1. **True or False:**

(1) Amdahl’s law doesn’t apply to parallel computers.

(1) It is possible to design a flawless architecture.

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

(1) Multiprocessors are "free".

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

(1) An architecture with flaws cannot be successful.

(1) Linear speedups are needed to make multiprocessors cost-effective.

(1) Scalability is almost free.

2. (3) What is the goal of the memory heirarchy? What two principles make it work?

3. (4) What do the following acronyms stand for:

SMT

COMA

NUMA

CMP

4. (2) In your own words, what does Amdahl's law say?

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

6. (4) List 4 of the 5 Miss Penalty Reduction Techniques.

7. (3) What is a victim cache, and how does it work?

8. (3) What is Cache Coherence, and why is it necessary?

9. (4) Describe the difference between shared memory and message passing machines. Include the impact on design, cost, speed, and programming model.

10. (12) Dopey is a computer which has a CPI of 1.0 when all memory accesses hit in the cache. The only data accesses are loads and stores, which total 30% of the instructions. If the miss penalty is 40 clock cycles and the miss rate is 3%, what is the memory stall time in terms of instruction count?

11. (26) Assume that the CPI of a processor with a perfect cache is 2.0, the clock cycle time is 1 ns, and there are 1.5 memory references per instruction, and the cache has a block size of 64 bytes. System A (Happy) uses a direct mapped cache, while System B (Sneezy) uses a two-way set associative one. Since a set-associative cache requires an extra multiplexor/selector, the cycle time of Sneezy is 1.25 longer than that of Happy. The cache miss penalty is 75ns for both systems.

a.) Calculate the average memory access time and CPU performance for each processor. Assume the hit time is 1 clock cycle, the miss rate of a direct-mapped 64KB cache is 1.4%, and the miss rate for a 64KB two-way set associative cache is 1.0%. Which system is faster?

b.) Now assume that Sneezy is modified to support Out Of Order exection, and also has it's set-associative cache replaced with a direct-mapped one. The clock cycle time stays the same as it did (1.25 times that of Happy). Assume 30% of the 75ns miss penalty can be hidden by overlapping it with other instructions; that is, the average CPU memory stall time is now 52.5ns. How does the performance of Sneezy now compare to that of Happy?

12. (15) Consider the following description of a computer (Grumpy) and its cache performance:

Block size = 1 word
Memory bus width = 1 word
Miss rate = 3%
Memory accesses per instruction = 1.2
Cache miss penalty = 64 cycles
Average cycles per instruction (ignoring cache misses) = 2

If we change the block size to 2 words, the miss rate falls to 2%, and a 4-word block has a miss rate of 1.2%. What is the improvement in performance if memory is interleaved two ways and four ways, versus doubling the width of memory and the bus? Assume it takes 4 clock cycles to send an address, the access time is 56 clock cyles per word, and it takes 4 clock cycles to send a word of data.

13. (12) Assume a directory-based cache coherence protocol. The directory currently has information that indicates that processor P1 has the data in "exclusive" mode. If the directory now gets a request for the same cache block from processor P1, what could this mean? What should the directory controller do?