

Sistemas Operativos: Concurrency

Locks

Pedro F. Souto (pfs@fe.up.pt)

March 18, 2020

Roadmap

Locks Usage

Implementation

Using locks to prevent race conditions

Question Is there a systematic way to prevent race conditions?

Answer The closest is the concept of *monitor*

- ▶ Use *abstract data types* (ADT), a kind of class, to structure your code;
- ▶ Add a lock to ensure mutual exclusion in the execution of the functions of the ADT:
 - ▶ The first step in each function is to acquire that lock
 - ▶ The last step in each function is to release that lock

The Counter ADT

```
1  typedef struct __counter_t {
2      int value;
3  } counter_t;
4
5  void init(counter_t *c) {
6      c->value = 0;
7  }
8
9  void increment(counter_t *c) {
10     c->value++;
11 }
12
13 void decrement(counter_t *c) {
14     c->value--;
15 }
16
17 int get(counter_t *c) {
18     return c->value;
19 }
```

The Counter Monitor

```
1  typedef struct __counter_t {
2      int          value;
3      pthread_mutex_t lock;
4  } counter_t;
5
6  void init(counter_t *c) {
7      c->value = 0;
8      Pthread_mutex_init(&c->lock, NULL);
9  }
10
11 void increment(counter_t *c) {
12     Pthread_mutex_lock(&c->lock);
13     c->value++;
14     Pthread_mutex_unlock(&c->lock);
15 }
16
17 void decrement(counter_t *c) {
18     Pthread_mutex_lock(&c->lock);
19     c->value--;
20     Pthread_mutex_unlock(&c->lock);
21 }
```

Monitors

- ▶ This is not the only way to design thread-safe programs
- ▶ Actually, the use of monitors may raise some modularity issues
 - ▶ Namely deadlocks as we will see later
- ▶ Furthermore, the performance may not be the best
 - ▶ If code uses several ADTs, may need to acquire the locks on all of them
- ▶ Nevertheless, this is a rather useful pattern
 - ▶ `java.util.concurrent` is a Java package, i.e. library, that provides thread-safe versions of many classes in the `java.util` package
 - ▶ Its implementation relies on the concept of monitor
- ▶ Problem 2 of this week's problem set asks you to develop a thread-safe stack using this pattern

Roadmap

Locks Usage

Implementation

Implementation Goals

Mutual exclusion that is the main purpose of locks

Fairness threads should have a fair chance to acquire a lock

- ▶ This is not always desired
- ▶ But at least threads should not starve

Performance There many facets of this

Without contention

With contention On uniprocessors vs.
multiprocessors/multicores

Approaches

- ▶ Software-based solutions
 - ▶ More of an intellectual exercise: they are not efficient
- ▶ Controlling Interrupts
- ▶ Atomic read-modify-write instructions

Controlling Interrupts

Idea Disable interrupts upon entering a critical section and re-enable them upon exit

```
void lock() {  
    DisableInterrupt();  
}  
void unlock() {  
    EnableInterrupt();  
}
```

Rationale By disabling interrupts, the thread will execute without interruption the critical section.

Advantage This is a simple solution that works

Controlling Interrupts: Issues

Requires the OS to trust applications If the application does not call `unlock()` the OS will not be able to regain control

- ▶ The only way out is to restart the system

Reduces responsiveness while the interrupts are disabled, the system cannot respond to interrupts

- ▶ May lead to lost interrupts, e.g. of the timer

Only works on uniprocessors

- ▶ Nothing prevents threads executing on other cores from entering interfering critical sections
- ▶ Furthermore, disabling interrupts operates on a single core;

Low performance

- ▶ On modern HW, interrupt-related instructions are slower than atomic read-modify-write instructions

Concluding On modern OSs controlling interrupts is done only at the kernel, mostly to prevent interrupt handling code from being interrupted.

Attempt to Implement Locks with Ordinary Instructions

```
1  typedef struct __lock_t { int flag; } lock_t;
2
3  void init(lock_t *mutex) {
4      // 0 -> lock is available, 1 -> held
5      mutex->flag = 0;
6  }
7
8  void lock(lock_t *mutex) {
9      while (mutex->flag == 1) // TEST the flag
10         ; // spin-wait (do nothing)
11     mutex->flag = 1;          // now SET it!
12 }
13
14 void unlock(lock_t *mutex) {
15     mutex->flag = 0;
16 }
```

Issues

Performance threads must busy wait

Correctness i.e. does not ensure mutual exclusion always – can you see why?

Attempt to Implement Locks with Ordinary Instructions

Thread 1

call `lock()`

while (`flag == 1`)

interrupt: switch to Thread 2

`flag = 1; // set flag to 1 (too!)`

Thread 2

call `lock()`

while (`flag == 1`)

`flag = 1;`

interrupt: switch to Thread 1

Issues

Solution Use atomic read-modify-write instructions

Atomic Exchange (Test-And-Set)

- ▶ This is the simplest atomic read-modify-write instruction: in the Intel32 ISA it is known as `xchg`

```
int TestAndSet(int *old_ptr, new) {  
    int old = *old_ptr;  
    *old_ptr = new;  
    return old;  
}
```

- ▶ Remzi called it `TestAndSet()` because it allows testing the value and possibly modifying it (but there are instructions that do exactly that, Remzi calls them `CompareAndSwap()`)
- ▶ The key is that this instruction is atomic:
 - ▶ The reading and the modification of the memory whose address is `old_ptr` is done in an indivisible way

Question How can we use this instruction to solve the problem with our last attempt?

