

Distributed Local Multi-Aggregation and Centrality Approximation

Benjamin Dissler
disslerb@ethz.ch
ETH Zurich

Stephan Holzer*
holzer@mit.edu
MIT

Roger Wattenhofer
wattenhofer@ethz.ch
ETH Zurich

May 24, 2016

Abstract

We study local aggregation and graph analysis in distributed environments using the message passing model. We provide a flexible framework, where each of the nodes in a set S —which is a subset of all nodes in the network—can perform a large range of common aggregation functions in its k -neighborhood. We study this problem in the CONGEST model, where in each synchronous round, every node can transmit a different (but short) message to each of its neighbors. While the k -neighborhoods of nodes in S might overlap and aggregation could cause congestion in this model, we present an algorithm that needs time $\mathcal{O}(|S| + k)$ even when each of the nodes in S performs a different aggregation on its k -neighborhood. The framework is not restricted to aggregation-trees such that it can be used for more advanced graph analysis. We demonstrate this by providing efficient approximations of centrality measures and approximation of minimum routing cost trees.

*Part of the work was done at ETH Zurich. At MIT the author was supported by the following grants: AFOSR Contract Number FA9550-13-1-0042, NSF Award 0939370-CCF, NSF Award CCF-1217506, NSF Award number CCF-AF-0937274.

1 Introduction

Data aggregation and analysis is one of the most basic tasks at the heart of many distributed systems and the question of aggregating and analyzing information and networks itself as efficient as possible arises daily. The result of this is a huge body of work ranging from theoretical to practical aspects focusing on optimizing e.g. speed, space, bandwidth, energy, fault-tolerance and accuracy. As already [21] stated, “the database community classifies aggregation functions into three categories: *distributive* (max, min, sum, count), *algebraic* (plus, minus, average, variance), and holistic (median, k^{th} smallest, or k^{th} largest value). Combinations of these functions are believed to support a wide range of reasonable aggregation queries.”. Often one is also interested in computing a combination of these such as, e.g., “What is the average of the 10% largest values?” [21].

However, most of this work considers the case in which only one node in the network aggregates information (and then often broadcasts it to all other nodes). In reality, many nodes of a large network are independently interested in aggregating data at the same time

and restricted to their local neighborhood. That is, all nodes in a subset $S \subseteq V$ want to perform a (possibly different) aggregation, where V is the set of nodes in the network. For example 1) a few nodes in a network want to know if certain data is stored in their neighborhood, or 2) cars participating in a vehicular ad hoc network (VANET) want to aggregate information on traffic, safety warnings and parking spots from their local neighborhood, and 3) nodes who turned idle search for the busiest nodes in their respective neighborhoods to take work from them. A simple approach is to just let all nodes in S perform their aggregation at the same time—however, this might lead to congestion and a worst case runtime of $\mathcal{O}(|S| \cdot k)$.

In this paper we present a general framework for local multi-aggregation that allows a time- and message-optimal implementation of a wide range of aggregation functions. We do this in a general setting, where nodes in S aggregate data from their k -neighborhood using shortest paths and achieve a runtime of $\mathcal{O}(|S| + k)$ to do so.

We show how to perform aggregation and graph analysis by aggregating through all possible shortest paths between any pair of nodes (not only using one path as it is usually done in an aggregation-tree). This is a powerful tool which enables us to provide an efficient approximation of Betweenness Centrality. We also present efficient approximations of Closeness Centrality and Minimum Routing Cost Trees.

1.1 Our Contribution

We provide a framework for local multi-aggregation that takes care of scheduling messages sent between nodes in an efficient way. One only needs to specify how nodes process incoming messages. That is designing algorithms depending on the aggregation-function at hand, which transform received messages into new messages to be sent. Using our framework, one can aggregate information not only by using a tree, but using all possible shortest paths from a root node to any other node. This has advantages for advanced computations as we demonstrate later. Thus we show two variations of the algorithm:

1. only one (shortest) r - t -path is needed for each $(r, t) \in S \times N_k(r)$ to perform the aggregation.
2. all shortest r - t -paths are needed for each $(r, t) \in S \times N_k(r)$ to perform the aggregation.

The last version is e.g. of interest for computing betweenness centrality measures, which is a measure that depends on all shortest paths, not only a single one. ¹.

¹In the algorithm we approximate the number of all shortest paths starting in a certain sampled set S of nodes. Note that it does not seem possible to approximate centrality measures without performing these $|S|$ independent aggregations.

To perform $|S|$ independent (possibly different) aggregations in the k -neighborhood of each of the nodes in S , our algorithm takes time $\mathcal{O}(|S| + k)$, which we show to be optimal (see Remark 3.3). As an example of an aggregation-function, which can be plugged into our framework, we consider computing the maximum value stored in the k -neighborhood of each of the S nodes. Further aggregation functions can be implemented in a similar way. Different root nodes in S can perform different aggregations at the same time. Based on this, we show how to approximate centrality measures and minimum routing cost trees and obtain the following theorems:

Theorem 1.1. *Algorithm 1 DLGCOMP computes $|S|$ valid directed leveled graphs with depth k , in time $\mathcal{O}(|S| + k)$. And Algorithm 2 DLGAGR aggregates information through $|S|$ directed leveled graphs with depth k , in time $\mathcal{O}(|S| + k)$.*

Theorem 1.2. *The tree variations of Algorithms 1 DLGCOMP and 2 DLGAGR compute $|S|$ trees T_r in time $\mathcal{O}(|S| + D_\omega)$ and aggregate information through these $|S|$ trees, in time $\mathcal{O}(|S| + D_\omega)$.*

Theorem 1.3. *Algorithm 3 computes an estimation of the betweenness centrality for at least \hat{n} nodes. For all those nodes holds: for $\frac{1}{n^\varepsilon} < \varepsilon' < \frac{1}{2}$ and $c \geq 0$, if the betweenness centrality of a node v is n^2/t for some constant $t \geq 1$, then with probability at least $1 - 2\varepsilon'$ its betweenness centrality can $(\times, \frac{1}{\varepsilon'})$ -approximated with computation of $|S| = t\varepsilon'$ samples of source nodes. Algorithm 3 runs in time $\mathcal{O}\left(|S| + D_\omega \cdot \max\left(1, \log \frac{|S|}{D_h}\right)\right)$.*

Theorem 1.4. *In a weighted graph $G = (V, E, \omega)$, when the weights ω denote the distance between two nodes, Algorithm 6 approximates closeness centrality of all nodes with an inverse additive error of εD_ω and runs in $\mathcal{O}\left(\frac{\log(n)}{\varepsilon^2} + D_\omega\right)$ time.*

Theorem 1.5. *Algorithm 7 MRCT computes a $(\times, 2)$ -approximation of the S -MRCT problem in time $\mathcal{O}(|S| + D_\omega)$, when $\omega(e)$ is uniform or corresponds to the time needed to transmit a message through edge e .*

Theorem 1.6. *Algorithm MRCTRAND computes a $(\times, 2 + \varepsilon)$ -approximation of the S -MRCT problem in time $\mathcal{O}(D_\omega + \log(n))$, when $\omega(e)$ is uniform or corresponds to the time needed to transmit a message through edge e .*

Algorithm and proof of Theorem 1.1 is stated in Section 3.1. Algorithm and proof of Theorem 1.2 is stated in Section 3.2. Algorithm and proof of Theorem 1.3 is stated in Section 4.1. Algorithm and proof of Theorem 1.4 is stated in Section 4.2. Algorithm and proof of Theorem 1.6 is stated in Section 4.3.

1.2 Related Work

(Local multi-)aggregation The authors of [25] survey the general area of data-aggregation in wireless sensor networks. It is known [2] that local multi-aggregation is useful in generalizations of Zone-Based Intrusion Detection Systems mentioned in. These systems detect local changes of the topology (assuming the graph does not change too fast) and react to it. The authors of [19] perform local multi-aggregation in sensor networks to detect basic changes of the local environment, such as variations of temperature or concentration of toxic gas. If the value in one node passes a threshold, it starts a local aggregation. They also provide applications to fire fighters and emergency services in case of a flood. In [19] changes of areal objects are detected by the cooperation of sensors in the locality of the changes and later reported to a central node.

Centrality Classical applications of Closeness Centrality [4, 5] and Betweenness Centrality [14] are in the area of social network analysis [28]. E.g. Closeness Centrality indicates how fast information spreads and Betweenness Centrality indicates how easy it can be manipulated by certain nodes. However, these findings can often be transferred to networks based on electric devices. One example are simple routing protocols, where malicious nodes with high centrality have the power to manipulate large parts of the communication. Eppstein and Wang showed in [12] a fast approximation algorithm to derive the closeness centrality of a node by sampling only a few nodes S and computing shortest paths from S to all nodes. Brandes [7] was able to provide a fast algorithm for computing betweenness centrality recursively, which Bader et al. [3] extended to an even faster approximation using an adaptive sampling technique.

We continue by reviewing a few applications in network analysis and design independent of social networks: For distance approximation, when using the landmark method [23, 24], it is mentioned in [1] that a modification of [26] chooses landmarks using local closeness centrality. According to [27], this is implemented e.g. in the Trinity Graph Engine. A relation to network flows is established e.g. in [6] for Closeness as well as Betweenness Centrality. The authors of [9] propose a routing scheme based on Betweenness Centrality.

MRCT Distributed approximations of MRCT in the CONGEST model were studied in [15, 20, 10]. The authors of [20] showed how to use results of [13] to obtain a randomized $\mathcal{O}(\log n)$ -approximation of the MRCT in time $\mathcal{O}(D_{sp} \cdot \log^2(n))$. Observe that this might be $\mathcal{O}(n \log^2 n)$ even in a graph with hop diameter $D_h = 1$. In [15] two algorithms were presented for unweighted graphs and weighted graphs, where the weight of an edge corresponds to the traversal time of the edge. A deterministic one, that computes a 2-approximation in linear time $\mathcal{O}(n)$ and a randomized one, that computes w.h.p. an $2 + \varepsilon$ -approximation in time $\mathcal{O}(D + \log n)$ (and variations thereof). A lower bound of $\Omega(\sqrt{n} + D)$ for any randomized ε -error algorithm that computes a $[poly\ n]$ -approximation was shown in [10] for weighted graphs. For a more detailed overview on work related to MRCT computations in various computational models, we refer the reader to [15].

Shortest Paths At the heart of our algorithms is an algorithm to compute $|S|$ directed layered graphs rooted in nodes $S \subseteq V$ in parallel. This algorithm is a modification of the S -Shortest Paths algorithm of [17]. Furthermore our modification extends this S -Shortest Paths algorithm to the weighted setting and to compute shortest paths only up to a certain distance that can be provided as input. Note that the k -source detection algorithm by Lenzen and Peleg [22] can be extended to weighted graphs in a similar way and could be modified to compute directed layered graphs as well, but we developed our method that resulted in this paper independently and simultaneously [11, 16].

2 Model, Definitions and Problems

Model: Our network is represented by an undirected graph $G = (V, E)$. Nodes V correspond to processors (computers, sensor nodes or routers), two nodes are connected by an edge from set E if they can communicate directly with each other. We denote the number of nodes of a graph by n , and the number of its edges by m . Furthermore we assume that each node has a unique ID in the range of $\{1, \dots, 2^{\mathcal{O}(\log n)}\}$, i.e. each node can be represented by $\mathcal{O}(\log n)$ bits. Initially the nodes know only G the IDs of their immediate neighborhood, the number of nodes in G and the node with the lowest ID further denoted as node 1.²

²Note that for computing the schedule in our framework we do not need to know n and the node with smallest ID. These assumptions are only necessary for our applications such as centrality computation. The runtime of these applications depends on the diameter of the graph such that computing these values does not affect the runtime.

An unweighted shortest path in G between two nodes u and v is a $u - v$ -path consisting of the minimum number of edges contained in any $u - v$ -path. Denote by $d(u, v)$ the unweighted distance between two nodes u and v in G , which is the length of an unweighted shortest $u - v$ -path in G . We also say u and v are $d(u, v)$ hops apart. By $\omega : E \rightarrow \mathbb{N}$ we denote a graph's weight function and by $\omega(e)$ the non-negative weight of an edge in G . For an edge $e = (u, v)$ we often write $\omega(u, v)$ to denote $\omega(e)$. For arbitrary nodes u and v , by $\omega(u, v) := \min_{\{P|P \text{ is } u-v\text{-path in } G\}} \sum_{e \text{ is edge in } P} \omega(e)$ we define the weighted distance between u and v , that is the weight of a shortest weighted path connecting u and v . In case we consider a subgraph H of a graph G , we denote by d_H and ω_H the distances with respect to using only edges in H .

In our model we consider only symmetric edge weights, i.e. an edge (u, v) has in both directions the same weight $\omega(u, v) = \omega(v, u)$. We denote a graph with weighted edges as $G = (V, E, \omega)$. By $N_k(v)$ we denote the k -neighborhood of a node v . That is in an unweighted graph all nodes in G that can be reached from v using k hops. And in a weighted graph all nodes that can be reached from v using paths with at most weight k . We use the convention that $v \in N_k(v)$. Given a set $S \subset V$, set $N_k(S) := \cup_{v \in S} N_k(v)$ denotes the k -neighborhood of S . During the paper we denote the degree $|N_1(v) - 1|$ of a node v by $\text{deg}(v)$.

