

Semaphores and Monitors: High-level Synchronization Constructs

Synchronization Constructs

- ◆ Synchronization
 - Coordinating execution of multiple threads that share data structures
- ◆ Past few lectures:
 - Locks: provide mutual exclusion
 - Condition variables: provide conditional synchronization
- ◆ Today: Historical perspective
 - Semaphores
 - ✦ Introduced by Dijkstra in 1960s
 - ✦ Main synchronization primitives in early operating systems
 - Monitors
 - ✦ Alternate high-level language constructs
 - ✦ Proposed by independently Hoare and Hansen in the 1970s

Semaphores

- ◆ Study these for history and compatibility
 - Don't use semaphores in new code
- ◆ A non-negative integer variable with two atomic and isolated operations

Semaphore → P() (*Passeren*; wait)
 If *sem* > 0, then decrement *sem* by 1
 Otherwise "wait" until *sem* > 0 and then decrement

Semaphore → V() (*Vrijgeven*; signal)
 Increment *sem* by 1
 Wake up a thread waiting in P()

- ◆ We assume that a semaphore is *fair*
 - No thread *t* that is blocked on a P() operation remains blocked if the V() operation on the semaphore is invoked infinitely often
 - In practice, FIFO is mostly used, transforming the set into a queue.

Key idea of Semaphores vs. Locks

- ◆ Locks: Mutual exclusion only (1-exclusion)
- ◆ Semaphores: k-exclusion
 - $k = 1$, equivalent to a lock
 - ✦ Sometimes called a mutex, or binary semaphore
 - $k = 2+$, up to k threads at a time
- ◆ Many semaphore implementations use "up" and "down", rather than Dutch names (P and V, respectively)
 - 'cause how many programmers speak Dutch?
- ◆ Semaphore starts at k
 - Acquire with down(), which decrements the count
 - ✦ Blocks if count is 0
 - Release with up(), which increments the count and never blocks

Important properties of Semaphores

- ◆ Semaphores are *non-negative* integers
- ◆ The *only* operations you can use to change the value of a semaphore are P()/down() and V()/up() (except for the initial setup)
 - P()/down() can block, but V()/up() never blocks
- ◆ Semaphores are used both for
 - Mutual exclusion, and
 - Conditional synchronization
- ◆ Two types of semaphores
 - Binary semaphores: Can either be 0 or 1
 - General/Counting semaphores: Can take any non-negative value
 - Binary semaphores are as expressive as general semaphores (given one can implement the other)

How many possible values can a binary semaphore take?

- ◆ How many possible values can a binary semaphore take?
 - A. 0
 - B. 1
 - C. 2
 - D. 3
 - E. 4

Using Semaphores for Mutual Exclusion

- Use a *binary semaphore* for mutual exclusion


```
Semaphore = new Semaphore(1);
```

```
Semaphore→P():  
Critical Section;  
Semaphore→V();
```
- Using Semaphores for producer-consumer with bounded buffer


```
int count;  
Semaphore mutex;  
Semaphore fullBuffers;  
Semaphore emptyBuffers;
```

Use a separate semaphore for each constraint

Coke Machine Example

- Coke machine as a shared buffer
- Two types of users
 - Producer: Restocks the coke machine
 - Consumer: Removes coke from the machine
- Requirements
 - Only a single person can access the machine at any time
 - If the machine is out of coke, wait until coke is restocked
 - If machine is full, wait for consumers to drink coke prior to restocking
- How will we implement this?
 - How many lock and condition variables do we need?
 - A. 1 B. 2 C. 3 D. 4 E. 5

Revisiting Coke Machine Example

```
Class CokeMachine{  
...  
int count;  
Semaphore new mutex(1);  
Semaphores new fullBuffers(0);  
Semaphores new emptyBuffers(numBuffers);  
}
```

```
CokeMachine::Deposit(){  
emptyBuffers→P();  
mutex→P();  
Add coke to the machine;  
count++;  
mutex→V();  
fullBuffers→V();  
}
```

```
CokeMachine::Remove(){  
fullBuffers→P();  
mutex→P();  
Remove coke from to the machine;  
count--;  
mutex→V();  
emptyBuffers→V();  
}
```

Does the order of P matter? Order of V matter?

