CZ4102 High Performance Computing

Assignment 2

Due on 20 October 2006, 11am.

1. Derive the number of different ways that a d-dimensional hypercube can be labeled. **(5 marks)**

2. The distance between nodes u and v in a hypercube is the length of the shortest path from u to v. Given a d-dimensional hypercube and a designated source node s, how many nodes are of distance i from the node s where $0 \leq i \leq d$? Derive your answer. **(5 marks)**

3. The dual of all-to-all broadcast is all-to-all reduction, in which each node is the destination of an all-to-one reduction. For example, consider the scenario where p nodes have a vector of p elements each, and the ith node (for all i such that $0 \leq i < p$) gets the sum of the ith elements of all the vectors. Propose a good parallel algorithm to perform all-to-all reduction on a hypercube with addition as the associative operator. If each message contains m words and t_{add} is the time to perform one addition, how much time does your algorithm take (in terms of m, p, t_{add}, t_s and t_w)? **(15 marks)**

   Hint: In all-to-all broadcast, each node starts with a single message and collects p such messages by the end of the operation. In all-to-all reduction, each node starts with p distinct messages but ends up with a single message.

4. The string matching problem is to find all occurrences of a particular substring called the pattern, in another string called the text. Propose a good parallel algorithm to solve the string matching problem with the use of **split**, **send**, **receive** and **gather** functions, and use diagrams in your illustrations. MPI program is not needed. You have to assume the existence of multiple matches, i.e., the substring can appear more than one time in the text. **(15 marks)**

5. Assume n is a multiple of p where $p \leq n$.

$$\sum_{i=1}^{n} i = \sum_{i=1}^{n/p} i + \sum_{i=(n/p)+1}^{2n/p} i + \sum_{i=(2n/p)+1}^{3n/p} i + \ldots + \sum_{i=((p-1)n/p)+1}^{n} i$$

   If p is the number of processors labeled as 0, 1, …, p-1, write a clear MPI program making use of these processors to compute the value of the summation when the value of n is to be indicated at the command line. Your program should also cater for the case where n is not a multiple of p. Processor 0 should perform a sum reduction and print the reduced value on the screen, and also print $n(n+1)/2$ for the ease of verification.
   (30 marks)

P.T.O. for Question 6.

6. The value of the definite integral $\int_0^1 \frac{4}{1 + x^2}\, dx$ is π . We can use numerical integration to compute π by approximating the area under the curve. A simple way to do this is called the rectangle rules as shown below.

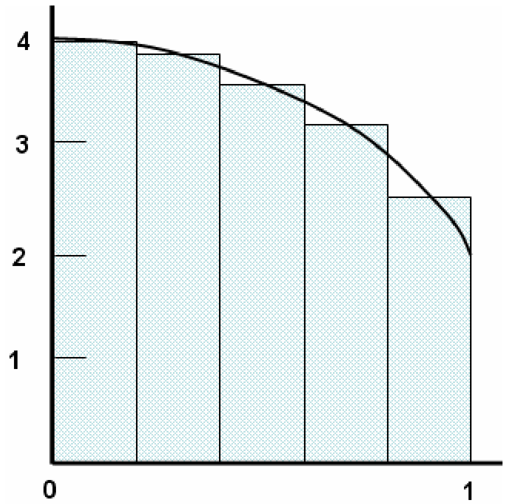

In this method we divide the interval [0, 1] into k subintervals (k = 5 in the above figure) of equal size. Subsequently we find the height of the curve at the mid point of each of these subintervals. With these heights we can construct k rectangles. The area of the rectangles approximates the area under the curve. As k increases, the accuracy of the estimate also increases.

Write a clear MPI program to compute the value of π using the rectangle rule with 1,000,000 intervals. No command-line input is required. Benchmark your MPI program with the following table:

No. of Processors (p)	1	4	8	12	16	20
Elapsed Time (μ sec)						
Speedup ($\frac{T_1}{T_p}$)	-					

(30 marks)

Hand in your source codes named as **string.c** and **pi.c** respectively for Questions 5 and 6 by email attachment to scitaysc@nus.edu.sg. Comment your name on the first line of source files. Also print the source codes in hardcopies together with the answers to Questions 1 to 4.

(This individual assignment carries 15% of the final score.
Please hand in your work on time. We practice the honour system.)