

# The Arboricity Captures the Complexity of Sampling Edges

Talya Eden\*      Dana Ron†      Will Rosenbaum‡

February 22, 2019

## Abstract

In this paper, we revisit the problem of sampling edges in an unknown graph  $G = (V, E)$  from a distribution that is (pointwise) almost uniform over  $E$ . We consider the case where there is some a priori upper bound on the arboricity of  $G$ . Given query access to a graph  $G$  over  $n$  vertices and of average degree  $d$  and arboricity at most  $\alpha$ , we design an algorithm that performs  $O\left(\frac{\alpha}{d} \cdot \frac{\log^3 n}{\varepsilon}\right)$  queries in expectation and returns an edge in the graph such that *every* edge  $e \in E$  is sampled with probability  $(1 \pm \varepsilon)/m$ . The algorithm performs two types of queries: degree queries and neighbor queries. We show that the upper bound is tight (up to poly-logarithmic factors and the dependence in  $\varepsilon$ ), as  $\Omega\left(\frac{\alpha}{d}\right)$  queries are necessary for the easier task of sampling edges from any distribution over  $E$  that is close to uniform in total variational distance. We also prove that even if  $G$  is a tree (i.e.,  $\alpha = 1$  so that  $\frac{\alpha}{d} = \Theta(1)$ ),  $\Omega\left(\frac{\log n}{\log \log n}\right)$  queries are necessary to sample an edge from any distribution that is pointwise close to uniform, thus establishing that a  $\text{poly}(\log n)$  factor is necessary for constant  $\alpha$ . Finally we show how our algorithm can be applied to obtain a new result on approximately counting subgraphs, based on the recent work of Assadi, Kapralov, and Khanna (ITCS, 2019).

---

\*School of Electrical Engineering, Tel Aviv University, Tel Aviv, Israel

†School of Electrical Engineering, Tel Aviv University, Tel Aviv, Israel

‡Max Planck Institute for Informatics, Saarbrücken, Germany

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Problem definition . . . . .	1
1.2	Results . . . . .	2
1.3	Discussion of the results . . . . .	3
1.4	An application to approximately counting subgraphs . . . . .	4
1.5	A high-level presentation of the algorithm and lower bounds . . . . .	4
1.5.1	The algorithm . . . . .	4
1.5.2	The lower bounds . . . . .	6
1.6	Related work . . . . .	7
1.7	Organization . . . . .	8
<b>2</b>	<b>The Algorithm</b>	<b>8</b>
2.1	Decomposing graphs of bounded arboricity . . . . .	8
2.2	Algorithm description . . . . .	9
<b>3</b>	<b>General Lower Bound</b>	<b>12</b>
<b>4</b>	<b>Lower Bound for Trees</b>	<b>13</b>
4.1	Details of how $\mathcal{P}$ answers queries and constructs $\widehat{T} \in \mathcal{T}$ . . . . .	14
4.1.1	The partially labeled trees $T_t$ and $T'_t$ . . . . .	14
4.1.2	Answering queries . . . . .	15
4.1.3	The final (fully-labeled) tree $\widehat{T} \in \mathcal{T}$ . . . . .	17
4.2	Completing the analysis . . . . .	17

# 1 Introduction

Let  $G = (V, E)$  be a graph over  $n$  vertices and  $m$  edges. We consider the problem of sampling an edge in  $G$  from a pointwise almost uniform distribution over  $E$ . That is, for each edge  $e \in E$ , the probability that  $e$  is returned is  $(1 \pm \varepsilon)/m$ , where  $\varepsilon$  is a given approximation parameter. An algorithm for performing this task has random access to the vertex set  $V = \{1, \dots, n\}$  and can perform queries to  $G$ . The allowed queries are (1) *degree queries* denoted  $\text{deg}(v)$  (what is the degree,  $d(v)$ , of a given vertex  $v$ ) and (2) *neighbor queries* denoted  $\text{nbr}(v, i)$  (what is the  $i^{\text{th}}$  neighbor of  $v$ ).<sup>1</sup> We refer to this model as *the uniform vertex sampling model*.

Sampling edges almost uniformly is a very basic sampling task. In particular it gives the power to sample vertices with probability approximately proportional to their degree, which is a useful primitive. Furthermore, there are sublinear algorithms that are known to work when given access to uniform edges (e.g., [2]) and can be adapted to the case when the distribution over the edges is almost uniform (see Section 1.4 for details). An important observation is that in many cases it is crucial that the sampling distribution is pointwise-close to uniform rather than close with respect to the Total Variation Distance (henceforth TVD) – see the discussion in [15, Sec. 1.1].

Eden and Rosenbaum [15] recently showed that  $\Theta^*(\sqrt{m}/d)$  queries are both sufficient and necessary for sampling edges almost uniformly, where  $d = 2m/n$  denotes the average degree in the graph. (We use the notation  $O^*$  to suppress factors that are polylogarithmic in  $n$  and polynomial in  $1/\varepsilon$ .) The instances for which the task is difficult (i.e., for which  $\Omega(\sqrt{m}/d)$  queries are necessary), are characterized by having very dense subgraphs, i.e., a subgraph with average degree  $\Theta(\sqrt{m})$ . Hence, a natural question is whether it is possible to achieve lower query complexity when some a priori bound on the density of subgraphs is known. A well studied measure for bounded density (“everywhere”) is the graph *arboricity* (see Definition 1.2 below). Indeed there are many natural families of graphs that have bounded arboricity such as graphs of bounded degree, bounded treewidth or bounded genus, planar graphs, graphs that exclude a fixed minor and many other graphs. In the context of social networks, preferential attachment graphs and additional generative models exhibit bounded arboricity [3, 6, 5], and this has also been empirically validated for many real-world graphs [18, 16, 24].

We describe a new algorithm for sampling edges almost uniformly whose runtime is  $O^*(\alpha/d)$  where  $\alpha$  is an upper bound on the arboricity of  $G$ . In the extremal case that  $\alpha = \Theta(\sqrt{m})$ , the runtime of our algorithm is the same as that of [15] (up to poly-log factors). For smaller  $\alpha$ , our algorithm is strictly faster. In particular for  $\alpha = O(1)$ , the new algorithm is *exponentially* faster than that of [15]. We also prove matching lower bounds, showing that for all ranges of  $\alpha$ , our algorithm is query-optimal, up to polylogarithmic factors and the dependence in  $1/\varepsilon$ .

Furthermore, while not as simple as the algorithm of [15], our algorithm is still easy to implement and does not incur any large constants in the query complexity and running time, thus making it suitable for practical applications.

## 1.1 Problem definition

In order to state our results precisely, we define the notion of pointwise-closeness of probability distributions (cf. [15]) and arboricity of a graph.

**Definition 1.1.** *Let  $D$  be a fixed probability distribution on a finite set  $X$ . We say that a probability distribution  $\widehat{D}$  is **pointwise  $\varepsilon$ -close to  $D$**  if for all  $x \in X$ ,*

$$|\widehat{D}(x) - D(x)| \leq \varepsilon D(x), \quad \text{or equivalently} \quad 1 - \varepsilon \leq \frac{\widehat{D}(x)}{D(x)} \leq 1 + \varepsilon.$$

---

<sup>1</sup>If  $i > d(v)$  then a special symbol, e.g.  $\perp$ , is returned.

If  $D = U$ , the uniform distribution on  $X$ , then we say that  $\widehat{D}$  is **pointwise  $\varepsilon$ -close to uniform**.

For the sake of conciseness, from this point on, unless explicitly stated otherwise, when we refer to an edge sampling algorithm, we mean an algorithm that returns edges according to a distribution that is pointwise-close to uniform.

**Definition 1.2.** Let  $G = (V, E)$  be an undirected graph. A forest  $F = (V_F, E_F)$  (i.e., a graph containing no cycles) with vertex set  $V_F = V$  and edge set  $E_F \subseteq E$  is a **spanning forest** of  $G$ . We say that a family of spanning forests  $F_1, F_2, \dots, F_k$  **covers**  $G$  if  $E = \bigcup_{i=1}^k E_{F_i}$ . The **arboricity** of  $G$ , denoted  $\alpha(G)$ , is the minimum  $k$  such that there exists a family of spanning forests of size  $k$  that covers  $G$ .

An edge-sampling algorithm is given as input an approximation parameter  $\varepsilon \in (0, 1)$  and a parameter  $\alpha$  which is an upper bound on the arboricity of  $G$ . The algorithm is required to sample edges according a distribution that is pointwise  $\varepsilon$ -close to uniform. To this end the algorithm is given query access to  $G$ . In particular we consider the aforementioned uniform vertex sampling model.

## 1.2 Results

We prove almost matching upper and lower bounds on the query complexity of sampling an edge according to a distribution that is pointwise-close to uniform when an upper bound on the arboricity of the graph is known. The first lower bound stated below (Theorem 2) holds even for the easier task of sampling from a distribution that is close to uniform in TVD.

**Theorem 1.** *There exists an algorithm  $\mathcal{A}$  that for any  $n, m, \alpha$ , and graph  $G = (V, E)$  with  $n$  nodes,  $m$  edges, and arboricity at most  $\alpha$ , satisfies the following. Given  $n$  and  $\alpha$ ,  $\mathcal{A}$  returns an edge  $e \in E$  sampled from a distribution  $\widehat{U}$  that is pointwise  $\varepsilon$ -close to uniform using  $O\left(\frac{\alpha}{d} \cdot \frac{\log^3 n}{\varepsilon}\right)$  degree and neighbor queries in expectation.*

In Section 1.5.1 we provide a high-level presentation of the algorithm referred to in Theorem 1 and shortly discuss how it differs from the algorithm in [15] (for the case that there is no given upper bound on the arboricity).

We next state our lower bound, which matches the upper bound in Theorem 1 up to a polylogarithmic dependence on  $n$  (for constant  $\varepsilon$ ).

**Theorem 2.** *Fix  $\varepsilon \leq 1/6$  and let  $n, m$  and  $\alpha$  be parameters such that  $\alpha = \alpha(n) \leq \sqrt{m}$  and  $m \leq n\alpha$ . Let  $\mathcal{G}_{n,m}^\alpha$  be the family of graphs with  $n$  vertices,  $m$  edges and arboricity at most  $\alpha$ . Then any algorithm  $\mathcal{A}$  that for any  $G \in \mathcal{G}_{n,m}^\alpha$  samples edges in  $G$  from a distribution that is  $\varepsilon$ -close to uniform in total variation distance—and in particular, any distribution that is pointwise  $\varepsilon$ -close to uniform—requires  $\Omega(\alpha/d)$  queries in expectation.*

When  $\alpha$  is a constant, then (assuming that  $m = \Omega(n)$ ) the lower bound in Theorem 2 is simply  $\Omega(1)$ , while Theorem 1 gives an upper bound of  $O(\log^3 n)$  (for constant  $\varepsilon$ ). We prove that an almost linear dependence on  $\log n$  is necessary, even for the case that  $\alpha = 1$  (where the graph is a tree).

**Theorem 3.** *Fix  $\varepsilon \leq 1/6$ , and let  $\mathcal{T}_n$  be the family of trees on  $n$  vertices. Then any algorithm  $\mathcal{A}$  that for any  $G \in \mathcal{T}_n$  samples edges in  $G$  from a distribution that is pointwise  $\varepsilon$ -close to uniform requires  $\Omega\left(\frac{\log n}{\log \log n}\right)$  queries in expectation.*

We note that both of our lower bounds also hold when the algorithm is also given access to pair queries  $\text{pair}(u, v)$  (is there an edge between  $u$  and  $v$ ).

### 1.3 Discussion of the results

**The arboricity captures the complexity of sampling edges.** The simplest algorithm for sampling edges uniformly is based on rejection sampling. Namely, it repeats the following until an edge is output: Sample a uniform vertex  $u$ , flip a coin with bias  $d(u)/d_{\max}$ , where  $d_{\max}$  the maximum degree, and if the outcome is HEADS, then output a random edge  $(u, v)$  incident to  $u$ . The expected complexity of rejection sampling grows like  $d_{\max}/d$ , that is, linearly with the maximum degree. As noted earlier, Eden and Rosenbaum [10] show that this dependence on the maximum degree is not necessary (for approximate sampling), as  $O^*(\sqrt{m}/d)$  queries and time always suffice, even when the maximum degree is not bounded (e.g., is  $\Theta(n)$ ). However, if the maximum degree is bounded, and in particular if  $d_{\max} = o(\sqrt{m})$ , then rejection sampling has lower complexity than the [10] algorithm (e.g., in the case that the graph is close to regular, so that  $d_{\max} = O(d)$ , we get complexity  $O(1)$ ).