Definition 2.1. *The weighted eccentricity $\text{ecc}_\omega(u)$ of a node u is the largest weighted distance to any other node in the graph, that is $\text{ecc}_\omega(u) := \max_{v \in V} \omega(u, v)$. The unweighted (hop) eccentricity $\text{ecc}_h(u)$ of a node u is the largest hop distance to any other node in the graph, that is $\text{ecc}_h(u) := \max_{v \in V} d(u, v)$.*

Definition 2.2. *The weighted diameter $D_\omega(G) := \max_{u \in V} \text{ecc}_\omega(u)$ of a graph G is the maximum weighted distance between any two nodes of the graph. The unweighted diameter (or hop diameter) $D_h(G) := \max_{u \in V} \text{ecc}_h(u)$ of a graph G is the maximum number of hops between any two nodes of the graph.*

Observe that $D_h = D_\omega$ for unweighted graphs. In weighted graphs it is true that $1 \leq D_h \leq n$ as well as $D_h \leq D_\omega$.

In this paper we need a graph substructure to describe shortest paths from a root node r to all other nodes. Such a representation can be found e.g. in a tree. However, for applications of our framework to e.g. calculating the Betweenness Centrality of a node in Section 4.1, we need to consider all shortest paths between two nodes, not just a single one as provided by a tree. Therefore we need to define a graph structure which clearly indicates all possible shortest paths between a root node $r \in V$ and all other nodes $v \in V$.

Definition 2.3 (Tree T_r). *Given a node r , we denote the spanning tree of G that results from performing a breadth-first search (BFS_r) starting at r by T_r .*

Definition 2.4 (Directed Levelled Graph (DLG) \mathcal{L}_r , unweighted). *Given a node $r \in V$ of a graph $G = (V, E)$, partition all nodes $v \in V$ in disjoint subsets V_0, V_1, \dots, V_{D_h} with $v \in V_i$ if $d(r, v) = i$. Add an directed edge (u, v) to a set \vec{E}_r for every edge $(u, v) \in E$ with $u \in V_i$ and $v \in V_{i+1}$. That is, every shortest path from r to any other node v in G is a directed path in $\mathcal{L}_r = (V, \vec{E}_r)$. We say \mathcal{L}_r is rooted in r . A graphic is provided in Figure 1.*

In the weighted case, we partition all nodes $v \in V$ in disjoint subsets $V_0, V_1, \dots, V_{D_\omega}$ with $v \in V_i$ if $\omega(r, v) = i$. Add an directed edge (u, v) to set \vec{E}_r for every edge $(u, v) \in E$ with $u \in V_i$ and $v \in V_{i+\omega(u, v)}$. That is, every shortest weighted path from r to any other node v in G is a directed weighted path in $\mathcal{L}_r = (V, \vec{E}_r, \omega)$. A graphic is provided in Figure 2.

Definition 2.5 (parent, children, ancestors and leaves in a DLG). *Given a node v , we define a parent of v in a DLG \mathcal{L}_r as any neighbor u_p of v connected by a directed edge $(u_p, v) \in \vec{E}_r$ to*

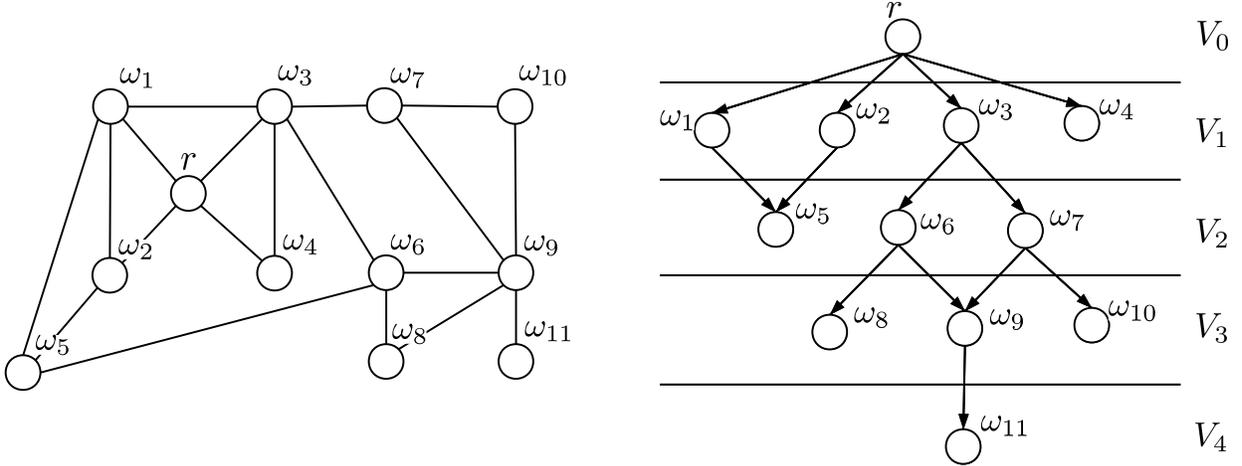


Figure 1: $G = (V, E)$ on the left side and its leveled graph \mathcal{L}_r rooted in r on the right.

v . As a node can have several parents, we often consider a set of parents, formally defined as: $P_r(v) = \{u \in V : \{u, v\} \in E \text{ and } \omega(r, v) = \omega(r, u) + \omega(u, v)\}$. Children are all neighbors u_c of v connected by a directed edge $(v, u_c) \in \vec{E}_r$ to v . A leaf node in a DLG \mathcal{L}_r is a node without any children.

Definition 2.6 (approximation). Given an optimization problem P , denote by sol_{OPT} the value of the optimal solution for P and by sol_A the value of the solution of an algorithm A for P . Let $\varepsilon \geq 0$. We say A is a $(+, \varepsilon)$ -approximation (additive approximation) with additive error ε for P if $|sol_{OPT} - sol_A| \leq \varepsilon$ for any input. Let $\delta \geq 0$. Like in [12] we say A is an inverse additive approximation with inverse additive error δ for P if $\left| \frac{1}{sol_{OPT}} - \frac{1}{sol_A} \right| \leq \delta$ for any input. We say A is (\times, ρ_1) -approximative (a one-sided multiplicative approximation) for P if $(\rho_1 \cdot sol_{OPT}) \leq sol_A \leq sol_{OPT}$ and $0 < \rho_1 < 1$ or if $sol_{OPT} \leq sol_A \leq (\rho_1 \cdot sol_{OPT})$ and $\rho_1 > 1$ for any input.

Like in [12], the inverse additive error is used in the closeness centrality approximation, see Section 4.2.

Fact 2.7. For any node $u \in V$ we know that $ecc_\omega(u) \leq D_\omega(G) \leq 2 \cdot ecc_\omega(u)$.

2.1 Problem Statements

We start by formally stating the problems we consider in this paper. First we state local multi-aggregation, which is not only of interest by itself but turns out to be at the heart of the other problems. Centrality-measures are graph-properties that are based on the number of shortest paths. Besides classical applications in social network analysis [8, 28], they have a number of applications in design and analysis of distributed networks such as [1, 6, 9]. Minimum Routing Cost Spanning Trees can be used to minimize the average cost of communication in a network while keeping a sparse routing-structure [18, 29].

Definition 2.8. (local multi-aggregation). In the problem of local multi-aggregation, we are given a subset S of nodes V in a graph G . Each of the nodes in G contains (several) values and each node in S wants to evaluate a (possibly different) aggregation-function based on (some of) these values stored in nodes of its k -neighborhood. We consider two variations:

1. only one (shortest) s - t -path is needed for each $(s, t) \in S \times N_k$ to perform the aggregation.

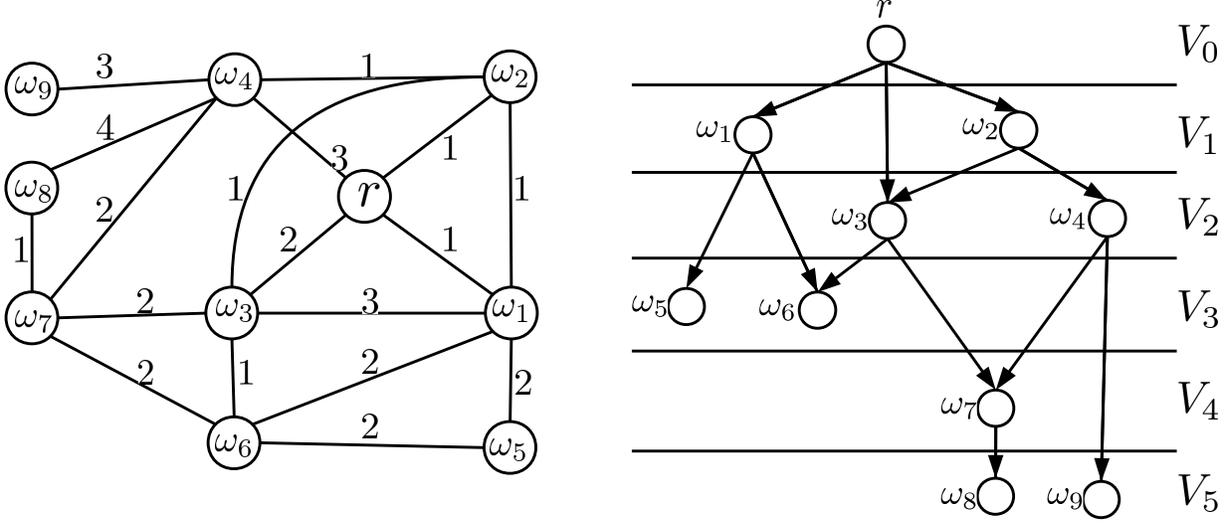


Figure 2: $G = (V, E, \omega)$ on the left side and its leveled graph \mathcal{L}_r rooted in r on the right

2. all shortest s - t -paths are needed for each $(s, t) \in S \times N_k$ to perform the aggregation.

Betweenness centrality is a measure for centrality of a node $u \in V$, which is based on the number of shortest paths in a graph G node u is part of. In a graph $G = (V, E)$, let σ_{st} denote the number of shortest paths from s to t and let $\sigma_{st}(v)$ denote the number of shortest paths from s to t that go through the node v for $s, t, v \in V$. Closeness centrality is another measure to identify important nodes in a graph. The closeness centrality of a node is the inverse of the average distance to all other nodes. At the heart of our solutions for these problems is the S -SP problem and we define:

Definition 2.9 (Betweenness Centrality, as stated in [14]). $BC(v) := \sum_{s \neq v \neq t \in V} \frac{\sigma_{st}(v)}{\sigma_{st}}$.

Definition 2.10 (Closeness Centrality, as stated in [5]). $CC(v) := \frac{n-1}{\sum_{i=1}^n d(v, i)}$.

Definition 2.11. (*Minimum Routing Cost Tree (MRCT) as stated in [30]*). The MRCT problem [30] is defined as follows. The routing cost $RC(H)$ of a subgraph H of G is the sum of the routing costs of all pairs of vertices in the tree, i.e., $RC(H) := \sum_{u, v \in V} \omega_H(u, v)$. Our goal is to find a spanning tree with minimum routing cost. Given a subset S of the vertices V in G , in the S -MRCT problem [15], our goal is to find a subtree T of G that spans S and has minimum routing cost with respect to S . That is a tree T such that $RC_S(T) := \sum_{u, v \in S} \omega_T(u, v)$ is minimized. Here, $RC_S(T)$ denotes the routing cost of T with respect to S .

Definition 2.12 (S -SP). Let $G = (V, E)$ be a graph. In the S -Shortest Paths (S -SP) problem, we are given a set $S \subseteq V$ and need to compute the shortest path lengths between any pair of nodes in $S \times V$ such that in the end each node in V knows its distance to all nodes in S .

3 Local Multi-Aggregation

Many distributed algorithms to compute network properties can be stated in a way that they distribute and/or aggregate information through a tree or a directed leveled graph (DLG). For example computing the max-value of the k -neighborhood (using a tree, see the text below), the

sum of all values in the k -neighborhood (using a tree, see the text below) and computing centrality measures (using a DLG, see Sections 4.1 and 4.2).

By extending the algorithms of [17] and [15], we present an algorithm which can aggregate information in the k -Neighborhood of a set of root nodes $r \in S$. Note that if information from the whole graph is needed, k can be set to D_ω .

All known exact computations of these properties need to compute a tree or DLG for every node $v \in V$, to evaluate v 's dependency on other nodes. Often the exact computation can be approximated by evaluating only the dependencies of a subset $S \subseteq V$ of nodes and therefore computing only $|S|$ trees or DLGs with depth k in respect to the root nodes $r \in S$. We provide in Section 3.1 an algorithm to compute $|S|$ such DLGs rooted in the nodes of S in time $\mathcal{O}(|S| + k)$. Furthermore we provide an algorithm to aggregate information on these previously computed DLGs, again with a time complexity of $\mathcal{O}(|S| + k)$. Both algorithms are able to execute additional algorithms, \mathcal{A}_{agr} and \mathcal{A}'_{agr} respectively within each node, to distribute and aggregate information about the desired network property, where the network property description replaces the index agr . These functions \mathcal{A}_{agr} and \mathcal{A}'_{agr} can be chosen depending on the aggregating problem at hand. A possible application with choices of \mathcal{A}_{max} and \mathcal{A}'_{max} to compute the max-value in the k -neighborhood of each root node $r \in S$ is shown in Example 3.1.

