NOTE: It is my policy to give a failing grade in the course to any student who either gives or receives aid on any exam or quiz.

INSTRUCTIONS:

- **Answer all questions on your Scantron answer sheet.**
- **Be sure *your* name and student ID number are printed on your answer sheet.**
- **Be sure the Exam ID printed on your answer sheet is 2486.**
- **Unless otherwise indicated, all questions count equally.**

1. What is the Exam ID printed on your answer sheet?
   A. 2486
   B. 3158
   C. 6739
   D. 4856
   E. 1577

Note: If your answer to question 1 is not "A," see Dr. Vickery before proceeding!

2. What does the *enable* input to a decoder do?
   A. It makes the decoder work like a multiplexer.
   B. It makes the decoder work like a flip-flop.
   C. It makes the decoder into an encoder.
   D. If it is false, all the outputs are false.
   E. If it is true, all the outputs are true.

3. What is a decoder used for in the design of a memory system with multiple RAM chips?
   A. It uses some of the address bits to select one of the rows of RAM chips.
   B. It uses some of the data-in bits to select one of the rows of RAM chips.
   C. It uses some of the data-out bits to select one of the rows of RAM chips.
   D. It is used to decide whether to do a read operation or a write operation.
   E. It is used to reset the system.

4. A 128M x 32 memory system is to be constructed using 64M x 32 RAM chips. Which of the choices best describes how this system would be built?
   A. 1 RAM chip and a 4x1 decoder.
   B. A column of 8 RAM chips, plus a 3x8 decoder to select one chip.
   C. A column of 4 RAM chips, plus a 2x4 decoder to select one chip.
   D. A column of 2 RAM chips, plus a 1x2 decoder to select one chip.
   E. 1 RAM chip, and no decoder.

5. A 64M x 64 memory system is to be constructed using 64M x 8 RAM chips. Which of the following choices best describes how this system would be built?
   A. One row of 64 RAM chips and an 8 x 1 decoder.
   B. One row of 16 RAM chips and no decoder.
   C. One row of 8 RAM chips and no decoder.
   D. One row of 4 RAM chips and no decoder.
   E. One row of 2 RAM chips and no decoder.

6. How many address wires connect to a 128M x 32 memory?
   A. 7
   B. 6
   C. 17
   D. 16
   E. 27

7. What does CS stand for, and what is it used for on a RAM chip?
   A. It stands for Computer Science, and it's not used on a RAM chip.
   B. It stands for Column Strobe, and it's used to read all the data from a column of RAM chips at the same time.
   C. It stands for Chip Select and it's used to control whether the chip responds to its other inputs.
   D. It stands for Chip Select and it's used when copying data from one chip to another.
   E. It stands for Copy Section and it's used when copying data from one chip to another.

8. "Memory Hierarchy" refers to an ordering of storage devices according to which of the following parameters?
   A. Access time, cost per bit, and capacity.
   B. Importance, volatility, and associativity.
   C. Bits per word, words per block, and volatility.
   D. Capacity, access time, and importance.
   E. Cost per bit, bits per word, and words per block.

9. What is the idea of having a split cache (Harvard Architecture)?
   A. Because the Ivy League has a separate funding mechanism.
   B. Because the CPU makes different patterns of reads and writes of instructions compared to data.
   C. Because the cache must be kept separate from the main memory.
   D. Because the cache must be kept separate from the CPU.
   E. Because it's better to put the ALU and the controller in different parts of a RAM chip.

10. Define a set associative cache design.
    A. A block of main memory may be loaded into any cache line.
    B. Each block of main memory is mapped to exactly one cache line.
    C. A block of main memory is mapped to a group cache lines.
    D. When the CPU writes a word, the data is immediately written both to cache and to main memory.
    E. When the CPU reads a word, the cache writes part of a block to memory.