Spin-Lock Implementation with Atomic TestAndSet

```
1  typedef struct __lock_t {
2      int flag;
3  } lock_t;
4
5  void init(lock_t *lock) {
6      // 0 indicates that lock is available, 1 that it is
7      lock->flag = 0;
8  }
9
10 void lock(lock_t *lock) {
11     while (TestAndSet(&lock->flag, 1) == 1)
12         ; // spin-wait (do nothing)
13 }
14
15 void unlock(lock_t *lock) {
16     lock->flag = 0;
17 }
```

Analysis

Entering a CS when the lock is not held

Entering a CS when the lock is held by another thread

Spin-Lock with Atomic TestAndSet: Evaluation

Correctness as long as the scheduler is preemptive

Fairness there is no guarantee

- ▶ Depends on the scheduler

Performance

Uniprocessor Poor

- ▶ While a thread spins waiting, other threads cannot run

Multiprocessor May be good, if:

- ▶ There is low contention
- ▶ Lock is held by a thread on another core/processor
- ▶ Critical section is short

Spin-Lock Performance Issues

- ▶ If the thread inside a CS is preempted/interrupted
 - ▶ The thread trying to enter the CS will be forced to spin
 - ▶ Possibly until it is preempted
- ▶ The higher the contention, i.e. the higher the number of threads trying to enter the CS, the worse the performance.

Solution (first try) Yield the CPU if lock is held by another thread

Lock with `yield()`

```
1 void init() {
2     flag = 0;
3 }
4
5 void lock() {
6     while (TestAndSet(&flag, 1) == 1)
7         yield(); // give up the CPU
8 }
9
10 void unlock() {
11     flag = 0;
12 }
```

Issues

Performance On high contention

- ▶ Too many `yield()`'s

Fairness depends on the scheduler

Locks: Avoiding Busy-Waiting

- ▶ Need more control on which thread gets the lock
 - ▶ Requires OS support
- ▶ E.g. Solaris offers two system-calls:

`park()` similar to `sleep()`

`unpark()` kind of wakeup

Idea Before "parking" add the thread to a queue of threads waiting for the lock

- ▶ Upon unlocking, "unpark" the thread at the head of the queue

Lock with queues and park ()

```
13 void lock(lock_t *m) {
14     while (TestAndSet(&m->guard, 1) == 1)
15         ; //acquire guard lock by spinning
16     if (m->flag == 0) {
17         m->flag = 1; // lock is acquired
18         m->guard = 0;
19     } else {
20         queue_add(m->q, gettid());
21         m->guard = 0;
22         park();
23     }
24 }
25
26 void unlock(lock_t *m) {
27     while (TestAndSet(&m->guard, 1) == 1)
28         ; //acquire guard lock by spinning
29     if (queue_empty(m->q))
30         m->flag = 0; // let go of lock; no one wants it
31     else
32         unpark(queue_remove(m->q)); // hold lock (for next t
33     m->guard = 0;
34 }
```

- But this has a race-condition known as **lost-wakeup**

Fixing the lost-wakeup or wakeup/waiting race

- ▶ Use the `setpark()` system call:

- ▶ It tells the OS that the thread is about to call `park()` (may be `prepare_park()` would be clearer)

```
queue_add(m->q, gettid());
setpark(); // tell the kernel thread is about
m->guard = 0;
park();
```

- ▶ If there is a call to `unpark()` between `setpark()` and `park()`, the latter returns immediately

- ▶ `park()/unpark()/setpark()` are OS-specific

- ▶ Linux offers the `futex()` system call with the same purpose:

`futex_wait(address, expected)` blocks the thread, if the value @ address is expected (otherwise, does not block);

`futex_wake(address)` wakes one thread that is waiting on address

mutex_lock() with futex (lowlevellock.h)

```
1 void mutex_lock (int *mutex) {
2     int v;
3     /* Bit 31 was clear, we got the mutex (this is the fastpath) */
4     if (atomic_bit_test_set (mutex, 31) == 0)
5         return;
6     atomic_increment (mutex);
7     while (1) {
8         if (atomic_bit_test_set (mutex, 31) == 0) {
9             atomic_decrement (mutex);
10            return;
11        }
12        /* We have to wait now. First make sure the futex value
13         we are monitoring is truly negative (i.e. locked). */
14        v = *mutex;
15        if (v >= 0)
16            continue;
17        futex_wait (mutex, v);
18    }
19 }
20
21 void mutex_unlock (int *mutex) {
22     /* Adding 0x80000000 to the counter results in 0 if and only if
23     there are not other interested threads */
24     if (atomic_add_zero (mutex, 0x80000000))
25         return;
26
27     /* There are other threads waiting for this mutex,
28     wake one of them up. */
29     futex_wake (mutex);
30 }
```