Low-Congestion Shortcuts without Embedding*

Bernhard Haeupler
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, USA
haeupler@cs.cmu.edu

Taisuke Izumi
Nagoya Institute of Technology
Nagoya, Aichi Prefecture
466-8555, Japan
t-izumi@nitech.ac.jp

Goran Zuzic
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, USA
gzuzic@cs.cmu.edu

ABSTRACT

Distributed optimization algorithms are frequently faced with solving sub-problems on disjoint connected parts of a network. Unfortunately, the diameter of these parts can be significantly larger than the diameter of the underlying network, leading to slow running times. Recent work by [Ghaffari and Haeupler; SODA'16] showed that this phenomenon can be seen as the broad underlying reason for the pervasive $\Omega(\sqrt{n} + D)$ lower bounds that apply to most optimization problems in the CONGEST model. On the positive side, this work also introduced low-congestion shortcuts as an elegant solution to circumvent this problem in certain topologies of interest. Particularly, they showed that there exist good shortcuts for any planar network and more generally any bounded genus network. This directly leads to fast $O(D \log^{O(1)} n)$ distributed optimization algorithms on such topologies, e.g., for MST and Min-Cut approximation, given that one can efficiently construct these shortcuts in a distributed manner.

Unfortunately, the shortcut construction of [Ghaffari and Haeupler; SODA'16] relies heavily on having access to a bounded genus embedding of the network. Computing such an embedding distributedly, however, is a hard problem - even for planar networks. No distributed embedding algorithm for bounded genus graphs is in sight.

In this work, we side-step this problem by defining a slightly restricted and more structured form of shortcuts and giving a novel construction algorithm which efficiently finds a shortcut which is, up to a logarithmic factor, as good as the best shortcut that exists for a given network. This new construction algorithm directly leads to an $O(D \log^{O(1)} n)$ -round algorithm for solving optimization problems like MST for any topology for which good restricted shortcuts exist -

without the need to compute any embedding. This includes the first efficient algorithm for bounded genus graphs.

CCS Concepts

• **Mathematics of computing** → *Network optimization*;
• **Theory of computation** → *Distributed computing models*; *Graph algorithms analysis*;

Keywords

low-congestion shortcuts; planar graphs; bounded-genus graphs; MST; CONGEST model

1. INTRODUCTION

1.1 Background and Motivation

Consider the problem of finding the minimum spanning tree (MST) on a distributed network with n independent processing nodes. The network is abstracted as a graph $G = (V, E)$ with n nodes and diameter D . The nodes communicate by synchronously passing $O(\log n)$ -bit messages to each of its direct neighbors. The goal is to design algorithms (protocols) that minimize the number of synchronous message passing rounds before the nodes collaboratively solve the optimization problem.

The message-passing setting we just described is a model called CONGEST [21]. The MST problem can be solved in such a setting using $O(\sqrt{n} \log^* n + D)$ rounds of communication [13]. Moreover, and perhaps more surprisingly, this bound was shown to be the best possible (up to polylogarithmic factors). Specifically, there are graphs in which one cannot do any better than $\tilde{\Omega}(\sqrt{n} + D)$ [1, 3, 22]. While clearly no algorithm can solve any global network optimization problem faster than $\Omega(D)$, the $\tilde{\Omega}(\sqrt{n})$ factor is harder to discern. To make matters worse, the $\tilde{\Omega}(\sqrt{n} + D)$ lower bound was shown to be far reaching. It applies to a multitude of important network optimization problems including MST, minimum-cut, weighted shortest-path, connectivity verification and so on [1].

While this bound precludes the existence of more efficient algorithms in the general case, it was not clear does it hold for special families of graphs. This question is especially important because any real-world application on huge networks should exploit the special structure that the network provides. The mere existence of “hard” networks for which one cannot design any fast algorithm might not be a limiting factor.

*This work was supported in part by KAKENHI No. 15H00852 and 16H02878 as well as NSF grants CCF-1527110 “Distributed Algorithms for Near Planar Networks” and NSF-BSF grant “Coding for Distributed Computing”.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PODC'16, July 25 - 28, 2016, Chicago, IL, USA

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-3964-3/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2933057.2933112>

In the first result that utilizes network topology to circumvent the lower bound, Haeupler and Ghaffari designed an $O(D \log^{O(1)} n)$ -round distributed MST algorithm for planar graphs [7]. Note that this algorithm offers a huge advantage over older results for planar graphs with small diameters.

They achieve this by introducing an elegant abstraction for designing distributed algorithms named **low-congestion shortcuts**. Their methods could in principle be used to achieve a similar result for genus-bounded graphs, but their presented algorithms have a major technical obstacle: they require a surface embedding of the planar/genus bounded graph to construct the low-congestion shortcuts. While computing a distributed embedding for planar graphs has a complex $O(D \log^{O(1)} n)$ -round solution [6], this remains an open problem for genus-bounded graphs [7].

This paper side-steps the issue by vastly simplifying the construction of low-congestion shortcuts. We define a more structured version of low-congestion shortcuts called **tree-restricted shortcuts** and propose a simple and general distributed algorithm for finding them. The algorithm is completely oblivious to any intricacies of the underlying topology and finds universally near-optimal tree-restricted shortcuts. As a simple consequence of our construction technique we get a $O(gD \log^{O(1)} n)$ -round algorithm for genus g graphs, which is a novel result. We believe that this simplicity makes the algorithm usable even in practice.

1.2 A Brief Overview of Low-Congestion Shortcuts

We now give a short introduction to the general low-congestion shortcuts as defined in [7]. Consider the following scenario, which is a recurring theme throughout distributed approaches for many network optimization problems:

A graph G is partitioned into a number of disjoint individually-connected parts $P_1, P_2, \dots, P_N$, and we need to compute a (typically simple) function for each of the parts in isolation.

A classical example for such a scenario is the 1926 algorithm of Boruvka [20] for computing Minimum Spanning Tree (MST): starting with a trivial partition of each node being its own part, in every iteration each part computes the minimum-weighted outgoing edge and merges with the part incident to this edge. After $O(\log n)$ iterations, we arrive at the MST, where n is the number of nodes in G .

A key concern in designing a distributed version of Boruvka's algorithm is finding good communication schemes that allow each part to collaborate with other nodes inside the same part and without interference from other parts. While a natural solution would be to allow communication only inside the same part, this could take a long time. The problem appears when the diameter of a part in isolation is much larger than the diameter D of the original graph G .

To overcome this issue, Ghaffari and Haeupler [7] introduced **low-congestion shortcuts**: each part P_i is given a subgraph of extra edges H_i that it can use to more efficiently communicate within itself. More precisely, each part P_i is associated with a **shortcut subgraph** H_i and is permitted to use $G[P_i] + H_i$ for communication.

To measure the quality of a shortcut, we characterize it with two quality parameters: **congestion** and **dilation**. A shortcut has congestion c and dilation d if *i*) the diameter of every subgraph $G[P_i] + H_i$ is at most d , and *ii*) every

edge is assigned to at most c different subgraphs $G[P_i] + H_i$. Given a shortcut with congestion c and dilation d , we can solve problems such as MST and Min-Cut approximation in $O((c + d) \log^{O(1)} n)$ rounds [7]. Therefore, designing a distributed algorithm can be reduced to finding good-quality shortcuts.

