

Property Testing of Planarity in the CONGEST model*

Reut Levi[†]Moti Medina[‡]Dana Ron[§]

Abstract

We give a distributed algorithm in the CONGEST model for property testing of planarity with one-sided error in general (unbounded-degree) graphs. Following Censor-Hillel et al. (DISC 2016), who recently initiated the study of property testing in the distributed setting, our algorithm gives the following guarantee: For a graph $G = (V, E)$ and a distance parameter ϵ , if G is planar, then every node outputs **accept**, and if G is ϵ -far from being planar (i.e., more than $\epsilon \cdot |E|$ edges need to be removed in order to make G planar), then with probability $1 - 1/\text{poly}(n)$ at least one node outputs **reject**. The algorithm runs in $O(\log |V| \cdot \text{poly}(1/\epsilon))$ rounds, and we show that this result is tight in terms of the dependence on $|V|$.

Our algorithm combines several techniques of graph partitioning and local verification of planar embeddings. Furthermore, we show how a main subroutine in our algorithm can be applied to derive additional results for property testing of cycle-freeness and bipartiteness, as well as the construction of spanners, in minor-free (unweighted) graphs.

1 Introduction

Planarity is an important and well studied property of graphs. In the setting of centralized algorithms, there are several algorithms that run in linear time for deciding whether a graph is planar (e.g., [36, 33, 4]). In the context of distributed algorithms in the CONGEST [42] (and even LOCAL [41]) model, the number of rounds must be at least linear in the diameter of the graph (for any deterministic or randomized one-sided error algorithm). One begging question is whether there exists an algorithm (in the CONGEST model) for deciding planarity whose round complexity matches (or is not too far) from this lower bound.¹ Another question is whether there exists a natural relaxation of this decision problem, which allows to obtain round complexity that does not depend on the diameter.

In this work we address the latter question, by considering the relaxation of *property testing* in the distributed setting. In all that follows, unless explicitly stated otherwise, when we refer to distributed algorithms, we mean in the CONGEST model. Following Censor-Hillel et al. [6],

*This article extends work presented at PODC 2018 [39].

[†]Weizmann Institute of Science, reut.levi@weizmann.ac.il. This project has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No. 819702). It was also supported by ERC-CoG grant 772839.

[‡]Department of Electrical & Computer Engineering, Ben-Gurion University of the Negev, medinamo@bgu.ac.il.

[§]School of Electrical Engineering, Tel Aviv University, danaron@tau.ac.il. This research was partially supported by the Israel Science Foundation grants No. 671/13 and 1146/18.

¹We note that Ghaffari and Haeupler [23] consider a different, but related question of finding a planar embedding of a planar graph. They give a distributed algorithm for this problem using $O(D \cdot \min\{\log n, D\})$ rounds (where D is the diameter and n is the number of nodes).

who recently initiated the study of distributed property testing, we require the following from the algorithm. Let $G = (V, E)$ be a graph over n nodes and m edges. If G is planar, then all nodes should output **accept**, while if G is ϵ -far from being planar (i.e., more than $\epsilon \cdot m$ edges should be removed in order to make G planar), then at least one node should output **reject**. The algorithm is allowed a bounded error probability, where if it errs only on graphs that are ϵ -far from being planar, then it is said to have *one-sided error*.²

Our main result is a distributed one-sided error property testing algorithm for planarity that runs in $O(\log(n)\text{poly}(1/\epsilon))$ rounds and succeeds with probability $1 - 1/\text{poly}(n)$. We also show that $\Omega(\log n)$ rounds are necessary for any such algorithm and constant ϵ (even if the algorithm is allowed a constant error probability), implying that our result is tight up to the dependence on ϵ .

In the context of (centralized) property testing, there is a line of work on two-sided error testing of planarity in bounded-degree graphs [3, 32, 40, 35]. The best known algorithm [35] performs polynomial number of queries in d and ϵ , where d is the degree bound (and succeeds with high constant probability). The best one-sided error testing algorithm [34], for bounded-degree graphs, has query complexity $n^{1/2+o(1)}$, which is almost optimal. There is no known sublinear testing algorithm for unbounded-degree graphs. Note that in contrast, in the distributed setting we are able to obtain an optimal (in terms of the dependence on n) algorithm that has one-sided error and works for unbounded-degree graphs.

In addition to our main result, we show that under the promise that G is planar (and more generally, minor-free³ for any fixed minor), we can use our techniques to obtain other distributed property testing algorithms as well as an algorithm for the construction of spanners.⁴

In the next two subsections we discuss our results in more detail.

1.1 A high-level description of our algorithm for testing planarity

The algorithm works in two stages. The goal of the first stage (which is deterministic) is to partition the nodes of G into parts for which the following holds: (1) Each part is connected and has diameter $\text{poly}(1/\epsilon)$; (2) The total number of edges between parts is at most $\epsilon m/2$. In the course of this stage, some node(s) may obtain evidence that the graph is not planar, and output **reject**. This evidence is in the form of messages received that are not consistent with the execution of the algorithm on any planar graph. Conditioned on this stage completing successfully, if G is ϵ -far from being planar, then the subgraph induced by at least one of the parts in the partition is $(\epsilon/2)$ -far from being planar. The goal of the second stage is to search for evidence in each part to non-planarity, by exploiting the fact that the diameter of each part is small.

We next give some more details about each stage of the algorithm. Before doing so we recall several notions and basic facts. A *forest decomposition* of a graph is a partition of its edges into forests. The *arboricity* of a graph is the minimum number of forests into which its edges can be partitioned. Any planar graph has arboricity at most 3, and if we perform any sequence of contractions of edges on a planar graph, then we obtain a planar graph.

²Observe that if the algorithm has one-sided error and a constant error probability, then its error probability can be reduced to δ , for a given parameter δ , at a multiplicative cost of $\log(1/\delta)$ in the number of rounds. We are able to obtain error probability $1/\text{poly}(n)$ without this extra cost.

³Recall that a graph H is a *minor* of a graph G if H is isomorphic to a graph that can be obtained by zero or more edge contractions on a subgraph of G . We say that a graph G is *H -minor free* (or *excludes H as a minor*) if H is not a minor of G . We say that G is “minor-free” if it is H -minor free for some fixed H of constant size.

⁴A spanner of a graph G is a (sparse) subgraph of G that maintains distances up to a multiplicative factor, s , which is called the *stretch* factor, and the spanner is referred to as an *s-spanner*.

The first stage. The first stage consists of $t = O(\log(1/\epsilon))$ phases. At the start of phase i , the nodes are partitioned into k_i parts, denoted $P_i^1, \dots, P_i^{k_i}$ where each part is connected and has diameter at most 4^{i-1} . Let us denote this partition by \mathcal{P}_i . In the initial partition, \mathcal{P}_1 , each part simply consists of a single node. For each phase, let \mathcal{G}_i denote the auxiliary weighted graph that results from contracting each part P_i^j into a single node, which we denote by $v(P_i^j)$. The weight of an edge $(v(P_i^j), v(P_i^{j'}))$ is the number of edges in G with one endpoint in P_i^j and the other in $P_i^{j'}$.

Each phase starts by emulating the (deterministic distributed) forest decomposition algorithm of Barenboim and Elkin [2] on \mathcal{G}_i (ignoring the weights). This algorithm works in $O(\log(n))$ rounds and gives the following guarantee. If \mathcal{G}_i has arboricity at most α , then the algorithm provides a forest decomposition into at most 3α forests. On the other hand, if the algorithm fails in defining such a decomposition, then at least one node (in \mathcal{G}_i and hence in G) has evidence that \mathcal{G}_i has arboricity greater than α .

Following the above *forest decomposition step*, the algorithm executes a *merging step* (based on the clustering method of Czygrinow, Hańćkowiak, and Wawrzyniak [10]). In this step, parts of \mathcal{P}_i are merged, thus creating \mathcal{P}_{i+1} . This merging procedure satisfies the following: (1) The maximum diameter of the parts in \mathcal{P}_{i+1} is at most a constant factor larger than the maximum diameter of the parts in \mathcal{P}_i ; (2) The number of edges in G between parts in \mathcal{P}_{i+1} (the total weight of edges in \mathcal{G}_{i+1}) is a constant fraction of the number of edges between parts in \mathcal{P}_i (the total weight of edges in \mathcal{G}_i). The latter crucially relies on the bounded arboricity of \mathcal{G}_i , which is ensured by the forest decomposition step.

Thus, the following central feature of (the first stage of) our algorithm emerges. Though we do not have a promise that the underlying graph is planar (indeed, it may be far from being planar), we are able to build on algorithms that work under the promise that the graph is planar. Namely, in each phase we verify that \mathcal{G}_i has bounded arboricity, by running an algorithm that works under the promise that the underlying graph has bounded arboricity. Failure of this algorithm is detected by at least one node in the graph. On the other hand, if this algorithm succeeds, then we are ensured that we shall make the desired progress in the transformation from \mathcal{G}_i to \mathcal{G}_{i+1} (in terms of the decrease in the total weight of edges).

The second stage. Assume the first stage completed successfully (where this always holds if G is planar), and let $\mathcal{P} = (P^1, \dots, P^k)$ be the final partition of the nodes (i.e., $\mathcal{P} = \mathcal{P}_t$). Recall that the subgraph induced by each part is connected, and has diameter $\text{poly}(1/\epsilon)$. Furthermore, the first stage ensures the following for each part P^j : (1) There is a designated *root* vertex in P^j , denoted r^j , where each node in P^j knows the identity of r^j ; (2) There is an underlying spanning tree in P^j , rooted at r^j , where each node in P^j knows which of its incident edges is also incident to its parent in the tree, and which edges are incident to its children.

The second stage consists of two steps. In the first step, the (deterministic distributed) planar embedding algorithm of Ghaffari and Haeupler [23] is emulated on the subgraph induced by each part P^j , denoted G^j . The planar embedding algorithm works under the promise that G^j is planar, and when it completes, each node in G^j has a circular ordering over its incident edges that corresponds to a planar embedding (what is known as a *combinatorial embedding*). Since the diameter of each part is $\text{poly}(1/\epsilon)$, the number of rounds performed by this algorithm is $\text{poly}(1/\epsilon)$. If this step fails in determining an ordering for all nodes (in the aforementioned number of rounds), then this constitutes evidence that G^j is not planar. However, it is possible that an ordering is determined though G^j is not planar.

Hence, the second step in this stage is aimed at detecting non-planarity of some G^j given the ordering provided by the [23] algorithm.⁵ More precisely, as noted previously, if G is ϵ -far from being planar, then at least one G^j is $(\epsilon/2)$ -far from being planar. Using the ordering of edges incident to each node in G^j together with a BFS tree rooted at r^j , denoted T_B^j (which can be constructed in $\text{poly}(1/\epsilon)$ rounds), we define a certain condition on each of the non-tree edges of T_B^j . We show that if G^j is far from being planar, then there are relatively many non-tree edges in G^j that violate this condition, while if G^j is planar, then no non-tree edge violates this condition. Furthermore, given a violating edge, it is possible to detect violation in $\text{poly}(1/\epsilon)$ rounds.⁶ Hence, by sampling $\Theta(\log(n)/\epsilon)$ non-tree edges in each G^j and running the detection procedure on each, a violation is detected with probability $1 - 1/\text{poly}(1/n)$.

