

Concurrency

- Exam 1
- 26 - Concurrency
- 27 - Overview, and POSIX threads (pthreads)
- 28 - Locks
- 29 - Concurrent Data Structures
- 30 Condition Variables
- 31 - Semaphores
- 32 - Common Problems
- 33 - Event-Based Concurrency

228

Concurrency

- Exam 1
- 26 - Concurrency
- 27 - Overview, and POSIX threads (pthreads)
- 28 - Lock Details
- 29 - Concurrent Data Structures
- 30 Condition Variable Details
- 31 - Semaphores
- 32 - Common Problems
- 33 - Event-Based Concurrency

229

Locks *the basic idea*

- Ensure that any critical sections executes atomically
 - Canonical update of a shared variable:

```
balance = balance + 1;
```

- Use locks:

```
1  lock_t mutex; // some globally-allocated lock 'mutex'  
2  ...  
3  lock(&mutex);  
4  balance = balance + 1;  
5  unlock(&mutex);
```

230

Locks *the basic idea*

- Lock variables hold the *lock state*:
 - *unlocked* (or available, or free)
 - no thread holds the lock
 - *locked* (or acquired or held)
 - exactly one thread holds the lock
 - presumably in the critical section

231

Locks *semantics*

- `lock()`
 - acquired if no other thread holds it
 - enter *critical section*
 - calling thread is now the lock's *owner*
 - other threads prevented from entering the critical section
 - assuming proper lock discipline

```
1 pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;
2
3 pthread_mutex_lock(&lock);
4 balance = balance + 1;
5 pthread_mutex_unlock(&lock);
```

- using many mutexes increases concurrency
- efficient locks require help from *hardware* and the *OS*

232

Goal and Metrics *for locks*

- Mutual exclusion:
 - does it work?
 - *correctness*
- Fairness:
 - can threads *starve*?
 - do they get a fair share?
- Performant:
 - how much overhead

233

Lock implementation *controlling interrupts*

- Half-century-old approach:
 - *disable interrupts* for critical sections
 - even for single-processors:

```
1. void lock() {
2.     DisableInterrupts();
3. }
4. void unlock() {
5.     EnableInterrupts();
6. }
```

- Problems:
 - requires *trust* in applications
 - greedy program might not enable interrupts until done
 - not sufficient for multiprocessors
 - expensive
 - is the above implementation correct and complete?

234

Lock implementation *do we really need hardware?*

- First attempt:
 - use a *flag* to show if lock held:

```
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. }
```

- This code has problems:
 - correctness
 - efficiency

235

Lock implementation *do we really need hardware?*

- Correctness: (no mutual exclusion)

Thread1	Thread2
<pre>call lock() while (flag == 1) interrupt: switch to Thread 2</pre>	<pre>call lock() while (flag == 1) flag = 1; interrupt: switch to Thread 1</pre>
<pre>flag = 1; // set flag to 1 (too!)</pre>	

- Performance:
 - spin-waiting
 - not doing useful work
 - might be actively *preventing* the lock from being released

236

Peterson's algorithm *do we really need hardware?*

```
int flag[2];
int turn;

void init() {
    // indicate you intend to hold the lock w/ 'flag'
    flag[0] = flag[1] = 0;
    // whose turn is it? (thread 0 or 1)
    turn = 0;
}

void lock() {
    // 'self' is the thread ID of caller
    flag[self] = 1;
    // make it other thread's turn
    turn = 1 - self;
    while ((flag[1-self] == 1) && (turn == 1 - self))
        ; // spin-wait while it's not your turn
}

void unlock() {
    // simply undo your intent
    flag[self] = 0;
}
```

doesn't work w/ relaxed consistency models

already had hardware support when this written

still important for understanding synchronization

237

Test-and-Set *hardware support*

- single atomic hardware instruction
- pseudocode:

```
1. int TestAndSet(int *ptr, int new) {
2.     int old = *ptr; // fetch old value at ptr
3.     *ptr = new;    // store 'new' into ptr
4.     return old;    // return the old value
5. }
```

- *returns* old value to be tested
- *simultaneously* updates value to new

238

Test-and-Set *making a spin lock*

```
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,
7.     // 1 that it is held
8.     lock->flag = 0;
9. }
10.
11. void lock(lock_t *lock) {
12.     while (TestAndSet(&lock->flag, 1) == 1)
13.         ; // spin-wait
14. }
15.
16. void unlock(lock_t *lock) {
17.     lock->flag = 0;
18. }
```

