

Low-Congestion Shortcuts for Graphs Excluding Dense Minors

Mohsen Ghaffari*

ETH Zurich

Zurich, Switzerland

ghaffari@inf.ethz.ch

Bernhard Haeupler†

Carnegie Mellon University and ETH Zurich

Zurich, Switzerland

haeupler@cs.cmu.edu

ABSTRACT

We prove that any n -node graph G with diameter D admits shortcuts with congestion $O(\delta D \log n)$ and dilation $O(\delta D)$, where δ is the maximum edge-density of any minor of G . Our proof is simple and constructive with a $\tilde{\Theta}(\delta D)$ -round¹ distributed construction algorithm. Our results are tight up to logarithmic factors and generalize, simplify, unify, and strengthen several prior results. For example, for graphs excluding a fixed minor, i.e., graphs with constant δ , only a $\tilde{O}(D^2)$ bound was known based on a very technical proof that relies on the Robertson-Seymour Graph Structure Theorem.

A direct consequence of our result is this: many graph families, including any minor-excluded ones, have near-optimal $\tilde{\Theta}(D)$ -round distributed algorithms for many fundamental communication primitives and optimization problems in the standard synchronous message-passing model with logarithmic message sizes, i.e., the CONGEST model. These problems include minimum spanning tree, minimum cut approximation, and shortest-path approximations.

CCS CONCEPTS

- Theory of computation → Distributed algorithms.

KEYWORDS

distributed graph algorithms, low-congestion shortcuts, minor-excluded graphs, planar graphs, congestion, dilation, minimum spanning tree

ACM Reference Format:

Mohsen Ghaffari and Bernhard Haeupler. 2021. Low-Congestion Shortcuts for Graphs Excluding Dense Minors. In *Proceedings of the 2021 ACM*

*Supported in part by funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 853109), and the Swiss National Foundation (project grant 200021-184735).

†Supported in part by NSF grants CCF-1527110, CCF-1618280, CCF-1814603, CCF-1910588, NSF CAREER award CCF-1750808, a Sloan Research Fellowship, the Swiss National Foundation (project grant 200021-184735), and funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (ERC grant agreement 949272).

¹We use \tilde{O} -notation to suppress polylogarithmic factors in n , e.g., $\tilde{O}(f(n, D, \delta)) = O(f(n, D, \delta) \log^{O(1)} n)$.

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 ACM 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 '21, July 26–30, 2021, Virtual Event, Italy

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8548-0/21/07...\$15.00

<https://doi.org/10.1145/3465084.3467935>

Symposium on Principles of Distributed Computing (PODC '21), July 26–30, 2021, Virtual Event, Italy. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3465084.3467935>

The full version of this paper is available on arXiv, at this link: <https://arxiv.org/abs/2008.03091>.

1 INTRODUCTION AND RELATED WORK

Low-congestion shortcuts are graph-theoretic objects whose quality captures the distributed complexity of a wide range of fundamental graph problems and communication primitives. In this paper, we provide nearly-tight shortcuts for all graphs excluding (dense) minors. Our results significantly strengthen, simplify, generalize and unify several prior results on shortcuts for restricted graph classes. Our results also directly imply simple near-optimal distributed algorithms for a number of well-studied graph problems in excluded minor graphs and many other graph families.

1.1 Background and Definition

Model: We work with the standard synchronous message passing model of computation on networks and distributed systems. The communication network is abstracted as an n -node undirected graph $G = (V, E)$ with $m = |E|$ edges. In each communication round, each node can send an $O(\log n)$ -bit message to each of its neighbors in G . At the beginning, nodes do not know the topology of G . At the end, each node should know its own part of the output, e.g., which of its edges are in the computed *minimum spanning tree*. This model is sometimes referred to as the CONGEST model [29].

The Motivation for Shortcuts: The $\tilde{\Omega}(\sqrt{n})$ -round lower bound for global distributed graph problems is well-known by now. Concretely, it is known that on worst-case general graphs, for a strikingly wide-range of global graph problems — including minimum spanning tree, minimum cut, maximum flow, single-source shortest path, etc — any distributed algorithm needs a round complexity of $\tilde{\Omega}(\sqrt{n})$ [3, 30, 33]. This holds even for any (non-trivial) approximation of these problems, and even on graphs with a small, e.g., logarithmic, diameter. On the other hand, the lower bound graphs are carefully crafted pathological topologies which do not occur in practice. Indeed, many real-world networks have small (polylogarithmic) diameters (e.g., the Facebook graph with billions of nodes has average distance below 5 and a diameter of around 50) and seem to allow for exponentially faster optimization algorithms with $\tilde{O}(D)$ -round complexities. Low-congestion shortcuts [7] were introduced as a graph-theoretic notion to capture and exploit this phenomenon and allow a more fine-grained study of the complexity of global problems and how this complexity relates to the structure of the network topology. In various graph families, algorithms

based on low-congestion shortcuts have obtained complexities that are considerably below the $\tilde{\Omega}(\sqrt{n})$ barrier— in fact, often obtaining near-optimal $\tilde{O}(D)$ -round algorithms².

The Definition of Shortcuts: Suppose that the set V of vertices is partitioned into disjoint subsets V_1, V_2, \dots, V_k , known as *parts*, such that the subgraph $G[V_i]$ induced by each part V_i is connected. We call a collection of subgraphs H_1, H_2, \dots, H_k a *shortcut with congestion c and dilation d* if we have the following two properties: (A) $\forall i \in [1, k]$, the diameter of the subgraph $G[V_i] + H_i$ is $O(d)$, and (B) each edge $e \in E$ is in $O(c)$ many subgraphs H_i .

Application of Shortcuts: Shortcuts organically lead to fast distributed algorithms. For instance, let us consider the minimum spanning tree problem. Suppose that we are in a graph family for which, we have a distributed algorithm that can compute a shortcut with congestion c and dilation d in T rounds (upon receiving the partition V_1, V_2, \dots, V_k). In general, we refer to $c + d$ as the *quality of the shortcut*. Then, we can use this algorithm to obtain a distributed algorithm for the minimum spanning tree problem, with a round complexity $\tilde{O}(c+d+T)$. This follows directly from Boruvka's 1926 approach [28]. A number of other graph problems can also be solved distributedly using shortcuts, with a similar complexity. This includes min-cut computation [7], single-source shortest-paths approximation [16], and many more [2, 10, 21].

Of course, the question that remains is this:

What is the existential shortcut quality $c + d$ of important graph families, and what is the corresponding construction time T ?

Graph Minors and Minor Density: Before proceeding to known results and our contribution, let us briefly recall the definition of minors and their density. A graph H is a minor of graph G if H can be obtained from G by contracting edges and deleting edges and vertices. Equivalently $H = (V', E')$ is a minor of graph $G = (V, E)$ if there is a mapping $\text{map}_{H,G}$ from vertices in H to disjoint connected subsets of vertices in G , each inducing a connected subgraph, such that for every edge $\{u', v'\} \in E'$ there exists a $\{u, v\} \in E$ with $u \in \text{map}_{H,G}(u')$ and $v \in \text{map}_{H,G}(v')$.

