# SOLUTIONS TO THE SECOND 3360/6310 QUIZ

Jehan-François Pâris Spring 2016

In some implementations of monitors, the signal operation is replaced by a operation.

In some implementations of monitors, the signal operation is replaced by a <u>notify</u> operation.

In Java, a monitor procedure must start with the two keywords public and

In Java, a monitor procedure must start with the two keywords public and synchronized

You should use streams rather than datagrams when you have to transmit \_\_\_\_\_\_ amounts of data.

You should use streams rather than datagrams when you have to transmit <u>large/huge</u> amounts of data.

The main disadvantage of the round robin scheduling policy is that it generates too many at heavy loads.

The main disadvantage of the round robin scheduling policy is that it generates too many context switches at heavy loads.

The main disadvantage of atomic transactions is that they are \_\_\_\_\_

The main disadvantage of atomic transactions is that they are <u>very costly to implement</u>

A signal operation has no effect if there are no other monitor procedures

- A signal operation has no effect if there are no other monitor procedures
   waiting on the signaled condition
- A signal operation has no effect if there are no other monitor procedures waiting for that signal

Unlike virtual circuits, streams do not preserve

Unlike virtual circuits, streams do not preserve message boundaries

 Priority-based scheduling algorithms typically assign \_\_\_\_\_\_ priorities to I/O-bound processes than to CPU-bound processes.

 Priority-based scheduling algorithms typically assign <u>higher</u> priorities to I/O-bound processes than to CPU-bound processes.

guarantee that each request will either be fully executed or have no effect.

- <u>Atomic transactions</u> guarantee that each request will either be fully executed or have no effect.
- <u>All-nothings semantics</u> guarantee that each request will either be fully executed or have no effect.
- Transactions semantics guarantee that each request will either be fully executed or have no effect.

# Second question

- Consider the function:
  - double\_decrement(int \*pa, int \*pb){
     (\*pa) -= 1; (\*pb) -= 1;
    }//double\_decrement
  - and assume the following calling sequence:

double\_decrement(&alpha, &alpha)

# Second question

- What will be the value of alpha after the call assuming that the call was:
  - A regular procedure call?
  - A remote procedure call?

# Second question

What will be the value of alpha after the call assuming that the call was:

• A regular procedure call?

alpha is decremented twice

• A remote procedure call?

alpha get assigned twice 3 -1 = 2

#### Alternate second question

- Consider the function:
  - double\_decrement(int \*pa, int \*pb){
     (\*pa) -= 1; (\*pb) -= 1;
    }//double\_decrement
  - and assume the following calling sequence:
    - int alpha = 5; // 5 not 3 double\_decrement(&alpha, &alpha)

#### Alternate second question

- What will be the value of alpha after the call assuming that the call was:
  - A regular procedure call?
  - A remote procedure call?

## Alternate second question

- What will be the value of alpha after the call assuming that the call was:
  - A regular procedure call?

- alpha is decremented twice
- A remote procedure call?

alpha get assigned twice 5 -1 = 4

Explain why *lottery scheduling* is inherently *starvation free*?

- Explain why *lottery scheduling* is inherently *starvation free*?
  - New tickets will not be issued until all processes that have thickets have used theirs

What should a process do when it receives a retransmission of a message it has previously received and acknowledged ?

- What should a process do when it receives a retransmission of a message it has previously received and acknowledged ?
  - It should acknowledge the retransmission (because otherwise the message will be resent a second time).

What is the difference between a *blocking receive* and a *non-blocking receive*?

- What is the difference between a blocking receive and a non-blocking receive?
  - A blocking receive waits until the process receives a message while a nonblocking receive never waits.

When should we worry about *little-endians* and *big-endians*?

- When should we worry about *little-endians* and *big-endians*?
  - We should worry about the byte ordering inside words each time we exchange data between two different machines.

What is the main disadvantage of spin locks?

- What is the main disadvantage of *spin locks*?
  - □ <u>Spin locks</u>
    - Waste processor cycles
    - Generate additional context switches unless the waiting thread runs on a different core than the thread inside the critical section.

Why do most experienced programmers put their signal operations at the end of their monitor procedures?

- Why do most experienced programmers put their signal operations at the end of their monitor procedures?
  - Because a thread performing a signal will lose its control of the monitor if a tread waiting on the condition catches the signal.

Explain why *non-preemptive* schedulers are *bad*.

- Explain why *non-preemptive* schedulers are *bad*.
  - Non-preemptive scheduling policies do not prevent CPU-bound processes from monopolizing the CPU.

What is the difference between a blocking send and a non-blocking send?

- What is the difference between a blocking send and a non-blocking send?
  - A blocking send waits until the message it sends had been delivered while a nonblocking send returns as soon as the message has been accepted for delivery.

When should we worry about the byte ordering inside computer words?

- When should we worry about the byte ordering inside computer words?
  - We should worry about the byte ordering inside words each time we exchange data between two different machines.

- When should we use spin locks instead of semaphore calls?
  - Spin locks are better for short waits whenever the waiting process and the process that holds the lock are on different core/on a multi-core computer

#### Fourth question

Consider the following system V scheduler:						
#ts_q	ts_tq	ts_slp…	ts_max	ts_lw	LE\	/EL
800	U	1	4000	1	#	0
400	V	2	2000	2	#	1
200	1	W	1000	Υ	#	2
100	2	Χ	500	Ζ	#	3

• Give correct values for the missing parameters:

#### Fourth question

Consider the following system V scheduler: #ts\_q... ts\_tq... ts\_slp... ts\_max... ts\_lw... LEVEL 1 2 W U 1 800 4000 # 0 V 2000 2 # 1 400 1 Y # 2 200 1000 2 X Z # 3 500 100

Give correct values for the missing parameters: U = V = 0 W = X = 3 Y = Z = 3

# Fifth question

- An interstate bus that can carry 80 passengers has a single door that lets passengers get in or out one by one at scheduled stops. Add semaphores to the following two functions to ensure that
  - The bus will never have more than 80 passengers and
  - Passengers will not collide with each other when getting through the bus door.

## Fifth question

- semaphore notfull = 80 ; semaphore door = 1 ;
- get\_in(){
   <u>P(&freeseats); // in that order</u>
   <u>P(&door); // in that order</u>
   go\_through\_door();
   <u>V(&door);</u>

} //get in

## Fifth question

get\_out(){

P(&door); go\_through\_door(); V(&door); // in any order V(&freeseats); // in any order } //get\_out

# Alternate fifth question

- An interstate bus that can carry 30 passengers has a single door that lets passengers get in or out one by one at scheduled stops. Add semaphores to the following two functions to ensure that
  - The bus will never have more than 30 passengers and
  - Passengers will not collide with each other when getting through the bus door.

# Alternate fifth question

- semaphore notfull = 80 ; semaphore door = 1 ;
- get\_in(){
   <u>P(&freeseats); // in that order</u>
   <u>P(&door); // in that order</u>
   go\_through\_door();
   <u>V(&door);</u>

}\_//get\_in

# Alternate fifth question

get\_out(){

P(&door); go\_through\_door(); V(&door); // in any order V(&freeseats); // in any order

} //get\_out