# Simple and Fast Distributed Computation of Betweenness Centrality

Pierluigi Crescenzi[*]
Université de Paris
IRIF, CNRS

Pierre Fraigniaud[†]
Université de Paris
IRIF, CNRS

Ami Paz[‡]
Faculty of Computer Science
Univarsity of Vienna

## Abstract

Betweenness centrality is a graph parameter that has been successfully applied to network analysis. In the context of computer networks, it was considered for various objectives, ranging from routing to service placement. However, as observed by Maccari et al. [INFOCOM 2018], research on betweenness centrality for improving protocols was hampered by the lack of a usable, fully distributed algorithm for computing this parameter. We resolve this issue by designing an efficient algorithm for computing betweenness centrality, which can be implemented by minimal modifications to any distance-vector routing protocol based on Bellman-Ford. The convergence time of our implementation is shown to be proportional to the diameter of the network.

## 1 Introduction

Betweenness centrality [15] is a measure of "importance" attributed to every node of a graph. Roughly, the betweenness centrality of a node $v$ is the sum, taken over all pairs $(s, t)$ of source-target nodes, of the ratio between the number of shortest paths from $s$ to $t$ passing through $v$, and the total number of shortest paths from $s$ to $t$. Thus, a node with high betweenness centrality belongs to relatively many shortest paths, while a node with low betweenness centrality belongs to relatively few shortest paths. Betweenness centrality has been successfully applied to network analysis: In social networks, a node with high betweenness centrality is plausibly an influential node; in computer networks, a node with high betweenness centrality might be efficiently used for storing relevant resources, but may also cause severe damage to the communications in case of failure or malfunction.

Consequently, betweenness centrality and its variants have been used for optimizing the behavior of communication networks and computer networks [30], whether it is for wireless mesh networks design [29], routing [13], link-sensing [36], resource placement [42] and allocation [54], topology control [50], transmission rates optimization [4], or security [35]. For instance, the optimal frequency at which the incident links must be sensed at each node is known to be inversely proportional to the square root of the betweenness centrality of the node [36].

Nevertheless, as observed by Maccari et al. [37], network optimization techniques based on betweenness centrality suffer from two main issues. First, even if every node has access to information about the whole network, which is the case in link-state protocols, the computation of the betweenness centrality may require excessive computational resources; hence, various

---

heuristics and random sampling techniques were proposed in order to reduce the computation time [2, 3, 10, 17, 28, 34, 38–40, 47–49]. Second, there are no known efficient algorithms for computing betweenness centrality in the context of distance-vector protocols. Existing distributed algorithms for computing betweenness centrality or all-pairs shortest paths are either designed for models that are too weak compared to real-world networks supporting distance-vector protocols (e.g., CONGEST model) [23], or dedicated to restricted classes of network topology (e.g., DAGs or trees) [51–53], or they exchange an amount of information between nodes that exceed the capacity of distance-vector protocols [55]. Actually, even the elegant and practical algorithm by Maccari et al. [37] for computing a related measure called load centrality [9, 19] exchanges slightly more information between nodes than one would expect from a distance-vector protocol.

## 1.1 Our Results

We describe a simple and fast distributed algorithm for computing betweenness centrality. Specifically, our algorithm enables every node $v$ to compute its own centrality $\mathsf{bc}_v$.

Our algorithm is simple in the sense that it can be implemented by minimal modifications to distance-vector protocols based on Bellman-Ford. Concretely, Bellman-Ford asks every node $v$ to send to each of its neighbors a pair of values $(t, d)$ for every target-node $t$, where $d$ is the current distance from $v$ to $t$, as perceived by $v$. Our algorithm simply asks every node to send to each neighbor a quadruple of values $(t, d, s, b)$ for every target-node $t$, where $s$ is the current estimation at $v$ of the number of shortest paths from $v$ to $t$, and $b$ is the current contribution of $t$ to the betweenness centrality of $v$.

Our algorithm is fast in the sense that it converges in a number of distance-vector phases proportional to the diameter of the network. Moreover, the amount of computations performed at each node $v$ upon reception of a message from a neighbor $u$ related to a target $t$ is (amortized) constant, i.e., independent of the size of the network.

We have performed an extensive set of simulations confirming both the correctness analysis of our algorithm, and its efficiency. We have considered different scenarios, including weighted and unweighted networks, and various topologies generated by synthetic models (grids, Erdös-Rényi, etc.) or extracted from real-world networks.

The main outcome of this paper is that betweenness centrality can be efficiently computed at every node in a distributed manner, even in the context of distance-vector protocols. As a consequence, there is no obstacle for using betweenness centrality for optimizing the functionality of networks, as far as computing this parameter at run-time by the network itself is concerned. This resolves a question left open in the work of Maccari et al. [37].

## 1.2 Related Work

In addition to the aforementioned contributions to computing or approximating betweenness centrality, significant effort has been made by the distributed computing community for efficiently computing related graph measure, the main one being all-pairs shortest paths (APSP). The vast majority of the results in the distributed setting were however derived under the so-called CONGEST model [43], in which failure-free processing nodes perform in a sequence of synchronously rounds, and the message traversing any link at any round must carry a constant number of words only (i.e., the messages are of $O(\log n)$ bits in $n$-node networks). In contrast, distance-vector protocols can be viewed as exchanging $\Theta(n)$ words along each edge at each phase of communication between neighbors, each word consisting of a pair $(t, d)$ of values. Nevertheless, it is worth mentioning previous contributions [23, 26, 46], computing betweenness centrality in unweighted graphs (directed or not) in $O(n)$ rounds under the CONGEST model. Some of these results also apply to weighted digraphs, but to DAGs only [46].

---

**Algorithm 1** Distributed Bellman-Ford at node $v$

---

1: **function** INIT
2:     **forall** $t \in V$ **do** $D[t] \leftarrow +\infty$                     $\triangleright$ *D is distance vector*
3:     $D[v] \leftarrow 0$                                          $\triangleright$ $\mathsf{dist}(v,v) = 0$

4: **function** SEND
5:     **loop**                          $\triangleright$ *periodic updates are sent to neighbors*
6:         **forall** $t \in V$ **do** send $(t, D[t])$ to every neighbor $u$

7: **function** RECEIVE(message $(t,d)$ from neighbor $u$)
8:     **if** $d + w(\{u,v\}) < D[t]$ **then**
9:         $D[t] \leftarrow d + w(\{u,v\})$

---

Computing all-pairs shortest paths in the CONGEST model was intensively studied. In unweighted graphs, several $O(n)$-round algorithms exist [24, 32, 44], and these were improved by an $O(n/\log n)$-rounds algorithm [27], which is tight [16]. In the weighted case, an $\tilde{O}(n)$-rounds algorithm was recently presented [7], almost matching the $\Omega(n)$ lower bound [11].

Finally, the computation of betweenness centrality was studied in the context of dynamic graphs [5, 6, 22, 31]. These works makes it possible to maintain the betweenness centralities of the nodes in a changing system, without having to recompute everything from scratch when a change occurs (e.g., adding or removing an edge). However, these algorithms are centralized and not distributed. Note that in the distance-vector model, maintaining even just the distances in a changing network (APSP) is far from being trivial (see, e.g., [45, Section 4.2.2]).

## 2 Definitions

Let $G = (V, E)$ be a connected undirected graph with positive edge-weights. The weight of an edge $e \in E$ is denoted by $w(e) > 0$. A path between two distinct nodes $s, t \in V$ is a sequence $v_0, \ldots, v_k$ with $k \geq 0$, $v_0 = s$, $v_k = t$, and $\{v_{i-1}, v_i\} \in E$ for every $i = 1, \ldots, k$. The length of such a path is $\sum_{i=1}^{k} w(\{v_{i-1}, v_i\})$. A path with minimum length between $s$ and $t$ is called a shortest path between $s$ and $t$. The distance $\mathsf{dist}(s,t)$ between two nodes $s$ and $t$ is the length of a shortest path between $s$ and $t$. The weighted diameter of $G = (V, E)$ is defined as $\max_{s,t \in V} \mathsf{dist}(s,t)$. The weighted diameter is however not reflecting the convergence time of routing protocols based on Bellman-Ford. The complexity of the latter is indeed related to the hop-diameter of $G$ defined as follows. For any two nodes $s \neq t$, let $P_1, \ldots, P_\ell$ be the $\ell \geq 1$ shortest paths between $s$ and $t$ in $G$. Let $\mathsf{minhop}(s,t)$ be the minimum, taken over all $i = 1, \ldots, \ell$, of the number of edges of $P_i$. The hop-diameter of $G$ is set as

$$\mathsf{diam}(G) = \max_{s,t \in V} \mathsf{minhop}(s,t).$$

The standard distributed version of Bellman-Ford enables every node $v$ to compute $\mathsf{dist}(v,t)$ for every $t \in V$ — see Algorithm 1. Here and later, we use different notations to distinguish the defined value (e.g. $\mathsf{dist}$), and the value computed by the algorithm (e.g. $D$). This algorithm converges in $\mathsf{diam}(G)$ phases, where a phase is defined as the time required by every node $v$ to send all the messages $(t, D[t])$, $t \in V$, to all its neighbors, receive all the messages $(t,d)$, $t \in V$, sent by each of its neighbors, and process all these messages. Indeed, a simple induction on $h \geq 0$ enables to show that, for every $t \in V$, every node $v$ with $\mathsf{minhop}(v,t) \leq h$ has computed $\mathsf{dist}(v,t)$ correctly after $h$ rounds.

For any two vertices $s \in V$ and $t \in V$, let $\sigma_{s,t}$ denote the number of shortest paths from $s$ to $t$ in $G$ (with $\sigma_{s,s} = 1$), and let $\sigma_{s,t}(v)$ denote the number of shortest paths from $s$ to $t$ passing through node $v$ (with $\sigma_{s,s}(s) = 1$).