Example 3.1. (Computing the max-value of each k -Neighborhood of a set of nodes $r \in S$, $S \subseteq V$, with our proposed algorithm.) Each node $r \in S$ starts Algorithm 1 DLGCOMP. While spanning the DLGs \mathcal{L}_r rooted in each node $r \in S$ in Algorithm 1 DLGCOMP no algorithm \mathcal{A}_{max} needs to be specified. In Algorithm 2 DLGAGR the maximum node value needs to be aggregated. Therefore we define algorithm \mathcal{A}'_{max} as follows: \mathcal{A}'_{max} initializes on each node $u \in V$ for each root node r a variable $max_v[r]$ to store the highest value received from any child in \mathcal{L}_r , first the node's own value $val(u)$ is stored in $max_v[r]$ for every $r \in S$. On execution of \mathcal{A}'_{max} with message msg_c and ID r from a child, \mathcal{A}'_{max} stores the value of msg_c in $max_v[r]$ if $val(msg_c) > max_v[r]$. Then $max_v[r]$ is stored in msg_p and sent to the parents of u . After the execution of Algorithm 1 DLGCOMP and 2 DLGAGR each root node r stores in $max_v[r]$ the max-value of its own k -Neighborhood.

Or \mathcal{A}_{sum} and \mathcal{A}'_{sum} to compute the sum of all values in the k -neighborhood of each root node $r \in S$ is provided in Example 3.2.

Example 3.2. (Computing the sum of all values in each k -Neighborhood of a set of nodes $r \in S$, $S \subseteq V$, with our proposed algorithm.) Each node $r \in S$ starts Algorithm 1 DLGCOMP. While spanning the DLGs \mathcal{L}_r rooted in each node $r \in S$ in Algorithm 1 DLGCOMP no algorithm \mathcal{A}_{max} needs to be specified. In Algorithm 2 DLGAGR the sum of all node values in the k -Neighborhood needs to be aggregated. Therefore we define algorithm \mathcal{A}'_{sum} as follows: \mathcal{A}'_{sum} initializes on each node $u \in V$ for each root node r a variable $sum_v[r]$ to store the sum of all values received from any child in \mathcal{L}_r plus its own value $val(u)$. First the node's own value $val(u)$ is stored in $sum_v[r]$ for every $r \in S$. On execution of \mathcal{A}'_{sum} with message msg_c and ID r from a child, \mathcal{A}'_{sum} adds the value of msg_c to $sum_v[r]$. Then $sum_v[r]$ is stored in msg_p and sent to (only) one parent of u . If node u has more than one parent, the message is sent to the parent node with the lowest ID. This results in the Tree Variation in Section 3.2. After the execution of Algorithm 1 DLGCOMP and 2 DLGAGR each root node r stores in $sum_v[r]$ the sum of all values in its own k -Neighborhood.

Remark 3.3 (Time optimality of Algorithms 1 DLGCOMP and 2 DLGAGR). In a setting where edge-weights correspond to the transmission time through the corresponding edge, any solution for an algorithm to aggregate information from all nodes $u \in V$ to one root node r needs at least $\Omega(D_\omega)$ time. Now consider the graph of Figure 3. There we have $|S|$ root nodes $1, \dots, |S|$ connected to a chain of $D - 1$ nodes and assume this chain determines the diameter D_ω . Assume each of the chain nodes stores $|S|$ different values that are encoded using $\mathcal{O}(\log n)$ bits each. Assume each root node i computes an aggregation function based on the i 'th value of each chain node. Due to

congestion, the time until all $|S|$ values of the chain node at distance D_ω to the root nodes arrives at the corresponding root nodes is $|S| + D_\omega$. This can be extended to $|S| + k$ in the unweighted case with $k < D$.

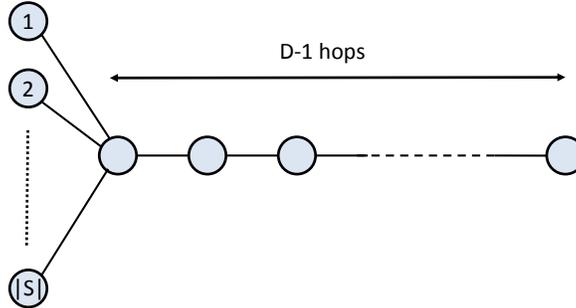


Figure 3: Graph based on the construction used in Remark 3.3.

3.1 Algorithm for Local Multi-Aggregation

First we describe Algorithm 1 DLGCOMP which computes $|S|$ times a DLG with depth k , one for each node in S , and can distribute/broadcast information (depending on the aggregation function at hand) along these computed DLGs. In Algorithm 1 DLGCOMP we extend the S -SP Algorithm of [17], which computes shortest paths. (Note that alternatively one could also modify the k -source detection algorithm by Lenzen and Peleg [22].) In Algorithm 2 DLGAGR, we aggregate information efficiently through the computed DLGs. Algorithm 1 DLGCOMP is designed as a subroutine and has two input parameters $|S|$ and k as well as a function \mathcal{A}_{agr} . Parameter $|S|$ is the number of DLGs \mathcal{L}_r rooted in $r \in S$ which Algorithm 1 DLGCOMP has to compute. Parameter k denotes the depth of the k -neighborhood.

And function \mathcal{A}_{agr} specifies an algorithm executed along with the computation of the $|S|$ DLGs. Further a node $r \in S$ knows that it is in S (e.g. as it is sampled by some procedure or interested in performing an aggregation by itself). Lets consider first only the special case where we compute a single DLG \mathcal{L}_r rooted in node r . We denote by r -message any messages sent during an execution of Algorithm 1 DLGCOMP that belongs to the construction of the DLG rooted in node r (later we keep this term for different roots $r \in S$). Node r starts with sending an r -message to all its neighbors w and in the next time slot, those neighbors w send an r -message to their neighbors, except to r . A DLG \mathcal{L}_r continues to grow as follows: in time slot t all nodes u at distance $\omega(r, u) = t$ receive a r -message from their neighbors w with distance $\omega(r, w) = t - 1$. If the first r -message is received from neighbor w_1 of node u , w_1 is considered a parent of u in \mathcal{L}_r . However, if in the same time slot u receives further r -messages from different neighbors w_i , those are considered parents too. One time slot after receiving the first r -message(s), node u computes an r -message and sends it to all neighbors w_i which are not considered a parent of u .

Now we say a r -message received by u is considered *valid*, if it is sent from a parent of u in \mathcal{L}_r . Each r -message contains the ID of the root node r and the weighted distance $\omega_r := d_\omega(r, u)$ from r to the receiving node u .

To simulate a weighted graph, the messages get delayed in Line 14 corresponding to the corresponding edge weight $\omega(u, v)$, if a message is to be sent from node u to v . For unweighted graphs, the variable l_i (Line 14) contains the lowest ID in L_i which is not in L_{delay} .

While computing \mathcal{L}_r , the ID r and some additional information is stored and sent along with the r -messages. Each node stores time t and parent w from which it received a r -message (Line 25

and 26) . This is needed to efficiently execute Algorithm 2 DLGAGR later on. A node u stores in schedule $\tau[r][h]$ and in $parent_in_T_{r_h}$ arrival times and parent nodes u_h if $u \in N_k(r)$ (Lines 23 to 27 are only executed if $u \in N_k(r_i)$, i.e. if $\omega_i \leq k$, where ω_i was received in a r_i -message). Nonetheless, r_i -messages are still forwarded, even if the distance ω_i in a r_i -messages is larger than k . This is needed to ensure that r -messages are properly delayed.

Furthermore each node u executes algorithm \mathcal{A}_{agr} (Line 27), and can be used to distribute information in \mathcal{L}_r . Algorithm \mathcal{A}_{agr} is executed with parameters (r, msg_p, u_h) each time u receives a valid r -message from neighbor $u_h \in \{u_1, \dots, u_{deg(u)}\}$. Message msg_p is the part of the r -message computed by algorithm \mathcal{A}_{agr} on a parent node u_h of u in \mathcal{L}_r . After the execution for every parent u_h of u in \mathcal{L}_r with the corresponding msg_p , algorithm \mathcal{A}_{agr} computed and stored in $msg_c[r]$ the message which is sent to the children of u in \mathcal{L}_r (Line 15). Furthermore algorithm \mathcal{A}_{agr} is executed once on every node for definition and initialization of additional global variables on the node(Line 11).

With Algorithm 1 DLGCOMP, multiple leveled graphs $\mathcal{L}_{r_1}, \dots, \mathcal{L}_{r_{|S|}}$ start growing in the same time slot. One \mathcal{L}_r for each $r \in S$. This could lead to congestions. To prevent that, every time two Algorithm 1 DLGCOMP messages cross³, the Algorithm 1 DLGCOMP message with the higher ID is delayed one time slot. This is done in Line 32 by putting the higher ID into queue L_{delay} or by re-transmitting the message in the next time slot through the same edge (u_h, u) again, respectively (due to the if-statement in Line 21).

Similar as in the proof of the S-SP algorithm [17], we show that the total delay of any Algorithm 1 DLGCOMP is at most $\mathcal{O}(|S|)$ when $|S|$ many DLGs are constructed in parallel. And we show that despite of the delays we still construct valid leveled graphs.

In Algorithm 2 DLGAGR, information gets aggregated according to Algorithm \mathcal{A}'_{agr} . To do this, the computed leveled graphs of each root node r_j get processed in a bottom-up fashion. Algorithm 2 DLGAGR has three inputs: the number of DLGs $|S|$, the depth k of the k -neighborhood, which are needed to bound the runtime, and an algorithm \mathcal{A}'_{agr} , which can consist of an initialization part and a computation part. The initialization part is executed once on every node $u \in V$ starting the computation(loop of Algorithm 2 DLGAGR, Line 2). In the loop in time slot $t = |S| + k - \tau[r_j][h]$ a node u sends the information computed by \mathcal{A}'_{agr} to its parent $u_h \in \{u_1, \dots, u_{deg(u)}\}$ in \mathcal{L}_{r_j} . Here schedule $\tau[r_j][h]$ is the time when u received a valid r_j -message from neighbor u_h . The schedule τ was computed before in Algorithm 1 DLGCOMP(Line 25).

The computation part of $\mathcal{A}'_{agr}(r, msg_{c,u_h}, u_h)$ is executed on a node $u \in V$ each time u received a message msg_{c,u_h} from a child node u_h in \mathcal{L}_r . When u has received all messages msg_{c,u_h} form all children u_h in \mathcal{L}_r and u has executed \mathcal{A}'_{agr} once for each of those messages, \mathcal{A}'_{agr} computed and stored in $msg_p[r]$ the information that is subsequently sent to all parents of u in \mathcal{L}_r .

Lemma 3.4. *Algorithm 1 DLGCOMP computes $|S|$ valid directed leveled graphs, in time $\mathcal{O}(|S| + D_\omega)$ for $k = D_\omega$.*

Proof. The proof is very similar to the proofs for the S-SP algorithm in [17]. For the *correctness* we state a slightly adapted version of Lemma 12 in [17].

Correctness: First, assume that $S = \{r\}$ and consider the computation of \mathcal{L}_r . At time t , each node u at weighted distance t from r receives a r -message from all of u 's neighbors u_h that are at distance $t - \omega(u_h, u)$ to r . All edges incident to neighbors u_h that sent such a message are added to DLG \mathcal{L}_r , directed from u_h to u .

Now consider the case, where the set $S \supset \{r\}$ contains at least two nodes. We analyze how the computations of other nodes affect the computation of \mathcal{L}_r . We say that a r -message has ‘reached’ a node u through edge (u_h, u) from a neighbor u_h in time slot t , if the message was successfully received and has been removed from the delay queue L_{delay} in time slot t , or it was successfully

³That means both messages are received in the same time slot by the same node or both are sent in the same time slot through the same edge (u_h, u) , one from u to u_h and one in the opposite direction.

