

CS 202: Advanced Operating Systems

Synchronization, Memory Consistency, and Cache Coherence *(some cache coherence slides adapted from Ian Watson; some memory consistency slides from Sarita Adve)*

Classic Example


- › Suppose we have to implement a function to handle withdrawals from a bank account:

```
withdraw (account, amount) {  
    balance = get_balance(account);  
    balance = balance – amount;  
    put_balance(account, balance);  
    return balance;  
}
```

- › Now suppose that you and your father share a bank account with a balance of \$1000
- › Then you each go to separate ATM machines and simultaneously withdraw \$100 from the account

Interleaved Schedules

- The problem is that the execution of the two threads can be interleaved:



- What is the balance of the account now?

Shared Resources




- Problem: two threads accessed a **shared resource**
 - Known as a **race condition** (remember this buzzword!)
- Need mechanisms to control this access
 - So we can reason about how the program will operate
- Our example was updating a shared bank account
- Also necessary for synchronizing access to **any shared data structure**
 - Buffers, queues, lists, hash tables, etc.

When Are Resources Shared?



› Local variables?

- › Not shared: refer to data on the stack
- › Each thread has its own stack
- › Never pass/share/store a pointer to a local variable on the stack for thread T1 to another thread T2



› Global variables and static objects?

- › **Shared:** in static data segment, accessible by all threads

› Dynamic objects and other heap objects?

- › **Shared:** Allocated from heap with malloc/free or new/delete

How Interleaved Can It Get?



How contorted can the interleavings be?

- ▶ We'll assume that the only atomic operations are reads and writes of individual memory locations
 - ▶ Some architectures don't even give you that!
- ▶ We'll assume that a **context switch can occur at any time**
- ▶ We'll assume that **you can delay a thread as long as you like as long as it's not delayed forever**

```
..... get_balance(account);
```

```
balance = get_balance(account);
```

```
balance = .....
```

```
balance = balance - amount;
```

```
balance = balance - amount;
```

```
put_balance(account, balance);
```

```
put_balance(account, balance);
```

Mutual Exclusion

- **Mutual exclusion** to synchronize access to shared resources
 - This allows us to have larger atomic blocks
 - What does atomic mean?
- Code that uses mutual called a **critical section**
 - Only one thread at a time can execute in the critical section
 - All other threads are forced to wait on entry
 - When a thread leaves a critical section, another can enter
 - Example: sharing an ATM with others
- What requirements would you place on a critical section?

Using Locks

```
withdraw (account, amount) {  
    acquire(lock);  
    balance = get_balance(account);  
    balance = balance - amount;  
    put_balance(account, balance);  
    release(lock);  
    return balance;  
}
```

**Critical
Section**

acquire(lock);
balance = get_balance(account);
balance = balance - amount;

acquire(lock);

put_balance(account, balance);
release(lock);

balance = get_balance(account);
balance = balance - amount;
put_balance(account, balance);
release(lock);

Using Test-And-Set

- Here is our lock implementation with test-and-set:

```
struct lock {  
    int held = 0;  
}  
void acquire (lock) {  
    while (test-and-set(&lock->held));  
}  
void release (lock) {  
    lock->held = 0;  
}
```

```
//Atomic!!!  
test-and-set(lock) {  
    boolean initial = lock;  
    lock = true;  
    return initial;  
}
```

- When will the while return? What is the value of held?
- Does it satisfy critical region requirements? (mutex, progress, bounded wait, performance?)

Another solution: Disabling Interrupts

- › Another implementation of acquire/release is to disable interrupts:

```
struct lock {  
}  
void acquire (lock) {  
    disable interrupts;  
}  
void release (lock) {  
    enable interrupts;  
}
```

- › Note that there is no state associated with the lock
- › Can two threads disable interrupts simultaneously?

On Disabling Interrupts



- › Disabling interrupts blocks notification of external events that could trigger a context switch (e.g., timer)
- › In a “real” system, this is only available to the kernel
 - › Why?
- › Disabling interrupts is insufficient on a multiprocessor
 - › Back to atomic instructions
- › Like spinlocks, only want to disable interrupts to implement higher-level synchronization primitives
 - › Don't want interrupts disabled between acquire and release


Summarize Where We Are



- Goal: Use **mutual exclusion** to protect **critical sections** of code that access **shared resources**
- Method: Use locks (spinlocks or disable interrupts)
- Problem: Critical sections can be long

Spinlocks:

- Threads waiting to acquire lock spin in test-and-set loop
- Wastes CPU cycles
- Longer the CS, the longer the spin
- Greater the chance for lock holder to be interrupted
- Memory consistency model causes problems (out of scope of this class)



Disabling Interrupts:

- Should not disable interrupts for long periods of time
- Can miss or delay important events (e.g., timer, I/O)
- Doesn't work for multiprocessor

Semaphore



Higher-Level Synchronization



- › Locks so far inefficient when critical sections are long
 - › Spinlocks – inefficient
 - › Disabling interrupts – can miss or delay important events
- › Instead, we want synchronization mechanisms that
 - › Block waiters
 - › Leave interrupts enabled inside the critical section
- › Plan:
 - › **Semaphores**: binary (mutex) and counting
 - › Use them to solve common synchronization problems

