1. **True or False:**

(1) Semiconductor memory will not replace magnetic disks in desktop and server computer systems in the near future.

(1) Computer systems achieve 99.999% availability ("five nines"), as advertised.

(1) Computer components fail suddenly, with little warning.

(1) Amdahl's law applies to parallel computers.

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

(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. (10) What do the following acronyms stand for:

SMT	SMP	MTTF	SRAM
COMA	ATM	RAID	DIMM
NUMA	MPP	DSM	ROM
CMP	WAN	TLB	CD
MIMD	SIMD	DDR SDRAM	PCI

Short Answers:

4. (3) What is false sharing?

5. (3) What is a cluster?

6. (3) What is the difference between connection-oriented and connectionless communication?

7. (3) Give three techniques used to reduce cache hit time.

8. (4) Explain the interaction between prefetching and increased block size in instruction caches.

9. (4) What is Cache Coherence, and why is it necessary?

10. (4) What is a victim cache, and how does it work?

11. (4) Briefly describe interleaved memory and how it works.

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

13. (12) Given the following data:

Hit time for direct-mapped L2 cache = 10 clock cycles
Local miss rate for direct-mapped L2 cache = 25%
Local miss rate for 2-way set associateive L2 cache = 20%
Miss penalty for L2 cache = 100 clock cycles

What is the impact of second-level cache associativity on the miss penalty?

14. (12) Assume that an L2 cache has a block size four times that of an L1. Show how a miss for an address that causes a replacement in L1 and L2 can lead to a violation of the inclusion property.

15. (12) Given a disk with the following parameters:

Average seek time = 5ms
Transfer rate = 40MB/sec
Rotation speed = 10,000 RPM
Controller overhead = 0.1ms

Assuming there is no queueing delay, what is the average time to read or write a 512-byte sector?

What is the time assuming the average seek time is 3 times that of the measured seek time?

16. (12) In this question you are to calculate the time to transfer 1000GB of data using 25 8mm tapes via an overnight delivery service, vs sending the data by FTP over the internet. Make the following 4 assumptions:

The tapes are picked up at 4PM Pacific time and delivered 3000 miles away at 10AM Eastern time (7AM Pacific)

On electronic route A the slowest link is a T3 line, which transfers at 45M bits/sec

On electronic route B the slowest link is a 100M bit/sec Ethernet

You can use half of the capacity of the slowest link between the two sites

How long does it take to electronically transfer the information via route A?

Route B?

Using Overnight delivery?