While the pervasive $\Omega(\sqrt{n} + D)$ lower bound clearly implies we cannot find shortcuts with congestion + dilation = $O(D)$ on general graphs, this might not be the case on specific families of graphs. For example, planar graphs always offer $O(D \log^{O(1)} D)$ congestion and dilation shortcuts, thus bypassing the $\tilde{\Omega}(\sqrt{n} + D)$ lower bound [7].

1.3 Our Contribution

Roughly speaking, there are two challenges in the design of shortcut-based algorithms. Let $\mathcal{G}$ be the target class we want to design distributed algorithms. The first challenge is to identify the (small) values of c and d such that $\mathcal{G}$ has shortcuts with congestion c and dilation d . This is purely a graph-theoretic problem. The second challenge is to convert the existential result proved by the first challenge to the constructive result, i.e., we must design a distributed algorithm constructing efficient shortcuts for that class. This is a distinct problem of theoretical distributed computing.

A natural idea of lowering the barrier of the algorithm design is to invent a generic algorithm which finds a congestion c and dilation d shortcut for the best (or approximately best) c and d , which provides an automatic conversion of the existential result to the constructive one. Unfortunately, the known construction for planar graphs [7] is far from such a generality: as we already mentioned, it strongly depends on the distributed planar embedding algorithm [6], and thus not applicable to any other graph class. This is also a primary reason why the construction for planar graphs in [7] cannot be extended even for bounded genus graphs.

The primary contribution of this paper is to present a new generic algorithm for constructing shortcuts, which partially resolve the issue mentioned above. While our algorithm does not cover all existential results, it is applicable to an interesting and useful case – **tree-restricted shortcuts**, which is a class of shortcuts newly defined in this paper. The details of our contribution are summarized as follows:

- In Section 4 we introduce a new class of shortcuts, called tree-restricted shortcuts, which only use edges of some fixed spanning tree $T \subseteq G$. More precisely, $H_i \subseteq T$ for each part P_i . We introduce a new quality parameter called **block parameter**, which is defined as an upper bound on the number of connected components of H_i that intersect P_i (over all $i \in [N]$). Note that these components are subtrees of T . The block parameter can be seen as a stronger version of dilation and will often be used instead. In Section 4.3 we propose deterministic algorithms for broadcast, convergecast, and leader election (for all parts in parallel) utilizing tree-restricted shortcuts, which are, simpler and faster compared with the general-case randomized algorithms shown in [7].
- In Section 5 we present a generic algorithm for constructing tree-restricted shortcuts. Let T be a spanning tree of G with depth D and assume there exists a tree-restricted shortcut on $T \subseteq G$ with congestion c and block parameter b . We describe an algo-

rithm that constructs a tree-restricted shortcut with congestion $O(c \log N)$ and block parameter $O(b)$ in $O(D \log n \log N + bD \log N + bc \log N)$ CONGEST rounds. It is also possible to run our algorithm in the environment where the system is not aware of the value of b and/or c with extra $\log(bc)$ factor, as described in Section 6.

- An important consequence of our algorithm is to provide the first distributed algorithm constructing a good shortcut for genus- g graphs. Fortunately, the known result for genus- g graphs exhibits the existence of tree-restricted shortcuts with congestion $O(gD \log D)$ and block parameter $O(\log D)$ for an arbitrary BFS tree T of depth D . Thus in Section 4.4 we can obtain a distributed algorithm constructing a tree-restricted shortcut with congestion $O(gD \log D \log N)$ and block $O(\log D)$ for graphs with genus at most g . For bounded genus graphs (i.e. $g = O(1)$), the algorithms based on our shortcut construction achieves the near-optimal time complexity (up to a polylogarithmic factor). According to the very recent unpublished result that is still in preparation, a similar result is obtained for graphs with bounded pathwidth and treewidth.

2. RELATED WORK

The complexity theoretic issues in the design of distributed graph algorithms for the CONGEST model have received much attention in the last decade, and got an extensive progress for many problems: Minimum-spanning tree [5, 12, 13, 22], Maximum flow [8], Minimum Cut [9, 19], Shortest paths and Diameter [4, 10, 11, 15, 16, 17, 18], and so on. Most of those problems have $\Theta(\sqrt{n} + D)$ -round upper and lower bounds for some sort of approximation guarantee [1, 2, 9, 15, 22]. The guarantee of exact results sometimes yield a nearly linear-time bound [4]. Note that almost all lower bounds above holds for small diameter graphs. Thus, in any case, the general lower bound is more expensive than the universal lower bound of $\Omega(D)$ rounds.

On the positive side, distributed algorithms typically use a variety of ideas. In an effort to unify them in an elegant framework, Ghaffari and Haeupler introduced **low-congestion shortcuts** [7]. Specifically, their ideas can be turned into a very short and clean $O((D + \sqrt{n}) \log n)$ round MST algorithm for general graphs. Furthermore, low-congestion shortcuts can serve as a simple explanation of the pervasive $\tilde{\Omega}(D + \sqrt{n})$ lower bound. However, the main contribution of their techniques is a $\tilde{O}(D)$ -round algorithm for planar graphs. To the best of our knowledge, it is the first attempt that considers a non-trivial popular graph class.

3. PRELIMINARIES

In this section, we formally define the CONGEST model and then recap the definitions of low-congestion shortcuts from [7].

3.1 CONGEST Model

We work in the classical CONGEST model [21]. In this setting, a network is given as a connected undirected graph $G = (V, E_G)$ with diameter D . Initially, nodes only know their immediate neighbors and they collaborate to compute some global function of the graph like the MST. Communication occurs in synchronous rounds; during a round, each

node can send $O(\log n)$ bits to each of its neighbors (note that the nodes also know some polynomially tight bound on n , otherwise sending $O(\log n)$ bits does not make sense). The nodes always correctly follow the protocol and never fail. The goal is to design protocols that minimize the resource of time - the number of rounds before the nodes compute the solution.

We now precisely formalize the notion of solving a problem in this model, e.g. how is the input and output given. While the formalization is specifically given for the MST, any other problem is completely analogous. All nodes synchronously wake up in the first round and start executing some given protocol. Every node initially only knows its immediate neighbors and the weight of each of its incident edges. After a specific number of rounds, all nodes must simultaneously output $i)$ the weight of the computed MST τ $ii)$ for each edge e incident to it, a 0/1 bit indicating if $e \in \tau$.

