COSC 4330 FIRST MIDTERM SEPTEMBER 24, 2003

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

1. What will the following program print out? $(2 \times 5 \text{ points})$

```
main() {
    printf("Hello!\n");
    fork();
    fork();
    printf("Goodbye!\n");
}
```

Answer: The program will print out _one_ "Hello!" and _four_ "Goodbye!"

- 2. For each of the statements below, indicate in one sentence whether the statement is true or false (2 points), and why (3 points).
 - (a) It makes no sense to have memory protection on a single user system.

False, without memory protection, any misbehaving process could crash your system.

(b) A process waiting on a system call completion is in the *suspended state*.

False, a process waiting on a system call completion is in the waiting state.

(c) Most modern operating systems have a *layered kernel*.

False, its is too difficult to partition a kernel into meaningful layers that can be stacked on the top of each other.

(d) Processes often interrupt themselves.

True, they do it every time they do a system call.

Consider a process waiting for the CPU for 20 ms, getting the CPU for 10 ms, doing a system call and getting again the CPU for 20 ms. List in order the five states the process will have visited. (5×4 points) (Hint: The same state(s) may be visited several times.)

Answer: The five states are _ready_ then _running_ then _waiting__ then _ready_ and finally _running_.

4. How does Amoeba manage its threads? (5 points) What is the main advantage of this solution? (5 points)

Amoeba threads are managed at user-level but the kernel is aware of them: when a thread executes a blocking system call, the kernel returns control to the <u>thread</u> <u>scheduler</u> of the task; this thread scheduler can either schedule another thread or return control to the OS.

The main advantage of the solution is that threads can do blocking system calls without blocking all threads in the same address space.

- **5.** Advantage and disadvantages: You will get no credit if you answer mentions a disadvantage when an advantage is asked and vice versa. (6×5 points)
 - (a) What is the major advantage of *monolithic kernels* over *microkernels*?

They are faster because they require two context switches per system call instead of four.

(b) What is the major disadvantage of *delaying writes*?

Data can be lost if the system crashes before the data are actually written to disk.

(c) What is the major advantage of *dual-mode CPUs*?

They prevent user processes from accessing directly disk drives and other peripherals, forcing them to go through the kernel.

(d) What is the major disadvantage of *the master-slave* organization for multiprocessor operating systems?

They require are kernel tasks to be performed on a single processor that can often become a bottleneck.

(e) What is the major advantage of *timer interrupts*?

They can be used to prevent one process from monopolizing the CPU when other processes are waiting.

(f) What is the major advantage of *vectorized interrupts*?

They ensure that more urgent interrupts will be handled before less urgent ones.

6. What is the main difference between *hard real-time* and *soft real-time* applications? (4 points) Give examples of both. (2×3 points)

In a hard real-time application, missed deadlines can have catastrophic effects while they will not in soft real-time applications.

Industrial process control and fly-by-wire systems are examples of hard real-time applications. DVD players are soft real-time applications: a missed deadline will only result in a few bad frames on the screen.