An important parameter throughout this paper is the *minor density* $\delta(G)$ of a graph G which is defined as:

$$\delta(G) = \max \left\{ \frac{|E'|}{|V'|} \mid H = (V', E') \text{ is a minor of } G \right\}.$$

It is known that the minor density $\delta(G)$ is (up to a small polylogarithmic factor) the same as the size of the largest complete minor in G , i.e., its *complete-graph minor size* $r(G) = \max\{r \mid K_r \text{ is a minor of } G\}$.

LEMMA 1.1 ([34]). $\forall G : \frac{r(G)-1}{2} \leq \delta(G) \leq 8r(G)\sqrt{\log_2 r(G)}$, i.e., $\delta(G) = \tilde{\Theta}(r(G))$.

Note that if a graph family \mathcal{G} is closed under taking minors or equivalently excludes a fixed minor H of size s then every $G \in \mathcal{G}$

²One can draw parallels between the role that shortcuts have turned out to play for distributed algorithms of global graph problems with the role that *separators* play as a key algorithmic tool in sequential algorithms for various graph families, such as planar [23, 26], bounded genus [11], and minor-excluded [1].

has $r(G) \leq s$ and therefore also a constant minor density $\delta(G) = O(s\sqrt{\log s})$.

1.2 Our Contribution

We first briefly summarize known results on shortcuts (see Section 1.3 for further details). Ghaffari and Haeupler [6, 7] provided shortcuts of quality $\tilde{O}(D)$ and construction time $\tilde{O}(D)$ for planar graphs. This was later extended to graphs with bounded genus, bounded treewidth, and bounded pathwidth [15], and with improved construction algorithms [12, 14]. Haeupler, Li, and Zuzic [17] gave shortcuts for excluded minor graphs, with quality $\tilde{O}(D^2)$, using elaborate arguments building on the Graph Structure Theorem of Robertson and Seymour [31, 32]. While excluded minor graphs are vastly more general, encompassing all previous graph classes for constant bounds on the above graph parameters, the question whether near-optimal $\tilde{O}(D)$ shortcuts and optimization algorithms for excluded minor topologies are possible remained open.

Existential Results for Graph with Minor Density δ : In this work, we resolve this question in the positive. We also significantly strengthen, simplify, generalize, and unify all the above results by giving a near-optimal existential shortcut guarantee for *any* graph G , which depends only on its minor density $\delta(G)$:

THEOREM 1.2. *Any n -node graph G with diameter D and minor density $\delta(G) = \delta$ admits shortcuts with dilation $O(\delta D)$ and congestion $O(\delta D \log n)$.*

Besides the important quadratic quantitative improvement from $\tilde{O}(D^2)$ to a near-optimal $\tilde{O}(D)$ for excluded minor graphs, i.e., graphs with constant δ , our proof is significantly simpler than that of [17]. Instead of the technical proof in [17] which uses the powerful Graph Structure Theorem [31, 32], we give a short and elementary proof. Our proof even provides small concrete constants³ whereas the Robertson-Seymour Graph Structure Theorem is known to often lead to tower-type dependencies on the size k of the excluded minor [24] and “galactic algorithms” [22].

Even more importantly, our result applies to *any* graph and graph family even if $\delta(G)$ is large or growing. For example it implies that graphs with sub-polynomial minor density or expansion (see [27] for definitions and treatment of such more inclusive graph families) have sub-polynomial shortcuts and fast $O(n^{o(1)})$ -round optimization algorithms, which still drastically improve over the $\tilde{\Omega}(\sqrt{n})$ -lower bound. Our results are the first that apply to graph classes strictly more general than minor-closed or excluded-minor graph families.

We complement Theorem 1.2 with a matching lower bound which shows that the linear dependency of Theorem 1.2 on the minor-density $\delta(G)$ is necessary and optimal:

LEMMA 1.3. *For any δ, D there is a graph G with diameter D and $\delta(G) \leq \delta$ and a collection of parts for which the quality of the best shortcut in G is $\Omega(\delta D)$.*

Lastly, having the minor density $\delta(G)$ as a parameter in Theorem 1.2 also directly implies often optimal results on shortcuts for

³No attempt was made to optimize the explicit constants in the 8δ and $8\delta D$ bounds of our main result Theorem 3.1, or any other explicit constants in this paper. Somewhat better constants are likely possible.

other graph parameters in a simple and uniform way. For example, [7] showed that any genus- g network admits shortcuts of quality $\tilde{O}(gD)$. In [15] this dependency of the shortcut quality on the genus was improved to an optimal bound of $\tilde{O}(\sqrt{g}D)$ via sophisticated topological arguments including cutting a genus- g graph along fundamental cycles of a shortest-path tree and cleverly combining planar shortcuts for pieces of parts that are cut by these cycles. We obtain the same result as a trivial corollary of Theorem 1.2, given that $\delta(G) = O(\sqrt{g})$ for any genus- g graph.

COROLLARY 1.4. *Any genus- g graph with n nodes and diameter D admits shortcuts with congestion $O(\sqrt{g}D \log n)$ and dilation $O(\sqrt{g}D)$.*

Bounds for other graph parameters follow similarly as simple corollaries of our main theorem. This includes tight $\tilde{O}(kD)$ shortcuts for k -pathwidth and k -treewidth graphs, matching the results in [15] (since $\delta(G) = O(k)$ for such graphs). In contrast to prior works like [15], we do not require a completely different proof specific to the graph parameter at hand to obtain these bounds.

Distributed Construction and Applications: In order to be algorithmically useful we also need fast distributed constructions for the new existential shortcut guarantees. We achieve this by proving strong additional structural guarantees for our shortcuts, namely tree-restrictedness and a small block number. Haeupler, Hershkovitz, Izumi, and Zuzic [12, 14] showed that these structural guarantees are strong enough for a simple uniform shortcut constructions to find a $\tilde{O}(Q)$ -quality shortcut in only $\tilde{O}(Q)$ rounds whenever a quality- Q shortcut with such structure exists.

THEOREM 1.5. *There exists a randomized distributed algorithm which, for any n -node m -edge graph G with diameter D and minor density δ , computes a shortcut of quality $\tilde{O}(\delta D)$ in $\tilde{O}(\delta D)$ rounds with high probability. There also exists a $\tilde{O}(\delta^2 D)$ -round deterministic algorithm. Both algorithms use only $\tilde{O}(m)$ messages.*

This, together with all the algorithms that are built on top of the low-congestion shortcut framework, shows that a wide range of fundamental communication primitives and global graph problems – e.g., minimum spanning tree, min-cut, shortest path approximation, etc – can be solved in $\tilde{O}(D\delta)$ rounds, in graphs that do not have a minor of density δ . This vastly widens the range of graph families for which we now know the correct round complexity for basic global graph problems (up to logarithmic factors). We state just two such corollaries as examples.