**Definition 1** *The* betweenness centrality *[15] of node $v$ is*

$$\mathsf{bc}_v = \frac{1}{(n-1)(n-2)} \sum_{s \neq v, t \neq v} \frac{\sigma_{s,t}(v)}{\sigma_{s,t}}.$$

Note that, since $G$ is connected, $\sigma_{s,t} > 0$ for every two vertices $s, t \in V$, and thus $\mathsf{bc}_v$ is well defined for every $v \in V$.

As mentioned earlier in the text, betweenness centrality is not only a central notion in the context of network analysis, but can also be used for various optimization scenarios in the context of computer networks. In particular, Maccari and Cigno [36] focused on optimizing the frequency of HELLO messages for link-sensing in wireless networks, and showed that the optimal frequency $f(v)$ at which every node $v$ must sense its neighbors is

$$f(v) \approx \sqrt{\frac{\mathsf{deg}_v}{\mathsf{bc}_v}}, \tag{1}$$

where $\mathsf{deg}_v$ denotes the degree of $v$, i.e., its number of neighbors. Unfortunately, the above formula can hardly be practically used in absence of an efficient way of computing $\mathsf{bc}_v$ at every node $v$. To overcome this, Maccari et al. [37] have proposed to replace the use of betweenness centrality by the use of load centrality [9], defined as follow.

Assume that one unit of flow is pushed from $s$ to $t$ in $G$ according to the following rule. For any node $v$, let $\mathsf{out}_{s,t}(v)$ denote the set of edges $e = \{v, v'\}$ incident to $v$ such that there exists a shortest path from $s$ to $t$ traversing the edge $e$ from $v$ to $v'$. If a fraction $r$ of the flow from $s$ to $t$ is received by a node $v \neq t$, and if $|\mathsf{out}_{s,t}(v)| = k > 0$, then $v$ forwards a fraction $r/k$ of flow along each of the edges in $\mathsf{out}_{s,t}(v)$. Let $\theta_{s,t}(v)$ denote the amount of flow from $s$ to $t$ traversing $v$.

**Definition 2** *The* load centrality *[9] of node $v$ is*

$$\mathsf{lc}_v = \sum_{s,t} \theta_{s,t}(v).$$

A distributed algorithm for computing load centrality has been described and analyzed in [37]. It was shown to be implementable by minimal modifications to the Bellman-Ford algorithm. Roughly, instead of every node $v$ sending just a pair $(t, D[t])$ for every node $t$, where $D[t]$ is the current estimation of the distance between $v$ and $t$ as perceived by $v$, every node $v$ sends a tuple

$$(t, D[t], \mathsf{NH}[t], \ell, L[t])$$

where $\mathsf{NH}[t]$ is the list of "next hops" to $t$, i.e., the set of neighbors of $v$ on a shortest path from $v$ to $t$ (corresponding to $\mathsf{out}_{v,t}(v)$), $\ell$ is the overall flow passing through $v$ to reach $t$, and $L[t]$ is the load centrality of $t$, as far as $v$ knows. Note that this tuple can be significantly larger than a pair $(t, D[t])$ as $\mathsf{NH}[t]$ can potentially contain many entries. Instead, our algorithm for computing the betweenness centrality exchanges messages with a bounded number of entries. Note also that the algorithm in [37] converges in a number of phases related to

$$\mathsf{Diam}(G) = \max_{s,t \in V} \mathsf{maxhop}(s,t),$$

where $\mathsf{maxhop}(s,t)$ is the maximum, taken over all shortest paths between $s$ and $t$, of the number of edges of each path. (Recall that the hop-diameter of $G$ is defined as $\mathsf{diam}(G) = \max_{s,t \in V} \mathsf{minhop}(s,t)$.) Running for $\mathsf{Diam}(G)$ rounds seems unavoidable, by the definition of the load centrality, because a positive fraction of the flow from $s$ to $t$ is indeed shipped via a shortest path of length $\mathsf{maxhop}(s,t)$ between $s$ and $t$. Nevertheless [37] the practical performance of the algorithm computing load centrality remains close to $\mathsf{diam}(G)$ — this is mainly due to the fact that there are typically few shortest paths between any two nodes in real-world weighted networks, all with very similar numbers of edges.

In the next section, we show that betweenness centrality can also be computed distributively, by minimal modifications of Bellman-Ford. In particular, our algorithm enables to compute exactly the optimal frequency $f(v)$ of each node $v$, as described in Eq. (1). The performance of our algorithm is also shown to be close to the diameter of the network.

# 3 Distributed Computation of BC

This section describes our distributed algorithm for computing betweenness centrality. This algorithm is in essence a distributed implementation of the dynamic programming algorithm by Brandes [8]. Note that this latter algorithm is well suited for a distributed implementation because a node $v$ does not need to know the number of shortest paths from $s$ to $t$ in order to compute the contribution of the pair $s$ and $t$ to its betweenness centrality (see Lemma 1).

For the ease of presentation, we view the undirected graph $G = (V, E)$ as a directed graph where every edge $\{u, v\}$ is replaced by two symmetric arcs $(u, v)$ and $(v, u)$, both with the same weight as the edge $\{u, v\}$. Also, we extend the definition of $\sigma_{s,t}(v)$ from nodes to arcs, by denoting, for every arc $(u, v) \in E$, $\sigma_{s,t}(u, v)$ as the number of shortest paths from $s$ to $t$ traversing the arc $(u, v)$, i.e., traversing the edge $\{u, v\}$ from $u$ to $v$. The following two facts directly follow from the definitions.

**Fact 1** *If $\sigma_{s,t}(v) \neq 0$, i.e., if $v$ belongs to a shortest path from $s$ to $t$, then $\sigma_{s,t}(v) = \sigma_{s,v} \cdot \sigma_{v,t}$. Similarly, if the arc $(u, v)$ belongs to a shortest path from $s$ to $t$, then $\sigma_{s,t}(u, v) = \sigma_{s,u} \cdot \sigma_{v,t}$.*

Let $v \in V$. For every $t \in V$, let $\mathsf{NH}_v(t)$ be set of neighbors of $v$ that belong to some shortest paths from $v$ to $t$, and, for every $s \in V$, let $\mathsf{PH}_v(s)$ be the set of neighbors $u$ of $v$ such that $v$ belongs to some shortest paths from $s$ to $u$. ($\mathsf{NH}$ and $\mathsf{PH}$ holds for "next hop" and "previous hop", respectively). Note that, since $G$ is undirected, for every $u, v, w \in V$, $u \in \mathsf{PH}_v(w)$ if and only if $v \in \mathsf{NH}_u(w)$.

**Fact 2** *For every $t \neq v$, we have $\sigma_{v,t} = \sum_{u \in \mathsf{NH}_v(t)} \sigma_{u,t}$. If $v$ is on a shortest path from $s$ to $t$, then $\sigma_{v,t} = \sum_{u \in \mathsf{PH}_v(s)} \sigma_{v,t}(v, u)$.*

Finally, for every $s \in V$, let $\mathsf{bc}_v(s)$ be the contribution of the source $s$ to $\mathsf{bc}_v$, that is,

$$\mathsf{bc}_v(s) = \sum_{t \neq v} \frac{\sigma_{s,t}(v)}{\sigma_{s,t}}.$$

By definition, we have

$$\mathsf{bc}_v = \frac{1}{(n-1)(n-2)} \sum_{s \neq v} \mathsf{bc}_v(s)$$

for every node $v$. The following result is central in our distributed implementation of Brandes algorithm.

**Algorithm 2** Computing betweenness centrality at node $v$

1: **function** INIT
2:     **for all** $t \in V$ **do**
3:         $D[t] \leftarrow +\infty$                                   ▷ *D is a distance vector*
4:         $\text{NH}[t] \leftarrow \emptyset$                                 ▷ *next-hop vector*
5:         $\text{PH}[t] \leftarrow \emptyset$                                ▷ *previous-hop vector*
6:         **for all** $u \in N[v]$ **do**
7:             $B[u,t] \leftarrow 0$                           ▷ *eventually* $\text{bc}_u(t)$
8:             $S[u,t] \leftarrow 0$                           ▷ *eventually* $\sigma_{u,t}$
9:     $S[v,v] \leftarrow 1$                                      ▷ $\sigma_{v,v} = 1$
10:    $D[v] \leftarrow 0$                                     ▷ $\text{dist}(v,v) = 0$

11: **function** SEND
12:     **loop**                             ▷ *periodic updates are sent to neighbors*
13:         **for all** $t \in V$ **do**
14:             send $(t, D[t], S[v,t], B[v,t])$ to all $u \in N(v)$

15: **function** RECEIVE(message $(t,d,s,b)$ from $u \in N(v)$)
16:     $\text{NH}[t] \leftarrow \text{NH}[t] \smallsetminus \{u\}$                    ▷ *initialization: u is removed*
17:     $\text{PH}[t] \leftarrow \text{PH}[t] \smallsetminus \{u\}$                ▷      *from both* NH *and* PH
18:     **if** $d + w(\{u,v\}) < D[t]$ **then**
19:         $D[t] \leftarrow d + w(\{u,v\})$                   ▷ *Bellman-Ford update*
20:     **else if** $d + w(\{u,v\}) = D[t]$ **then**
21:         $\text{NH}[t] \leftarrow \text{NH}[t] \cup \{u\}$               ▷ *u is placed (back) in* NH
22:     **else if** $d - w(\{u,v\}) = D[t]$ **then**
23:         $\text{PH}[t] \leftarrow \text{PH}[t] \cup \{u\}$               ▷ *u is placed (back) in* PH
24:     $S[u,t] \leftarrow s$
25:     $B[u,t] \leftarrow b$
26:     **if** $t \neq v$ **then** $S[v,t] \leftarrow \sum_{x \in \text{NH}[t]} S[x,t]$
27:     $B[v,t] \leftarrow S[v,t] \cdot \sum_{x \in \text{PH}[t]} \frac{B[x,t]+1}{S[x,t]}$
28:     $C \leftarrow \sum_{x \neq v} B[v,x]$                            ▷ *eventually* $C = \text{bc}_v$

**Lemma 1 ([8, Theorem 6])** *For every $s \neq v$, we have*

$$\text{bc}_v(s) = \sigma_{s,v} \sum_{u \in \text{PH}_v(s)} \frac{\text{bc}_u(s) + 1}{\sigma_{s,u}}.$$

    Our algorithm requires that each node $v \in V$ computes and maintains the sets $\text{NH}_v(t)$ and $\text{PH}_v(t)$ for every node $t \in V$. This is achieved by applying simple modifications to the function RECEIVE in Bellman-Ford algorithm, as described in Algorithm 2, Lines 16-23. In this latter algorithm, and in the remaining of the text, we denote by $N(v)$ the set of neighbors of $v \in V$ in $G = (V, E)$, that is,

$$N(v) = \{u \in V : \{u,v\} \in E\},$$

and we denote by $N[v] = N(v) \cup \{v\}$ the closed neighborhood of node $v$.