We note that, as shown in [14, 13], the algorithm of Elkin and Neiman [12] can be adapted to obtain with high probability a partition of the nodes into parts of diameter $O(\log(n)/\epsilon)$ such that the number of edges between parts is at most ϵm . Replacing Stage I in our algorithm with this procedure (and running Stage II on each part), results in a testing algorithm that runs in $O(\log^2(n)\text{poly}(1/\epsilon))$ (while our algorithm runs in $O(\log(n)\text{poly}(1/\epsilon))$ rounds).

1.2 Implications and applications for minor-free graphs

Suppose we have a promise that G is planar, or more generally, minor-free for any constant size minor. In such a case, the first stage of our algorithm always ensures that the nodes of G are partitioned into parts with diameter $\text{poly}(1/\epsilon)$ such that the number of edges between parts is at most $\epsilon m/2$. Such a partition can be used for testing properties such as cycle-freeness and bipartiteness (and more generally, hereditary properties that can be tested in a number of rounds that is linear (or even polynomial) in the diameter). Thus, for cycle-freeness and bipartiteness we obtain a *deterministic* testing algorithm that runs in $O(\log(n)\text{poly}(1/\epsilon))$ rounds. Such a partition can also be used to obtain $\text{poly}(1/\epsilon)$ -spanners (for unweighted graphs) (deterministically) in $O(\log(n)\text{poly}(1/\epsilon))$ rounds.

We also show how to modify the partition algorithm so as to obtain a tradeoff between the round-complexity and the success probability of the algorithm. More precisely, with probability at least $1 - \delta$, the modified algorithm gives the same guarantee as above for the partition in $O(\log(1/\epsilon)(\log^*(n) + \log(1/\delta)))$ rounds. The complexity of the testing algorithms and spanner construction algorithm are improved accordingly (see Corollaries 20 and 21, respectively). Finally, if constant success probability suffices, then it is possible to remove the dependence on n completely.

The testing results can be compared with the $\Omega(\log n)$ lower bound of Censor-Hillel et al. [6] for distributed testing of these properties on general (bounded-degree) graphs (with constant success probability). The spanner result can be compared to the recent result of Elkin and Neiman [12]. They provide a k -round distributed algorithm for general (unweighted) graphs that with probability $1 - \delta$ constructs a $(2k - 1)$ -spanner with $O(n^{1+1/k}/\delta)$ edges. In order to obtain an *ultra-sparse spanner*, namely, a spanner of size $n(1 + o(1))$ (with probability $1 - o(1)$), it is necessary to set $k = \omega(\log n)$. In our context, of minor-free graphs, we can obtain an ultra-sparse spanner deter-

⁵It was communicated to us by one of the authors of [23] that their algorithm can be modified so as to detect if the underlying graph is not planar [28]. For the sake of a self-contained presentation, we rely on version of the algorithm as provided in [23] (which works under the promise that the graph is planar), and check that the ordering computed by [23] is consistent with a planar embedding.

⁶We note that a previous version of the definition of this condition, given in the conference version of this paper [39], contained an error, which is fixed in the current version.

ministically for any $\epsilon = o(1)$, which allows to construct ultra-sparse spanners with stretch s for any $s = \omega(1)$.

1.3 Related work

Distributed Property Testing. As noted previously, the study of distributed property testing was initiated by Censor-Hillel et al. [6]. In particular, they designed and analyzed distributed property testing algorithms for: triangle-freeness, cycle-freeness, and bipartiteness. As noted previously, they also proved a logarithmic lower bound for the latter two properties. Fraigniaud et al. [22] studied distributed property testing of excluded subgraphs of size 4 and 5.

Since the appearance of the above papers, there was a fruitful line of research in distributed property testing for various properties, mainly focusing on properties of whether a graph excludes a sub-graph, e.g., triangle-freeness, cycle-freeness, subgraphs of constant size, tree-freeness, clique freeness [21, 19, 20, 14, 13, 18]. In [17] the problem of testing the conductance of the input graph was studied, and a two-sided error tester is given.

Brakerski and Patt-Shamir [5] considered a related problem in the distributed setting. They show how to find a subset of vertices that is ϵ -close to being a clique if there is a subset that is ϵ^3 -close to being a clique.

Distributed Algorithms with a promise in the CONGEST model. There is a large variety of distributed algorithms that work under a promise that the graph is planar (similarly, excludes a fixed minor, has bounded arboricity, and more). See e.g. [2, 38, 37, 1, 23, 24, 29, 30, 25, 26, 31].

Centralized Property Testing. Most works in property testing on H -minor freeness (and related properties) focus on two-sided error testers. The problem of testing general H -minor freeness was studied by Benjamini, Schramm and Shapira in [3]. They showed that every minor-closed property of bounded degree graphs is testable with two-sided error with query complexity which is independent of the size of the graph. Subsequent work [32, 40] improved the dependence on the proximity parameter, ϵ . Yoshida and Ito [43] provided a tester with two-sided error for outerplanarity whose query complexity is $\text{poly}(1/\epsilon)$. Eden et al. [11] studied tolerant testing of bounded arboricity in the general graph model and showed almost tight bounds in terms of the dependence on n and m .

Czumaj et al. [8] studied property testing of H -minor freeness with one-sided error. They proved that for H which is a forest, H -minor freeness can be tested in query complexity which is independent on n . For any H that contains a cycle they showed a lower bound of $\Omega(\sqrt{n})$ queries. They also provided almost matching upper bounds for any H which is a simple cycle. Recently, Fichtenberger et al. [16] provided a $\tilde{O}(n^{2/3})$ -query tester for outerplanarity, and other properties that can be characterized by forbidden minors, with one-sided error.

Czumaj et al. [9] studied testing bipartiteness under the promise the input graph is planar. While in general, the query complexity of testing bipartiteness is $\Omega(\sqrt{n})$, even for bounded degree graphs, they showed that under the promise of planarity, bipartiteness can be tested in time which is independent of n , even if the maximum degree is unbounded.

2 The algorithm for testing planarity

In this section we establish the following theorem.

Theorem 1 *There exists a distributed one-sided error property testing algorithm for planarity that runs in $O(\log(n) \cdot \text{poly}(1/\epsilon))$ rounds in the CONGEST model.*

As noted in the introduction (Section 1.1), our algorithm works in two stages: a *Partition stage* and a *Planarity testing stage*. In what follows we describe and analyze each in detail, where we refer to the first as Stage I and the second as Stage II.

2.1 A detailed description and analysis of Stage I

Recall from the description in the introduction (Section 1.1) that Stage I consists of $t = O(\log(1/\epsilon))$ phases. Each phase is associated with a partition of the nodes, where in the initial partition each node is a singleton part, and in general, each part is connected. We denote by $\mathcal{P}_i = (P_i^1, \dots, P_i^{k_i})$ the partition associated with Phase i . Each partition \mathcal{P}_i defines an auxiliary weighted graph, denoted \mathcal{G}_i . Each part P_i^j corresponds to a node in \mathcal{G}_i , denoted $v(P_i^j)$, and the weight of an edge $(v(P_i^j), v(P_i^{j'}))$ is the number of edges with one endpoint in P_i^j and the other in $P_i^{j'}$. Each part in \mathcal{P}_{i+1} is a union of several parts in \mathcal{P}_i .

Each phase consists of two steps: A *forest decomposition step* and a *merging step*. We first describe these steps in terms of the auxiliary graphs $\{\mathcal{G}_i\}_{i=1}^{t+1}$ (so that each part P_i^j is viewed as a single node, which may be denoted by v or u rather than $v(P_i^j)$). We then explain how they are emulated on G .

Sending message up and down trees. Both in G and in the auxiliary graphs $\{\mathcal{G}_i\}_{i=1}^{t+1}$, we consider distributing information on trees (where each node knows which of its incident edges is incident to its parent in the tree and which of these edges are incident to its children). For such a tree T , let $r(T)$ denote its root. When we say that $r(T)$ sends a message *down the tree*, we mean that $r(T)$ sends the message to its children in T and they send it to their children, and so on, until the message reaches all nodes in T . In some cases the message may be augmented/modified as it goes down the tree. When we say that a node $v \in T$ sends a message *up the tree*, we mean that v sends the message to its parent in T , which in turn send it to its parent, and so on, until it reaches the root $r(T)$. Here too a message may be modified as it goes up the tree. In particular, if several nodes simultaneously send different messages up the tree, then this may cause congestion, and we explain how this is addressed whenever it arises.

2.1.1 The forest decomposition step on \mathcal{G}_i

This step correspond to the forest decomposition algorithm of Barenboim and Elkin [2]. Their algorithm works under the promise that the underlying graph \mathcal{G}_i has arboricity at most α where in our case α is set to 3 (the bound on the arboricity of planar graphs). The algorithm ignores the weights on the edges of \mathcal{G}_i , and proceeds as follows. Initially all nodes in \mathcal{G}_i are *active*. For $s = \Theta(\log n)$ rounds, each active node u does the following. If u has at most 3α active neighbors (in the current round), then u sends a message to all its neighbors that it becomes *inactive* in the next round. As shown in [2] (and is not hard to verify), if \mathcal{G}_i has arboricity at most α , then in each round a constant fraction of the nodes become inactive. Since the number of rounds is $\Theta(\log n)$ (and the number of nodes in \mathcal{G}_i is at most n), if \mathcal{G}_i has arboricity at most α , all the nodes are inactive by the termination of the algorithm. In other words, if some node remains active after s rounds, then this serves as evidence that \mathcal{G}_i has arboricity larger than α .

If the process terminates successfully (i.e., all nodes become inactive), then it is possible to define a forest decomposition into at most 3α forests. That is, it is possible to orient the edges of \mathcal{G}_i so that each node in \mathcal{G}_i has at most 3α outgoing edges, one to each of its parents in the different forests (where no cycles are formed). Specifically, consider a node u that becomes inactive in round ℓ , and let v_1, \dots, v_d , for $d \leq 3\alpha$, be its active neighbors at the start of round ℓ . For each v_q such that v_q remains active in round $\ell + 1$, we orient the edge $\{u, v_q\}$ from u to v_q . For each v_q such that v_q also becomes inactive in round ℓ , we orient the edge $\{u, v_q\}$ from the node with the smaller id to the node with the larger id (so that no directed cycles are formed). We emphasize that each node knows the orientation of its incident edges.

2.1.2 The merging step: from \mathcal{G}_i to \mathcal{G}_{i+1}

Assuming the forest decomposition step completed successfully (and hence each node in \mathcal{G}_i has at most 3α outgoing edges pointing to its parents in the forest decomposition), the merging step consists of the following sub-steps. Sub-step 2 is as in [10].