COROLLARY 1.6. *There is a distributed algorithm that, for any n -node m -edge graph G with diameter D and minor density δ , computes a minimum spanning tree in $\tilde{O}(\delta D)$ rounds with high probability (or $\tilde{O}(\delta^2 D)$ rounds deterministically) using $\tilde{O}(m)$ messages.*

COROLLARY 1.7. *There is a randomized distributed algorithms that, for any n -node m -edge graph G with diameter D and minor density δ , computes an exact minimum cut in G in $\tilde{O}(\delta^{O(1)} D)$ rounds with high probability using $\tilde{O}(m)$ messages.*

Similar results can be obtained for sub-graph connectivity, single-source shortest paths approximations [16], as well as several fundamental communication primitives like multiple unicasts or partwise-aggregation (see Section 2).

1.3 Prior Work on Shortcuts

We next overview the known results about the existential quality and construction time of shortcuts.

Let us start with general graphs. It is easy to see that any n -node graph G , whose diameter is at most D , admits shortcuts of quality $D + \sqrt{n}$: Let T be a BFS of the graph G . Define $H_i = \emptyset$ for any each part with $|V_i| \leq \sqrt{n}$ and $H_i = T$ for any other part. Moreover, this $D + \sqrt{n}$ bound is nearly-optimal [7] in general graphs. This $D + \sqrt{n}$ bound is implicitly the underlying reason for the seminal $\tilde{O}(D + \sqrt{n})$ round minimum spanning tree algorithm of Kutten and Peleg [4, 19], although there are more fine-grained aspects there to avoid extra logarithmic factors. Moreover, this $\tilde{O}(D + \sqrt{n})$ round complexity is nearly optimal for solving MST in general graphs, as mentioned before, because of the carefully crafted lower bound graphs of [3, 30, 33].

Ghaffari and Haeupler [7] showed that any planar graph with diameter at most D admits shortcuts of quality $O(D \log D)$, and that this bound is nearly tight, almost matching a lower bound of $\Omega(D \frac{\log D}{\log \log D})$. They also provided a distributed algorithm for constructing such shortcuts in $\tilde{O}(D)$ rounds, assuming a planar embedding of the graph. Combined with the distributed planar embedding algorithm provided by [6], this led to an $\tilde{O}(D)$ round distributed algorithm for MST, and some other graph problems, thus exhibiting a family of graphs in which one can go below the notorious $\tilde{O}(D + \sqrt{n})$ lower bound [33]. They [7] also showed that graphs embeddable on a surface of genus g admit shortcuts of quality $O(gD \log D)$.

Haeupler, Izumi, and Zuzic [14] showed that one can construct shortcuts with a quality similar to those of [7], even without a planar embedding. Haeupler, Izumi, and Zuzic [15] showed that graphs of treewidth or pathwidth k admit shortcuts of quality $\tilde{O}(kD)$, and they also improved the genus dependency to $\tilde{O}(\sqrt{g}D)$. Haeupler, Hershkovitz, and Wajc [12] improved and extended the results of [14] by showing that one can obtain algorithms that are near-optimal in both time and message complexity, using shortcuts.

Generalizing the span of shortcuts much further, Haeupler, Li, and Zuzic [17] gave shortcuts for excluded minor graphs. Concretely, they showed that any graph that does not have a clique of constant size as a minor admits a shortcut of quality $\tilde{O}(D^2)$. This result is fairly involved and it builds on the Graph Structure Theorem of Robertson and Seymour [31, 32]. Using the approach of [12, 14], all these shortcuts mentioned above [7, 14, 15, 17] can be constructed in a round complexity matching their quality, up to logarithmic factors.

Ghaffari, Kuhn, and Su [8] provided shortcut constructions for well-connected graphs. In particular, they showed that any graph where the lazy random walk has mixing time T_{mix} admits shortcuts of quality $T_{mix} \cdot \text{poly}(\log n)$ and they presented distributed algorithms for constructing shortcuts of quality $T_{mix} \cdot 2^{O(\sqrt{\log n \log \log n})}$ in $T_{mix} \cdot 2^{O(\sqrt{\log n \log \log n})}$ rounds. These quality and construction time bounds were both improved later to $T_{mix} \cdot 2^{O(\sqrt{\log n})}$, by Ghaffari and Li [9].

Finally, Kitamura et al. [18] showed that any k -chordal graph admits shortcuts of quality $O(kD)$, which can be determined in even $O(1)$ rounds, and that graphs of diameter 4 and 3 admit shortcuts

of quality and construction time $\tilde{O}(n^{1/4})$ and $\tilde{O}(n^{1/3})$ respectively. These essentially match the respective lower bounds [25, 33].

2 PRELIMINARIES: DEFINITIONS OF SHORTCUTS

In many distributed algorithms for global graph problems, the problem boils down to the following natural part-wise aggregation task. For instance, this exactly captures the problem of identifying minimum-weight outgoing edges in Boruvka's classic approach to MST [28].

DEFINITION 2.1. (The Part-wise Aggregation Problem) Consider a network graph $G = (V, E)$ and suppose that the vertices are partitioned into disjoint parts P_1, P_2, \dots, P_k such that the subgraph induced by each part is connected. In the part-wise aggregation problem, the input is a value x_v for each node $v \in V$. The output is that each node $u \in P_i$ should learn an aggregate function of the values held by vertices in its part P_i , e.g., $\sum_{v \in P_i} x_v$, $\min_{v \in P_i} x_v$, or $\max_{v \in P_i} x_v$. Alternatively, exactly one node in each part has a message and it should be delivered to all nodes of the part.

Typically, if we can solve the part-wise aggregation problem in a network in T time, we can obtain algorithms for various fundamental graph problems with a round complexity of $\tilde{O}(D + T)$ [7]. Thus, we want fast algorithms for part-wise aggregation.

The point to emphasize in the part-wise aggregation problem is that, it is possible that the diameter of the subgraph induced by each part P_i is quite large, much larger than the diameter of the base graph G . For instance, the former can be up to $\Theta(n)$ while the latter is just 2 (considering a wheel graph with one part being all nodes except the center). Hence, to obtain fast algorithms for part-wise aggregation, we would like to allow some parts to use edges of G which are outside the part. Of course, we should limit the number of parts that try to use each single given edge, as that would cause congestion and would slow down the solution. This naturally brings us to the concept of *low-congestion shortcuts*, as defined in [7].

DEFINITION 2.2. (Shortcuts) Given a part-wise aggregation problem – i.e., a graph $G = (V, E)$ where vertices are partitioned in disjoint parts P_1, P_2, \dots, P_k , each of which induces a connected subgraph – we call a collection of subgraphs H_1, H_2, \dots, H_k a shortcut with *congestion* c and *dilation* d if we have the following two properties: (I) for each $i \in [1, k]$, the diameter of the subgraph $G[P_i] + H_i$ is at most $O(d)$, and (II) each edge e is in at most $O(c)$ many of the subgraphs H_i . We refer to $Q = c + d$ as the *quality* of the shortcut.