Algorithm 1 DLGCOMP computes $|S|$ DLGs in the k -Neighborhood of the root nodes $r \in S$ (executed by node u)

Input: $|S|, k, \mathcal{A}_{agr}$

parameters passed to $\mathcal{A}_{agr}(v, msg_p, u_h)$: message msg_p of a parent u_h of u in DLG \mathcal{L}_v

```

1:  $L := \emptyset$ ;
2:  $L_{delay} := \emptyset$ ;
3:  $\tau := \{\infty, \infty, \dots, \infty\} \times [n]$ ; //  $\tau[v][i]$  = time when message of DLG  $\mathcal{L}_v$  reaches  $u$ , for each parent
    $w_1, \dots, w_{\max} \in P_v[u]$  enumerated in  $i$ 
4:  $sp := \{0, 0, \dots, 0\}$ ; //  $sp[v]$  := number of parents in shortest paths from  $v$  to  $u$ 
5:  $\omega := \{\infty, \infty, \dots, \infty\}$ ; //  $\omega[v] = \omega(v, u)$ ,  $\omega[\emptyset] := \infty$ 
6: if  $u \in S$  then
7:    $L := \{u\}$ ;
8:    $\omega[u] := 0$ ;
9:    $\tau[u][1] := 0$ ;
10:  $L_1, \dots, L_{deg(u)} := L$ ;
11: initialize algorithm  $\mathcal{A}_{agr}$ ;
12: for  $t = 0, \dots, |S| + k$  do
13:   for  $i = 1, \dots, deg(u)$  do
14:      $l_i := \min \{v \in L_i : \tau[v][sp[v]] + \omega(u, v) \geq t \text{ and } v \notin L_{delay}\}$ ; // smallest ID that is not delayed and
       ready to be scheduled
15:     within one time slot:
       send  $(l_1, \omega[l_1] + \omega(u, u_1), msg_c[l_1])$  to neighbor  $u_1$ , receive  $(r_1, \omega_1, msg_{p,1})$  from  $u_1$ ;
       send  $(l_2, \omega[l_2] + \omega(u, u_2), msg_c[l_2])$  to neighbor  $u_2$ , receive  $(r_2, \omega_2, msg_{p,2})$  from  $u_2$ ;
       ...
       send  $(l_{deg(u)}, \omega[l_{deg(u)}] + \omega(u, l_{deg(u)}), msg_c[l_{deg(u)}])$  to neighbor  $u_{deg(u)}$ , receive
        $(r_{deg(u)}, \omega_{deg(u)}, msg_{p,deg(u)})$  from  $u_{deg(u)}$ ;
16:    $R := \{r_i : r_i < l_i \text{ and } i \in \{1, \dots, deg(u)\} \setminus L$ ;
17:    $s := \min(L_{delay})$ ;
18:   if  $s \leq \min(R)$  and  $s < \infty$  then
19:      $L_{delay} := L_{delay} \setminus \{s\}$ ;
20:   for  $i = 1, \dots, deg(u)$  do
21:     if  $r_i < l_i$  then //  $T_{l_i}$ 's message will be delayed due to  $T_{r_i}$ .
22:       if  $\omega_i \leq \omega[r_i]$  then
23:         if  $\omega_i \leq k$  then
24:            $sp[r_i] = sp[r_i] + 1$ ;
25:            $\tau[r_i][sp[r_i]] := t$ ;
26:            $parent\_in\_T_{r_i}[sp[r_i]] := \text{neighbor } u_i$ ;
27:           execute algorithm  $\mathcal{A}_{agr}(r_i, msg_{p,i})$ ;
28:         if  $r_i \notin L$  then
29:            $\omega[r_i] = \omega_i$ ;
30:            $L := L \cup \{r_i\}, L_1 := L_1 \cup \{r_i\}, L_2 := L_2 \cup \{r_i\}, \dots, L_{i-1} := L_{i-1} \cup \{r_i\}, L_{i+1} := L_{i+1} \cup$ 
            $\{r_i\}, \dots, L_{deg(u)} := L_{deg(u)} \cup \{r_i\}$ ;
31:           if  $\min(R) < r_i$  or  $s < r_i$  then
32:              $L_{delay} := L_{delay} \cup \{r_i\}$ ;
33:         else
34:            $L_i := L_i \setminus \{r_i\}$ ;
35:       else
36:          $L_i := L_i \setminus \{l_i\}$ ; //  $T_{l_i}$ 's message was successfully sent to neighbor  $u_i$ .

```

Algorithm 2 DLGAGR computes aggregation function (executed by node u)

Input: $|S|, k, \mathcal{A}'_{agr}$

parameters passed to $\mathcal{A}'_{agr}(v, msg_c, u_h)$: message msg_c of a child u_h of u in DLG \mathcal{L}_v

```

1: initialize algorithm  $\mathcal{A}'_{agr}$ ; // and use same variables as Algorithm 1 DLGCOMP
2: for  $t = 0, \dots, |S| + k$  do
3:   within one time slot:
       foreach  $v \in L, i \in \{1, \dots, sp[v]\}$  such that  $t = |S| + k - \tau[v][i]$  do
           send  $(v, msg_p[v])$  to  $parent\_in\_T_v[i]$ ;
           receive  $(v_1, msg_{c,u_1})$  from neighbor  $u_1$ ;
           receive  $(v_2, msg_{c,u_2})$  from neighbor  $u_2$ ;
           ...
           receive  $(v_{deg(u)}, msg_{c,u_{deg(u)}})$  from neighbor  $u_{deg(u)}$ ;
4:   for  $i = 1, \dots, deg(u)$  do
5:     if  $v_i \neq \text{empty}$  then
6:       execute algorithm  $\mathcal{A}'_{agr}(v_i, msg_{c,u_i})$ ;

```

received in time slot t and not put into the queue L_{delay} at all. It turns out that the first r -messages of \mathcal{L}_r 's computation which reach u are transmitted through the edges at distance $t - \omega(u_h, u)$ as in the case before. Consider two neighbors $u_1, u_2 \in N_1(u)$ of u ; we can ignore the case that u has only one neighbor since this case trivially satisfies our claim. A r -message containing weighted distance $\omega(r, u_j)$, has reached u through the edge (u_1, u) earlier than through edge (u_2, u) if and only if $\omega(r, u_1) < \omega(r, u_2)$. We show this by proving that the set of messages with lower IDs which delay the r -messages is the same for both paths (r, u_1, u) and (r, u_2, u) : Assume that the computation of \mathcal{L}_i is delaying the r -message sent through (r, u_1, u) at some point. Then the i -message reaches, in case the i -message is coming from u 's direction, node u earlier than the r -message. Even if an i -message and a r -message are transmitted in the same time slot to u (through (u_1, u) and (u_2, u) , respectively), the i -message delays the r -message in the node by putting it into the queue L_{delay} , and the i -message reaches u earlier. Thus the i -message also delays the r -message running through path (r, u_2, u) , if it did not already delayed it earlier.

Runtime: In case $S = \{r\}$ contains only one element, the computation of DLG \mathcal{L}_r , which is not delayed by other computations, takes at most D'_ω time steps. Because in time step t all nodes $u \in V$ at distance $\omega(r, u) = t$ receive a r -message, as a consequence in time step D_ω every node $v \in V$ with distance $\omega(r, v) \leq D_\omega$ receives or has received a r -message. Since there are no nodes $u \in V$ with $\omega(r, v) > D_\omega$, the computation of \mathcal{L}_r stops after D'_ω time steps.

We showed that a r -message can be delayed at most one time slot by another i -message with $i < r$, and therefore by at most $|S|$ time slots by all other $|S|$ DLG-constructions. Thus Algorithm 1 DLGCOMP runs in $\mathcal{O}(|S| + D_\omega)$. \square

Lemma 3.5. *Algorithm 2 DLGAGR aggregates information through $|S|$ directed leveled graphs, in time $\mathcal{O}(|S| + D_\omega)$ for $k = D_\omega$.*

Proof. Runtime: In Algorithm 2 DLGAGR, a node u sends only messages to parent nodes in all DLGs $\mathcal{L}_r \forall r \in S$. The schedule to send the messages is the same as in Algorithm 1 DLGCOMP stored in τ , but reversed. Therefore the runtime of Algorithm 2 DLGAGR is bound by the runtime of Algorithm 1 DLGCOMP, which is $\mathcal{O}(|S| + D_\omega)$. The complexity of executing algorithms \mathcal{A}_{agr} and \mathcal{A}'_{agr} inside of a node causes no additional communication. Since we are only interested in communication complexity, the total runtime is $\mathcal{O}(|S| + D_\omega)$.

Correctness: Since we use the schedule τ of Algorithm 1 DLGCOMP in reverse order to send messages, it is guaranteed that a node u receives all messages from all children in a DLG \mathcal{L}_r before

u sends the first message to any parent in \mathcal{L}_r . □

Lemma 3.6. *Algorithms 1 DLGCOMP and 2 DLGAGR aggregate information in $\mathcal{O}(|S| + k)$ time for $k \leq D_\omega$.*

Proof. Runtime: In the k -Neighborhood restricted Algorithm 1 DLGCOMP a r -message which is not delayed by other IDs, needs k time slots to reach all nodes $u \in N_k(r)$. We showed in Lemma 3.4 that a r -message can be delayed at most $|S|$ time slots due to other ID's. Hence, to reach all nodes in its k -neighborhood an r -message needs at most $|S| + k$ time slots. For aggregating information Algorithm 2 DLGAGR uses schedule τ , which was computed in the k -Neighborhood restricted Algorithm 1 DLGCOMP and therefore is bound by the same number of time slots $|S| + k$.

Correctness: As long as a distance ω_i in a r -message is less or equal to k , the message is processed the same way as in the restricted Algorithm 1 DLGCOMP ($k = D_\omega$). When ω_i is larger than k , then r -messages with ID r do no longer extend the DLG \mathcal{L}_r , but just delay all i -message with higher ID ($i > r$) the same way as in the restricted Algorithm 1 DLGCOMP. As a consequence after the execution of the k -Neighborhood restricted Algorithm 1 DLGCOMP, all nodes $u \in N_k(r)$ have the same information of DLG \mathcal{L}_r stored as if the restricted Algorithm 1 DLGCOMP has been executed. All nodes $v \notin N_k(r)$ that still send r -messages do not contribute to \mathcal{L}_r 's aggregation but are necessary to ensure that the other tree's computations are delayed and trees are constructed correctly.

By executing Algorithm 2 DLGAGR a node u sends aggregation messages $msg_p[v]$ back to a root node r if and only if $u \in N_k(r)$, as mentioned above. Therefore Algorithm 2 DLGAGR just aggregates information of the neighborhood $N_k(r)$ to a root node r . □

Proof of Theorem 1.1. Follows by combining Lemma 3.4 and 3.5 and Lemma 3.6. □

3.2 Tree Variation

In Algorithm 2 DLGAGR we can aggregate information along all shortest r - u -paths between a root node $r \in S$ and a node $u \in V$. For some applications it is desirable to aggregate information only along one shortest r - u -path, as for example in max-value/average-value aggregation or in a S -Shortest Paths to approximate problems such as closeness centrality. That means a node u sends a message only to one parent in a DLG \mathcal{L}_r while executing Algorithm 2 DLGAGR. The result is the same as when aggregating information through a tree T_r rooted in r (instead of a DLG \mathcal{L}_r). For completeness we show an adaptation of Algorithms 1 DLGCOMP and 2 DLGAGR which computes and aggregates along trees instead of DLGs. We provide a detailed description of the Tree Variation in Appendix 4.4.

4 Applications

4.1 Betweenness Centrality

Brandes showed in [7], an algorithm for computing betweenness centrality recursively. Let $\delta_{st}(v) := \sigma_{st}(v)/\sigma_{st}$ be the ratio of all shortest path between s and t that contain node v compared to all shortest paths between s and t . We denote by $\delta_{s*}(v)$ the dependency of a node $s \in V$ on another node $v \in V$, defined by $\delta_{s*}(v) := \sum_{t \in V, t \neq v \neq s} \delta_{st}(v)$. The dependency can be calculated recursively as Brandes [7] stated:

Theorem 4.1. *(Recursive Betweenness Centrality Dependency, Brandes [7] Theorem 6)*

$$\delta_{s*}(v) = \sum_{w: v \in P_s(w)} \frac{\sigma_{sv}}{\sigma_{sw}} \cdot (1 + \delta_{s*}(w))$$

Bader et al. showed in [3] an adaptive sampling algorithm, which approximates the betweenness centrality with significantly reduced number of single-source shortest path computations for a node with high centrality when using Brandes' recursive algorithm.

Theorem 4.2. (Bader et al [3] Theorem 3). For $0 < \varepsilon < \frac{1}{2}$, if the centrality of a vertex v is n^2/t for some constant $t \geq 1$, then with probability $\geq 1 - 2\varepsilon$ its betweenness centrality can be estimated to within a factor of $1/\varepsilon$ with $t\varepsilon$ samples of source vertices.

This algorithm can be adapted efficiently in a distributed setting by using Algorithms 1 DLGCOMP and 2 DLGAGR. We define an Algorithm 3BCSETUP to perform multiple rounds as suggested by Bader in [3]. In each round, Algorithms 1 DLGCOMP and 2 DLGAGR are executed with algorithms \mathcal{A}_{BC} and \mathcal{A}'_{BC} as defined in Algorithm 4 and 5. With \mathcal{A}_{BC} and \mathcal{A}'_{BC} we state the Algorithms \mathcal{A}_{agr} and \mathcal{A}'_{agr} of our framework specified for betweenness centrality approximation (BC) being the function agr . In contrast to Bader, our algorithm samples not just the betweenness centrality dependency of one node in each round, but of multiple nodes. Furthermore, Bader's algorithm concentrates on one node $v \in V$ and stops sampling if $BC(v)$ can be approximated within an error $(\times, 1 + \frac{1}{\varepsilon})$ with probability $\geq 1 - 2\varepsilon$. Our Algorithm 3 BCSETUP considers all nodes in the graph and stops after at least \hat{n} nodes have been approximated within a similar multiplicative error $(\times, 1 + \frac{1}{\varepsilon'})$ with probability $\geq 1 - 2\varepsilon'$ (where $1/n^c < \varepsilon' = \varepsilon - 1/n^c < \varepsilon < 1/2$).

