

CSE 451: Operating Systems

Section 5: Synchronization

Project 2a is due on Wednesday

- * Review the `stthread` code that we give you
 - * If a function is there, then there is probably a good reason why
- * Remember to write more test cases
- * `test-burgers`: your shell may limit the number of threads you can run
 - * `ulimit -u`
- * Has anybody found bugs in the 64-bit code yet?

4/12/2012

2

Synchronization

4/12/2012

3

Synchronization support

- * Processor level:
 - * Disable / enable interrupts
 - * Atomic instructions
- * Operating system / library level:
 - * Special variables: mutexes, semaphores, condition variables
- * Programming language level:
 - * Monitors, Java synchronized methods

4/12/2012

4

Disabling / enabling interrupts

```

Thread A:          Thread B:
  disable_irq()    disable_irq()
  critical_section()  critical_section()
  enable_irq()     enable_irq()

```

- * Prevents context-switches during execution of critical sections
- * Sometimes necessary
- * Many pitfalls

4/12/2012

5

Processor support

- * Atomic instructions:
 - * test-and-set
 - * compare-exchange (x86)
- * Use these to implement higher-level primitives
 - * E.g. test-and-set on x86 (given to you for part 5) is written using compare-exchange
 - * You'll use test-and-set when implementing mutexes and condition variables (part 5)

4/12/2012

6

Processor support

- * Test-and-set using compare-exchange:

```

compare_exchange(lock_t *x, int y, int z):
  if(*x == y)
    *x = z;
    return y;
  else
    return *x;
}

test_and_set(lock_t *lock) {
  ???
}

```

4/12/2012

7

Processor support

- * Test-and-set using compare-exchange:

```

compare_exchange(lock_t *x, int y, int z):
  if(*x == y)
    *x = z;
    return y;
  else
    return *x;
}

test_and_set(lock_t *lock) {
  compare_exchange(lock, 0, 1);
}

```

4/12/2012

8

Preemption

- * You will need to use test-and-set and interrupt disabling in part 5 of project 2 (preemption)
- * You can start thinking about this while completing the code for the first 2 parts: where are the critical sections where I shouldn't be interrupted?

4/12/2012

9

Semaphores

- * Semaphore = a special variable
 - * Manipulated atomically via two operations
 - * P (wait): tries to decrement semaphore
 - * V (signal): increments semaphore
 - * Has a *queue* of waiting threads
 - * If execute wait() and semaphore is available, continue
 - * If not, block on the waiting queue
 - * signal() unblocks a thread on queue

4/12/2012

10

Mutexes

- * What is a mutex?
 - * A binary semaphore (semaphore initialized with value 1)
- * Why use a mutex rather than a low-level lock?
 - * Threads wait on a mutex by *blocking*, rather than *spinning*

4/12/2012

11

How *not* to implement mutexes

- * Definitely not like this:

```
void pthread_user_mutex_lock(
    pthread_mutex_t lock) {
    while (lock->held) { ; }
    lock->held = true;
}
```

- * And also not like this:

```
void pthread_user_mutex_lock(
    pthread_mutex_t lock) {
    while (lock->held) {
        yield();
    }
    lock->held = true;
}
```

4/12/2012

12

Condition variables

- * Let threads block until a certain *event* or *condition* occurs (rather than polling)
- * Associated with some logical *condition* in a program:

```
pthread_mutex_lock(lock);
while (x <= y) {
    pthread_cond_wait(cond, lock);
}
pthread_mutex_unlock(lock);
```

4/12/2012

13

Condition variables

- * Operations:
 - * wait: sleep on wait queue until event happens
 - * signal: wake up *one* thread on wait queue
 - * broadcast: wake up *all* threads on wait queue
- * signal or broadcast is called explicitly by the application when the event / condition occurs

4/12/2012

14

Condition variables

```
pthread_cond_wait(pthread_cond_t cond,
pthread_mutex_t lock)
```

- * Should do the following *atomically*:
 - * Release the lock (to allow someone else to get in)
 - * Add current thread to the waiters for cond
 - * Block thread until awoken (by signal/broadcast)
- * So, application must acquire `lock` before calling `wait()`!
- * Read man page for `pthread_cond_[wait|signal|broadcast]`

