Instructions:

- For multiple choice questions, circle the letter of the one best choice unless the question explicitly states that it might have multiple correct answers (in which case there still might be just one correct choice).
- There is no penalty for guessing.
- Place drawings on the back of any exam sheet.
  - Be sure to put the question number next to your drawing
  - Use pencil rather than ink.
- Unless otherwise indicated, all questions count equally.

1. Complete this timing diagram for a positive edge-triggered D flip-flop. Note that Q is assumed to be initially *true.* Use vertical lines to show how changes in Q line up with changes in the inputs.

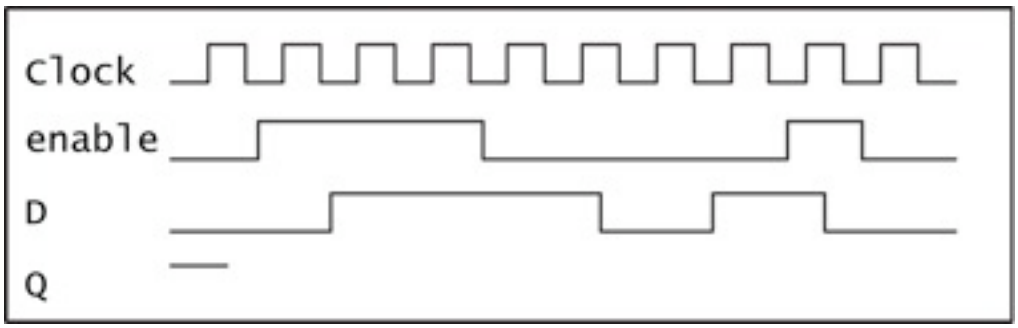

2. Draw a single symbol that represents the entire MIPS register file. Label all inputs and outputs meaningfully, and indicate the number of wires for each input and output groups. *Answer on the back of an exam sheet.*
3. Describe in English how the target address of a *beq* instruction is computed:

4. Write the Verilog expression corresponding to the answer to Question 3:

5. Draw diagrams for the R, I, and J MIPS instruction formats. Show the names and number of bits in each field of each format; be sure to indicate which format is which:

6. What *binary property* do all MIPS instruction addresses have in common with each other?

7. What determines the minimum allowable period for the MIPS single-cycle datapath clock? Be as specific as you can:

8. Name the two elements of the MIPS single-cycle datapath that are connected to the datapath clock:

9. Explain what the *RegDst* control signal and its associated multiplexer do and why they are needed:

10. A single-cycle processor design with a 1 GHz clock is converted to a ten-stage, perfectly balanced, pipelined design. What will be the clock *period* of the new design?
    A. 0.1 GHz
    B. 10 GHz
    C. 1000 picoseconds
    D. 100 picoseconds
    E. 10 nanoseconds
11. What is instruction *latency*? (*If you choose the correct answers for this question and the next one, but reverse them, you will get half credit for both.*)
    A. The time it takes to fetch an instruction from memory.
    B. The time from starting to fetch an instruction to completing its execution.
    C. The rate at which instructions are executed.
    D. The frequency of the clock
    E. The period of the clock
12. What is instruction *throughput?*
    A. The time it takes to fetch an instruction from memory.
    B. The time from starting to fetch an instruction to completing its execution.
    C. The rate at which instructions are executed.
    D. The frequency of the clock
    E. The period of the clock

13. What is the unit of measure for instruction throughput?
    A. Instructions per second
    B. Megahertz
    C. Nanoseconds
    D. Number of instructions
    E. Number of seconds
14. What is the advantage of pipelining?
    A. It increases instruction latency
    B. It increases instruction throughput
    C. it decreases instruction latency
    D. It decreases instruction throughput
    E. It makes the CPU more expensive
15. In the pipelined design, why is there a separate adder for computing PC+4 instead of just using the ALU? (Bonus: name the type of hazard the extra adder eliminates.)

16. Write a sequence of two MIPS assembly language instructions that would cause a pipeline *data hazard*, and describe (briefly) the mechanism for eliminating that hazard.

