

Test Question 1. (1 point per blank) For each of the following characteristics, which style of parallel architecture—shared memory (S), message passing (M), or data parallel (D)—has most of it, and which has least/fewest? In some cases, more than one answer may be correct.

In each of the blanks below, write “S”, “M”, and/or “D”.

Characteristic	Which style has most of these?	Which style has least of these?
Remote memory references		
Interprocessor messages		
Threads within a single process		
Processors (full CPUs)		
Need for lock operations		
Copying of words used by multiple threads or processes		
Index registers shared by multiple processors or PEs		
Network interfaces integrated with memory controller		
Special instructions to move data between memories of different processors		
Nonblocking sends implemented by DMA		

Test Question 2. We have a problem consisting of $4/5$ parallelizable work and $1/5$ serial work.

(a) (4 points) What is the largest amount of speedup that can be obtained (if we can arbitrarily increase the number of processors)?

(b) (5 points) For what value of N is the speedup half of the answer to part (a)?

(c) (4 points) What would be the efficiency if 20 processors were deployed to work on the problem?

(d) (7 points) Suppose that a program is composed of a serial phase and a parallel phase.

- The whole program runs for 1 time unit.
- The serial phase runs for time s , and the parallel phase for time $1-s$.

(i) What is the efficiency of the computation in terms of N , the number of processors used?

(ii) What is the limit of this efficiency as $p \rightarrow \infty$?

Homework Problem 1. (25 points) Assume we have an MIMD computer system with 16 independent processors accessing a shared memory through an interconnection network. Let the time to perform one add operation on one processor be T_a . This means the system can perform 16 additions in time T_a , 32 in time $2T_a$, and so on. In the following parts, ignore the time necessary for memory access, communication, and synchronization.

(a) Plot the time required to perform n additions, where n ranges from 1 to 64. Mark time along the vertical axis, and n along the horizontal. Mark at least 4 points on the graph, and explicitly specify these points in a table of values of n against the corresponding values of time.

(b) Plot the speedup S_n resulting from the use of the 16-processor system as compared with a one-processor system that uses a processor of the same type, for n additions, as in part (a) above. Mark speedup S_n along the vertical axis and n along the horizontal. Mark at least 4 points and explicitly specify these in a table. Of values of n against corresponding values of time and the speedup S_n .

(c) Now consider a more specific problem. We want to sum the elements of a vector of length 64. That is, form a single number by adding all the values from a 1-dimensional array with 64 elements. If a binary-tree-structured algorithm is used (that is, by splitting the set of numbers to be added into two disjoint sets of equal size, and adding corresponding elements of each set), how much time is required for the entire addition process? Note that data dependencies have to be taken into consideration for this algorithm.