Principles of Distributed Computing 04/25, 2018

Lecture 9

Lecturer: Mohsen Ghaffari *Scribe:*

In this lecture, we discuss a rather general method for solving various a wide range of *local* problems. The key concept in our discussion will be *network decompositions* first introduced by [ALGP89], also known as *low-diameter graph decomposition* [LS91].

0.1 Definition and Applications

Let us start with defining this concept.

Definition 1. *(Weak Diameter Network Decomposition) Given a graph $G = (V, E)$, a $(\mathcal{C}, \mathcal{D})$ weak diameter network decomposition of G is a partition of G into vertex-disjoint graphs G_1, G_2, ..., $G_{\mathcal{C}}$ such that for each $i \in \{1, 2, \ldots, \mathcal{C}\}$, we have the following property: the graph G_i is made of a number of vertex-disjoint and mutually non-adjacent clusters X_1, X_2, ..., X_ℓ, where each two vertices $v, u \in X_j$ have distance at most D in graph G. We note that we do not bound the number ℓ. We refer to each subgraph G_i as one* block *of this network decomposition.*

Definition 2. *(Strong Diameter Network Decomposition) Given a graph $G = (V, E)$, a $(\mathcal{C}, \mathcal{D})$ strong diameter network decomposition of G is a partition of G into vertex-disjoint graphs G_1, G_2, ..., $G_{\mathcal{C}}$ such that for each $i \in \{1, 2, \ldots, \mathcal{C}\}$, we have the following property: each connected component of G_i has diameter at most D.*

Notice that a strong diameter network decomposition is also a weak diameter network decomposition.

Network decompositions can be used to solve a wide range of *local* problems. To see the general method in a concrete manner, let us go back to our beloved $(\Delta + 1)$-coloring problem.

Theorem 3. *Provided an $(\mathcal{C}, \mathcal{D})$ weak-diameter network decomposition of a graph G, we can compute a $\Delta + 1$ coloring of G in $O(\mathcal{C}\mathcal{D})$ rounds.*

Proof. We will color graphs G_1, G_2, ..., $G_{\mathcal{C}}$ one by one, each time considering the coloring assigned to the previous subgraphs. Suppose that vertices of graphs graphs G_1, G_2, ..., G_i are already colored using colors in $\{1, 2, \ldots, \Delta + 1\}$. We explain how to color G_{i+1} in $O(D)$ rounds. Consider the clusters X_1, X_2, ..., X_ℓ of G_{i+1} and notice their two properties: (1) they are mutually non-adjacent, (2) for each cluster X_j, its vertices are within distance D of each other (where distances are according to the base graph G). For each cluster X_j, let node $v_j \in X_j$ who has the maximum identifier among nodes of X_j be the leader of X_j. Notice that leaders of clusters X_1, X_2, ..., X_ℓ can be identified in $O(D)$ rounds (why?). Then, let v_j aggregate the topology of the subgraph induced by X_j as well as the colors assigned to nodes adjacent to X_j in the previous graphs G_1, G_2, ..., G_i. This again can be done in $O(D)$ rounds, thanks to the fact that all the relevant information is within distance $D + 1$ of v_j. Once this information is gathered, node v_j can compute a $(\Delta + 1)$-coloring for vertices of X_j, while taking into account the colors of neighboring nodes of previous graphs, using a simple greedy procedure. Then, node v_j can report back these colors to nodes of X_j. This will happen for all the clusters X_1, X_2, ..., X_ℓ in parallel, thanks to the fact that they are non-adjacent and thus, their coloring choices does not interfere with each other. □

0.2 Randomized Algorithm for Network Decomposition

Theorem 4. *There is a randomized* LOCAL *algorithm that computes a $(\mathcal{C}, \mathcal{D})$ weak-diameter network decomposition of any n-node graph G, for $\mathcal{C} = O(\log n)$ and $\mathcal{D} = O(\log n)$, in $O(\log^2 n)$ rounds, with high probability*[1]*.*

As we see in the exercises of this class, the two key parameters $\mathcal{C}$ and $\mathcal{D}$ are nearly optimal and one cannot improve them simultaneously and significantly.

Network Decomposition Algorithm: Suppose that we have already computed subgraphs $G_1, \ldots, G_i$ so far. We now explain how to compute a subgraph $G_{i+1} \subseteq G \setminus (\cup_{j=1}^{i} G_j)$, in $O(\log n)$ rounds, which would satisfy the properties of one block of a weak diameter network decomposition.

