Instructions

- Use pencil, if you have one.
- For multiple-choice questions, circle the letter of the *one best choice* unless the question specifically says to select "all" correct choices.
- There is no penalty for guessing, so answer all questions.
- Place drawings where indicated in the question; be sure to put the question number next to your drawing; use pencil rather than ink.
- No calculators or other electronic aids allowed.

1. Which instruction above lets you know that you won't be asked to compute the square root of π on this exam?

   A. 3.14159
   B. 1.77245
   C. The one that says there is no penalty for guessing.
   D. The one that says to use pencil.
   E. The one that says calculators are not allowed.

2. What are the inputs to the adder that computes the Branch Target Address?

   A. 3 and 5
   B. The outputs of data memory and the ALU
   C. The two outputs from the register file
   D. The output from the PC+4 adder and the immediate field of the instruction multiplied by 4 and sign extended
   E. The PC and *RegDst*

3. Circle *all* the values computed by the ALU. (*Hint: there is more than one letter that has to be circled.)*

   A. PC+4
   B. The Branch Target Address
   C. Memory addresses for *lw* and *sw* instructions
   D. The difference between two registers being compared for *beq* instructions
   E. The result for *R-type* instructions

4. From where do the *RR1* and *RR2* inputs to the register file come?

   A. RD1 and RD2
   B. The *rs* and *rt* fields of the instruction being executed
   C. R[rs] and R[rt]
   D. Data Memory
   E. The ALU

5. From where do the ALU inputs come?

   A. The *rs* and *rt* fields of the instruction being executed.
   B. The DataIn and DataOut outputs of data memory.
   C. The RD1 output of the registers and the multiplexer that selects between the RD2 output of the register and the sign extended immediate field of the instruction being executed.
   D. The RD1 and RD2 outputs of the register file.
   E. The outputs of the ALU and Data Memory

6. What are the values of *MemRead, MemWrite,* and *RegWrite* (in that order) for *lw* instructions?

   A. 0, 0, 1
   B. 0, 1, 0
   C. 0, 1, 1
   D. 1, 0, 0
   E. 1, 0, 1

7. Where do the signals named *R-type, lw,* and *sw* (among others) come from?

   A. The op code decoder in the control unit
   B. The register file
   C. The PC
   D. The ALU
   E. Data Memory

8. What is the only instruction for which the *MemWrite* control signal is true?

   A. *add*
   B. *slt*
   C. *lw*
   D. *sw*
   E. *beq*

9. Which statement describes a memory *read* operation?

   A. The data on the DataIn lines is stored in the location specified by the address lines.
   B. The data on the address lines is stored in the location specified by the DataIn lines.
   C. The data at the location specified by the address lines becomes available on the DataOut lines.
   D. The data on the DataOut lines is made available on the address lines.
   E. The address of the data replaces the data of the address.

10. What is the only instruction for which the *MemToReg* control signal is true?

    A. *add*
    B. *slt*
    C. *lw*
    D. *sw*
    E. *beq*

11. What do *sw, beq*, and *j (jump)* instructions all have in common with one another?

    A. They do not use the PC
    B. They do not use the ALU
    C. They do not use Instruction Memory
    D. They do not use Data Memory
    E. They do not write to the Register File

12. The time interval from starting to fetch and instruction until its execution completes is called:

    A. *rate*
    B. *throughput*
    C. *latency*
    D. *frequency*
    E. *leading edge*

13. What characteristic of instruction processing would be measured in *instructions per second?*

    A. *period*
    B. *throughput*
    C. *latency*
    D. *trailing edge*
    E. *flip-flop*

14. A single-cycle datapath (like the first one described in Chapter 5) is observed to execute one million instructions per second. What is the clock *frequency*?

    A. 1 MHz
    B. 1 GHz
    C. 1 μsec
    D. 1 nsec
    E. 1 cm

15. How long does it take to execute a billion instructions if a processor uses one clock cycle per instruction and the clock period is 0.5 nsec?

    A. 0.5 sec
    B. 1.0 sec
    C. 2.0 sec
    D. 5.0 sec
    E. 1 G

16. If a single-cycle datapath requires one second to process one instruction, how long will it take to execute one million instructions?

    A. One minute
    B. One hour
    C. One day
    D. One second
    E. One million seconds

17. If a single-cycle datapath requires one second to process one instruction, and is re-implemented as a perfectly balanced ten-stage pipeline, how long will the new datapath take to process *one instruction*?

    A. One second
    B. One tenth of a second
    C. Ten seconds
    D. Two seconds
    E. Half a second

18. If a single-cycle datapath requires one second to process one instruction, and is re-implemented as a perfectly balanced ten-stage pipeline, how long will the new datapath take to process *one million instructions*? (Pick the best choice, even though it is not exactly right.)

    A. 10,000,000 seconds
    B. 1,000,000 seconds
    C. 100,000 seconds
    D. 10,000 seconds
    E. 1,000 seconds

19. Would the actual answer to the previous question be less or more than the value given, and why?

    A. It would be less (faster) because pipelines can take advantage of spatial locality.
    B. It would be more (slower) because hazards reduce pipeline performance.
    C. It would be less (faster) because of multicore microarchitecture.
    D. It would be more (slower) because of the need for additional cooling modules.
    E. It would be less (faster) because of the overlap in processing between pipeline stages.