4.1.1 Algorithm for Betweenness Centrality

In Algorithm 3 BCSETUP, the idea is to select in multiple rounds multiple sample nodes $s \in S$ and calculate the betweenness centrality dependency $\delta_{s*}(u)$ on all other nodes $u \in V$ for each $s \in S$. The algorithm stops if more than \hat{n} nodes in V with high betweenness centrality are found, \hat{n} is an input parameter. Set S_i is the set of nodes which are sampled in round i , all sets S_i are disjoint and form together the set $S = S_1 \cup S_2 \cup \dots \cup S_{max}$.

During the algorithm each node $u \in V$ stores the sum of all the dependencies $\delta_{s*}(u)$ of the nodes $s \in S$ sampled so far in variable s_{BC} . Furthermore the number k_i of sample nodes $s \in S_1 \cup S_2 \cup \dots \cup S_i$ that were sampled so far (in any of rounds $\{1, 2, \dots, i\}$), need to be stored to indicate whether the node itself is sampled ($is_{sampled}$).

All centralized communication to and from node 1 uses tree T_1 . In round i (Line 13 to 24) node 1 calculates the probability p_s (according to the proof in Theorem 1.3)

with which each node gets sampled. The value of p_s is broadcasted in the network and each node decides if itself is a sample node and reports back to node 1 if so. In Line 23, if a node in T_1 receives multiple messages in one time slot, it sends the sum of the message value to its parent in T_1 . The number of sample nodes $|S_i|$ gathered by node 1 is then broadcasted again, this is needed to determine the runtime $|S_i| + D_\omega$ in Algorithms 1 DLGCOMP and 2 DLGAGR and to maintain the number of sample nodes k .

Next, all nodes simultaneously start to execute Algorithm 1 DLGCOMP with algorithm \mathcal{A}_{BC} , whose pseudocode is specified in Algorithm 4. In round i Algorithm 1 DLGCOMP computes a DLG \mathcal{L}_s for every sample node $s \in S_i$. And \mathcal{A}_{BC} ensures that while computing the DLGs all nodes $u \in V$ compute the total number of shortest u - s -paths (tsp) to all DLG root nodes s . Therefore \mathcal{A}_{BC} needs to initialize the array $total_{sp}$ on each node $u \in V$, which stores at $total_{sp}[s]$ the number of shortest path between u and s (Lines 1 to 4). A sample node s initializes the distance to itself with 1

During the execution of Algorithm 1 DLGCOMP in node u , algorithm \mathcal{A}_{BC} processes messages (msg_p) from parents $u_h \in \{u_1, \dots, u_{deg(u)}\}$ in DLG \mathcal{L}_s , which contains the value tsp from u_h to s . Algorithm \mathcal{A}_{BC} then updates the corresponding entry $total_{sp}[s]$ by adding the value tsp received from parent u_h . Entry $total_{sp}[s]$ stores the correct value after all messages of parents in \mathcal{L}_s have

Algorithm 3 BCSETUP: Distributed Betweenness Centrality Approximation (executed by node u)

Input: τ_c : threshold criterion, \hat{n} : number of nodes which have to reach the threshold τ_c

```

1: global  $s_{BC} := 0$ ; // sum of dependency scores  $\delta_{v^*}(u)$ ,  $v \in V \setminus \{u\}$ 
2: global  $k := 0$ ; // number of samples so far
3: global  $is_{sampled} := 0$ ;
4:  $BC(u) := \perp$ ; // approximated betweenness centrality of  $u$ , initialized as undefined
5: if  $u = 1$  then
6:   estimate  $D'_h$  and  $D'_\omega$  by generating a spanning tree  $T_1$  and using Fact 2.7;
7:   broadcast  $D'_h$  and  $D'_\omega$  on  $T_1$ ;
8:    $n_{BC} := 0$ ; // number of high BC nodes found so far
9: else
10:  wait until value of  $D'_h$  and  $D'_\omega$  received;
11:  for  $T := \log(D'_h), \log(D'_h) + 1, \dots, \log(n)$  do
12:     $t := 2^T$ ;
13:    if  $u = 1$  then
14:      broadcast  $p_s := \frac{t}{n}$  on  $T_1$ ;
15:      wait for responses  $r_i$  for  $2D'_\omega$  time slots;
16:       $|S_i| := \sum_i r_i$ ; // number of samples in this round
17:      broadcast  $|S_i|$  through  $T_1$ ;
18:    else
19:      receive value of  $p_s$ ;
20:      if  $is_{sampled} = 0$  then
21:        set  $gets_{sampled} := 1$  with probability  $p_s$ ;
22:        if  $gets_{sampled} = 1$  then
23:          send (“1”) to node 1 by  $T_1$ ; // accumulate messages
24:          receive value of  $|S|$ ;
25:          start Algorithm 1 DLGCOMP with  $(|S_i|, D'_\omega, \mathcal{A}_{BC}())$ ;
26:          start Algorithm 2 DLGAGR with  $(|S_i|, D'_\omega, \mathcal{A}'_{BC}())$ ;
27:           $k := k + |S_i|$ ;
28:          if  $gets_{sampled} = 1$  then
29:             $is_{sampled} := 1$ ;
30:             $gets_{sampled} := 0$ ;
31:          if  $s_{BC} \geq n \cdot \tau_c$  and  $BC(u) = \perp$  then
32:             $BC(u) := \frac{n \cdot s_{BC}}{2k}$ ; // divided by two, since G is a undirected graph
33:            send (“reached threshold”, 1) to node 1 by using paths in  $T_1$ ; // accumulate messages
34:          if  $u = 1$  then
35:            wait to receive message(s)  $M$  for up to  $D'_\omega$  time slots;
36:             $n_{BC} := n_{BC} + |M|$ ;
37:            if  $n_{BC} \geq \hat{n}$  then
38:              broadcast “terminate”;

```

Algorithm 4 $\mathcal{A}_{BC}()$

```
1: // INITIALIZATION
2: global  $total_{sp} := \{0, 0, \dots, 0\}$ ; //  $total_{sp}[v] :=$  total shortest paths from  $v$  to  $u$ 
3: if  $gets_{sampled} = 1$  then
4:    $total_{sp}[u] := 1$ ;

5: // COMPUTATION,  $f(v, msg_p)$ :  $v$  root node ID of the message,  $msg_p$  message of a parent in
    $\mathcal{L}_v$ 
6:  $total_{sp}[v] := total_{sp}[v] + msg_p$ ;
7:  $msg_c[v] := total_{sp}[v]$ ; // output
```

arrived. Node u then forwards $total_{sp}[s]$ during the execution of Algorithm 1 DLGCOMP in Line 15 to all its children in \mathcal{L}_s (encoded in message $msg_c[s]$).

In Algorithm 2 DLGAGR each node $u \in V$ adds the dependencies $\delta_{s^*}(u)$ of all sampled nodes $s \in S_i$ on itself to the sum s_{BC} . Each time a child node w in a DLG \mathcal{L}_s sends a message to u , algorithm \mathcal{A}'_{BC} , whose pseudocode is specified in Algorithm 5, gets executed. The message ($msg_c = (\delta_c, tsp_c)$) of child w in \mathcal{L}_s contains the dependency $\delta_{s^*}(w)$ and the tsp from w to s . The dependencies $\delta_{s^*}(u)$ are then calculated recursively with the formula in Theorem 4.1 and stored in $\delta_{BC}[s]$. Entry $\delta_{BC}[s]$ stores the correct value after all messages of children in \mathcal{L}_s have arrived. Node u then forwards $(\delta_{BC}[s], total_{sp}[s])$ during the execution of Algorithm 2 DLGAGR in Line 3 to all its parents in \mathcal{L}_s (in message $msg_p[s]$). In a leaf node l of DLG \mathcal{L}_s , Algorithm \mathcal{A}'_{BC} never gets executed and therefore forwards the initialization value of $\delta_{BC}[s] = 0$ to its parents in \mathcal{L}_s .

Algorithm 5 $\mathcal{A}'_{BC}()$

```
1: // INITIALIZATION
2: global  $\delta_{BC} := \{0, \dots, 0\}$ ; //  $\delta_{BC}[s] :=$  dependency of the source vertex  $s$  on  $u$ .

3: // COMPUTATION,  $g(v, msg_c)$ :  $v$  root node ID of the message,  $msg_c$  message of a child in  $\mathcal{L}_v$ 
4:  $(\delta_c, tsp_c) := msg_c$ ;
5:  $\delta := (total_{sp}[v_i]/tsp_c) \cdot (1 + \delta_c)$ ;
6:  $\delta_{BC}[v_i] := \delta_{BC}[v_i] + \delta$ ;
7:  $s_{BC} := s_{BC} + \delta$ ;
8:  $msg_p[v] = (\delta_{BC}[v], total_{sp}[v])$ ; // output
```

After the execution of Algorithm 2 DLGAGR each node checks in Line 31 if s_{BC} passed the threshold $n \cdot \tau_c$ and reports back to node 1 if this is the case. Node 1 adds the number of passed threshold to n_{BC} , a variable that indicates how many nodes passed the threshold so far. In [3] it is stated that $\tau_c = 5$ yields a good trade off between computation (time) and approximation quality. To prevent congestion when multiple nodes reached the threshold and therefore multiple messages are sent back to node 1, we use the sum-aggregation of our framework as seen in Example 3.2.

If not enough nodes reached the threshold so far, i.e. $n_{BC} < \hat{n}$, another round with additional sample nodes gets initiated by node 1.

Remark 4.3. *Note, that in an undirected graph the calculated value of $BC(u)$ has to be divided by 2, since for every node pair $(s, t) \in V \times V$ the betweenness dependency of the shortest s - t -paths on u is computed twice, once as a part of $\delta_{s^*}(u)$ and once as a part of $\delta_{t^*}(u)$.*

Lemma 4.4. *Algorithm 3 BCSETUP computes $\delta_{s*}(u)$ for each node $u \in V$ and all samples $s \in S_i$ in time $\mathcal{O}(|S_i| + D_\omega)$ in round i , with a multiplicative approximation guarantee in the range of $[1 - \frac{1}{n^c}, 1 + \frac{1}{n^c}]$.*

Proof. Runtime: The computation of broadcasting and aggregating p_s and $|S_i|$ in Lines 13 to 24 takes $\mathcal{O}(D_\omega)$ time slots. The execution of Algorithms 1 DLGCOMP and 2 DLGAGR is bounded when using our framework by $|S_i| + D'_\omega$, see Theorem 1.1. Therefore the runtime of Algorithm 3 BCSETUP is bound by $\mathcal{O}(|S_i| + D_\omega)$.

Correctness: As defined in Theorem 4.1, to calculate $\delta_{s*}(u)$, it suffices that a node u needs to know the number of shortest u - s -paths σ_{us} and for every child u_h of u in \mathcal{L}_s the number of u_h - s -paths $\sigma_{u_h s}$ and the child's $\delta_{s*}(u_h)$. The number of shortest u - s -paths of a node u can be obtained by adding the number of shortest paths of all parents in \mathcal{L}_s together, which is done in \mathcal{A}_{BC} along the execution of Algorithm 1 DLGCOMP. A leaf u_l in \mathcal{L}_s has a betweenness dependency of zero per definition. Therefore Algorithm 2 DLGAGR and \mathcal{A}'_{BC} can start aggregating the missing parameters in the leaves of \mathcal{L}_s . At the end of the execution of Algorithm 2 DLGAGR in Line 26 of Algorithm 3 BCSETUP, each node $u \in V$ has the information to calculate $\delta_{s*}(u)$.

While executing Algorithm 1 DLGCOMP to compute a $DLG \mathcal{L}_s$, a node u sends the total number of shortest paths (tsp) between s and u to its children. As the graph in Figure 4 shows, this number can be as large as $\mathcal{O}(2^n)$ and therefore $\mathcal{O}(n)$ bits are needed to transmit the exact number. In order to use not more than $\mathcal{O}(\log n)$ bits per message, we approximated tsp by using a $(4 + c) \log n$ bits: $(3 + c) \log(n)$ bit for the most significant bits b of tsp and $\log(n)$ bits to encode the exponent E of tsp , that is $tsp = b \cdot 2^E$. Therefore the approximation of tsp has a multiplicative approximation guarantee in the range of $[1 - \delta, 1 + \delta]$ for the betweenness centrality with $|\delta| \leq \frac{1}{n^c}$. We start with the proof for the case of an unweighted graph. And state later in Lemma 4.6 that it holds too for weighted graphs.

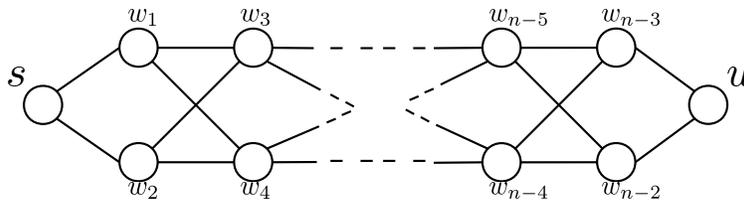


Figure 4: Graph with $2^{\frac{n-2}{2}}$ shortest s - u -paths

Let $tsp_k(u)$ be the exact and $\overline{tsp}_k(u)$ the approximated number of total shortest paths from a sample node s to u , where k is the hop distance of u to s .