1. Each node u in \mathcal{G}_i selects its outgoing edge (u, v) that has the highest weight. Let \mathcal{F}_i denote the forest induced by the selected edges.
2. Select a set of “shallow” subtrees of \mathcal{F}_i , denoted \mathcal{T}_i , as follows:
 - (a) Obtain a coloring χ of \mathcal{F}_i using colors in $\{1, 2, 3\}$ by running the distributed algorithm of Cole and Vishkin [7], and Goldberg, Plotkin and Shannon [27].
 - (b) Mark the edges of \mathcal{F}_i according to the rules defined next (where if a node marks an incident edge, it notifies the other endpoint).
 - For each node u such that $\chi(u) = 1$, u marks its outgoing edge (assuming such exists) if the weight of this edge is greater or equal to the sum of the weights of all its incoming edges. Otherwise, u marks all its incoming edges.
 - For each node u such that $\chi(u) = 2$, u marks its outgoing edge if the other endpoint is colored 3 and the weight of this edge is greater or equal to the sum of the weights of all its incoming edges whose other endpoint is colored 3. Otherwise it marks all these incoming edges.
 - (c) Let $\bar{\mathcal{T}}_i$ be the set of subtrees induced by the marked edges.
3. For each subtree $T \in \bar{\mathcal{T}}_i$, let $w_0(T)$ denote the total weight of edges that go from an even level in T up to an odd level (referred to as “even edges”), and let $w_1(T)$ denote the total weight of the remaining edges in T (referred to as “odd edges”). The root of T , denoted $r(T)$, obtains $w_0(T)$ and $w_1(T)$ (by sending a message down the tree so that each node learns its level, and receiving message sent up the tree in which weights of even/odd edges are summed). If $w_0(T) \geq w_1(T)$, then $r(T)$ sends the message ‘0’ down the tree, and otherwise it sends ‘1’.
4. If the message sent down the tree is ‘0’, then all even edges are contracted, and otherwise all odd edges are contracted.

Observe that each node in \mathcal{G}_{i+1} corresponds to a star subgraph in \mathcal{G}_i .

Let $w(\mathcal{G}_i)$ denote the total weight of edges in \mathcal{G}_i (and similarly define $w(\mathcal{F}_i)$ and $w(\mathcal{T}_i)$).

Claim 1 *The merging step runs in $O(\log^* n)$ rounds (on \mathcal{G}_i) and $w(\mathcal{G}_{i+1}) \leq (1 - \frac{1}{12\alpha}) \cdot w(\mathcal{G}_i)$.*

Proof: First observe that by the definition of \mathcal{F}_i (in Sub-step 1), $w(\mathcal{F}_i) \geq w(\mathcal{G}_i)/3\alpha$. By [7], the number of rounds performed in Sub-step 2a is $O(\log^* n)$ (and hence $O(\log^* n)$ bounds the number of rounds performed in all of Sub-step 2). By the analysis presented in [10, Section 2], $w(\mathcal{T}_i) \geq w(\mathcal{F}_i)/2$, and the height of each tree T in \mathcal{T}_i is at most 10. Therefore, Sub-step 3 runs in a constant number of rounds (and the same is true of Sub-step 4). Finally, by the choice of which edges to contract (in Sub-step 3), the weight of the contracted edges is at least $w(\mathcal{T}_i)/2$. The claim follows. ■

2.1.3 Successful completion of Stage I

Before turning to the emulation of Stage I on G , we introduce one definition and two claims, whose correctness does not depend on the details of the emulation.

Definition 2 *We say that Stage I completes successfully if the forest decomposition step in each phase terminates with no remaining active node.*

Claim 3 *If G is planar, then Stage I always completes successfully. If G is ϵ -far from being planar, then either Stage I does not complete successfully, or $w(\mathcal{G}_{t+1}) \leq \epsilon m/2$.*

Proof: The first part of the claim follows immediately from the fact that the arboricity of planar graphs is at most 3 and that any minor of a planar graph is planar. The second part of the claim follows from the fact that $w(\mathcal{G}_1) = m$, Claim 1, and the setting of the number of phases $t = \Theta(\log(1/\epsilon))$. ■

Claim 4 *For each phase i and part P_i^j , the subgraph induced by P_i^j is connected and has diameter at most 4^i .*

Proof: We prove the claim by induction on i . The claim trivially holds for $i = 1$. To establish the induction step, consider a single merging step in which the nodes $v(P_i^{j_1}), \dots, v(P_i^{j_s})$ in \mathcal{G}_i all merge with $v(P_i^{j_0})$ (that is, the edges $\left\{ \left(v(P_i^{j_q}), v(P_i^{j_0}) \right) \right\}_{q=1}^s$ were contracted). Clearly, the subgraph induced by $\bigcup_{q=0}^s P_i^{j_q}$ (which corresponds to a part in \mathcal{P}_{i+1}) is connected. As for the diameter of this subgraph, by the induction hypothesis, it is at most $3 \cdot 4^i + 2 \leq 4^{i+1}$. ■

Conditioned on Stage I completing successfully, let $\mathcal{P} = (P^1, \dots, P^k)$ denote the final partition (i.e., $\mathcal{P} = \mathcal{P}_{t+1}$). For each $j \in [k]$, let $G^j = G(P^j)$ denote the subgraph induced by P^j . As a corollary of Claim 4 we get:

Corollary 5 *Each G^j is connected and has diameter $\text{poly}(1/\epsilon)$.*

2.1.4 Preliminaries for the emulation

For each phase $i \in [t]$ and part $P_i^j \in \mathcal{P}_i$, let G_i^j denote the subgraph induced by P_i^j . We say that a node $u \in P_i^j$ is a *boundary* node, if at least one of its neighbors belongs to a part $P_i^{j'}$ for $j' \neq j$.

In Section 2.1.6 we establish the following lemma.

Lemma 6 For every $i \in [t]$ and every $P_i^j \in \mathcal{P}_i$, there is a unique root node $r_i^j \in P_i^j$, such that each node in P_i^j knows the identity of r_i^j . Furthermore, there is a spanning tree of G_i^j , rooted at r_i^j and denoted T_i^j , for which the following holds. Each node u in G_i^j knows which of its incident edges is incident to its parent in T_i^j and which of these edges is incident to its children.

2.1.5 Emulating the forest decomposition step

We assume that at the start of each phase i , the conditions in Lemma 6 hold. We refer to each round in the forest decomposition algorithm described in Section 2.1.1 (on \mathcal{G}_i) as a *super-round*. In the forest decomposition step of each phase i , we have $s = O(\log n)$ super-rounds in which each active node $v(P_i^j)$ in \mathcal{G}_i needs to determine if it is still active in the next super-round and to send a corresponding message to its neighbors. Each super-round is emulated by several rounds (on G), as described next.

The root of P_i^j , r_i^j , plays the role of $v(P_i^j)$ as follows. For each super-round ℓ , at the start of which $v(P_i^j)$ is active, if r_i^j determines in the course of this super-round that $v(P_i^j)$ should remain active in the next super-round ($\ell+1$), then it sends a message ('Active', r_i^j) down the tree T_i^j . Each boundary node in P_i^j also sends this message to its neighbors outside of P_i^j . The process by which r_i^j determines in super-round ℓ whether $v(P_i^j)$ remains active or not is defined as follows (where this process is also executed one super-round after $v(P_i^j)$ becomes inactive so that r_i^j can learn which neighbors of $v(P_i^j)$ also became inactive in super-round ℓ).

At the beginning of each super-round ℓ , each boundary node u in P_i^j that received in the previous super-round messages of the form ('Active', $r_i^{j'}$) for $j' \neq j$, does the following. If u received more than 3α such messages with *distinct* root ids, then it sends a message 'Active' up the tree (meaning that $v(P_i^j)$ should remain active since it has more than 3α active neighbors). Otherwise, for each $r_i^{j'}$ such that u received a message ('Active', $r_i^{j'}$), u sends a message ('Active', $r_i^{j'}, x$) to its parent, where x indicates how many messages ('Active', $r_i^{j'}$) it received. These messages go up the tree, where if a node u receives the message 'Active', then this is the single message it passes on. If u did not receive 'Active' but it received more than 3α messages ('Active', $r_i^{j'}, x$) with distinct root ids, then it also sends 'Active' up the tree. Otherwise, for each $r_i^{j'}$, let the messages u received with this root id be ('Active', $r_i^{j'}, x_1$), ..., ('Active', $r_i^{j'}, x_q$). Then u sends its parent a message ('Active', $r_i^{j'}, \sum_{p=1}^q x_p$). Finally, if r_i^j received the message 'Active' or if it received more than 3α messages with distinct root ids, then r_i^j determines that $v(P_i^j)$ remains active in the next super-round. Otherwise, it determines that $v(P_i^j)$ becomes inactive. In the latter case, not only that the out-edges of $v(P_i^j)$ can be determined by r_i^j , their weights can be determined as well (this is the role of the third parameter in the messages going up the tree).

The total number of rounds (on G) sufficient for emulating a single super-round on \mathcal{G}_i is hence upper bounded by the maximum diameter of parts in \mathcal{P}_i , which by Claim 4 is $\text{poly}(1/\epsilon)$ (times 3α , which is a constant).

If after all $O(\log(n))$ super rounds there is some r_i^j such that $v(P_i^j)$ is still active, then r_i^j outputs reject (implying that Stage I did not complete successfully).

2.1.6 Emulating the merging step

In this subsection we explain how to emulate all sub-steps in the merging step, and establish Lemma 6 by induction on i . The base of the induction, $i = 1$ is trivial, since in \mathcal{P}_1 each node belongs to a singleton part, and in one round each node learns the identity of all its neighbors.

Determining the heaviest out-edge (Sub-step 1). For a part P_i^j , consider the super-round ℓ in the forest decomposition algorithm in which $v(P_i^j)$ becomes inactive. By the end of this super-round, r_i^j obtained the ids of the roots, $r_i^{j_1}, \dots, r_i^{j_q}$ where $q \leq 3\alpha$, such that $v(P_i^{j_1}), \dots, v(P_i^{j_q})$ are the active neighbors of $v(P_i^j)$ in \mathcal{G}_i at the start of super-round ℓ . Furthermore, for each of the corresponding parts $P_i^{j_p}$, the root r_i^j obtained the number of edges between P_i^j and $P_i^{j_p}$, so that it knows the weight of $(v(P_i^j), v(P_i^{j_p}))$ in \mathcal{G}_i .

It remains to determine which of these edges is an outgoing edge of $v(P_i^j)$. If ℓ is not the final round, then in the next round, r_i^j learns which of nodes $v(P_i^{j_1}), \dots, v(P_i^{j_q})$ remained active in round $\ell + 1$. For each such node $v(P_i^{j_p})$, the edge $(v(P_i^j), v(P_i^{j_p}))$ is an outgoing edge of $v(P_i^j)$, and for each $v(P_i^{j_p})$ that also became inactive in super-round ℓ , the direction of the edge is determined by the ids of r_i^j and $r_i^{j_p}$. If ℓ is the last round, then either some node in \mathcal{G}_i remained active, causing the corresponding root node in G to reject, or all nodes in \mathcal{G}_i became inactive, so that edge directions are determined by root ids.

