(1) (4) What is the primary reason why very few modern machines have a stack-based instruction set?

(2) (4) List Flynn's 4 different catagories of parallel processors.

(3) (4) List at least 4 different addressing modes, and describe (in words or pictures) where the operand is stored.

(4) (24) For each of the following questions, state whether or not they are true, and give a **brief** justification/explanation for your answer. (Emphasis on **brief**, just so I know you are not guessing.)

1) There is such a thing as a typical program.

2) An architecture with flaws can be successful.

3) It is possible to design a flawless architecture.

4) You can predict cache performance of Program A by analyzing Program B.

5) Amdahl's law doesn't apply to parallel computers.

6) Linear speedups are not needed to make multiprocessors cost-effective.

7) Multiprocessors are "free".

8) Scalability is expensive.

9) The rated MTTF of disks is 140 years, so disks practically never fail.

10) Components fail fast.

11) Computer systems do not achieve 99.999% availability as advertised.

12) Semiconductor memory will soon replace magnetic disks in desktops and servers.

(5) (28) Assume two different machines with the same processor and main memory, but different cache organizations. Assume that both processors run at 4GHz, have a CPI of 1, and have a cache read miss time of 200ns. Further, assume that writing a 32-bit word to main memory requires 200ns and writing a full cache line requires 400 ns. Both organizations feature unified first-level caches, and each cache has a total data capacity of 64KB.

The cache on system A is two-way set associative and has 32-byte blocks. It is write through and does not allocate a block on a write miss.

The cache on system B is direct mapped and also has 32-byte blocks. It is write back, and allocates a block on a write miss.

a) Describe a program that makes system A run as fast as possible relative to system B's speed. (You do not have to write a program, but you do have to make it clear what the program would have to do in order to cause the difference in performance.)

b) How much faster is the program in part (a) on system A compared to system B?

c) Describe a program that makes system B run as fast as possible relative to system A's speed. (You do not have to write a program, but you do have to make it clear what the program would have to do in order to cause the difference in performance.)

d) How much faster is the program in part (b) on system B compared to system A?

(6) (7) As caches increase in size, line sizes often increase as well. Explain the interaction between prefetching and (increasing) line size in instruction caches.

(7) (7) Explain the interaction between prefetching and (increasing) line size in data caches.

(8) (12) One of the design challenges for Simultaneous MultiThreaded processors (SMTs) is ensuring that cache conflicts generated by the simultaneous execution of multiple threads do not cause significant performance degradation (often referred to as destructive cache interference.) However, a simulation study of the online transaction-processing workload, OTLP, on an 8-threaded SMT processor shows a decrease in instruction cache miss rate from 14% to 9%. When is such a decrease, called constructive cache interference, likely to happen?

(9) (10) What is the average time to read or write a 512 byte sector for a disk that has a seek time of 5ms, a transfer rate of 40MB/sec, rotates at 10,000 RPM, and has a controller overhead of .1 ms? What percent of the time is the disk transferring data in this example?