Chapter I:
1) Discuss the characteristics of CPU design that an architect must consider. Include in the discussion the role of design complexity and technology trends. Identify the important technology trends that a designer should remember.
ANSWER: Ch. 1, Part I, pages 3-6.

2) Discuss the impact of the learning curve and volume on cost trends. Describe the components in the IC cost equation, including cost of the die and die yield, and their relationship to cost.
ANSWER Ch. 1, Part I, pages 7-11.

3) List the type of benchmarks available and describe their advantages and disadvantages. Why are benchmark suites better ?
ANSWER Ch. 1, Part II, pages 3-4.

4) What is an important design principle addressed by Amdahl’s law ? Explain Amdahl’s law in words and give and explain the expression for Speedup (which expresses Amdahl’s law). Explain why it expresses the law of diminishing returns.
ANSWER Ch. 1, Part III, pages 1-5.

5) Give the CPU performance equation. Even though it is true that, for example, a 10% improvement in any one of the components of this equation leads to a 10% improvement in overall CPU performance, explain why is it difficult to improve any one of these components in isolation?
ANSWER Ch. 1, Part III, pages 6-8.

6)Write the expression for MIPS. Describe the problems associated with using MIPS as a metric to compare the performance of CPUs.
ANSWER Ch. 1, Part III, page 9.

Chapter II:
1) Describe the various types of instruction set architectures. Why are GPR ISAs so popular ?
ANSWER Ch. 2, Part I, pages 1-3.

2) Describe the four ISA metrics and how they relate to each other and to GPR ISAs.
ANSWER Ch. 2, Part I, page 4.

3) Explain the impact of addressing modes on CPU design, performance and programs. Which addressing modes do you expect to be most expensive with respect to these characteristics ?
ANSWER Ch. 2, Part I, pages 6-8.

4) Identify the important addressing modes and how they are used. What ISA parameter is relevant to two of these addressing modes and why is it important ?
ANSWER Ch. 2, Part I, pages 9-10.

5) Identify instructions that cause a change in control and describe the possible addressing modes. For conditional branches, describe three ways to test the condition.
ANSWER Ch. 2, Part II, pages 3-6.

6) Discuss the trade-offs between variable encoding and fixed encoding.
ANSWER Ch. 2, Part II, page 8.

7) Describe the structure of recent optimizing compilers and four desirable instruction set properties that help the compiler writer.
ANSWER Ch. 2, Part III, pages 2-4.

Chapter III:
1) Describe the problems that must be addressed when converting the non-pipelined DLX to a pipelined machine.
ANSWER Ch. 3, Part II, page 2.

2) Why is it true that pipelining decreases execution time but increases clock cycle time. What feature of instruction irregularity can be taken advantage of in a non-pipelined machine that is not possible in a pipelined machine ?
ANSWER Ch. 3, Part II, pages 3-4.

3) Define a pipeline hazard and name and describe three types of pipeline hazards.
ANSWER Ch. 3, Part II, page 5.

4) Derive the pipeline speedup assuming ideal CPI is 1, the clk cycle time is unchanged (ignore pipeline overhead) and the unpipelined CPI is equal to the depth of the pipeline. Do the same but include pipeline overhead.
ANSWER Ch. 3, Part II, pages 8-9.

5) Describe the two most common structural hazards. What can be done to prevent them from causing stalls ?
ANSWER Ch. 3, Part III, pages 1-3.

6) Give two reasons why we would allow structural hazards.
ANSWER Ch. 3, Part III, page 4.

7) In the standard DLX pipeline, why do instructions that use the result of a previous instruction cause a data hazard ? Describe your solution using a pipeline diagram.
ANSWER Ch. 3, Part III, 5-7.

8) Discuss the three types of data hazards. Describe two methods of avoiding hazards.
ANSWER Ch. 3, Part III, pages 8-15

9) With respect to DLX, explain why a stall is necessary when a load instruction is followed by an ALU instruction that uses its result. What type of instruction following a load does NOT need to stall ?
ANSWER Ch. 3, Part III, pages 16-17.