Let each node v pick a random radius r_u from an geometric distribution with parameter ε, for a desired (free parameter) constant $\varepsilon \in (0, 1)$. That is, for each integer $y \geq 1$, we have $\Pr[r_u = y] = \varepsilon(1-\varepsilon)^{y-1}$. We will think of the vertices within distance r_u of u as the *ball of node* u. Now for each node v, let $Center(v)$ be the node u^* among nodes u such that $dist_G(u, v) \leq r_u$ that has the smallest identifier. The is, $Center(v) = u^*$ is the smallest-identifier node whose ball contains v. Define the clusters of G_i by letting all nodes with the same center define one cluster, and then discarding nodes who are at the boundary of their cluster. That is, any node v for which $dist_G(v, u) = r_u$ where $u = Center(v)$ remains unclustered.

There are two properties to prove: one that the clusters have low diameter, and second, that after $\mathcal{C}$ iterations, all nodes are clustered. In the following two lemmas, we argue that with high probability, each cluster has diameter $O(\log n/\varepsilon)$ and after $\mathcal{C} = O(\log_{1/\varepsilon} n)$ iterations, all nodes are clustered.

Lemma 5. *With high probability, the maximum cluster diameter is at most $O(\log n/\varepsilon)$. Hence, this clustering can be computed in $O(\log n/\varepsilon)$ rounds, with high probability.*

Proof. The proof is simple and is left as an exercise. □

Lemma 6. *For each node v, the probability that v is not clustered — that v is on the boundary of its supposed cluster and thus it gets discarded — is at most ε.*

Proof. Notice that

$$\begin{aligned} &\Pr\ [v \text{ is not clustered }] = \\ \sum_{u \in V} &\Pr\ [v \text{ is not clustered } | Center(v) = u] \cdot \Pr[Center(v) = u] \end{aligned}$$

For each vertex u, let $before(u)$ denote the set of all vertices whose identifier is less than that of u. Define the following events

- $\mathcal{E}_1 = (r_u = dist_G(v, u))$.
- $\mathcal{E}_2 = (r_u \geq dist_G(v, u))$.
- $\mathcal{E}_3 = (\forall u' \in before(u), r_{u'} < dist_G(v, u'))$.

[1] Throughout, we will use the phrase *with high probability to indicate that an event happens with probability at least $1 - \frac{1}{n^c}$, for a desirably large but fixed constant $c \geq 2$.*

We have

$$\begin{aligned}
&\Pr \quad [v \textit{ is not clustered } | Center(v) = u] \\
&= \quad \Pr[\mathcal{E}_1 \cap \mathcal{E}_3 \,|\, \mathcal{E}_2 \cap \mathcal{E}_3] \\
&= \quad \frac{\Pr[\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3]}{\Pr[\mathcal{E}_2 \cap \mathcal{E}_3]} \\
&= \quad \frac{\Pr[\mathcal{E}_1 \cap \mathcal{E}_3]}{\Pr[\mathcal{E}_2 \cap \mathcal{E}_3]} \\
&= \quad \frac{\Pr[\mathcal{E}_3] \cdot Pr[\mathcal{E}_1|\mathcal{E}_3]}{\Pr[\mathcal{E}_3] \cdot Pr[\mathcal{E}_2|\mathcal{E}_3]} \\
&= \quad \frac{Pr[\mathcal{E}_1]}{Pr[\mathcal{E}_2]} = \varepsilon,
\end{aligned}$$

where in the penultimate equality, we used the property that the event $\mathcal{E}_3$ is independent of events $\mathcal{E}_1$ and $\mathcal{E}_2$, and the last equality follows from the probability distribution function of the exponential distribution (recall that this is exactly the *memoryless property* of the exponential distribution). Hence, we can now go back and say that

$$\begin{aligned}
&\Pr \quad [v \textit{ is not clustered }] \\
&= \quad \sum_{u \in V} \Pr[v \textit{ is not clustered } | Center(v) = u] \cdot \Pr[Center(v) = u] \\
&= \quad \sum_{u \in V} \varepsilon \cdot \Pr[Center(v) = u] = \varepsilon.
\end{aligned}$$

□

Corollary 7. *After $\mathcal{C} = O(\log_{1/\varepsilon} n)$ iterations, all nodes are clustered, with high probability.*

References

[ALGP89] Baruch Awerbuch, M Luby, AV Goldberg, and Serge A Plotkin. Network decomposition and locality in distributed computation. In *Foundations of Computer Science, 1989., 30th Annual Symposium on*, pages 364–369. IEEE, 1989.

[LS91] Nathan Linial and Michael Saks. Decomposing graphs into regions of small diameter. In *Proceedings of the Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '91, pages 320–330, 1991.