For the sake of the following sub-steps, it will be convenient to designate, for each selected outgoing edge $(v(P_i^j), v(P_i^{j'}))$, a single edge (u, v) in G such that $u \in P_i^j$ and $v \in P_i^{j'}$. To this end, let $(v(P_i^j), v(P_i^{h(i,j)}))$ be the heaviest outgoing edge of $v(P_i^j)$. The root, r_i^j send a message with the id of $r_i^{h(i,j)}$ down the tree. Each node $u \in P_i^j$ that has a neighbor in $P_i^{h(i,j)}$ sends its id up the tree, where if a node receives more than one “candidate” node id from its children, then it sends the minimum id among them. In this manner, r_i^j obtains the id of a single node $u_i^j \in P_i^j$ that has a neighbor, v_i^j , in $P_i^{h(i,j)}$, and it can notify u_i^j that it “in charge” of the outgoing edge of P_i^j (by sending an appropriate message down the tree).

Selecting (marking) shallow subtrees (Sub-step 2). In order to emulate this sub-step it is first necessary to emulate the coloring algorithm of Cole and Vishkin [7], and Goldberg, Plotkin and Shannon [27]. The important observation is that in this algorithm, whenever a node $v(P_i^j)$ in \mathcal{F}_i sends a message to its children in \mathcal{F}_i (where a message is always of size $O(\log n)$ bits), it sends the same message. Hence, this can be emulated by simply sending (broadcasting) this message from r_i^j down the tree T_i^j . Once the message reaches the boundary nodes of P_i^j they also send it to their neighbors in G , and the message can go up the trees of the parts corresponding to the children of $v(P_i^j)$. Sending a message from $v(P_i^j)$ to its parent in \mathcal{F}_i is similar (and even simpler, since there is a single parent).

In order to emulate the marking of edges, each root r_i^j needs to gather information regarding the number of edges between P_i^j and parts $P_i^{j'}$ such that $j = h(i, j')$, for the different color classes. This information can be easily gathered by sending appropriate messages up the tree, and summing edge counts that correspond to the same color.

Deciding if to contract even or odd edges in each tree (Sub-step 3). The emulation of this part is also simple. For each tree $T \in \mathcal{T}_i$, first messages should be sent down T , so that each $v(P_i^j)$ can learn its level. Each such message from $v(P_i^j)$ to its children in T , $v(P_i^{j1}), \dots, v(P_i^{jq})$ is emulated by sending a message down T_i^j from r_i^j , and then up the trees T_i^{jp} . In a similar manner messages are sent up the tree T , summing up the weights of even and odd edges, and then the bit ‘0’ or ‘1’ is sent down T .

Contracting edges (Sub-step 4). Once a root r_i^j corresponding to a node $v(P_i^j)$ learns that the edge $(v(P_i^j), v(P_i^{h(i,j)}))$ should be contracted, it sends a message down the tree T_i^j notifying all nodes that $r_i^{h(i,j)}$ is their new root. When this message reaches u_i^j (the node in charge of the edge $(v(P_i^j), v(P_i^{h(i,j)}))$), it makes v_i^j (its neighbor in $P_i^{h(i,j)}$) its parent, and sends a message up T_i^j that each edge on the path to r_i^j should flip its orientation. The induction step for Lemma 6, follows.

Emulation cost. The total number of rounds (on G) sufficient for emulating the merging step (from \mathcal{G}_i to \mathcal{G}_{i+1}) is hence upper bounded by $O(\log^*(n))$ (by Claim 1) times the maximum diameter of parts in \mathcal{P}_i , which by Claim 4 is $\text{poly}(1/\epsilon)$.

2.2 A detailed description and analysis of Stage II

Assume Stage I completes successfully. By Claim 3, this always holds when G is planar, and by the definition of planarity, each subgraph G^j is planar. On the other hand, if G is ϵ -far from being planar, then by Claim 3, at least one subgraph G^{j*} is $(\epsilon/2)$ -far from being planar. That is, if for each $j \in [k]$ we let $m(G^j)$ denote the number of edges in G^j , then the number of edges that need to be removed from G^{j*} in order to make it planar is at least $(\epsilon/2)m(G^{j*})$.

2.2.1 Preliminary preprocessing rounds

Stage II begins with several preliminary rounds of *basic information gathering*, where we build on Corollary 5 and Lemma 6. Specifically, we use the fact that each G^j is connected, has diameter $\text{poly}(1/\epsilon)$ and has a designated root node, r^j , that is known to all nodes in G^j .

- In the first $\text{poly}(1/\epsilon)$ rounds, for each $j \in [k]$, the nodes in each G^j construct a BFS tree, rooted at r^j and denoted T_B^j . This is done simply as follows. The root r^j sends a message $(r^j, r^j, 0)$ to all its neighbors (indicating that it belongs to level 0 in the BFS tree rooted at r^j). Once a node u in G^j receives a message (r^j, v, s) from a neighbor v , u notifies v that it is v ’s child in the tree, and sends a message $(r^j, u, s + 1)$ to all its neighbors (this is of course done only once, upon receiving this first such message). When this process terminates, each node in G^j knows which of its incident edges is incident to its parent in T_B^j , which edges are incident to its children in T_B^j , and which are edges in G^j that do not belong to T_B^j (which we refer to as *non-tree edges*). Each edge in G^j is assigned to its higher-level endpoint (breaking ties by ids in the case of edges with both endpoints in the same level).
- In the next $\text{poly}(1/\epsilon)$ rounds, for each $j \in [k]$, the root r^j obtains the number of nodes $n(G^j)$ in G^j and the number of edges, $m(G^j)$. This is done simply by sending the corresponding information up the tree T_B^j . Namely, to obtain $n(G^j)$, each node sends its parent the number of nodes in its subtree (once it obtains the number of nodes in the subtrees of its children).

Similarly, to obtain $m(G^j)$, each node sends its parent the number of edges assigned to nodes in its subtree. Once $n(G^j)$ and $m(G^j)$ are computed by r^j , it can broadcast them (down T_B^j) to all nodes in G^j (in another $\text{poly}(1/\epsilon)$ rounds).

2.2.2 Planarity testing within each G^j

If $m(G^j) > 3n(G^j) - 6$, then r^j rejects. Otherwise, the distributed planar embedding algorithm of Ghaffari and Haeupler [23] is executed on G^j . Recall that if G^j is planar, then this algorithm computes for each node a circular ordering of its incident edges (known as a *combinatorial embedding*), such that there exists a planar (geometric) embedding of G^j that is *consistent* with all edge orderings. If some node v does not obtain an ordering of its incident edges in G^j (within the allotted number of rounds: $O(D(G^j) + \min(\log(n(G^j)), D(G^j)))$ where $D(G^j)$ is the diameter of G^j), then it rejects.

As noted in the introduction, it was communicated to us by one of the authors of [23] that their algorithm can be modified so as to detect if any G^j is not planar [28]. For the sake of a self-contained presentation, we rely on version of the algorithm as provided in [23] (which works under the promise that G^j is planar). Therefore, it remains to show how to verify (efficiently) that the ordering of edges incident to each node is indeed consistent with a planar embedding. To this end, we first introduce several notations and definitions.

Let $\tau = \{\tau_u\}_{u \in V}$ denote the ordering of edges in G^j (as computed by the algorithm of Ghaffari and Haeupler [23]), which is with respect to the clockwise order of the embedding, and let $\tilde{\tau}$ be the same ordering of edges as τ but with respect to the counter-clockwise order. In particular for each $u \in V$, and for every three edges $(u, v_1), (u, v_2), (u, v_3)$, $\tilde{\tau}_u$ indicates whether (u, v_2) is between (u, v_1) and (u, v_3) in counter-clockwise order.⁷

Using the ordering $\tilde{\tau}$, together with the BFS tree T_B^j , each node in G^j associates a label with each of its incident edges (and in particular those incident to its children in T_B^j). Specifically, r^j arbitrarily labels one of its incident edges e by ‘1’, and the remaining edges are labeled consecutively according to their order with respect to e . For each other node u in the tree, if the circular counter-clockwise order of its incident edges is $(u, v_1), (u, v_2), \dots, (u, v_k)$ where v_1 is its parent in the BFS tree, then u labels (u, v_j) by j , and we denote this by $\ell(u, v_j) = j$ (indeed each edge has two labels, one from each of its endpoint, so that $\ell(v_j, u)$ may differ from $\ell(u, v_j)$).⁸

This labeling of edges is then used to induce a labeling on the nodes of G^j in the natural manner: the label of a node u , denoted $\ell(u)$ is the concatenation of the labels of the edges on the path in T_B^j from r^j down to the node, where for each edge we use the label associated by the parent node. This labeling can be computed in $\text{poly}(1/\epsilon)$ rounds, by distributing the label information down the tree, starting from r^j . This labeling of nodes defines a lexicographic order on the nodes.⁹

We next introduce several notions regarding cycles and violating edges. See Figure 1a for an illustration of the notions introduced in Definitions 7 and 8.

⁷Note that for any subgraph H of G^j , the ordering $\tilde{\tau}$ defines a combinatorial embedding of H as well.

⁸The reason that we use the counter-clockwise order rather than the clockwise order is so that the edges in the tree will be in the standard, left-to-right order.

⁹That is, for two (different) nodes u and v , let $\ell(u) = \sigma_1, \dots, \sigma_p$ and $\ell(v) = \sigma'_1, \dots, \sigma'_q$, where without loss of generality, $p \leq q$ (and for $u = r^j$ we have $p = 0$). Let i be the maximum index such that $\sigma_1, \dots, \sigma_i = \sigma'_1, \dots, \sigma'_i$, where if no such index exists, then $i = 0$. If $i = p$, then $\ell(u) < \ell(v)$. Otherwise, $\ell(u) < \ell(v)$ if $\sigma_{i+1} < \sigma'_{i+1}$, and $\ell(u) > \ell(v)$ otherwise.

