

**Wednesday 9 December 1998**

---

**1. Compare interrupts to subroutines. (12 points)**

**a. In what ways are they similar?**

In both cases, control transfers from one point in the program to another piece of code. When that code finishes, control transfers back to the instruction following the one it was executing before.

Also, in both cases, some mechanism (either hardware or software) must save and restore the PC, SP, registers, etc.

**b. In what ways are they different?**

Subroutines are called at a known point in the program — at locations of the programmer's choosing. Interrupts may occur at any time, usually due to some piece of hardware (e.g., timer, I/O device).

**2. Compare microprogrammed and hardwired control. (12 points)**

**a. Which is faster, and why?**

Hardwired control is faster, because the hardware is simpler and there is no extra level of instruction decoding and execution.

**b. Which is more flexible, and why?**

Microprogrammed control is more flexible, because it can be changed by changing the microprogram stored in ROM. (Note that a hardwired controller can also be changed, but changing the hardware on the chip is much more difficult, as it requires designing a new chip.)

Also, microprogrammed control is more flexible because it allows the emulation of functionality not available in the hardware (for example, multiplication on a CPU without a multiplier).

3. Given the figure of the Chapter 9 simple machine shown at the right, explain how it is used to do the following: (24 points)

**a. Fetch an instruction**

Cycle 1  
 PC -> bus3, load MAR  
 clear MD, set MS, set Rd, set Sz

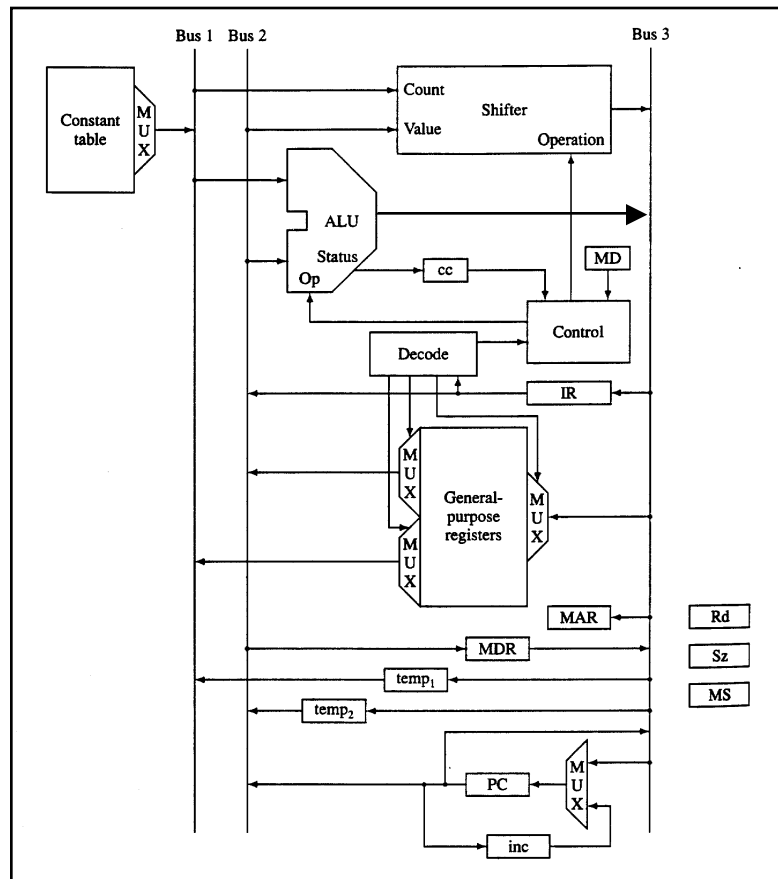
Cycle 2...n-1  
 while (MD == 0) do nothing

Cycle n  
 MDR -> bus3, load IR  
 inc -> PC, load PC  
 clear MS

(Note — you could also explain these same steps textually)

**b. Execute an “add” instruction**

Cycle n+1  
 Reg[src1] -> bus1  
 Reg[src2] -> bus2  
 select ALU add operation  
 ALU -> bus3  
 load Reg[dest]



4. How do the level 1 cache, level 2 cache, and main memory compare in terms of size and access time? (12 points)

Level 1 cache is small (128KB-512KB) and very fast (1-2 clock cycles access time).

Level 2 cache is larger (1MB) and moderately fast (6-8 clock cycles access time).

Main memory is very large (32MB-64MB) and very slow (50-75 clock cycles access time).

5. Briefly summarize one of the two methods for avoiding a write/read data hazard. (15 points)

Explain either register interlocks or data forwarding. See Lecture 32 for details...

6. Compared to a SIMD machine, why is a MIMD machine better suited for control parallelism, and less well suited for massive data parallelism. (10 points)

Name: \_\_\_\_\_

A MIMD machine is better suited to control parallelism because each PE can be executing a different instruction, which works well when one set of PEs execute the “then” part of a conditional and another set of PEs execute the “else” part. A MIMD machine is less well suited to massive data parallelism because it lacks the huge number of PEs typically present in a SIMD machine.

**7. One of the principle items supported by an operating system is the process.  
(15 points)**

**a. List some reasons why a process may be created.**

User logs on, user starts a program, OS creates a process to provide a service, or a program starts another program

**b. How does a process compare to a program?**

A process is one instance of a program in execution. Many processes can be executing the same program.

*Note — there was apparently some confusion as to whether or not Lecture 34 (Introduction to Operating Systems) was supposed to be on the exam. I’m not sure why people thought it was not included, unless they were looking at an earlier version of the class schedule. However, the class schedule posted 1 December clearly indicates that Exam 4 covers Lectures 27-34, and I announced the lectures that would be on the exam all week during the week after Thanksgiving. Accordingly, this question does count, but I graded it very generously.*