In Equation (1) the number of shortest paths from a node u to itself is set to 1. This is done for simplicity—note that the actual number is 0—as it ensures that the number of shortest paths from u to any neighbor is (at least) 1. And since there is only one shortest path between two nodes which are one hop apart, consisting of the edge between them, Equation (2) follows. The exact computation of $tsp(u)$, which is dependent of the values received from parents of u , is shown in Equation (3). It is the sum of all shortest path from s to all parents of u .

$$tsp_0(s) = 1 \tag{1}$$

$$tsp_1(u) = 1 \tag{2}$$

$$tsp_k(u) = \sum_{v \in P(u)} tsp_{k-1}(v) \tag{3}$$

We proof by induction a lower bound for the approximation error at a node u dependent of the hop distance k from u to the root node s .

$$\textbf{Claim:} \quad \overline{tsp}_k(u) \geq (1 - \varepsilon)^{k-1} \cdot tsp_k(u) \quad (4)$$

$$\textbf{Base:} \quad \overline{tsp}_1(u) \geq (1 - \varepsilon)^0 \cdot tsp_0(u) = 1 \cdot 1 = 1 \quad (5)$$

$$\begin{aligned} \textbf{Step:} \quad \overline{tsp}_{k+1}(u) &\geq (1 - \varepsilon) \cdot \left(\sum_{v \in P(u)} \overline{tsp}_k(v) \right) \\ &\geq (1 - \varepsilon) \cdot \left(\sum_{v \in P(u)} (1 - \varepsilon)^{k-1} tsp_k(v) \right) \\ &= (1 - \varepsilon)^k \cdot \left(\sum_{v \in P(u)} tsp_k(v) \right) \\ &= (1 - \varepsilon)^k \cdot tsp_{k+1}(u) \quad \square \quad (6) \end{aligned}$$

From Equation (4) we derive a maximum lower error bound independent of k for all nodes $u \in V$, as shown in Equation (7), using the fact that $k \leq D_h$.

$$\overline{tsp}_k(u) \geq (1 - \varepsilon)^{D_h} \cdot tsp_k(u) \quad \forall 1 \leq k \leq D_h \quad (7)$$

Now let $dep(u)$ be the exact and $\overline{dep}(u)$ the approximated betweenness dependency of a sample node s on u , where k is the hop distance of u to s . Let us assume tsp got approximated at every node with the maximum error $(1 - \varepsilon_m) = (1 - \varepsilon)^{D_h}$. All nodes without any children in DLG_s (leaves), e.g. nodes at maximum hop distance of s , have no betweenness dependency of s , since there are no shortest paths going *through* them, this proves Equation (8). Then the exact recursive computation of dep is shown in Equation (9).

$$dep_{D_h}(u) = dep_{leaf}(u) = 0 \quad (8)$$

$$dep_k(u) = \sum_{w:u \in P(w)} \frac{tsp_k(u)}{tsp_{k+1}(w)} \cdot (1 + dep_{k+1}(w)) \quad (9)$$

Now we derive error bounds by evaluating the term $\frac{\overline{tsp}_i}{tsp_{i+1}}$. The maximum errors occur at $\frac{\overline{tsp}_i}{tsp_{i+1}} = \frac{tsp_i \cdot (1 - \varepsilon_m)}{tsp_{i+1}}$ (i) and at $\frac{\overline{tsp}_i}{tsp_{i+1}} = \frac{tsp_i}{tsp_{i+1} \cdot (1 - \varepsilon_m)}$ (ii). With (i) and (ii) we show by induction an upper and lower bound for the multiplicative error of \overline{dep}_k in Equations (10) to (15).

Proof by Induction on the lower approximation error bound of \overline{dep} :

$$\textbf{Claim:} \quad \overline{dep}_k(u) \geq (1 - \varepsilon_m)^{D_h - k} \cdot dep_k(u) \quad (10)$$

$$\begin{aligned} \textbf{Base:} \quad \overline{dep}_{D_h-1}(u) &= \sum_{w:u \in P(w)} \frac{\overline{tsp}_{D_h-1}(u)}{tsp_{D_h}(w)} \cdot (1 + 0) \\ &\geq \sum_{w:u \in P(w)} \frac{tsp_{D_h-1}(u) \cdot (1 - \varepsilon_m)}{tsp_{D_h}(w)} = (1 - \varepsilon_m) \cdot dep_{D_h-1}(u) \quad (11) \end{aligned}$$

$$\begin{aligned}
\text{Step: } \quad \overline{dep}_{k-1}(u) &= \sum_{w:u \in P(w)} \frac{\overline{tsp}_{k-1}(u)}{\overline{tsp}_k(w)} \cdot (1 + \overline{dep}_k(w)) \\
&\geq \sum_{w:u \in P(w)} \frac{tsp_{k-1}(u) \cdot (1 - \varepsilon_m)}{tsp_k(w)} \cdot \left(1 + (1 - \varepsilon_m)^{D_h - k} \cdot dep_k(w)\right) \\
&\geq (1 - \varepsilon_m)^{D_h - (k-1)} \cdot \sum_{w:u \in P(w)} \frac{tsp_{k-1}(u)}{tsp_k(w)} \cdot (1 + dep_k(w)) \\
&= (1 - \varepsilon_m)^{D_h - (k-1)} \cdot dep_{k-1} \quad \square \quad (12)
\end{aligned}$$

Proof by Induction on the upper approximation error bound of \overline{dep} :

$$\text{Claim: } \quad \overline{dep}_k(u) \leq (1 - \varepsilon_m)^{k - D_h} \cdot dep_k(u) \quad (13)$$

$$\begin{aligned}
\text{Base: } \quad \overline{dep}_{D_h-1}(u) &= \sum_{w:u \in P(w)} \frac{\overline{tsp}_{D_h-1}(u)}{\overline{tsp}_{D_h}(w)} \cdot (1 + 0) \\
&\leq \sum_{w:u \in P(w)} \frac{tsp_{D_h-1}(u)}{tsp_{D_h}(w) \cdot (1 - \varepsilon_m)} = (1 - \varepsilon_m)^{-1} \cdot dep_{D_h-1}(u) \quad (14)
\end{aligned}$$

$$\begin{aligned}
\text{Step: } \quad \overline{dep}_{k-1}(u) &= \sum_{w:u \in P(w)} \frac{\overline{tsp}_{k-1}(u)}{\overline{tsp}_k(w)} \cdot (1 + \overline{dep}_k(w)) \\
&\leq \sum_{w:u \in P(w)} \frac{tsp_{k-1}(u)}{tsp_k(w) \cdot (1 - \varepsilon_m)} \cdot \left(1 + (1 - \varepsilon_m)^{k - D_h} \cdot dep_k(w)\right) \\
&\leq (1 - \varepsilon_m)^{(k-1) - D_h} \cdot \sum_{w:u \in P(w)} \frac{tsp_{k-1}(u)}{tsp_k(w)} \cdot (1 + dep_k(w)) \\
&= (1 - \varepsilon_m)^{(k-1) - D_h} \cdot dep_{k-1} \quad \square \quad (15)
\end{aligned}$$

From the upper and lower approximation error bounds in a node u with a certain hop distance k to the root node, we can derive in Equations (10) and (13) a maximum error range for any node at any distance independent of k , as shown in Equation (16).

$$\begin{aligned}
(1 - \varepsilon_m)^{D_h - 1} \cdot dep_k(u) &\leq \overline{dep}_k(u) \leq (1 - \varepsilon_m)^{1 - D_h} \cdot dep_k(u) \\
(1 - \varepsilon)^{D_h^2} \cdot dep_k(u) &\leq \overline{dep}_k(u) \leq (1 - \varepsilon)^{-D_h^2} \cdot dep_k(u) \quad \forall 1 \leq k \leq D_h \quad (16)
\end{aligned}$$

By using the fact, that $D_h \leq n$ we derive in Equations (17) to (19) a tight bound for $|\delta| \leq \frac{1}{n}$.

$$(1 + \delta) \geq (1 - \varepsilon)^{D_h^2} \underbrace{\geq}_{\varepsilon = \frac{1}{n^{3+c}} \leq \frac{1}{n^c \cdot D_h^2}} \left(1 - \frac{1}{n^c \cdot D_h^2}\right)^{D_h^2} \approx 1 - \frac{1}{n^c} \quad (17)$$

$$(1 + \delta) \leq (1 - \varepsilon)^{-D_h^2} \underbrace{\leq}_{\varepsilon = \frac{1}{n^{3+c}} \leq \frac{1}{n^c \cdot D_h^2}} \left(1 - \frac{1}{n^c \cdot D_h^2}\right)^{-D_h^2} \approx 1 + \frac{1}{n^c} \quad (18)$$

$$-\frac{1}{n^c} \leq \delta \leq \frac{1}{n^c} \quad \square \quad (19)$$

Corollary 4.5. *For betweenness centrality on unweighted graphs we obtain an approximation of the total shortest paths which is guaranteed to be in the range of a $[1 - \delta, 1 + \delta]$ multiplicative factor, where $|\delta| \leq \frac{1}{n^c}$, when assuming that the total shortest path error is multiplicatively bounded to be smaller than $\frac{1}{n^{3+c}}$.*

Lemma 4.6. *The approximation guarantee of a factor within the range $[1 - \delta, 1 + \delta]$ with $|\delta| \leq \frac{1}{n^c}$ of the betweenness centrality for unweighted graphs holds for weighted graphs, too.*

To prove this, we need to introduce the shortest path diameter D_{sp} .

Definition 4.7 (Shortest path diameter). $D_{sp}(G) := \max_{u,v \in V} \{|P| \mid P \text{ is a shortest weighted } u\text{-}v\text{-path}\}$ of a graph G is the maximum number of hops of a shortest weighted path between any two nodes of the graph.

Note that $1 \leq D_h \leq D_{sp} \leq n$.

Proof. The maximum hop distance between any leaf l in a DLG and its root node r can be longer than D_h . Because a shortest weighted path is not bounded by D_h hops, but by D_{sp} as defined in Definition 4.7. That changes the maximum exponents in Equations (7), (10) and (13) from D_h to D_{sp} . Since D_{sp} is smaller than n too, the error bounds in Equations (17) and (18) do not change. \square

Proof of Theorem 1.3. Runtime: The computation of D_h and D_ω in Lines 5 to 10 takes $\mathcal{O}(D_\omega)$ time slots. Each round takes $\mathcal{O}(|S_i| + D_\omega)$ time as shown in Lemma 4.4. And for sampling $|S|$ nodes we need $\log \frac{|S|}{D_h}$ rounds, since we start with D_h sample nodes and double this value in every round. This leads to a total time complexity of $\mathcal{O}\left(|S| + D_\omega \log \frac{|S|}{D_h}\right)$.

Correctness: Due to Theorem 4.2, the betweenness centrality $BC(v)$ of a node v can be estimated with probability $\geq 1 - 2\varepsilon$ by sampling the dependency $\delta_{s^*}(v)$ of $t\varepsilon = |S|$ root nodes $s \in S$, if $BC(v) \geq \frac{n^2}{t}$. The approximation error given in Theorem 4.2 is $(\times, 1 + \frac{1}{\varepsilon})$.

As shown in Lemma 4.4, Algorithm 3 BCSETUP approximates those dependencies $\delta_{s^*}(v)$ within a with a multiplicative approximation guarantee in the range of $[1 - \frac{1}{n^c}, 1 + \frac{1}{n^c}]$. In Line 31 and 37, Algorithm 3 BCSETUP ensures that $s_{BC} \geq n \cdot \tau_c$ root nodes have been sampled to approximate the value of $BC(v)$ of a node v such that the betweenness centralities of at least \hat{n} nodes are approximated close enough before terminating. Combining both approximation errors and noting that $\frac{1}{n^c} < \frac{1}{\varepsilon}$, leads to a total error of $(\times, 1 + \frac{1}{\varepsilon'})$. \square

Corollary 4.8. *Let $BC(v)$ be the betweenness centrality of node v . Then $\mathcal{O}\left(\frac{n^2}{BC(v)} + D_\omega \log \frac{n}{BC(v) \cdot D_h}\right)$ rounds of communication are necessary to compute an $(\times, 1 + \frac{1}{\varepsilon'})$ -approximation of $BC(v)$. Furthermore, in the same time our Algorithm 3 BCSETUP finds w.h.p all other nodes with higher betweenness centrality than v .*

4.2 Closeness Centrality

Eppstein and Wang showed in [12] a fast approximation algorithm to derive the closeness centrality of a node by sampling only a few nodes S and computing shortest paths from S to all other nodes (S -SP) instead of calculating all pairs shortest paths (APSP). These ideas can be implemented efficiently in a distributed setting by using Algorithm 1 DLGCOMP of our proposed framework. The Algorithm 6 first computes T_1 with a breath-first-search algorithm starting in node 1, then the value of ε gets broadcasted through T_1 in the network. When a node has received ε , it joins S as a sample node with probability $\frac{\log(n)}{n \cdot \varepsilon^2}$. The number of nodes joining set S are collected

and broadcasted in Lines 7 to 14. In Line 8, if a node $u \in V$ receives multiple messages in one time slot, then in the next time slot node u sends the sum of the values of all messages to its parent in T_1 . Now, in line 15, each node executes Algorithm 1 DLGCOMP. No algorithm \mathcal{A}_{agr} or \mathcal{A}'_{agr} specified for Closeness Centrality is needed, as Algorithm 1 DLGCOMP already computes the required distances. After Algorithm 1 DLGCOMP is executed, every node knows its distances $\omega[r_i]$ to all sample nodes $r_i \in S$ and therefore the number of samples $|S|$. Using those distances, every node u can approximate its closeness centrality CC_u .

