

## Reminders

- n Homework 3 due Monday, Oct. 25
  - n Synchronization
- n Project 2 parts 1,2,3 due Tuesday, Oct. 26
  - n Threads, synchronization
- n Today:
  - n Project 2 continued (parts 2,3)
  - n Synchronization

1

## Project 2 Part 1 Questions

- n Anything about writing user threads?
- n Recall that:
  - n `sthread_create` doesn't immediately run the new thread
  - n `sthread_exit` can ignore its `ret` argument
- n Recall how stacks are allocated
  - n `sthread_new_ctx` – creates a new stack and makes it ready to run after first context switch
  - n `sthread_new_blank_ctx` – create new stack but don't initialize. Suitable to use as `old` parameter to `switch()`.
- n Read `.h` files for function specs!

2

## Synchronization

- n Why do we need it?
  - n ensure correct and efficient cooperation
  - n Prevent race conditions
- n How?
  - n Protect code in critical sections
    - n Allow at most one process/thread in critical section
    - n Maintain fairness & progress
    - n Don't make deadlocks

3

## Synchronization Solutions

**High-level**

- n Monitors
- n Java synchronized method

**OS-level support**

- n Special variables – mutexes, semaphores, condition vars
- n Message passing primitives

**Low-level support**

- n Disable/enable interrupts
- n Atomic instructions

+ Software algorithms ...

4

## Disabling/Enabling Interrupts

<p>Thread A:</p> <pre> disable_interrupts() critical_section() enable_interrupts()           </pre>	<p>Thread B:</p> <pre> disable_interrupts() critical_section() enable_interrupts()           </pre>
---	---

- n Prevents context-switches during execution of CS
- n In Linux: `cli()`, `sti()`
- n Sometimes necessary
  - n E.g. to prevent further interrupts during interrupt handling
- n Problems?

5

## Hardware support

- n Atomic instructions:
  - n Test and set
  - n Swap
  - n Compare-exchange (x86)
  - n Load-linked store conditional (MIPS, Alpha, PowerPC)
- n Use these to implement higher-level primitives
  - n E.g. test-and-set on x86:
 

```

int atomic_test_and_set(lock_t *l) {
    int val;
    __asm {
        mov edx, dword ptr [l] ; Get the pointer to l
        mov ecx, l             ; load l into the cmpxchg source
        mov eax, 0             ; load 0 into the accumulator
                                ; if l == 0 then
        lock cmpxchg dword ptr [edx], ecx ; l = 1 (and eax = 0)
                                ; else
                                ; l = 1 (and eax = 1)
        mov val, eax           ; set eax to be the return val
    }
    return val;
}
          
```

6

## Software algorithms

- book, p. 193
- Example algorithm for two processes 0 and 1,  $P_i$  does this:
 

```
while(1) {
    while(turn != i) ;
    <critical section>
    turn = 1-i;
}
```
- What's wrong with it?

7

## Hyman's Algorithm

```
bool flag[2];
int turn;
void Protocol(int id) {
    while(true) {
        1 flag[id] = true;
        2 while(turn != id) {
        3     while(flag[1-id])
        4         ; /* Spin */
        5     turn = id;
        }
        6 <Critical Section>
        7 flag[id] = false;
        8 <rest of code>
    }
}
```

- Two processes: P0, P1
- flag initialized to {false, false}, turn to 0
- "Elegant"
- Wrong – why?**

8

## Hyman's Algorithm

```
bool flag[2];
int turn;
void Protocol(int id) {
    while(true) {
        1 flag[id] = true;
        2 while(turn != id) {
        3     while(flag[1-id])
        4         ; /* Spin */
        5     turn = id;
        }
        6 <Critical Section>
        7 flag[id] = false;
        8 <rest of code>
    }
}
```

Timeline:	
P1	P0
	1   flag[1]=true
2	2
3	3
4	4
5	5   1   flag[0]=true
6	6   2
7	7   6
8	8   5   turn=1
	9   6   mutual exclusion has been violated

- P1 executes 1, 2, 3
- P0 executes 1, 2, 6
- P1 executes 5, 6

9

## Semaphore review

- Semaphore = a special *variable*
  - Manipulated atomically via two operations:
    - P (wait)
    - V (signal)
  - To access critical section:
    - P(sem)
    - <critical section>
    - V(sem)
- Has a counter = number of available resources
- Has a queue of waiting threads
  - If execute wait() and semaphore is free, continue
  - If not, block on that waiting queue
- signal() unblocks a thread if it's waiting