3.2 Low-Congestion Shortcuts

Let G be a undirected graph along with a node partition $\mathcal{P} = (P_1, P_2, \dots, P_N)$. **Low-congestion shortcuts** intuitively augment each part P_i with extra edges H_i that may be used to communicate within a part more efficiently. With a small abuse of notation, in the following we use the symbol H_i to indicate both the edge set and the subgraph induced by the set. As communication for part P_i occurs on $G[P_i] + H_i$, it is natural to try to minimize the diameters of those subgraphs. Hence we define **dilation** as an upper bound on the diameter of any shortcut subgraph $G[P_i] + H_i$. On the other hand, assigning an edge to almost every part will lead to over-congestion on that edge. Therefore, we define another quality measure of a shortcut, **congestion**, as an upper bound on the number of shortcut subgraphs $G[P_i] + H_i$ that contain any edge e .

Definition 1. Let $G = (V, E_G)$ be an undirected graph with vertices subdivided into **disjoint and connected** subsets $\mathcal{P} = (P_1, P_2, \dots, P_N)$, $P_i \subseteq V$. In other words, $G[P_i]$ is connected and $P_i \cap P_j = \emptyset$ for $i \neq j$. The subsets P_i are called **parts**. We define a **shortcut** $\mathcal{H}$ as a tuple of N **shortcut subgraphs** $(H_1, H_2, \dots, H_N)$, $H_i \subseteq G$. A shortcut is characterized by the following parameters:

- $i)$ $\mathcal{H}$ has congestion c if each edge $e \in E_G$ is used in at most c different subgraphs $G[P_i] + H_i$, i.e. $\forall e \in E_T : |\{i : e \in H_i\}| \leq c$.
- $ii)$ $\mathcal{H}$ has dilation d if the diameter of any subgraph $G[P_i] + H_i$ is at most d .

The parameters determine the efficiency of communications facilitated by the shortcut. For example, Ghaffari and Haeupler show in [7] that one can solve the Minimum Spanning Tree and Min-Cut problems in $O((\text{congestion} + \text{dilation}) \log^{O(1)} n)$, given an efficient algorithm for finding shortcuts with parameters congestion and dilation. Note that congestion and dilation are traditional parameters that are extensively used in routing [7, 14].

4. TREE-RESTRICTED SHORTCUTS

In this section, we define tree-restricted shortcuts: a narrower notion of shortcuts which are $i)$ simpler to work with,

Figure 1: Illustration of a T -restricted shortcut subgraph with respect to the part P_i , composed of block components b_1, b_2, b_3 and b_4 .

ii) often equally powerful as the general shortcuts, *iii)* offer deterministic routing schemes and, most importantly, *iv)* can be efficiently constructed on any graph that contains them. Following the definitions, we rephrase the results of [7] in our new terms, showcase an efficient deterministic routing scheme on them, and finally state our main result and show its applications.

4.1 Definition

Tree-restricted shortcuts are shortcuts with the additional property that any shortcut subgraph H_i is restricted to some spanning tree T . The user of the shortcut can typically fix any tree T , so a cogent choice would be the BFS tree because of its optimal depth.

Definition 2. Let $\mathcal{H} = (H_1, H_2, \dots, H_N)$ be a T -restricted shortcut on the graph $G = (V, E_G)$ with respect to the parts $\mathcal{P} = (P_i)_{i=1}^N$. Given a rooted spanning tree $T = (V, E_T) \subseteq G$ we say that a shortcut $\mathcal{H}$ is **T -restricted** if for each $i \in [N]$, $H_i \subseteq E_T$ i.e. every edge of H_i is a tree edge of T .

Congestion and dilation are still well-defined for tree-restricted shortcuts. However, it is more convenient to use an alternative **block parameter** in place of dilation. The block parameter upper-bounds the number of connected components of each H_i that intersects P_i . Note that, while $G[P_i]$ and therefore $G[P_i] + H$ are connected, H_i by itself might not be. The intersection property ensures that we do not count components that have no vertices in P_i .

Definition 3. Let $\mathcal{H} = (H_1, H_2, \dots, H_N)$ be a T -restricted shortcut on the graph $G = (V, E_G)$ with respect to the parts $\mathcal{P} = (P_i)_{i=1}^N$. Fix a part P_i and consider the connected components of the spanning subgraph (V, H_i) . If such a connected component intersects P_i we call it a **block component**. Furthermore, we define the **block parameter** b of $\mathcal{H}$ to be any upper bound to the number of block components for all parts.

A block parameter implies a bound on dilation, hence the block parameter can be seen as a stronger measure of quality. Lemma 1 argues that a block parameter of b implies the dilation of $b(2 \cdot \text{depth}(T) + 1)$. The Lemma also suggests that it is often beneficial to fix T to a BFS tree of G , thereby having asymptotically minimal depth. In that case, the depth of T is at most the diameter of the original graph, namely D . For this reason throughout this paper we denote the diameter of G and the depth of T by the same symbol D .

LEMMA 1. *Let T be a spanning tree with depth D and let $\mathcal{H} = (H_i : i \in [N])$ be a T -restricted shortcut with congestion c and block parameter b with respect to parts $\mathcal{P} = (P_i : i \in [N])$. Then the dilation of $\mathcal{H}$ is at most $b(2D + 1)$.*

PROOF. Fix $i \in [N]$. If we contract every block component of H_i into a supernode and remove all other nodes, supergraph will contain $b' \leq b$ supernodes and will be connected (because $G[P_i]$ is connected). Hence its diameter is $b' - 1 \leq b - 1$. Every supernode consists of a block component of diameter $2D$, so the diameter of H_i is at most $2bD + b - 1 < b(2D + 1)$. $\square$

Distributed representation of a tree-restricted shortcut:

Before we proceed to describe algorithms for shortcut routing and construction, we quickly specify here more precisely how a shortcut is represented distributively, i.e., what information regarding the shortcut any node is supposed to know in order to make the various routing algorithms on top of a shortcut efficient.

Formally, we say that a T -restricted shortcut $\mathcal{H}$ is computed when each node v knows *i)* T -depth each of its neighbors and itself *ii)* subset of incident edges that are tree edges of T *iii)* all the part IDs that can use v 's parent edge as well as the depth of their respective block component root. For the sake of clarity, the described construction algorithms do not go into details about the computation of each of those properties. However, they can be easily augmented to compute them explicitly.

4.2 Shortcuts on Genus-Bounded and Planar Graphs

Tree-restricted shortcuts are particularly useful on genus-bounded (e.g. planar) graphs. In particular, we can reinterpret the low-congestion result of Haeupler and Ghaffari [7] using our notation.

THEOREM 1 (HAEUPLER AND GHAFFARI [7]). *Let G be a graph with genus g and diameter D , and let T be any tree with depth D (e.g. BFS tree). There exists a T -restricted shortcut with congestion $O(gD \log D)$ and block parameter $O(\log D)$.*

The paper originally also provided a $O(D \log D)$ upper bound on the dilation of the shortcut. However, this bound can be implicitly recovered from Lemma 1 and block parameter $O(\log D)$. Note that the Theorem proves only the existence of such shortcuts. While the original paper does describe an algorithm that can in principle be used to compute them, it requires an embedding of G on a surface of genus g . It is an open problem to compute such an embedding efficiently in the CONGEST model [7].