**Lemma 2** *Algorithm 2 enables every node $v \in V$ to compute the sets $\text{NH}_v(t)$ and $\text{PH}_v(t)$ for every node $t \in V$. More specifically, at node $v$, for every node $t \in V$, $\text{NH}[t] = \text{NH}_v(t)$ after $\text{maxhop}(v, t) + 1$ phases, and $\text{PH}[t] = \text{PH}_v(t)$ after $\text{Diam}(G) + 2$ phases.*

*Proof.* Let $v \in V$. The proof is by induction on $h = \mathsf{maxhop}(v, t)$. For $h = 0$, the fact that $D[v]$ remains zero throughout the execution of the algorithm guarantees that $\mathrm{NH}[v]$ remains empty throughout the execution, as desired. Let $h > 0$, and let us assume that, for every $t \in V$ with $\mathsf{maxhop}(v, t) < h$, $\mathrm{NH}[t] = \mathsf{NH}_v(t)$ after $h$ phases. Let $t \in V$ with $\mathsf{maxhop}(v, t) = h$. Since $\mathsf{minhop}(x, y) \leq \mathsf{maxhop}(x, y)$ for every two nodes $x, y \in V$, we have $D[t] = \mathsf{dist}(v, t)$ after $h$ phases. Let $u \in \mathsf{NH}_v(t)$, and let us consider the updates occurring at phase $h + 1$. Since $\mathsf{maxhop}(u, t) < \mathsf{maxhop}(v, t)$, the value $d$ in the tuple $(t, d, s, b)$ sent by $u$ to $v$ at phase $h + 1$ satisfies $d = \mathsf{dist}(u, t)$. Therefore, the equality $d + w(\{u, v\}) = D[t]$ holds at $v$, resulting to $u$ being added to $\mathrm{NH}[t]$, as desired. That is, for every $u \in \mathsf{NH}_v(t)$, Algorithm 2 guarantees that $u \in \mathrm{NH}[t]$ at $v$, after $h + 1$ phases. This completes the proof of the induction for the next hops vector.

The proof is similar for the previous hops vector. For the base case $h = 0$, observe that

$$\mathsf{PH}_v(v) = \{u \in N(v) : w(\{u, v\}) = \mathsf{dist}(u, v)\}.$$

It follows that, for every $u \in \mathsf{PH}_v(v)$, the value $d$ in the tuple $(t, d, s, b)$ sent by $u$ to $v$ at phase 2 satisfies $d = \mathsf{dist}(u, v)$. Therefore, the equality $d - w(\{u, v\}) = 0 = D[v]$ holds at $v$, resulting in $u$ being added to $\mathrm{PH}[t]$, as desired. Now, let $h > 0$, and let us assume that, for every $t \in V$ with $\mathsf{maxhop}(v, t) < h$, $\mathrm{PH}[t] = \mathsf{PH}_v(t)$ after $h + 1$ phases. Let $t \in V$ with $\mathsf{maxhop}(v, t) = h$, and let $u \in \mathsf{PH}_v(t)$. Since $u \in \mathsf{PH}_v(t)$, $\mathsf{maxhop}(u, t) \leq h + 1$, for which it follows that the distance to $t$ is correctly set at $u$ after $h + 1$ phases. Therefore, at phase $h + 2$, the value $d$ in the tuple $(t, d, s, b)$ sent by $u$ to $v$ satisfies $d = \mathsf{dist}(u, t)$, and thus the equality $d - w(\{u, v\}) = D[t]$ holds at $v$, resulting in $u$ being added to $\mathrm{PH}[t]$. This completes the proof of the induction for the previous hops vector.

The proof completes by noticing that, as $D[t] = \mathsf{dist}(v, t)$ remains stable after $\mathsf{minhop}(v, t)$ phases, a node $u$ added to $\mathrm{NH}[t]$ or to $\mathrm{PH}[t]$ after the due number of phases stays in this set for the remaining phases of the algorithm. $\qquad\square$

We now claim that, in Algorithm 2, every node $v \in V$ computes $\sigma_{v,t}$ and $\mathsf{bc}_v(t)$ for every node $t \in V$, at Lines 24-27. We treat $\sigma_{v,t}$ and $\mathsf{bc}_v(t)$ separately, as the former is somehow computed top-down, while the latter is computed bottom-up.

**Lemma 3** *Algorithm 2 enables every node $v \in V$ to compute $\sigma_{v,t}$ for every node $t \in V$. More specifically, at node $v$, for every node $t \in V$, we have $S[v, t] = \sigma_{v,t}$ after $\mathsf{Diam}(G) + 1$ phases.*

*Proof.* Let $v \in V$. The proof is by induction on $h = \mathsf{maxhop}(v, t)$. The statement holds for $h = 0$, i.e., for $t = v$, as $S[v, v] = 1$ in the function INIT, and $S[v, v]$ is not modified by the function RECEIVE (cf. the test performed at Line 26). Now, let $h > 0$, and assume that for every $t$ such that $\mathsf{maxhop}(v, t) < h$, $S[v, t] = \sigma_{v,t}$ after $h$ phases. Consider phase $h + 1$. By Lemma 2, we have $\mathrm{NH}[t] = \mathsf{NH}_v(t)$ at Line 26. Moreover, by induction, since $\mathsf{maxhop}(x, t) < \mathsf{maxhop}(v, t)$ for every $x \in \mathsf{NH}_v(t)$, we have $S[x, t] = \sigma_{x,t}$ for every such node $x$ after at most $h$ phases. During phase $h + 1$, node $v$ receives all the values $s = S[x, t]$ from these nodes, and thus, once it has received all of them, the update of Line 26 yields $S[v, t] = \sigma_{v,t}$, by Fact 2. $\qquad\square$

**Lemma 4** *Algorithm 2 enables every node $v \in V$ to compute $\mathsf{bc}_v(s)$ for every node $s \in V \setminus \{v\}$. More specifically, at node $v$, for every node $s \in V \setminus \{v\}$, $B[v, s] = \mathsf{bc}_v(s)$ after $2\,\mathsf{Diam}(G) + 1$ phases.*

*Proof.* Let $v \in V$. The proof is by induction on $h = \mathsf{Diam}(G) - \mathsf{maxhop}(v, s)$, that is, we show that, for every $s$ such that $\mathsf{Diam}(G) - \mathsf{maxhop}(v, s) \leq h$, $B[v, s] = \mathsf{bc}_v(s)$ after $\mathsf{Diam}(G) + h + 2$ phases.
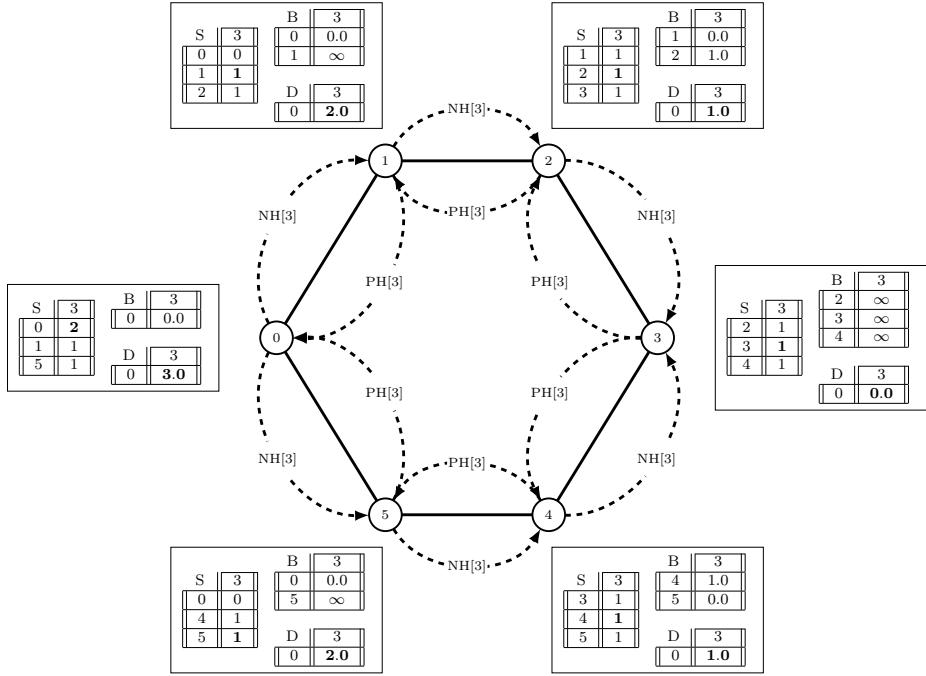
Figure 1: The state of Algorithm 2 executed on a cycle of 6 nodes, at the end of the fourth send/receive phase of the algorithm.