- requires a *preemptive scheduler*, even for single processor

239

Test-and-Set *Goal and Metrics*

- **Mutual exclusion:** yes
 - does it work?
 - *correctness*
- **Fairness:** no
 - can threads *starve*?
 - do they get a fair share?
- **Performant:** not usually
 - on single CPU often quite bad
 - may be ok if:
 - *#threads* about the same as *#processors*

240

Compare-and-Swap *hardware support*

- **Test whether** `*ptr == expected_value`
 - *if so*: update `*ptr` with `expected_value`
 - *always*: return actual value from prior to instruction
- **pseudocode:**

```
1. int CompareAndSwap(int *ptr, int expected, int new) {
2.     int actual = *ptr;
3.     if (actual == expected)
4.         *ptr = new;
5.     return actual;
6. }
```

- **Spin lock using compare-and-swap:**

```
1. void lock(lock_t *lock) {
2.     while (CompareAndSwap(&lock->flag, 0, 1) == 1)
3.         ; // spin
4. }
```

- **vs test-and-set?**
 - more powerful

241

Load-Linked Stores *hardware support*

- pseudocode:

```
1. int LoadLinked(int *ptr) {
2.     return *ptr;
3. }
4.
5. int StoreConditional(int *ptr, int value) {
6.     if (*ptr not updated since the LoadLinked to this address) {
7.         *ptr = value;
8.         return 1; // success!
9.     } else {
10.        return 0; // failed to update
11.    }
12. }
```

- only succeeds if *no intervening store to same address*
 - success: 1 is returned, and update *ptr to value
 - failure: 0 is returned, no change to *ptr
- vs test-and-set?
 - more powerful
- can be efficient for hardware

242

Locks *so much spinning*

- Hardware-based spin locks are simple and correct
 - they can also be very inefficient....
- Address with OS support:
 - instead of spinning, just *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. }
```

243

Using queues *sleeping instead of spinning*

- Use a queue to track threads waiting to enter a lock
 - `park()` : put calling thread to sleep
 - `unpark(threadID)` : wake specific thread

```
1. typedef struct __lock_t { int flag; int guard; queue_t *q; } lock_t;
2.
3. void lock_init(lock_t *m) {
4.     m->flag = 0;
5.     m->guard = 0;
6.     queue_init(m->q);
7. }
8.
9. void lock(lock_t *m) {
10.    while (TestAndSet(&m->guard, 1) == 1)
11.        ; // acquire guard lock by spinning
12.    if (m->flag == 0) {
13.        m->flag = 1; // lock is acquired
14.        m->guard = 0;
15.    } else {
16.        queue_add(m->q, gettid());
17.        m->guard = 0;
18.        park();
19.    }
20. }
21. ...
```

244

Using queues *sleeping instead of spinning*

```
void unlock(lock_t *m) {
    while (TestAndSet(&m->guard, 1) == 1)
        ; // acquire guard lock by spinning
    if (queue_empty(m->q))
        m->flag = 0; // let go of lock; no one wants it
    else
        unpark(queue_remove(m->q)); // hold lock (for next thread!)
    m->guard = 0;
}
```

245

Using queues *sleeping instead of spinning*

- There is a race between waking up and waiting
 - Think of releasing a lock in T_A just before T_B calls `park()`
 - T_B could sleep forever...
- Solaris solves by adding a third system call: `setpark()`
 - indicates that a thread is *about to park*
 - if a thread is interrupted, and another thread calls `unpark()` before park actually happens, the `park()` returns immediately

```
1.     queue_add(m->q, gettid());
2.     setpark(); // new code
3.     m->guard = 0;
4.     park();
```

246

Concurrency

- Exam 1
- 26 - Concurrency
- 27 - Overview, and POSIX threads (`pthread`s)
- 28 - Locks
- 29 - Concurrent Data Structures
- 30 Condition Variables
- 31 - Semaphores
- 32 - Common Problems
- 33 - Event-Based Concurrency

247

Mutual Exclusion *mechanism summary*

- Mechanisms:
 - disabling interrupts
 - pretty much all we need if single core
 - but
 - privileged instruction
 - need to *trust* thread
 - not efficient
 - doesn't work on multiprocessors
 - atomic instructions
 - test-and-set
 - set memory location to value, returning old value
 - compare-and-swap
 - store at memory location only if it equals specific value
 - load-linked store
 - load from memory location
 - store new value to same location (only if it has not **been updated**)