(1) **True or False:**

(1) DRAM and Disk access times are rapidly converging.

(1) Symmetric Shared-memory Multiprocessors are by far the most popular organization.

(1) Measuring performance on multiprocessors using linear speedup instead of execution time is a bad idea.

(1) The use of large, multilevel caches can substatially reduce the memory bandwidth demands of a processor.

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

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

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

(1) It is not necessary to simulate very many instructions in order to get an accurate performance measure of the memory heirarchy.

(1) A program's locality behavior is varies over the run of an entire program.

(1) Communication is a significant problem for parallel processor systems.

(1) Coherence and Consistency are complementary.

(1) The instruction set architecture does not impact the implementability of a virtual machine monitor.

(1) Using a write-through cache simplifies the implementation of a cache coherence protocol.

(1) Write-back caches can use the same snooping scheme for both cache reads and writes

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

(3) (6) What do the following acronyms stand for:

OLTP	VMM	TPC
UMA	COMA	DSM

(4) (2) What type of parallelism do SIMD machines exploit?

(5) (2) What type of parallelism do MIMD machines exploit?

(6) (4) What are the four classifications of cache misses?

(7) (1) Which type of cache miss can be changed by altering the mapping scheme?

(8) (1) Which type of cache miss can be reduced by using longer lines?

(9) (1) Which type of cache miss can be increased by using longer lines?

The questions on this page are extremely short answer - write down the appropriate number(s) corresponding to the answers (the list of answers is on the last page of the test.) The first one is done for you.

(10) (0) Taking this test has made my brain **3** (or **3-Hurt**) (or just **Hurt**)

(11) (1) Relaxing the requirement that Writes complete before Reads yields a model called

(12) (1) Relaxing the requirement that Writes complete before Writes yields a model called

(13) (1) Which type of operation is necessary in order to support synchronization?

(14) (1) Which benchmarks are most affected by the window size?

(15) (1) Which benchmarks are most affected by the accuracy of the branch predictor?

(16) (1) Alias analysis has the most impact on which benchmark?

(17) (4) List 4 of the techniques the book says are useful for reducing the Miss Rate.

(18) (4) List 4 of the techniques the book says are useful for reducing the Hit Time.

(19) (4) List 4 of the techniques the book says are useful for reducing the Miss Penalty.

(20) (2) Coherent caches provide _____ and _____ of shared data items.

(21) (1) A _____ protocol invalidates other copies on a write.

(22) The processor with the sole copy of a cache block is normally called the _____ of the cache block.

(23) (1) What hardware structure is necessary in order to make a snooping protocol work?

(24) (1) Often the decision is made to snoop only on the tags in the L2 cache. This requires the L1 cache entries to follow the _____ property.

Very Short Answers:

(25) (2) The authors claim there are two main reasons MIMD machines are so popular. What are they?

(26) (2) Give an example of a SIMD type of machine.

(27) (2) Why have SIMD architectures experienced a rebirth in popularity?

(28) (4) How does instruction-level differ from thread-level parallelism?

(29) (4) What are the two biggest challenges in parallel processing? In other words, what two things are keeping parallel processors from being the dominant architecture?

(30) (4) What pair of instructions are used to implement a lock in RISC systems (as described in the text)? Describe how this pair works together in order to accomplish the goal.

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

(32) (8) Snooping is one main approach to providing coherence - state what the other main approach is, and briefly outline how each of them work.

(33) (4) Draw the basic structure of a centralized shared-memory processor

(34) (4) Draw the basic structure of a distributed memory processor

Here are answers to go with the questions on pages 3 and 4. Feel free to tear this page off of your test.

(1) Bigger
(2) Smaller
(3) Hurt
(4) Total Store Ordering
(5) Partial Store Ordering
(6) Weak Ordering
(7) Integer
(8) Float
(9) Equally affected
(10) Atomic Synchronization Primitive
(11) Implement Synchronization
(12) Modified
(13) Unmodified
(14) Shared
(15) Unshared
(16) Exclusive
(17) Valid
(18) Invalid
(19) Expensive
(20) Owner
(21) Write Validate
(22) Write Invalidate
(23) Read Validate
(24) Read Invalidate
(25) Broadcast Medium
(26) Nuclear
(27) Atomic
(28) Small and simple caches
(29) Larger block size
(30) Bigger caches
(31) Way prediction
(32) Trace caches
(33) Higher associativity
(34) Multilevel caches
(35) Pipelined caches
(36) Non-blocking caches
(37) Multibanked caches
(38) Compiler optimizations
(39) Victim cache
(40) Priority to Read Misses
(41) Critical Word First/Early Restart
(42) Merging Write Buffer
(43) Avoiding Address Translation when Indexing Cache
(44) Hardware Prefetching
(45) Software prefetching
(46) Small and simple caches
(47) Data Level Parallelism
(48) Instruction Level Parallelism
(49) Thread Level Parallelism
(50) Multithreading
(51) Migration
(52) Replication
(53) Broadcast Medium
(54) Owner
(55) Renter
(56) Inclusion
(57) Exclusion
(58) Buzz word A
(59) Other Buzz Words