

This exam is **closed book**. You can have **one** page of notes. UH expels cheaters

1. For each of the statements below, indicate in one sentence whether the statement is true or false (2 points), **and why** (3 points).

- (a) Most of the UNIX kernel is written in *assembly language*.

False, most of the UNIX kernel is written in C.

- (b) *Kernel-supported threads* are easier to use than *user-level threads*.

True, as programmers can use blocking system calls without fearing to block all the other threads sharing the same address space.

- (c) *Big monolithic kernels* are faster than *microkernels*.

True, as system calls will only require two context switches instead of four.

- (d) In a *multiprogramming system*, there can be *many programs* in the system but only *one process*.

False, there can be many processes. Systems that can only have one process at a time called uniprogramming systems.

- (e) In the *round-robin* scheduling policy, the priority of a process is function of its past CPU usage.

False, all processes have the same priority.

- (f) Multiprogramming generally results in a better utilization of the system resources than uniprogramming.

True, as they use much more efficiently the CPU.

2. *Advantage and disadvantages: you will get no credit if you answer mentions a disadvantage when an advantage is asked and vice versa.* (4×5 points)

- (a) What is the major advantage of *time-sharing* systems over *batch* systems?

They provide interactive access to the computer.

- (b) What is the major advantage of *DMA controllers*?

They move data between the main memory and the disk without any CPU intervention, thus speeding up data transfers.

- (c) What is the major disadvantage of *non-preemptive scheduling policies*?

The scheduler cannot preempt processes that monopolize the CPU while other processes are waiting.

- (d) What is the major disadvantage of *not having privileged instructions*?

User processes can directly access disk units, thus preventing any kind of access control.

3. Complete the following fragment of code to ensure that the standard error of the process is redirected to the pipe `thispipe`. (2×5 points)

```
int thispipe[2];
pipe(thispipe);
_close(2);_____
_dup(thispipe[1]);__
close(thispipe[0]);close(thispipe[1]);
```

4. Consider the System V scheduler defined by the following table

#ts_quantum	ts_tqexp	ts_slpret	ts_maxwait	ts_lwait	LEVEL
800	0	1	3000	1	# 0
400	0	2	2000	2	# 1
200	1	3	1000	3	# 2
100	2	3	500	3	# 3

and the four processes P, Q, R and S with the following priorities:

p(P) = 0, p(Q) = 1, p(R) = 2, p(S) = 3

What will be the *new priorities* of these four processes after the four following events? (4×5 points)

- (a) Process P waits in the ready queue for 4000 ms.
- (b) Process Q is preempted after having used the CPU for 200 ms.
- (c) Process R returns to the ready queue after completion of a system call.
- (d) Process S releases the CPU after 90 ms.

p(P) = __1__, p(Q) = __1__, p(R) = __3__, p(S) = __3__

5. Why is `fork()` a *very expensive system call*? (10 points)

Because it must create a new address space and copy the whole contents of the parent address space into this new address space.

6. Describe a bad thing that will happen when a time-sharing system does not have enough main memory. (10 points)

The kernel will start swapping out ready processes, which it will have to bring back into main memory when it will be their turn to run. This will create a lot of overhead.