

EE380 Spring 2004 Sample Final Exam

This sample exam is derived from the final in Spring 2002. Note that it is only a sample; although general topic coverage will be similar, the precise material that you will be questioned on will be somewhat different.

1. We discussed several different types of network topologies for parallel processing. Which of the following statements about networks is *false*?

- Direct Connection* (aka, fully connected) networks offer very good latency and bandwidth, but do not scale well
- Channel Bonding* refers to the concept of using multiple lower-speed networks in parallel to achieve the performance of a more-expensive higher-speed network
- Flat Neighborhood Networks* (FNNs) make use of the fact that a PC can have more than one NIC
- For communication-intensive parallel computing, networks built using *switches* generally outperform those built using *hubs*
- Most applications that run well on a *cluster* also will run well on a *grid* or *farm*

2. Which of the following four statements about the memory hierarchy is *false*?

- L1 cache generally is faster to access than L2 cache
- Temporal Locality refers to the concept of objects with addresses near each other being referenced over a short period of time
- Larger *cache line* sizes take better advantage of Spatial Locality
- A high *cache hit ratio* usually is a good thing
- None of the above four statements is false

3. Which of the following I/O mechanisms requires the least hardware support?

- Polled
- Interrupt-driven
- DMA

4. Why might one use a larger line size in an instruction cache than in a data cache?

Larger line sizes are better for spatial locality and code tends to have much higher spatial locality than data.

5. Consider the two MIPS subset implementations shown at the back of this test. Which of the following four statements about how pipelining changes the architecture is *false*?

- The ALU used to add 4 to the PC could be the same circuit in both implementations
 - The ALU used for operations like **and** and **sllt** could be the same circuit in both implementations
 - The Data Memory (data cache) module could be the same circuit in both implementations
 - The Instruction Memory (instruction cache) module could be the same circuit in both implementations
- ☞ None of the above four statements is false; in fact, all of the modules can be the same circuits in both implementations because pipelining only adds buffers, changes/adds some datapaths, and modifies the control logic

6. Consider the pipelined MIPS subset implementation shown at the back of this test. Given this design, briefly explain the performance problem that is caused by the **Beq** instruction.

The branch in branch-if-equal might, or might not, be taken; thus, it causes a control hazard in that we don't know which instruction should enter the pipeline immediately after **Beq**. Further, if the branch is taken, the target address might not have been computed quickly enough to avoid inserting a pipeline bubble.

7. In the supercomputing field, there is an old joke that the highest bandwidth communication medium is a station wagon full of magnetic tapes. Modernizing the joke, a typical minivan could hold about 10,000 tape cartridges, each of which stores about 10GB of data. Assume that the total time taken to write the tapes, drive to the other computer, and read the tapes is about 100,000 seconds (just under 28 hours). Which of the following more conventional networks is closest to the same bandwidth? (Hint: b/s isn't the same as B/s.)

- 9600 baud modem (phone line connection)
 - 10 Mb/s Ethernet
 - 100 Mb/s Fast Ethernet
 - 1 Gb/s Gigabit Ethernet
- ☞ 10 Gb/s Ethernet
(Note b is bits, B is Bytes)

8. Consider executing each of the following code sequences on the **pipelined** MIPS implementation given at the end of this test (do not assume any datapaths not shown in the diagram). Incidentally, both code sequences produce the same final results. Which of the following statements *best describes* the execution times you would expect to observe?

(A) `Add $t0,$t1,$t2`
`Add $t5,$t6,$t7`
`Add $t3,$t0,$t4`

(B) `Add $t0,$t1,$t2`
`Add $t3,$t0,$t4`
`Add $t5,$t6,$t7`

- (A) would be faster than (B)
 (B) would be faster than (A)
 (A) would take the same number of clock cycles as (B)

9. Consider executing each of the following code sequences on the **pipelined** MIPS implementation given at the end of this test. Also consider executing them on this design with value **forwarding** logic and datapaths added. Which of the following statements *best describes* how the forwarding logic would alter the execution times?

