

1. Which of the following does *not* describe a way in which the MIC processor differs from the host CPUs on Stampede?
  - a. There is no L3 cache on the MIC coprocessors on Stampede
  - b. Host CPUs do not implement out of order instruction processing
  - c. Each MIC core is capable of executing more threads in hardware than a corresponding CPU core
  - d. Each MIC core has a wider vector unit than a corresponding host CPU core
  
2. The maximum vector width used by the SIMD vector units on the MIC coprocessors currently installed on Stampede?
  - a. 245-bit
  - b. 512-bit
  - c. 1024-bit
  
3. True or false: Running code on hardware that has double the vector width will always double an application's performance. This is why the MIC performance is usually twice that of the host CPUs.
  - True
  - False
  
4. True or False: Because the MIC coprocessors are based on the x86 architecture, executables that are compiled for the host CPU will also run on the MIC coprocessor, but not the other way around.
  - True
  - False
  
5. Which of the following helps explain why it is good to ensure that each MIC core executes at least two threads or processes concurrently?
  - a. Instructions from given thread can only execute every other clock cycle, using at most 50% of a core's capacity.
  - b. Because each MIC core processes instructions out of order, increasing the thread count gives the core more opportunity to avoid instruction pipeline stalls.
  - c. Because each MIC core contains two vector units, this avoids leaving one of them idle.
  - d. It is faster for two threads to share a core's local L2 cache rather than rely on the MIC's cache coherency to pass data between the L2 caches of different cores.
  
6. What does it mean when we say that the MIC's L2 cache is coherent?
  - a. The L2 cache is as fast as the L1 cache.
  - b. No two cores can have different values for the same segment of memory in their L2 caches.
  - c. Data read into or out of the L2 cache into memory is accessed in a unit-stride pattern.

7. Which of the following is *not* true with compiler-assisted offload
- a. Source code directives may be used to identify sections of code to offload to a coprocessor
  - b. Compiler flags may be used to automatically identify sections of code that are good candidates to offload
  - c. Code blocks intended for offload may contain OpenMP threading directives
8. Which of the following best describes how offloading works at runtime.
- a. A "fat binary" contains both host code and offload code. The system sends offload code and data to/from the coprocessor as necessary when an offload section is encountered.
  - b. Two versions of the executable are created, one for the host, and one for the MIC. When the program is run, both programs run in their respective environment and transfer data as necessary.
9. What type of thread affinity spreads out threads across cores, but also guarantees that if multiple threads execute on the same core, they will have consecutive thread number?
- a. compact
  - b. scatter
  - c. balanced
10. Which of the following programming paradigms allows the user to balance the workload on MIC coprocessors by adjusting the number of processes the MIC, as well as the number of threads per process?
- a. Offload
  - b. Hybrid
  - c. Distributed