

OPERATING SYSTEMS

by
Marwa Yusuf

Lecture 7
Sunday 15-11-2020

Chapter 2 (2.3.4 to 2.3.5)

Processes and Threads

Sleep and Wakeup

- Both solutions require **busy waiting**.
 - Waste CPU time.
 - **Priority inversion problem**: lower priority process in its critical region.
- Use **sleep** and **wakeup** instead.

Producer-Consumer (Bounded-Buffer)

- Shared buffer between 2 processes (for simplification).
- Producer *sleeps* waiting for an empty slot (awakened by Consumer), Consumer *sleeps* waiting for a non-empty slot (awakened by producer).
- Leads to race condition.

Producer-Consumer (cont.)

```
#define N 100
int count = 0;
```

```
void producer(void)
{
    int item;
```

```
    while (TRUE) {
        item = produce_item( );
        if (count == N) sleep( );
        insert_item(item);
        count = count + 1;
        if (count == 1) wakeup(consumer);
    }
```

```
}
```

```
void consumer(void)
```

```
{
```

```
    int item;

    while (TRUE) {
        if (count == 0) sleep( );
        item = remove_item( );
        count = count - 1;
        if (count == N - 1) wakeup(producer);
        consume_item(item);
    }
```

```
}
```

```
/* number of slots in the buffer */
/* number of items in the buffer */
```

```
/* repeat forever */
/* generate next item */
/* if buffer is full, go to sleep */
/* put item in buffer */
/* increment count of items in buffer */
/* was buffer empty? */
```

```
/* repeat forever */
/* if buffer is empty, got to sleep */
/* take item out of buffer */
/* decrement count of items in buffer */
/* was buffer full? */
/* print item */
```

Producer-Consumer (cont.)

- Producer

- Consumer

read count

Insert item

count++

wakeup (consumer)

count == 0? → sleep

...

sleep

Producer-Consumer (cont.)

- Producer

- Consumer

read count

Insert item

count++

Is consumer awake? → **wakeup**

waiting bit = 1

wakeup (consumer)

count == 0? → **wakeup**
waiting bit == 1? don't sleep

...

sleep

Producer-Consumer (cont.)

- What about more processes? 16, 32 or more? Shall we keep adding more **wakeup waiting bits**?

Semaphores (by Dijkstra)

- **Semaphore** variables to count the number of wakeups.
- **Up** and **down** (generalizations of wakeup and sleep)
 - **Down** checks semaphore value, if > 0 , decrement and continue. Otherwise, sleep.
 - Checking, changing and sleep are done as an atomic operation.
 - **Up** increments semaphore, if some process is sleeping waiting for it, it is waked up and it completed its **down** operation.
 - Incrementing and waking up also is atomic.
 - In the original paper **P** for **down** and **V** for **up**.

Semaphores (by Dijkstra)

```
#define N 100
typedef int semaphore;
semaphore mutex = 1;
semaphore empty = N;
semaphore full = 0;

void producer(void)
{
    int item;

    while (TRUE) {
        item = produce_item();
        down(&empty);
        down(&mutex);
        insert_item(item);
        up(&mutex);
        up(&full);
    }
}
```

```
void consumer(void)
{
    int item;

    while (TRUE) {
        down(&full);
        down(&mutex);
        item = remove_item();
        up(&mutex);
        up(&empty);
        consume_item(item);
    }
}
```

```
/* number of slots in the buffer */
/* semaphores are a special kind of int */
/* controls access to critical region */
/* counts empty buffer slots */
/* counts full buffer slots */
```

```
/* TRUE is the constant 1 */
/* generate something to put in buffer */
/* decrement empty count */
/* enter critical region */
/* put new item in buffer */
/* leave critical region */
/* increment count of full slots */
```

```
/* infinite loop */
/* decrement full count */
/* enter critical region */
/* take item from buffer */
/* leave critical region */
/* increment count of empty slots */
/* do something with the item */
```

Semaphores (by Dijkstra)

- **Binary Semaphores.**
- To make **up** and **down** operations atomic:
 - Make them as system calls where OS disables interrupts during them.
 - In multiple CPUs case, protect them with a lock variable using TSL or XCHG.
- Note that as the operations are short (few instructions) there is no problem in both solutions.
- Semaphores were used for **synchronization** (full and empty) and for **mutual exclusion** (mutex).