Given a c -congestion d -dilation shortcut for a part-wise aggregation problem, we can solve the part-wise aggregation problem in $O(c + d \log n) = \tilde{O}(Q)$ rounds, using the *random delays* technique [5, 13, 20]. This makes the shortcut quality Q the dominant parameter which determines the round complexity of shortcut-based algorithms (modulo the shortcut's construction time).

Tree-Restricted Shortcuts, and their Block-Number: In many graph families [7, 15, 17], shortcuts can be chosen to come from one low-depth tree of the original graph, which provides a particularly clean and simple structure which can be utilized for efficient shortcut constructions. In particular, one can fix any rooted breadth

first search tree T – or any other low-depth spanning tree – and restrict each H_i to only include edges from T . Note that only edges with descendants⁴ in P_i are useful for short-cutting P_i . A particular simple way to construct H_i is to specify a set S of b edges in T and define H_i to be all edges with descendants in the forest $T \setminus S$. We say such an H_i has *block number* at most b since the graph $(P_i \cup V(H_i), H_i)$ has at most b connected components (of diameter $O(D)$). More generally we define tree-restricted shortcuts and their block number as follows:

DEFINITION 2.3. (Tree-Restricted Shortcuts and Block Number) Consider a part-wise aggregation setup in graph G with diameter D , parts P_1, \dots, P_k , a rooted tree T of G with depth D . We say the shortcut H_1, \dots, H_k is *tree-restricted* or *T-restricted* if all its edges are in T , i.e., if $\bigcup_i H_i \subseteq T$. Moreover, for any part P_i , we call the connected components of the graph $(P_i \cup V(H_i), H_i)$ the *blocks* of P_i . The *block number* b of a shortcut is the maximum block number of any part.

If a topology has shortcuts for any collection of parts and any choice of the tree T then we say it admits good shortcuts. We also introduce the concept of partial shortcuts which lead to slightly tighter bounds and simpler proofs:

DEFINITION 2.4. (Admitting Shortcuts) We say a topology G with diameter D admits *tree-restricted c -congestion b -block shortcuts*, if for any tree T with depth at most D and for any collection of node-disjoint connected parts P_1, \dots, P_k , there exists a T -restricted c -congestion b -block shortcut.

DEFINITION 2.5. (Admitting Partial Shortcuts) We say a topology G with diameter D admits *tree-restricted c -congestion b -block partial shortcuts*, if for any tree T with depth at most D and for any collection of node-disjoint connected parts P_1, \dots, P_k , there are at least $k/2$ of the parts with a T -restricted c -congestion b -block shortcut.

It is easy to see that a small block number directly implies a small dilation and that admitting partial shortcuts is essentially the same as admitting shortcuts – up to a $\tilde{O}(1)$ factor in the congestion.

OBSERVATION 2.6. Any b -block T -restricted shortcut in a graph with diameter D has dilation at most $b(2D + 1)$.

PROOF. For each part P_i , the graph $(P_i \cup V(H_i), H_i) \subseteq T$ is a forest with at most b connected components, each of them has a diameter of at most twice the depth of T which is D . These components are connected via edges of P_i , because P_i induces a connected subgraph. Hence, the diameter of $G[P_i] + H_i$, i.e., the dilation for this part, is at most $b(2D + 1)$. \square

OBSERVATION 2.7. Any n -node graph G that admits tree-restricted c -congestion b -block partial shortcuts also admits tree-restricted $c \log_2 n$ -congestion b -block shortcuts.

PROOF. For a collection of k parts and a tree T consider $\log_2 k$ iterations in which one takes a T -restricted c -congestion b -block partial shortcut for any still remaining parts. Given that such a partial shortcut defines sets of shortcut edges for at least half of the remaining parts, all parts will have a set of shortcut edges from T

⁴Here, as standard, we call node v a descendant of an edge $e \in T$ if the shortest path connecting v to the root of the tree passes through edge e .

in the end. Taking this as a full T -restricted shortcut might lead to a $(c \log_2 k)$ -congestion but leaves the block number of b for every part unaffected. \square

Efficient Shortcut Constructions for Tree-restricted Shortcuts

Efficient Shortcut Constructions for Tree-restricted Shortcuts: The main reason for Haeupler, Izumi and Zuzic [14] to introduce the concept of b -block tree-restricted shortcuts is that this additional structure can be used to obtain a very simple and efficient shortcut construction. This construction was further improved by Haeupler, Hershkowitz, and Wajc [12] to give a slightly faster, deterministic, and message optimal construction:

LEMMA 2.8 ([12, 14]). *There exists a simple distributed algorithm which, for any n -node m -edge D -diameter graph G which admits c -congestion b -block partial shortcuts with quality $Q = c + bD$, computes a quality- $\tilde{O}(Q)$ shortcut for any given collection of parts in $\tilde{O}(Q)$ rounds with high probability using $\tilde{O}(m)$ message (or deterministically in $\tilde{O}(bQ)$ rounds).*

By proving that the shortcuts from Theorem 1.2 can be chosen to be tree-restricted with a small $\delta(G)$ block number we get an efficient construction algorithm, and therefore Theorem 1.5, “for free”. We remark that proving this additional tree-restriction structure and thus having a fast construction algorithm is crucial for the algorithmic usability of shortcuts and generally quite hard. Indeed, efficiently constructing good general shortcuts for all graph families that admit them remains a major open problem. Even in the special case of well-connected graphs (with a small mixing time, e.g., expander or random graphs), for which shortcuts generally cannot be chosen to be tree-restricted, there is currently a $2^{O(\sqrt{\log n})}$ gap between the construction time and the shortcut quality [8, 9].

3 SHORTCUTS FOR GRAPHS WITH MINOR DENSITY δ

3.1 Main Result

We prove the following main result, which directly implies Theorem 1.2 and Theorem 1.5.

THEOREM 3.1. *Every G with diameter D and minor density $\delta = \delta(G)$ admits tree-restricted $8\delta D$ -congestion 8δ -block partial shortcuts.*

Indeed, using Observation 2.7 and Observation 2.6 the existence of a $8\delta D$ -congestion 8δ -block partial shortcut directly implies the existence of an $(8\delta D \log_2 n)$ -congestion $(8\delta(2D+1))$ -dilation shortcut and therefore Theorem 1.2. The constructive main theorem Theorem 1.5 directly follows from using Lemma 2.8 on the tree-restricted shortcuts of Theorem 3.1.

The general idea to prove Theorem 3.1 is to “run” the shortcut construction algorithm from [14] and prove that if it fails to find a sufficiently good tree-restricted shortcut, then G contains a minor with density exceeding $\delta(G)$.