Since the arboricity of a graph is both upper bounded by  $d_{\max}$  and by  $\sqrt{m}$ , our algorithm can be viewed as “enjoying both worlds”. Furthermore, our results can be viewed as showing that the appropriate complexity measure for sampling edges is not the maximum degree but rather the maximum average degree (recall that the arboricity  $\alpha$  measures the maximum density of any subgraph of  $G$  – for a precise statement, see Theorem 4).

**Approximately counting the number of edges in bounded-arboricity graphs.** As shown by Eden, Ron and Seshadhri [10],  $\Theta^*(\alpha/d)$  is also the complexity of estimating the number of edges in a graph when given a bound  $\alpha$  on the arboricity of the graph. Their algorithm improves on the previous known bound of  $O^*(\sqrt{m}/d)$  by Feige [17] and Goldreich and Ron [19] (when the arboricity is  $o(\sqrt{m})$ ). However, other than the complexity, our algorithm for sampling edges and the algorithm of [10] for estimating the number of edges do not share any similarities, in particular, as the result of [10] is allowed to “ignore” an  $\varepsilon$ -fraction of the graph edges.

Furthermore, while the complexity of sampling and of approximate counting of edges are the same (up to  $\log n$  and  $1/\varepsilon$  dependencies), we have preliminary results showing that for other subgraphs this is not necessarily the case. Specifically, there exist graphs with constant arboricity for which estimating the number of triangles can be done using  $O^*(1)$  queries, but pointwise-close to uniform sampling requires  $\Omega(n^{1/4})$  queries in expectation.

**On the necessity of being provided with an upper bound on the arboricity.** While our algorithm does not require to be given any bound on the average degree  $d$ , it must be provided with an upper bound  $\alpha$  on the arboricity of the given graph. To see why this is true, consider the following two graphs. The first graph consists of a perfect matching between its vertices, so that both its average degree and its arboricity are 1. For  $\tilde{\alpha} > 1$ , the second graph consists of a perfect matching over  $n - n/\tilde{\alpha}$  vertices and a  $\tilde{\alpha}$ -regular graph over the remaining  $n/\tilde{\alpha}$  vertices. This graph has an average degree of roughly 2, and arboricity  $\tilde{\alpha}$ . If an edge-sampling algorithm is not provided with an appropriate upper bound on the arboricity, but is still required to run in (expected) time that grows like the ratio between the arboricity and the average degree, then it means it can be used to distinguish between the two graphs. However, assuming a random labeling of the vertices of the two graphs, this cannot be done in time  $o(\tilde{\alpha})$ .

**Pointwise closeness vs. closeness with respect to the TVD.** The lower bound of Theorem 2 holds for sampling from a distribution that is close to uniform with respect to TVD, and *a fortiori* to sampling from pointwise almost uniform distributions. In contrast, the lower bound of Theorem 3 does not apply to sampling edges from a distribution that is  $\varepsilon$ -close to

uniform in TVD. Indeed, a simple rejection sampling procedure (essentially ignoring all nodes with degrees greater than  $1/\varepsilon$ ) can sample edges from a distribution that is  $\varepsilon$ -close to uniform in TVD using  $O(1/\varepsilon)$  queries in expectation. Thus, Theorem 3 gives a separation between the tasks of sampling from distributions that are pointwise-close to uniform versus close to uniform in TVD. The general upper and lower bounds of Theorems 1 and 2 show that the separation between the complexity of these tasks can be at most poly-logarithmic for any graph.

## 1.4 An application to approximately counting subgraphs

In a recent paper [2], Assadi, Kapralov, and Khanna made significant progress on the question of counting arbitrary subgraphs in a graph in sublinear time. Specifically, they provide an algorithm that estimates the number of occurrences of any arbitrary subgraph  $H$  in  $G$ , denoted by  $\#H$ , to within a  $(1 \pm \varepsilon)$ -approximation with high probability. The running time of their algorithm is  $O^*\left(\frac{m^{\rho(H)}}{\#H}\right)$ , where  $\rho(H)$  is the fractional edge cover of  $H$ .<sup>2</sup> Their algorithm assumes access to uniform edge samples in addition to degree, neighbor and pair queries. As noted in [2], their algorithm can be adapted to work with edge samples that are pointwise  $\varepsilon$ -close to uniform (where this is not true for edge samples that are only  $\varepsilon$ -close to uniform in TVD—e.g., when all the occurrences of  $H$  are induced by an  $\varepsilon$ -fraction of the edges). Invoking the algorithm of [2], and replacing each edge sample with an invocation of `Sample-edge` results in the following corollary.

**Corollary 1.** *Let  $G$  be a graph  $G = (V, E)$  with  $n$  nodes,  $m$  edges, and arboricity at most  $\alpha$ . There exists an algorithm that, given  $n, \alpha, \varepsilon \in (0, 1)$ , a subgraph  $H$  and query access to  $G$ , returns a  $(1 \pm \varepsilon)$  approximation of the number of occurrences of  $H$  in  $G$ , denoted  $\#H$ . The expected query complexity and running time of the algorithm are*

$$O^*\left(\min\left\{m, \frac{n\alpha \cdot m^{\rho(H)-1}}{\#H}\right\}\right) \quad \text{and} \quad O^*\left(\frac{n\alpha \cdot m^{\rho(H)-1}}{\#H}\right),$$

*respectively, where  $\rho(H)$  denotes the fractional edge cover of  $H$ , and the allowed queries are degree, neighbor and pair queries.*

Thus, by combining our result with [2], we extend the known results for approximately counting the number of subgraphs in a graph in the uniform vertex sampling model. Furthermore, for graphs in which  $m = \Theta(n\alpha)$ , we obtain the same query complexity and running time of [2] without the assumption that the algorithm has access to uniform edge samples.

## 1.5 A high-level presentation of the algorithm and lower bounds

While our results concern *undirected* graphs  $G = (V, E)$ , it will be helpful to view each edge  $\{u, v\} \in E$  as a pair of *ordered* edges  $(u, v)$  and  $(v, u)$ .

### 1.5.1 The algorithm

Sampling an (ordered) edge (almost) uniformly is equivalent to sampling a vertex with probability (almost) proportional to its degree. Hence we focus on the latter task. A single iterations of the algorithm we describe either returns a vertex or outputs FAIL. We show that the probability that it outputs FAIL is not too large, and that conditioned on the algorithm returning

---

<sup>2</sup>The fractional edge cover of a graph  $H = (V_H, E_H)$  is a mapping  $\psi : E_H \rightarrow [0, 1]$  such that for each vertex  $a \in V_H$ ,  $\sum_{e \in E_H, a \in e} \psi(e) \geq 1$ . The fractional edge-cover number  $\rho(H)$  of  $H$  is the minimum value of  $\sum_{e \in E_H} \psi(e)$  among all fractional edge covers  $\psi$ .

a vertex, each vertex  $v$  is returned with probability proportional to its degree up to a factor of  $(1 \pm \varepsilon)$ .

Our starting point is a structural decomposition result for graphs with bounded arboricity (Lemma 2.4). Our decomposition defines a partition of the graph's vertices into *levels*  $L_0, L_1, \dots, L_\ell$ . For parameters  $\theta$  and  $\beta$ ,  $L_0$  consists of all vertices with degree at most  $\theta$ , and for  $i > 0$ , level  $L_i$  contains all vertices  $v$  that do not belong to previous levels  $L_0, \dots, L_{i-1}$ , but have at least  $(1 - \beta)d(v)$  neighbors in these levels. We prove that for any graph with arboricity at most  $\alpha$ , for  $\theta = \Theta(\alpha \log n / \varepsilon)$  and  $\beta = \Theta(\varepsilon / \log n)$ , there exists such a partition into layers with  $\ell = O(\log n)$  levels. We stress that the algorithm does not actually construct such a partition, but rather we use the partition in our analysis of the algorithm.<sup>3</sup>

In order to gain intuition about the algorithm and its analysis, suppose that all vertices in  $L_0$  have degree exactly  $\theta$ , and that all edges in the graph are between vertices in consecutive layers. Consider the following *random walk* algorithm. The algorithm first selects an index  $j \in [0, \ell]$  uniformly at random. It then selects a vertex  $u_0$  uniformly at random. If  $u_0 \in L_0$ , then it performs a random walk of length  $j$  starting from  $u_0$  (otherwise it outputs FAIL). If the walk did not pass through any vertex in  $L_0$  (with the exception of the starting vertex  $u_0$ ), then the algorithm returns the final vertex reached.

First observe that for every  $u \in L_0$ , the probability that  $u$  is returned is  $\frac{1}{\ell+1} \cdot \frac{1}{n}$  (the probability that the algorithm selected  $j = 0$  and selected  $u$  as  $u_0$ ). This equals  $\frac{d(u)}{(\ell+1) \cdot \theta \cdot n}$  (by our assumption that  $d(u) = \theta$  for every  $u \in L_0$ ). Now consider a vertex  $v \in L_1$ . The probability that  $v$  is returned is at least  $\frac{1}{\ell+1} \cdot \frac{(1-\beta)d(v)}{n} \cdot \frac{1}{\theta} = \frac{(1-\beta)d(v)}{(\ell+1) \cdot \theta \cdot n}$  (the probability that the algorithm selected  $j = 1$ , then selected one of  $v$ 's neighbors  $u \in L_0$ , and finally selected to take the edge between  $u$  and  $v$ ). In general, our analysis shows that for every  $i$  and every  $v \in L_i$ , the probability that  $v$  is returned is at least  $\frac{(1-\beta)^i d(v)}{(\ell+1) \cdot \theta \cdot n}$ . On the other hand, we show that for every vertex  $v$ , the probability that  $v$  is returned is at most  $\frac{d(v)}{(\ell+1) \cdot \theta \cdot n}$ . By the choice of  $\theta$  and  $\beta$  we get that each vertex  $v$  is returned with probability in the range  $[(1 - \varepsilon)d(v)\rho(\varepsilon, n), d(v)\rho(\varepsilon, n)]$  for  $\rho(\varepsilon, n) = \Theta(\varepsilon / (\alpha n \log^2 n))$ . By repeating the aforementioned random-walk process until a vertex is returned— $\Theta\left(\frac{\alpha n}{m} \cdot \frac{\log^2 n}{\varepsilon}\right) = \Theta\left(\frac{\alpha}{d} \cdot \frac{\log^2 n}{\varepsilon}\right)$  times in expectation—we obtain a vertex that is sampled with probability proportional to its degree, up to  $(1 \pm \varepsilon)$ .

We circumvent the assumption that  $d(u) = \theta$  for every  $u \in L_0$  by rejection sampling: In the first step, if the algorithm samples  $u_0 \in L_0$ , then it continues with probability  $d(u)/\theta$  and fails otherwise. The assumption that all edges are between consecutive levels is not necessary for the analysis described above to hold. The crucial element in the analysis is that for every vertex  $v \notin L_i$ , where  $i > 0$ , at least  $(1 - \beta)$  of the neighbors of  $v$  belong to  $L_0, \dots, L_{i-1}$ . This allows us to apply the inductive argument for the lower bound on the probability that  $v$  is returned when we average over all choices of  $j$  (the number of steps in the random walk). For precise details of the algorithm and its analysis, see Section 2.

