

Computer Science 260  
Midterm 2  
October 24, 2007

Your name: \_\_\_\_\_

Write your name on your exam paper now and check to make sure that you have \_3\_ sheets printed front and back.

Part I: Multiple choice – circle the letters of the answer or answers that result in true statements. Please be careful, some of the lettered answers may be true by themselves but are not true in the context of the lead-in to the problem.

1. The BIOS code for a computer would typically be stored in
  - a) volatile read-only memory
  - b) non-volatile read-write memory
  - c) sequential read-only memory
  - d) non-volatile read-only memory
2. Why do computer designers organize the memory in a hierarchy?
  - a) Some data are more important than others
  - b) Memories higher in the hierarchy exert control over lower memories
  - c) Slower memory is less expensive so more of it can be provided
  - d) Having more fast memory contributes to good performance
3. Static RAM is often used for cache memory because
  - a) it is less expensive than dynamic RAM
  - b) it requires less chip area than dynamic RAM
  - c) it is faster than dynamic RAM
  - d) it is slower than dynamic RAM
4. The read cycle time of a memory characterizes
  - a) how fast the memory can respond to successive read operations
  - b) the time from when a read request is made to when the response is available
  - c) the latency of read requests
  - d) the time required between a write request and subsequent read request
5. As a programmer you need to be aware of byte alignment issues because
  - a) referring to unaligned data may cause the processor to take an exception (interrupt)
  - b) referring to unaligned data may be slow
  - c) a load of unaligned data may produce an incorrect value
  - d) a store of unaligned data may produce an incorrect value
6. A computer that uses 32-bit addresses to refer to bytes in storage can refer to at most
  - a) approximately 4 billion bytes of memory
  - b) approximately 2 billion bytes of memory
  - c) approximately 64 thousand bytes of memory
  - d) approximately 16 million words of memory
7. Content addressable memory

- a) is included in computers so application programmers can implement fast lookup tables
  - b) is used for translation look-aside buffers
  - c) is preferred for caches because it is so cheap
  - d) is an exotic technology that is not practical for use in computer systems
8. Compared to using physical memory directly, virtual memory
- a) may make programming more convenient
  - b) requires additional hardware in the form of a memory management unit
  - c) runs the risk that multiple processes may overwrite each others' memory
  - d) requires that physical memory addressable by a process be contiguous
9. If a computer uses 32-bit addresses and uses a paging system with 4Kbyte (4096 bytes) pages, how many bits are needed in each page table entry to represent the location of the virtual page in physical memory?
10. In the situation described in question 9, how many different pages exist within the 32-bit address space?
11. In the situation described in question 9, how many page table entries fit on a page? (You will

need to make some assumptions – clearly state your assumptions and show your work in calculating your answer.

12. Why is a translation look-aside buffer important for increasing the performance of a system using virtual memory?
13. Explain the LRU page replacement algorithm. What problem does it solve? What prediction about program behavior underlies its use?

14. Suppose that a computer has a 1MByte cache with access time 1ns (1 nanosecond) and 1GByte ( $2^{30}$  bytes) main memory with access time 40ns and on-disk page storage with access time 10ms (10 milliseconds). If the hit rate in the 1MByte cache is 40% and 0.01% (i.e. 1 in 10,000) of all memory references cause a page fault requiring a disk access, what is the average time required for a memory reference?

15. Define

- a) Direct mapped cache
- b) 2-way set associative cache
- c) Fully associative cache

16. Define the following two terms. Include in your definitions the appropriate units for each.

- a) Latency
- b) Throughput

17. The book says that parallel interfaces are faster than serial interfaces, yet in recent years, as disk speeds have increased the industry has moved to serial interfaces. Why?
18. What is DMA? How does it contribute to overall system performance?
19. What is interrupt-driven I/O? How does it contribute to overall system performance? What are some issues that come up in programming a system using interrupt-driven I/O?
20. List the program steps (in the order they are executed) used for polling I/O for output. As a specific example consider the display output in the SPIM simulator. You don't have to use precise addresses or even exact MIPS instructions – just say what has to happen.

Extra space for answers or scratch work