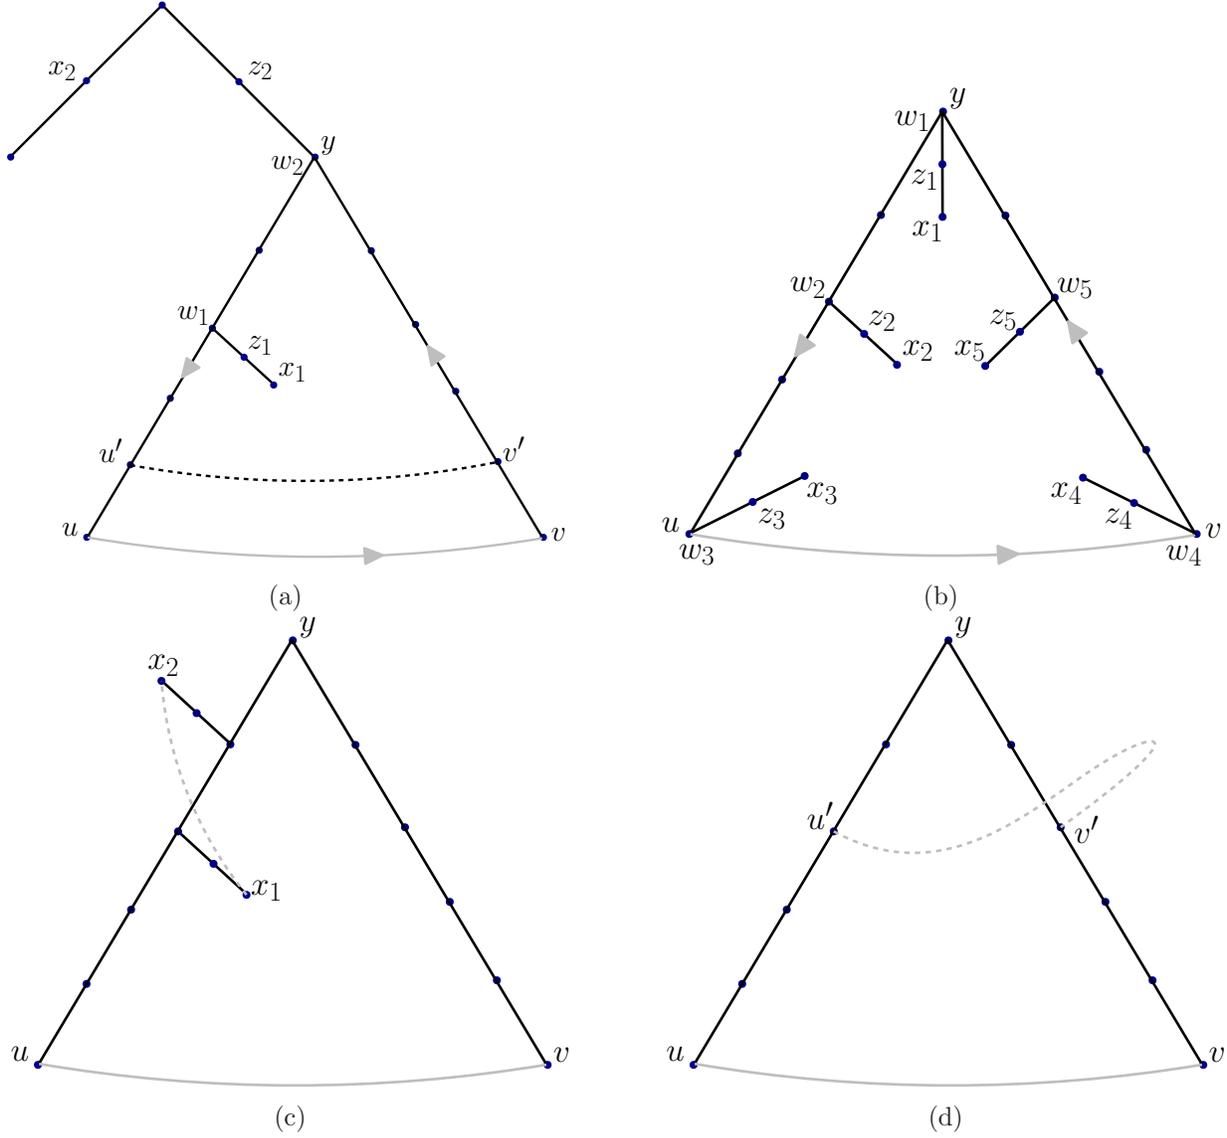


Figure 1: In all subfigures, the BFS tree edges are depicted as black edges. The non-tree edge (u, v) is depicted by a gray edge - forming the simple cycle $C(u, v)$. The vertices of $C(u, v)$ are ordered as depicted by the red arrows, i.e., $(y, \dots, w_1, \dots, u, v, \dots, y)$.

Two cases are considered in Sub-figure 1a: (1) x_1 is inside, and (2) x_2 is outside. Moreover, the edge (u', v') is also inside $C(u, v)$ w.r.t. u' and v' . The shortest paths from $C(u, v)$ to these x 's are along the BFS tree.

In Sub-figure 1b we give examples of nodes x_i that are inside $C(u, v)$: (1) x_1 satisfies: $\ell(u) < \ell(x_1) < \ell(v)$ and x_1 is not a descendant of u (similarly for x_2 and x_5); (2) x_3 satisfies: (a) $\ell(u) < \ell(x_3) < \ell(v)$, (b) it is a descendant of u and (c) $\ell(u, z_3) > \ell(u, v)$; (3) x_4 is a descendant of v and $\ell(v, z_4) < \ell(v, u)$.

In Sub-figures 1c and 1d we give examples of violating edges: (1) in Sub-figure 1c, $x_1 \in I(u, v)$ while $x_2 \in O(u, v)$, hence the edge (x_1, x_2) is a violating edge. In Sub-figure 1d the edge (u', v') is inside $C(u, v)$ w.r.t. u' but is outside $C(u, v)$ w.r.t. v' , hence it is a violating edge.

Definition 7 Let H be a subgraph of G^j , and let C be a cycle in H , where we consider a fixed ordering of the vertices on the cycle: $C = (x_1, \dots, x_k, x_1)$. For a vertex x_i on the cycle and an edge (x_i, y) , we say that (x_i, y) is *inside* C with respect to x_i (and $\tilde{\tau}$), if (x_i, y) is between (x_i, x_{i+1}) and (x_i, x_{i-1}) in the (counter-clockwise) ordering of edges incident to x_i as defined by $\tilde{\tau}$ (where for $i = 1$, $x_{i-1} = x_k$ and similarly for $i = k$, $x_{i+1} = x_1$). Otherwise, (x_i, y) is *outside* C (with respect to x_i and $\tilde{\tau}$).

Definition 8 Let (u, v) be a non-tree edge in G^j (with respect to the BFS tree T_B^j) where $\ell(u) < \ell(v)$. Let y be the least common ancestor of u and v in T_B^j , and let $C(u, v)$ be the simple cycle consisting of (in this order), the path from y to u in T_B^j , the edge (u, v) , and the path from v to y in T_B^j .

Let x be a node that does not belong to $C(u, v)$, and consider the node on C , w , for which the length of the path from u to w in T_B^j is minimized. We say that x is *inside* $C(u, v)$ if the first edge (w, z) on this path is inside $C(u, v)$ with respect to w and $\tilde{\tau}$ (as defined in Definition 7). Otherwise x is *outside* $C(u, v)$.

Observe that by Definition 8 and the definition of the labeling ℓ , a node x is inside $C(u, v)$ if and only if one of the following conditions holds (see Figure 1b).

- $\ell(u) < \ell(x) < \ell(v)$ and x is not a descendant of u ;
- $\ell(u) < \ell(x) < \ell(v)$, x is a descendant of u and $\ell(u, z) > \ell(u, v)$ where z is the ancestor of x that is a child of u .
- x is a descendant of v and $\ell(v, z) < \ell(v, u)$ where z is the ancestor of x that is a child of v .

Using Definitions 7 and 8 we define violations between edges.

Definition 9 Let (u, v) and (u', v') be two non-tree edges in G^j (with respect to the BFS tree T_B^j). We say that (u', v') is *in violation* with respect to (u, v) if one of the following conditions hold.

- u' and v' are both on $C(u, v)$, and the edge (u', v') is on different sides of $C(u, v)$ with respect to its two endpoints (see Figure 1d).
- u' is inside $C(u, v)$ and v' is outside of $C(u, v)$ (see Figure 1c).
- u' is inside (outside) of $C(u, v)$, v' is on $C(u, v)$, and the edge (u', v') is outside (respectively, inside) $C(u, v)$. (An illustration for this case is very similar to that shown in Figure 1c, and is hence omitted.)

We say that (u, v) is a **violating edge**, if there exists at least one non-tree edge (u', v') that is violating with respect to (u, v) .

Claim 10 If G^j is planar, then there are no violating edges in G^j .

Proof: Consider a planar embedding of G^j that is consistent with the ordering $\tilde{\tau}$. Recall that $\tilde{\tau}$ is the output of the algorithm of Ghaffari and Haeupler [23], and that by the correctness of the algorithm, such an embedding exists. For each non-tree edge (u, v) consider the cycle $C(u, v)$ as defined in Definition 8. Let u' and v' be two nodes in G^j . If u' and v' are both on $C(u, v)$

and there is an edge between them, then either (u', v') is inside $C(u, v)$ both with respect to u' and with respect to v' or it is outside $C(u, v)$ with respect to both. Otherwise there is no planar embedding of $C(u, v) + (u', v')$ that is consistent with $\tilde{\tau}$. Similarly, if neither u' nor v' belongs to $C(u, v)$, suppose that u' is defined to be *inside* $C(u, v)$ according to Definition 8 and v' is defined to be *outside*, or vice versa. Since the planar embedding is consistent with $\tilde{\tau}$, the nodes u' and v' must be embedded on different sides of $C(u, v)$, and hence there cannot be an edge between u' and v' . Similar arguments hold for the case that one of the two nodes is on $C(u, v)$ and the other is inside/outside $C(u, v)$. ■

Recall that T_B^j is a BFS tree defined over G^j and that violations are defined with respect to a labeling ℓ that is induced by $\tilde{\tau}$. For a subgraph H of G^j , we say that H is *connected by* T_B^j if for every two nodes u and v in H , there is a path in H consisting only of edges belonging to T_B^j . We shall use the following notation in our proof that if G^j does not contain any violating edges, then it is planar (Claim 13).

Definition 11 *Let H be a subgraph of G^j that is connected by T_B^j . For any edge (u, v) in H that does not belong to the tree T_B^j where $\ell(u) < \ell(v)$, let $I_H(u, v)$ and $O_H(u, v)$ denote the subset of nodes in H that are inside and outside $C(u, v)$, respectively (according to Definition 8).*

- *Let $E_H^I(u, v)$ denote the set of edges that are either incident to $I_H(u, v)$ or edges for which both endpoints are on $C(u, v)$ and are inside $C(u, v)$ with respect to both endpoints.*
- *Similarly, let $E_H^O(u, v)$ denote the set of edges that are either incident to $O_H(u, v)$ or edges for which both endpoints are on $C(u, v)$ and are outside $C(u, v)$ with respect to both endpoints.*
- *Let $S_H^I(u, v)$ denote the subgraph consisting of $C(u, v)$ and $E_H^I(u, v)$.*
- *Similarly, let $S_H^O(u, v)$ denote the subgraph consisting of $C(u, v)$ and $E_H^O(u, v)$.*

We will also use the following lemma. In order to prove the lemma we apply a slight variant of Lemma 7.2 and Corollary 7.1 in [15]. Since the details are very similar to those appearing in [15], they are deferred to the appendix.

Lemma 12 *Let $C(u, v)$ be a cycle in H as defined in Definition 8. If both $S_H^I(u, v)$ and $S_H^O(u, v)$ (see Definition 11) are planar and have a planar embedding that is consistent with $\tilde{\tau}$, then H is planar and has a planar embedding that is consistent with $\tilde{\tau}$ as well.*

We are now ready to prove Claim 13, stated next.

Claim 13 *If G^j does not contain any violating edges, then it is planar.*

Proof: We show that for any subgraph H of G^j that is connected by T_B^j , if there are no violating edges in H , then there exists a planar embedding of H that is consistent with $\tilde{\tau}$.

We prove this claim by induction on the pair (t, h) , where t is the number of non-tree edges in H , and h is the total number of edges. The base cases are $t = 0$ and $t = 1$ (for any h), for which it is easy to see that H is always planar. (Observe that if $h \leq 3$, then $t \leq 1$.)

For the induction step, consider a subgraph H (connected by the edges of T_B^j) with $t \geq 2$ non-tree edges and h edges. Note that for any non-tree edge (u, v) , H is the union of $S_H^I(u, v)$ and $S_H^O(u, v)$, since there are no edges between nodes in $I_H(u, v)$ and $O_H(u, v)$. We consider two cases.

