University of Freiburg
Dept. of Computer Science
Prof. Dr. F. Kuhn
S. Faour, P. Schneider

Theory of Distributed Systems
Exercise Sheet 2

Due: Wednesday, 5th of May 2021, 12:00 noon

The CONGEST model

The CONGEST model is a **synchronous** message passing model where the **size** of each message is only allowed to contain bitstrings of length $O(\log n)$, where n is the number of nodes (do not confuse message size with message complexity). Further, we assume that the nodes have IDs in $\{1, \ldots, n\}$.

This means that each message may contain for example (the binary representation of) a constant number of integers $\leq n^c$ for some constant c (in particular IDs). However, a node can not send the IDs of all its neighbors in a single message, as the degree of the network may be large.

Exercise 1: Leader Election *(10 Points)*

a) Given a graph G, describe a deterministic algorithm in the CONGEST model such that every node learns the smallest ID in the graph and *terminates* after $O(D)$ rounds. You may *not* assume that nodes initially know D.

b) Analyze the message complexity of the algorithm. Show that your bound is tight.

Exercise 2: k-Selection Problem in Graphs *(10 Points)*

Given a graph G with n nodes that have pairwise distinct input values $\leq n^c$ for some constant c. In order to solve the k-selection problem in the distributed setting for some $k \leq n$, the k^{th}-smallest value in the graph needs to be announced by exactly one node.
Our goal is to describe a randomized distributed algorithm in the CONGEST model that always solves the k-selection problem with an expected runtime of $O(D \cdot \log n)$.

a) Assume a tree T of depth D. Describe an algorithm that computes in $O(D)$ rounds for every node v a value s_v which equals the size (number of nodes) of the subtree with root v.

b) Assume a tree T of depth D and root r in which each node is able to flip coins. Describe a method to choose a node from the tree uniformly at random (i.e., each node has the same probability to be chosen) in time $O(D)$.

   Hint: Use the algorithm from a).

c) Assume a tree T of depth D, where each node v a boolean b_v as input. Modify the algorithm of a) such that for every node v, the value s_v is equal to the number of nodes in the subtree rooted at v that have $b_v =$ True. Also, modify the algorithm from b) to choose uniformly at random a node among all nodes v with $b_v =$ True.

d) Describe a randomized algorithm that solves the k-selection problem with an expected runtime of $O(D \cdot \log n)$.

   Hint: Use the previous algorithms.