4.3 Routing on Tree-Restricted Shortcuts

In this section, we show how to use tree-restricted shortcuts to efficiently communicate within parts. The tree-restricted structure of the shortcut allows for simpler and more efficient routing methods than general shortcuts. The main reason for this is that distributed approaches for various network optimization problems often use broadcasting and converging as primitives. However, such tasks can be efficiently and deterministically solved on subtrees, even when multiple (non-disjoint) subtrees have to execute the task in parallel. Lemma 2 formalizes this statement.

LEMMA 2 (ROUTING ON TREES). *Let T be a tree of depth D . Given a family of subtrees such that any edge of T is contained in at most c subtrees, there is a simple deterministic algorithm that can perform a convergecast/broadcast on all of the subtrees in $O(D + c)$ CONGEST rounds.*

Specifically, for convergecasts, if multiple messages are scheduled over the same edge, the algorithm forwards the packet with the smallest depth of the subtree root, breaking ties with the smallest ID of the subtree.

PROOF. The convergecast and broadcasts operations are symmetric, so we will only prove the Lemma for convergecasts.

Let v be a node of T . We will prove that no message gets transmitted along v 's parent edge after $h_v + c$ rounds where h_v is the height of v (distance to the farthest leaf in its subtree). Note that any message that gets transmitted along v 's parent edge must belong to a subtree that contains that edge. Let $\mathcal{S}_v = (s_1, s_2, \dots, s_k)$ be the tuple of subtrees that contain v 's parent edge, ordered by their priority (as described in the statement). In particular, we say that s_i has priority i in node v . The congestion condition stipulates that $k \leq c$.

We will prove by induction that for $i \in [k]$ the message associated with s_i will be transmitted no later than after $h_v + i$ rounds. The claim clearly holds for the leafs of T . Note that *i*) the relative priority-ordering between elements $\mathcal{S}_v$ is unchanged in any node of T *ii*) any subtree $s \notin \mathcal{S}_v$ that contains any child of v will have lower priority than any subtree in $\mathcal{S}_v$.

Fix s_i . By the time $h_v + i$, all of the messages corresponding to $\{s_1, \dots, s_{i-1}\}$ will be sent by the induction hypothesis, so it is sufficient to argue that at time $h_v + i$, v has received messages corresponding to s_i from all of its children contained in s_i . But this is exactly the induction hypothesis as for any child w , its height $h_w \leq h_v - 1$ and the priority of s_i is at most i or $w \notin s_i$. Hence v will send the message corresponding to s_i in round h_v or before. $\square$

Convergecasting/broadcasting on a tree is helpful in tree-restricted shortcut routing because we can intuitively envision each shortcut subgraph H_i as a family of subtrees (in our notation: block components). This communication within each block component will be the main building-block of primitives that operate on entire parts in parallel.

THEOREM 2 (ROUTING ON TREE-RESTRICTED SHORTCUTS). *On a T -restricted shortcut with congestion c and block parameter b there are deterministic distributed algorithms for:*

- i) *Electing a leader for each of the parts in parallel.*
- ii) *Convergecasting $O(\log n)$ -bit messages to the leader of each part in parallel.*
- iii) *Broadcasting a $O(\log n)$ -bit message from the leader of each part in parallel.*

Each algorithm takes $O(b(D + c))$ CONGEST rounds.

PROOF. All of these algorithms have a common flavor: for each part we perceive its shortcut subgraph as a supergraph of at most b supernodes where each supernode is a block component. We proceed to describe each of the algorithms on the supergraph and implicitly assume that intra-block communication happens after each step of the algorithm.

Communication within block components can be done in parallel using Lemma 2: all the nodes of a block component convergecast the relevant information to the block-root and subsequently the block-root broadcasts the result back.

Electing a leader for each part is performed by electing a leader for each supernode (block component) and broadcasting the leader to all neighborhood supernodes for b steps. Every supernode keeps the smallest leader ID ever seen as its current leader. After b rounds all the supernodes have the same leader. The algorithm requires $O(b(D + c))$ rounds as each of the b broadcasting steps is followed by an $O(D + c)$ intra-block communication step.

Broadcasting/convergecasting from/to the leader can be done by building a BFS tree from the leader-supernode. We can utilize the standard distributed BFS algorithm on the supergraph requiring $O(b)$ steps. The algorithm similarly requires $O(b(D + c))$ rounds as each of the $O(b)$ BFS steps is followed by an $O(D + c)$ intra-block communication step. $\square$

We also state a simple technical lemma that will be needed for the tree-restricted shortcut construction.

LEMMA 3. *Given a T -restricted shortcut with congestion c , a deterministic distributed algorithm can find all parts whose designated shortcut subgraph has at most b' block components. The algorithm executes in $O(b'(D + c))$ rounds.*

PROOF. Similarly to the proof of Theorem 2, we consider the supergraph of each shortcut subgraph for each part. We need to find all parts whose supergraphs have at most b' supernodes.

Each supernode broadcasts its leader for exactly b' rounds and every supernode keeps the minimum ID as their current leader. Subsequently, each leader r (there may be multiple ones as we have not bounded the block parameter) tries to build a BFS tree comprised of all the nodes that believe r is the leader. We can detect the existence of multiple leaders as in that case each BFS tree will contain two neighboring supernodes in different BFS trees and report failure. If this is not the case (all the supernodes of a part belong to the same BFS tree), we can convergecast the number of supernodes back to the root and subsequently broadcasts their count back. $\square$

Comparison with routing on general shortcuts: Ghafari and Haeupler [7] give a method for routing on general shortcuts in $O(\text{dilation} \cdot \log n + \text{congestion})$ rounds that is *i*) randomized and *ii*) assumes a leader is already elected for each part. They describe a process of leader election via a complicated randomized bootstrapping process that takes $O(\text{dilation} \cdot \log^2 n + \text{congestion} \cdot \log n)$ rounds. We contrast those results with our current tree-restricted shortcut routing where leader election is essentially no more difficult than broadcast/convergecasting and the routing is simpler and deterministic. The downside is that non-tree-restricted shortcuts sometimes offer better quality guarantees and therefore better performance.

4.4 Main Result and Applications

The main contribution of the paper is to introduce a general framework for finding good-quality shortcuts in graphs where the only assurance is that they exist. In other words, no topology assumption is assumed.

THEOREM 3. *Let G be a graph with a spanning tree $T \subseteq G$ such that there exists a T -restricted shortcut with congestion c and block parameter b . There exists a distributed algorithm that finds a T -restricted shortcut with congestion $O(c \log N)$ and block parameter $3b$ with high probability. The shortcut can be found in $O(D \log n \log N + bD \log N + bc \log N)$ rounds.*

We note that the Theorems 1 and 3 immediately give a novel result: an algorithm for constructing shortcuts on bounded genus graphs.

COROLLARY 1. *Given a genus- g graph with diameter D and N parts there is a distributed algorithm that computes a tree-restricted shortcut with congestion $O(gD \log D \log N)$ and block parameter $O(\log D)$ in $O(gD \log^2 D \log N)$ rounds.*

Next, we explain how to use tree-restricted shortcuts to distributively compute the Minimum Spanning Tree (MST) on genus- g graphs. Similarly to [7], we incorporate the shortcuts into the classic 1926 algorithm of Boruvka [20].

