



UNIT II

PROCESS SYNCHRONIZATION - THE CRITICAL-SECTION PROBLEM -SYNCHRONIZATION HARDWARE





sns
INSTITUTIONS

Background

- 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 consumer-producer 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.



sns
INSTITUTIONS

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++;  
}
```



sns
INSTITUTIONS

Consumer

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



sns
INSTITUTIONS

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



sns
INSTITUTIONS

Race Condition

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

S0: producer execute $\text{register1} = \text{count}$ {register1 = 5}

S1: producer execute $\text{register1} = \text{register1} + 1$ {register1 = 6}

S2: consumer execute $\text{register2} = \text{count}$ {register2 = 5}

S3: consumer execute $\text{register2} = \text{register2} - 1$ {register2 = 4}

S4: producer execute $\text{count} = \text{register1}$ {count = 6}

S5: consumer execute $\text{count} = \text{register2}$ {count = 4}

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

Peterson's Solution

- Two process solution
- Assume that the LOAD and STORE instructions are atomic; that is, cannot be interrupted.
- The two processes share two variables:
 - int **turn**;
 - Boolean **flag[2]**
- The variable **turn** indicates whose turn it is to enter the critical section.
- The **flag** array is used to indicate if a process is ready to enter the critical section.
flag[i] = true implies that process P_i is ready!



sns
INSTITUTIONS

Algorithm for Process P_i

do {

flag[i] = TRUE;

turn = j;

while (flag[j] && turn == j);

critical section

flag[i] = FALSE;

remainder section

} while (TRUE);



sns
INSTITUTIONS

Synchronization Hardware

- Many systems provide hardware support for critical section code
- **Uniprocessors** – could disable interrupts
 - Currently running code would execute without preemption
 - Generally too inefficient on multiprocessor systems
 - Operating systems using this not broadly scalable
- Modern machines provide special atomic hardware instructions
 - **Atomic = non-interruptable**
 - Either test memory word and set value or swap contents of two memory words



sns
INSTITUTIONS

Solution to Critical-section Problem Using Locks

```
do {  
    acquire lock  
    critical section  
    release lock  
    remainder section  
} while (TRUE);
```



sns
INSTITUTIONS

TestAndSet Instruction

- Definition:

```
boolean TestAndSet (boolean *target)
{
    boolean rv = *target;
    *target = TRUE;
    return rv;
}
```



sns
INSTITUTIONS

Solution using TestAndSet

- Shared boolean variable lock, initialized to false.

- **Solution:**

```
do {  
    while ( TestAndSet (&lock ))  
        ; // do nothing  
        // critical section  
    lock = FALSE;  
        // remainder section  
} while (TRUE);
```



sns
INSTITUTIONS

Swap Instruction

- Definition:

```
void Swap (boolean *a, boolean *b)
{
    boolean temp = *a;
    *a = *b;
    *b = temp;
}
```




sns
INSTITUTIONS

Solution using Swap

- Shared Boolean variable lock initialized to FALSE; Each process has a local Boolean variable key

- **Solution:**

```
do {
```

```
    key = TRUE;
```

```
    while ( key == TRUE)
```

```
        Swap (&lock, &key );
```

```
        //  critical section
```

```
    lock = FALSE;
```

```
        //  remainder section
```

```
    } while (TRUE);
```



sns
INSTITUTIONS

Bounded-waiting Mutual Exclusion with TestAndSet()

```
do {  
    waiting[i] = TRUE;  
    key = TRUE;  
    while (waiting[i] && key)  
        key = TestAndSet(&lock);  
    waiting[i] = FALSE;  
    // critical section  
    j = (i + 1) % n;  
    while ((j != i) && !waiting[j])  
        j = (j + 1) % n;  
    if (j == i)  
        lock = FALSE;  
    else  
        waiting[j] = FALSE;  
    // remainder section  
} while (TRUE);
```