University of Freiburg
Dept. of Computer Science
Prof. Dr. F. Kuhn
Z. Parsaeian, G. Schmid

Distributed Graph Algorithms
Sample Solution Exercise Sheet 8

Motivation. We saw that network decompositions can help us to solve many interesting problems like $\Delta+1$-coloring, or MIS *deterministically*. Are there other problems for which we can obtain a fast deterministic algorithm using network decomposition?
We will see that network decompositions allow us to solve any problem efficiently that can be solved efficiently in the so called SLOCAL model.

The SLOCAL-Model

In the SLOCAL-Model we are asked to solve a problem Π on a graph $G = (V, E)$. Π is a problem which requires nodes to output values from a set Σ. We process nodes in an arbitrary (adversarially chosen) order. When node v is processed, v can read its r-hop neighborhood and it computes and locally stores its output y and potentially additional information. When reading its r-hop neighborhood, v also reads all the information that has been locally stored by the previously processed nodes there. We call the parameter r the locality of an SLOCAL algorithm.

Exercise 1: $\Delta + 1$-coloring and MIS in SLOCAL

Give locality 1 algorithms for $\Delta + 1$-coloring and MIS in the SLOCAL Model.

Sample Solution

We need to answer how a node v behaves given a partial solution to the problems. As we are dealing with locality 1, this means that to deecide v only knows the partial solution on its immediate neighbors. We give the following two simple rules which solve the problems.

$\Delta + 1$-coloring: Each node picks the smallest color that is not used by any neighbor.

MIS: If no neighbor is in the independent set, then v joins the independent set, else v does not join the independent set.

Exercise 2: Coloring and SLOCAL

Let Π be a problem that admits a locality T SLOCAL algorithm and assume you are given a distance T-coloring[1] of G with C colors.
Show that Π can be solved in $O(C)$ rounds of the LOCAL model.

[1] In a distance T coloring, is a coloring such that any two vertices within distance T of each other are assigned different colors. Equivalently, two nodes of the same color have distance $> T$, so what is normally understood as a proper coloring is a distance 1 coloring.

Sample Solution

We process the nodes color class by color class. In the iteration for color i, every node of color i locally simulates the T-locality SLOCAL algorithm and fixes its output.
This is well defined because any two nodes u, v of the same color have distance strictly larger than T. Hence, the radius-T neighborhood of a node v contains no other node of color i, and therefore the output computed by v does not depend on the order in which nodes of the same color are processed.
To formalize this, let $V_1, \dots, V_C$ denote the sets of nodes with colors $1, \dots, C$, respectively. For each i, fix an arbitrary ordering S_i of the nodes in V_i, and let

$$S = S_1 \circ S_2 \circ \cdots \circ S_C$$

be the concatenation of these sequences. Running the SLOCAL algorithm on G according to the order S yields a correct solution by assumption.
Now observe that within each set V_i, all nodes have pairwise distance greater than T. Consequently, the relative order of nodes inside S_i is irrelevant: none of the outputs produced by nodes in V_i can influence the computation of another node in V_i. Therefore, processing all nodes of color i simultaneously in one LOCAL round produces the same outcome as processing them sequentially according to S_i in the SLOCAL model.
Repeating this for all C color classes yields the same output as the SLOCAL execution on S and takes $O(C)$ rounds in the LOCAL model.

Exercise 3: Distance Colorings

What is the minimum number of colors required for a distance T-coloring in a Δ-regular tree? What is the locality of distance T-coloring in SLOCAL?

Sample Solution

We do a sum over the levels of the Δ-regular tree

$$\begin{aligned} N &= 1 + \Delta + \Delta \cdot (\Delta - 1) + \Delta \cdot (\Delta - 1)^2 + \ldots + \Delta \cdot (\Delta - 1)^{T-1} \\ &= 1 + \Delta \sum_{i=0}^{T-1} (\Delta - 1)^i \end{aligned}$$

we have

$$\Delta^T < 1 + \Delta \sum_{i=0}^{T-1} (\Delta - 1)^i < 1 + \sum_{i=1}^{T} \Delta < 2\Delta^T$$

So the number of colors is $\Theta(\Delta^T)$.
Such a coloring can be computed in the LOCAL model, by computing a $\Delta + 1$ coloring of the power graph G^T. Note that the maximum degree of G^T is $\Theta(\Delta^T)$ aswell. So we may either use the random color trial algorithm to obtain such a coloring in $O(\log n)$ rounds, or use the $O(\Delta + \log^* n)$ algorithm, which will take $O(\Delta^T + \log^* n)$ rounds.

Exercise 4: Network Decomposition and SLOCAL with locality 1

Let Π be a problem that admits a locality 1 SLOCAL algorithm and assume you are given a (c, d)-network decomposition.
Show that Π can be solved in $O(c \cdot d)$-rounds of the LOCAL model.

Sample Solution

We process the clusters color by color. For a fixed color, all clusters of this color are pairwise non-adjacent and can therefore be handled in parallel.
Consider a single cluster C. Since the cluster has weak diameter at most d, one designated node of C can gather the entire induced subgraph on C, together with all edges incident to C, in $O(d)$ rounds. Using this information, the designated node fixes an arbitrary ordering S_C of the nodes in C and locally simulates the locality-1 SLOCAL algorithm on this sequence.
Because clusters of the same color are non-adjacent, these simulations are independent and do not interfere with each other. Hence, all clusters of a fixed color can be processed simultaneously in $O(d)$ rounds.
As in Exercise 2, we can conceptually concatenate the sequences S_C of all clusters of one color, and then concatenate these sequences over all colors. The result is a global ordering of the nodes of G. Running the SLOCAL algorithm according to this ordering produces a correct solution, and the parallel LOCAL execution described above yields the same outcome.
Since there are c colors and each color class requires $O(d)$ rounds, the total running time is $O(c \cdot d)$ rounds in the LOCAL model.

Exercise 5: Network Decomposition and SLOCAL with larger localities

Let Π be a problem that admits a locality T SLOCAL algorithm and assume you have a LOCAL algorithm that computes a $(\log n, \log n)$-network decomposition in $O(\log^2 n)$ rounds.
Show that Π can be solved in $O(T \cdot \log^2 n)$ rounds in the LOCAL model.
Hint: Exercise 4 was easy, because of the fact that same colored clusters of the network decomposition are non-adjacent, but if $T > 1$, we need a stronger separation guarrantee, i.e. that clusters are further away from each other.

Sample Solution

We run the network decomposition algorithm on G^T. We get that clusters of the same color are non-adjacent in G^T and have diameter d in G^T. As a result, clusters of the same color have distance $> T$ in G and have diameter $T \cdot d$ in G.
We can now follow the same procedure as in Exercise 4 and get a running time of $O(T \cdot d \cdot c) = O(T \cdot \log^2 n)$.