Implementing Semaphores

```
Semaphore::P() {  
if (0 > atomic_dec(&value)) {  
Put TCB on wait queue for semaphore;  
Switch(); // dispatch a ready thread  
atomic_inc(&value);  
}  
}
```

Does this work?

```
Semaphore::V() {  
int notify = atomic_inc(&value);  
// atomic_inc returns new value  
if (notify <= 0) {  
Move a waiting thread to ready queue;  
}  
}
```

value:
1..k = Resource available
0 = All resources used, no waiters
<0 = -1 * number of waiters

Implementing Semaphores

```
Semaphore::P() {  
while (0 > atomic_dec(&value)) {  
Put TCB on wait queue for semaphore;  
Switch(); // dispatch a ready thread  
atomic_inc(&value);  
}  
}
```

```
Semaphore::V() {  
int notify = atomic_inc(&value);  
// atomic_inc returns new value  
if (notify <= 0) {  
Move a waiting thread to ready queue;  
}  
}
```

value:
1..k = Resource available
0 = All resources used, no waiters
<0 = -1 * number of waiters

The Problem with Semaphores

- Semaphores are used for dual purpose
 - Mutual exclusion
 - Conditional synchronization
- Difficult to read/develop code
- Waiting for condition is independent of mutual exclusion
 - Programmer needs to be clever about using semaphores

```
CokeMachine::Deposit(){  
emptyBuffers→P();  
mutex→P();  
Add coke to the machine;  
count++;  
mutex→V();  
fullBuffers→V();  
}
```

```
CokeMachine::Remove(){  
fullBuffers→P();  
mutex→P();  
Remove coke from to the machine;  
count--;  
mutex→V();  
emptyBuffers→V();  
}
```

Introducing Monitors	
<ul style="list-style-type: none"> ◆ Separate the concerns of mutual exclusion and conditional synchronization ◆ What is a monitor? <ul style="list-style-type: none"> ➢ One lock, and ➢ Zero or more condition variables for managing concurrent access to shared data ◆ General approach: <ul style="list-style-type: none"> ➢ Collect related shared data into an object/module ➢ Define methods for accessing the shared data ◆ Monitors first introduced as programming language construct <ul style="list-style-type: none"> ➢ Calling a method defined in the monitor automatically acquires the lock ➢ Examples: Mesa, Java (synchronized methods) ◆ Monitors also define a programming convention <ul style="list-style-type: none"> ➢ Can be used in any language (C, C++, ...) 	13

Critical Section: Monitors	
<ul style="list-style-type: none"> ◆ Basic idea: <ul style="list-style-type: none"> ➢ Restrict programming model ➢ Permit access to shared variables only within a critical section ◆ General program structure <ul style="list-style-type: none"> ➢ Entry section <ul style="list-style-type: none"> ◆ "Lock" before entering critical section ◆ Wait if already locked, or invariant doesn't hold ◆ Key point: synchronization may involve wait ➢ Critical section code ➢ Exit section <ul style="list-style-type: none"> ◆ "Unlock" when leaving the critical section ◆ Object-oriented programming style <ul style="list-style-type: none"> ➢ Associate a lock with each shared object ➢ Methods that access shared object are critical sections ➢ Acquire/release locks when entering/exiting a method that defines a critical section 	14