For the base case $h = 0$, let $s \in V$ such that $\mathsf{maxhop}(v, s) = \mathsf{Diam}(G)$. (If there are no such $s$, then the base case holds trivially for $v$). By Lemma 2, after $\mathsf{Diam}(G) + 2$ phases, we have $\mathrm{PH}[s] = \mathrm{PH}_v(s)$ at Line 27. Therefore, $\mathrm{PH}[s] = \emptyset$, because $\mathrm{PH}_v(s) = \emptyset$ as $v$ is a node at maximum hop-distance from $s$. As a consequence, $B[v, s]$ is correctly set to $0 = \mathsf{bc}_v(s)$ at Line 27.

Now, let $h \geq 0$, and let us assume that, for every $s$ such that $\mathsf{Diam}(G) - \mathsf{maxhop}(v, s) \leq h$, $B[v, s] = \mathsf{bc}_v(s)$ after $\mathsf{Diam}(G) + h + 2$ phases. Let $s$ such that $\mathsf{Diam}(G) - \mathsf{maxhop}(v, s) = h + 1$. (Again, if there are no such $s$, then the induction step $h \to h + 1$ holds trivially for $v$). By Lemma 2, after $\mathsf{Diam}(G) + 2$ phases, we have $\mathrm{PH}[s] = \mathrm{PH}_v(s)$ at Line 27. Moreover, by Lemma 3, we have $S[v, s] = \sigma_{v,s}$ and, for every $x \in \mathrm{PH}_v(s) = \mathrm{PH}[s]$, the equality $S[x, s] = \sigma_{x,s}$ holds. Furthermore, for every $x \in \mathrm{PH}_v(s)$, we have $\mathsf{maxhop}(x, s) > \mathsf{maxhop}(v, s)$, from which it follows by induction that $B[x, s] = \mathsf{bc}_x(s)$ after $\mathsf{Diam}(G) + h + 2$ phases. It then follows from Lemma 1 that the computation performed at Line 27 guarantees that $B[v, s] = \mathsf{bc}_v(s)$, which completes the induction step. As we only need to compute $\mathsf{bc}_v(s)$ for $s \neq v$, we need only to consider $h < \mathsf{Diam}(G)$, and the lemma follows. $\qquad \square$

Since $\mathsf{bc}_v$ does not depend on $\mathsf{bc}_v(v)$ but only on $\mathsf{bc}_v(s)$ for $s \in V \setminus \{v\}$, Lemma 4 immediately implies that Algorithm 2 enables every node $v \in V$ to compute $\mathsf{bc}_v$ at Line 28, after $2\,\mathsf{Diam}(G) + 1$ phases, up to renormalization by $1/((n - 1)(n - 2))$. The following theorem summarizes the results in this section.

**Theorem 1** *Algorithm 2 enables every node to compute its betweenness centrality in any network $G$ after $2\,\mathsf{Diam}(G) + 1$ phases.*

## 3.1 An example of execution of Algorithm 2

We now describe few steps of the execution of Algorithm 2 on an unweighted cycle of six nodes, whose diameter is thus $\mathsf{Diam} = 3$. In particular, we show how the algorithm computes the
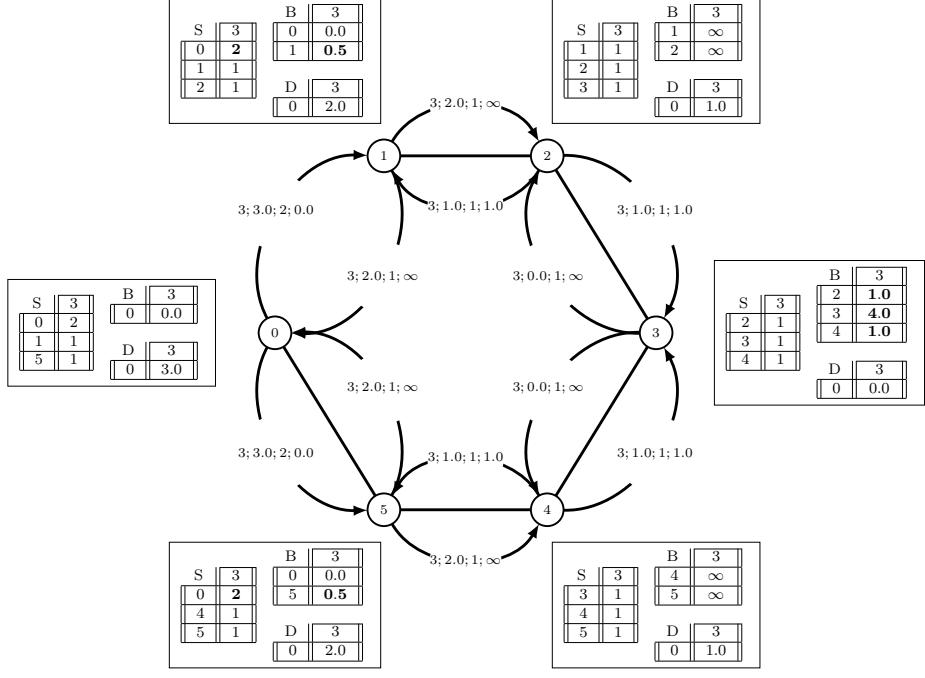
Figure 2: The state of Algorithm 2 executed on a cycle of 6 nodes, at the end of the fifth send/receive phase of the algorithm.

contribution of node $t = 3$ to the betweenness centrality of all nodes, after the first 4 phases have been completed (that is, after the correct value of the sets NH and PH and of the matrix S has been computed). In Figure 1 we show the state of all nodes at this moment of the algorithm execution. For example, node 5 now knows that its set of next hops towards node 3 contains only node 4, while its set of previous hops towards node 3 contains only node 0 (which is the only neighbor of node 5 which admits a shortest path towards node 3 passing through node 5). Note that all nodes have also correctly computed their distance from node 3 (in the case of node 5, for example, this distance is 2.0), and the number of shortest paths connecting them to node 3 (in the case of node 5, for example, this number is 1). In the figure, all these correct values are shown in boldface.

In Figure 2, we show the messages sent by each node to its neighbors concerning node $t = 3$ in the fifth send/receive phase (these are the messages sent by the function SEND defined in lines 11–14 of Algorithm 2). For instance, node 5 sends to both node 0 and node 4 the message $(3, 2.0, 1, \infty)$, telling them that its distance from node 3 is 2.0, that there is only one shortest path from it to node 3, and that the current contribution of node 3 to its betweenness centrality is still unknown. On the other hand, suppose that node 5 receives first the message $(3, 3.0, 2, 0.0)$ from node 0 and then the message $(3, 1.0, 1, 1.0)$ from node 4. Even if node 0 is taken out of PH[3] by node 5 at line 16 of Algorithm 2, this node is reinserted in PH[3] by node 5 at line 23 of Algorithm 2, since the distance of node 5 from node 3 (that is, the value D[3]) is equal to $d = 3.0$ minus 1 (remember that the graph is unweighted). Hence, the set PH[3] of node 5 does not change. After receiving the message $(3, 3.0, 2, 0.0)$ from node 0, node 5 updates (at lines 24–25 of Algorithm 2) S[0, 3] (which becomes 2) and B[0, 3] (which remains 0.0). Since the set NH[3] contains only node 4, the value S[5, 3] is then set equal to S[4, 3] = 1 (line 26 of Algorithm 2). At line 27 of Algorithm 2, the value B[5, 3] is set equal to $1 \cdot \frac{0+1}{2} = 0.5$. Note that, at this moment, node 5 has correctly computed the number of shortest paths connecting node 0 to node 3, and the contribution of node 3 to its betweenness centrality (these values are
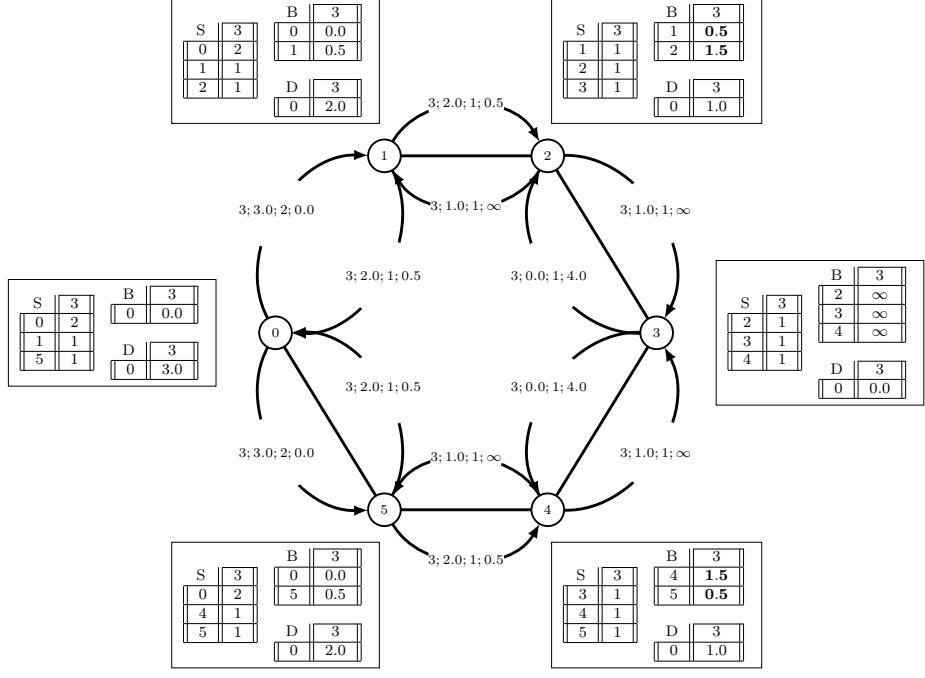
Figure 3: The state of Algorithm 2 executed on a cycle of 6 nodes, at the end of the sixth send/receive phase of the algorithm.

shown in bold in the figure). It is easy to verify that the arrival of the message from node 4 does not change the state of node 5. Similarly, node 1 (which is symmetric to node 5) and node 3 update their state (once again, the updates are shown in bold in the figure).

