Distributed algorithms for networks – Retake exam

- **Duration: 2h.**
- **Documents, calculators and dictionaries allowed.**
- **Exercises are independent.**
- **Any unjustified answer will bear no point, but presenting intuitions is encouraged.**
- **You may answer in English or French.**
- **In all exercises, the identifiers of the nodes of a graph of n nodes are assumed to be in the range $[1, n^c]$, for some constant $c > 1$.**

1 Weak c-coloring in cycles

Let $c \geqslant 3$ and G be a graph. A *weak c-coloring* of G is an assignment of colors to the vertices of G such that each vertex has at least one neighbor colored differently.

1. Give a LOCAL algorithm computing a weak 3-coloring in a cycle in $O(\log^* n)$ rounds.
2. Show that, given a weak c-coloring of a cycle, one can compute a *proper* 3-coloring in $O(c)$ rounds.
3. Deduce that there is no LOCAL algorithm producing a weak c-coloring of a cycle in $o(\log^* n)$ rounds.

2 Computing the number of nodes

We consider the following problem: every node of the network G has to output the number of nodes in G.

1. Prove that in the LOCAL model this problem requires $\Omega(n)$ rounds.
2. Prove that this problem can be solved in the CONGEST model in $O(n)$ rounds.

3 Improved universal certification

1. Recall briefly why every property can be locally certified with $O(n^2)$ bit certificates.
2. Assuming that we consider only *d-regular* graphs (i.e. where all the vertices have degree d), show that every property can be certified with $O(dn \log n)$ bit certificates.
3. Let $\mathcal{P} = \{p_1, \dots, p_d\}$ be a set of d *pieces*. We choose uniformly independently at random d subsets $S_1, \dots, S_d$ of $k \leq d$ pieces. Show that the probability that some piece was not picked in $S_1 \cup \cdots \cup S_d$ is at most e^{-k}.
4. Deduce that, in a d-regular graph, every property can be certified with $O(n \log^2 n)$ bit certificates, when allowing each vertex to see the certificates of its neighborhood at distance at most 2.

4 Rake and compress revisited

We consider the following labeling algorithm of an n-vertex tree T:

- Initially $i = 1$ and every vertex is unlabeled.
- At every step i, we consider the subgraph induced by the unlabeled nodes. We label with i the nodes of degree one and those of degree 2 that belong to a path of at least three vertices of degree 2.
- Increase i by one, and repeat until all the nodes are labeled.

1. Let us denote by V_i the subset of V unlabeled at the beginning of step i. Give a lower bound on the proportion of vertices of V_i labeled at step i.
2. Deduce an upper bound N on i for which the algorithm stops. What is the round complexity of the algorithm in the LOCAL model?

A *light labeling* of a tree T is a labeling of the vertices of T such that every node labeled i has at most 2 neighbors labeled at least i and at most 1 neighbor labeled at least $i + 1$.

3. Prove that the algorithm described above outputs a light labeling.

Given a 3-coloring of a tree T and a light labeling of T, we construct a MIS S (initialized as an empty set) as follows. From $i = N$ to $i = 1$:

- Add in S all the the vertices x with one neighbor y labeled at least $i + 1$ such that $y \notin S$.
- Complete greedily S into a MIS of the subgraph labeled by indices at least i.

4. Explain how you can use the 3-coloring of T to perform the second item in a constant number of rounds.
5. Prove by contradiction that $V \setminus S$ only contains connected components of size 1 or 2.
6. Deduce an algorithm which, given a tree T, computes an MIS whose deletion yields connected components of size 1 or 2. What is its round complexity?