10) Give an example code sequence that shows how compiler scheduling can eliminate the stall due to a load followed by an immediate use by an ALU instruction. What resource limits the amount of scheduling that can be performed ?
ANSWER Ch. 3, Part IV, pages 1-2.

11) Describe why control hazards occur. Give a pipeline diagram showing the stalls that result from a control hazard under the standard DLX. In general, what must be done to reduce the stalls? Describe the change in the datapath of the DLX that reduces the branch penalty to 1 cycle.
ANSWER Ch. 3, Part IV, pages 5-8

12) Describe four static prediction schemes that can be used to reduce branch penalties.
ANSWER Ch. 3, Part IV, pages 11-13.

13) Explain the method of filling the branch delay slot with an instruction preceding the branch, from the branch target and from the fall through. Which method is always a win ? How does adding a nullifying branch instruction increase the flexibility of filling the branch delay slot?
ANSWER Ch. 3, Part IV, pages 14-15.

14) Explain how having accurate prediction information at compile time can help in scheduling data hazards. Give an example.
ANSWER Ch. 3, Part V, page 4.

15) List four methods that compiler can use to reduce branch stall penalties.
ANSWER Ch. 3, Part V, pages 5-8.

16) Why do exceptions make pipelining difficult ? Describe the five characteristics of exceptions, that can be used to distinguish between them. Which type is the most difficult to deal with ?
ANSWER Ch. 3, Part V, pages 9-13

17) Describe the sequence of events that must occur when an instruction causes a page fault in MEM and must be restarted. For each of the pipeline stages IF, ID, EX and MEM, identify one type of exception that may occur.
ANSWER Ch. 3, Part V, page 14-16.

18) Distinguish between precise and imprecise exceptions. Give an example 2 instruction code sequence and a pipeline diagram that shows an exception occurring “out-of-order”. How can this be solved ?
ANSWER Ch. 3, Part V, pages 15-18.

19) Explain why precise interrupts are easy for the integer DLX in terms of instruction commit. Explain why imprecise interrupts are possible for some CISC machines.
ANSWER Ch. 3, Part V, pages 20-21.

20) Why is it not realistic to require FP units to execute in one clock cycle ? Define latency and initiation interval.
ANSWER Ch. 3, Part VI, pages 1-4.

21) Given the pipeline characteristics (shown on slide 4), identify and explain the additional complications introduced when FP functional units are added to DLX.
ANSWER Ch. 3, Part VI, pages 4-6.

22) Assume only a single write port exists. Give an instruction sequence that causes a stall due to a structural hazard on the write port. Give two solutions to this problem and discuss the advantages and disadvantages of each.
ANSWER Ch. 3, Part VI, pages 8-9.

23) Give an instruction sequence that causes a WAW hazard. Explain why it seems unreasonable for this to occur. Give two solutions to this problem and discuss the advantages and disadvantages of each.
ANSWER Ch. 3, Part VI, page 10.

24) Assuming all hazards are detected in ID in DLX, explain the hazards and the detection logic required to avoid the hazards.
ANSWER Ch. 3, Part VI, pages 11-12.

25) Give a code sequence that shows an imprecise exception using the DLX latencies. Briefly describe the four methods of handling these exceptions.
ANSWER Ch. 3, Part VII, pages 1-7.

26) Discuss four guidelines for designing instruction sets and the complications that arise due to hazards and exceptions if these guidelines are not followed.
ANSWER Ch. 3, Part VII, pages 8-9.

Chapter IV:
1) Define ILP and why it is necessary to exploit ILP across basic block.
ANSWER Ch. 4, Part I, pages 1-6.

2) Explain loop-level parallelism.
ANSWER Ch. 4, Part I, page 7.

3) Define pipeline scheduling. Explain how this is done in the context of a simple loop given on slide 9.
ANSWER Ch. 4, Part I, pages 8-11.

4) Explain how loop unrolling can be used to improve pipeline scheduling.
ANSWER Ch. 4, Part I, pages 12-15.

5) Define the relationship between dependencies, hazards and stalls. Identify three types of dependencies.
ANSWER Ch. 4, Part II, page 1.