The first case is that there exists a non-tree edge, (u, v) , such that both $E_I(u, v)$ and $E_O(u, v)$ are non-empty. Observe that both $S_H^I(u, v)$ and $S_H^O(u, v)$ are connected by edges of T_B^j . Since there are no violating edges in H , there are also no violating edges in each of these two subgraphs. We can therefore apply the induction hypothesis (since the number of edges in each subgraph is strictly smaller than in H), and infer that both $S_H^I(u, v)$ and $S_H^O(u, v)$ have planar embeddings, and furthermore, that each of these embeddings is consistent with $\tilde{\tau}$. By Lemma 12, the claim follows.

The second case is that there exists a non-tree edge, (u, v) , such that $E_I(u, v)$ is empty. In this case we remove (u, v) and consider a planar embedding that is consistent with $\tilde{\tau}$ of the resulting graph H' . By the induction hypothesis such an embedding of H' exists. Now, we claim that it is possible to add (u, v) to this embedding and obtain a planar embedding of H that is consistent with $\tilde{\tau}$. To verify this, observe that the inside of $C(u, v)$ with respect to $\tilde{\tau}$ is empty and therefore it is possible to add the edge (u, v) , in a manner that is consistent with $\tilde{\tau}$, without crossing any edges of H' .

If both cases do not occur, then it is implied that either there are no cycles (namely, $t = 0$) or that for all non-tree edge, (u, v) , $E_O(u, v)$ is empty. The latter implies that there is a single non-tree edge in the graph (i.e., $t = 1$). This completes the proof of the claim. ■

As a corollary of Claim 13 we get:

Corollary 14 *If G^j is γ -far from planarity, then there exist at least $\gamma \cdot m(G^j)$ violating edges in G^j .*

Given Claims 13 and 10, the algorithm proceeds as follows. First r^j broadcasts the labels of $s = \Theta(\log n/\epsilon)$ non-tree edges of G^j that are selected uniformly, independently, at random. Such a selection can be performed in $\log n \cdot \text{poly}(1/\epsilon)$ rounds. In particular, each node can decide independently for each of the non-tree edges assigned to it whether it is selected (by flipping a coin with bias $\Theta((\log n/\epsilon)/\tilde{m}^j)$ for each of these edges, where \tilde{m}^j is the total number of non-tree edges in G^j). The selected edges (i.e., pairs of node labels) are sent up the tree, where if the number of selected edges is significantly larger than the expected number, then the algorithm fails (this happens with probability $1/\text{poly}(n)$). Once r^j obtains such a sample of non-tree edges, it broadcasts the labels of these edges to all nodes in the tree. Each node in the tree can now check whether any of the non-tree edges assigned to it is in violation with any one of the sampled edges, and reject based on such a violation.

We have thus completed establishing Theorem 1.

3 A lower bound

Recall that for a fixed graph H , H is a *minor* of a graph G if H is isomorphic to a graph that can be obtained by zero or more edge contractions on a subgraph of G . We say that a graph G is *H -minor free* (or *excludes H as a minor*) if H is not a minor of G . For a family \mathcal{H} of (constant-size) graphs, we say that a graph G is *\mathcal{H} -minor free* if it is H -minor free for every $H \in \mathcal{H}$. In particular, planar graphs are $\{K_{3,3}, K_5\}$ -minor free.

In this section we establish the following theorem, which extends a result of Censor-Hillel et al. [6, Theorem 7.3] for K_3 -minor freeness (cycle-freeness).¹⁰

¹⁰To be precise, the graphs in the lower-bound construction of Censor-Hillel have a constant degree, while the

Theorem 2 *Let \mathcal{H} be a fixed family of constant-size graphs where at least one $H \in \mathcal{H}$ contains a cycle. Any distributed one-sided error algorithm for testing \mathcal{H} -minor freeness must run in $\Omega(\log n)$ rounds (for constant ϵ).*

Our proof of Theorem 2 is very similar to the proof of Theorem 7.3 in [6]. We also build on a lower bound proof of Czumaj et al. [8] for one-sided error testing of minor-freeness. Similarly to [6], we use the probabilistic method to establish that for any constant k and any number of nodes n , there exist graphs G over n nodes for which the following hold: (1) G is ϵ -far from being K_k -minor free for $\epsilon = \epsilon(k)$; (2) G contains no cycles of length $\log(n)/c$ for a sufficiently large constant $c = c(k)$. Theorem 2 directly follows by setting k to be the minimum size of $H \in \mathcal{H}$ that contains a cycle and observing that for any one-sided error algorithm that runs in less than $\log(n)/c$ rounds, when executed on G , all nodes must accept.

In order to construct such graphs, we first select a graph \tilde{G} distributed according to $\mathcal{G}(n, p)$ for $p = \Theta(1/n)$, and prove that it is far from being K_k -free with high probability. We then show that by removing a relatively small number of edges, the resulting graph, G , has no short cycles, and remains far from being K_k -free.

Claim 15 *Let \tilde{G} be a graph selected according to $\mathcal{G}(n, p)$ for $p = 1000k^2/n$. With probability $1 - 2^{-\Omega(n)}$, the graph \tilde{G} has at most $2000k^2n$ edges and is ϵ -far from K_k -minor freeness for $\epsilon = 1/(50k^2)$.*

Proof: Since the expected number of edges in \tilde{G} , denoted $m(G)$, is $p \cdot \binom{n}{2} < pn^2 = 1000k^2n$, by a multiplicative Chernoff bound, the probability that $m(G) > 2000k^2n$ is at most $e^{-pn^2/3} = 2^{-\Omega(n)}$. From this point on we condition on the event that $m(G) \leq 2000k^2n$.

We say that \tilde{G} is *well connected* if for every two disjoint subsets C_1 and C_2 of nodes such that $|C_1|, |C_2| \geq n/3k$, the number of edges with one endpoint in C_1 and the other in C_2 is greater than ϵm . We next establish the following subclaim: With probability $1 - 2^{-\Omega(n)}$, the graph \tilde{G} is well connected. For any two subsets C_1 and C_2 of nodes such that $|C_1|, |C_2| \geq n/3k$, the expected number of edges between them is at least $p(n/3k)^2 \geq 100n$. Once again by a multiplicative Chernoff bound, the probability that there are less than $50n$ edges between the two sets is at most e^{-5n} . The number of such pairs of subsets is upper bounded by 3^n , and so the probability that for some such pair of subsets there are less than $50n$ edges between them, is upper bounded by $2^{-\Omega(n)}$. Setting $\epsilon = 1/(50k^2)$, the subclaim follows. From this point we also condition on the event that \tilde{G} is well connected.

The remainder of the argument follows [8, Proof of Claim 6.2]. Consider an arbitrary partition of the nodes in \tilde{G} into k equal size subsets, U_1, \dots, U_k , and let \tilde{G}_i be the subgraph induced by U_i . We claim that each \tilde{G}_i contains a connected component of size at least $n/3k$. To verify this, let W_i^1, \dots, W_i^t be the connected components of \tilde{G}_i . Assume, contrary to the claim, that each connected component contains less than $n/3k$ nodes. But this means that there exists a subset of indices $J \subset [t]$ such that both $W_i = \bigcup_{j \in J} W_i^j$ and $W_i' = \bigcup_{j \in [t] \setminus J} W_i^j$ contain at least $n/3k$ nodes each. But since \tilde{G} is well connected, there must be an edge between some node in W_i and some node in W_i' , and we get a contradiction. We thus have, for each part U_i , a connected component, $W_i^{j(i)}$ of size at least $n/3k$. Using once again the assumption that \tilde{G} is well connected, we get that

graphs in our lower bound construction do not necessarily have a constant degree. However, we can easily modify the construction so that the graphs have a constant degree, in the same manner as in [6, Theorem 7.3].

for each pair $(W_i^{j(i)}, W_{i'}^{j(i')})$, there are more than ϵm edges between $W_i^{j(i)}$ and $W_{i'}^{j(i')}$. This implies that \tilde{G} is ϵ -far from being K_k -minor free. ■

The next claim uses the same argument as in [6, Lemma 7.7].

Claim 16 *Let \tilde{G} be as defined in Claim 15 and let G be a graph resulting from \tilde{G} by removing a single edge from each cycle in \tilde{G} whose length is less than $\log(n)/c(k)$, where $c(k) = \Theta(\log k)$. With probability at least $1/2 - 2^{-\Omega(n)}$, the graph G is ϵ -far from K_k -minor freeness for $\epsilon = 1/(100k^2)$.*

Proof: Let S be a fixed set of ℓ nodes. The probability (over the choice of G) that there is a cycle over S is at most $\ell! \cdot p^\ell$. Therefore, the expected number of cycles of length at most ℓ is upper bounded by $\binom{n}{\ell} \cdot \ell! \cdot p^\ell < (1000k^2)^\ell$, and with probability at least $1/2$ it is at most twice this number. If we set $\ell = \log(n)/\log(1000k^2)$ then by Claim 15 and a union bound over all “bad” events, Claim 16 follows. ■

4 A partitioning algorithm for minor-free graphs and applications

In what follows, when we use the term “a distributed partitioning algorithm”, we mean an algorithm that gives the following guarantee. Upon completion, there is a partition $\mathcal{P} = (P^1, \dots, P^k)$ of the nodes such that for each $j \in [k]$, the subgraph induced by P^j is connected, and there is a designated node $r^j \in P^j$ such that all nodes in P^j know the id of r^j . We first note that Stage I of our testing algorithm (described in Section 2.1) implies the next theorem.

Theorem 3 *There exists a deterministic distributed partitioning algorithm in the CONGEST model for which the following holds. For an edge-cut parameter $\epsilon \in (0, 1)$, the algorithm runs in $O(\text{poly}(1/\epsilon) \log n)$ rounds, the diameter of each part is $\text{poly}(1/\epsilon)$, and if G is minor-free, then the total number of edges between parts is at most ϵn .*

We show that by modifying the algorithm referred to in Theorem 3, we obtain a tradeoff between the round complexity and the success probability, as stated next.

Theorem 4 *There exists a distributed partitioning algorithm in the CONGEST model for which the following holds. For an edge-cut parameter $\epsilon \in (0, 1)$ and a confidence parameter $\delta \in (0, 1)$, the algorithm runs in $O(\text{poly}(1/\epsilon)(\log(1/\delta) + \log^* n))$ rounds, the diameter of each part is $\text{poly}(1/\epsilon)$, and if G is minor-free, then with probability at least $1 - \delta$, the total number of edges between parts is at most ϵn .*

Remark 1 *If one is willing to settle for constant success probability (i.e., constant δ), then the round complexity of Theorem 4 can be improved to be only $\text{poly}(1/\epsilon)$.*

In what follows we prove Theorem 4.

Recall that in the algorithm for testing planarity described in Section 2, the source of the dependence on $\log(n)$ in the round complexity was due to the forest decomposition step in each phase of the partition stage. This step served to verify that each \mathcal{G}_i has constant arboricity, as well as to allow for each node in \mathcal{G}_i to select its heaviest outgoing edge in the corresponding forest decomposition (when the arboricity is bounded as required).