Algorithm 6 \mathcal{A}_{CC} Approximation (executed by node u)

Input: ε

```

1: if  $u = 1$  then
2:   compute tree  $T_1$  and estimate  $D'_\omega$  by using Fact 2.7;
3:   broadcast values of  $\varepsilon$  and  $D'_\omega$  on  $T_1$ ;
4: else
5:   wait until values of  $\varepsilon$  and  $D'_\omega$  received;
6:  $u$  joins  $S$  with probability  $\frac{\log(n)}{n \cdot \varepsilon^2}$ ;
7: if  $u \in S$  then
8:   send ("1") to node 1 by  $T_1$ ; // accumulate messages
9: if  $u = 1$  then
10:  wait for responses  $r_i$ ;
11:   $|S| := \sum_i r_i$ ;
12:  broadcast  $|S|$  through  $T_1$ ;
13: else
14:  receive value of  $|S|$ ;
15: execute Algorithm 1 DLGCOMP with  $(|S|, D'_\omega, \perp)$ ;
16:  $CC_u := 1 / \sum_{i=1}^{|S|} \frac{n \cdot \omega[r_i]}{|S| \cdot (n-1)}$  //**  $\omega[r_i]$ ; is the distance of  $u$  to sample node  $r_i$ 

```

Lemma 4.9. *Algorithm 6 approximates closeness centrality of all nodes with an inverse additive error of εD_h in time $\mathcal{O}\left(\frac{\log(n)}{\varepsilon^2} + D_h\right)$ in unweighted graphs.*

Proof. As stated by Epstein and Wang in [12], by using $k = \mathcal{O}\left(\frac{\log(n)}{\varepsilon^2}\right)$ samples, we can approximate the closeness centrality w.h.p. with an inverse additive error of εD_h . As shown in Theorem 1.1, Algorithm 1 DLGCOMP takes $\mathcal{O}(|S| + D_\omega)$ time to complete. Computing, estimating and/or broadcasting T_1 , D'_ω and ε takes in each case $\mathcal{O}(D_\omega)$ time. Algorithm 6 therefore takes $\mathcal{O}\left(\frac{\log(n)}{\varepsilon^2} + D_h\right)$ time for unweighed graphs, since $D_h = D_\omega$ in unweighted graphs. \square

Proof of Theorem 1.4. In the proof of Epstein and Wang in [12], all D_h can be substituted by D_ω and the proof is still valid. Hence, by using $k = \mathcal{O}\left(\frac{\log(n)}{\varepsilon^2}\right)$ samples we can approximate the closeness centrality w.h.p. with an inverse additive error of εD_ω .

As shown in Lemma 4.9 Algorithm 6 takes $\mathcal{O}(|S| + D_\omega)$ time to complete. Thus, by sampling $k = \mathcal{O}\left(\frac{\log(n)}{\varepsilon^2}\right)$ root nodes, Algorithm 6 needs time $\mathcal{O}\left(\frac{\log(n)}{\varepsilon^2} + D_\omega\right)$. \square

4.3 Minimum Routing Cost Tree

The S-MRCT-Algorithm in [15] is executed in a similar fashion as our proposed Algorithms 1 DLGCOMP and 2 DLGAGR, as we generalized the aggregation concept of the S-MRCT-Algorithm [15].

Therefore the S-MRCT-Algorithm as stated in [15] can be simulated with the Tree Variation of our proposed framework, as stated in Algorithms 9 TREECOMPUTING and 10 TREEAGGREGATING .

The S-MRCT-Algorithm in [15] computes a $(\times, 2)$ -approximation of the S-Minimum Routing Cost Spanning Tree (S-MRCT) problem. It relies on the recursive computation of routing costs as stated in Lemma 4.10. The number of nodes of the set S which are in the subtree of a child c_i is denoted as $|Z_{(u,c_i)}^2(T)| = |\{v \in S : v \in T|_{c_i}\}|$ and $|Z_{(u,c_i)}^1(T)| = |\{v \in S : v \notin T|_{c_i}\}|$ denotes the number of all other nodes in set S .

Lemma 4.10 (Lemma 5.6 in [15]).

$$rc_S(T, u) = \sum_{i=1}^{deg(u)-1} rc_S(T, c_i) + \sum_{i=1}^{deg(u)-1} \omega(u, c_i) \cdot |Z_{(u,c_i)}^1(T)| \cdot |Z_{(u,c_i)}^2(T)|$$

Lemma 4.10 in words: the routing cost $rc_S(T, u)$ of a subtree $T|_u$ rooted in u of tree T is the sum of two sums. First, the sum of the routing cost $rc_S(T, c_i)$ of all subtrees $T|_{c_i}$ rooted in the children c_i of u . And second, the sum of the routing cost off all edges (u, c_i) from u to all its children c_i . The product $|Z_{(u,c_i)}^1(T)| \cdot |Z_{(u,c_i)}^2(T)|$ denotes the number of v - w -paths in T which include the edge (u, c_i) , for $v, w \in S$. Note that we set the routing cost $rc_S(T, u_l)$ to be 0 for a node u_l which is a leaf in T . And $rc_S(T, r) := RC_S(T)$ for the root node r of tree T .

The idea of the S-MRCT-Algorithm in [15] is, in a first step to compute in parallel $|S|$ trees T_r rooted in r for each node $r \in S$. And in a second step to compute the routing cost $RC_S(T_r)$ for each tree T_r in parallel. Therefore the algorithm aggregates the routing costs $rc_S(T_r, u)$ of the subtrees $T_r|_u$ in tree T_r in a bottom-up fashion with the formula stated in Lemma 4.10. Finally the algorithm selects the tree T_r with the minimum routing cost $RC_S(T_r)$ and broadcasts the result in the network.

We use Algorithm 7 MRCT and Algorithm \mathcal{A}'_{MRCT} (whose pseudocode is stated in Algorithm 8) to transfer the S-MRCT-Algorithm from [15] into our framework. The subalgorithm \mathcal{A}'_{MRCT} denotes the algorithm \mathcal{A}'_{agr} of our framework, specified for the S-MRCT approximation. Since we are only interested in Trees and not in DLGs, the Tree Variation involving Algorithms 9 TREECOMPUTING and 10 TREEAGGREGATING is used. Assume each node $u \in V$ knows whether it is in set S and u knows $|S|$. Further the weight of an edge denotes the routing cost for that edge.

Algorithm 7 MRCT starts with estimating and distributing D'_ω in the graph (Line 1 to 5). This is needed to determine the runtime of Algorithms 9 TREECOMPUTING and 10 TREEAGGREGATING . Algorithm 9 TREECOMPUTING executed in Line 6 is only needed to compute the trees T_r for Algorithm 10 TREEAGGREGATING for each node $r \in S$. Therefore we do not need algorithm \mathcal{A}_{MRCT} for the execution. During the aggregating in Line 7, algorithm \mathcal{A}'_{MRCT} executed inside Algorithm 10 TREEAGGREGATING computes the routing costs for the trees T_r by applying the recursive formula stated above. Algorithm \mathcal{A}'_{MRCT} therefore uses the recursive relation of the routing cost stated in Lemma 4.10.

In \mathcal{A}'_{MRCT} Line 2 to 6, each node $u \in V$ is initialized with zero routing cost $rc_S[r] = 0$ (per root node $r \in S$) and $z[r]$ the number of root nodes $s \in S$ in subtree $T_r|_u$ is set to 1 if $u \in S$, and to 0 otherwise. In Line 7, the message $msg_p[r]$ is initialized in case node u is a leaf in tree T_r , and therefore is sent to the parent of u in T_r before \mathcal{A}'_{MRCT} is executed the first time. If u is not a leaf in in tree T_r , then the values of $rc_S[r]$, $z[r]$ and $msg_p[r]$ will change during the execution of Algorithm 10 TREEAGGREGATING .

Back in Algorithm 7 MRCT in Lines 8 to 11, after the execution of Algorithms 9 TREECOMPUTING and 10 TREEAGGREGATING , all nodes $r \in S$ store the routing cost $rc_S[r]$ of their tree T_r in $RC_S(r)$, all other nodes $v \notin S$ store ∞ . Then, in Line 12 to 14 the tree with the smallest routing cost is selected and broadcast.

Proof of Theorem 1.5. Note that our Algorithms compute exactly the same trees as Algorithm 1 in [15], as received messages (which determine the parent) from neighbors are forwarded based on the same rule (prioritizing neighbor IDs) in both algorithms. As a consequence Algorithm 7 MRCT computes the same results as the S-MRCT algorithm in [15] by aggregating along the trees computed by Algorithm 9 TREECOMPUTING. Lines 1 to 5, 13 and 14 take time $\mathcal{O}(D_\omega)$ to estimate, broadcast and aggregate along tree T_1 . Algorithms 9 TREECOMPUTING and 10 TREEAGGREGATING are bound by time $\mathcal{O}(|S| + D_\omega)$, as seen before. Hence, Algorithm 7 MRCT runs in time $\mathcal{O}(|S| + D_\omega)$. Theorem 2 in [15] states that the algorithm that we transfer into our setting computes a $(\times, 2)$ -approximation to an MRCT. Note that [15] assumes that $\omega(e)$ is uniform or corresponds to the time needed to transmit a message through edge e . \square

Proof of Theorem 1.6. Algorithm MRCTRAND is very similar to Algorithm 7 MRCT, only the set of root nodes is different. Like in [15] we sample a set $S' \subset S$ of size $\mathcal{O}(\log n)$ uniformly at random and then execute Algorithm 7 MRCT on set S' . In the end we choose the tree with smallest routing cost as the approximation. Theorem 3 in [15] states this yields a $(\times, 2 + \varepsilon)$ -approximation to an MRCT. From the proof of Theorem 1.5 we know that the runtime of our implementation is $\mathcal{O}(|S'| + D_\omega) = \mathcal{O}(D_\omega + \log n)$. Note that [15] assumes that $\omega(e)$ is uniform or corresponds to the time needed to transmit a message through edge e . \square

Algorithm 7 MRCT Approximation (executed by node u)

```

1: if  $u = 1$  then
2:   estimate  $D'_\omega$  by generating a spanning tree  $T_1$  and using Fact 2.7;
3:   broadcast  $D'_\omega$  through  $T_1$ ;
4: else
5:   receive value of  $D'_\omega$ ;
6: execute Algorithm 9 TREECOMPUTING with  $(|S|, D'_\omega, \perp)$ ;
7: execute Algorithm 10 TREEAGGREGATING with  $(|S|, D'_\omega, \mathcal{A}'_{MRCT}())$ ;
8: if  $u \in S$  then
9:    $RC_S(T_u) = rc_S[u]$ ;
10: else
11:    $RC_S(T_u) = \infty$ ;
12: if  $u = 1$  then
13:   aggregate  $v := \arg \min_{v \in V} RC_S(T_v)$  via  $T_1$ ;
14:   broadcast “ $T_v$  is  $(\times, 2)$ -approximation for S-MRCT”;

```

References

- [1] C. C. Aggarwal and H. Wang. A survey of clustering algorithms for graph data. In C. C. Aggarwal and H. Wang, editors, *Managing and Mining Graph Data*, volume 40 of *Advances in Database Systems*, pages 275–301. Springer, Berlin & Heidelberg, Germany, 2010.
- [2] T. Anantvalee and J. Wu. A survey on intrusion detection in mobile ad hoc networks. In *Wireless Network Security*, pages 159–180. Springer, 2007.
- [3] D. A. Bader, S. Kintali, K. Madduri, and M. Mihail. Approximating betweenness centrality. In A. Bonato and F. Chung, editors, *Algorithms and Models for the Web-Graph*, volume 4863 of *LNCS*, pages 124–137. 2007.