**On the relation to [15].** We briefly discuss the relation between our algorithm for bounded-arboricity graphs, which we denote by  $\mathcal{A}_{ba}$  and the algorithm presented in [15] (for the case that no upper bound is given on the arboricity), which we denote by  $\mathcal{A}_{ua}$ . The algorithm  $\mathcal{A}_{ua}$  can be viewed as considering a partition of the graph vertices into just two layers according to a degree threshold of roughly  $\sqrt{m}$ . It performs a random walk similarly to  $\mathcal{A}_{ba}$ , but where the walk has either length 0 or 1. This difference in the number of layers and the length of

<sup>3</sup>This decomposition is related to the forest decomposition of Barenboim and Elkin [4]. The main difference, which is essential for our analysis, is that the partition we define is based on the number of neighbors that a vertex has in lower levels *relative* to its degree, while in [4] the partition is based on the absolute number of neighbors to higher levels.

the walk, is not only quantitative. Rather, it allows  $\mathcal{A}_{ua}$  to determine to which layer does a vertex belong simply according to its degree. This is not possible in the case of  $\mathcal{A}_{ba}$  (with the exception of vertices in  $L_0$ ). Nonetheless, despite the apparent “blindness” of  $\mathcal{A}_{ba}$  to the layers it traverses in the random walk, we can show the following: Choosing the length of the random walk uniformly at random and halting in case the walk returns to  $L_0$ , ensures that each vertex is output with probability approximately proportional to its degree.

### 1.5.2 The lower bounds

**The lower bound of  $\Omega\left(\frac{\alpha}{d}\right)$  for general  $\alpha$ .** In order to prove Theorem 2, we employ the method of [14] (which builds upon the paradigm introduced in [7]) based on communication complexity. The idea of the proof is to reduce from the two-party communication complexity problem of computing the disjointness function. The reduction is such that (1) any algorithm that samples edges from an almost-uniform distribution reveals the value of the disjointness function with sufficiently large probability, and (2) every allowable query can be simulated in the two-party communication setting using little communication.

**The lower bound of  $\Omega\left(\frac{\log n}{\log \log n}\right)$  for  $\alpha = 1$ .** As opposed to the proof of the lower bound for general  $\alpha$ , in the case of  $\alpha = 1$  we did not find a way to employ the communication-complexity method (which tends to result in compact and “clean” proofs). Instead, we design a direct, albeit somewhat involved, proof from first-principles.

Specifically, in order to prove Theorem 3, we consider a complete tree in which each internal vertex has degree  $\log n$  (so that its depth is  $\Theta\left(\frac{\log n}{\log \log n}\right)$ ). We then consider the family of graphs that correspond to all possible labelings of such a tree. As noted in Section 1.5.1, sampling edges almost uniformly is equivalent to sampling vertices with probability approximately proportional to their degree. In particular, in our construction, the label of the root should be returned with probability approximately  $\log n/n$ . We show that any algorithm that succeeds in returning the label of the root of the tree with the required probability must perform  $\Omega\left(\frac{\log n}{\log \log n}\right)$  queries.

To this end we define a *process*  $\mathcal{P}$  that interacts with any algorithm  $\mathcal{A}$ , answering  $\mathcal{A}$ ’s queries while constructing a uniform random labeling of the vertices and edges in the tree. The vertices of the tree are assigned random labels in  $[n]$ , and for each vertex  $v$  in the tree, its incident edges are assigned random labels in  $[d(v)]$ . We say that  $\mathcal{A}$  *succeeds*, if after the interaction ends,  $\mathcal{A}$  outputs the label of the root of the tree, as assigned by  $\mathcal{P}$ .

Let  $L = \frac{\log n}{C \cdot \log \log n}$  be the lower bound we would like to prove, where  $C$  is a sufficiently large constant (so that in particular,  $L$  is a (small) constant fraction of the depth of the tree). Intuitively,  $\mathcal{A}$  would like to “hit” a vertex at depth at most  $L$  and then “walk up the tree” to the root. There are two sources of uncertainty for  $\mathcal{A}$ . One is whether it actually hits a vertex at depth at most  $L$ , and the second is which edges should be taken to go up the tree. Our lower bound argument mainly exploits the second uncertainty, as we sketch next.

The process  $\mathcal{P}$  starts with an unlabeled tree (of the aforementioned structure), and assigns labels to its vertices and edges in the course of its interaction with  $\mathcal{A}$ . Recall that  $\mathcal{P}$  answers the queries of  $\mathcal{A}$  while constructing a uniform labeling. Therefore, whenever  $\mathcal{A}$  asks a query involving a *new* label (i.e., that has not yet appeared in its queries or answers to them), the vertex to which this label is assigned, should be uniformly selected among all vertices that are not yet labeled. We shall say that a vertex is *critical* if its depth is at most  $L$ . As long as no critical vertex is hit,  $\mathcal{A}$  cannot reach the root. This implies that if no critical vertex is hit in the course of its queries, then the probability that  $\mathcal{A}$  succeeds is  $O(1/n)$ .

While the probability of hitting a critical vertex is relatively small, it is not sufficiently small to be deemed negligible. However, suppose that  $\mathcal{A}$  hits a critical vertex  $u$  at depth  $\Delta < L$ ,

which occurs with probability  $(\log n)^\Delta/n$ . Then, conditioned on this event, each of the  $(\log n)^\Delta$  edge-labeled paths from  $u$  is equally likely to lead to the root, and the labels of vertices on these paths are uniformly distributed, thus intuitively conveying no information regarding the “right path”.

A subtlety that arises when formalizing this argument is the following. Suppose that in addition to hitting a critical vertex  $u$ ,  $\mathcal{A}$  hits another vertex,  $v$ , which is not necessarily critical, but is at distance less than  $L$  from  $u$  (and in particular has depth at most  $2L$ , which we refer to as *shallow*). Then, a path starting from  $v$  might meet a path starting from  $u$ , hence adding a conditioning that makes the above argument (regarding uniform labelings) imprecise. We address this issue by upper bounding the probability of such an event (i.e., of hitting both a critical vertex and a shallow vertex), and accounting for an event of this type as a success of  $\mathcal{A}$ .

## 1.6 Related work

Some of the works presented below were already mentioned earlier in the introduction, but are provided in this subsection for the sake of completeness.

The work most closely related to the present work is the recent paper of Eden and Rosenbaum [15]. In [15], the authors proved matching upper and lower bounds of  $\Theta^*(\sqrt{m}/d)$  for the problem of sampling an edge from an almost uniform distribution in an arbitrary graph using degree, neighbor, and pair queries.

The problem of sampling edges in a graph is closely related to the problem of estimating  $m$ , the number of edges in the graph. In [17], Feige proved an upper bound of  $O^*(\sqrt{m}/d)$  for obtaining a  $(2 + \varepsilon)$ -factor multiplicative approximation of  $m$  using only degree queries,<sup>4</sup> and shows that it is not possible to go below a factor of 2 with a sublinear number of degree queries. In [19], Goldreich and Ron showed that  $\Theta^*(\sqrt{m}/d)$  queries are necessary and sufficient to obtain a  $(1 + \varepsilon)$ -factor approximation of  $m$  if neighbor queries are also allowed.

Several works prove matching upper and lower bounds on the query complexity of counting the number of triangles [8], cliques [12], and star-graphs of a given size [20] using degree, neighbor, and pair queries (when the latter are necessary). Eden Ron and Seshadhri devised algorithms for estimating the number of  $k$ -cliques [11] and moments of the degree distribution [10] whose runtimes are parameterized by the arboricity  $\alpha$  of the input graph (assuming a suitable upper bound for  $\alpha$  is given to the algorithm as input). These algorithms outperform the lower bounds of [12] and [20] (respectively) in the case where  $\alpha \ll \sqrt{m}$ . In [9], Eden, Levi, and Ron described an efficient algorithm for distinguishing graphs with arboricity at most  $\alpha$  from those that are far from any graph with arboricity  $3\alpha$ .

Two recent works [1, 2] consider a query model that allows uniform random edge sampling in addition to degree, neighbor, and pair queries. In this model, Aliakbarpour et al. [1] described an algorithm for estimating the number of star subgraphs. In the same model, Assadi et al. [2] devised an algorithm that relies on uniform edge samples as a basic query to approximately count the number of instances of an arbitrary subgraph in a graph. The results in [1] and [2] imply that uniform edge samples afford the model strictly more power: the sample complexity of the algorithm of [1] outperforms the lower bound of [20] for the same task, and the sample complexity of the algorithm of [2] outperforms the lower bound of [12] for estimating the number of cliques. (The results of [20, 12] are in the uniform vertex sampling model.)

---

<sup>4</sup>To be precise, Feige [17] shows that, given a lower bound  $d_0$  on the average degree,  $O(\sqrt{n/d_0}/\varepsilon)$  degree queries are sufficient. If such a lower bound is not provided to the algorithm, then a geometric search can be performed, as shown in [19].

## 1.7 Organization

The rest of the paper is organized as follows. We describe and analyze our main algorithm, thereby proving Theorem 1, in Section 2. The lower bounds of Theorems 2 and 3 are proven in Sections 3 and 4, respectively.

## 2 The Algorithm

In this section we describe an algorithm that samples an edge  $e$  from an arbitrary graph  $G$  with arboricity at most  $\alpha$ , according to a pointwise almost uniform distribution. Theorem 1 follows from our analysis of the algorithm. In what follows, for integers  $i \leq j$ , we use  $[i, j]$  to denote the set of integers  $\{i, \dots, j\}$ , and for a vertex  $v$  we let  $\Gamma(v)$  denote its set of neighbors.

As noted in the introduction, sampling edges from a uniform distribution is equivalent to sampling vertices proportional to their degrees. Indeed, if each vertex  $v$  is sampled with probability  $d(v)/2m$ , then choosing a random neighbor  $w \in \Gamma(v)$  uniformly at random returns the (directed) edge  $e = (v, w)$  with probability  $1/2m$ . Thus, it suffices to sample each vertex  $v \in V$  with probability (approximately) proportional to its degree.

### 2.1 Decomposing graphs of bounded arboricity

Before describing the algorithm, we describe a decomposition of a graph  $G$  into *layers* depending on its arboricity. We begin by recalling the following characterization of arboricity due to Nash-Williams [22].

**Theorem 4** (Nash-Williams [22]). *Let  $G = (V, E)$  be a graph. For a subgraph  $H$  of  $G$ , let  $n_H$  and  $m_H$  denote the number of vertices and edges, respectively, in  $H$ . Then*

$$\alpha(G) = \max_H \{ \lceil m_H / (n_H - 1) \rceil \}, \quad (1)$$

where the maximum is taken over all subgraphs  $H$  of  $G$ .

Another folklore result is the connection between the arboricity of a graph and its degeneracy. The degeneracy of a graph, denoted  $\delta$ , is the minimum value such that for every subgraph  $H \subseteq G$ , every vertex  $v \in H$  has degree at least  $\delta$  (in  $H$ ).

**Corollary 2.1.** *For every graph  $G$ ,  $\delta \leq 2\alpha(G)$ , where  $\delta$  is the degeneracy of  $G$ .*

**Definition 2.2.** *Let  $G = (V, E)$  be a graph, and  $\theta \in \mathbf{N}$ ,  $\beta \in (0, 1)$  parameters. We define a  $(\theta, \beta)$ -**layering** in  $G$  to be the sequence of non-empty disjoint subsets  $L_0, L_1, \dots, L_\ell \subseteq V$  defined by*

$$L_0 = \{v \in V \mid d(v) \leq \theta\} \quad (2)$$

and for  $i \geq 1$ ,

$$L_{i+1} = \{v \notin L_0 \cup L_1 \cup \dots \cup L_i \mid |\Gamma(v) \cap (L_0 \cup \dots \cup L_i)| \geq (1 - \beta)d(v)\}.$$

That is,  $L_0$  consists of all vertices of degree at most  $\theta$ , and a vertex  $v$  is in  $L_{i+1}$  if  $i$  is the smallest index for which a  $(1 - \beta)$ -fraction of  $v$ 's neighbors resides in  $L_0 \cup L_1 \cup \dots \cup L_i$ . We say that  $G$  admits a  $(\theta, \beta)$ -**layered partition of depth  $\ell$**  if we have  $V = L_0 \cup L_1 \cup \dots \cup L_\ell$ .