Remember Condition Variables	
<ul style="list-style-type: none"> ◆ Locks <ul style="list-style-type: none"> ➢ Provide mutual exclusion ➢ Support two methods <ul style="list-style-type: none"> ◆ Lock::Acquire() – wait until lock is free, then grab it ◆ Lock::Release() – release the lock, waking up a waiter, if any ◆ Condition variables <ul style="list-style-type: none"> ➢ Support conditional synchronization ➢ Three operations <ul style="list-style-type: none"> ◆ Wait(): Release lock; wait for the condition to become true; reacquire lock upon return (Java wait()) ◆ Signal(): Wake up a waiter, if any (Java notify()) ◆ Broadcast(): Wake up all the waiters (Java notifyAll()) ➢ Two semantics for implementation of wait() and signal() <ul style="list-style-type: none"> ◆ Hoare monitor semantics ◆ Hansen (Mesa) monitor semantics 	15

So what is the big idea?	
<ul style="list-style-type: none"> ◆ (Editorial) Integrate idea of condition variable with language <ul style="list-style-type: none"> ➢ Facilitate proof ➢ Avoid error-prone boiler-plate code 	16

Coke Machine – Example Monitor	
<pre> Class CokeMachine{ ... Lock lock; int count = 0; Condition notFull, notEmpty; } </pre>	<p>Does the order of acquire/while(){wait} matter?</p> <p>Order of release/signal matter?</p>
<pre> CokeMachine::Deposit(){ lock→acquire(); while (count == n) { notFull.wait(&lock); } Add coke to the machine; count++; notEmpty.signal(); lock→release(); } </pre>	<pre> CokeMachine::Remove(){ lock→acquire(); while (count == 0) { notEmpty.wait(&lock); } Remove coke from to the machine; count--; notFull.signal(); lock→release(); } </pre>
17	

Monitors: Recap	
<ul style="list-style-type: none"> ◆ Lock acquire and release: often incorporated into method definitions on object <ul style="list-style-type: none"> ➢ E.g., Java's synchronized methods ➢ Programmer may not have to explicitly acquire/release ◆ But, methods on a monitor object do execute under mutual exclusion ◆ Introduce idea of condition variable 	18

◆ Every monitor function should start with what?

- A. wait
- B. signal
- C. lock acquire
- D. lock release
- E. signalAll

19

Hoare Monitors: Semantics

- ◆ Hoare monitor semantics:
 - Assume thread *T1* is waiting on condition *x*
 - Assume thread *T2* is in the monitor
 - Assume thread *T2* calls *x.signal*
 - *T2* gives up monitor, *T2* blocks!
 - *T1* takes over monitor, runs
 - *T1* gives up monitor
 - *T2* takes over monitor, resumes
- ◆ Example


```

      T1                                T2
      fn1(...)
      ...
      x.wait // T1 blocks  → fnA(...)
      ...
      // T1 resumes ← x.signal // T2 blocks
      Lock→release();
      → T2 resumes
      
```

20

Hansen (Mesa) Monitors: Semantics

- ◆ Hansen monitor semantics:
 - Assume thread *T1* waiting on condition *x*
 - Assume thread *T2* is in the monitor
 - Assume thread *T2* calls *x.signal*; wake up *T1*
 - *T2* continues, finishes
 - When *T1* get a chance to run, *T1* takes over monitor, runs
 - *T1* finishes, gives up monitor
- ◆ Example:


```

      fn1(...)
      ...
      x.wait // T1 blocks  → fnA(...)
      ...
      // T1 resumes ← x.signal // T2 continues
      // T1 finishes // T2 finishes
      
```

21

Tradeoff

<h4>Hoare</h4> <ul style="list-style-type: none"> ◆ Claims: <ul style="list-style-type: none"> ➢ Cleaner, good for proofs ➢ When a condition variable is signaled, it does not change ➢ Used in most textbooks ◆ ...but <ul style="list-style-type: none"> ➢ Inefficient implementation ➢ Not modular - correctness depends on correct use and implementation of signal 	<h4>Hansen</h4> <ul style="list-style-type: none"> ◆ Signal is only a hint that the condition may be true <ul style="list-style-type: none"> ➢ Need to check condition again before proceeding ➢ Can lead to synchronization bugs ◆ Used by most systems (e.g., Java) ◆ Benefits: <ul style="list-style-type: none"> ➢ Efficient implementation ➢ Condition guaranteed to be true once you are out of while!
--	---

```

CokeMachine::Deposit(){
lock→acquire();
if (count == n){
    notFull.wait(&lock); }
Add coke to the machine;
count++;
notEmpty.signal();
lock→release();
}
      
```

```

CokeMachine::Deposit(){
lock→acquire();
while (count == n){
    notFull.wait(&lock); }
Add coke to the machine;
count++;
notEmpty.signal();
lock→release();
}
      
```