The execution of the next two phases of Algorithm 2 are shown in Figure 3 and 4, respectively. In particular, at the end of the seventh phase, each node $v$ has correctly computed, for each of its neighbors $u$, the number of shortest path connecting $u$ to node 3, and the contribution of node 3 to the betweenness centrality of $v$.

# 4 Implementation of Algorithm 2

In this section, we present a more practical implementation of Algorithm 2, in which every node performs only a constant number of elementary operations upon reception of every message. The instructions performed at Lines 26 and 27 of Algorithm 2 may, indeed, consume a large time, if $\mathsf{NH}_v(t)$ or $\mathsf{PH}_v(t)$ are large (note that the operations on the array NH and PH can be implemented in (amortized) constant time using appropriate data structures).

For this reason, we modify the function RECEIVE in order to accumulate values for the computation of $\sigma_{v,t}$ and $\mathsf{bc}_v(t)$ instead of summing up a potentially large number of values. This is done by the modified RECEIVE function described in Algorithm 3. In addition, the variable $C$ must be initialized to 0 in INIT, while the auxiliary array $A$, which stores the contribution of $u \in N(v)$ to $\mathsf{bc}_v(t)$, does not need to be initialized.
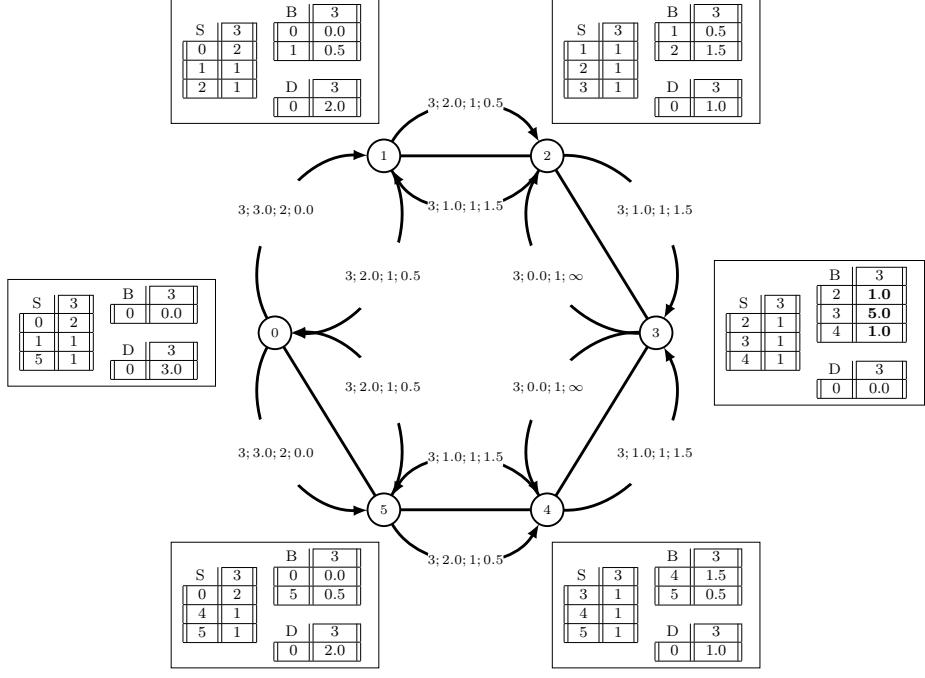
Figure 4: The state of Algorithm 2 executed on a cycle of 6 nodes, at the end of the seventh and last send/receive phase of the algorithm.

Algorithm 3 aims at replacing the sums in the instructions

$$S[v,t] \leftarrow \sum_{x \in \text{NH}[t]} S[x,t]$$

$$B[v,t] \leftarrow S[v,t] \cdot \sum_{x \in \text{PH}[t]} \frac{B[x,t]+1}{S[x,t]}$$

$$C \leftarrow \sum_{x \neq v} B[v,x]$$

in Algorithm 2 by a bounded number of operations. The instruction $S[v,t] \leftarrow \sum_{x \in \text{NH}[t]} S[x,t]$ is replaced by removing $S[u,t]$ from $S[v,t]$ whenever $u$ must be removed from $\text{NH}[t]$ (cf. Line 5), and adding $S[u,t]$ to $S[v,t]$ whenever $u$ must be added to $\text{NH}[t]$ (cf. Line 15). Similarly, the instruction $B[v,t] \leftarrow S[v,t] \cdot \sum_{x \in \text{PH}[t]} \frac{B[x,t]+1}{S[x,t]}$ is replaced by removing from $B[v,t]$ the previously added contribution of $u$ to $B[v,t]$ stored in $A[u,t]$ (cf. Line 8), whenever $u$ must be removed from $\text{PH}[t]$, and adding the contribution $S[v,t] \cdot \frac{B[u,t]+1}{S[u,t]}$ to $B[v,t]$ whenever $u$ must be added to $\text{PH}[t]$ (cf. Line 22), after having saved it into $A[u,t]$. Finally, the instruction $C \leftarrow \sum_{x \neq v} B[v,x]$ is replaced by exchanging the old value of $B[v,t]$ with the (potentially) new value computed at Line 22 — cf. Lines 2 and 23.

## 5 Experimental Results

We have implemented a Java simulator in order to analyze the global and the local convergence times, when applying Algorithm 3 to lattice networks, to networks generated by different random graph models, and to real-world networks. The simulator uses a virtual clock and, for each node, triggers a send-event once per virtual time unit (also called *phase*). Unlike [37], no random jitter

11

**Algorithm 3** Implementation of Function RECEIVE of Algorithm 2, with constant number of elementary operations per message: instructions performed at node $v$

---

1: **function** RECEIVE(message $(t, d, s, b)$ from $u \in N(v)$)
2:     **if** $t \neq v$ **then** $C \leftarrow C - B[v, t]$
3:     **if** $u \in \mathrm{NH}[t]$ **then**
4:         $\mathrm{NH}[t] \leftarrow \mathrm{NH}[t] \smallsetminus \{u\}$
5:         **if** $t \neq v$ **then** $S[v, t] \leftarrow S[v, t] - S[u, t]$
6:     **if** $u \in \mathrm{PH}[t]$ **then**
7:         $\mathrm{PH}[t] \leftarrow \mathrm{PH}[t] \smallsetminus \{u\}$
8:         $B[v, t] \leftarrow B[v, t] - A[u, t]$
9:     $S[u, t] \leftarrow s$
10:    $B[u, t] \leftarrow b$
11:    **if** $d + w(\{u, v\}) < D[t]$ **then**
12:       $D[t] \leftarrow d + w(\{u, v\})$
13:    **else if** $d + w(\{u, v\}) = D[t]$ **then**
14:       $\mathrm{NH}[t] \leftarrow \mathrm{NH}[t] \cup \{u\}$
15:       **if** $t \neq v$ **then** $S[v, t] \leftarrow S[v, t] + S[u, t]$
16:    **else if** $d - w(\{u, v\}) = D[t]$ **then**
17:       $\mathrm{PH}[t] \leftarrow \mathrm{PH}[t] \cup \{u\}$
18:       **if** $S[u, t] \neq 0$ **then**
19:          $A[u, t] \leftarrow S[v, t] \cdot \frac{B[u,t]+1}{S[u,t]}$
20:       **else**
21:          $A[u, t] \leftarrow 0$
22:       $B[v, t] \leftarrow B[v, t] + A[u, t]$
23:    **if** $t \neq v$ **then** $C \leftarrow C + B[v, t]$

---

has been added to the events scheduled by nodes, so the simulator analyzes the algorithm in the case of perfect synchronization. Each simulation ends when all nodes converge to steady state, that is, when the $C$ values do not change anymore. In order to perform the global convergence study, during each simulation we also record the phases at which the $D$ and the $S$ values do no change anymore. Finally, for the local convergence study we also record, for each node $v$, the last phase in which the $C$ value of $v$ changed. The output of our algorithms is compared with the output of the Python NetworkX library [21].

## 5.1 The datasets

We have performed our convergence study on the following networks (in the rest of this section, we will present the results relative only to some of these networks, since the results on the others are very similar).

- Hypercubes of several different dimensions, grids with several different widths and heights, and complete binary trees of different heights.

- Random Erdös-Rényi graphs with different diameters [14], random Barabási-Albert graphs with different numbers of links for each new node added to the graph [1], and random geometric graphs with different communication ranges [18].

- Several real-world networks, such as an e-mail graph, whose edges indicate e-mail interchanges between members of the Univeristy Rovira i Virgili (Tarragona, Spain) [20];
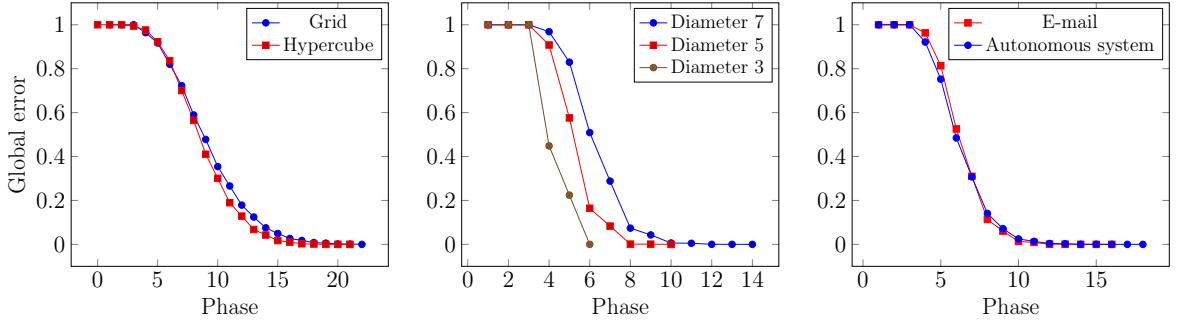
Figure 5: The (mean) global error as a function of time in (left) a $7 \times 6$ grid and a hypercube of dimension 11 (where, hence, diameter is 11), (center) Erdös-Renyi graphs with 500 nodes and different diameters (20 samples for each diameter), and (right) an e-mail network with 1133 nodes and in an autonomous system network with 3011 nodes.

several autonomous system networks, which are communication networks of "who-talks-to-whom" obtained from the BGP logs [33]; a road graph, which contains a large portion of the road network of the city of Rome, Italy, from 1999 (vertices correspond to intersections between roads, edges correspond to roads or road segments, and weights correspond to distances) [12]; and a co-authorship (both unweighted and weighted) graph [41].