LEMMA 4. *Given a genus- g graph with n nodes and diameter D , there is a distributed algorithm that computes the Minimum Spanning Tree in $O(gD \log^2 D \log^2 n)$ rounds.*

For completeness we give a brief proof outline:

PROOF. Boruvka's algorithm runs in $O(\log n)$ phases. Each phase starts with a partition of the graph into connected parts and a computed MST for each part. Initially, the algorithm starts with the trivial partition in which each node is in its own part. At each phase, each part P_i suggests a merge along the minimum-weighted edge going out of P_i . It is well-known that all such edges belong to some MST. By computing a tree-restricted shortcut for each part in $O(gD \log^2 D \log n)$ rounds and using our convergecast algorithm on it in $O(gD \log^2 D)$ rounds we can compute the min-weight outgoing edges from each part. The only slight technical difficulty that remains is to assign IDs to parts which have merged. While we can communicate efficiently within each part, using the previously computed shortcuts, many parts could chain together to form a new part. This can be avoided by restricting the merge shapes to be star graphs: each part can become a **head** or **tail** with probability $\frac{1}{2}$ and we are only allowed to merge tails to heads. The number of phases remains $O(\log n)$ as every minimum-weighted outgoing edge will be used for merging with probability at least $\frac{1}{4}$, thus reducing the expected number of parts by a constant. $\square$

5. CONSTRUCTING TREE-RESTRICTED SHORTCUTS

In this section, describe an algorithmic framework that solves the problem of finding near-optimal tree-restricted shortcuts.

5.1 Overview over the Algorithmic Framework

Our algorithm **FindShortcut** uses two separate subroutines:

- **Core:** This subroutine finds a good-quality shortcut with respect to at least a constant fraction of the parts. As a prerequisite, we must compute and fix a tree T with depth D such there exists a T -restricted shortcut with congestion c and block parameter b . Note that we only assume its existence.

LEMMA 5. *Let T be a spanning tree with depth D and assume there exists a T -restricted shortcut with congestion c and block parameter b . The subroutine **CoreFast** finds a T -restricted shortcut $\mathcal{H}' = (H'_i)_{i=1}^N$ with the following properties:*

- The congestion of $\mathcal{H}'$ is at most $8c$ with high probability.*
- There exists a subset of parts $\mathcal{P}' \subseteq \mathcal{P}$ with size at least $|\mathcal{P}'| \geq \frac{N}{2}$ such that the shortcut subgraphs corresponding to parts in $\mathcal{P}'$ have block parameter $3b$.*

The subroutine takes $O(D \log n + c)$ CONGEST rounds to execute. Upon completion, each node knows for each of its incident edges which parts are they assigned to in $\mathcal{H}'$.

We divide out the exposition of the core subroutine in two versions: a deterministic and simpler **CoreSlow** requiring $O(D \cdot c)$ rounds; and a randomized **CoreFast** requiring $O(D \log n + c)$ rounds. We note that the **CoreFast** subroutine is the only randomized building block of our framework. Therefore we can replace it with a deterministic (albeit slower) version at a cost of an addition $\frac{c}{\log n}$ factor.

- **Verification:** This subroutine is used to check which of the T -restricted shortcut subgraphs found by the core subroutine have sufficiently small block parameter (in particular, at most $3b$).

LEMMA 6. *Given a tree T with depth D and a tentative T -restricted shortcut $\mathcal{H}'$ with congestion c , the deterministic subroutine **Verification** finds all parts $\mathcal{P}' \subseteq \mathcal{P}$ whose designated shortcuts have at most b block components. The subroutine takes $O(b(D + c))$ CONGEST rounds to execute. Upon completion, each node knows whether its part is in the set $\mathcal{P}'$ or not.*

We call the parts whose designated shortcut subgraphs have this property as **good** and the rest as **bad**. **FindShortcut** runs the core subroutine followed by a verification step after which parts that have been marked as good are removed. This is repeated until no more bad parts remain.

5.2 FindShortcut Algorithm

Before we dive into the **FindShortcut** subroutine we must fix a spanning tree T . As the depth of T determines the efficiency of our framework, we can choose T to be a BFS tree rooted at any node of the graph G . This choice ensures that the depth of T is *i)* asymptotically optimal and *ii)* bounded by the diameter of G . For this reason throughout this paper we denote the diameter of G and the depth of T by the same symbol D .

Computing a BFS tree T in our distributed CONGEST model is a standard subroutine and can be computed in $O(D)$ rounds. Henceforth we assume that a tree T with depth D is computed.

FindShortcut subroutine: We run the **CoreFast** subroutine that computes a shortcut $\mathcal{H}'$ with congestion $8c$, but possibly an unacceptably large block parameter. The next step is to run the **Verification** subroutine that finds all parts whose computed shortcut subgraphs have at most $3b$ block components. We call those parts **good** and fix their

computed shortcut subgraphs. The subroutine is iteratively repeated until all the parts have been marked as good. The check can be executed via a $O(D)$ convergecast on the entire tree T .

We next prove Theorem 3 properties of **FindShortcut** using Lemma 5 and Lemma 6.

THEOREM 4 (RESTATED THEOREM 3). *Let G be a graph with a spanning tree $T \subseteq G$ such that there exists a T -restricted shortcut with congestion c and block parameter b . There exists a distributed algorithm that finds a T -restricted shortcut with congestion $O(c \log N)$ and block parameter $3b$ with high probability. The shortcut can be found in $O(D \log n \log N + bD \log N + bc \log N)$ rounds.*

PROOF. Let $\mathcal{H}'$ be the set of all shortcut subgraphs that have been marked as good through the lifetime of the entire subroutine. As any shortcut subgraph in $\mathcal{H}'$ has block parameter $3b$ and congestion $8c$ w.h.p., it only remains to show that the algorithm terminates.

By Lemma 5 in each iteration we find a shortcut with congestion $8c$ and block parameter $3b$ for at least a half of the parts that have not been marked as good, w.h.p. This implies that after $O(\log N)$ iterations all the parts are marked as good. This further implies that the congestion of $\mathcal{H}'$ is $O(c \log N)$ as the congestion of the union of partial shortcuts is at most the sum of congestion of individual partial shortcuts.

Finally, the number of rounds is at most $O(\log N)$ times the combined number of rounds of the **CoreFast** and **Verification** subroutines, namely $O(\log N \cdot (D \log n + c + bD + bc)) = O(D \log N \log n + bD \log N + bc \log N)$. $\square$

5.3 Warm-up: An $O(D \cdot c)$ Version of the Core Subroutine

In this section, we explain a simpler and deterministic, but slower version of the core subroutine named **CoreSlow** that takes $O(D \cdot c)$ rounds. This is improved in the next section where we present a $O(D \log n + c)$ -round version of the same subroutine.

On a high level, the subroutine takes each part P_i and tries to assign all the T -ancestors of nodes in P_i to its shortcut subgraph. This may, however, lead to large congestions on some edges. We mitigate that issue by declaring an edge **unusable** if more than $2c$ parts try to use it. This ensures the congestion is $2c$. The process provably leads to a constant fraction parts having both small congestion and small block parameter.

