1. Which of the following statements is **TRUE?**
   a. OpenMP programs cannot run on UMA multiprocessors with snoopy cache coherence hardware.
   b. OpenMP programs cannot run on NUMA multiprocessors with directory-based cache coherence.
   c. Both "a" and "b" statements above are false.
   d. Both "a" and "b" statements above are true.
   e. OpenMP can only be run on centralized memory multiprocessors, while MPI can only be run on distributed memory multiprocessors.

2. Which of the following statements about the MSI coherence protocol is **TRUE?**
   a. The Invalid state means that the cache block has not been read or written by the CPU.
   b. The Shared state means that the cache block is in more than one cache.
   c. The Shared state means that the cache block has been read in the cache but has not been modified.
   d. The Exclusive state means that the cache block has been read in one and only one cache.
   e. None of the above is TRUE.

3. A CPU write hits a block in shared state and the block is in the cache of another processor. Which of the following statements is **TRUE?**
   a. The other processor writes back its block and invalidates the block in its cache.
   b. The other processor invalidates the block in its cache without writing the block back.
   c. The other processor does nothing.
   d. The CPU writes its cache without sending any bus request since it is a hit to the cache.
   e. None of the above.

4. A CPU write misses and the block is in exclusive-modified state in another processor. Which of the following statements is **TRUE?**
   a. The other processor writes back its block and invalidates the block in its cache.
   b. The other processor writes back the block and changes its state to Shared.
   c. The other processor does nothing.
   d. The CPU that misses sends an invalidate bus request.
   e. None of the above.

5. A CPU read misses a block in shared state and the block is in shared state in the cache of another processor. Which of the following statements is **TRUE?**
   a. The other processor writes back its block and invalidates the block in its cache.
   b. The other processor writes back its block without invalidating the block in its cache.
   c. The other processor invalidates its block without writing it back.
   d. The other processor does nothing.
   e. None of the above.

6. In a multiprocessor system that uses directory-based cache coherence, a CPU write misses a block in shared state and the miss block is in shared state in two other caches. Assuming that the block home directory is in a fourth node. Which of the following statements is **TRUE?**
   a. There will be 2 cache coherence messages.
   b. There will be 3 cache coherence messages.
   c. There will be 4 cache coherence messages.
   d. There will be 5 cache coherence messages.
   e. None of the above.

7. A shared memory distributed multiprocessor computer consists of 8 nodes. Each node contains 2-way set associative cache with total capacity 16K bytes and block size of 32 bytes. Each node contains 512 Kbytes (2^{19} bytes) of DRAM and a directory for maintaining directory-based cache coherence. Which of the following statements is **TRUE**?
   a. Each node directory consists of 256 entries.
   b. Each node directory consists of 512 entries.
   c. Each node directory consists of 16K entries.
   d. Each node directory consists of 64K entries.
   e. Directory size depends on the number of tag bits it uses.

8. A cache has a total capacity of 16K bytes. It is implemented as 2-way set associative, with block size of 32 bytes. The physical address on the machine consists of 32 bits. Which of the following statements is **TRUE**?
   a. Number of set index bits = 7 and number of tag bits = 20
   b. Number of set index bits = 8 and number of tag bits = 19
   c. Number of set index bits = 9 and number of tag bits = 18
   d. Number of set index bits = 8 and number of tag bits = 20
   e. None of the above is TRUE.

9. Which of the following statements about memory consistency is **TRUE**:
   a. The memory consistency model defines how loads from one processor can be ordered relative to stores from a different processor in a shared memory system.
   b. The memory consistency model does not allow a processor to reorder its loads and stores to global memory.
   c. Memory consistency is something that only programmers need to be concerned about to write correct programs.
   d. The memory consistency model is determined by the cache coherence hardware.
   e. None of the above statements is TRUE.

10. Which of the following statements about processor consistency is **TRUE**:
   a. Processor consistency does not allow a CPU to issue loads to memory according to sequential order.
   b. Processor consistency allows a CPU to issue stores to memory in any order.
   c. Processor consistency allows a CPU to issue loads and stores to memory in any order.
   d. Processor consistency does not allow loads from different CPUs to issue to memory in arbitrary order.
   e. Processor consistency does not allow stores from different CPUs to issue to memory in arbitrary order.

11. Which of the following statements about weak memory consistency is **FALSE**:
   a. All memory operations from the CPU must complete before a fence instruction later in the program executes.
   b. The memory operations from the CPU cannot issue to memory until an earlier fence in the program executes.
   c. All memory operations from all the CPUs must complete before the fence executes.
   d. Fence instructions from the same CPU cannot be reordered.
   e. Two fence instructions from two different CPUs could be in any order.

12. Which of the following statements about weak memory ordering is **FALSE**?
   a. All memory operations from the CPU must complete before a later fence instruction executes.
   b. The memory operations from the CPU cannot issue to memory until an earlier fence executes.
   c. Stores from the same CPU must be issued in program order.
   d. Fence instructions from the same CPU cannot be reordered.
   e. Loads between two fences can be issued in any order.

13. Which of the following statements about memory consistency in shared memory multiprocessors is **TRUE?**
   a. Memory consistency defines how loads are ordered relative to loads and stores from the same thread.
   b. Memory consistency defines how loads from one thread are ordered relative to stores from another thread.
   c. Memory consistency defines how loads and stores from one thread are ordered relative to loads and stores from another thread.
   d. Memory consistency ensures that loads and stores from different threads to a shared variable in memory are properly synchronized.
   e. None of the above.

14. Which of the following statements about memory consistency is **FALSE?**

   We need to know about memory consistency in order to:
   a. Write correct shared memory parallel programs.
   b. Write correct message passing parallel programs.
   c. Design distributed shared memory multiprocessors hardware.
   d. Understand performance of OpenMP parallel programs.
   e. Design various cache optimizations correctly.