## 5.2 Analysis of global convergence time

We start by studying the global convergence over time. To this end, at each phase $P$, we study the $\ell_2$-norm of the distance between the current betweenness centrality computed by the nodes, $C = C(P)$, and the actual betweenness centrality, $\mathsf{bc}$. This is normalized by the actual betweenness centrality, which gives the following global-error formula:

$$\frac{\|\mathsf{bc} - C\|_2}{\|\mathsf{bc}\|_2} = \frac{\sqrt{\sum_{v \in V}(\mathsf{bc}_v - C[v])^2}}{\sqrt{\sum_{v \in V}(\mathsf{bc}_v)^2}}$$

### 5.2.1 Unweighted networks: lattices, Erdös-Renyi, e-mail

Our first experiments are with unweighted, synthetic networks. We start with a grid and a hypercube (Figure 5 (left)), both with diameter 11. Grids and hypercubes with different dimensions present very similar behaviours, and so do binary trees. In the first 3 phases, the betweenness centrality of all nodes is 0, since the list PH is empty. This phenomena is normal and predictable, and reoccurs in all our experiments, both in unweighted and weighted graphs.

As predicted by our analysis, the algorithm converges on both networks in roughly $2\,\mathsf{Diam}(G)$ time. In an unweighted network, the estimate of $\mathsf{bc}$ made by each node can only increase: as time passes, each node learns about more shortest paths it belong to, increases its estimate of the betweenness centrality, and thus makes it more accurate, until converging to the right values. In these examples, the values computed by the nodes after $\mathsf{Diam} + 2$ phases already give a relatively low error (less than 10%).

Next, we study the convergence on Erdös-Renyi graphs. We average over 20 randomly-generated graphs with a given diameter, for diameters $3, 5$ and $7$ (Figure 5 (center)). While slightly more cluttered, the results are similar to the ones observed in the lattices above. The betweenness centrality is not updated in the first 3 phases, and then rapidly and monotonically decreases, until full convergence in $2\,\mathsf{Diam}$ phases. Other random graphs present similar behavior.
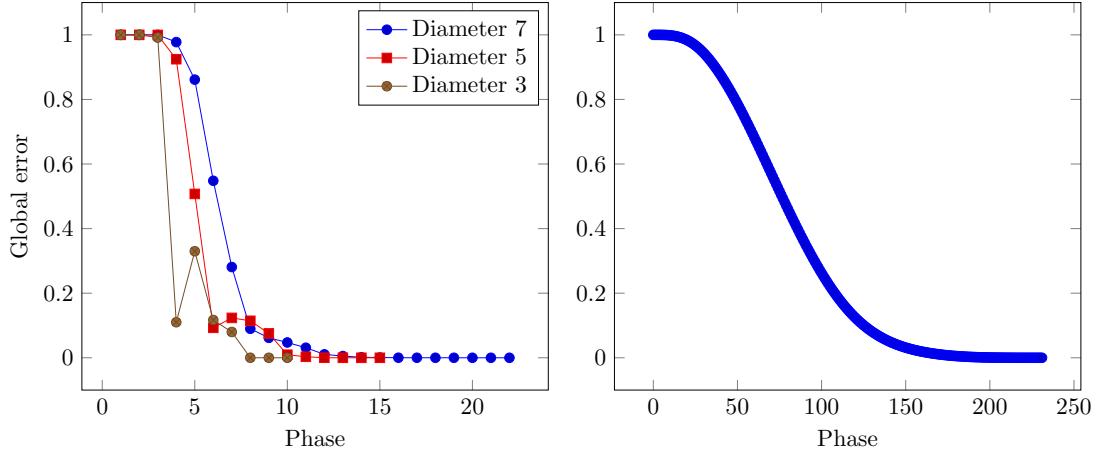
Figure 6: The (mean) global error as a function of time in randomly weighted Erdös-Renyi with 500 nodes (20 samples per diameter, where the noted diameter is the one of the underlying, unweighted graphs), and in a road network with 3353 vertices.

The last two convergence analyses of unweighted graphs are the ones of the e-mail network and of the autonomous system network (Figure 5 (right)). The convergence patterns are of similar nature, and are actually slightly more rapid then in the previous experiments — the convergence time is still $2\,\mathsf{Diam}+1$, but the error after $\mathsf{Diam}$ phases is smaller.

### 5.2.2 Weighted networks: Erdös-Renyi, road network

We now turn to describe our experiments with weighted Erdös-Renyi graphs (Figure 6 (left)). We build the graphs as before, and then assign random weights form the set $\{1, 2, 5\}$ to each edge, with expected weight 2. The shapes of the convergence curves are generally similar to the unweighted case, with one interesting difference: the global error is no longer monotonically decreasing, and, instead, it sometimes increases, but always before $\mathsf{Diam}$ phases. This phenomena is explained by the fact that the algorithm first finds paths of the least number of hops, but then updates them to paths with more hops but less weight. When such an update occurs, $S[v, t]$ may decrease, as many few-hops paths are replaced by a few (or one) lighter paths. This causes $B[v, t]$ to temporarily decrease, and the error to increase. This phenomena is moderated when the diameter grows larger.

A real world, large scale weighted network we study is composed of a large portion of the road network of Rome, represented by a weighted graph (Figure 6 (right)). In this network, whose diameter $\mathsf{Diam}(G)$ is very high (roughly 120 hops), we observe a very smooth convergence pattern, reminding the clean convergence patterns of the unweighted lattices (Figure 5 (left)). This could be an artifact of the large diameter or of the larger number of different weights.

## 5.3 Analysis of local convergence time

A different perspective on the convergence time is given by considering the number of nodes which converged completely over time. We consider two values: $T_D$, the time it takes for the Bellman-Ford algorithm to converge locally, i.e., until the distances are correctly computed; and, $T_C$, the time it takes for the betweenness centrality value to converge. A third value of interest is the convergence time of $S[v, t]$, the number of shortest path. However, in unweighted graphs our algorithm always computes this value exactly one round after the convergence of
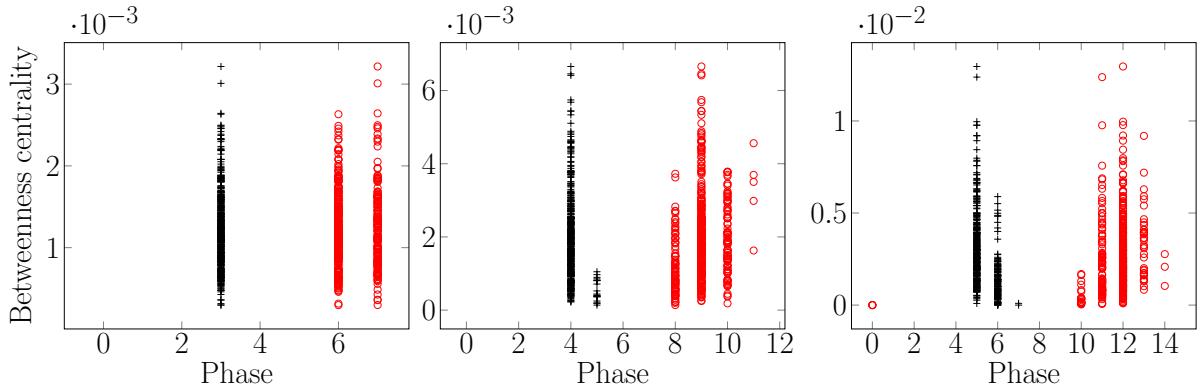
Figure 7: $T_D$ and $T_C$ for all the nodes of three Erdös-Renyi graphs with diameter 3, 5, and 7, respectively.

the length of the shortest paths ($T_D$). In weighted graphs, experiments show that this is almost always true as well, so we do not include this value in our plots.

### 5.3.1 Unweighted networks: Erdös-Renyi and e-mail

We study $T_D$ and $T_C$ for all the nodes of three unweighted Erdös-Renyi graphs (Figure 7), and for the e-mail and the autonomous system networks (Figure 8). All these networks present a somewhat similar behaviour. As predicted, all distances converge after Diam phases, and the betweenness centrality values converge after $2\,\mathsf{Diam}+1$ phases. The distribution of the convergence times of $T_D$ reassembles a Poisson distribution: most of the nodes, and especially nodes with a large betweenness centrality, correctly compute all their distances relatively fast — these are central nodes. Few, peripheral nodes, take more time to compute their distances correctly, and these nodes have lower betweenness centrality. The *eccentricity* of a node is the maximal distance from it to all other nodes, and it is always between $\mathsf{diam}\,/2$ and $\mathsf{diam}$. The time $T_D$ it takes a node to find all its distances to other nodes is thus exactly its eccentricity (up to a small additive constant), which better explains why $T_D$ is always between $\mathsf{diam}\,/2$ and $\mathsf{diam}$. Note that in unweighted graphs, $\mathsf{Diam} = \mathsf{diam}$.