**Notation 2.3.** *For a fixed  $i$ , we denote  $L_{\leq i} = L_0 \cup L_1 \cup \dots \cup L_i$ , and similarly for  $L_{< i}$ ,  $L_{\geq i}$ , and  $L_{> i}$ . We use the notation  $d_i(v)$  to denote  $|\Gamma(v) \cap L_i|$  and similarly for  $d_{\leq i}(v)$  and  $d_{\geq i}(v)$ .*

**Lemma 2.4.** *Suppose  $G$  is a graph with arboricity at most  $\alpha$ . Then  $G$  admits a  $(\theta, \beta)$ -layered partition of depth  $\ell$  for  $\theta = 4\alpha \lceil \log n \rceil / \varepsilon$ ,  $\beta = \varepsilon / 2 \lceil \log n \rceil$ , and  $\ell \leq \lceil \log n \rceil$ .*

*Proof.* For each  $i$ , let  $W_i = V \setminus (L_0 \cup L_1 \cup \dots \cup L_{i-1})$  be the set of vertices not in levels  $0, 1, \dots, i-1$ . Let  $m(W_i)$  denote the number of edges in the subgraph of  $G$  induced by  $W_i$ . For any fixed  $i$  and  $v \in W_{i+1}$ , we have  $d_{<i}(v) < (1 - \beta)d(v)$  because  $v \notin L_{\leq i}$ . Therefore,  $v$  has at least  $\beta d(v) > \beta\theta$  neighbors in  $W_i$ . Summing over vertices  $v \in W_{i+1}$  gives

$$m(W_i) = \frac{1}{2} \sum_{v \in W_i} d_{\geq i}(v) \geq \frac{1}{2} \sum_{v \in W_{i+1}} d_{\geq i}(v) > \frac{1}{2} |W_{i+1}| \cdot \beta\theta. \quad (3)$$

On the other hand, since  $G$  has arboricity at most  $\alpha$ , Theorem 4 implies that

$$m(W_i) \leq \alpha |W_i|. \quad (4)$$

Combining Equations (3) and (4), we find that

$$\frac{|W_{i+1}|}{|W_i|} \leq \frac{2\alpha}{\beta\theta} = \frac{1}{2},$$

where the inequality is by the choice of  $\beta$  and  $\theta$ . Therefore,  $\ell \leq \lceil \log n \rceil$ , as required.

To see that  $V = L_0 \cup L_1 \cup \dots \cup L_\ell$  (i.e.,  $W_{\ell+1} = \emptyset$ ) suppose to the contrary that there exists  $v \in W_{\ell+1}$ . Then every  $v \in W_{\ell+1}$  has at least  $\beta d(v) \geq \beta\theta = 4\alpha$  neighbors in  $W_\ell$ . By Corollary 2.1, this implies that  $\alpha(G) \geq 2\alpha$ , which is a contradiction.  $\square$

## 2.2 Algorithm description

The algorithm exploits the structure of graphs  $G$  with arboricity at most  $\alpha$  described in Lemma 2.4. More precisely, as the algorithm does not have direct access to this structure, the structure is used explicitly only in the analysis of the algorithm. Let  $L_0, L_1, \dots, L_\ell$  be a  $(\theta, \beta)$ -layered partition of  $V$  with  $\theta = 4\alpha \lceil \log n \rceil / \varepsilon$ ,  $\beta = \varepsilon / \log n$ , and  $\ell = \log n$ . Vertices  $v \in L_0$  are sampled with probability exactly proportional to their degree using a simple rejection sampling procedure, `Sample-a-leaf`( $G, \theta$ ). In order to sample vertices in layers  $L_i$  for  $i > 0$ , our algorithm performs a random walk starting from a random vertex in  $L_0$  chosen with probability proportional to its degree. Specifically, the algorithm `Sample-edge`( $G, \alpha$ ) chooses a length  $j$  to the random walk uniformly in  $[0, \ell]$ . The subroutine `Random-walk`( $G, \theta, j$ ) performs the random walk for  $j$  steps, or until a vertex  $v \in L_0$  is reached in some step  $i > 0$ . If the walk returns to  $L_0$ , the subroutine aborts and does not return any vertex. (This behavior ensures that samples are not too biased towards vertices in lower layers.) Otherwise, `Random-walk` returns the vertex at which the random walk halts. Our analysis shows that the probability that the random walk terminates at any vertex  $v \in V$  is approximately proportional to  $d(v)$  (Corollary 2.8), although `Sample-edge` may fail to return any edge with significant probability. Finally, we repeat `Sample-edge` until it successfully returns a vertex.

**Definition 2.5.** *We let  $P_j[v]$  denote the probability that `Random-walk` returns  $v$ , when invoked with parameters  $G, \theta$  and  $j \in [0, \ell]$ . We also let  $P_{\leq j}[v] \stackrel{\text{def}}{=} \sum_{i=0}^j P_i[v]$  and similarly for  $P_{\geq j}$ .*

**Lemma 2.6.** *Let  $\ell$  be as set in Step 1 of `Sample-edge` and let  $P_{\leq j}$  be as defined in Definition 2.5. For all  $v \in V$ ,  $P_{\leq \ell}[v] \leq \frac{d(v)}{n\theta}$ .*

*Proof.* We argue by induction on  $j$  that for any  $j \in [0, \ell]$ ,  $P_{\leq j}[v] \leq d(v)/n\theta$ . For the case  $j = 0$ , it is immediate from the description of `Random-walk` and `Sample-a-leaf` that  $P_0[v] = d(v)/(n\theta)$

**Sample-edge**( $G, \alpha, \varepsilon$ )

1. Let  $\theta = 4\alpha \lceil \log n \rceil / \varepsilon$  and let  $\ell = \lceil \log n \rceil$ .
2. Choose a number  $j \in [0, \ell]$  uniformly at random.
3. Invoke **Random-walk**( $G, \theta, j$ ) and let  $v$  be the returned vertex if one was returned. Otherwise, **return FAIL**.
4. Sample a uniform neighbor  $w$  of  $v$  and **return**  $e = (v, w)$ .

**Random-walk**( $G, \theta, j$ )

1. Invoke **Sample-a-leaf**( $\theta$ ) and let  $v_0$  be the returned vertex if one was returned. Otherwise, **return FAIL**.
2. For  $i = 1$  to  $j$  do
  - (a) Sample a random neighbor  $v_i$  of  $v_{i-1}$ .
  - (b) If  $v_i \in L_0$  then **return FAIL**.
3. **Return**  $v_j$ .

**Sample-a-leaf**( $G, \theta$ )

1. Sample a vertex  $u \in V$  uniformly at random and query for its degree.
2. If  $d(u) > \theta$  **return FAIL**.
3. **Return**  $u$  with probability  $\frac{d(u)}{\theta}$ , and with probability  $1 - \frac{d(u)}{\theta}$  **return FAIL**.

if  $v \in L_0$  and  $P_0[v] = 0$  otherwise. Further, for  $v \in L_0$ , due to Step 2b,  $P_i(v) = 0$  for all  $i > 0$ , so that the lemma holds for all  $v \in L_0$ . Now suppose that for all  $v \in V$  we have  $P_{\leq j-1}(v) \leq d(v)/n\theta$ . Then for any fixed  $v \notin L_0$  we compute

$$\begin{aligned} P_{\leq j}[v] &= \sum_{i=1}^j P_i[v] = \sum_{i=1}^j \sum_{u \in \Gamma(v)} P_{i-1}[u] \frac{1}{d(u)} = \sum_{u \in \Gamma(v)} \frac{1}{d(u)} \sum_{i=0}^{j-1} P_i[u] \\ &= \sum_{u \in \Gamma(v)} \frac{1}{d(u)} P_{\leq j-1}[u] \leq \sum_{u \in \Gamma(v)} \frac{1}{d(u)} \frac{d(u)}{n\theta} = \frac{d(v)}{n\theta}. \end{aligned}$$

The second equality holds by the definition of **Random-walk**, and the one before the last inequality holds by the inductive hypothesis.  $\square$

**Lemma 2.7.** *Let  $\ell$  be as set in Step 1 of **Sample-edge**. For every  $j \in [\ell]$ ,  $v \in L_j$  and  $k \in [j, \ell]$ , we have  $P_{\leq k}[v] \geq \frac{(1-\beta)^j d(v)}{n\theta}$ .*

*Proof.* We prove the claim by induction on  $j$ . For  $j = 0$  and  $k = 0$ , by the description of **Random-walk** and **Sample-a-leaf**, for every  $v \in L_0$ ,

$$P_0[v] = \frac{d(v)}{n\theta}. \quad (5)$$

For  $j = 0$  and  $0 < k \leq \ell$ ,

$$P_{\leq k}[v] = \sum_{i=0}^k P_i[v] = P_0[v] + \sum_{i=1}^k P_i[v] = \frac{d(v)}{n\theta}, \quad (6)$$

where the last equality is due to Step 2b in Random-walk.

For  $j = 1$  and  $1 \leq k \leq \ell$ , for every  $v \in L_1$ , according to Step 2b in the procedure Random-walk,  $P_0[v] = 0$ . Also, for every  $u \notin L_0$ ,  $P_0[u] = 0$ , since by Step 2 in Sample-a-leaf it always holds that  $v_0$  is in  $L_0$ . Therefore,

$$P_1[v] = \sum_{u \in \Gamma(v)} P_0[u] \cdot \frac{1}{d(u)} = \sum_{u \in \Gamma(v) \cap L_0} P_0[u] \cdot \frac{1}{d(u)} = \sum_{u \in \Gamma(v) \cap L_0} \frac{d(u)}{n\theta} \cdot \frac{1}{d(u)} = \frac{d_0(v)}{n\theta}, \quad (7)$$

where the second to last inequality is by Equation (5). By the definition of  $L_1$ , for every  $v \in L_1$ ,  $d_0(v) \geq (1 - \beta)d(v)$ , and it follows that  $P_{\leq k}[v] \geq P_1[v] \geq (1 - \beta)d(v)/(n\theta)$ .

We now assume that the claim holds for all  $i \leq j - 1$  and  $k \in [i, \ell]$ , and prove that it holds for  $j$  and for every  $k \in [j, \ell]$ . By the induction hypothesis and the definition of  $L_j$ , for any  $v \in L_j$  we have

$$\begin{aligned} P_{\leq k}[v] \geq P_{\leq j}[v] &= \sum_{u \in \Gamma(v)} P_{\leq j-1}[u] \cdot \frac{1}{d(u)} \geq \sum_{i=0}^{j-1} \sum_{u \in \Gamma(v) \cap L_i} P_{\leq j-1}[u] \cdot \frac{1}{d(u)} \\ &\geq \sum_{i=0}^{j-1} \sum_{u \in \Gamma(v) \cap L_i} \frac{(1 - \beta)^i d(u)}{n\theta} \cdot \frac{1}{d(u)} \geq \frac{(1 - \beta)^{j-1} d_{\leq j-1}(v)}{n\theta} \geq \frac{(1 - \beta)^j d(v)}{n\theta}. \end{aligned}$$

Hence, the claim holds for every  $j \in [\ell]$  for every  $k \in [j, \ell]$ .  $\square$

**Corollary 2.8.** *For any graph  $G$  with arboricity at most  $\alpha$ , the procedure Sample-edge when invoked with  $G$ ,  $\alpha$  and  $\varepsilon$ , returns each edge in the graph with probability in the range  $\left[\frac{1 - \varepsilon/2}{\rho}, \frac{1}{\rho}\right]$  for  $\rho = n\theta(\ell + 1)$ ,  $\theta = 4\alpha \lceil \log n \rceil / \varepsilon$  and  $\ell = \lceil \log n \rceil$ .*

*Proof.* Consider a specific edge  $e^* = (v^*, w^*)$ , and let  $i$  be the index such that  $v^* \in L_i$ . By the description of the procedure Sample-edge, the procedure Random-walk is invoked with an index  $j$  that is chosen uniformly in  $[0, \ell]$ . Hence, the probability that  $v^*$  is returned by Random-walk in Step 3 is

$$\Pr[v = v^*] = \frac{1}{\ell + 1} \sum_{j=0}^{\ell} P_j[v] = \frac{1}{\ell + 1} P_{\leq \ell}[v].$$