11. Define a write through cache design.
    A. A block of main memory may be loaded into any cache line.
    B. Each block of main memory is mapped to exactly one cache line.
    C. A block of main memory is mapped to a group cache lines.
    D. When the CPU writes a word, the data is immediately written both to cache and to main memory.
    E. When the CPU reads a word, the cache writes part of a block to memory.

12. Define a fully associative cache design.
    A. A block of main memory may be loaded into any cache line.
    B. Each block of main memory is mapped to exactly one cache line.
    C. A block of main memory is mapped to a group cache lines.
    D. When the CPU writes a word, the data is immediately written both to cache and to main memory.
    E. When the CPU reads a word, the cache writes part of a block to memory.

13. Define a direct mapped cache design.
    A. A block of main memory may be loaded into any cache line.
    B. Each block of main memory is mapped to exactly one cache line.
    C. A block of main memory is mapped to a group cache lines.
    D. When the CPU writes a word, the data is immediately written both to cache and to main memory.
    E. When the CPU reads a word, the cache writes part of a block to memory.

14. What is a "valid bit"?
    A. A condition code bit that tells whether the ALU function code is one of the defined values or not.
    B. A cache housekeeping bit that tells whether there was a hit or not.
    C. A cache housekeeping bit that tells whether a line has been modified since being loaded from memory.
    D. A bit used by device controllers to tell the PIC when a data transfer has completed.
    E. A cache housekeeping bit that indicates whether a cache line actually has data from main memory in it or not.

15. What is a "dirty bit"?
    A. A condition code bit that tells whether the ALU function code is one of the defined values or not.
    B. A cache housekeeping bit that tells whether there was a hit or not.
    C. A cache housekeeping bit that tells whether a line has been modified since being loaded from memory.
    D. A bit used by device controllers to tell the PIC when a data transfer has completed.
    E. A cache housekeeping bit that indicates whether a cache line actually has data from main memory in it or not.

16. A Computer has 128MB of main memory. It has a set associative cache with 128 sets with 4 lines per set (4-way set associative). Each cache line holds 32 bytes. How many bits in an address select the byte within a line?
    A. 4
    B. 5
    C. 6
    D. 7
    E. 8

17. For the computer in Question 16, how many address bits select the set?
    A. 5
    B. 6
    C. 7
    D. 8
    E. 9

18. For the computer in Question 16, how many address bits are used for the tags?
    A. 7
    B. 9
    C. 11
    D. 13
    E. 15

19. For the computer in Question 16, how many bytes of memory are there in the cache, not counting the housekeeping information?
    A. 2^9
    B. 2^{11}
    C. 2^{14}
    D. 2^{17}
    E. 2^{20}

20. A computer has a main memory access time of 350 nsec, a cache access time of 10 nsec, and a particular program running on this computer has a hit ratio of 0.95. What is the effective access time for this program?
    A. 350 nsec.
    B. 10 nsec.
    C. 0.95 nsec.
    D. 27 nsec.
    E. 350 psec.

21. What are the three groups of wires in a memory or I/O bus?
    A. Control, Data, and Arbitration
    B. Interrupts, Vectors, and Control
    C. Input, Output, and Data
    D. Data, Address, and Control
    E. Data, Vectors, and Arbitration

22. What is a device controller?
    A. A program that responds to interrupts.
    B. Hardware that acts as an interface between a device and a bus.
    C. A program that manages buffers in the runtime library.
    D. Hardware that controls the state of the CPU's IRQ pin.
    E. Hardware that manages the state of the condition code bits.

23. What is an interrupt vector?
    A. A list of device controllers.
    B. Housekeeping bits in the cache.
    C. A list of addresses of software routines.
    D. The wires going into the PIC from the device controllers.
    E. The wire going from the PIC to the CPU's IRQ pin.