22

Problems with Monitors

Nested Monitor Calls

- ◆ What happens when one monitor calls into another?
 - What happens to `CokeMachine::lock` if thread sleeps in `CokeTruck::Unload`?
 - What happens if truck unloader wants a coke?

```

CokeMachine::Deposit(){
lock→acquire();
while (count == n){
    notFull.wait(&lock); }
truck→unload();
Add coke to the machine;
count++;
notEmpty.signal();
lock→release();
}
      
```

```

CokeTruck::Unload(){
lock→acquire();
while (soda.atDoor() != coke){
    cokeAvailable.wait(&lock); }
Unload soda closest to door;
soda.pop();
Signal availability for soda.atDoor();
lock→release();
}
      
```

23

More Monitor Headaches

The priority inversion problem

- ◆ Three processes (*P1*, *P2*, *P3*), and *P1* & *P3* communicate using a monitor *M*. *P3* is the highest priority process, followed by *P2* and *P1*.
- ◆ 1. *P1* enters *M*.
- ◆ 2. *P1* is preempted by *P2*.
- ◆ 3. *P2* is preempted by *P3*.
- ◆ 4. *P3* tries to enter the monitor, and waits for the lock.
- ◆ 5. *P2* runs again, preventing *P3* from running, subverting the priority system.
- ◆ A simple way to avoid this situation is to associate with each monitor the priority of the highest priority process which ever enters that monitor.

24

Comparing Semaphores and Monitors	
<pre>CokeMachine::Deposit(){ emptyBuffers→P(); mutex→P(); Add coke to the machine; count++; mutex→V(); fullBuffers→V(); }</pre>	<pre>CokeMachine::Deposit(){ lock→acquire(); while (count == n) { notFull.wait(&lock); } Add coke to the machine; count++; notEmpty.notify(); lock→release(); }</pre>
<pre>CokeMachine::Remove(){ fullBuffers→P(); mutex→P(); Remove coke from to the machine; count--; mutex→V(); emptyBuffers→V(); }</pre>	<pre>CokeMachine::Remove(){ lock→acquire(); while (count == 0) { notEmpty.wait(&lock); } Remove coke from to the machine; count--; notFull.notify(); lock→release(); }</pre>
<p>Which is better? A. Semaphore B. Monitors</p>	

Other Interesting Topics
<ul style="list-style-type: none"> ◆ Exception handling <ul style="list-style-type: none"> ➢ What if a process waiting in a monitor needs to time out? ◆ Naked notify <ul style="list-style-type: none"> ➢ How do we synchronize with I/O devices that do not grab monitor locks, but can notify condition variables. ◆ Butler Lampson and David Redell, “Experience with Processes and Monitors in Mesa.”

Summary
<ul style="list-style-type: none"> ◆ Synchronization <ul style="list-style-type: none"> ➢ Coordinating execution of multiple threads that share data structures ◆ Past lectures: <ul style="list-style-type: none"> ➢ Locks → provide mutual exclusion ➢ Condition variables → provide conditional synchronization ◆ Today: <ul style="list-style-type: none"> ➢ Semaphores <ul style="list-style-type: none"> ◆ Introduced by Dijkstra in 1960s ◆ Two types: binary semaphores and counting semaphores ◆ Supports both mutual exclusion and conditional synchronization ➢ Monitors <ul style="list-style-type: none"> ◆ Separate mutual exclusion and conditional synchronization