By Lemma 2.6,  $P_{\leq \ell}[v] \leq \frac{d(v)}{n\theta}$ , and by Lemma 2.7,  $P_{\leq \ell}[v^*] \geq \frac{(1 - \beta)^\ell d(v^*)}{n\theta}$ , where the probability is over the random coins of the procedures Sample-edge and Random-walk. Hence,

$$\Pr[v = v^*] \in [(1 - \beta)^\ell, 1] \cdot \frac{d(v^*)}{n\theta(\ell + 1)},$$

implying that for  $\rho = n\theta(\ell + 1)$ ,

$$\Pr[(v^*, w^*) \text{ is the returned edge}] \in [(1 - \beta)^\ell, 1] \cdot \frac{1}{n\theta(\ell + 1)} \in \left[\frac{1 - \varepsilon/2}{\rho}, \frac{1}{\rho}\right], \quad (8)$$

where the last inequality is by the setting of  $\beta = \varepsilon/2 \lceil \log n \rceil$ .  $\square$

*Proof of Theorem 1.* Consider the algorithm that repeatedly calls `Sample-edge`( $G, \alpha$ ) until an edge  $e$  is successfully returned. For a single invocation of `Sample-edge` and fixed edge  $e$  let  $A_e$  denote the event that `Sample-edge` returns  $e$ . By Corollary 2.8 we have that  $\Pr[A_e] \geq (1 - \varepsilon)/n\theta(\ell + 1)$ . Further, for any edge  $e' \neq e$  the events  $A_e$  and  $A_{e'}$  are disjoint, so we bound

$$\Pr[\text{Sample-edge returns an edge}] = \Pr\left[\bigcup_{e \in E} A_e\right] = \sum_{e \in E} \Pr[A_e] \geq \frac{(1 - \varepsilon)m}{n\theta(\ell + 1)}.$$

The expected number of iterations until `Sample-edge` succeeds is the reciprocal of this probability, so

$$\mathbf{E}[\# \text{ invocations until success}] \leq \frac{n\theta(\ell + 1)}{(1 - \varepsilon)m} = O\left(\frac{n\alpha}{m\varepsilon} \cdot \log^2 n\right).$$

Since each invocation of `Sample-edge` uses  $O(\log n)$  queries, the expected number of queries before an edge is returned is  $O\left(\frac{n\alpha}{\varepsilon m} \cdot \log^3 n\right)$ .

Finally, when conditioned on a successful invocation of `Sample-edge`, Corollary 2.8 implies that for any  $e, f \in E$  the probabilities  $p_e, p_f$  of returning  $e$  and  $f$ , respectively, satisfy

$$1 - \varepsilon/2 \leq \frac{p_e}{p_f} \leq \frac{1}{1 - \varepsilon/2} \leq 1 + \varepsilon.$$

Therefore, the induced distribution  $P$  over edges returned by a successful invocation of `Sample-edge` is pointwise  $\varepsilon$  close to uniform, which gives the desired result.  $\square$

### 3 General Lower Bound

In this section, we prove Theorem 2, thereby showing that the upper bound implied by the algorithm in Section 2 is tight for all values of  $\alpha$ , up to a poly-logarithmic factor in  $n$ . The lower bound is robust, as it applies to the easier problem of sampling edges from a distribution that is almost uniform with respect to TVD.

In order to prove a general query lower bound for edge sampling (Theorem 2), we employ the method of [14] based on communication complexity. The proof is a natural generalization of Theorem C.3 in the full version [13]. The idea of the proof is to construct an embedding of the two-party disjointness function into  $\mathcal{G}_n^\alpha$  in such a way that (1) any algorithm that samples edges from an almost-uniform distribution reveals the value of the disjointness function with sufficiently large probability, and (2) every allowable query can be simulated in the two-party communication setting using little communication. The following definitions and theorem are taken directly from [14].

**Definition 3.1.** Let  $\mathcal{X} \subseteq \{0, 1\}^N \times \{0, 1\}^N$ . Suppose  $f : \mathcal{X} \rightarrow \{0, 1\}$  is an arbitrary (partial) function, and let  $g$  be a Boolean function on  $\mathcal{G}_n^\alpha$ . Let  $\mathcal{E} : \{0, 1\}^N \times \{0, 1\}^N \rightarrow \mathcal{G}_n^\alpha$ . We call the pair  $(\mathcal{E}, g)$  an **embedding** of  $f$  if for all  $(x, y) \in \mathcal{X}$  we have  $f(x, y) = g(\mathcal{E}(x, y))$ .

**Definition 3.2.** Let  $q : \mathcal{G}_n^\alpha \rightarrow \{0, 1\}^*$  be a query and  $(\mathcal{E}, g)$  an embedding of  $f$ . We say that  $q$  has **communication cost** at most  $B$  and write  $\text{cost}_\mathcal{E}(q) \leq B$  if there exists a (zero-error) communication protocol  $\Pi_q$  such that for all  $(x, y) \in P$  we have  $\Pi_q(x, y) = q(\mathcal{E}(x, y))$  and  $|\Pi_q(x, y)| \leq B$ .

**Theorem 5.** Let  $\mathcal{Q}$  be a set of allowable queries,  $f : \mathcal{X} \rightarrow \{0, 1\}$ , and  $(\mathcal{E}, g)$  an embedding of  $f$ . Suppose that each query  $q \in \mathcal{Q}$  has communication cost  $\text{cost}_\mathcal{E}(q) \leq B$ . Suppose  $\mathcal{A}$  is an algorithm that computes  $g$  using  $T$  queries (in expectation) from  $\mathcal{Q}$ . Then the expected query complexity of  $\mathcal{A}$  is  $T = \Omega(R(f)/B)$ .

We prove Theorem 2 by applying Theorem 5, and taking  $f$  to be the *disjointness function* defined by  $\text{disj}(x, y) = 1$  if  $\sum_{i=1}^N x_i y_i = 0$  and  $\text{disj}(x, y) = 0$  otherwise, where  $x, y \in \{0, 1\}^N$ . The following fundamental result gives a lower bound on the communication complexity of  $\text{disj}$ .

**Theorem 6** ([21, 23]). *The randomized communication complexity of the disjointness function is  $R(\text{disj}) = \Omega(N)$ . This result holds even if  $x$  and  $y$  are promised to satisfy  $\sum_{i=1}^N x_i y_i \in \{0, 1\}$ —that is, Alice’s and Bob’s inputs are either disjoint or intersect on a single point.*

*Proof of Theorem 2.* We employ the method of reduction from communication of the disjointness function described above. We first describe an embedding  $\mathcal{E}$  of the disjointness function  $\text{disj}$ . Fix an arbitrary graph  $H$  on  $n'$  vertices with arboricity  $\alpha < n'$  and  $m'$  edges. Let  $K$  be an arbitrary graph with arboricity at most  $\alpha$ ,  $2m'$  edges, and  $2m'/\alpha$  vertices (for example, taking  $K$  to be an  $\alpha$ -regular graph on  $2m'/\alpha$  vertices). We set  $n = 2n'$  and  $m = m'$  or  $3m'$  depending on the input of the disjointness function.

Let  $N = n'\alpha/2m'$ . Given  $x, y \in \{0, 1\}^N$ , we define  $G(x, y) = (V, E)$  as follows. We partition  $V = W_0 \cup W_1 \cup \dots \cup W_N$  where  $|W_0| = n'$  and  $|W_i| = 2m'\alpha/n'$  for every  $i > 0$ . The edge set  $E$  is determined as follows:  $W_0$  is an isomorphic copy of  $H$ ; for  $i \geq 1$ ,  $W_i$  is isomorphic to  $K$  if  $x_i = y_i = 1$ , and is a set of isolated vertices otherwise. We define the (partial) function  $g : \mathcal{G}_n^\alpha \rightarrow \{0, 1\}$  defined by  $g(G) = 1$  if and only if every  $W_i$  is a set of isolated vertices. It is straightforward to verify that the pair  $(\mathcal{E}, g)$  where  $\mathcal{E}(x, y) = G(x, y)$  is an embedding of  $\text{disj}$  in the sense of Definition 3.1. Further, each degree, neighbor, and pair query to  $G(x, y)$  can be simulated by Alice and Bob using  $O(1)$  communication. Therefore, by Theorem 5 and the communication lower bound for  $\text{disj}$  (Theorem 6), any algorithm that distinguishes  $g(G) = 1$  and  $g(G) = 0$  with probability bounded away from  $1/2$  requires  $\Omega(N) = \Omega(n\alpha/m)$  queries in expectation.

Finally, we show that an algorithm  $\mathcal{A}$  that samples edges from a distribution that is  $\varepsilon$ -close to uniform in TVD can be used to compute  $g$  on the image of  $\mathcal{E}$ . The key observation is that when  $\text{disj}(x, y) = 0$ , a  $2/3$  fraction of  $G(x, y)$ ’s edges lie in some  $W_i$  for  $i \geq 1$ , while if  $\text{disj}(x, y) = 1$ , all of  $G$ ’s edges are in  $W_0$ . Since  $\mathcal{A}$  samples edges from a  $\varepsilon$ -close to uniform distribution for  $\varepsilon \leq 1/6$ , in the case  $\text{disj}(x, y) = 0$ ,  $\mathcal{A}$  must return an edge in some  $W_i$  for  $i \geq 1$  with probability at least  $2/3 - 1/6 = 1/2$ .

Consider the following algorithm: run  $\mathcal{A}$  twice independently to sample edges  $e_1$  and  $e_2$  from  $G = G(x, y)$  using  $2q$  queries in expectation. If either  $e_i$  is not in  $W_0$ , output  $\text{disj}(x, y) = 0$ , otherwise output  $\text{disj}(x, y) = 1$ . The algorithm is always correct on instances where  $\text{disj}(x, y) = 1$ . On instances where  $\text{disj}(x, y) = 0$ , it only fails in the case where  $e_1$  and  $e_2$  both lie in  $W_0$ . By the computation in the previous paragraph, this event occurs with probability at most  $1/2^2 = 1/4$ . Therefore, the expected runtime of  $\mathcal{A}$  satisfies  $2q = \Omega(n\alpha/m)$ , as desired.  $\square$

## 4 Lower Bound for Trees

In this section we prove Theorem 3, which asserts a query lower bound of  $\Omega\left(\frac{\log n}{\log \log n}\right)$  for any algorithm  $\mathcal{A}$  that samples edges in a tree from a pointwise almost uniform distribution.

In order to prove Theorem 3 we consider a family of labeled trees  $\mathcal{T}$ , where all the trees in the family have the same topology, and differ only in their vertex and edge labeling. The underlying tree  $T$  is defined as follows: For  $k = \log n$ , each internal vertex in the tree is incident to  $k$  edges, so that the root  $r$  has  $k$  children, and every other internal vertex has  $k - 1$  children. Hence, the depth of the tree is  $D = \Theta\left(\frac{\log n}{\log \log n}\right)$ . In each labeled tree belonging to  $\mathcal{T}$ , the vertices are given pairwise distinct labels in  $[n]$ . Similarly, for each non-leaf vertex  $u \in V$  the (ordered) edges  $(u, v)$  incident with  $u$  are given pairwise distinct labels in  $[k]$ . We take  $\mathcal{T}$  to be the set

of all possible labelings constructed in this way. Since the trees differ only by their labeling, we will sometimes refer to  $\mathcal{T}$  as a family of *tree labelings*.

Again we note that sampling edges from a pointwise almost uniform distribution is equivalent to sampling vertices with probability (approximately) proportional to their degree. We prove Theorem 3 by showing that any algorithm that returns the root  $r$  of  $T$  with probability  $\Omega\left(\frac{k}{n}\right) = \Omega\left(\frac{\log n}{n}\right)$  requires  $\Omega\left(\frac{\log n}{\log \log n}\right)$  queries. More formally, we will prove the following proposition.

**Proposition 4.1.** *There exists a constant  $C > 1$  such that for any algorithm  $\mathcal{A}$  using at most  $L = \frac{\log n}{C \cdot \log \log n}$  queries, we have*

$$\Pr[\mathcal{A} \text{ returns } r \text{ on input } \widehat{T}] \leq \frac{\log n}{n \log \log n},$$

where  $\widehat{T}$  is chosen according to the uniform distribution in  $\mathcal{T}$ .