Semaphores



- › Semaphores are an **abstract data type** that provide mutual exclusion to critical sections
 - › Block waiters, interrupts enabled within critical section
 - › Described by Dijkstra in THE system in 1968
- › Semaphores are integers that support two operations:
 - › `sem_wait(sem_t *s)`: wait if the value is zero; otherwise, decrement
 - › Also P(), after the Dutch word for test, or down()
 - › `sem_post(sem_t *s)`: if one thread waiting, wake it; otherwise, increment
 - › Also V() after the Dutch word for increment, or up()
 - › That's it! No other operations – not even just reading its value – exist
- › Semaphore safety property: the semaphore value is always greater than or equal to 0

Semaphore Types



- › Semaphores come in two types
- › **Mutex** semaphore (or **binary** semaphore)
 - › Represents single access to a resource
 - › Guarantees mutual exclusion to a critical section
- › **Counting** semaphore (or **general** semaphore)
 - › Multiple threads pass the semaphore determined by count
 - › mutex has count = 1, counting has count = N
 - › Represents a resource with many units available
 - › or a resource allowing some unsynchronized concurrent access (e.g., reading)

Protecting a critical region



```
1  sem_t m;  
2  sem_init(&m, 0, X); // initialize 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)

A parent waiting for its child

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

Figure 31.6: A Parent Waiting For Its Child

Producer/ Consumer (Bounded Buffer)

```
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: The Put And Get Routines

```
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();          // Line C2
18         sem_post(&empty);     // Line C3
19         printf("%d\n", tmp);
20     }
21 }
22
23 int main(int argc, char *argv[]) {
24     // ...
25     sem_init(&empty, 0, MAX); // MAX are empty
26     sem_init(&full, 0, 0);    // 0 are full
27     // ...
28 }
```

Figure 31.10: Adding The Full And Empty Conditions

Second Attempt (Add Mutual Exclusion)



```
1 void *producer(void *arg) {
2     int i;
3     for (i = 0; i < loops; i++) {
4         sem_wait(&mutex);          // Line P0 (NEW LINE)
5         sem_wait(&empty);          // Line P1
6         put(i);                    // Line P2
7         sem_post(&full);           // Line P3
8         sem_post(&mutex);          // Line P4 (NEW LINE)
9     }
10 }
11
12 void *consumer(void *arg) {
13     int i;
14     for (i = 0; i < loops; i++) {
15         sem_wait(&mutex);          // Line C0 (NEW LINE)
16         sem_wait(&full);           // Line C1
17         int tmp = get();           // Line C2
18         sem_post(&empty);          // Line C3
19         sem_post(&mutex);          // Line C4 (NEW LINE)
20         printf("%d\n", tmp);
21     }
22 }
```

Third Attempt

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


Any more improvement?

Reader-Writer Lock



```
1  typedef struct _rwlock_t {
2      sem_t lock;          // binary semaphore (basic lock)
3      sem_t writelock;     // allow ONE writer/MANY readers
4      int   readers;       // #readers 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) // first reader gets writelock
17         sem_wait(&rw->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) // last reader lets it go
25         sem_post(&rw->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 }
```

Dining Philosophers



```
while (1) {  
    think();  
    get_forks(p);  
    eat();  
    put_forks(p);  
}  
  
int left(int p) { return p; }  
int right(int p) { return (p + 1) % 5; }  
  
void get_forks(int p) {  
    sem_wait(&forks[left(p)]);  
    sem_wait(&forks[right(p)]);  
}  
  
void put_forks(int p) {  
    sem_post(&forks[left(p)]);  
    sem_post(&forks[right(p)]);  
}
```

What is wrong?

How to fix it?

Solution: Break Dependency



```
void get_forks(int p) {  
    if (p == 4) {  
        sem_wait(&forks[right(p)]);  
        sem_wait(&forks[left(p)]);  
    } else {  
        sem_wait(&forks[left(p)]);  
        sem_wait(&forks[right(p)]);  
    }  
}
```

Can you think of another solution?

Semaphore Summary



- Semaphores can be used to solve any of the traditional synchronization problems
- However, they have some drawbacks
 - They are essentially shared global variables
 - Can potentially be accessed anywhere in program
 - No connection between the semaphore and the data being controlled by the semaphore
 - Used both for critical sections (mutual exclusion) and coordination (scheduling)
 - Note that I had to use comments in the code to distinguish
 - No control or guarantee of proper usage
- Sometimes hard to use and prone to bugs
 - Another approach: Use programming language support

Overview



- › Before we talk deeply about synchronization
 - › Need to get an idea about the memory model in shared memory systems
 - › Is synchronization only an issue in multi-processor systems?
- › What is a shared memory processor (SMP)?
- › Shared memory processors
 - › Two primary architectures:
 - › Bus-based/local network shared-memory machines (small-scale)
 - › Directory-based shared-memory machines (large-scale)