17. Define the term pipeline *control hazard*:

18. What is *delayed branching*; what does it accomplish?
    A. Branches take longer than other instructions; so other instructions can catch up with them.
    B. Branches take longer than other instructions; so the CPU has time to ask the user whether to take the branch or not.
    C. Branches are not executed until the CPU knows whether the next memory access is *lw* or *sw*; this prevents the wrong memory operation from being performed.
    D. Branches are executed as a group only after all the other instructions in the program have completed; this avoids messy details that make it hard to execute the most important instructions as efficiently as possible.
    E. The next instruction after a conditional branch instruction is executed unconditionally; this avoids a control hazard.

19. What is a pipeline *bubble*?
    A. The equivalent of a *nop* (no operation) instruction that gets inserted into the instruction stream to prevent a pipeline hazard.
    B. An instruction that floats to the first position after a conditional branch instruction to fill the delayed branch slot.
    C. Executing instructions in a different sequence from their order in the instruction stream in order to avoid structural hazards.
    D. An instruction that floats to a point earlier in the instruction stream in order to avoid sinking.
    E. A mechanism for minimizing cache misses by performing memory writes as soon as they occur.
20. In the MIPS pipeline developed in the book, where do the write register number and write data come from?
    A. From the IF/ID pipeline register
    B. From the ID/EX pipeline register
    C. From the EX/MEM pipeline register
    D. From the MEM/WB pipeline register
    E. The register number comes from the ALU and the data comes from Instruction Memory
21. Which pipeline register(s) hold the values of `R[rs]` and `R[rt]` in it/them? Circle all correct letters.
    A. IF/ID
    B. ID/EX
    C. EX/MEM
    D. MEM/WB
    E. WB/IF
22. Which pipeline register(s) hold the value computed by the ALU? Circle all correct letters.
    A. IF/ID
    B. ID/EX
    C. EX/MEM
    D. MEM/WB
    E. WB/IF
23. Which instruction(s) cause data from data memory to be written to the register file? Circle all correct letters.
    A. lw
    B. sw
    C. add
    D. beq
    E. addi
24. What is the *hit ratio*?
    A. The proportion of instruction fetch operations that actually succeed in fetching a new instruction
    B. The number of instructions divided by the number of memory accesses
    C. The number of read operations divided by the sum of the number of read and number of write operations
    D. The number of clock cycles divided by the number of instructions executed
    E. The number of memory accesses that are satisfied by using data in cache divided by the total number of memory accesses

25. Tell how having multiple words per memory block/cache line helps improve the hit ratio.

26. Tell how the phenomenon of *locality* improves the hit ratio.

27. What would be the average access time of a memory system with one level of cache, a main memory access time of 20 nsec, a cache access time of 200 psec, and a hit ratio of 0.9?
    A. 2.18 sec
    B. 2.18 msec
    C. 2.18 μsec
    D. 2.18 nsec
    E. 2.18 psec

The following questions (28-31) are based on a memory system that has 4 gigabytes of byte addressable memory; 4 bytes per word; 16 words per cache line; 64 kilobytes of cache data; and a 1-way associative (direct mapped) cache.

28. Draw a diagram that shows the relevant fields of a memory address; show the number of bits in each field.

29. How many *cache lines* are there? Show your work.

30. Re-draw your answer to Question 28, but for a 2-way associative (set size = 2) cache.

31. How many *cache sets* are there for the 2-way associative design?

32. List one *advantage* and one *disadvantage* of increasing the degree of associativity (set size) of a cache system.

33. Define each of the following components of a hard disk drive in English. Do not draw a diagram.
    A. Platter

    B. Surface

    C. Head

    D. Track

    E. Sector

34. Name and define the three major components of disk latency (access time):
    A.

    B.

    C.

35. What would be the average rotational delay of a 15,000 RPM disk?
    A. 2 sec
    B. 2 msec
    C. 2 μsec
    D. 2 nsec
    E. 2 psec
36. Have you completed your online course evaluations yet or, if not, will you do so by tomorrow's deadline?
    A. Yes