If however, there is a promise that G is H -minor free, for any fixed H of constant size h , then the arboricity of every \mathcal{G}_i is upper bounded by a constant $c(h)$. Therefore, there is no need to

perform this arboricity verification step. Furthermore, as we show below, instead of selecting the heaviest outgoing edge in a forest decomposition, it suffices to select a random incident edge, where the probability to select an edge is a function of its weight.

As discussed above, the algorithm referred to in Theorem 4 is a modified version of Stage I of the planarity testing algorithm. It too runs in $\log(1/\epsilon)$ phases, where in Phase i it coarsens the partition \mathcal{P}_i and obtains the partition \mathcal{P}_{i+1} (initially, \mathcal{P}_1 is the partition into singleton parts). The first difference is that the forest decomposition step is not executed. The second difference is in the choice of an incident edge for each node in \mathcal{G}_i . We next describe how the choice of an incident edge is performed as well as the merge decision. In Section 4.1 we explain how this choice and decision are emulated on G . Once a decision to merge is made, the emulation of the merge is performed as described in Section 2.1.6.

Recall that the algorithm in Section 2.1.2, which runs on \mathcal{G}_i , consists of 4 Sub-steps. Sub-step 1 is the only sub-step that is modified: Instead of selecting the heaviest out-going edge, each node in \mathcal{G}_i randomly selects an edge as described next. The resulting graph, i.e., the graph induced on the selected edges, is guaranteed to be a directed pseudo-forest: the edge selected by each node is its only out-edge, and if an edge is selected by both endpoints then it is oriented as the out-edge of the node of lower id. We note that for Sub-steps 2a and 2b we only rely on the fact that \mathcal{F}_i is a directed pseudo-forest. In order to show that the marking process in Sub-step 2b results in a graph which is a forest we prove Claim 19. This claim is required for the correctness of Sub-step 3 of the algorithm.

Edge Selection. Let α be the arboricity of \mathcal{G}_i (which is a constant since G is minor-free). Each node in \mathcal{G}_i draws one of its incident edges with probability that is proportional to its weight. Namely, for a node $u \in \mathcal{G}_i$ and an edge (u, v) , the probability that u draws (u, v) is $\frac{w(u, v)}{w(v)}$ where $w(v) = \sum_{(y, v) \in E(\mathcal{G}_i)} w(y, v)$. This is repeated $s = \Theta(\log(1/\delta))$ times, and then each node selects the edge of maximum weight over the s trials. We call this *weighted-edge selection*.

We prove the following lemma.

Lemma 17 *With probability at least $1 - \delta$, the total weight of the edges selected in \mathcal{G}_i is at least $\frac{w(\mathcal{G}_i)}{16\alpha}$.*

Proof: Consider a forest decomposition of \mathcal{G}_i into α forests. Orient the edges from children to parents so that the out-degree of each node is at most α . Let $w_{\text{out}}(v)$ denote the weight of the out-going edges incident to v . Observe that $w(\mathcal{G}_i) = \sum_{v \in V(\mathcal{G}_i)} w_{\text{out}}(v) = \frac{1}{2} \sum_{v \in V(\mathcal{G}_i)} w(v)$. Let U denote the set of nodes, v , such that $w_{\text{out}}(v) \geq w(v)/4$. Then,

$$\sum_{v \in U} w_{\text{out}}(v) = \sum_{v \in V} w_{\text{out}}(v) - \sum_{v \notin U} w_{\text{out}}(v) \geq w(\mathcal{G}_i) - \sum_{v \notin U} w(v)/4 \geq w(\mathcal{G}_i)/2. \quad (1)$$

Let $v \in U$ and $i \in [s]$. Define $w(v, i)$ to be the weight of the edge that v drew in trial i . Then $\text{Ex}[w(v, i)] \geq w_{\text{out}}(v)/(4\alpha)$. To verify this, observe that with probability at least $1/4$, v draws one of its out-edges, and conditioned on that, with probability at least $1/\alpha$, the heaviest out-edge is drawn. By linearity of expectation and Equation (1), for each trial $i \in [s]$, $\text{Ex}[\sum_{v \in U} w(v, i)] \geq w(\mathcal{G}_i)/(8\alpha)$. We claim that for every $i \in [s]$, with probability at least $1 - 1/(16\alpha - 1)$, $\sum_{v \in U} w(v, i) \geq w(\mathcal{G}_i)/(16\alpha)$. Assume otherwise and obtain a contradiction:

$$\text{Ex} \left[\sum_{v \in U} w(v, i) \right] < \frac{w(\mathcal{G}_i)}{16\alpha - 1} + \left(1 - \frac{1}{16\alpha - 1} \right) \frac{w(\mathcal{G}_i)}{16\alpha} = \frac{w(\mathcal{G}_i)}{8\alpha}.$$

Thus, the probability that in all s trials we get that $\sum_{v \in U} w(v, i) < w(\mathcal{G}_i)/(16\alpha)$ is at most $(1/(16\alpha - 1))^s$. Since $s = \Theta(\log(1/\delta))$, we obtain that with probability at least $1 - \delta$, there exists a trial $i \in [s]$ such that $\sum_{v \in U} w(v, i) \geq w(\mathcal{G}_i)/(16\alpha)$. From the fact that $\sum_{v \in U} \max_i(w(v, i)) \geq \max_i(\sum_{v \in U} w(v, i))$, we obtain the desired result. ■

Claim 18 *With probability at least $1 - \delta$, $w(\mathcal{G}_{i+1}) \leq (1 - \frac{1}{64\alpha}) \cdot w(\mathcal{G}_i)$.*

Proof: The proof follows from Lemma 17 and the same analysis as in Claim 1. ■

Claim 19 *For an input graph that is a directed pseudo-forest, the graph resulting from the marking process in Sub-step 2b is a tree.*

Proof: We first note that given a directed pseudo-forest, the only cycle that might exist in the graph has to be a directed cycle. Assume towards contradiction that there exists a directed cycle in the marked graph. We claim that it must contain a node that is colored by 1. To verify this, observe that any vertex that is colored by 2, cannot have a marked outgoing edge and a marked incoming edge, such that for both edges, the other endpoint is colored by 3. But since every node that is colored by 1 can only have either out-going edges or in-going edges, we reach a contradiction. ■

4.1 Emulation of the weighted-edge selection

In what follows, when we say that an edge $(u, v) \in E(G)$ is incident to a part P_i^j , we mean that $u \in P_i^j$ and $v \in P_i^{j'}$ for $j' \neq j$. In order to draw an edge incident to $v(P_i^j)$ in \mathcal{G}_i with probability proportional to its weight (i.e., emulate the drawing of edges in the weighted-edge selection), we run a procedure for uniformly selecting an edge in G incident to P_i^j . If the selected edge in G is (u, v) where $v \in P_i^{j'}$, then the corresponding drawn edge in \mathcal{G}_i is $(v(P_i^j), v(P_i^{j'}))$.

This uniform selection is implemented as follows. First, each node sends a message to all its neighbors with the id of the root of its part. Following this round, each node u on the boundary of P_i^j , knows the set of incident edges (u, v) such that $v \notin P_i^j$. We denote this set by $E_{i,\text{out}}(u)$ and let $d_{i,\text{out}}(u) = |E_{i,\text{out}}(u)|$. Provided with this information, u selects, uniformly at random, one edge $e \in E_{i,\text{out}}(u)$ and sends its parent (in the tree T_i^j) the message $(e, d_{i,\text{out}}(u))$. In each consecutive round, if a node v received the messages $(e_1, d_1), \dots, (e_s, d_s)$ from its children, then it does the following. It sets $d = \sum_{p=1}^s d_p$, and selects one of the edges e_p with probability d_p/d . It then sends the message (e_p, d) to its parent. For the sake of consistency of the description, r_i^j sends messages to itself. At the end of this process (after $\text{poly}(1/\epsilon)$ rounds), r_i^j has a single edge, denoted $e_i^j = (u_i^j, v_i^j)$ for $u_i^j \in P_i^j$ and $v_i^j \in P_i^{j'}$, that is uniformly distributed among the edges incident to P_i^j .

Theorem 4 now follows from Claim 18, and the fact that the total cost of the weighted-edge selection is linear in the diameter of the parts P_i^j , which is $\text{poly}(1/\epsilon)$, times the number of repetitions which is $O(\log(1/\delta))$.

4.2 Applications of the partitioning algorithm for minor-free graphs

As a corollary of Theorems 3 and 4 we get the following.

Corollary 20 *There is a deterministic algorithm and a randomized algorithm in the CONGEST model for testing the following properties on minor-free graphs: cycle-freeness and bipartiteness. The deterministic algorithm runs in $O(\text{poly}(1/\epsilon)\log n)$ rounds. The randomized algorithm has one-sided error, runs in $O(\text{poly}(1/\epsilon)(\log(1/\delta) + \log^* n))$ rounds and has success probability $1 - \delta$.*

We note that similar statements can be derived for any hereditary property that can either be verified or (property) tested in a number of rounds that is polynomial in the diameter.

Proof: For both properties, first the algorithm of Theorem 3 or Theorem 4 is run with the edge-cut parameter set to slightly below ϵ (the distance parameter for property testing). Let $\mathcal{P} = (P^1, \dots, P^k)$ be the resulting partition, and let G^j denote the subgraph induced by P^j . By both theorems, if G is ϵ -far from having the property in question, then (with probability 1 or with probability $1 - \delta$) at least one subgraph G^j does not have the property (while if G has the property, then every G^j has the property). Therefore, it suffices to verify the property on each G^j . To this end, an algorithm for finding a BFS tree is executed on each G^j . In the case of cycle freeness, each node now checks whether it has any incident non-tree edges in G^j , and in the case of bipartiteness it checks whether there is any such edge that closes an odd-length cycle. ■

The proof of the following corollary is similar.

Corollary 21 *There is a deterministic algorithm and a randomized algorithm in the CONGEST model that, given $\epsilon \in (0, 1)$, construct an $O(\text{poly}(1/\epsilon))$ -spanner of any unweighted, minor-free graph. The spanner has $(1 + O(\epsilon))n$ edges with probability 1 in the deterministic algorithm, and with probability $1 - \delta$ in the randomized algorithm. The round complexity of the deterministic algorithm is $O(\text{poly}(1/\epsilon)\log n)$ and of the randomized algorithm is $O(\text{poly}(1/\epsilon)(\log(1/\delta) + \log^* n))$.*

Acknowledgements

We would like to thank Mohsen Ghaffari and Merav Parter for helpful information. We would also like to thank the anonymous PODC reviewers for their helpful comments.

References

- [1] Saeed Akhoondian Amiri, Stefan Schmid, and Sebastian Siebertz. A local constant factor MDS approximation for bounded genus graphs. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 227–233. ACM, 2016.
- [2] Leonid Barenboim and Michael Elkin. Sublogarithmic distributed MIS algorithm for sparse graphs using Nash-Williams decomposition. *Distributed Computing*, 22(5-6):363–379, 2010.
- [3] Itai Benjamini, Oded Schramm, and Asaf Shapira. Every minor-closed property of sparse graphs is testable. In *Proceedings of the 40th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 393–402, 2008.
- [4] John M. Boyer and Wendy J. Myrvold. On the cutting edge: simplified $O(n)$ planarity by edge addition. *Journal of Graph Algorithms and Applications*, 8(3):241–273, 2004.
- [5] Zvika Brakerski and Boaz Patt-Shamir. Distributed discovery of large near-cliques. *Distributed Computing*, 24(2):79–89, 2011.