Preliminaries: As standard, assume we fixed a spanning tree T of depth D such that G has a T -restricted shortcut with congestion c and block parameter b . During the execution of the algorithm some of the edges will be marked as **unusable**. Furthermore, we say that an tree edge $e \in E_T$ **can see** a node $v \in V$ if v is in the subtree of e and no edge on the unique simple path between the lower endpoint of e and v is unusable. Analogously, an edge can see a part ID i if it can see any node in P_i .

Outline of the CoreSlow subroutine: Initially, no edge is unusable. Process the (tree) edges of T in order of decreasing depth (bottom to top). An edge e is assigned to all the parts P_i such that e can see some node $v \in P_i$, but only if it would be assigned to at most $2c$ such parts. If this is not the case (more than $2c$ shortcut subgraphs would

contain e), we mark this edge e as **unusable** and proceed without assigning any part to it.

Detailed description of the CoreSlow subroutine:

Each node v maintains a list L_v of part IDs that its T -parent edge can see. All the L_v 's are initially empty. The subroutine runs in $\text{depth}(T)$ phases where in the k th phase all the nodes at depth $\text{depth}(T) - k$ update L_v in parallel and send L_v to its T -parent. The update for a node v works by first receiving $L_{v'}$ for all its T -children v' . We assign the union of all received lists and the singleton part ID of v (if any) to L_v . If $|L_v| \leq 2c$, we assign the parent edge of v to all the parts in L_v and transmit L_v to its parent (potentially requiring $2c$ rounds). Otherwise, if $|L_v| > 2c$, we declare the parent edge as unusable.

A direct implementation of this would lead to a subroutine that takes $O(D \cdot c)$ rounds in the CONGEST model. Each of the D levels must propagate at most $2c$ part IDs to their parent nodes. However, this bottleneck can be improved by random sampling, as we show in the next section with the subroutine **CoreFast**. A high-level pseudocode is given in Algorithm 1.

Algorithm 1 CoreSlow

1. At time k each node v at depth $\text{depth}(T) - k$ does the following in parallel:
 - (a) if v is an element of P_i , set $L_v \leftarrow \{i\}$, otherwise $L_v \leftarrow \emptyset$
 - (b) receive all the part IDs from v 's children and assign their union to L'
 - (c) $L_v \leftarrow L_v \cup L'$
 - (d) if $|L_v| > 2c$, mark v 's parent edge as unusable
 - (e) otherwise (serially) send all the part IDs of L_v up to v 's parent node
 2. For each node v :
 - (a) if the parent edge e of v is marked as unusable, e will not be assigned to any shortcut subgraphs
 - (b) otherwise e will be assigned to all $H_i, \forall i \in L_v$
-

LEMMA 7. *Let T be a spanning tree with depth D and assume there exists a T -restricted shortcut with congestion c and block parameter b . The subroutine **CoreSlow** finds a T -restricted shortcut $\mathcal{H}' = (H'_1, H'_2, \dots, H'_N)$ with the following properties:*

- i) *The congestion of $\mathcal{H}'$ is at most $2c$.*
- ii) *There exists a subset of parts $\mathcal{P}' \subseteq \mathcal{P}$ with size at least $|\mathcal{P}'| \geq \frac{N}{2}$ such that the shortcut subgraphs corresponding to parts in $\mathcal{P}'$ have block parameter $3b$.*

The subroutine takes $O(D \cdot c)$ CONGEST rounds to execute. Upon completion, each node knows for each of its incident edges which parts are they assigned to in $\mathcal{H}'$.

PROOF. Let $\mathcal{H} = (H_i)$ be any T -restricted shortcut with congestion c and block parameter b and let $\mathcal{H}' = (H'_i)$ be the shortcut computed by **CoreSlow**. We call $\mathcal{H}$ the **canonical** shortcut and $\mathcal{H}'$ the **computed** shortcut.

By construction, the congestion of $\mathcal{H}'$ is $2c$ as any edge that would be assigned to more than $2c$ shortcut subgraphs is marked as unusable. Hence we proved property *i*).

Let $U \subseteq E_T$ be the set of unusable edges marked by the subroutine. In this paragraph we find an upper bound for $|U|$. Consider **blaming** a part P_i for congesting an unusable edge $e \in U$ when $e \notin G[P_i] + H_i$ and e can see P_i , i.e. edge e was not in the canonical shortcut subgraph H_i , but e was congested by part P_i (and ultimately declared unusable). Each part can be blamed at most b times because each block component can only be blamed for the first unusable edge in his T -tree path towards the T -root. Furthermore, if e is unusable, it takes at least $2c - c$ different block components (from different parts) to be blamed for congesting e . Therefore $|U| \leq N \frac{b}{c}$.

We say that a part P_i **missed** an edge e when $e \in G[P_i] + H_i$ and $e \in U$ (consequently $e \notin H'_i$). Furthermore, call a part **bad** if it missed at least $2b$ edges and **good** otherwise. Note that if a part P_i is good, the block parameter of H'_i is at most $2b + \text{blockParameter}(\mathcal{H}) = 3b$. This is because each missed edge induces a new block component in $\mathcal{H}'$ (more precisely, we can identify each block component of $\mathcal{H}'$ by either an unique block component of $\mathcal{H}$ or an unique missed edge $e \in U$). Consequently, it is sufficient to prove that the subroutine finds at least $\frac{1}{2}N$ good parts.

As any unusable edge is contained in at most c canonical shortcut subgraphs and for a part to be bad we need at least $2b$ edges to be missed, we have that the number of bad parts is at most $|U| \frac{c}{2b} \leq \frac{1}{2}N$. Hence, the subroutine finds at least $\frac{1}{2}N$ good shortcuts.

The number of rounds the subroutine takes is $O(D \cdot c)$: on each of the D depths of the tree T all the nodes in parallel must send the part IDs trying to use its parent edge up the tree. A node can send up to $2c$ IDs, each requiring a round for its transmission. $\square$

5.4 A Faster $O(D \log n + c)$ Version of the Core Subroutine

In this section, we describe a faster version of the core subroutine named **CoreFast**. On a high level, we lower the running time of **CoreSlow** by estimating the number of parts trying to use an edge by random sampling. In particular, each part becomes **active** with probability p and we declare an edge unusable when $\Omega(c \cdot p)$ active parts try to use that edge.

Preliminaries: In addition to the preliminaries of **CoreSlow** we need shared randomness between all the nodes within a part. In other words, all the nodes of the same part must have access to the same seeds for a pseudorandom generator. This can be done by sharing $O(\log^2 n)$ random bits among all the nodes of G in $O(D + \log n)$ rounds, as described in [7].