**Remark 4.2.** *One may wonder if we can increase the lower bound  $L$  in Proposition 4.1 by setting the upper bound on the probability of returning  $r$  in the proposition to  $\frac{\log n}{C' \log \log n}$  for a constant  $C' > 1$ . Indeed we can obtain a lower bound of  $L = \frac{\log n}{C \cdot W(\log n)}$  where  $W(\cdot)$  is the Lambert  $W$  function (which in particular satisfies  $W(x) = \frac{\ln x}{W(x)}$  so that it is asymptotically slightly smaller than  $\log(\cdot)$ ). However we believe the difference is negligible and hence not worth the slightly less natural construction.*

Let  $\mathcal{A}$  be any algorithm that returns a vertex  $v$  in its input graph. We describe a process  $\mathcal{P}$  that interacts with  $\mathcal{A}$  and answers its queries while constructing a uniformly random labeled tree  $\widehat{T}$  in  $\mathcal{T}$  on the fly. We say that  $\mathcal{A}$  *succeeds*, if  $\mathcal{A}$  outputs the label of the root of  $\widehat{T}$ . The algorithm  $\mathcal{A}$  is allowed  $Q \leq L$  queries. For the sake of the analysis, if certain ‘‘bad’’ events occur,  $\mathcal{P}$  will reveal the entire labeling  $\widehat{T}$  to  $\mathcal{A}$ . Conditioned on a bad event, we simply bound the probability that  $\mathcal{A}$  succeeds by 1.

## 4.1 Details of how $\mathcal{P}$ answers queries and constructs $\widehat{T} \in \mathcal{T}$

### 4.1.1 The partially labeled trees $T_t$ and $T'_t$

The process  $\mathcal{P}$  starts with an initially unlabeled tree, denoted  $T$  (where each internal vertex has degree  $k = \log n$  and its depth is  $D = \Theta\left(\frac{\log n}{\log \log n}\right)$ ). We denote the set of vertices of  $T$  by  $V$  and the set of ordered edges by  $E$ . We also let  $\Delta(u)$  denote the depth of vertex  $u$  in  $T$ .

Following each query of  $\mathcal{A}$ ,  $\mathcal{P}$  labels some of the vertices/edges in  $T$ . After it answers the last query of  $\mathcal{A}$ ,  $\mathcal{P}$  determines all labels of yet unlabeled vertices and edges in  $T$ , thus obtaining the final labeled tree  $\widehat{T}$ . The labeling decisions made by  $\mathcal{P}$  ensure that  $\widehat{T}$  is uniformly distributed in  $\mathcal{T}$  for any algorithm  $\mathcal{A}$ .

We denote the partially labeled tree that  $\mathcal{P}$  holds after the first  $t$  queries  $q_1, \dots, q_t$  of  $\mathcal{A}$ , by  $T_t$ . Thus,  $T_t = (V, E, \pi_t, \tau_t)$ , where  $\pi_t : V \rightarrow [n] \cup \{\perp\}$  is the partial labeling function of the vertices and  $\tau_t : E \rightarrow [k] \cup \{\perp\}$  is the partial labeling function of the edges. The symbol  $\perp$  stands for ‘unlabeled’. The partial labelings  $\pi_t$  and  $\tau_t$  have the following properties.

- *Distinctness:* For each pair of distinct vertices  $u, v \in V$ , if  $\pi_t(u) \in [n]$  and  $\pi_t(v) \in [n]$ , then  $\pi_t(u) \neq \pi_t(v)$ , and an analogous property holds for labels of edges incident to a common vertex.
- *Consistency:* For each vertex  $u \in V$ , if  $\pi_{t-1}(u) \in [n]$ , then  $\pi_t(u) = \pi_{t-1}(u)$ , and an analogous property holds for edge labels.

We let  $V_t = \{u \in V : \pi_t(u) \in [n]\}$  denote the set of labeled vertices after the first  $t$  queries and  $\Pi_t = \{\pi_t(u) : u \in V_t\}$  denotes the set of labels of these vertices. (We set  $V_0 = \emptyset$  and  $\Pi_0 = \emptyset$ .) In order to distinguish between vertices in  $T_t$  and their labels, we shall use  $u, v, w$  for the former, and  $x, y, z$  for the latter. Note that if  $x \in \Pi_t$ , then  $\pi_t^{-1}(x)$  is well defined.

In addition to the partially labeled tree  $T_t$ , the process  $\mathcal{P}$  maintains an auxiliary tree  $T'_t = (V', E', \pi'_t, \tau'_t)$  of depth  $L = \frac{\log n}{C \cdot \log \log n}$  in which each internal vertex has degree  $k$ . The role of this tree will become clear shortly. For now we just say that in  $T'_t$ ,  $\mathcal{P}$  maintains labeling information concerning vertices and edges that are “close to the root of  $T$ ” (without yet determining their exact identity).

Initially  $T'_t$  is unlabeled, so that  $\pi'_0(u') = \perp$  for every  $u' \in V'$  and  $\tau'_0((u', v')) = \perp$  for every  $(u', v') \in E'$ . The labeling functions  $\pi'_t$  and  $\tau'_t$  maintain the distinctness and consistency properties as defined for  $\pi_t$  and  $\tau_t$ . Furthermore, the labels of vertices in  $T'_t$  are distinct from those in  $T_t$ . We let  $\Pi'_t$  denote the set of labels of vertices in  $T'_t$  (i.e.,  $\Pi'_t = \{\pi'_t(u') : u' \in V'\} \setminus \{\perp\}$ ).  $\mathcal{P}$  maintains the invariant that in  $T'_t$ , the labeled vertices form a connected subtree (rooted at the root  $r'$  of  $T'_t$ ). This is as opposed to  $T_t$ , in which the labeled vertices may correspond to several connected components.

We view  $\mathcal{P}$  as *committing* to the labels of vertices in  $V_t$ , in the sense that these will be the labels of the corresponding vertices in the final labeled tree  $\widehat{T} \in \mathcal{T}$ . On the other hand, for labeled vertices in  $T'_t$ , there is only a “partial commitment.” That is,  $\mathcal{P}$  commits to the (partial) labeling of  $T'$  constructed during the interaction with  $\mathcal{A}$ , but does not commit to a correspondence between vertices and edges in  $T'$  and those in  $\widehat{T}$  until after  $\mathcal{A}$  is finished with its queries. Only at the end of the interaction does  $\mathcal{P}$  choose a random embedding of  $T'$  into  $\widehat{T}$ . The embedding induces a partial labeling on  $\widehat{T}$  as follows: First, the root  $r'$  of  $T'$  is mapped to a vertex  $u_* \in V$  at depth at most  $L$ , and  $u_*$  is labeled by  $\pi'_t(r')$ . Then, the labels of the children of  $r'$  are assigned to neighbors of  $u_*$ , and so on, until all the vertices in  $T'_t$  are mapped to the vertices of  $T$ .

#### 4.1.2 Answering queries

We refer to a query  $q_t$  involving a label  $x_t \notin \Pi_{t-1} \cup \Pi'_{t-1}$  as a *new label* query. In order to simplify the presentation, we assume that for any query involving a new label,  $\mathcal{A}$  first performs a degree query on  $x_t$ , and following this query we have  $x_t \in \Pi_t \cup \Pi'_t$ . A lower bound on the number of queries performed by  $\mathcal{A}$  under this assumption translates to the same lower bound up to a factor of three in the standard model.

We also assume that if a neighbor query  $(x_t, i_t)$  was answered by  $y_t$ , where  $y_t$  is a label yet unobserved by  $\mathcal{A}$ , then  $\mathcal{P}$  provides  $\mathcal{A}$  with the degree of the vertex labeled by  $y_t$  as well as the label of the edge from this vertex to the vertex labeled by  $x_t$ . Similarly, if a pair query  $\text{pair}(x_t, y_t)$  is answered positively, then  $\mathcal{P}$  returns the labels of the edges between the vertices labeled by  $x_t$  and  $y_t$ , respectively. Clearly, any lower bound that holds under these “augmented answers” holds under the standard query model. It follows that for each neighbor query  $q_t = \text{nbr}(x_t, i_t)$  we have that  $x_t \in \Pi_{t-1} \cup \Pi'_{t-1}$  and similarly for each pair query  $q_t = \text{pair}(x_t, y_t)$ , we have that  $x_t, y_t \in \Pi_{t-1} \cup \Pi'_{t-1}$ . On the other hand, for each degree query  $q_t = \text{deg}(x_t)$ , we may assume that  $x_t \notin \Pi_{t-1} \cup \Pi'_{t-1}$  (since  $\mathcal{A}$  is provided with the degrees of all vertices with labels in  $\Pi_{t-1} \cup \Pi'_{t-1}$ ).

Let  $\mathcal{T}_t$  denote the set of all labeled trees in  $\mathcal{T}$  that are consistent with  $T_t$  and  $T'_t$ . That is,  $\mathcal{T}_t$  is the family of all labeled trees in  $\mathcal{T}$  that are consistent with  $T$  and with an embedding of  $T'$  into  $T$  such that the root of  $T'$  is mapped to a vertex of  $T$  whose depth is at most  $L$ . In order to simplify the analysis we shall elaborately describe how  $\mathcal{P}$  answers each query. However, as will be evident from our description, it holds that  $\mathcal{P}$ 's answer to the  $t^{\text{th}}$  query  $q_t$  can be viewed as choosing a random labeled tree  $\widehat{T}_t \in \mathcal{T}_{t-1}$ , and answering  $q_t$  according to  $\widehat{T}_t$ . (Observe that

$\widehat{T}_t$  is only used to answer  $q_t$  and is later discarded.) Therefore, at the end of the interaction  $\mathcal{P}$  generates a uniform labeled tree  $\widehat{T} \in \mathcal{T}$ .<sup>5</sup>

$\mathcal{P}$  answers the  $t^{\text{th}}$  query  $q_t$  of  $\mathcal{A}$  as follows:

**Degree queries.** Recall that by assumption, for each degree query  $q_t = \text{deg}(x_t)$  we have  $x_t \notin \Pi_t \cup \Pi'_t$ . When a degree query is made,  $\mathcal{P}$  first decides whether the depth of the vertex to be labeled by  $x_t$  is larger than  $2L$ , at most  $L$ , or in between. (Recall that  $L = \frac{\log n}{C \cdot \log \log n}$ .) For each vertex  $u \in V$ , we say that  $u$  is *deep* if  $\Delta(u) > 2L$ , *shallow* if  $\Delta(u) \leq 2L$ , and *critical* if  $\Delta(u) \leq L$ . We denote the number of deep, shallow, and critical vertices in  $T$  by  $n_d, n_s$ , and  $n_c$ , respectively. For any  $t \leq Q$ , we denote the number of deep, shallow, and critical vertices that have been labeled by  $\mathcal{P}$  before  $\mathcal{A}$ 's  $t^{\text{th}}$  query by  $\ell_d(t)$ ,  $\ell_s(t)$ , and  $\ell_c(t)$ , respectively. The count  $\ell_s(t)$  also includes the number of labels assigned to the auxiliary tree  $T'$ , as these labels will eventually be assigned to shallow vertices.

If the  $t^{\text{th}}$  query performed by  $\mathcal{A}$  is a degree query,  $\mathcal{P}$  first flips a coin with *bias*  $p(t)$  (i.e., the probability that the result is *heads*) defined by

$$p(t) \stackrel{\text{def}}{=} \frac{n_s - \ell_s(t)}{n - \ell_d(t) - \ell_s(t)}.$$

That is,  $p(t)$  is the fraction of unlabeled vertices in  $T_{t-1}$  that are shallow. If the coin flip turns out *tails* (with probability  $1 - p(t)$ ), then we say that the outcome of the query is *deep*. In this case  $\mathcal{P}$  selects an unlabeled deep vertex  $u$  uniformly at random, sets  $\pi_t(u) = x_t$ , and returns the degree of  $u$ .

If the coin flip turns out *heads*, then we say that the outcome of the query is *shallow*. In this case  $\mathcal{P}$  first checks if the outcome of any previous degree query was shallow. If a previous degree query was shallow,  $\mathcal{P}$  completes the labeling of  $T$  as follows. It first selects a random embedding of  $T'_t$  into  $T_t$ , then it selects uniformly at random a shallow unlabeled vertex to be the preimage of  $x$  and finally it completes the labeling of  $T_t$  uniformly at random to a labeled tree  $\widehat{T}$ . After the completion of the labeling, it gives  $\mathcal{A}$  the label of the root of  $\widehat{T}$ , thereby allowing  $\mathcal{A}$  to succeed. In this case we say that  $\mathcal{A}$  *succeeds trivially*.