15. Which of the following hardware mechanisms is not needed to implement processor consistency model?
   a. A mechanism to order stores to memory.
   b. A mechanism to prevent overlapping memory writes.
   c. A mechanism to prevent late cache invalidates.
   d. A mechanism to prevent non-blocking memory reads.
   e. A load buffer to snoop stores of other threads.

16. Which of the following statement about SIMD computers is **False?**
   a. All processing units execute the same instruction at any given clock cycle
   b. Each processing unit can operate on a different data element.
   c. Vector processors are examples of SIMD architectures.
   d. SIMD is rarely used in modern multicore processors.
   e. One of the above is False.

17. A computer features functional units that operate on data via multiple instruction streams. If one conceivable use of this computer is to apply multiple cryptography algorithms to crack a single coded message, which of the following statements is **TRUE?**

    a. This computer must be a SIMD computer.
    b. This computer must be message passing computer.
    c. This computer is MISD or MIMD computer.
    d. This computer must be a MISD computer.
    e. None of the above.

18. Which of the following computers does not exploit parallel execution?

    a. SIMD computer.
    b. UMA multiprocessor.
    c. NUMA multiprocessor.
    d. SISD computer.
    e. MISD computer.

19. Which of the following statements is **False**?

    a. OpenMP specifies nothing about parallel I/O.
    b. In OpenMP, a parallel region is a block of code that will be executed by multiple threads.
    c. In OpenMP, a parallel region is a block of code that does not contain any shared data objects.
    d. It is illegal to branch out of a parallel region.
    e. In OpenMP, work-sharing constructs do not launch new threads

20. Which of the following statements about Message Passing computers is **FALSE**:

    a. A processor can only access local memory directly.
    b. Message Passing computers are non-uniform access computers (NUMA).
    c. Remote memory is accessed by passing messages.
    d. MPI programming environment is used with Message Passing computers.
    e. Message Passing Computers can be classified as Multiple Instruction Multiple Data computers.

21. A data cache uses MSI snoopy cache coherence protocol. A store instruction from the same CPU hits a block in exclusive state. Which of the following statements is **TRUE?**

    a. The other processors invalidate their blocks that have the same address.
    b. This processor writes its block and places an invalidate transaction on the external bus.
    c. This processor writes back the block to memory.
    d. All of the above.
    e. None of the above.

22. A cache has a total capacity of 32K bytes. It is implemented as 4-way set associative, with block size of 32 bytes. The physical address on the machine consists of 32 bits. Which of the following statements is **TRUE**?

    a. Number of byte offset bits = 6, number of set index bits = 8, number of tag bits = 19
    b. Number of byte offset bits = 5, number of set index bits = 8, number of tag bits = 19
    c. Number of byte offset bits = 5, number of set index bits = 9, number of tag bits = 18
    d. Number of byte offset bits = 5, number of set index bits = 7, number of tag bits = 20
    e. Number of byte offset bits = 2, number of set index bits = 8, number of tag bits = 19

23. Which of the following statements about parallel programming is **FALSE**:

    a. You can have multiple parallel regions in an OpenMP program.
    b. OpenMP can run on shared bus multiprocessor that uses snoopy cache coherence.
    c. Both OpenMP and message passing programs require thread synchronization to access shared data.
    d. A programmer can write a set of library functions to send and receive messages between OpenMP threads.
    e. You need cache coherence hardware to execute OpenMP programs efficiently.

24. Which of the following statements about parallel programing is **FALSE**?

    Parallel overhead refers to:

    a. The amount of time required for coordinating parallel tasks.
    b. The extra cost of the parallel hardware.
    c. The extra time of starting and terminating tasks.
    d. The extra time of data synchronization.
    e. The execution time of parallel programing language libraries and operating system.

25. Which of the following statements about parallel computers is **FALSE**?

In shared memory parallel computers:

a. Multiple processors can operate independently but share the same memory resources.
b. Changes in a memory location effected by one processor are visible to all other processors.
c. There is no concept of global address space across all processors.
d. Computers are classified as ***UMA*** and ***NUMA***, based upon memory access times.
e. Symmetric Multiprocessor (SMP) machines are shared memory parallel computers.

26. Which of the following statements about shared memory parallel computers is **FALSE**?

a. Global address space provides a user-friendly programming perspective to memory.
b. Data sharing between tasks can be fast due to data migration and replication in the caches.
c. Adding more CPUs can geometrically increase traffic on the shared memory-CPU path.
d. Easy to program since programmer is not responsible for synchronization to ensure "correct" access to memory.
e. If caches are featured, they are cache coherent systems.

27. Which of the following statements about distributed memory parallel computers is **FALSE**?

a. Distributed memory systems require a communication network to connect inter-processor memory.
b. Since they have distributed memory, they require directory based cache coherence hardware.
c. Each Processor has its own local memory.
d. The programmer is required to explicitly define how and when data is communicated.
e. Changes a processor makes to its local memory have no effect on the memory of other processors.

28. Which of the following statements about programming with OpenMP is **TRUE**?

a. Critical Directive is a “parallel region” directive.
b. Critical Directive is a “work sharing” directive.
c. Critical Directive is a synchronization directive.
d. Critical directive is a scheduling directive.
e. None of the above is TRUE.

29. An application is 80% parallel and 20% serial. If a multicore computer with 4 cores is used to parallelize this application, the maximum overall speedup will be:

a. 2.
b. 2.5.
c. 3.
d. 3.5.
e. 4.

30. An application is 80% parallel and 20% serial. If a parallel computer is used to parallelize this application, the maximum possible overall speedup will be:

a. 2.
b. 2.5.
c. 3.
d. 4.
e. 5.