Outline of the CoreFast subroutine: Each part becomes **active** with probability $p = \frac{\gamma \log n}{2c}$ where $\gamma > 0$ is sufficiently large constant. We basically do the same **CoreSlow** subroutine, but instead of propagating all part IDs of L_v , we propagate only the active ones and declare an edge **unusable** if at least $4c \cdot p = \Omega(\log n)$ (active) part IDs want to use it. Hence by a standard Chernoff bound argument we can claim with high probability that *i*) we never propagate more than $O(\log n)$ part IDs through an edge *ii*) each unusable edge has at least $2c$ part IDs trying to use that edge and *iii*) each usable (non-congested) edge has at most

$8c$ part IDs. After determining which edges are unusable in $O(D \log n)$, **CoreFast** must nevertheless find the complete set of part IDs that can use each edge. This is a tree routing problem where each message (part ID) has to be routed up the tree T until the first unusable edge. No message needs to travel more than D edges and no edge needs to transmit more than $8c$ different part IDs w.h.p. Hence this routing can be done in $O(D + c)$ using Lemma 2.

Detailed description of the CoreSlow subroutine:

Due to shared randomness, each part independently becomes **active** with probability $p = \frac{\gamma \log n}{2c}$ (all the nodes within the part agree on this label). Similarly as in **CoreSlow**, each node v maintains a list $\tilde{L}_v$ of active part IDs that its (T) parent edge can see. All the lists $\tilde{L}_v$ are initially empty. The subroutine runs in $\text{depth}(T)$ phases where in the k th phase all the nodes at depth $\text{depth}(T) - k$ try to update $\tilde{L}_v$ in parallel and send $\tilde{L}_v$ to its T -parent. The update for a node v works by first receiving $\tilde{L}_{v'}$ for all its T -children v' . We assign the union of all received lists and the singleton part ID of v (if any) to $\tilde{L}_v$. If $|\tilde{L}_v| \leq 4c \cdot p$, we assign the parent edge of v to all the parts in $\tilde{L}_v$ and transmit $\tilde{L}_v$ to its parent (requiring $O(\log n)$ rounds). This finalizes the first part of the subroutine where we determine all unusable edges. It remains to forward the complete set of part IDs (and not just the sampled ones) that can use some edge e to the endpoints of e . This is a classic tree routing problem where no route has its length larger than D and no edge intersects more than $8c$ paths w.h.p. Lemma 2 provides a method to route all part IDs in at most $O(D + c)$ rounds. Note that any two part IDs whose routes share an edge have the same endpoint (lowest unusable ancestor edge), so any routing priority between the messages gives the aforementioned $O(D + c)$ bound w.h.p. A high-level pseudocode is given in Algorithm 2

LEMMA 8 (RESTATED LEMMA 5). *Let T be a spanning tree with depth D and assume there exists a T -restricted shortcut with congestion c and block parameter b . The subroutine **CoreFast** finds a T -restricted shortcut $\mathcal{H}' = (H'_i)_{i=1}^N$ with the following properties:*

- i) *The congestion of $\mathcal{H}'$ is at most $8c$ with high probability.*
- ii) *There exists a subset of parts $\mathcal{P}' \subseteq \mathcal{P}$ with size at least $|\mathcal{P}'| \geq \frac{N}{2}$ such that the shortcut subgraphs corresponding to parts in $\mathcal{P}'$ have block parameter $3b$.*

The subroutine takes $O(D \log n + c)$ CONGEST rounds to execute. Upon completion, each node knows for each of its incident edges which parts are they assigned to in $\mathcal{H}'$.

PROOF. This proof extensively utilizes methods used in the proof of Lemma 7. For completeness, we redefine all of the used terminologies and reprove all of the intermediate results.

Let $\mathcal{H} = (H_i)$ be any T -restricted shortcut with congestion c and block parameter b and let $\mathcal{H}' = (H'_i)$ be the shortcut computed by **CoreFast**. We call $\mathcal{H}$ the **canonical** shortcut and $\mathcal{H}'$ the **computed** shortcut.

As $4c \cdot p = \Omega(\log n)$, a standard Chernoff bound argument demonstrates that any edge that is not marked as unusable can see at most $8c$ different part IDs w.h.p. Hence, the congestion of $\mathcal{H}'$ is $8c$ w.h.p.

Algorithm 2 CoreFast

1. Each part becomes active with probability $p = \frac{\gamma \log n}{2c}$
 2. At time k each node v at depth $\text{depth}(T) - k$ does the following in parallel:
 - (a) if v is an element of P_i and P_i is active, set $\tilde{L}_v \leftarrow \{i\}$, otherwise $\tilde{L}_v \leftarrow \emptyset$
 - (b) receive all the active part IDs from v 's children and assign their union to L'
 - (c) $\tilde{L}_v \leftarrow \tilde{L}_v \cup L'$
 - (d) if $|\tilde{L}_v| \geq 4c \cdot p$, mark v 's parent edge as unusable
 - (e) otherwise send all the part IDs $\tilde{L}_v$ up to v 's parent node
 3. Each node v initializes Q_v with its part ID (or $\emptyset$ if not in any part)
 4. Each node v does the following in parallel:
 - (a) add all received IDs to the Q_v
 - (b) if parent edge of v is not unusable and $\exists i \in Q_v$ that was never forwarded
 - i. forward minimum such i along the parent edge
 5. Each part ID in Q_v can use the parent edge of v , unless it is unusable
-

Let $U \subseteq E_T$ be the set of unusable edges marked by the subroutine. In this paragraph we find an upper bound for $|U|$. Consider **blaming** a part P_i for congesting an unusable edge $e \in U$ when $e \notin G[P_i] + H_i$ and e can see P_i , i.e. edge e was not in the canonical shortcut subgraph $G[P_i] + H_i$, but e was congested by part P_i (and ultimately declared unusable). We can similarly argue via a Chernoff bound that each unusable edge $e \in U$ can see at least $2c$ parts, hence we blame at least $2c - \text{congestion}(\mathcal{H}) = c$ parts for congesting e . Each part can be blamed at most b times because each block component can only be blamed for the first unusable edge in his T -tree path towards the T -root. Furthermore, if e is unusable, it takes at least $2c - c$ different block components (from different parts) to be blamed for congesting e . Therefore $|U| \leq N \frac{b}{c}$.

We say that a part P_i **missed** an edge e when $e \in G[P_i] + H_i$ and $e \in U$ (consequently $e \notin H'_i$). Furthermore, call a part **bad** if it missed at least $2b$ edges and **good** otherwise. Note that if a part P_i is good, the block parameter of H'_i is at most $2b + \text{blockParameter}(\mathcal{H}) = 3b$. This is because each missed edge induces a new block component in $\mathcal{H}'$ (more precisely, we can identify each block component of $\mathcal{H}'$ by either an unique block component of $\mathcal{H}$ or an unique missed edge $e \in U$). Consequently, it is sufficient to prove that the subroutine finds at least $\frac{1}{2}N$ good parts.

As any unusable edge is contained in at most c canonical shortcut subgraphs and for a part to be bad we need at least $2b$ edges to be missed, we have that the number of bad parts is at most $|U| \frac{c}{2b} \leq \frac{1}{2}N$. Hence, the subroutine finds at least $\frac{1}{2}N$ good shortcuts.