PROOF OF THEOREM 3.1. Let T be any rooted spanning tree in G of depth at most D . Let $\mathcal{P} = \{P_1, \dots, P_k\}$ be a collection of connected node-disjoint parts. We set our desired congestion to be $c = 8\delta D$. For a tree edge $e \in T$, let v_e be the endpoint of e that is further away from the root.

Defining overcongested edges: Initially, let $O = \emptyset$. We process tree edges in order of decreasing depths, level by level. For any edge $e \in T$, let $I_e \subseteq \mathcal{P}$ be the parts that have a non-empty intersection with the descendants of v_e in $T \setminus O$. If $|I_e| \geq c$ we say e is **overcongested** and we add e to O .

The bipartite graph B : We define the bipartite graph $B = (O \cup \mathcal{P}, E')$ whose node set consists of edge-nodes corresponding to overcongested edges on the one side and part-nodes corresponding to the parts from \mathcal{P} on the other side. The edges $E' = \{(e, P_i) \mid e \in O, P_i \in I_e\} \subseteq O \times \mathcal{P}$ of B indicate which part contributed to which edge being overcongested. We associate every edge $(e, P_i) \in B$ with some representative node $r_{(e, P_i)} \in P_i$, that is a descendant of v_e and can be reached from v_e via $T \setminus O$. A schematic illustration of this setup is given in Figure 1. Let $R_i \subseteq P_i$ be the set of representative nodes in P_i . Note that $|R_i|$ is equal to the degree of node $P_i \in B$. The degree of any edge-node $e \in B$ is $|I_e| \geq c$, as we only have overcongested edges represented in B .

Next, we argue that one of the following two cases applies: either (I) there exists a good partial shortcut, or (II) graph G has a minor of density exceeding δ – contradicting the assumption that $\delta = \delta(G)$.

(I) Either there exists a good partial shortcut: If at least half of all parts have a degree of at most 8δ in B , then defining the shortcut H_i for any such part P_i to be all ancestor edges of P_i in the forest $T \setminus O$ identifies a c -congestion 8δ -block partial shortcut.

(II) Or there is a dense minor $B_{\mathcal{P}'}$ in G : If we are not in case (I), then at least half of all parts in B have degree at least 8δ in B . In this case, the average degree among part-nodes in B is at least 4δ , because at least half of the parts have degree at least 8δ . Moreover, the average degree among edge-nodes (and in fact even their minimum degree) in B is at least c .

Let \mathcal{P}' be a random subset of \mathcal{P} in which each part is included independently at random with probability $\frac{1}{4D}$. We define a subgraph $B_{\mathcal{P}'} = (V_{\mathcal{P}'}, E_{\mathcal{P}'})$ of the bipartite graph B which is also a minor of G , as follows. The part-nodes in $B_{\mathcal{P}'}$ are exactly the $P_i \in \mathcal{P}'$ and in the minor-mapping $\text{map}_{B_{\mathcal{P}'}, G}$ such a node is mapped to the vertex set of P_i . All edges $e \in O$ with $v_e \notin \bigcup_{P_i \in \mathcal{P}'} P_i$ are edge-nodes in $B_{\mathcal{P}'}$. The vertex set $\text{map}_{B_{\mathcal{P}'}, G}(e)$ in G of such an edge-node $e \in B_{\mathcal{P}'}$ is exactly the vertices in the connected component containing v_e in the forest $(T \setminus O) \setminus (\bigcup_{P_i \in \mathcal{P}'} P_i)$.

To define which edges $(e, P_i) \in B$ are in $B_{\mathcal{P}'}$, we say (e, P_i) is **potentially present** if the tree path between v_e and the representative $r_{(e, P_i)} \in P_i$, including the deeper endpoint v_e but excluding the representative node $r_{(e, P_i)} \in P_i$, does not contain any node from $\bigcup_{P_j \in \mathcal{P}'} P_j$. We say edge (e, P_i) is **actually present** and add (e, P_i) to $B_{\mathcal{P}'}$ if it is both potentially present and $P_i \in \mathcal{P}'$. Note that $B_{\mathcal{P}'}$ is indeed a minor of G under the mapping function $\text{map}_{B_{\mathcal{P}'}, G}$ since the vertex sets corresponding to nodes in $B_{\mathcal{P}'}$ are disjoint and connected in G and edges in $B_{\mathcal{P}'}$ are a subset of the edges produced when contracting these vertex sets in G .

Density of the graph $B_{\mathcal{P}'} = (V_{\mathcal{P}'}, E_{\mathcal{P}'})$: Let k be the number of edges in B . Every edge $(e, P_i) \in B$ has a probability of at least $1 - (1 - \frac{1}{4D})^D \geq \frac{3}{4}$ to be potentially present. This probability is independent from the $\frac{1}{4D}$ probability for P_i to be in \mathcal{P}' . Hence, $\mathbb{E}[|E_{\mathcal{P}'}|] \geq \frac{3k}{16D}$.

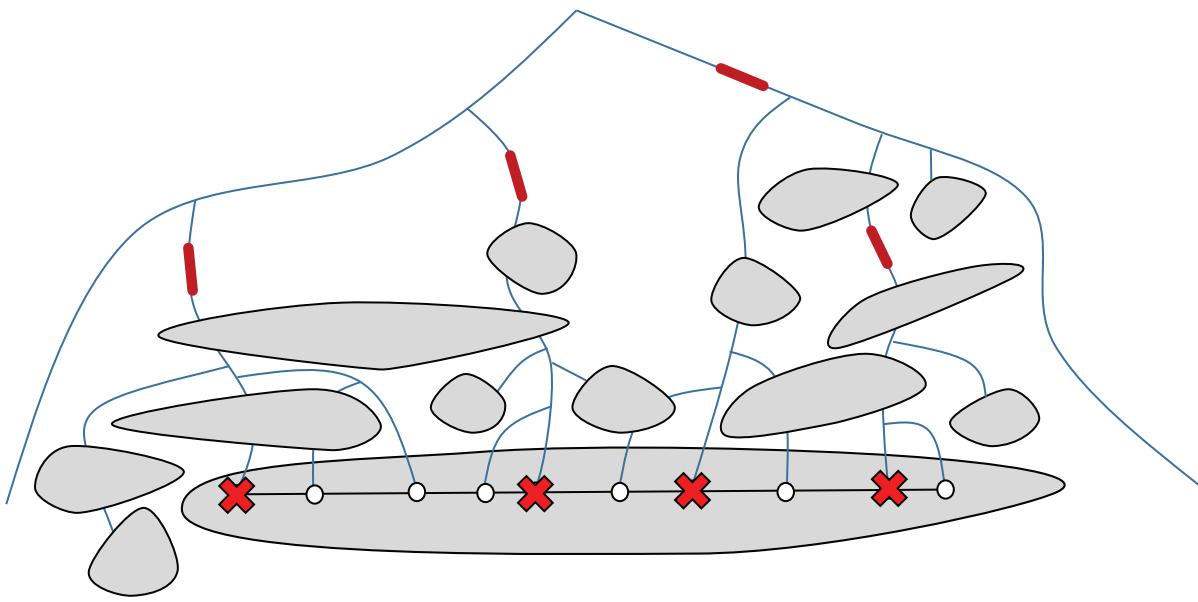


Figure 1: A schematic illustration of a tree T (indicated in blue), the one part that we are focusing on (indicated as a gray area with white nodes in it) along with a few others (indicated just as gray areas), and the overcongested edges (indicated in red) that have descendants in this part, which we call representatives (indicates with red cross marks). When the blue tree goes through a part (a gray area), we indicate that this section of the tree has some vertex of that part.

