CS 5334 Spring 2016 Name ___________________________

Shirley Moore, Instructor

Homework due March 31

Analysis of Collective Communication Costs

This assignment involves both group and individual work. For the group part, you will be assigned to a group to carry out the analysis of a specific collective communication operation. Each group will give a brief (< 10 minutes) presentation to the class on the due date describing their operation and its analysis. For the individual part, you should do your own work, although you may discuss the problems with other class members.

Group Part:

For your assigned operation, you should prepare a presentation on the following:

- An explanation of the semantics of the operation and a description of an efficient algorithm on a hypercube architecture
- A detailed analysis of the communication costs of the algorithm
- A theoretical (as opposed to experimental) analysis of how the algorithm scales
  - with increasing message size while holding the number of processors constant
  - with increasing numbers of processors while holding the message size constant

The collective operations we will analyze (one per group) are the following:

1) One-to-all broadcast and all-to-one reduction (Module 3.1)
2) All-to-all broadcast (Module 3.2)
3) All-to-all reduction (Module 3.2)
4) All-reduce and prefix sum (Module 3.2)
5) Gather and scatter (Module 3.3)
6) All-to-all personalized communication and optimal one-to-all broadcast and all-to-one reduction (Module 3.4)

Individual Part:

1. Design an optimal k-to-k broadcast operation for a hypercube network. For this operation, you are given two disjoint sets of k processors (2k < p) within the ensemble of p processors. Initially, each of the k processors in the first group has a unique data item of size m (the size is identical across all k processors). At the end of the operation, each of the k processors in the second group must have ALL of the data from **<u>all</u>** processors in the first group of k processors (i.e., at the end of the operation, each of the k processors in the second group must have the same k x m data items, from all processors in the first group). Give a detailed description of your algorithm. Derive the runtime for your algorithm and justify its optimality.

2. You are given an n x n matrix that is block-row partitioned across p processors. That is, each processor has a block of n/p rows. What is the optimal parallel time to perform the transpose of the matrix on a hypercube? The transposed matrix should also be block-row partitioned after the operation. Include both local operations and communication operations in the runtime expression.