

Homework/Lab 6

Practice Only, No Submission Required

You need not submit your answers. These problems are provided as practice with parallel algorithms. You should make sure you understand these problems before you take the final exam.

1. Show the steps needed to do a bitonic sort on this initial sequence. The final sequence should be in non-decreasing order.

8 42 23 1 44 16 3 10 4 5 99 49 77 7 14 35

2. How many processors can the bitonic sort make use of in the previous example?
3. We looked at two parallel formulations of Dijkstra's Algorithm for finding all-pairs shortest paths. Which formulation would likely be better for a network of processors connected by a slow network? Why?
4. Message packing, the grouping of a number of small messages into a single larger message, can improve the performance of a message passing application. Why?
5. We have seen the IBM SP architecture (clusters of nodes containing one or more processors connected by a high-speed switch network) including the largest examples today: ASCI Blue Pacific and ASCI White at Lawrence Livermore National Labs. The IBM Blue Gene project, first proposed in December 1999, is the next generation of this architecture. Its goal is to achieve a the petaflop (one quadrillion operations per second).

Here is IBM's description of the computer: "Blue Gene will consist of more than one million processors, each capable of one billion operations per second (1 gigaflop). Thirty-two of these ultra-fast processors will be placed on a single chip (32 gigaflops). A compact two-foot by two-foot board containing 64 of these chips will be capable of 2 teraflops, making it as powerful as the 8000-square foot ASCI computers. Eight of these boards will be placed in 6-foot-high racks (16 teraflops), and the final machine (less than 2000 sq. ft.) will consist of 64 racks linked together to achieve the one petaflop performance."

Suppose you are developing software to run on today's supercomputers (consider the ASCI-class computers we discussed) but you would also like to be sure your software will be able to run efficiently on Blue Gene when it becomes available. What factors would influence your software design and how would you approach the problem? Do you think it is reasonable to expect software which is efficient on today's supercomputers to be efficient on the next generation of supercomputers without significant modification?