6) Define data dependency. Describe two ways to overcome data dependencies. Which method can be applied using hardware and which using software (compiler).
ANSWER Ch. 4, Part II, pages 2-5.

7) Define name dependency. Describe the two types.
ANSWER Ch. 4, Part II, pages 6-8.

8) Define control dependency. Is preserving control dependence critical ? What two properties are critical ?
ANSWER Ch. 4, Part II, pages 9-11.

9) Name four methods that can be used to avoid control stalls.
ANSWER Ch. 4, Part II, page 14.

10) Define dynamic scheduling. List 2 advantages of dynamic scheduling over static scheduling. One method of handling out-of-order execution is to split ID. Explain.
ANSWER Ch. 4, Part III, page 1-3.

11) Describe how scoreboarding works. Include a description of IS, RD, EX and WB with respect to data hazards.
ANSWER Ch. 4, Part III, pages 4-9.

12) Describe how the scoreboard distinguishes between RAW and WAR hazards.
ANSWER Ch. 4, Part III, page 15.

13) Discuss the 4 limitations of scoreboarding.
ANSWER Ch. 4, Part III, pages 16-17.

14) Describe the method used in Tomasulo's approach to eliminate WAW and WAR hazards. In scoreboarding, we had to wait for these to clear. Describe two other differences between scoreboarding and Tomasulo's approach.
ANSWER Ch. 4, Part IV, pages 1-2.

15) Describe how Tomasulo's approach works. Include a description of IS, EX and WB with respect to data hazards.
ANSWER Ch. 4, Part IV, pages 4-7.

16) Discuss the two major advantages of Tomasulo's approach. Describe what makes dynamic loop unrolling and memory disambiguation possible.
ANSWER Ch. 4, Part IV, pages 9-10.

17) Discuss how accuracy and latency of dynamic branch prediction schemes affect performance.
ANSWER Ch. 4, Part V, page 2.

18) Explain the branch prediction buffer scheme. Describe a situation in which a 2-bit predictor can perform better than a single bit predictor. What is the idea behind an n-bit predictor.
ANSWER Ch. 4, Part V, pages 4-6.

19) Describe a 2-level correlating branch predictor. Give a code sequence on which this scheme can provide a performance improvement over other static prediction and a branch prediction buffer.
ANSWER Ch. 4, Part V, pages 8-12.

20) Describe an (m,n) predictor. Suggest why it may be able to achieve better prediction results than a 2-level predictor. Describe how the Global History Buffer might be implemented.
ANSWER Ch. 4, Part V, page 13.

21) Explain why it is possible for a Branch-Target Buffer to reduce stalls even further than it is possible under branch-prediction buffers and correlated branch schemes, possibly to zero.
ANSWER Ch. 4, Part V, page 14.

22) Describe the Branch-Target Buffer scheme. Can this scheme be combined with branch prediction ? Explain why it is not efficient to include the prediction bits in the branch-target buffer cache.
ANSWER Ch. 4, Part V, pages 15-16.

23) Suggest why the branch misprediction penalties may be higher under the branch-target buffer scheme than in other schemes.
ANSWER Ch. 4, Part V, page 17.

24) Explain how branch folding is different from branch-target buffers. How does this effect the way unconditional branches are handled ?
ANSWER Ch. 4, Part V, page 18.

25) Explain the effect of misprediction rate and misprediction penalties on dynamic branch prediction schemes.
ANSWER Ch. 4, Part V, page 19.

26) Explain why more ambitious compiler and hardware scheduling techniques are required for superscalar machines.
ANSWER Ch. 4, Part VI, pages 4-5.

27) Describe the general instruction format for a VLIW machine. Why is it important that the compiler unroll loops and schedule code across basic blocks ?
ANSWER Ch. 4, Part VI, page 8.

28) Explain why there are limits on the number of instructions simultaneously issued by multi-issue processors in terms of latency, memory bandwidth and design complexity.
ANSWER Ch. 4, part VI, pages 10-12.

29) Discuss the limitations specific to superscalar and VLIW machines.
ANSWER Ch. 4, part VI, page 13.