In the case of the Erdös-Renyi graphs, when the random network has small diameter (3 and 5 in our examples) and is very dense, each node lies at least on one shortest path, and they all have non-zero betweenness centrality. In networks of larger diameter (Erdös-Renyi with diameter 7, e-mail network and autonomous system network), we notice that some nodes do not lie on any shortest path — these nodes' initial estimate of their betweenness centrality, 0, is correct, and thus they also have $T_C = 0$, appearing in the lower-left corner of the figures.

The convergence time of $T_C$ reassembles a normal distribution: nodes with low betweenness centrality may take longer or shorter times than nodes with very high betweenness centrality to compute their betweenness centrality values. This phenomena can also be explained by studying the eccentricity of the nodes: we have found that in general, nodes with large betweenness centrality have low eccentricities, of roughly $\mathsf{Diam}\,/2$ which causes faster convergence. For example, in the case of the autonomous system network, the average eccentricity of all nodes is approximately 6.6, while the average eccentricity of the ten nodes with highest betweenness centrality is roughly 5.5 (note that this network has diameter 9, so no node can have eccentricity smaller than 5).

Thus, after the computation of all distances and numbers of shortest paths, which takes roughly Diam phases, nodes of high betweenness centrality take only $\mathsf{Diam}\,/2$ more phases to compute their betweenness centrality. Nodes with low betweenness centrality may have
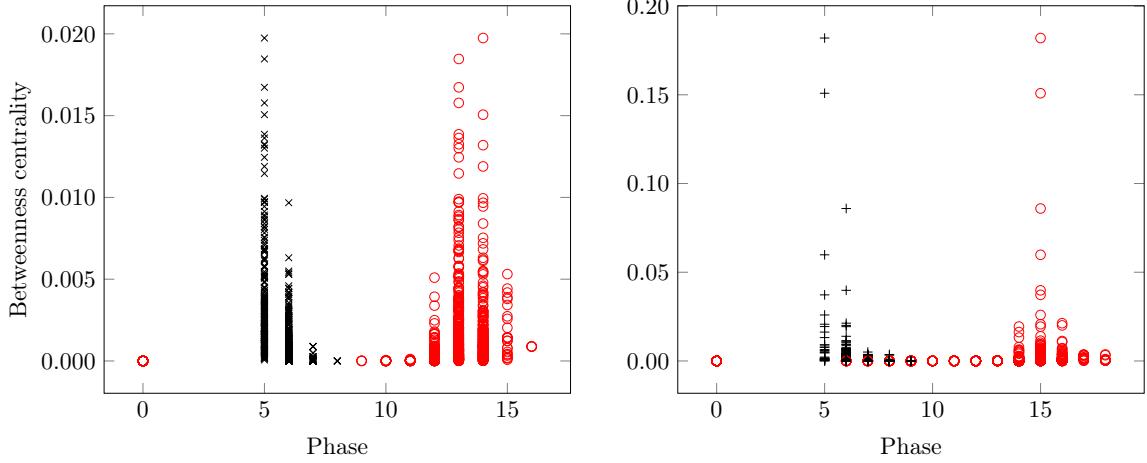
15

Figure 8: $T_D$ and $T_C$ for all the nodes of the e-mail network (left) and of the autonomous system network (right).

different eccentricities, which explains the wider scatter of their convergence times. To better understand the betweenness centrality convergence times $T_C$, we also check the *number of nodes* that converge in each phase, as shown, for the case of the autonomous system network, in Figure 9. In this figure, we omit roughly 65% of the nodes, that have bc = 0 and converge in 0 phases. The rest of the nodes present a normal distribution, where most nodes converge in roughly $\frac{3}{2}$ Diam phases.
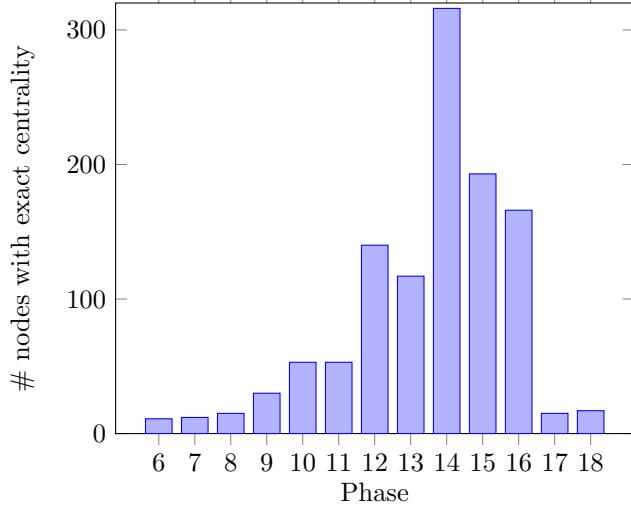


Figure 9: The number of nodes whose $C$ value converges to the exact centrality value as a function of the phase, in the case of the autonomous system network.

### 5.3.2 Weighted networks: road network

Finally, we revisit the weighted graph representing the road network of Rome. The scatter of betweenness centrality values and convergence times of $T_D$ and $T_C$ give a distribution very similar to the ones of the unweighted networks (see Figure 10 (left)). It can be observed that the convergence times of the distances, $T_D$, of nodes with high betweenness centrality is low — roughly Diam $/2$ — due to their low eccentricity. Other node take also time proportional
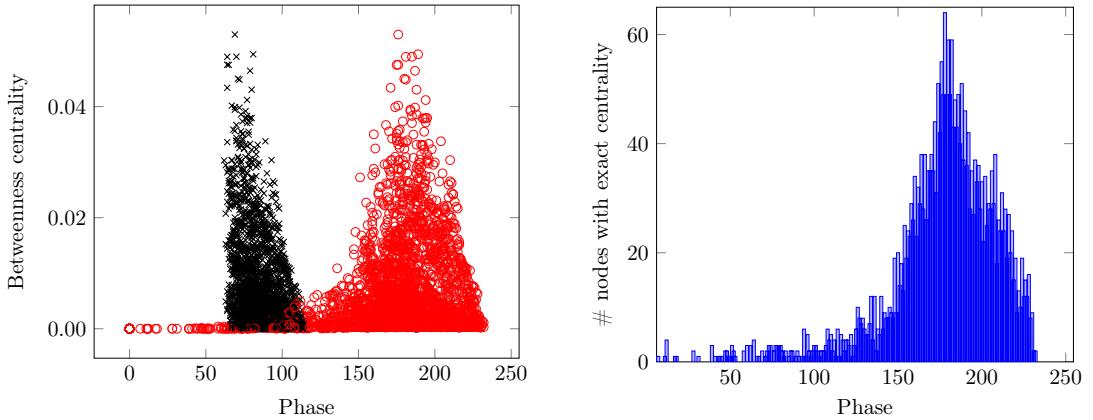
16

Figure 10: $T_D$ and $T_C$ for all the nodes (left) and number of nodes whose $C$ value converges to the exact centrality value as a function of time (right), in the case of the road network.

to their eccentricity, which is between diam $/2$ and diam (this is a graph theoretic fact). The convergence time of the betweenness centrality, $T_C$, is again scattered around $\frac{3}{2}$ Diam, with nodes of high betweenness centrality indeed converging in roughly this time, while other nodes may also converge faster or slower.

A unique phenomena observed here is that of nodes with betweenness centrality of 0, who's $T_C$ is *not* 0. This is a result of paths going through these nodes, which are not shortest paths (in terms of weight), but seem to be so before enough hops are considered. All these nodes have their $T_C$ less than diam, since in this time all shortest path distances are correct.

Once again, we also check how many nodes converge in each phase (see Figure 10 (right)). Here, roughly half the nodes converge in 0 rounds, with bc $= 0$. The rest of the nodes present a behaviour similar to before — most of them take roughly $\frac{3}{2}$ Diam phases to converge, while a few take longer or shorter, but never more than $2$ Diam $+1$.

## 6    Conclusion

In this paper, we have shown that betweenness centrality can be easily computed with only minimal modifications to the classical Bellman-Ford algorithm, and, hence, be integrated into distance-vector routing protocols. We have also presented an experimental analysis of the global and local convergence time of the proposed algorithm.

An interesting question left open by our work is designing bandwidth-adaptive distributed algorithms for betweenness centrality, assuming that the graph has a polynomial number of shortest paths (if this not the case, which would cause the $S$-values to have a linear number of bits, we could use a floating point representation with $O(\log n)$ bits to approximate the $S$-values, as described in [25]). Specifically, let CONGEST($B$) be the variant of the CONGEST model described in Section 1.2, in which at most $B$ words of $O(\log n)$ bits each can be sent through each link at each round. The distributed algorithm for weighted graphs described in this paper applies to the CONGEST($n$) model, and converges in $O($Diam$)$ rounds. The distributed algorithms for unweighted graphs described in [23, 26, 46] apply to the CONGEST(1) model, and converge in $O(n)$ rounds. We thus leave open the question of computing betweenness centrality of weighted graphs in $O(n/B + $Diam$)$ rounds in the CONGEST($B$) model, for every $1 \leq B \leq n$.

17

# References

[1] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47–97, 2002.

[2] David A. Bader, Shiva Kintali, Kamesh Madduri, and Milena Mihail. Approximating betweenness centrality. In *5th International Workshop on Algorithms and Models for the Web-Graph (WAW)*, pages 124–137, 2007.

[3] Miriam Baglioni, Filippo Geraci, Marco Pellegrini, and Ernesto Lastres. Fast exact computation of betweenness centrality in social networks. In *International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 450–456, 2012.

[4] Luca Baldesi, Leonardo Maccari, and Renato Lo Cigno. On the use of eigenvector centrality for cooperative streaming. *IEEE Communications Letters*, 21(9):1953–1956, 2017.

