Mastermath midterm examination
Parallel Algorithms

Teacher: Rob H. Bisseling, Utrecht University

October 23, 2019

Each of the four questions is worth 10 points. Total time 120 minutes. Motivate you answers!

1. (a) [5 pt] Describe the superstep structure of a BSP algorithm.

   (b) [5 pt] Give an example of a balanced 4-relation and an unbalanced 4-relation for 4 processors.

2. Let $\mathbf{x}$ be a bit vector of length n, i.e. $x_i \in \{0, 1\}$ for $0 \leq i < n$. We want to compute the *parity bit* of $\mathbf{x}$, which is 0 if the total number of 1-bits is even and 1 if it is odd. (The parity bit is sometimes added to a bitstring to help detect a single-bit error.)

   (a) [4 pt] Assume that the input vector $\mathbf{x}$ is distributed by the block distribution over p processors with $n \bmod p = 0$. Give an efficient BSP algorithm for processor $P(s)$ for the computation of the parity bit of $\mathbf{x}$. Only processor $P(0)$ needs to know the result.

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

   (c) [3 pt] Now assume that $\mathbf{x}$ is a long random bit string, where every bit has the same probability $d > 0$ of being 1 and hence $1 - d$ of being 0. Can you optimise the communication in your algorithm to exploit this property? What is the expected gain in BSP cost?

3. Assume we have an $n \times n$ matrix A which is distributed by the square cyclic distribution with $p = M^2$ processors. Assume that $n \bmod M = 0$. Let k, r be integers with $0 \leq k < r < n$ and $\sigma : \{0, 1, \ldots, n-1\} \rightarrow \{0, 1, \ldots, n-1\}$ the permutation that swaps k and r.

   (a) [1 pt] Give the matrix P_σ as defined for this permutation σ.

(b) [2 pt] Describe the meaning of the matrix $P_\sigma A P_\sigma$.

(c) [4 pt] Give an efficient BSP algorithm for processor $P(s,t)$ for the transformation of the matrix A into $P_\sigma A P_\sigma$ on a BSP computer with a relatively high synchronisation cost. Use the notation we have learned to express algorithms.

(d) [3 pt] Analyse the BSP cost of your algorithm.

4. Let A be a tall-and-skinny $n \times k$ matrix with $k \ll n$ and B a $k \times n$ matrix. Assume that we have $p = M^2$ processors, and that $n \bmod M = 0$. Assume that A is available in processor column $P(*,0)$, distributed by the block distribution over the M processors of the processor column. Assume that the matrix A is really skinny, with $k \ll M$. It is also really tall, with $n \gg M$. Similarly, assume that B is available in processor row $P(0,*)$, also distributed by a block distribution over M processors.

(a) [6 pt] Design an efficient BSP algorithm for the computation of the matrix $C = AB$. The output may be distributed, using a distribution of your own choice. You can describe the algorithm in words (instead of using the full notation we learned).

(b) [4 pt] Analyse the BSP cost of your algorithm.