- [6] Keren Censor-Hillel, Eldar Fischer, Gregory Schwartzman, and Yadu Vasudev. Fast distributed algorithms for testing graph properties. In *Proceedings of the 30th International Symposium on Distributed Computing (DISC)*, pages 43–56, 2016.
- [7] Richard Cole and Uzi Vishkin. Deterministic coin tossing with applications to optimal parallel list ranking. *Information and Control*, 70(1):32–53, 1986.
- [8] Artur Czumaj, Oded Goldreich, Dana Ron, C. Seshadhri, Asaf Shapira, and Christian Sohler. Finding cycles and trees in sublinear time. *Random Structures and Algorithms*, 45(2):139–184, 2014.
- [9] Artur Czumaj, Morteza Monemizadeh, Krzysztof Onak, and Christian Sohler. Planar graphs: Random walks and bipartiteness testing. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 423–432, 2011.
- [10] Andrzej Czygrinow, Michał Hańćkowiak, and Wojciech Wawrzyniak. Fast distributed approximations in planar graphs. In *Proceedings of the 22nd International Symposium on Distributed Computing (DISC)*, pages 78–92, 2008.
- [11] Talya Eden, Reut Levi, and Dana Ron. Testing bounded arboricity. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)*, pages 2081–2092, 2018.
- [12] Michael Elkin and Ofer Neiman. Efficient algorithms for constructing very sparse spanners and emulators. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)*, pages 652–669. Society for Industrial and Applied Mathematics, 2017.
- [13] Guy Even, Orr Fischer, Pierre Fraigniaud, Tzlil Gonen, Reut Levi, Moti Medina, Pedro Montealegre, Dennis Olivetti, Rotem Oshman, Ivan Rapaport, and Ioan Todinca. Three notes on distributed property testing. In *Proceedings of the 41st International Symposium on Distributed Computing (DISC)*, pages 15:1–15:30, 2017.
- [14] Guy Even, Reut Levi, and Moti Medina. Faster and simpler distributed algorithms for testing and correcting graph properties in the congest-model. *CoRR*, abs/1705.04898, 2017.
- [15] Shimon Even. *Graph Algorithms*. Cambridge University Press, 2011. 2nd edition.
- [16] Hendrik Fichtenberger, Reut Levi, Yadu Vasudev, and Maximilian Wötzel. A sublinear tester for outerplanarity (and other forbidden minors) with one-sided error. In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*, pages 52:1–52:14, 2018.
- [17] Hendrik Fichtenberger and Yadu Vasudev. A two-sided error distributed property tester for conductance. In *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018, August 27-31, 2018, Liverpool, UK*, pages 19:1–19:15, 2018.
- [18] Orr Fischer, Shay Gershtein, and Rotem Oshman. On the multiparty communication complexity of testing triangle-freeness. In *Proceedings of the 2017 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 111–120. ACM, 2017.
- [19] Orr Fischer, Tzlil Gonen, and Rotem Oshman. Distributed property testing for subgraph-freeness revisited. *CoRR*, abs/1705.04033, 2017.

- [20] Pierre Fraigniaud, Pedro Montealegre, Dennis Olivetti, Ivan Rapaport, and Ioan Todinca. Distributed subgraph detection. *CoRR*, abs/1706.03996, 2017.
- [21] Pierre Fraigniaud and Dennis Olivetti. Distributed detection of cycles. In *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures*, pages 153–162. ACM, 2017.
- [22] Pierre Fraigniaud, Ivan Rapaport, Ville Salo, and Ioan Todinca. Distributed testing of excluded subgraphs. In *Proceedings of the 30th International Symposium on Distributed Computing (DISC)*, volume 9888 of *LNCS*, pages 342–356. Springer, 2016.
- [23] Mohsen Ghaffari and Bernhard Haeupler. Distributed algorithms for planar networks I: Planar embedding. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 29–38, 2016.
- [24] Mohsen Ghaffari and Bernhard Haeupler. Distributed algorithms for planar networks II: low-congestion shortcuts, MST, and min-cut. pages 202–219, 2016.
- [25] Mohsen Ghaffari and Christiana Lymouri. Simple and near-optimal distributed coloring for sparse graphs. In *Proceedings of the 41st International Symposium on Distributed Computing (DISC)*, 2017.
- [26] Mohsen Ghaffari and Merav Parter. Near-optimal distributed DFS in planar graphs. In *Proceedings of the 41st International Symposium on Distributed Computing (DISC)*, pages 21:1–21:16, 2017.
- [27] Andrew V Goldberg, Serge A Plotkin, and Gregory E Shannon. Parallel symmetry-breaking in sparse graphs. *SIAM Journal on Discrete Mathematics*, 1(4):434, 1988.
- [28] Bernhard Haeupler. Private communication.
- [29] Bernhard Haeupler, Taisuke Izumi, and Goran Zuzic. Low-congestion shortcuts without embedding. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 451–460. ACM, 2016.
- [30] Bernhard Haeupler, Taisuke Izumi, and Goran Zuzic. Near-optimal low-congestion shortcuts on bounded parameter graphs. In *Proceedings of the 30th International Symposium on Distributed Computing (DISC)*, pages 158–172. Springer, 2016.
- [31] Bernhard Haeupler, Jason Li, and Goran Zuzic. Minor excluded network families admit fast distributed algorithms. *arXiv preprint arXiv:1801.06237*, 2018.
- [32] Avinatan Hassidim, Jonathan A. Kelner, Huy N. Nguyen, and Krzysztof Onak. Local graph partitions for approximation and testing. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 22–31, 2009.
- [33] John Hopcroft and Robert E. Tarjan. Efficient planarity testing. *Journal of the ACM*, 21(4):549–568, 1974.

- [34] Akash Kumar, C. Seshadhri, and Andrew Stolman. Finding forbidden minors in sublinear time: A $n^{1/2+o(1)}$ -query one-sided tester for minor closed properties on bounded degree graphs. In *Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 509–520, 2018.
- [35] Akash Kumar, C. Seshadhri, and Andrew Stolman. Random walks and forbidden minors II: a $\text{poly}(d \epsilon^{-1})$ -query tester for minor-closed properties of bounded degree graphs. In *Proceedings of the 51st Annual ACM Symposium on the Theory of Computing (STOC)*, pages 559–567, 2019.
- [36] Abraham Lempel, Shimon Even, and Israel Cederbaum. An algorithm for planarity testing of graphs. *Theory of Graphs, International Symposium, Rome, July 1966*, 1967.
- [37] Christoph Lenzen, Yvonne-Anne Pigolet, and Roger Wattenhofer. Distributed minimum dominating set approximations in restricted families of graphs. *Distributed computing*, 26(2):119–137, 2013.
- [38] Christoph Lenzen and Roger Wattenhofer. Minimum dominating set approximation in graphs of bounded arboricity. In *Proceedings of the 24th International Symposium on Distributed Computing (DISC)*, pages 510–524, 2010.
- [39] Reut Levi, Moti Medina, and Dana Ron. Property testing of planarity in the CONGEST model. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 347–356, 2018.
- [40] Reut Levi and Dana Ron. A quasi-polynomial time partition oracle for graphs with an excluded minor. *ACM Trans. Algorithms*, 11(3):24:1–24:13, 2015.
- [41] Nathan Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992.
- [42] David Peleg. Distributed computing. *SIAM Monographs on discrete mathematics and applications*, 5, 2000.
- [43] Yuichi Yoshida and Hiro Ito. Testing outerplanarity of bounded degree graphs. *Algorithmica*, 73(1):1–20, 2015.

A Proof of Lemma 12

In order to prove Lemma 12, we introduce several notions. These notions are very similar to those defined in [15, Chap. 7], except that there it was assumed that the graphs in question have vertex connectivity at least two. Here we do not make this assumption.

Let H be a connected graph and let C be a simple cycle in H . Consider the connected components of the graph resulting from removing the nodes in C . For each such connected component D , let $A(D)$ denote the subset of nodes on C that neighbor nodes in D . We refer to $A(D)$ as the *attachment nodes* of D on C . Let $B(D)$ denote the subgraph induced by nodes of D and $A(D)$, not including edges of C . If $|A(D)| \geq 2$, then we refer to $B(D)$ as a *bridge*, and if $|A(D)| = 1$, then it is a *half-bridge*. We also refer to edges between pairs of nodes on C as bridges. Two bridges B and B' are said to *interlace* if one of the following holds:

1. There are two attachments of B , x and y , and two attachments of B' , w and z , such that all four are distinct, and appear on C in the order (say, clockwise) x, w, y, z .
2. There are three attachments common to B and B' . That is, $|A(B) \cap A(B')| \geq 3$.

The next lemma is a slight modification of Lemma 7.2 in [15] (where here the graph H in question is not assumed to have vertex connectivity at least two).

Lemma 22 *Let B_1, \dots, B_s be the set of bridges and half-bridges of a graph H with respect to a simple cycle C . Suppose that $C + B_i$ is planar for every $1 \leq i \leq s$ and that no two bridges in the set interlace. Then $C + B_1 + \dots + B_s$ can be embedded in the plane so that all the bridges and half-bridges are inside C .*

Proof: We prove the claim by induction on the number of nodes in $C + B_1 + \dots + B_s$. The base case is three nodes (there is just a cycle C and no bridges). For the induction step, as in the proof of Lemma 7.2 in [15], since no two bridges interlace, there must be at least one bridge, B_i for which the following holds. If we consider the attachments of B_i in clockwise order, a_1, \dots, a_k (for $k \geq 2$), then there is no other bridge B_j with attachments (strictly) after a_1 and before a_k on C . ■

As a corollary of Lemma 22 we obtain:

Corollary 23 *Let B_1, \dots, B_s be the set of bridges and half-bridges of a graph H with respect to a simple cycle C . Suppose that $C + B_i$ is planar for every $1 \leq i \leq s$ and that the set of bridges can be partitioned into two subsets, such that within each subset no two bridges interlace. Then H is planar.*

Building on Corollary 23 we are now ready to prove Lemma 12.

Proof of Lemma 12: Consider the bridges of H with respect to $C(u, v)$. They can be partitioned into two pairwise noninterlacing subsets: one corresponding the bridges inside $C(u, v)$ and one corresponding to bridges outside $C(u, v)$ (recall that if the bridges interlace then we cannot embed them in one side of $C(u, v)$). Therefore H is planar. It remains to show that there exists a planar embedding of H which is consistent with $\tilde{\tau}$. For each vertex x in $C(u, v)$, we have the circular order of its edges in the planar embedding of $S_H^I(u, v)$ and $S_H^O(u, v)$, respectively. We can merge this pair of orders into a single order that is consistent with $\tilde{\tau}$ simply by concatenating them (the only edges that these orders have in common are the edges on $C(u, v)$, which are the first and last edges in both orders). Other vertices of H are either in $S_H^I(u, v)$ or $S_H^O(u, v)$ (but not in both), so the ordering remains consistent. □