The nodes in $B_{P'}$ on the other hand consist of (A) at most some $\frac{k}{c}$ edge-nodes, given that they have degree at least c in B , each of which is included in $B_{P'}$ with probability $1 - \frac{1}{4D}$ and (B) at most some $\frac{k}{4\delta}$ part-nodes, with average degree more than 4δ in B , each of which is included in $B_{P'}$ with probability $\frac{1}{4D}$. Hence,

$$\mathbb{E}[|V_{P'}|] < \frac{k}{8\delta D} + \frac{k}{4\delta} \cdot \frac{1}{4D} = \frac{3k}{16D} \frac{1}{\delta}.$$

Therefore, by linearity of expectation, we can conclude that

$$\mathbb{E}[|E_{P'}| - \delta|V_{P'}|] = \mathbb{E}[|E_{P'}|] - \delta\mathbb{E}[|V_{P'}|] > \frac{3k}{16D} - \delta \frac{3k}{16D} \frac{1}{\delta} = 0,$$

which implies that $\Pr[|E_{P'}| - \delta|V_{P'}| > 0] > 0$. That is, with a non-zero probability⁵, the minor $B_{P'}$ in G has density exceeding δ , giving the desired contradiction. \square

We remark that the above proof of Theorem 3.1 can easily be made constructive directly. A trivial implementation would lead to a deterministic $O(\delta^2 D^2)$ -round algorithm. Using the sampling idea from [12, 14] to identify overcongested-edges one can speed this up to $\tilde{O}(\delta D)$. Overall, this would save a $\Theta(\log n)$ factor in the quality of the computed (partial) shortcut. One could also make the algorithm certifying, i.e., output a dense minor if a (partial) shortcut of desired quality cannot be found. For example, one can obtain an algorithm which when run on a graph G with tree T of depth D and

⁵With a slightly more careful argument, we can show that there is $\Omega(1/D)$ probability to find a minor of density exceeding δ , but for our existence proof, just a positive probability suffices.

a collection of parts terminates in $\tilde{O}(\delta D)$ rounds for some $\delta \leq \delta(G)$ and outputs both an 8δ -block $8\delta D$ -congestion partial shortcut and a $(\delta - 1)$ -dense bipartite minor, which explains/certifies why no better shortcut was found.

3.2 Optimality of the Main Result

Next we prove that our main result Theorem 3.1 is existentially optimal up to only small explicitly given *constant* factors in the congestion and block number of partial shortcuts. This directly implies the slightly weaker Lemma 1.3 for full shortcuts presented in Section 1.2.

LEMMA 3.2. *For every $\delta', D' \in \mathbb{N}$ with $5 \leq \delta' \leq D'/2$ there exists a topology G (with $O(\delta D)$ nodes) and a set of node-disjoint paths such that:*

- G has diameter D' and every minor of G has density less than δ' and
- the best partial shortcut quality for the set of paths in G is at least $\frac{(\delta'-3)D'}{6} = \Theta(\delta'D')$, i.e., any partial shortcut has either congestion at least $\Omega(\delta'D')$ or dilation at least $\Omega(\delta'D')$.

PROOF. Let $\delta = \delta' - 2$, $k = \lfloor \frac{D'}{2\delta} \rfloor$, $D = k\delta$. Note that $k \geq 2$, $\delta' \geq 3$, $D \in [6, \lfloor \frac{D'}{2} \rfloor]$. The topology $G = (V, E)$ is made of one special path of length $(\delta - 1)k + 1$ at the top, along with $(\delta - 1)D + 1$ many paths of length $(\delta - 1)D + 1$ at the bottom, known as rows. In every D^{th} column, every D^{th} row is connected to a node in the top path, such

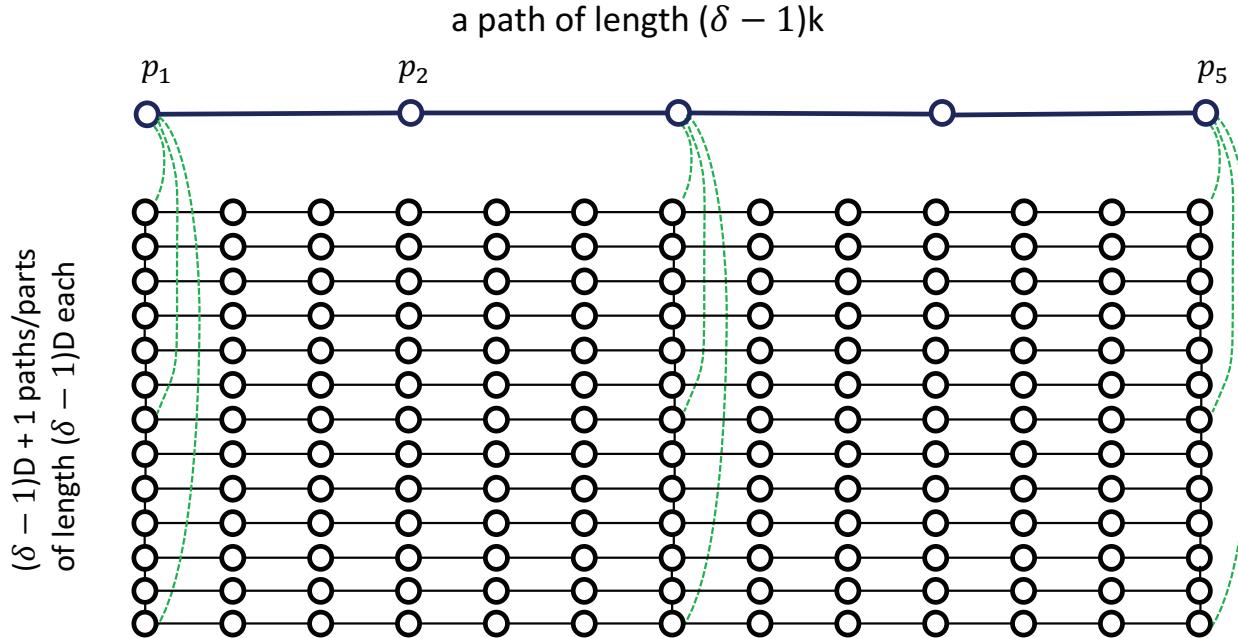


Figure 2: The lower bound graph for Lemma 3.2 with $\delta' = 3$, $k = 2$ and $D' = 6$. The special path of length at most D is at the top, and below are the $\Theta(\delta D)$ paths/parts of length $\Theta(\delta D)$. Every D steps on these parts there is a column connecting all paths and (green) connections every D steps on this column to the special path on top.