Otherwise (the outcome of this query is shallow, and there was no previous shallow outcome),  $\mathcal{P}$  flips another coin, this time with bias

$$p'(t) \stackrel{\text{def}}{=} \frac{n_c - \ell_c(t)}{n_s - \ell_s(t)} = \frac{n_c}{n_s - \ell_s(t)},$$

That is,  $p'(t)$  is the fraction of unlabeled shallow vertices that are also critical, where the equality is justified as follows. Given that there was no previous degree query whose outcome was shallow, it holds that for every  $v \in V_t$ ,  $\Delta(v) > 2L - t$ , implying that no critical vertex could have been reached through neighbor queries. Therefore,  $\ell_c(t) = 0$ . If this coin flip results in *tails* (with probability  $1 - p'(t)$ ), then we say that the outcome of the query is *not critical*. In this case  $\mathcal{P}$  picks a uniformly random unlabeled vertex  $u$  that is shallow, but not critical (i.e., with  $L < \Delta(u) \leq 2L$ ), assigns  $\pi_t(u) = x_t$ , and returns the degree of  $u$  (which is necessarily  $k$ ). Finally, in the case where the second coin flip is *heads*, i.e., the outcome is *critical*,  $\mathcal{P}$  does the following. It labels the root  $r'$  of  $T'$  with  $x_t$ , i.e., sets  $\pi'_t(r') = x_t$ , so that now  $\Pi'_t = \{x_t\}$ . In this subcase  $\mathcal{P}$  returns that the answer to this degree query is  $k$ .

---

<sup>5</sup>The described process is equivalent to the following. Consider all possible answers to the query  $q_t$  and weigh each answer  $a_t$  according to the fraction of trees in  $\mathcal{T}_{t-1}$  for which  $\text{ans}(q_t) = a_t$ . That is, for every possible answer  $a_t$ ,  $\Pr[a_t] = \frac{|\mathcal{T}_{t-1}(a_t)|}{|\mathcal{T}_{t-1}|}$ , where  $\mathcal{T}_{t-1}(a_t)$  is the subset of labeled trees in  $\mathcal{T}_{t-1}$  in which  $\text{ans}(q_t) = a_t$ .

Observe that the above process is equivalent to choosing a labeled tree  $\widehat{T}_t \in \mathcal{T}_{t-1}$  uniformly at random, and answering according to  $\pi^{-1}(x_t)$  in  $\widehat{T}_t$ . Then  $\mathcal{T}_t \subseteq \mathcal{T}_{t-1}$  is taken to be the sub-family of labelings consistent with this query answer.

**Neighbor queries.** First consider the case that  $q_t$  is a neighbor query  $\text{nbr}(x_t, i_t)$ , and recall that by our assumption on  $\mathcal{A}$ ,  $x_t \in \Pi_{t-1} \cup \Pi'_{t-1}$ . If  $x_t \in \Pi_t$ , then let  $u = \pi_{t-1}^{-1}(x_t)$ . We may assume without loss of generality that there is no edge  $(u, v)$  incident to  $u$  such that  $\tau_{t-1}((u, v)) = i$  (or else  $\mathcal{A}$  does not need to perform this query, as its answer is implied by answers to previous queries). We may also assume that if  $u$  is a leaf, so that it has degree 1, then  $i_t = 1$ .

Now  $\mathcal{P}$  selects uniformly at random an edge  $(u, v) \in E$  such that  $\tau_{t-1}((u, v)) = \perp$ . If  $\pi_{t-1}(v) \in [n]$ , then  $\mathcal{P}$  sets  $\tau_t((u, v)) = i_t$ , and lets  $\tau_t((v, u))$  be a uniformly selected  $j$  in  $[k] \setminus \{\tau_{t-1}((v, w))\}$ . Otherwise ( $\pi_{t-1}(v) = \perp$ ), it sets  $\pi_t(v)$  to a uniformly selected label in  $[n] \setminus (\Pi_{t-1} \cup \Pi'_{t-1})$  and then proceeds as in the case that  $\pi_{t-1}(v) \in [n]$ . Finally,  $\mathcal{P}$  returns  $\pi_t(v)$  as well as the degrees of  $u$  and  $v$  and  $\tau_t((v, u))$ .

If  $x_t \in \Pi'_t$ , then for  $u' = (\pi'_{t-1})^{-1}(x_t)$ ,  $\mathcal{P}$  selects uniformly at random an edge  $(u', v') \in E'$  such that  $\tau_{t-1}((u', v')) = \perp$ . Here it will always be the case that  $\pi'_{t-1}(v') = \perp$  (since the labeling of  $T'$  is always done from the root down the tree), so that  $\mathcal{P}$  sets  $\pi'_t(v')$  to a uniformly selected label in  $[n] \setminus (\Pi_{t-1} \cup \Pi'_{t-1})$ . It then sets  $\tau'_t((u', v')) = i_t$ , and lets  $\tau'_t((v', u'))$  be a uniformly selected  $j$  in  $[k]$ .

**Pair queries.** Next consider the case that  $q_t$  is a pair query  $\text{pair}(x_t, y_t)$ . Recall that  $x_t, y_t \in \Pi_{t-1} \cup \Pi'_{t-1}$  by assumption. We may assume without loss of generality that neither  $y_t$  was an answer to a previous neighbor query  $(x_{t'}, i_{t'})$  such that  $x_{t'} = x_t$ , nor that  $x_t$  was an answer to a previous neighbor query  $(x_{t'}, i_{t'})$  such that  $x_{t'} = y_t$  (or else  $\mathcal{A}$  does not need to perform this query, as its answer is implied by answers to previous queries).

If  $x_t, y_t \in \Pi_{t-1}$ , then let  $u = \pi_{t-1}^{-1}(x_t)$  and  $v = \pi_{t-1}^{-1}(y_t)$ . If  $u$  and  $v$  are neighbors in  $T$ , then  $\mathcal{P}$  returns a positive answer. It also selects random labels for  $(u, v)$  and  $(v, u)$  (among those labels not yet used for edges incident to  $u$  and  $v$ , respectively), updates  $\tau_t$  accordingly, and returns these labels to  $\mathcal{A}$ . If  $u$  and  $v$  are not neighbors in  $T$ , then  $\mathcal{P}$  returns a negative answer.

If  $x_t \in \Pi_{t-1}$  and  $y_t \in \Pi'_{t-1}$  (or vice versa), then  $\mathcal{P}$  returns a negative answer, since for any embedding of  $T'$  into  $T$  it cannot be the case that  $(\pi'_{t-1})^{-1}(y_t)$  will be mapped to a neighbor of  $\pi_{t-1}^{-1}(x)$ . This is true since for every  $x \in \Pi_{t-1}$ ,  $\Delta(\pi^{-1}(x)) > 2L - t_1$  and every vertex  $u' = (\pi')^{-1}(y)$  for  $y \in \Pi'_{t-1}$  will be mapped to a vertex in  $T$  whose depth is at most  $L + t_2$ , for  $t_1, t_2$  such that  $t_1 + t_2 = t$ . Hence, for any embedding of  $T'$  into  $T$ , the preimages of  $x_t$  and  $y_t$  cannot be neighbors. Finally, we may assume that  $\mathcal{A}$  does not perform a pair query  $(x_t, y_t)$  where  $x_t, y_t \in \Pi'_{t-1}$ , since the answer to this query is implied by previous queries.

### 4.1.3 The final (fully-labeled) tree $\widehat{T} \in \mathcal{T}$

Following the last query of  $\mathcal{A}$ , the process  $\mathcal{P}$  selects a labeled tree  $\widehat{T} \in \mathcal{T}_Q$  uniformly at random. Observe that if  $\Pi'_Q = \emptyset$ , then each labeled tree in  $\mathcal{T}_Q$  simply corresponds to a possible completion of the labeling of  $T_Q$  to a complete labeling of  $T$ . If  $\Pi'_Q \neq \emptyset$ , then the trees in  $\mathcal{T}_Q$  can be viewed as trees obtained by the following two step process. First choose a uniform embedding  $T'$  into  $T$  from all possible embeddings that map the root of  $T'$  to a critical vertex in  $T$ . Second, complete the labeling to a complete labeling of  $T$ .

## 4.2 Completing the analysis

In order to complete the analysis, we must argue that any  $\mathcal{A}$  succeeds when interacting with  $\mathcal{P}$  is small.

To bound the probability that  $\mathcal{A}$  succeeds, we consider three events:

- $\mathcal{E}_1$ : There was no (degree) query whose outcome was critical.
- $\mathcal{E}_2$ : There was a degree query whose outcome was critical, and no other degree query whose outcome was shallow.
- $\mathcal{E}_3$ : There was more than one degree query whose outcome was shallow.

The three events above are exhaustive. The following three claims bound the probability that  $\mathcal{A}$  succeeds in each case.

**Claim 4.3.** *Let  $\mathcal{A}$  be an algorithm that performs  $Q \leq L$  queries interacting with  $\mathcal{P}$ . Then  $\Pr[\mathcal{A} \text{ succeeds} \mid \mathcal{E}_1] = O(1/n)$ , where the probability is taken over both the randomness of  $\mathcal{P}$  and the randomness of  $\mathcal{A}$ .*

*Proof.* We consider two subcases. First, suppose  $\mathcal{A}$  returns the label  $x$  of a vertex  $v$  that was labeled in  $T_Q$  (i.e., before  $\mathcal{P}$  completed the partial labeling formed during the interaction, so that  $x \in \Pi_Q$ ). Since none of  $\mathcal{A}$ 's queries had a critical outcome, the root  $r$  was unlabeled in  $T_Q$ . Therefore, in this case,  $\mathcal{A}$  succeeds with probability 0.

Now consider the case that  $\mathcal{A}$  returns a label  $x \notin \Pi_Q$ . Since  $\mathcal{P}$  chooses the label of  $r$ ,  $\pi(r)$ , uniformly from  $[n] \setminus \Pi_Q$ , the probability that  $\mathcal{A}$  succeeds in this case is  $1/(n - |\Pi_Q|) \leq 1/(n - Q) = O(1/n)$ , which gives the desired result.  $\square$

**Claim 4.4.** *Let  $\mathcal{A}$  be any algorithm that performs  $Q \leq L$  queries interacting with  $\mathcal{P}$ . Then*

$$\Pr[\mathcal{A} \text{ succeeds} \wedge \mathcal{E}_2] = \Pr[\mathcal{A} \text{ succeeds} \mid \mathcal{E}_2] \cdot \Pr[\mathcal{E}_2] = O\left(\frac{Q}{n}\right),$$

where the probability is taken over both the randomness of  $\mathcal{P}$  and the randomness of  $\mathcal{A}$ .

*Proof.* We first observe that

$$\Pr[\mathcal{E}_2] \leq \frac{Q \cdot n_c}{n - Q} = O\left(\frac{Q \cdot n_c}{n}\right). \quad (9)$$

Indeed, the outcome of the  $t^{\text{th}}$  query is critical with probability

$$(1 - p(t)) \cdot p'(t) = \frac{n_s - \ell_s(t)}{n - \ell_d(t) - \ell_s(t)} \cdot \frac{n_c}{n_s - \ell_s(t)} = \frac{n_c}{n - |\Pi_t|},$$

so the expression in Equation (9) follows by taking a union bound over  $t = 1, 2, \dots, Q$ .

We now turn to bound that probability that  $\mathcal{A}$  succeeds conditioned on the event  $\mathcal{E}_2$ . First note that if  $\mathcal{A}$  outputs a label not in  $\Pi'_Q$ , then the probability that it succeeds (in particular conditioned on  $\mathcal{E}_2$ ) is  $O(1/n)$ , as in the proof of Claim 4.3.

Now assume that  $\mathcal{A}$  outputs a label  $x \in \Pi'_Q$ . Recall that  $\mathcal{P}$  completes the labeling of  $T$  by choosing a labeled tree in  $\mathcal{T}_Q$  (which is the family of labeled trees that result from embedding  $T'_Q$  into  $T_Q$  uniformly at random and then completing the labeling of  $T$  uniformly at random). Recall that the embedding is done by mapping the root  $r'$  of  $T$  into a uniformly selected critical vertex  $u_*$  in  $T$ . In order to complete the proof of Claim 4.4, we build on the next claim.

