Consumer calls `sem_wait()` to wait for data
 - No data available, consumer blocks and yields the CPU
 - Still has mutex (the lock)
 - Producer tries to acquire mutex (the lock)
 - Producer becomes stuck in **deadlock**
 - Consumer is waiting for data, and will never release the mutex

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.19

MULTITHREAD P/C W/ SEMAPHORES

- Lock should only protect `put()`, `get()`

```
1 sem_t empty;
2 sem_t full;
3 sem_t mutex;
4
5 void *producer(void *arg) {
6     int i;
7     for (i = 0; i < loops; i++) {
8         sem_wait(&empty);           // line p1
9         sem_wait(&mutex);           // line p1.5 (MOVED MUTEX HERE...)
10        put(i);                     // line p2
11        sem_post(&mutex);           // line p2.5 (... AND HERE)
12        sem_post(&full);            // line p3
13    }
14 }
15 (Cont.)
```

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.20

MULTITHREAD P/C W/ SEMAPHORES - 2

- Try it out...

```
(Cont.)
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 }
```

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.21

CONCURRENT DATA STRUCTURES

- Concurrent data structures ideally will:
 - Ensure atomicity of writes
 - Enable multiple synchronous reads
 - As long as elements being read are not concurrently changed
- Concurrent linked list, use a **Reader-Writer Lock**
 - Insert
 - Has traditional critical section which **must not** be multiply entered
 - Read
 - Should support concurrent reads if not being changed
 - Semaphores: *good for tracking concurrent reads*

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.22

CONCURRENT LIST WITH SEMAPHORES

- Multiple readers can acquire a lock
 - Writer must wait until all readers finish

```
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     -
```

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.23

CONCURRENT LIST WITH SEMAPHORES - 2

```
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 }
```

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.24

READER-WRITER LOCK

- Fairness problem
 - With many readers, it becomes difficult for a writer to obtain the lock
 - One improvement is to prevent more readers from reading once a writer is waiting for the lock
 - **How could we implement this improvement?**

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma


L10.25

DINING PHILOSOPHERS PROBLEM

- Classic computer science problem
- Possible job interview question

Philosopher's

1. Think
2. Pick up forks (wait if not available)
3. Eat
4. Put down forks




October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.26

DINING PHILOSOPHERS - 2

- P- Philosopher
- f- fork (eating utensil)
- Key challenges
 - There is no **deadlock**
 - No philosopher starves
 - Concurrency is high
 - Forks get used as much as possible
 - Philosophers have plenty of eating opportunities



October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.27

DINING PHILOSOPHERS - 3

Philosophers:


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

Fork helper functions

```
// helper functions
int left(int p) { return p; }

int right(int p) {
    return (p + 1) % 5;
}
```

- Fork on left: left(P1)
- Fork on right: right(P1)



October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.28

DINING PHILOSOPHERS - 4

- If we just protect the forks with semaphores:

```
void getforks() {
    sem_wait(forks[left(p)]);
    sem_wait(forks[right(p)]);
}

void putforks() {
    sem_post(forks[left(p)]);
    sem_post(forks[right(p)]);
}
```

- Try this:

October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.29


DINING PHILOSOPHERS - 5

Complete the table

Philosopher	LEFT	RIGHT
P0		
P1		
P2		
P3		
P4		

```
void getforks() {
    think();
    getforks();
    eat();
    putforks();
}

void putforks() {
    sem_post(forks[left(p)]);
    sem_post(forks[right(p)]);
}
```



October 24, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L10.30

DINING PHILOSOPHERS - 5

■ **DEADLOCK: All Philosophers Starve!**

Philosopher	LEFT	RIGHT
P0	acquires f0	waits for f1
P1	acquires f1	waits for f2
P2	acquires f2	waits for f3
P3	acquires f3	waits for f4
P4	acquires f4	

```

void getforks() {
    sem_wait(&f0);
    sem_wait(&f1);
}
void putforks() {
    sem_post(&f0);
    sem_post(&f1);
}
    
```

October 24, 2016 TCSS422: Operating Systems [Fall 2016] Institute of Technology, University of Washington - Tacoma L10.31

ALTERNATE PHILOSOPHER

- We need another approach to acquiring forks
- Consider which fork philosophers grab first
- What if we have a **alternate-handed philosopher?**

```

void getforks() {
    if (p == 4) {
        sem_wait(&f0);
        sem_wait(&f1);
    } else {
        sem_wait(&f1);
        sem_wait(&f0);
    }
}
    
```

- **Solves the Dining Philosopher's problem !!!**
- Remember that one philosopher grabs a different fork

October 24, 2016 TCSS422: Operating Systems [Fall 2016] Institute of Technology, University of Washington - Tacoma L10.32

ALTERNATE PHILOSOPHER - 2

■ **P3 eats! Solves deadlock**

Philosopher	LEFT	RIGHT
P0	acquires f0	waits for f1
P1	acquires f1	waits for f2
P2	acquires f2	waits for f3
P3	acquires f3	acquires f4, eats...
P4		Waits for f0

```

void getforks() {
    if (p == 4) {
        sem_wait(&f0);
        sem_wait(&f1);
    } else {
        sem_wait(&f1);
        sem_wait(&f0);
    }
}
    
```

October 24, 2016 TCSS422: Operating Systems [Fall 2016] Institute of Technology, University of Washington - Tacoma L10.33

SEMAPHORE IMPLEMENTATION

- Semaphores can be built using locks and conditions
 - pthread_mutex_t
 - pthread_cond_t
- Linux implementation
 - Does not track negative counter values
 - Easier to implement
- Zemaphore

October 24, 2016 TCSS422: Operating Systems [Fall 2016] Institute of Technology, University of Washington - Tacoma L10.34

SEMAPHORE IMPLEMENTATION - 2

```

1 typedef struct __sem_t {
2     int value;
3     pthread_cond_t cond;
4     pthread_mutex_t lock;
5 } sem_t;
6
7 // only one thread can call this
8 void sem_init(sem_t *s, int value) {
9     s->value = value;
10    Cond_init(&s->cond);
11    Mutex_init(&s->lock);
12 }
13
14 void sem_wait(sem_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 sem_post(sem_t *s) {
23     Mutex_lock(&s->lock);
24     s->value++;
25     Cond_signal(&s->cond);
26     Mutex_unlock(&s->lock);
27 }
    
```

October 24, 2016 TCSS422: Operating Systems [Fall 2016] Institute of Technology, University of Washington - Tacoma L10.35

SEMAPHORES SUMMARY

- Provide one construct for both concurrency features
 - Binary semaphore: provides basic mutex lock
 - Ensures mutual exclusion in critical sections
 - Condition semaphore: Synchronize one or more threads which need to wait for something to happen
 - Allows fewer concurrency related variables in your code
 - Potentially makes code more ambiguous
- After seeing Locks, Conditions, and Semaphores, Which do you like better?

October 24, 2016 TCSS422: Operating Systems [Fall 2016] Institute of Technology, University of Washington - Tacoma L10.36

