

Quiz 4

Course: cs241 - System Programming, CS Department

Date: February 24, 2006

Netid:

Name:

UIN:

Note: Completion of quiz is an individual effort. The quiz takes 10 minutes. The student gets additional 5 points for taking the quiz. *Each question has ONLY ONE ANSWER!!!*

1. (1 Point) What is a signal?
 - a. convoy situation caused by $N-1$ I/O processes and 1 CPU-bound process
 - b. context switch between processes
 - c. software notification of an event to a process

2. (1 Point) What does a process signal mask contain?
 - a. Set of UNBLOCKED signals
 - b. Set of currently BLOCKED signals
 - c. Set of IGNORED signals

3. (2 Points) Consider the following piece of code with `sigset_t intmask` signal mask that includes signal `SIGINT (CTRL C)`, and with variables `int i, int y`.

```
sigprocmask(SIG_BLOCK, &intmask, NULL);
for (i=0; i<100; i++)
    y=i++;
sigprocmask(SIG_UNBLOCK, &intmask, NULL);
```

- What happens in this code?
- a. `SIG_INT` signal will be added to the signal mask `intmask` in the middle of the calculation of `y`
 - b. Calculation of `y` can be interrupted during its execution with signal `SIG_INT`
 - c. Calculation of `y` can not be interrupted during its execution with signal `SIG_INT`
4. (1 Point) The `sigaction` function allows the caller to:
 - a. examine or specify the action association with a specific signal
 - b. delete a signal from signal mask
 - c. change a signal from BLOCKED to UNBLOCKED

 5. (1 Point) A signal handler is useful for
 - a. Catching a SIGKILL signal
 - b. Changing a signal from BLOCKED to UNBLOCKED
 - c. Specifying an action to be taken for a specific signal

6. (1 Point) Let us assume two events (jobs) arriving each with the constant arrival rate of 10 processes per second. Let us assume one server (processor) with the constant service rate of 100 processes per second. What is the server utilization under these assumptions?
- 20%
 - 100%
 - 5%
7. (1 Point) A real-time scheduling policy, called the Earliest Deadline First, aims to schedule processes so that they meet their deadlines. To achieve this goal, the system can accept
- Unlimited number of processes
 - Only I/O bound processes
 - Only processes that pass schedulability condition
8. (1 Points) Let us assume the following snapshot of the system with the arrival order of the processes in the ready queue as follows: P1 first, P2 second, and P3 third.

Process ID	Burst Time	Arrival Time
P1	10	0
P2	5	0
P3	15	0

- and let us assume for the Round-Robin Scheduling policy that the time quantum is 5 time units.
- Under these assumptions, consider the following claim: **The process P2 finishes earlier when scheduled with Non-Preemptive Shortest Job First scheduling policy than with Round-Robin scheduling policy.** The claim is
- True
 - False
9. (1 Point) Let us consider a pre-emptive Shortest Job First scheduling where process A arrives at time 0 and needs to run for 1 hour. From the start time 0, other (short) processes will arrive every 1 minute and run for 2 minutes each. This situation will cause
- Deadlock for all processes
 - Starvation for process A
 - Starvation for the short processes
 - None of the above