Very Short Answer:

(1) (2) What are the 4 different types of control flow changes?

(2) (1) Is it possible to design a flawless architecture?

(3) (1) Peak performance does or does not track observed performance.

(4) (1) If I am more interested in code size than performance, I will choose which kind of instruction encoding?

(5) (1) Ignoring instruction set issues, What makes pipelining hard to implement?

(6) (1) Is MIPS an accurate measure for comparing performance among computers?

(7) (1) Out of Order completion makes supporting what very difficult?

(8) (1) Are wire delays or transistors more likely to be the most significant limit on clock frequency in the future?

Short Answers:

(9) (2) Most silicon dies are fairly small. Why are they not bigger?

(10) (2) Describe what the geometric mean is, and what the biggest drawback to using it is.

(11) (3) Write down the 3-term CPU performance equation developed in class.

(12) (8) List the 5 classes of benchmarks, give an example for each class, and tell me what the perfect benchmark would be.

(13) (6) There are 3 kinds of Hazards. List and give brief descriptions of all three.

(14) (6) There are 3 kinds of dependencies. List and give brief descriptions of all three.

(15) (10) Compare and contrast Tomasulo's algorithm with Scoreboarding. (Convince me you understand both - in other words, explain what they are, how they work, why they work, how they differ, how they are the same, etc.)

(16) (10) Compare and contrast Superscalar with VLIW. Describe each, and list the advantages and disadvantages of each approach.

(17) (10) Determine the total branch penalty for a branch target buffer assuming there is a 2 cycle penalty if the branch is not in the buffer and the branch is actually taken, and also a 2 cycle penalty if the branch is in the buffer but incorrectly predicted to be taken. Make the following assumptions about the prediction accuracy and hit rate:

60% of all branches are taken.
For instructions in the buffer, the prediction accuracy is 90%
For branches predicted taken, the hit rate in the buffer is 90%

(18) (10) A simple hardware implementation of the basic 5-stage RISC pipeline uses the EX stage hardware to calculate the branch instruction comparison, and does not deliver the target PC to the Instruction Fetch unit until the clock cycle in which the branch instruction reaches the MEM stage. Control hazard stalls can be reduced by adding hardware to resolve the branch condition in the ID stage, but doing so might reduce performance in certain circumstances. How does determining the branch outcome in the ID stage have the potential to increase data hazard stall cycles?

(19) (24) Consider the following three processors:

A) A simple 2-issue machine (Curly) running with a clock rate of 1 GHZ and achieving a pipeline CPI of 1.0. This processor has a cache system that yields 0.01 misses per instruction.

B) A deeply pipelined version of the same machine (Larry) with slightly smaller caches and a 1.2 GHZ clock rate. The pipeline CPI of the processor is 1.2, and the smaller caches yield (on average) 0.015 misses per instruction.

C) A speculative superscalar machine with a 64-entry instruction window (Moe). It achieves 50% of the ideal issue rate of 9 instructions/cycle. This processor has the smallest caches, which leads to 0.02 misses per instruction, but is able to hide 10% of the miss penalty on every miss because it does dynamic scheduling. This processor has an 800 MHZ clock.

Assume that the main memory time (which sets the miss penalty) is 100 ns. Determine the relative performance of these three processors.