(A) `Add $t0,$t1,$t2`
`Add $t0,$t0,$t3`
`Add $t0,$t0,$t4`

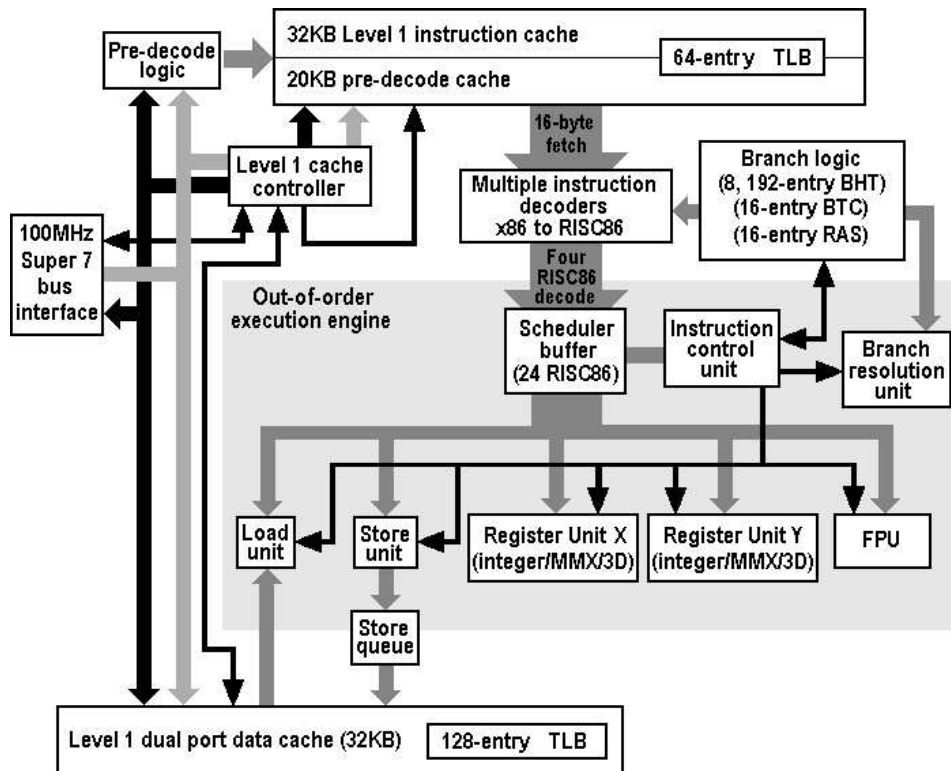
(B) `Add $t0,$t1,$t2`
`Add $t3,$t4,$t5`
`Add $t6,$t7,$t8`

- Neither (A) nor (B) is affected by forwarding
 (A) is not affected, (B) would be faster using forwarding
 (A) would be faster using forwarding, (B) is not affected
 Both (A) and (B) would be faster using forwarding
 The execution time improvements due to forwarding depend on the values in the registers, not on the instructions being executed; thus, it is impossible to say how execution times for (A) and (B) are affected

10. What is a **pipeline bubble**? Precisely what is in a pipeline stage that is “executing” a bubble?

A pipeline bubble is one or more pipeline stages that do not contain useful instructions; i.e., they contain NOPs or other side-effect-free instructions. Bubbles typically are introduced to accommodate various types of pipeline hazards.

11. The following diagram shows the internals of the AMD K6-2 processor. According to the diagram, which of the following techniques is *not* used in this design?



- Separate translation lookaside buffers for code and data
- Separate L1 caches for code and data
- Separate L2 caches for code and data
- Superscalar execution of two SWAR (*MMX* or *3DNow!*) instructions per clock cycle
- Superscalar execution of integer operations

12. Suppose that x is a value represented in 8-bit 2's complement binary notation. Oddly enough, the result of computing $-x$ cannot be represented in 8-bit 2's complement binary notation. What is the decimal value of x ?

2's complement numbers go one value more negative than positive, from $-128..+127$ for 8 bits. Since there is no $-(-128)$, x must be -128 .


13. You have two different ways to implement a particular system. In one version, the cache is on the processor chip; in the other the cache is slower but much larger because it is off chip. Suppose that the first design achieves 1 cycle access latency with a hit rate of 80% and the second design achieves 2 cycle latency with a hit rate of 99%. In either case, main memory has a latency of 10 cycles. Which design is faster and by how much?

This is really $1*80+10*20=280$ vs. $2*99+10*1=208$; the second design is faster.

14. Which of the following four statements about control logic in general is *false*?

- Outputs are enabled at the end of each clock cycle
- Inputs are latched at the end of each clock cycle
- The choice of how instructions are encoded as bit patterns can change the complexity of the control logic
- Although "random logic" implementations of control may result in faster clock rates than schemes using gate arrays, memory, or microcode, it is more difficult to make changes to fix bugs in random logic implementations
- None of the above four statements is false

15. In terms of A, B, and C, what does the gate shown below do?



A	B	C
0	0	Z
0	1	0
1	0	Z
1	1	1

16. What does it mean when we say a cache is “two-way associative”?

It means that the cache uses set associativity in which each line of data from memory can go into either of two line positions within the cache.

17. Consider the single-cycle implementation at the end of this test. Which of the following statements about how we said MIPS instructions could be implemented is *false*?
- Although the ALU operation control lines would be different, the `or` and `sub` instructions are otherwise handled identically
 - The `sw` instruction would set `ALUSrc` to select the lower input to the Mux
 - The Shift function unit is used by the `sle` instruction
 - The result of the Add ALU in the upper right of the diagram is meaningless garbage for an `add` instruction
 - All the above four statements are true.

18. Consider the single-cycle implementation at the end of this test. Which of the following control lines is a “don’t care” while executing a **Beq** instruction?

- PCSrc
- ALUSrc
- RegWrite
- MemtoReg
- ALU operation

19. Which of the following typically is **not** found inside the case of an Intel Pentium 4 or AMD Athlon PC?

- PCI bus
- DRAM
- CPU
- Interociter
- One or more fans

20. What is an instruction that loads an immediate value into the PC generally called?

Jump