The number of rounds the subroutine takes is $O(D \log n +$

$c)$: on each of the D depths of the tree T all the nodes in parallel must send the active part IDs that its parent edge can see. If an edge e is not unusable, a Chernoff bound proves that at most $O(c \cdot p) = O(\log n)$ active part IDs can be seen from e , hence the number of rounds for determining unusable edges is $O(D \log n)$ w.h.p.

Propagating the part IDs upwards along T described in Lemma 2 takes $O(D + c)$ rounds, bringing the total number of rounds to $O(D \log n + c)$. $\square$

5.5 Verification Subroutine

In this section, we describe the **Verification** subroutine. Given a tree-restricted shortcut with congestion c and possibly unbounded block parameter, it inspects each of the shortcut subgraphs in parallel and marks the ones that have at most b block components.

The subroutine runs precisely the algorithm described in Lemma 3 which we restate here.

LEMMA 9 (RESTATED LEMMA 3). *Given a T -restricted shortcut with congestion c , a deterministic distributed algorithm can find all parts whose designated shortcut subgraph has at most b' block components. The algorithm executes in $O(b'(D + c))$ rounds.*

The Lemma provides a direct method to implement the formal requirements of the **Verification** subroutine which we restate here for clarity.

LEMMA 10 (RESTATED LEMMA 6). *Given a tree T with depth D and a tentative T -restricted shortcut $\mathcal{H}'$ with congestion c , the deterministic subroutine **Verification** finds all parts $\mathcal{P}' \subseteq \mathcal{P}$ whose designated shortcuts have at most b block components. The subroutine takes $O(b(D + c))$ CONGEST rounds to execute. Upon completion, each node knows whether its part is in the set $\mathcal{P}'$ or not.*

6. SHORTCUT CONSTRUCTION IN THE CASE OF UNKNOWN PARAMETERS

The algorithm presented in Section 5 assume that upper-bound values on b and c are available. That is, each node must know those values before starting the algorithm. Fortunately, in our algorithm, the lack of that knowledge is not a problem. A key property of our construction algorithm is that it inherently includes termination detection, which allows us to use a simple doubling mechanism: We first start the first trial with a small estimated value of parameters, and if we fail the construction, the next trial is executed after doubling the values of parameters. This mechanism removes the requirement of the knowledge on b and/or c with a potential extra log-factor of the running time. It should be noted that utilizing this mechanism can yield much better shortcuts than the theoretical bound. For example, even for graphs with large genus g , the algorithm can find a good (i.e. $o(gD \log D \log N)$ -congestion) shortcut if it luckily exists.

References

- [1] 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. In *Proc. of the Symp. on Theory of Comp. (STOC)*, pages 363–372, 2011.

- [2] M. Elkin. Unconditional lower bounds on the time-approximation tradeoffs for the distributed minimum spanning tree problem. In *Proc. of the Symp. on Theory of Comp. (STOC)*, pages 331–340, 2004.
- [3] M. Elkin. An unconditional lower bound on the time-approximation trade-off for the distributed minimum spanning tree problem. *SIAM Journal on Computing*, 36(2):433–456, 2006.
- [4] S. Frischknecht, S. Holzer, and R. Wattenhofer. Networks cannot compute their diameter in sublinear time. In *Proc. of ACM-SIAM Symp. on Disc. Alg. (SODA)*, pages 1150–1162, 2012.
- [5] J. Garay, S. Kutten, and D. Peleg. A sub-linear time distributed algorithm for minimum-weight spanning trees. In *Proc. of the Symp. on Found. of Comp. Sci. (FOCS)*, 1993.
- [6] M. Ghaffari and B. Haeupler. Distributed algorithms for planar networks I: Planar embedding. Manuscript, 2015.
- [7] M. Ghaffari and B. Haeupler. Distributed algorithms for planar networks II: Low-congestion shortcuts, mst, and min-cut. In *Proc. of ACM-SIAM Symp. on Disc. Alg. (SODA)*, pages 202–219. SIAM, 2016.
- [8] M. Ghaffari, A. Karrenbauer, F. Kuhn, C. Lenzen, and B. Patt-Shamir. Near-optimal distributed maximum flow: Extended abstract. In *the Proc. of the Int’l Symp. on Princ. of Dist. Comp. (PODC)*, pages 81–90, 2015.
- [9] M. Ghaffari and F. Kuhn. Distributed minimum cut approximation. In *Proc. of the Int’l Symp. on Dist. Comp. (DISC)*, pages 1–15, 2013.
- [10] S. Holzer and R. Wattenhofer. Optimal distributed all pairs shortest paths and applications. In *the Proc. of the Int’l Symp. on Princ. of Dist. Comp. (PODC)*, pages 355–364, 2012.
- [11] T. Izumi and R. Wattenhofer. Time lower bounds for distributed distance oracles. In *Proc. of the International Conference on Principles of Distributed Systems*, pages 60–75, 2014.
- [12] M. Khan and G. Pandurangan. A fast distributed approximation algorithm for minimum spanning trees. *Distributed Computing*, 20(6):391–402, 2008.
- [13] S. Kutten and D. Peleg. Fast distributed construction of k -dominating sets and applications. In *the Proc. of the Int’l Symp. on Princ. of Dist. Comp. (PODC)*, pages 238–251, 1995.
- [14] F. T. Leighton, B. M. Maggs, and S. B. Rao. Packet routing and job-shop scheduling in $O(\text{congestion} + \text{dilation})$ steps. *Combinatorica*, 14(2):167–186, 1994.
- [15] C. Lenzen and B. Patt-Shamir. Fast routing table construction using small messages: Extended abstract. In *Proc. of the Symp. on Theory of Comp. (STOC)*, pages 381–390, 2013.
- [16] C. Lenzen and B. Patt-Shamir. Fast partial distance estimation and applications. In *the Proc. of the Int’l Symp. on Princ. of Dist. Comp. (PODC)*, pages 153–162, 2015.
- [17] C. Lenzen and D. Peleg. Efficient distributed source detection with limited bandwidth. In *the Proc. of the Int’l Symp. on Princ. of Dist. Comp. (PODC)*, pages 375–382, 2013.
- [18] D. Nanongkai. Distributed approximation algorithms for weighted shortest paths. In *Proc. of the Symp. on Theory of Comp. (STOC)*, pages 565–573, 2014.
- [19] D. Nanongkai and H.-H. Su. Almost-tight distributed minimum cut algorithms. In *Proc. of the Int’l Symp. on Dist. Comp. (DISC)*, pages 439–453, 2014.
- [20] J. Nešetřil, E. Milková, and H. Nešetřilová. Otakar boruvka on minimum spanning tree problem translation of both the 1926 papers, comments, history. *Discrete Math.*, 233(1):3–36, 2001.
- [21] D. Peleg. *Distributed Computing: A Locality-sensitive Approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.
- [22] D. Peleg and V. Rubinovich. A near-tight lower bound on the time complexity of distributed MST construction. In *Proc. of the Symp. on Found. of Comp. Sci. (FOCS)*, pages 253–, 1999.