Mastermath midterm examination
Parallel Algorithms. Retake.

Teacher: Rob H. Bisseling, Utrecht University

December 20, 2017

Each of the four questions is worth 10 points. Total time 90 minutes.

1. Explain the structure of a BSP algorithm.

2. Assume we have a set of p vectors $\mathbf{x}^0, \ldots, \mathbf{x}^{p-1}$. Each vector has length n, with $n \bmod p = 0$, and is distributed by the block distribution over the p processors of a parallel computer. (So every processor has a part of every vector.) Assume that the synchronisation cost of the computer is relatively low.

   (a) [5 pt] Give an efficient BSP algorithm for processor $P(s)$ for the computation of the norms of these vectors, in the notation we have learned. Here, the norm is the Euclidean norm defined by

   $$||\mathbf{x}|| = (\sum_{i=0}^{n-1} x_i^2)^{\frac{1}{2}}.$$

   On output, all processors need to know all the norms.

   (b) [3 pt] Analyse the BSP cost.

   (c) [2 pt] What would you change if the synchronisation cost would be very high? Where is the break-even point?

3. Let p be the number of processors of a parallel computer and n a fixed integer. Each processor $P(s)$ has a set A_s consisting of n integers with

$n \gg p$. Design an efficient parallel algorithm for $P(s)$ that counts the number of elements of the union

$$A = \bigcup_{t=0}^{p-1} A_t.$$

It is sufficient to describe the algorithm in words; no need for a program text. The output value has to be obtained by $P(0)$ only. Analyse the BSP cost of the algorithm.

4. The two-dimensional FFT is defined by performing a 1D FFT operation on every row of a matrix X followed by a 1D FFT operation on every column. We write this operation as FFT2D(X, n), where the matrix X has size $n \times n$. Assume that n and the number of processors p are powers of 2.

   (a) [4 pt] Design a parallel 2DFFT algorithm that uses the 1D sequential algorithm FFT$(\mathbf{x}, n)$, where $\mathbf{x}$ is a vector of length n. Analyse its BSP cost, using the fact that the 1D FFT costs $5n \log_2 n$ flops. Choose a suitable distribution for the input matrix X. The output matrix X has to be distributed in the same way as the input matrix. It is sufficient to describe the algorithm in words; no need for a program text. What are the restrictions on p for this case?

   (b) [4 pt] Design a parallel 2DFFT algorithm that uses the separate supersteps of the 1D parallel algorithm $\text{FFT}_{\text{par}}(\mathbf{x}, n, p)$ we learned. These supersteps are a computation superstep S_0, followed by a communication superstep S_1 that redistributes all the data, and a final computation superstep S_2. Again, choose a suitable distribution for the input matrix X. The output matrix X has to be distributed in the same way. What are the restrictions on p for this case?

   (c) [2 pt] Discuss the advantages and disadvantages of both methods.