

# Midterm examination Parallel Algorithms (WISM 459).

Teacher: Rob H. Bisseling, Utrecht University

September 26, 2012

*Each of the four questions is worth 10 points. Total time 60 minutes.*

1. Explain the four parameters of the BSP architecture.
2. Give two different examples of a 10-relation with a total communication volume (i.e., number of data words communicated) of 40 for 4 processors.
3. Let  $\mathbf{x}$  be a given vector of length  $n$ , which is distributed over  $p$  processors, with  $n \bmod p = 0$ . You may choose a suitable distribution. Give an efficient BSP algorithm for processor  $P(s)$  for the computation of the product

$$\alpha = x_0^n \cdot x_1^{n-1} \cdots x_{n-1}^1.$$

(Meaning every component  $x_i$  is raised to the power  $n - i$ , for  $i = 0, \dots, n - 1$ .) On output, every processor has to know  $\alpha$ . Analyse the BSP cost.

4. In a molecular dynamics simulation, we have  $n$  particles (atoms) distributed by the block distribution over  $p$  processors, with  $n \bmod p = 0$ . The particles are numbered  $0, \dots, n - 1$ . For each pair of particles  $(i, j)$ , the force  $F_{ij}$  that particle  $i$  exerts on  $j$  must be computed, based on the location of the two particles. Assume that computing one value  $F_{ij}$  costs 10 flops.

In the so-called *ring algorithm*, each processor keeps one block of  $n/p$  particles local, and moves a copy of one block around the ring. In superstep (0), all local interactions are computed. In superstep (1), all copies are moved to the next processor in the ring. In superstep (2), interactions between the local block and the currently held copy

of a block are computed. And so on. Formulate the algorithm in our notation and analyse its BSP cost.

By Newton's third law of motion, we do not need to compute  $F_{ji} = -F_{ij}$ . How much can you save by exploiting the symmetry?