[5] Elisabetta Bergamini and Henning Meyerhenke. Fully-dynamic approximation of betweenness centrality. In *23rd Annual European Symposium on Algorithms (ESA)*, pages 155–166, 2015.

[6] Elisabetta Bergamini, Henning Meyerhenke, and Christian Staudt. Approximating betweenness centrality in large evolving networks. In *Conference on Algorithm Engineering and Experiments (ALENEX)*, pages 133–146, 2015.

[7] Aaron Bernstein and Danupon Nanongkai. Distributed exact weighted all-pairs shortest paths in near-linear time. In *STOC*, 2018.

[8] U. Brandes. A faster algorithm for betweenness centrality. *J. of Mathematical Sociology*, 25(2):163–177, 2001.

[9] U. Brandes. On variants of shortest-path betweenness centrality and their generic computation. *Social Networks*, 30(2):136–145, 2008.

[10] Ulrik Brandes and Christian Pich. Centrality estimation in large networks. *International Journal of Bifurcation and Chaos*, 17(7):2303–2318, 2007.

[11] Keren Censor-Hillel, Seri Khoury, and Ami Paz. Quadratic and near-quadratic lower bounds for the congest model. In *DISC*, 2017.

[12] DIMACS. 9th DIMACS Implementation Challenge - Shortest Paths. `http://users.diag.uniroma1.it/challenge9/`. Accessed: 2019-07-26.

[13] Shlomi Dolev, Yuval Elovici, and Rami Puzis. Routing betweenness centrality. *J. ACM*, 57(4):25:1–25:27, 2010.

[14] P. Erdös and A. Rényi. On random graphs. i. *Publicationes Mathematicae*, 6:290–297, 1959.

[15] Linton C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, 40(1):35–41, 1997.

[16] Silvio Frischknecht, Stephan Holzer, and Roger Wattenhofer. Networks cannot compute their diameter in sublinear time. In *SODA*, 2012.

[17] Robert Geisberger, Peter Sanders, and Dominik Schultes. Better approximation of betweenness centrality. In *10th Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 90–100, 2008.

[18] Edgar Gilbert. Random plane networks. *Journal of the Society for Industrial and Applied Mathematics*, 9:533–543, 1961.

[19] K.L. Goh, B. Kahng, and D. Kim. Universal behavior of load distribution in scale-free networks. *Physical Review Letters*, 87(27):1–4, 2001.

[20] R. Guimera, L. Danon, A. Diaz-Guilera, F. Giralt, and A. Arenas. Self-similar community structure in a network of human interactions. *Physical Review E*, 68:065103(R), 2033.

[21] Aric Hagberg, Dan Schult, and Pieter Swart. NetworkX. Software for complex networks. `https://networkx.github.io/`, 2019.

[22] Takanori Hayashi, Takuya Akiba, and Yuichi Yoshida. Fully dynamic betweenness centrality maintenance on massive networks. *Proceedings of the VLDB Endowment*, 9(2):48–59, 2015.

[23] Loc Hoang, Matteo Pontecorvi, Roshan Dathathri, Gurbinder Gill, Bozhi You, Keshav Pingali, and Vijaya Ramachandran. A round-efficient distributed betweenness centrality algorithm. In *24th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*, pages 272–286, 2019.

[24] Stephan Holzer and Roger Wattenhofer. Optimal distributed all pairs shortest paths and applications. In *PODC*, 2012.

[25] Q. Hua, H. Fan, M. Ai, L. Qian, Y. Li, X. Shi, and H. Jin. Nearly optimal distributed algorithm for computing betweenness centrality. In *2016 IEEE 36th International Conference on Distributed Computing Systems (ICDCS)*, pages 271–280, 2016.

[26] Qiang-Sheng Hua, Haoqiang Fan, Ming Ai, Lixiang Qian, Yangyang Li, Xuanhua Shi, and Hai Jin. Nearly optimal distributed algorithm for computing betweenness centrality. *2016 IEEE 36th International Conference on Distributed Computing Systems (ICDCS)*, pages 271–280, 2016.

[27] Qiang-Sheng Hua, Haoqiang Fan, Lixiang Qian, Ming Ai, Yangyang Li, Xuanhua Shi, and Hai Jin. Brief announcement: A tight distributed algorithm for all pairs shortest paths and applications. In *SPAA*, 2016.

[28] Riko Jacob, Dirk Koschützki, Katharina Anna Lehmann, Leon Peeters, and Dagmar Tenfelde-Podehl. *Network Analysis*, chapter Algorithms for Centrality Indices, pages 62–82. LNCS 3418. Springer, 2005.

[29] Miray Kas, Sandeep Appala, Chao Wang, Kathleen M. Carley, L. Richard Carley, and Ozan K. Tonguz. What if wireless routers were social? *IEEE Wireless Communications*, 19(6):36–43, 2012.

[30] Dimitrios Katsaros, Nikos Dimokas, and Leandros Tassiulas. Social network analysis concepts in the design of wireless ad hoc network protocols. *IEEE Network*, 24(6):23–29, 2010.

[31] Nicolas Kourtellis, Gianmarco De Francisci Morales, and Francesco Bonchi. Scalable online betweenness centrality in evolving graphs. *IEEE Trans. Knowl. Data Eng.*, 27(9):2494–2506, 2015.

[32] Christoph Lenzen and David Peleg. Efficient distributed source detection with limited bandwidth. In *PODC*, 2013.

[33] Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. `http://snap.stanford.edu/data`, June 2014.

[34] Yeon-Sup Lim, Daniel S. Menasché, Bruno Ribeiro, Don Towsley, and Prithwish Basu. Online estimating the k central nodes of a network. In *IEEE Network Science Workshop*, pages 118–122, 2011.

[35] Leonardo Maccari and Renato Lo Cigno. Betweenness estimation in olsr-based multi-hop networks for distributed filtering. *J. Comput. Syst. Sci.*, 80(3):670–685, 2014.

[36] Leonardo Maccari and Renato Lo Cigno. Pop-routing: Centrality-based tuning of control messages for faster route convergence. In *35th IEEE International Conference on Computer Communications (INFOCOM)*, pages 1–9, 2016.

[37] Leonardo Maccari, Lorenzo Ghiro, Alessio Guerrieri, Alberto Montresor, and Renato Lo Cigno. On the distributed computation of load centrality and its application to DV routing. In *37th IEEE Conference on Computer Communications (INFOCOM)*, pages 2582–2590, 2018.

[38] Leonardo Maccari, Quynh Nguyen, and Renato Lo Cigno. On the computation of centrality metrics for network security in mesh networks. In *IEEE Global Communications Conference (GLOBECOM)*, pages 1–6, 2016.

[39] Arun S. Maiya and Tanya Y. Berger-Wolf. Online sampling of high centrality individuals in social networks. In *14th Pacific-Asia Conference on Advances in Knowledge Discovery and Data Mining (PAKDD)*, pages 91–98, 2010.

[40] Eduardo Montijano, Gabriele Oliva, and Andrea Gasparri. Distributed estimation and control of node centrality in undirected asymmetric networks. Technical report, arXiv abs/1903.09689, 2019.

[41] M.E.J. Newman. Scientific collaboration networks. ii. shortest paths, weighted networks, and centrality. *Physical Review E*, 64:016132, 2001.

[42] Panagiotis Pantazopoulos, Merkourios Karaliopoulos, and Ioannis Stavrakakis. Distributed placement of autonomic internet services. *IEEE Trans. Parallel Distrib. Syst.*, 25(7):1702–1712, 2014.

[43] David Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, 2000.

[44] David Peleg, Liam Roditty, and Elad Tal. Distributed algorithms for network diameter and girth. In *ICALP*, 2012.

[45] Larry L. Peterson and Bruce S. Davie. *Computer networks: a systems approach (3. ed.)*. Morgan Kaufmann, 2003.

[46] Matteo Pontecorvi and Vijaya Ramachandran. Distributed algorithms for directed betweenness centrality and all pairs shortest paths. Technical report, arXiv abs/1805.08124, 2018.

[47] Rami Puzis, Yuval Elovici, Polina Zilberman, Shlomi Dolev, and Ulrik Brandes. Topology manipulations for speeding betweenness centrality computation. *J. Complex Networks*, 3(1):84–112, 2015.

[48] Matteo Riondato and Evgenios M. Kornaropoulos. Fast approximation of betweenness centrality through sampling. *Data Min. Knowl. Discov.*, 30(2):438–475, 2016.

[49] Matteo Riondato and Eli Upfal. ABRA: approximating betweenness centrality in static and dynamic graphs with rademacher averages. *ACM Transactions on ACM Transaction on Knowledge Discovery from Data*, 12(5):61:1–61:38, 2018.

[50] Andrés Vázquez Rodas and Luis J. de la Cruz Llopis. A centrality-based topology control protocol for wireless mesh networks. *Ad Hoc Networks*, 24:34–54, 2015.

[51] W. Wang and C. Y. Tang. Distributed computation of node and edge betweenness on tree graphs. In *52nd IEEE Conf. on Decision and Control*, pages 43–48, 2013.

[52] W. Wang and C. Y. Tang. Distributed computation of classic and exponential closeness on tree graphs. In *American Control Conf. (ACC)*, pages 2090–2095, 2014.

[53] Keyou You, Roberto Tempo, and Li Qiu. Distributed algorithms for computation of centrality measures in complex networks. *IEEE Trans. Automat. Contr.*, 62(5):2080–2094, 2017.

[54] Polina Zilberman, Rami Puzis, and Yuval Elovici. On network footprint of traffic inspection and filtering at global scrubbing centers. *IEEE Trans. Dependable Sec. Comput.*, 14(5):521–534, 2017.

[55] Katharina Anna Zweig and Michael Kaufmann. Decentralized algorithms for evaluating centrality in complex network. Technical report, Universität Tübingen, Germany, 2003.