

CS 5513 Fall 2009 Quiz

This quiz will help assess whether you are prepared to take CS 5513. It covers a sampling of material that you should have learned in previous computer architecture and operating systems classes. Each question has one right answer. Hexadecimal numbers are preceded by "0x" as in C/C++/Java. If you don't know the answer to a question, please select (e). This quiz will not count toward your grade in CS 5513.

1. In operating systems, a *process* is:

- (a) An algorithm for maintaining printing facilities.
- (b) A method for compiling shell scripts to byte code.
- (c) An abstraction for the central processing unit.
- (d) An instance of a running program.
- (e) I don't know.

2. The hexadecimal number 0x25 in decimal is:

- (a) 25
- (b) 31
- (c) 37
- (d) 47
- (e) I don't know.

3. Which of these does not belong?

- (a) 8
- (b) 16
- (c) 1,536
- (d) 4,096
- (e) I don't know.

4. Which one of these is a volatile memory?

- (a) CD-ROM
- (b) Flash
- (c) ROM
- (d) DRAM
- (e) I don't know.

5. The 10th power of 2 is:

- (a) 256
- (b) 512
- (c) 1,024
- (d) 2,048
- (e) I don't know.

6. Which of these storage technologies costs the least per byte?

- (a) SRAM
- (b) DRAM
- (c) Flash
- (d) Hard disk
- (e) I don't know.

7. Tomasulo's algorithm is:

- (a) An optimal scheduler for jobs in an operating system.
- (b) A greedy algorithm for scheduling disk accesses.

- (c) An algorithm for out-of-order execution of CPU instructions.
- (d) A cache replacement algorithm that evicts the least recently used block.
- (e) I don't know.

8. The concept of a page table:

- (a) Relates to virtual memory.
- (b) Relates to instruction encoding.
- (c) Relates to symmetric multi-processors.
- (d) Relates to cache memory.
- (e) I don't know.

9. The sum of the binary numbers 100101 and 11010 in decimal is:

- (a) 32
- (b) 63
- (c) 64
- (d) 127
- (e) I don't know.

10. On a 32-bit byte-addressed little-endian machine the hexadecimal value 0xA1B2C3D4 is stored into address 0x100. What is the value of the byte stored in address 0x101?

- (a) 0xB2
- (b) 0xC3
- (c) 0x3D
- (d) 0x2C
- (e) I don't know.

11. In Motorola 6809 processor machine language, the byte sequence 0x27 0x10 means "if the result of a previous operation was zero, then jump 16 bytes forward." Given only this information, which of these facts can we deduce?

- (a) The 6809 is a big-endian machine.
- (b) The 6809 uses 16-bit addresses.
- (c) The 6809 uses condition codes.
- (d) The 6809 has index registers.
- (e) I don't know.

12. Which of these statements is true?

- (a) Direct-mapped caches usually have a lower miss rate than set-associative caches.
- (b) Direct-mapped caches have at least 2 blocks per set.
- (c) A direct-mapped cache will have more tag bits than a set-associative cache with the same capacity.
- (d) Direct-mapped caches do not need a replacement algorithm.
- (e) I don't know.

13. What is the IEEE 754 Standard?

- (a) It is a standard for determining voltage and clock rate in TTL design.
- (b) It is a standard for encoding character data as octet values.
- (c) It is a standard for representing floating point numbers in binary.

- (d) It is a standard reference instruction set from which several popular processors are derived.
(e) I don't know.

14. Which of these does not occur with a fully-associative cache?

- (a) Compulsory misses.
(b) Conflict misses
(c) Capacity misses
(d) Cache coherence
(e) I don't know.

15. How many inputs and outputs does a full-adder circuit have?

- (a) One and One
(b) Two and Two
(c) Three and Two
(d) Three and Three
(e) I don't know.

16. A page fault occurs when:

- (a) A read from the disk has a checksum error.
(b) Data is read into memory but stale values remain in the cache, as in a DMA transfer to a cached block.
(c) Address translation fails to map a memory reference to a physical memory location.
(d) Two processors attempt to share the same block of memory.
(e) I don't know.

17. TLB stands for:

- (a) Translation look-aside buffer.
(b) Test, load, and branch.
(c) Transfer loaded binary.
(d) Target lookahead byte.
(e) I don't know.

18. The performance improvement from pipelining comes from:

- (a) The inclusion of a non-blocking cache for resolving data hazards.
(b) The simpler decoding and increased number of registers in a RISC instruction set.
(c) The ability to increase the clock rate.
(d) The increased number of instructions executed per cycle.
(e) I don't know.

19. What is Moore's Law?

- (a) The number of devices on an integrated circuit doubles every 18 months.
(b) The maximum clock frequency of microprocessors doubles every 18 months.
(c) The maximum expected improvement to a computer system is limited by the number of instructions to which that improvement can be applied.
(d) The current between two points is inversely proportional to the resistance between them.
(e) I don't know.

20. What is spatial locality?

- (a) The use of data elements within close storage locations.
(b) A layout methodology for motherboard design.
(c) Placing pipeline stages on the chip in a physical arrangement that mirrors their logical functions.
(d) A network topology for interconnecting multiple cores on a single chip.
(e) I don't know.

21. Which of these architectural features is necessary for a pre-emptive multitasking operating system?

- (a) Cache memory.
(b) A memory management unit.
(c) Multiple cores.
(d) Interrupts.
(e) I don't know.

22. Consider a computer system with 32-bit addresses and 4GB of RAM. The data cache has 32-byte blocks. What information could you use to compute the size of a tag in the cache?

- (a) The size of the cache in bytes and the number of blocks per set.
(b) The number of page table entries and the number of sets.
(c) The size of the cache in bytes and the size of a page.
(d) All of the above.
(e) I don't know.

23. What is an example of an atomic instruction that is used to implement semaphores and locks?

- (a) Rotate-left-and-carry
(b) Clear-and-wait-for-interrupt
(c) Sign-extend
(d) Test-and-set
(e) I don't know.

24. Which of these is *not* a pipeline data hazard?

- (a) Read-after-write
(b) Read-after-read
(c) Write-after-write
(d) Write-after-read
(e) I don't know.

25. Suppose we increase the number of processors from p to $2p$. Is it possible that some reasonably written programs will run more than twice as fast under the new system? Why or why not?

- (a) No. Speedup is linear at best.
(b) No. Communication overhead increases with an increased number of processors.
(c) Yes. Some algorithms have a running time that is quadratic in the number of processors.
(d) Yes. Cache effects can explain a superlinear speedup.
(e) I don't know.