(1) (6) What are Benchmarks, and why are they important? List the 4 classes of benchmarks, give an example for each class, and tell me what the perfect benchmark would be.

(2) (8) Convince me you understand Tomasulo's algorithm. (In other words, explain what it is, how it works, why it works, etc.)

(3) (6) List 3 ways to support precise interrupts in an Out Of Order machine.

(4) (3) Describe in words what Amdahl's law says.

(5) (4) What are the advantages and disadvantages of stack machines?

(6) (5) There are 3 kinds of Hazards - Structural, and two others. List the other two, and give brief descriptions of all three.

(7) (5) Why is branch prediction important? What performance enhancing techniques have made it so? Give 1 static and 2 dynamic branch prediction approaches (in order of increasing effectiveness).

(8) (15) Implementations of floating point square root (FPSQR) vary significantly in performance. Suppose FPSQR is responsible for 20% of the execution time of a critical benchmark on a machine. One proposal is to add FPSQR hardware that will speed up this operation by a factor of 10. The other alternative is to make *all* floating point instructions run twice as fast. Floating point instructions account for 50% of the execution time of this benchmark. Compare these two design alternatives.

(9) (20) Which has the lower miss rate: a 16KB instruction cache with a 16KB data cache, or a 32KB unified cache? Assume a hit takes 1 clock cycle and a miss costs 50 clock cycles, and a load or store hit takes 1 extra clock cycle on a unified cache since there is only one cache port to satisfy two simultaneous requests. The Miss Rate for the 16KB instruction cache is .64%, the miss rate for the 16KB data cache is 6.47%, and the miss rate for the 32KB unified cache is 1.99%. Also, what is the average memory access time in each case? Assume write-through caches with a write buffer and ignore stalls due to the write buffer.

(10) (25) Consider the following three hypothetical (but not atypical) processors:

A) A simple DLX machine running with a clock rate of 300 MHZ and achieving a pipeline CPI of 1.1. This processor has a cache system that yields 0.03 misses per instruction.

B) A deeply pipelined version of DLX with slightly smaller caches and a 400 MHZ clock rate. The pipeline CPI of the processor is 1.5, and the smaller caches yield (on average) 0.035 misses per instruction.

C) A speculative superscalar with a 32-entry instruction window. It achieves 75% of the ideal issue rate of 8 instructions/cycle. This processor has the smallest caches, which leads to 0.05 misses per instruction, and has a 200 MHZ clock.

Assume that the main memory time (which sets the miss penalty) is 200 ns. Determine the relative performance of these three processors.