Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

FS 2018

Prof. R. Wattenhofer
Jakub Sliwinski

Principles of Distributed Computing
Exercise 14

1 Flow labeling schemes

In this exercise, we focus on flow labeling schemes. Let $G = \langle V, E, w\rangle$ be a weighted undirected graph where, for every edge $e \in E$, the weight $w(e)$ is integral and represents the capacity of the edge. For two vertices $u, v \in V$, the maximum flow possible between them (in either direction), denoted $\mathrm{flow}(u, v)$, can be defined in this context as follows. Denote by G' the multigraph obtained by replacing each edge e in G with $w(e)$ parallel edges of capacity 1. A set of paths P in G' is edge-disjoint if each edge (with capacity 1) appears in no more than one path $p \in P$. Let $\mathcal{P}_{u,v}$ be the collection of all sets P of edge-disjoint paths in G' between u and v. Then $\mathrm{flow}(u, v) = \max_{P\in\mathcal{P}_{u,v}}\{|P|\}$.

We consider the family $\mathcal{G}(n, \hat{\omega})$ of undirected capacitated connected n-vertex graphs with maximum (integral) capacity $\hat{\omega}$, and will find flow labeling schemes for this family. Given a graph $G = \langle V, E, w\rangle$ in this family and an integer $1 \leq k$, let us define the following relation:

$$R_k = \{(x, y)|x, y \in V,\ \mathrm{flow}(x, y) \geq k\}.$$

Question 1 Show that[1] for every $k \geq 1$, the relation R_k induces a collection of equivalence classes on V, $C_k = \{C_k^1, \ldots, C_k^{m_k}\}$, such that $C_k^i \cap C_k^j = \emptyset$ (if $i \neq j$) and $\bigcup_i C_k^i = V$. What is the relationship between C_k and C_{k+1}?

According to the solution of Question 1, given G, one can construct a tree T_G corresponding to its equivalence relations. The k'th level of T corresponds to the relation R_k. The tree is truncated at a node once the equivalence class associated with it is a singleton. For every vertex $v \in V$, denote by $t(v)$ the leaf in T_G associated with the singleton set $\{v\}$.

For two nodes x, y in a tree T rooted at r, we define the separation level of x and y, denoted $\mathrm{SepLevel}_T(x, y)$, as the depth of $z = lca(x, y)$, the least common ancestor of x and y. I.e., $\mathrm{SepLevel}_T(x, y) = \mathrm{dist}_T(z, r)$, the distance of z from the root.

Question 2 Show that if there exists a labeling scheme for distance in trees with labeling size $\mathcal{L}(\mathrm{dist}, T)$, then there is a labeling scheme for separation level with labeling size $\mathcal{L}(\mathrm{SepLevel}, T) \leq \mathcal{L}(\mathrm{dist}, T) + \lceil \log m\rceil$ where m is the number of nodes in the tree. Based on this result and Theorem 14.8 (there is an $O(\log^2 m)$ labeling scheme for distance in trees), show that $\mathcal{L}(\mathrm{flow}, \mathcal{G}(n, \hat{\omega})) = O(\log^2(n\hat{\omega}))$.

Question 3 Find a more careful design of the tree T_G which can improve the bound on the label size to $\mathcal{L}(\mathrm{flow}, \mathcal{G}(n, \hat{\omega})) = O(\log n \log \hat{\omega} + \log^2 n)$. Hint: i) consider all nodes of degree 2 in the tree T_G and weighted trees, ii) naturally extend the notion of separation level to weighted rooted trees.

[1] As a convention, $\mathrm{flow}(x, x) = \infty$.