10

## Synchronization in Project 2

- Part 2: write two synchronization primitives
- Implement mutex (binary semaphore)
  - How is it different from spinlock?
    - Need to keep track of lock state
    - Need to keep waiting threads on a queue
  - In **lock()**, may need to block current thread
    - Don't put on ready queue
    - Do run some other thread
  - For **unlock()**, need to take a thread off the waiting queue if available

11

## Condition Variable

- A "place" to let threads wait for a certain condition or event to occur while holding a lock (often a monitor lock).
- It has:
  - Wait queue
  - Three functions: *wait*, *signal*, and *broadcast*
    - wait* – sleep until condition becomes true.
    - signal* – event/condition has occurred. If wait queue nonempty, wake up *one* thread, o.w. *do nothing*
      - Do not run the woken up thread right away
      - FIFO determines who wakes up
    - broadcast* – just like *signal*, except wake up all threads (not just one).
- In part 2, you implement all of these

12

## Condition Variables 2

- More about `cond_wait(pthread_cond_t cond, pthread_mutex_t lock)`:
  - Called while holding `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
  - After woken up, a thread should reacquire its lock before continuing
- How are CVs different from semaphores?
- More info: `man pthread_cond_wait`
  - We follow the same spec for `wait`, `signal`, `bcast`

13

## Monitors: preview

- One thread inside at a time
- Lock + a bunch of condition variables (CVs)
- CVs used to allow other threads to access the monitor while one thread waits for an event to occur

14

## No preemption

- You get atomic critical sections for free
- However, you should understand what to do if you had preemption
  - Mark *critical sections* with comments
  - Describe appropriate protection that might apply (e.g. spinlock).

15

## Part 3 problem

- N cooks produce burgers & place on stack
- M students grab burgers and eat them
- Provide correct synchronization
  - Check with your threads and pthreads!
- Print out what happens!
- sample output (rough draft):
 

```
...
cook 2 produces burger #5
cook 2 produces burger #6
cook 3 produces burger #7
student 1 eats burger #7
student 2 eats burger #6
cook 1 produces burger #8
student 1 eats burger #8
student 1 eats burger #5
...
```

16

## Synchronization – Important Points

- Necessary when multiple threads have access to same data
- Can't use some primitives in interrupt handlers
  - Why? Which ones?
- Don't forget to release lock, semaphore, etc
  - Check all paths**
- Synchronization bugs can be very difficult to find
  - Read your code

17

## Homework questions?

- Sleeping barber problem
- Cigarette-smoker problem

18

## Sample synchronization problem

### Late-Night Pizza

- n A group of students study for cse451 exam
- n Can only study while eating pizza
- n Each student thread executes the following:
  - n while (1) {
    - n pick up a piece of pizza;
    - n study while eating the pizza;
- n If student finds pizza is gone, the student goes to sleep until another pizza arrives
- n First student to discover pizza is gone phones Pizza Hut and orders a new one.
- n Each pizza has S slides.

19

## Late-Night Pizza

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

20

## Semaphore solution

```

shared data:
semaphore pizza; (counting sema, init to 0, represent
                 number of available pizza resources)
semaphore deliver; (init to 1)
int num_slices = 0;
semaphore mutex; (init to 1) // guard updating of num_slices

Student {
while (diligent) {
P(pizza);
P(mutex);
num_slices--;
if (num_slices==0)
// took last slice
V(deliver);
V(mutex);
study();
}
}

DeliveryGuy {
while(employed) {
P(deliver);
P(mutex);
num_slices=S;
V(mutex);
for (int i=0,i<S,i++) {
V(pizza);
}
}
}

```

21

## Condition Variable Solution

```

int slices=0;
Condition order, deliver;
Lock mutex;
bool first = true;

Student() {
while(diligent) {
mutex.lock();
if( slices > 0 ) {
slices--;
}
else {
if(first) {
order.signal(mutex);
first = false;
}
deliver.wait(mutex);
}
mutex.unlock();
Study();
}
}

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

```

22