that all these nodes of the same column are connected to one node in the top node. See Figure 2 for an illustration. More formally, the graph is defined as follows:

- $V = \{p_i \mid i \in [(\delta-1)k+1]\} \cup \{v_{i,j} \mid i, j \in [(\delta-1)D+1]\}$
- The edges E are such that the p -nodes form a path of length $(\delta-1)k$. Moreover, for any $i \in [(\delta-1)D+1]$, the $v_{i,*}$ nodes form a path P_i of length $(\delta-1)D$. Also, for any $j \in [\delta]$, the nodes $v_{*,(j-1)D+1}$ form a path of length $(\delta-1)D$ of which every D^{th} node connects to $p_{(j-1)k+1}$, that is, $\{v_{(j'-1)D+1, (j-1)D+1}, p_{(j-1)k+1}\} \in E$ for every $j, j' \in [\delta]$.

We first argue that the graph has diameter at most D' . From every v node one can reach a p -node by going at most $\frac{D}{2}$ steps to the closest node in its P -path which is in the same column as of a p -node, then going at most $\frac{D}{2}$ steps up or down to a node that has a p -node neighbor, and then doing one more step to that p -node. From any p -node one can reach $p_{\frac{D}{2}}$ in at most $\frac{D}{2}$ steps. Overall the diameter of G is therefore at most $1.5D + 1 \leq D'$.

Now we argue that G has no minor with density δ' . Notice that G is planar after deleting the $\delta(\delta-1)$ edges between the p nodes and any P_i -path except P_1 . Any minor of G with $s \geq \delta+1$ nodes therefore has, according to Euler's formula, at most $3s - 6 + \delta(\delta-1)$ edges and thus an edge density of at most $\frac{3s-6+\delta(\delta-1)}{s} < 3 + \frac{\delta}{s}(\delta-1) < \delta + 2 = \delta'$.

Finally, we argue about the shortcut quality. We have $(\delta-1)D+1$ paths P_i , each of length $(\delta-1)D$. For each path P_i , the only way to shorten the distance between the two endpoints is to use the edges of the top-path. In fact, unless a part P_i has at least $1/2$

of the edges of the top path in its shortcut H_i , the dilation of its part, i.e., the diameter of $G[P_i] + H_i$ would be at least $\frac{1}{2}(\delta - 1)D$. Therefore, if the shortcut quality is smaller than $\frac{1}{2}(\delta - 1)D$, each part P_i needs to have at least half of the edges of the top path in its shortcut. But then, overall, the $(\delta - 1)k + 1$ edges of the top path appear at least $\frac{1}{2}((\delta - 1)D)((\delta - 1)k + 1)$ times, in total, in shortcuts. Hence, at least one edge has congestion at least $\frac{1}{2}(\delta - 1)D$. Lastly $\frac{1}{2}(\delta - 1)D \geq \frac{1}{6}(\delta' - 3)D'$. \square

We note that there is a $\Theta(\log n)$ factor gap in the congestion (but not the dilation) between the upper bound for full shortcuts in Theorem 1.2 and Lemma 1.3. This gap stems from the $\log_2 n$ loss of *Observation 2.7*, which goes from partial shortcuts to (full) shortcuts. A planar lower bound topology with $\Omega(D \log D)$ shortcuts given in [7] shows that the congestion-gap between partial shortcuts and full shortcuts can be at least $\Omega(\log D)$. Whether this is the maximal gap between partial and full shortcuts and whether Theorem 1.2 can be improved to shortcuts of quality independent of n , e.g., $O(\delta D \log D)$, are interesting questions – albeit not ones of particular importance to the algorithmic applications of shortcuts.

3.3 Shortcuts for other Graph Parameters and Algorithmic Applications

We finish by giving the few remaining technical details for the direct implications of our existential and algorithmic shortcut guarantees from Theorem 3.1 and Theorem 1.5, which are stated in Section 1.2.

As discussed in Section 1.2, our shortcut guarantees in terms

parameters of interest. Consider, for example, the following easy and existentially tight bounds on the minor density in terms of graph parameters considered in [15]:

LEMMA 3.3. *The following bounds on the minor density $\delta(G)$ of a graph G hold:*

- If G has genus, non-orientable genus, or Euler genus of g , then $\delta(G) = O(\sqrt{g})$.
- If G has treewidth or pathwidth at most k , then $\delta(G) \leq k$.

PROOF. Being embeddable into an orientable or non-orientable surface is a graph property that is closed under minor-operations. Moreover, An n -node graph G having a genus, non-orientable genus, or Euler genus of g implies that this graph has at most $3n + O(g)$ edges. Since any graph with density δ has at least $n \geq \delta$ nodes and thus at least $\min\{\delta n, \frac{\delta^2}{2}\} \geq \frac{\delta n}{2} + \frac{\delta^2}{4}$ edges, the (minor) density of any such graph is at most $O(\sqrt{g})$. Similarly, having treewidth or pathwidth at most k are graph properties closed under taking minors. Moreover, a graph of treewidth (or pathwidth) at most k and n nodes has less than kn edges and therefore its (minor) density is at most k . \square

Combining Theorem 3.1 with Lemma 3.3 now directly implies Corollary 1.4 and the following analogous corollary for treewidth- k graphs:

COROLLARY 3.4. *Any graph with n nodes, diameter D , and treewidth at most k admits shortcuts with congestion $O(kD \log n)$ and dilation $O(kD)$.*

Note that while completely different and highly nontrivial proofs specific to planar graphs, bounded genus graphs, and bounded treewidth graphs were given in [7, 15], we obtain the same existentially optimal results by simply plugging in bounds on the minor density $\delta(G)$ in terms of the desired graph parameter into Theorem 3.1. Constructive results of these shortcuts follow similarly from Theorem 1.5.

Algorithmic applications, such as, the fast distributed MST and minimum-cut algorithms claimed in Corollary 1.6 and Corollary 1.7 follow immediately and in a completely modular fashion from our new constructive shortcuts given in Theorem 1.5 and shortcut-based algorithms like the ones given in [7, 12].

PROOF OF COROLLARY 1.6 AND COROLLARY 1.7. The min-cut algorithm follows from a randomized algorithm given in [7] that computes a $(1 + \varepsilon)$ approximation in $O(Q \text{poly}(1/\varepsilon))$ time, with high probability, given a shortcut of quality Q which is can be constructed according to Theorem 1.5. To convert this into an exact algorithm we observe that the minimum degree of G and therefore also its min-cut is of size at most 2δ given that the density of G can be at most δ . Setting $\varepsilon = \frac{1}{4\delta}$ therefore implies an exact algorithm. The MST statement follows directly from Theorem 1.5 and Boruvka's distributed MST algorithm, as described in [7, 12]. \square

ACKNOWLEDGEMENTS

The second author thanks Anupam Gupta for suggesting the question that led to this paper.