4/12/2012

15

Example synchronization problem

- * Late-Night Pizza
 - * A group of students study for CSE 451 exam
 - * Can only study while eating pizza
 - * If a student finds pizza is gone, the student goes to sleep until another pizza arrives
 - * First student to discover pizza is gone orders a new one
 - * Each pizza has S slices

4/12/2012

16

Late-night pizza

- * Each student thread executes the following:

```
while (must_study) {
    pick up a piece of pizza;
    study while eating the pizza;
}
```

4/12/2012

17

Late-night pizza

- * Need to *synchronize* student threads and pizza delivery thread
- * Avoid deadlock
- * When out of pizza, order it *exactly once*
- * No piece of pizza may be consumed by more than one student

4/12/2012

18

Semaphore / mutex solution

- * Shared data:

```
semaphore_t pizza; //Number of
                  //available pizza
                  //resources;
                  //init to 0

semaphore_t deliver; //init to 1

int num_slices = 0;

mutex_t mutex; //protects accesses
               //to num_slices
```

4/12/2012

19

<pre>student_thread { while (must_study) { wait(pizza); acquire(mutex); num_slices--; if (num_slices==0) signal(deliver); release(mutex); study(); } }</pre>	<pre>delivery_guy_thread { while (employed) { wait(deliver); make_pizza(); acquire(mutex); num_slices=S; release(mutex); for (i=0;i<S;i++) signal(pizza); } }</pre>
--	---

20

Condition variable solution

*** Shared data:**

```
int slices=0;
bool has_been_ordered;
Condition order;           //an order has
                           //been placed

Condition deliver;        //a delivery has
                           //been made

Lock mutex;               //protects
                           //"slices";
                           //associated with
                           //both Condition
                           //variables
```

4/12/2012

21

```
Student() {
while(diligent) {
mutex.lock();
if (slices > 0) {
slices--;
}
else {
if(!has_been_ordered){
order.signal(mutex);
has_been_ordered =
true;
}
while (slices <= 0) {
deliver.wait(mutex);
}
slices--;
}
mutex.unlock();
Study();
}
}

DeliveryGuy() {
while(employed) {
mutex.lock();
order.wait(mutex);
makePizza();
slices = S;
has_been_ordered =
false;
mutex.unlock();
deliver.broadcast();
}
}
```

22

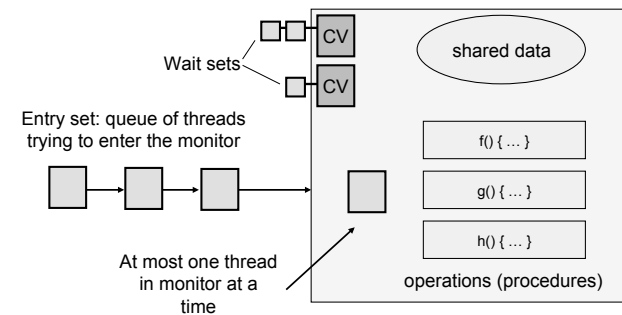
Monitors

- * An object that allows one thread inside at a time
- * Contain a lock and some condition variables
 - * Condition variables used to allow other threads to access the monitor while one thread waits for an event to occur

4/12/2012

23

Monitors



10/28/10

24

Midterm!

11/4/10

25

The kernel

- * Kernel mode vs user mode
 - * How these modes differ conceptually and from the CPU's point of view
 - * How we switch between the two
- * Interrupts

11/4/10

26

System calls

- * What they are
- * What they do
- * How they do it
- * What hardware is involved
- * Who uses them and when

11/4/10

27

Processes and threads

- * Kernel processes, kernel threads, and user threads
 - * How these differ from one another
- * Context switching
- * Process and thread states
- * fork, exec, wait

11/4/10

28

Synchronization

- * Critical sections
- * Locks and atomic instructions
- * Mutexes, semaphores, and condition variables
- * Monitors and how they are implemented
- * Ways to detect / avoid deadlock

11/4/10

29

Scheduling

- * Different scheduling algorithms and their tradeoffs
- * Average response time, various “laws”
- * Starvation
- * Cooperative vs. preemptive scheduling

11/4/10

30

Tips

- * Focus on lecture slides
- * Review textbook, section slides and project writeups to emphasize key concepts and fill in gaps
- * On Monday, when taking the exam:
 - * Arrive early
 - * Focus on key points
 - * Work quickly; finish easy problems first

11/4/10

31