**Claim 4.5.** *Consider embedding  $T'$  into  $T$  by choosing a critical vertex  $u_*$  in  $T$  uniformly at random and mapping  $r'$  to  $u_*$ . Then any vertex  $u'$  in  $T'$  is equally likely to be mapped to  $r$ .*

We defer the proof of the claim and now continue assuming its correctness. Given Claim 4.5, it immediately follows that

$$\Pr[\mathcal{A} \text{ succeeds} \mid \mathcal{E}_2] = \Pr[(\pi'_Q)^{-1}(x) \text{ will be mapped to } r \mid \mathcal{E}_2] = \frac{1}{n_c}. \quad (10)$$

Claim 4.4 follows by combining Equation (10) with Equation (9).  $\square$

We now prove Claim 4.5.

*Proof of Claim 4.5.* Fix any choice of  $u' \in T'$ . Let  $\Delta' = \Delta(u')$  be the depth of  $u'$  in  $T'$ , and let  $r' = u'_0, u'_1, \dots, u'_{\Delta'} = u'$  denote the sequence of vertices on the path from  $r'$  to  $u'$  in  $T'$ . In order for  $u'$  to be mapped to  $r$  in the embedding of  $T'$  into  $T$  (when  $r'$  is mapped to a randomly selected critical vertex  $u_*$  in  $T$ ), the following events must occur.

- First, the vertex  $u_*$  in  $T$  to which  $r'$  is mapped must be at depth  $\Delta'$  in  $T$ . This occurs with probability  $\frac{k \cdot (k-1)^{\Delta'-1}}{n_c}$ .
- Second, letting  $u_* = v_0, v_1, \dots, v_{\Delta'} = r$  be the sequence of vertices on the path from  $u_*$  to  $r$  in  $T$ , the following must hold. For each  $j \in [\Delta']$ , the child  $u'_j$  of  $u'_{j-1}$  must be mapped to the parent,  $v_j$ , of  $v_{j-1}$  (recall that  $u'_0 = r'$  and  $v_0 = u_*$ ). This occurs with probability  $k^{-1} \cdot (k-1)^{-(\Delta'-1)}$  as explained next. In the random embedding process, each of the  $k$  children of  $r' = u'_0$  is equally likely to be mapped to the parent of  $u_* = v_0$ , and for  $j > 1$ , conditioned on  $u'_1, \dots, u'_{j-1}$  being mapped to  $v_1, \dots, v_{j-1}$ , respectively,  $u'_j$  is mapped to  $v_j$  with probability  $\frac{1}{k-1}$ .

The claim follows since the events are independent.  $\square$

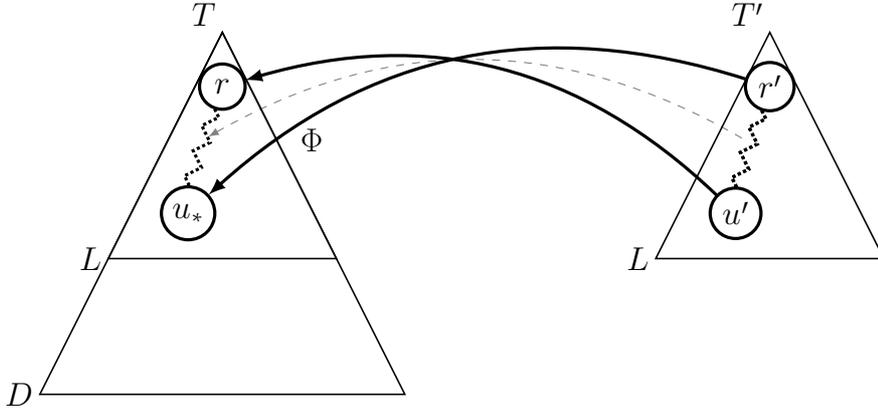


Figure 1: An illustration of a possible embedding  $\phi$  of  $T'$  into  $T$ , where  $r'$  is mapped by  $\phi$  to  $u_*$ . Denoting the depth of  $u_*$  by  $\Delta$ , we have that one of the vertices at depth  $\Delta$  in  $T'$  will be mapped by  $\phi$  to the root  $r$  of  $T$ .

Finally we consider the event  $\mathcal{E}_3$  (conditioned on which  $\mathcal{A}$  trivially succeeds).

**Claim 4.6.** For  $Q \leq L$ ,

$$\Pr[\mathcal{E}_3] = \frac{Q^2 k^{4L}}{n^2}.$$

*Proof.* The number of vertices at depth  $\Delta \geq 1$  is  $k(k-1)^{\Delta-1}$ , so that the number of shallow vertices is

$$n_s = 1 + k + k(k-1) + \dots + k(k-1)^{2L-1} = O(k^{2L}).$$

Thus, the probability that the outcome of any degree query is shallow is at most  $n_s/(n-Q) = O(k^{2L}/n)$ . Taking a union bound over all queries, the probability that at least two are shallow is  $O(Q^2 k^{4L}/n^2)$ .  $\square$

We are now ready to prove Proposition 4.1 which implies Theorem 3.

*Proof of Proposition 4.1.* Since the events  $\mathcal{E}_1, \mathcal{E}_2$  and  $\mathcal{E}_3$  are exhaustive, by Claims 4.3, 4.4 and 4.6, it holds for  $Q \leq L$  and  $L = \frac{\log n}{C \cdot \log \log n}$  that

$$\begin{aligned} \Pr[A \text{ succeeds}] &= \Pr[A \text{ succeeds} \mid \mathcal{E}_1] \cdot \Pr[\mathcal{E}_1] + \Pr[A \text{ succeeds} \mid \mathcal{E}_2] \cdot \Pr[\mathcal{E}_2] \\ &\quad + \Pr[A \text{ succeeds} \mid \mathcal{E}_3] \cdot \Pr[\mathcal{E}_3] \\ &= O\left(\frac{1}{n}\right) \cdot 1 + O\left(\frac{Q}{n}\right) + 1 \cdot O\left(\frac{Q^2 k^{4L}}{n^2}\right) = O\left(\frac{Q}{n}\right). \end{aligned}$$

Therefore, the proposition follows for  $Q = O\left(\frac{\log n}{\log \log n}\right)$ .  $\square$

## References

- [1] Maryam Aliakbarpour, Amartya Shankha Biswas, Themis Gouleakis, John Peebles, Ronitt Rubinfeld, and Anak Yodpinyanee. Sublinear-time algorithms for counting star subgraphs via edge sampling. *Algorithmica*, pages 1–30, 2017. URL: <http://dx.doi.org/10.1007/s00453-017-0287-3>, doi:10.1007/s00453-017-0287-3.
- [2] Sepehr Assadi, Michael Kapralov, and Sanjeev Khanna. A simple sublinear-time algorithm for counting arbitrary subgraphs via edge sampling. In *ITCS*, volume 124 of *LIPICs*, pages 6:1–6:20. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2019.
- [3] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.
- [4] Leonid Barenboim and Michael Elkin. Sublogarithmic distributed MIS algorithm for sparse graphs using Nash-Williams decomposition. *Distributed Computing*, 22(5-6):363–379, 2010. URL: <https://doi.org/10.1007/s00446-009-0088-2>, doi:10.1007/s00446-009-0088-2.
- [5] Reinhard Bauer, Marcus Krug, and Dorothea Wagner. Enumerating and generating labeled k-degenerate graphs. In *Proceedings of the Meeting on Algorithm Engineering & Experiments*, pages 90–98. Society for Industrial and Applied Mathematics, 2010.
- [6] Michael Baur, Marco Gaertler, Robert Görke, Marcus Krug, and Dorothea Wagner. Generating graphs with predefined k-core structure. In *Proceedings of the European Conference of Complex Systems*. Citeseer, 2007.
- [7] Eric Blais, Joshua Brody, and Kevin Matulef. Property testing lower bounds via communication complexity. *Computational Complexity*, 21(2):311–358, 2012.

- [8] Talya Eden, Amit Levi, Dana Ron, and C Seshadhri. Approximately counting triangles in sublinear time. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 614–633. IEEE, 2015.
- [9] Talya Eden, Reut Levi, and Dana Ron. *Testing bounded arboricity*, pages 2081–2092. URL: <https://epubs.siam.org/doi/abs/10.1137/1.9781611975031.136>, [arXiv:https://epubs.siam.org/doi/pdf/10.1137/1.9781611975031.136](https://arxiv.org/abs/https://epubs.siam.org/doi/pdf/10.1137/1.9781611975031.136), [doi:10.1137/1.9781611975031.136](https://doi.org/10.1137/1.9781611975031.136).
- [10] Talya Eden, Dana Ron, and C. Seshadhri. Sublinear Time Estimation of Degree Distribution Moments: The Degeneracy Connection. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, volume 80 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 7:1–7:13, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2017/7374>, [doi:10.4230/LIPIcs.ICALP.2017.7](https://doi.org/10.4230/LIPIcs.ICALP.2017.7).
- [11] Talya Eden, Dana Ron, and C. Seshadhri. Faster sublinear approximations of  $k$ -cliques for low arboricity graphs. *CoRR*, abs/1811.04425, 2018. URL: <http://arxiv.org/abs/1811.04425>, [arXiv:1811.04425](https://arxiv.org/abs/1811.04425).
- [12] Talya Eden, Dana Ron, and C. Seshadhri. On approximating the number of  $k$ -cliques in sublinear time. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 722–734, 2018. URL: <https://doi.org/10.1145/3188745.3188810>, [doi:10.1145/3188745.3188810](https://doi.org/10.1145/3188745.3188810).
- [13] Talya Eden and Will Rosenbaum. Lower bounds for approximating graph parameters via communication complexity. *CoRR*, abs/1709.04262, 2017. URL: <http://arxiv.org/abs/1709.04262>, [arXiv:1709.04262](https://arxiv.org/abs/1709.04262).
- [14] Talya Eden and Will Rosenbaum. Lower bounds for approximating graph parameters via communication complexity. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2018, August 20-22, 2018 - Princeton, NJ, USA*, pages 11:1–11:18, 2018. URL: <https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2018.11>, [doi:10.4230/LIPIcs.APPROX-RANDOM.2018.11](https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2018.11).
- [15] Talya Eden and Will Rosenbaum. On Sampling Edges Almost Uniformly. In Raimund Seidel, editor, *1st Symposium on Simplicity in Algorithms (SOSA 2018)*, volume 61 of *OpenAccess Series in Informatics (OASICS)*, pages 7:1–7:9, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2018/8300>, [doi:10.4230/OASICS.SOSA.2018.7](https://doi.org/10.4230/OASICS.SOSA.2018.7).
- [16] David Eppstein and Darren Strash. Listing all maximal cliques in large sparse real-world graphs. In *International Symposium on Experimental Algorithms*, pages 364–375. Springer, 2011.
- [17] Uriel Feige. On sums of independent random variables with unbounded variance and estimating the average degree in a graph. *SIAM Journal on Computing*, 35(4):964–984, 2006.

- [18] Gaurav Goel and Jens Gustedt. Bounded arboricity to determine the local structure of sparse graphs. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 159–167. Springer, 2006.
- [19] Oded Goldreich and Dana Ron. Approximating average parameters of graphs. *Random Struct. Algorithms*, 32(4):473–493, 2008.
- [20] Mira Gonen, Dana Ron, and Yuval Shavitt. Counting stars and other small subgraphs in sublinear-time. *SIAM Journal on Discrete Mathematics*, 25(3):1365–1411, 2011.
- [21] Bala Kalyanasundaram and Georg Schintger. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics*, 5(4):545–557, 1992.
- [22] C. St. JA. Nash-Williams. Edge-disjoint spanning trees of finite graphs. *Journal of the London Mathematical Society*, 1(1):445–450, 1961.
- [23] Alexander A. Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106(2):385–390, 1992.
- [24] Kijung Shin, Tina Eliassi-Rad, and Christos Faloutsos. Patterns and anomalies in k-cores of real-world graphs with applications. *Knowledge and Information Systems*, 54(3):677–710, 2018.