REFERENCES

- [1] Noga Alon, Paul Seymour, and Robin Thomas. 1990. A separator theorem for nonplanar graphs. *Journal of the American Mathematical Society* 3, 4 (1990), 801–808.
- [2] Michal Dory and Mohsen Ghaffari. 2019. Improved distributed approximations for minimum-weight two-edge-connected spanning subgraph. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing (PODC)*. 521–530.
- [3] Michael Elkin. 2004. 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)*. 331–340.
- [4] J.A. Garay, S. Kutten, and D. Peleg. 1993. A sub-linear time distributed algorithm for minimum-weight spanning trees. In *Proc. of the Symp. on Found. of Comp. Sci. (FOCS)*.
- [5] Mohsen Ghaffari. 2015. Near-Optimal Scheduling of Distributed Algorithms. In *the Proc. of the Int'l Symp. on Princ. of Dist. Comp. (PODC)*, to appear.
- [6] Mohsen Ghaffari and Bernhard Haeupler. 2016. Distributed Algorithms for Planar Networks I: Planar Embedding. In *the Proc. of the Int'l Symp. on Princ. of Dist. Comp. (PODC)*.
- [7] Mohsen Ghaffari and Bernhard Haeupler. 2016. Distributed Algorithms for Planar Networks II: Low-Congestion Shortcuts, MST, and Min-Cut. In *Proc. of ACM-SIAM Symp. on Disc. Alg. (SODA)*.
- [8] Mohsen Ghaffari, Fabian Kuhn, and Hsin-Hao Su. 2017. Distributed MST and routing in almost mixing time. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*. 131–140.
- [9] Mohsen Ghaffari and Jason Li. 2018. New Distributed Algorithms in Almost Mixing Time via Transformations from Parallel Algorithms. In *32nd International Symposium on Distributed Computing (DISC)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [10] Mohsen Ghaffari and Merav Parter. 2017. Near-optimal distributed DFS in planar graphs. In *31st International Symposium on Distributed Computing (DISC)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [11] John R Gilbert, Joan P Hutchinson, and Robert Endre Tarjan. 1984. A separator theorem for graphs of bounded genus. *Journal of Algorithms* 5, 3 (1984), 391–407.
- [12] Bernhard Haeupler, D Ellis Hershkowitz, and David Wajc. 2018. Round-and message-optimal distributed graph algorithms. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing*. 119–128.
- [13] Bernhard Haeupler, D Ellis Hershkowitz, and David Wajc. 2019. Near-Optimal Schedules for Simultaneous Multicasts. *arXiv cs.DS 2001.00072* (2019).
- [14] Bernhard Haeupler, Taisuke Izumi, and Goran Zuzic. 2016. Low-congestion shortcuts without embedding. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing (PODC)*. 451–460.
- [15] Bernhard Haeupler, Taisuke Izumi, and Goran Zuzic. 2016. Near-optimal low-congestion shortcuts on bounded parameter graphs. In *International Symposium on Distributed Computing (DISC)*. Springer, 158–172.
- [16] Bernhard Haeupler and Jason Li. 2018. Faster Distributed Shortest Path Approximations via Shortcuts. In *32nd International Symposium on Distributed Computing (DISC)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [17] Bernhard Haeupler, Jason Li, and Goran Zuzic. 2018. Minor excluded network graphs admit fast distributed algorithms. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing (PODC)*. 465–474.
- [18] Naoki Kitamura, Hirotaka Kitagawa, Yota Otachi, and Taisuke Izumi. 2019. Low-Congestion Shortcut and Graph Parameters. In *33rd International Symposium on Distributed Computing (DISC)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [19] Shay Kutten and David Peleg. 1995. Fast Distributed Construction of K -dominating Sets and Applications. In *the Proc. of the Int'l Symp. on Princ. of Dist. Comp. (PODC)*. 238–251.
- [20] Frank Thomson Leighton, Bruce M Maggs, and Satish B Rao. 1994. Packet routing and job-shop scheduling in (congestion+ dilation) steps. *Combinatorica* 14, 2 (1994), 167–186.
- [21] Jason Li and Merav Parter. 2019. Planar diameter via metric compression. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. 152–163.
- [22] Richard J Lipton and Kenneth W Regan. 2013. *People, Problems, and Proofs: Essays from Gödel's Lost Letter: 2010*. Springer Science & Business Media.
- [23] Richard J Lipton and Robert Endre Tarjan. 1979. A separator theorem for planar graphs. *SIAM J. Appl. Math.* 36, 2 (1979), 177–189.
- [24] Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. 2020. Efficient Graph Minors Theory and Parameterized Algorithms for (Planar) Disjoint Paths. In *Treewidth, Kernels, and Algorithms*. Springer, 112–128.
- [25] Zvi Lotker, Boaz Patt-Shamir, and David Peleg. 2006. Distributed MST for constant diameter graphs. *Distributed Computing* 18, 6 (2006), 453–460.
- [26] Gary L Miller. 1984. Finding small simple cycle separators for 2-connected planar graphs. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing (STOC)*. ACM, 376–382.
- [27] Jaroslav Nešetřil and Patrice Ossona De Mendez. 2012. *Sparsity: graphs, structures, and algorithms*. Vol. 28. Springer Science & Business Media.
- [28] Jaroslav Nešetřil, Eva Milková, and Helena Nešetřilová. 2001. Otakar Boruvka on minimum spanning tree problem translation of both the 1926 papers, comments, history. *Discrete Math.* 233, 1 (2001), 3–36.

- [29] David Peleg. 2000. *Distributed Computing: A Locality-sensitive Approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA.
- [30] David Peleg and Vitaly Rubinovich. 1999. A Near-Tight Lower Bound on the Time Complexity of Distributed MST Construction. In *Proc. of the Symp. on Found. of Comp. Sci. (FOCS)*. 253–.
- [31] Neil Robertson and Paul D. Seymour. 1986. Graph minors. II. Algorithmic aspects of tree-width. *Journal of algorithms* 7, 3 (1986), 309–322.
- [32] Neil Robertson and Paul D Seymour. 2003. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B* 89, 1 (2003), 43–76.
- [33] Atish Das Sarma, Stephan Holzer, Liah Kor, Amos Korman, Danupon Nanongkai, Gopal Pandurangan, David Peleg, and Roger Wattenhofer. 2012. Distributed verification and hardness of distributed approximation. *SIAM J. Comput.* 41, 5 (2012), 1235–1265.
- [34] Andrew Thomason. 2001. The extremal function for complete minors. *Journal of Combinatorial Theory, Series B* 81, 2 (2001), 318–338.