

CS 566 Parallel Processing, Fall 2012

Programming Assignment 1

Assigned: Sept 19, 2012

Due: Oct 9, 3:30pm

1 Goals

This assignment will introduce you to run very simple MPI-2 programs (use C) on the ARGO cluster. You will familiarize yourself with simple MPI-2 calls, and how to execute programs on clusters. Your goal is to create a large array A of n numbers at one node (let us call it the *source node*), distribute blocks of the array across the processors, and find the global maximum number of A .

Your goal is to solve this FIND_MAX problem using at least these two different logical topologies:

1. a ring of k processors
2. a hypercube of $k = 2^d$ processors

You can also try 2-D mesh with wraparound, and a tree.

Please observe the following points.

1. You may work in pairs if you so wish.
2. Remember to start early, because ARGO is a shared machine and there are users from all over campus. You submit jobs in batch mode, and may not get the result immediately. (During bad times, it takes a day or more for completion of even small jobs).
3. First get the basic programs correctly running. Then you can experiment with varying parameters, and more interesting variants of decomposing the workload and of MPI primitives.

2 Experiment

1. For each logical topology, write running code for each of the parallel formulations you design. You may use divide and conquer, or any other strategy for splitting your input array of numbers. Similarly for collecting and aggregating the partial results into a single result.

You are encouraged to experiment with multiple approaches for each topology. For example, for the ring, you can collect the locally computed maximums directly at the source. Or while collecting the local maximums, each intermediate node can calculate the maximum of the maximum values it receives and its own local maximum. Using different approaches will allow you to use different MPI primitives.

2. Measure the *timings* for each of your approaches.
If time permits, you can experiment with different values of n and k . The actual parameter ranges you vary will depend on various factors such as the workload on and the availability of the ARGO nodes, OS limits etc..

3 Input and Testing

1. The program will take as inputs: n , k , p . Here, p is the number of physical processors you use on ARGO. Whereas k is the number of logical processors. (In practice, you may be constrained by the *sysadmin policy* which does not allocate more than a handful of physical processors.) When $k > p$,

multiple logical processors need to be mapped to a physical processor. The array A is created by one process only. You can use a random number generator or any other mechanism.

2. While you debug your code, please work with small data sets, small number of logical nodes, and few physical processors i.e., small n , small k , small p .

4 Output

1. Submit a **hard copy of all the files including the code file**.
2. Submit a **hard-copy report describing your experiment**. The report must have the following sections/information.
 - (a) **Formulations:** Describe each of the formulations of your parallel FIND_MAX in about 1 page each. List clearly the MPI calls you used in the implementation of each formulation.
 - (b) **Parameter ranges:** For each of the formulations, describe the ranges of the parameters you experimented with.
 - (c) **Results:** Give the tables and plot the graphs showing the *timing* variations for each of the formulations, for the ranges of parameters you used. Remember to use appropriate scales for the graphs.
 - (d) **Analysis:** Analyse the results. Give your interpretation and reading of the results for each formulation. You should address questions such as the following:
 - Which formulation is better? under what (or all) circumstances? Why?
 - Can you measure the computation time versus communication time? How do the formulations compare with each other with respect to this breakup of the time overhead?Explain why you observe what you observe. In particular, any anomalous observations should be explained.
 - (e) **Lessons:** What insights you learned from this assignment.

5 Grading

The problem is reasonably well-formulated but the experimental approach is open-ended.

A non-running program may get zero credit. Your assignment will be judged on how *comprehensively* and *methodically* you have designed and run the experiment, and reported on the results. Your insights into the analysis of the results will also be considered in judging the assignment.

6 Reference Material/Chapters

Chapters 4 and 6 from the text. The sample MPI programs of the ACCC web site, and any other source you can find on the web.