CSE Ph.D. Qualifying Exam, Fall 2015: High-Performance Computing

Answer 3 out of 4 of the following questions. If you choose to answer all 4, your grade will be based on the questions with the 3 *lowest* scores.

Please write clearly and concisely, explain your reasoning, and show all work.

Note: Unless otherwise specified, assume a distributed memory model of computation where a message of m words can be sent in $O(\tau + \mu m)$ time, where τ denotes the latency and μ denotes the per word transfer time. You may assume each processor can send and receive one message within the same parallel communication step.

1. Assume that n is a power of 2. Let $A = a_0, a_1, \ldots, a_{n-1}$ and $B = b_0, b_1, \ldots, b_{n-1}$ be two sorted lists. Let

$$even(A) = a_0, a_2, a_4, \ldots, a_{n-2} \quad odd(A) = a_1, a_3, a_5, \ldots, a_{n-1}$$
$$even(B) = b_0, b_2, b_4, \ldots, b_{n-2} \quad odd(B) = b_1, b_3, b_5, \ldots, b_{n-1}$$

Let $C = c_0, c_1, \ldots, c_{n-1}$ be the result of merging $even(A)$ and $even(B)$ and $D = d_0, d_1, \ldots, d_{n-1}$ be the result of merging $odd(A)$ and $odd(B)$. Let $X = x_o, x_1, \ldots, x_{2n-1}$ be the result of merging A and B. Prove that

   (a) $x_0 = c_0$ and $x_{2n-1} = d_{n-1}$. (10%)

   (b) $x_{2i-1} = min(c_i, d_{i-1})$ and $x_{2i} = max(c_i, d_{i-1})$ for $1 \leq i \leq n-1$. (40%)

   (c) A comparator is a hardware circuit that takes two inputs and produces two outputs. The upper output is the minimum of the two inputs, and the lower output is the maximum of the two inputs. Using comparators, draw a parallel circuit to merge two sorted lists of 8 elements each exploiting the above properties. (50%)
   (**Note:** You can attempt this problem assuming (a) and (b) even if you are not able to prove them.)

2. (a) Let A and B be two arrays of size n distributed across p processors such that each processor has $\frac{n}{p}$ consecutive entries. Suppose array B contains a permutation of $0, 1, 2, \ldots, n-1$. We need to compute array C of size n such that $C[i] = A[B[i]]$. Which MPI communication primitives should be used here and what is the estimated run-time? (40%)

   (b) An undirected graph G consists of n vertices and m edges. The graph is described by an unordered list of m edge tuples – each of the form (u, v) to describe the edge connecting u and v. Assume the edge tuples are distributed evenly among p processors. Design a parallel algorithm to compute the degree of each vertex in the graph. (60%)

3. Consider a system of n particles where the i^{th} particle is located at position $\vec{r_i}$ and has mass m_i. The gravitational potentional at particle i due to j is given by the formula

$$V_{ij} = -\frac{G m_i m_j}{(||\vec{r_i} - \vec{r_j}||)}$$

The multibody problem is to compute the total potentional V_i for each particle using

$$V_i = \sum_{j=0, j \neq i}^{n} V_{ij}$$

(a) Design a parallel algorithm to solve this problem and analyze its run-time. (80%)

(b) What is the largest number of processors that can be used while being efficient? (20%)

4. Your friend developed a parallel architecture where each processor with rank i (i.e., P_i) can simultaneously communicate (send or receive) with the $\log p$ processors whose ranks differ from i in one bit position. For example, in a 8 processor system, processor 000 can send three different messages to 100, 010, and 001 simultaneously. The run time for sending $\log p$ such messages each of size k is estimated as $O(\tau + \mu k)$. Design a Broadcast algorithm for this machine such that a large message m can be broadcast in $O(\tau logp + \mu m)$ time. Assume p is a power of 2 and m is a multiple of $\log p$.