


Chapter 7: Process Synchronization

王振傑 (Chen-Chieh Wang)
ccwang@mail.ee.ncku.edu.tw

System Programming, Spring 2010

Producer-Consumer Problem

- ✦ Concurrent access to shared data may result in **data inconsistency**
- ✦ Maintaining **data consistency** requires mechanisms to ensure the orderly execution of cooperating processes
- ✦ Suppose that we wanted to provide a solution to the producer-consumer problem that fills **all** the buffers. We can do so by having an integer **count** that keeps track of the number of full buffers. Initially, count is set to 0. It is incremented by the producer after it produces a new buffer and is decremented by the consumer after it consumes a buffer.



Producer

```
while (true) {  
  
    /* produce an item and put in nextProduced */  
    while (count == BUFFER_SIZE)  
        ; // do nothing  
  
    buffer [in] = nextProduced;  
    in = (in + 1) % BUFFER_SIZE;  
    count++;  
  
}
```

3

System Programming, Spring 2010

Consumer

```
while (true) {  
    while (count == 0)  
        ; // do nothing  
  
    nextConsumed = buffer[out];  
    out = (out + 1) % BUFFER_SIZE;  
    count--;  
    /* consume the item in nextConsumed  
  
}
```

4

System Programming, Spring 2010

Race Condition

- ✦ `count++` could be implemented as

```
register1 = count
register1 = register1 + 1
count = register1
```

- ✦ `count--` could be implemented as

```
register2 = count
register2 = register2 - 1
count = register2
```

- ✦ Consider this execution interleaving with “count = 5” initially:

```
S0: producer execute register1 = count {register1 = 5}
S1: producer execute register1 = register1 + 1 {register1 = 6}
S2: consumer execute register2 = count {register2 = 5}
S3: consumer execute register2 = register2 - 1 {register2 = 4}
S4: producer execute count = register1 {count = 6}
S5: consumer execute count = register2 {count = 4}
```

5

Definitions

- ✦ **Synchronization**: using atomic operations to ensure cooperation between threads
 - For now, only loads and stores are atomic
 - We are going to show that its hard to build anything useful with only reads and writes
- ✦ **Mutual Exclusion**: ensuring that only one thread does a particular thing at a time
 - One thread *excludes* the other while doing its task
- ✦ **Critical Section**: piece of code that only one thread can execute at once. Only one thread at a time will get into this section of code.
 - Critical section is the result of mutual exclusion
 - Critical section and mutual exclusion are two ways of describing the same thing.

6

Solution to Critical-Section Problem

1. **Mutual Exclusion** - If process P_i is executing in its critical section, then no other processes can be executing in their critical sections
2. **Progress** - If no process is executing in its critical section and there exist some processes that wish to enter their critical section, then the selection of the processes that will enter the critical section next cannot be postponed indefinitely
3. **Bounded Waiting** - A bound must exist on the number of times that other processes are allowed to enter their critical sections after a process has made a request to enter its critical section and before that request is granted
 - Assume that each process executes at a nonzero speed
 - No assumption concerning relative speed of the N processes

7

System Programming, Spring 2010

Where are we going with synchronization?

Programs	Shared Programs
Higher-level API	Locks Semaphores Monitors Send/Receive
Hardware	Load/Store Disable Ints Test&Set Comp&Swap

- ✦ We are going to implement various higher-level synchronization primitives using atomic operations
 - Everything is pretty painful if only atomic primitives are load and store
 - Need to provide primitives useful at user-level

8

System Programming, Spring 2010