University of Freiburg
Dept. of Computer Science
Prof. Dr. F. Kuhn
A. Mályusz, M. Fuchs, Z. Parsaeian

Theory of Distributed Systems
Exercise Sheet 2

The CONGEST model

The CONGEST model is a **synchronous** message passing model where the **size** of each message is bounded, i.e., the bitstring representing a message consists of at most $O(\log n)$ bits, where n is the number of nodes (do not confuse message size with message complexity). Further, we assume that each node is initially equipped with a unique ID in $\{1, \dots, n^3\}$.

Note that sending such an ID requires at most $\lfloor \log_2 n^3 \rfloor + 1 = O(\log n)$ bits. However, a node can not send the IDs of all its neighbors in a single message, as the degree of the node could 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 (hence, the node with the smallest ID is our leader) and *terminates* after $O(D)$ rounds. You may *not* assume that nodes initially know D. Note that after the execution each node should know the ID of the leader as well as some parent node that points into the direction of the leader.

b) Analyze the message complexity of the algorithm in terms of n. Show that your bound is tight i.e., if $B(n)$ is your message complexity, give a family of graphs where your algorithm indeed has a message complexity of $\Omega(B(n))$.

Exercise 2: k-smallest ID problem *(10 Points)*

Given a graph G with n nodes that have ID's as described in the CONGEST model definition. In order to solve the k-smallest ID problem in the distributed setting for some $k \leq n$, the k^{th}-smallest ID in the graph needs to be announced by exactly one node.
Our goal is to describe a distributed algorithm in the CONGEST model that always solves the k-smallest ID problem with a runtime of $O(D \cdot \log n)$ rounds.
Note that our goal is to construct a deterministic algorithm, however, if you come up with an randomized algorithm that solves the problem in the same number of rounds (in expectation) you will also get full points.

a) Give an $O(D)$ round algorithm that computes a spanning tree T on a graph G, such that the root of T knows the minimum and the maximum ID of the nodes in G.

b) Assume the setting of a) where the root is given an additional value x. Give an $O(D)$ round algorithm that counts the number of nodes with $ID < x$, with $ID > x$ and $ID = x$. It is sufficient if the root node of the tree knows these values in the end.

c) Assume the setting from b) where each node v additionally has a boolean $b_v \in \{0, 1\}$ as input. In the following we call a node active if $b_v = 1$, else we call it inactive. Modify the algorithm of b) such that the root knows the number of **active** nodes that have $ID < x$, $ID > x$ and $ID = x$.

d) Describe an algorithm that solves the k-smallest ID problem in time $O(D \cdot \log n)$.