Algorithm 8 $\mathcal{A}'_{MRCCT}()$

```
1: // INITIALIZATION
2: global  $rc_S := \{0, \dots, 0\}$ ;
3: if  $u \in S$  then
4:    $z := \{1, \dots, 1\}$ ;
5: else
6:    $z := \{0, \dots, 0\}$ ;
7: for each  $v \in S$  do  $msg_p[v] = (rc_S[v], z[v])$ ; // output if  $u$  is a leaf

8: // COMPUTATION,  $g(v, msg_c, u_c)$ :  $v$  root node ID of the message,  $msg_c$  message of a child
    $u_c$  in  $T_v$ 
9:  $(r_c, z_c) := msg_c$ ;
10:  $rc_S[v] := rc_S[v] + r_c + \omega(u, u_c) \cdot z_c \cdot (|S| - z_c)$ ;
11:  $z[v] := z[v] + z_c$ ;
12:  $msg_p[v] = (rc_S[v], z[v])$ ; // output
```

- [4] A. Bavelas. Communication patterns in task-oriented groups. *The Journal of the Acoustical Society of America*, 22(6):725–730, 1950.
- [5] M. A. Beauchamp. An improved index of centrality. *Behavioral Science*, 10(2):161–163, 1965.
- [6] S. P. Borgatti. Centrality and network flow. *Social Networks*, 27(1):55–71, 2005.
- [7] U. Brandes. A faster algorithm for betweenness centrality. *The Journal of the Mathematical Sociology*, 25(2):163–177, 2001.
- [8] P. Carrington, J. Scott, and S. Wasserman, editors. *Models and methods in social network analysis*. Cambridge University Press, Cmabridge, UK, 2005.
- [9] E. M. Daly and M. Haahr. Social network analysis for routing in disconnected delay-tolerant manets. In E. Kranakis, E. M. Belding, and E. Modiano, editors, *Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc 2007, Montreal, Quebec, Canada, September 9-14, 2007*, pages 32–40, 2007.
- [10] A. Das Sarma, S. Holzer, L. Kor, A. Korman, D. Nanongkai, G. Pandurangan, D. Peleg, and R. Wattenhofer. Distributed verification and hardness of distributed approximation. *SIAM Journal on Computing (SICOMP)*, 1(2):283–323, 2005.
- [11] B. Dissler. Efficient multi-aggregation with applications to centrality computation. Semester thesis, ETH Zürich, Department of Information Technology and Electrical Engineering, Zürich, Switzerland, 2013.
- [12] D. Eppstein and J. Wang. Fast approximation of centrality. In *Proceedings of the twelfth annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 228–229, 2001.
- [13] J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In L. L. Larmore and M. X. Goemans, editors, *Proceedings of the 35th annual ACM Symposium on Theory of Computing, STOC 2003, San Diego, California, USA, June 9-11, 2003*, pages 448–455, 2003.
- [14] L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, 40(1):35–41, 1977.

- [15] A. Hochuli, S. Holzer, and R. Wattenhofer. Distributed approximation of minimum routing cost trees. In M. M. Halldórsson, editor, *Structural Information and Communication Complexity - 21st International Colloquium, SIROCCO 2014, Takayama, Japan, July 23-25, 2014. Proceedings*, volume 8576 of *Lecture Notes in Computer Science*, pages 121–136. Springer, 2014.
- [16] S. Holzer. *Distance Computation, Information Dissemination, and Wireless Capacity in Networks*, Diss. ETH No. 21444. Phd thesis, ETH Zurich, Zurich, Switzerland, 2013.
- [17] S. Holzer and R. Wattenhofer. Optimal distributed all pairs shortest paths and applications. In D. Kowalski and A. Panconesi, editors, *Proceedings of the 31st annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, PODC 2012, Funchal, Madeira, Portugal, July 16-18, 2012*, pages 355–364, 2012.
- [18] T. C. Hu. Optimum communication spanning trees. *SIAM Journal on Computing*, 3(3):188–195, 1974.
- [19] J. Jiang and M. Worboys. Detecting basic topological changes in sensor networks by local aggregation. In *Proceedings of the 16th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, GIS '08*, pages 4:1–4:10, New York, NY, USA, 2008. ACM.
- [20] M. Khan, F. Kuhn, D. Malkhi, G. Pandurangan, and K. Talwar. Efficient distributed approximation algorithms via probabilistic tree embeddings. In R. A. Bazzi and B. Patt-Shamir, editors, *Proceedings of the 27th annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, PODC 2008, Toronto, Ontario, Canada, August 18-21, 2008*, pages 263–272, 2008.
- [21] F. Kuhn, T. Locher, and R. Wattenhofer. Distributed selection: a missing piece of data aggregation. *Communications of the ACM*, 51(9):93–99, 2008.
- [22] C. Lenzen and D. Peleg. Efficient distributed source detection with limited bandwidth. In P. Fattourou and G. Taubenfeld, editors, *Proceedings of the 32nd annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, PODC 2013, Montreal, Quebec, Canada, July 22-24, 2013*, pages 375–382, 2013.
- [23] T. S. E. Ng and H. Zhang. Predicting internet network distance with coordinates-based approaches. In *Proceedings of the 21st annual Joint Conference of the IEEE Computer and Communications Societies, INFOCOM 2002, New York, New York, USA, June 23-27, 2002*, volume 1, pages 170–179, 2002.
- [24] M. Pias, J. Crowcroft, S. Wilbur, T. Harris, and S. Bhatti. Lighthouses for scalable distributed location. In M. F. Kaashoek and I. Stoica, editors, *Peer-to-Peer Systems II. Second International Workshop, IPTPS 2003 Berkeley, CA, USA, February 21-22, 2003 Revised Papers*, volume 2735 of *Lecture Notes in Computer Science*, pages 278–291. Springer, Berlin & Heidelberg, Germany, 2003.
- [25] R. Rajagopalan and P. Varshney. Data-aggregation techniques in sensor networks: a survey. *Communications Surveys & Tutorials, IEEE*, 8(4):48–63.
- [26] M. J. Rattigan, M. Maier, and D. Jensen. Graph clustering with network structure indices. In Z. Ghahramani, editor, *Machine Learning, Proceedings of the Twenty-Fourth International Conference (ICML 2007), Corvallis, Oregon, USA, June 20-24, 2007*, volume 227 of *ACM International Conference Proceeding Series*, pages 783–790. ACM, New York, New York, USA, 2007.

- [27] B. Shao, H. Wang, and Y. Li. Trinity: A distributed graph engine on a memory cloud. In K. A. Ross, D. Srivastava, and D. Papadias, editors, *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2013, New York, New York, USA, June 22-27, 2013*, pages 505–516, 2013.
- [28] S. Wasserman and K. Faust. *Social network analysis: Methods and applications*, volume 8 of *Structural Analysis in the Social Sciences*. Cambridge University Press, Cambridge, UK, 1994.
- [29] R. T. Wong. Worst-case analysis of network design problem heuristics. *SIAM Journal of Algebraic Discrete Methods*, 1(1):51–63, 1980.
- [30] B. Y. Wu, G. Lancia, V. Bafna, K.-M. Chao, R. Ravi, and C. Y. Tang. A polynomial-time approximation scheme for minimum routing cost spanning trees. *SIAM Journal on Computing*, 29(3):761–778, 1999.

Appendix

4.4 Tree Variation

In schedule τ initialized in Algorithm 1 (Line 3) we only store the arrival time of the first r -message per root node $r \in S$. Therefore the second dimension of the array τ representing the schedule can be omitted in both, Algorithm 1 and 2. The number of parents sp can be omitted too. And since we have to handle only one r -message per root node r , we rewrite Lines 22–34 of Algorithm 1 as stated in Algorithm 9.

In Algorithm 2 DLGAGR just a minor modification is required. The ‘*For each*’-loop in Line 3 needs to iterate over one parent $parent_in_T_r$ per tree T_r only, since in a tree a node has at most one parent.

The fully rewritten Tree Variation of Algorithms 1 DLGCOMP and 2 DLGAGR are stated in Algorithm 9 and 10.

Proof of Theorem 1.2. By modifying Algorithms 1 DLGCOMP and 2 DLGAGR, we only changed the space requirements inside the nodes $u \in V$ and did not change anything concerning the running time. Therefore the time complexity remains $\mathcal{O}(|S| + k)$ after the modifications. To show that the modification computes valid trees, note that a tree T_r is a subgraph of a DLG \mathcal{L}_r . A tree T_r can be computed from a DLG \mathcal{L}_r , by removing all but one edge to the parents at each node $u \in V$, except for the root node r , which has no parent node at all. This is done in Algorithm 9 by considering only the neighbor node u_p as a parent in tree T_r , which sends the first r -message. \square

Algorithm 9 computing $|S|$ trees (executed by node u)

Input: $|S|, D'_\omega, \mathcal{A}_{agr}$

passed parameters on execution of $\mathcal{A}_{agr}(v, msg_p[v])$: $msg_p[v]$: message of the parent in tree T_v , $msg_c[v]$: message that is sent to children in tree T_v after execution

```

1: // On node globally usable variables:
2:  $L := \emptyset, L_{delay} := \emptyset$ ;
3:  $\tau := [\infty, \infty, \dots, \infty]$ ; //  $\tau[v]$  = time when the first message of tree  $T_v$  reaches  $u$ 
4:  $\omega := \{\infty, \infty, \dots, \infty\}$ ; //  $\omega[v] = \omega(v, u)$ ,  $\omega[\emptyset] := \infty$ 
5: if  $u \in S$  then
6:    $L := \{u\}$ ;
7:    $\omega[u] := 0$ ;
8:    $\tau[u] := 0$ ;
9:  $L_1, \dots, L_{deg(u)} := L$ ;
10: initialize algorithm  $\mathcal{A}_{agr}$ ;
11: for  $t = 0, \dots, |S| + D'_\omega$  do
12:   for  $i = 1, \dots, deg(u)$  do
13:      $l_i := \min \{v \in L_i : \tau[v] + \omega(u, v) \geq t \wedge v \notin L_{delay}\}$ ;
14:   within one time slot:
     send  $(l_1, \omega[l_1] + \omega(u, u_1), msg_c[l_1])$  to neighbor  $u_1$ , receive  $(r_1, \omega_1, msg_{p,1})$  from  $u_1$ ;
     send  $(l_2, \omega[l_2] + \omega(u, u_2), msg_c[l_2])$  to neighbor  $u_2$ , receive  $(r_2, \omega_2, msg_{p,2})$  from  $u_2$ ;
     ...
     send  $(l_{deg(u)}, \omega[l_{deg(u)}] + \omega(u, l_{deg(u)}), msg_c[l_{deg(u)}])$  to neighbor  $u_{deg(u)}$ , receive
      $(r_{deg(u)}, \omega_{deg(u)}, msg_{p,deg(u)})$  from  $u_{deg(u)}$ ;
15:    $R = \{r_i : r_i < l_i \text{ and } i \in \{1, \dots, deg(u)\}\} \setminus L$ ;
16:    $s := \begin{cases} \infty & \text{if } L_{delay} = \emptyset \\ \min(L_{delay}) & \text{else} \end{cases}$ ;
17:   if  $s \leq \min(R)$  and  $s < \infty$  then
18:      $L_{delay} := L_{delay} \setminus \{s\}$ ;
19:   for  $i = 1, \dots, deg(u)$  do
20:     if  $r_i < l_i$  then //  $T_{l_i}$ 's message will be delayed due to  $T_{r_i}$ .
21:       if  $r_i \notin L$  then
22:          $\tau[r_i] := t$ ;
23:          $parent\_in\_T_{r_i} := \text{neighbor } i$ ;
24:         execute algorithm  $\mathcal{A}_{agr}(r_i, msg_{p,i})$ ;
25:          $\omega[r_i] = \omega_i$ ;
26:          $L := L \cup \{r_i\}, L_1 := L_1 \cup \{r_i\}, L_2 := L_2 \cup \{r_i\}, \dots, L_{i-1} := L_{i-1} \cup \{r_i\}, L_{i+1} := L_{i+1} \cup$ 
            $\{r_i\}, \dots, L_{deg(u)} := L_{deg(u)} \cup \{r_i\}$ ;
27:         if  $\min(R) < r_i$  or  $s < r_i$  then
28:            $L_{delay} := L_{delay} \cup \{r_i\}$ ;
29:       else
30:          $L_i := L_i \setminus \{r_i\}$ ;
31:     else
32:        $L_i := L_i \setminus \{l_i\}$ ; //  $T_{l_i}$ 's message was successfully sent to neighbor  $i$ .

```

Algorithm 10 Tree aggregation (executed by node u)

Input: $|S|, D'_\omega, \mathcal{A}'_{agr}$

passed parameters on execution of $\mathcal{A}'_{agr}(v, msg_c)$: msg_c : message of a child in tree T_v , $msg_p[v]$: message that is sent to parents in tree T_v after execution

- 1: **initialize** algorithm \mathcal{A}'_{agr} ;
 - 2: **for** $t = 0, \dots, |S| + D'_\omega$ **do**
 - 3: within one time slot:
 For each $v \in L$
 if $t = |S| + D'_\omega - \tau[v]$ **then**
 send $(v, msg_p[v])$ to *parent_in* T_v ;
 receive (v_1, msg_{c,u_1}) from neighbor u_1 ;
 receive (v_2, msg_{c,u_2}) from neighbor u_2 ;
 ...
 receive $(v_{deg(u)}, msg_{c,u_{deg(u)}})$ from neighbor $u_{deg(u)}$;
 - 4: **for** $i = 1, \dots, deg(u)$ **do**
 - 5: **if** $v_i \neq \perp$ **then**
 - 6: **execute** algorithm $\mathcal{A}'_{agr}(v_i, msg_{c,u_i})$;
-