

**Due Feb. 07 at 11:59pm Submit your assignment using CMS**

1. Which of the following should NOT be allowed in user mode? Briefly explain.

- a) Disable all interrupts.
- b) Read the time-of-day clock
- c) Set the time-of-day clock
- d) Perform a trap
- e) TSL (test-and-set instruction used for synchronization)

Which of the following components of a program state are shared across threads in a multi-threaded process ? Briefly explain.

- a) Register Values
- b) Heap memory
- c) Global variables
- d) Stack memory

2. After an interrupt occurs, hardware needs to save its current state (content of registers etc.) before starting the interrupt service routine. One issue is where to save this information. Here are two options:

- a) Put them in some special purpose internal registers which are exclusively used by interrupt service routine.
- b) Put them on the stack of the interrupted process.

Briefly discuss the problems with above two options.

3. You are asked to implement a online computer chess game server. The idea is that people can use the client-side GUI to connect to, and play against the game server. The server should be designed to support a large number of concurrent players. Suppose you can only use blocking I/O to read from network, would you choose user-level threads or kernel-level threads for the implementation? Briefly explain.

4. Consider a system running ten I/O-bound tasks and one CPU-bound task. Assume that the I/O-bound tasks issue an I/O operation once for every millisecond of CPU computing and that each I/O operation takes 10 milliseconds to complete. Also assume that the context-switching overhead is 0.1 millisecond and all processes are long-running tasks. What is the CPU utilization for a round-robin scheduler when:

- a) The time quantum is 1 millisecond
- b) The time quantum is 10 millisecond

5. Consider a system implementing multilevel queue scheduling. What strategy can a computer user employ to maximize the amount of CPU time allocated to the user's process? What if the system implements round-robin?

6. You are asked to consult on the design of a simple operating system called Linux. Linux is designed to support multiprogramming on a single-processor machine. However, the designer have observed that under various conditions the machine hangs because some of the concurrently executing processes either stop making forward progress, or make progress so slowly they might as well have stopped. Your job is to debug their kernel.

The Linux kernel job queue consists of an array of process control blocks with the following structure:

```
struct PCB {
    int process_id;
    enum { running, ready, waiting } status;
    int priority;
    struct PCB *next_ready;
    /* "next ready" points to the next process on the ready */
    /* queue, or is NULL if this is the last process */
    void *PC;
    /* program counter at which to restart execution */
    ...more stuff...
}
```

The scheduler also has a variable *ready\_queue* of type *struct PCB\** that is a pointer to the first process on the ready queue. The ready queue is a linked list formed by the next ready pointers, in which the status of every process is ready. The linked list is kept sorted in order of descending process priority.

a) To impress the Linux designers with your fitness for the consulting job, name two additional kinds of information you would expect to find in the section of the PCB labeled “more stuff.”

b) The Linux kernel implements preemptive priority scheduling, yet the system designers are finding that often processes are not preempted. They have determined that the problem is in the scheduler itself. The scheduler selects and dequeues a new process to run using the following code:

```
struct PCB *choose_process() {
    struct PCB *ret = ready_queue;
    ready_queue = ret->next_ready;
    return ret;
}
```

Because the scheduler is preemptive, sometimes running processes must be placed back on the ready queue when preempted; this is implemented by the following function which ensures the ready queue is kept sorted on priority:

```
struct PCB *reenqueue_running(PCB *current_proc) {
    current_proc->status = ready;
    struct PCB *p = ready_queue;
    struct PCB **prev = &ready_queue;
    while (p != NULL && p->priority > current_proc->priority) {
        prev = &p->next_ready;
        p = p->next_ready;
    }
    *prev = current_proc;
    current_proc->next_ready = p; // insert into linked list
}
```

Why isn't preemption working? Suggest a one-line fix.

c) Even after you fix preemptive scheduling, it is discovered that some processes are not scheduled for execution for long periods of time. When can this happen? What “term” is commonly used to describe this phenomenon?

d) Linux provides shared-memory segments to which multiple processes can read and write. To allow control of concurrent access to these segments, it provides mutexes with the following operations:

```
void lock(Mutex m); /* wait until mutex is available and then acquire it */
void unlock(Mutex m); /* release a mutex that is held */
```

The Linux printer server is implemented using shared memory; processes submit a print request by locking a mutex  $m$ , putting the request in a print-queue data structure, and releasing the mutex. This involves no system calls and all memory accessed is physically resident – there are no page faults. Note that a mutex is like a lock and is only released through an explicit call to unlock.

When high-priority and low-priority processes attempt to share access to the printer, in the presence of other medium-priority processes, it is found that the high-priority processes often block indefinitely on the initial call to  $\text{lock}(m)$ , while medium-priority processes continue to be scheduled and run. What is happening, and how might the scheduler be modified to address the problem?