24. Is it possible to do I/O transfers before the interrupt vector and/or the PIC have been initialized?
    A. Yes, by doing polled I/O.
    B. No, they both have to be initialized.
    C. The interrupt vector but not the PIC must be initialized.
    D. The PIC but not the interrupt vector must be initialized.
    E. Neither has to be initialized before doing interrupt-driven I/O.

25. What is EOI?
    A. A wire used by device controllers to indicate when they have finished causing interrupts.
    B. A wire used by the PIC to tell the CPU when an interrupt is finished.
    C. A control word output by an interrupt service routine to the PIC telling it an interrupt has been handled.
    D. The name of an ISA instruction a device driver executes when it wants the kernel to transfer a data buffer to the runtime library.
    E. How much money the computer manufacturer earned on his/her investment.

26. What condition(s) must be met before the CPU will acknowledge an interrupt request?
    A. The device controller must have signaled an interrupt request to the PIC.
    B. The PIC must have signaled an interrupt request to the CPU.
    C. The CPU must be between instructions.
    D. The CPU's interrupt enable bit must be true.
    E. All of the above.

27. What is the difference between a synchronous bus and an asynchronous bus?
    A. A synchronous bus doesn't have any data wires, but an asynchronous bus does.
    B. A synchronous bus has address wires, but an asynchronous bus does not.
    C. A synchronous bus has a clock wire, but an asynchronous bus does not.
    D. A synchronous bus can connect to more devices than an asynchronous bus.
    E. A synchronous bus cannot connect to memory, but an asynchronous bus can.

28. What is the advantage of DMA?
    A. Fewer interrupts and less bus traffic.
    B. More interrupts and less bus traffic.
    C. Fewer interrupts and more bus traffic
    D. More interrupts and more bus traffic.
    E. No interrupts and no bus traffic.

29. What does the PIC do when the CPU acknowledges an interrupt request?
    A. Tells the device controller to write data to the CPU.
    B. Tells the device controller to read data from the CPU.
    C. Writes the highest priority active interrupt level to the CPU.
    D. Reads the IRQ from the CPU and puts it in the application's buffer.
    E. Restores electrical power to the CPU and the DCs.

30. A disk has six recording surfaces. Which statement is true?
    A. A track can be read from it six times faster than from a disk with one recording surface.
    B. A track can be read from a disk with one recording surface six times faster than this disk.
    C. This is a 6X CD-RW drive.
    D. It has six read-write heads.
    E. It has six tracks per surface.

31. A disk spins at 6000 RPM and has ten sectors per track. How long would it take to read one sector from the disk to its controller?
    A. One second
    B. One millisecond
    C. One microsecond
    D. One nanosecond
    E. One picosecond

32. What is the bandwidth of a bus that has a clock speed of 100MHz and 8 data wires?
    A. 800MHz
    B. 800MB/s
    C. 800Mb/s
    D. 6400Mb/s
    E. 80KHz

33. A display controller is used to draw images on a monitor at a refresh rate of 100Hz. Each image has a resolution of 1024x2048 pixels, and each pixel can be any one of 2^{18} intensities. What is the bandwidth needed to transfer images to the controller in real time?
    A. $2^{100} * 2^{21} * 2^{18}$ bytes per second.
    B. $100 * 2^{21} * 2^{18}$ bits per second
    C. $100 * 2^{21} * 18$ bits per second
    D. $2^{100} + 2^{21} + 2^{18}$ bytes per second.
    E. 14 nanoseconds.

34. What is the purpose of a bus bridge?
    A. To connect two CPUs together.
    B. To connect two DCs together.
    C. To buffer data between a bus and the runtime library.
    D. To handle bandwidth differences between two busses.
    E. To connect the CPU to the chipset.

35. Which is faster, a PCI bus or a USB bus?
    A. A USB bus, because it is newer.
    B. A USB bus, because it transfers more bits per clock pulse.
    C. A PCI bus, because it has just one data wire.
    D. A PCI bus, because it has more data wires.
    E. Neither, they are two different names for the same thing.