20. What is the name of the type of pipeline hazard introduced by *branch* instructions?

    A. Structural
    B. Data
    C. Control
    D. Memory
    E. Latency

21. Which of the following describes *register forwarding*?

    A. Connecting the outputs of the register file to the inputs of the ALU without going through the ID/EX pipeline register.
    B. Getting a register operand from the EX/MEM or MEM/WB pipeline register instead of from the register file.
    C. Exchanging the positions of the IF/ID and ID/EX registers in the pipeline.
    D. Eliminating the EX/MEM pipeline register.
    E. Using the MEM/WB pipeline register as the input to the PC.

22. What is stored in the EX/MEM pipeline register?

    A. All the control and data information required by the MEM and WB stages of the pipeline
    B. A copy of the PC register and the flip-flops that are in the ALU
    C. A copy of PC register and the full-adders that are in the Instruction Memory
    D. The outputs of the Data Memory and the inputs to the Instruction Memory
    E. PC + 4

23. If a pipeline has *n* stages, how many places does the system clock connect to?

    A. 0
    B. 1
    C. $n - 1$
    D. n
    E. $n + 1$

24. How are a memory's *speed* and *access time* related?

    A. They are reciprocals of each other: the shorter the access time, the greater the speed
    B. They are the same thing
    C. There is no relationship between them
    D. Access time is a rate, and speed is an interval
    E. Speed depends on cost, but access time does not

25. A disk spins at 3600 RPM. What is its rotational *period*, to 3 decimal places, in seconds:

    _____ seconds.

26. Using current technology, which of the following is the most accurate value for the number of bits that can be stored in *one square inch* of a disk's surface?

    A. 600 bits
    B. 6,000 bits
    C. 60,000 bits
    D. 600,000 bits
    E. 600,000,000,000 bits

27. Which of the following statements is *not* true about the memory hierarchy?

    A. The closer to the CPU, the faster the memory.
    B. The closer to the CPU, the shorter the memory's access time.
    C. The closer to the CPU, the larger the capacity of the memory.
    D. The closer to the CPU, the smaller the capacity of the memory.
    E. The closer to the CPU, the higher the cost per bit of the memory.

28. What feature of cache design takes advantage of the fact that programs often execute instructions taken from successive locations in memory?

    A. When fetching an instruction causes a cache miss, not only the word containing the instruction, but an entire block of memory containing other nearby instructions is brought into cache.
    B. By increasing the set size, the penalty for cache misses can be eliminated.
    C. All instructions are loaded into cache when a program starts, so there are no cache misses until data accesses occur.
    D. Because the cache is slower than main memory, fetching instructions from main memory takes advantage of the lower cost per bit of the cache memory.
    E. The cache's tag bits eliminate the structural hazards caused by branch instructions.

29. If a processor referenced memory in a totally random fashion, what would be the probability of a cache hit once the cache is full?

    A. Size of main memory divided by size of cache
    B. Size of cache divided by size of main memory
    C. Size of main memory times size of cache
    D. Size of main memory minus size of cache
    E. Size of main memory plus size of cache

30. For this and the following questions, assume a computer uses a **byte-addressable main memory with a capacity of 2^{36} bytes**.
    How many *address bits* would this computer send to the memory system when performing a read or write operation?

    A. 2^{36} for both reading and writing
    B. 36 for both reading and writing
    C. 8 for both reading and writing
    D. 2^{36} for reading and 0 for writing
    E. 36 for reading and 0 for writing

31. For this and the following questions, assume the computer uses **256-byte blocks.**
    How many <u>blocks</u> of memory are there?

    A. $36 \div 8$
    B. $36 \div 2^8$
    C. $2^{36} \div 2^8$
    D. $2^{36} \times 2^8$
    E. 123

32. How large is a cache line?

    A. 2^8 bits
    B. 2^8 bytes
    C. 2^8 blocks
    D. 36 bits
    E. 36 bytes

33. For this and the following questions, assume the computer has **2^{16} cache lines**.
What is the *capacity* (*not* including tag and *v* bits) of the cache?

    A. 17 KB
    B. $2^{16} \times 2^{8}$ bytes
    C. 16 + 8 bytes
    D. $2^{36} \div 2^{16}$ bytes
    E. $2^{36} \div 2^{16}$ bits

34. How many *sets* are there if the cache is *direct mapped?*

    A. 1
    B. 16
    C. 2^{16}
    D. 2^{14}
    E. 2^{18}

35. How many *sets* are there if the cache is *4-way associative*?

    A. 1
    B. 16
    C. 2^{16}
    D. 2^{14}
    E. 2^{18}

36. How many *sets* are there if the cache is *fully associative?*

    A. 1
    B. 16
    C. 2^{16}
    D. 2^{14}
    E. 2^{18}

37. How many *bytes per word* are there?

    A. 2
    B. 4
    C. 8
    D. 16
    E. You didn't tell us.

38. How many address bits are needed for the *index field* using a *direct mapped* design?

    A. 0
    B. 14
    C. 16
    D. 18
    E. 2^{16}

39. How many address bits are needed for the *index field* using a *4-way set associative design*?

    A. 0
    B. 14
    C. 16
    D. 18
    E. 2^{16}

40. What is the purpose of *tag bits*?

A. They tell which block of memory occupies a cache line
B. They tell which cache line occupies a block of memory
C. They are the minimum amount of information that can be read from or written to a disk
D. They hold the op code of the instruction being executed
E. They compensate for the increased cost per bit of disks compared to registers.