

# Determinant-Preserving Sparsification of SDDM Matrices with Applications to Counting and Sampling Spanning Trees

David Durfee\*      John Peebles†      Richard Peng‡      Anup B. Rao§

May 3, 2017

## Abstract

We show variants of spectral sparsification routines can preserve the total spanning tree counts of graphs, which by Kirchhoff’s matrix-tree theorem, is equivalent to determinant of a graph Laplacian minor, or equivalently, of any SDDM matrix. Our analyses utilizes this combinatorial connection to bridge between statistical leverage scores / effective resistances and the analysis of random graphs by [Janson, Combinatorics, Probability and Computing ‘94]. This leads to a routine that in quadratic time, sparsifies a graph down to about  $n^{1.5}$  edges in ways that preserve both the determinant and the distribution of spanning trees (provided the sparsified graph is viewed as a random object). Extending this algorithm to work with Schur complements and approximate Choleksy factorizations leads to algorithms for counting and sampling spanning trees which are nearly optimal for dense graphs.

We give an algorithm that computes a  $(1 \pm \delta)$  approximation to the determinant of any SDDM matrix with constant probability in about  $n^2\delta^{-2}$  time. This is the first routine for graphs that outperforms general-purpose routines for computing determinants of arbitrary matrices. We also give an algorithm that generates in about  $n^2\delta^{-2}$  time a spanning tree of a weighted undirected graph from a distribution with total variation distance of  $\delta$  from the  $w$ -uniform distribution .

## 1 Introduction

The determinant of a matrix is a fundamental quantity in numerical algorithms due to its connection to the rank of the matrix and its interpretation as the volume of the ellipsoid corresponding of the matrix. For graph Laplacians, which are at the core of spectral graph theory and spectral algorithms, the matrix-tree theorem gives that the determinant of the minor obtained by removing one row and the corresponding column equals to the total weight of all the spanning trees in the graph [Kir47] . Formally on a weighted graph  $G$  with  $n$  vertices we have:

$$\det(\mathbf{L}_{1:n-1,1:n-1}^G) = \mathcal{T}_G,$$

where  $\mathbf{L}^G$  is the graph Laplacian of  $G$  and and  $\mathcal{T}_G$  is the total weight of all the spanning trees of  $G$ . As the all-ones vector is in the null space of  $\mathbf{L}^G$ , we need to drop its last row

---

\*Georgia Institute of Technology. email:ddurfee@gatech.edu

†Massachusetts Institute of Technology. email:jpeebles@mit.edu

‡Georgia Institute of Technology. email:rpeng@cc.gatech.edu

§Georgia Institute of Technology. email:anup.rao@gatech.edu

and column and work with  $\mathbf{L}_{1:n-1,1:n-1}^G$ , which is precisely the definition of SDDM matrices in numerical analysis [ST14]. The study of random spanning trees builds directly upon this connection between tree counts and determinants, and also plays an important role in graph theory [GRV09, AGM<sup>+</sup>10, FHHP11].

While there has been much progress in the development of faster spectral algorithms, the estimation of determinants encapsulates many shortcomings of existing techniques. Many of the nearly linear time algorithms rely on sparsification procedures that remove edges from a graph while provably preserving the Laplacian matrix as an operator, and in turn, crucial algorithmic quantities such as cut sizes, Rayleigh quotients, and eigenvalues. The determinant of a matrix on the other hand is the product of all of its eigenvalues. As a result, a worst case guarantee of  $1 \pm (\epsilon/n)$  per eigenvalue is needed to obtain a good overall approximation, and this in turn leads to additional factors of  $n$  in the number of edges needed in the sparse approximate.

Due to this amplification of error by a factor of  $n$ , previous works on numerically approximating determinants without dense-matrix multiplications [BDKZ15a, HAB14a, HMS15] usually focus on the log-determinant, and (under a nearly-linear running time) give errors of additive  $\epsilon n$  in the log determinant estimate, or a multiplicative error of  $\exp(\epsilon n)$  for the determinant. The lack of a sparsification procedure also led to the running time of random spanning tree sampling algorithms to be limited by the sizes of the dense graphs generated in intermediate steps [KM09, MST15, DKP<sup>+</sup>16].

In this paper, we show that a slight variant of spectral sparsification preserves determinant approximations to a much higher accuracy than applying the guarantees to individual edges. Specifically, we show that sampling  $\omega(n^{1.5})$  edges from a distribution given by leverage scores, or weight times effective resistances, produces a sparser graph whose determinant approximates that of the original graph. Furthermore, by treating the sparsifier itself as a random object, we can show that the spanning tree distribution produced by sampling a random tree from a random sparsifier is close to the spanning tree distribution in the original graph in total variation distance. Combining extensions of these algorithms with sparsification based algorithms for graph Laplacians then leads to quadratic time algorithms for counting and sampling random spanning trees, which are nearly optimal for dense graphs with  $m = \Theta(n^2)$ .

This determinant-preserving sparsification phenomenon is surprising in several aspects: because we can also show—both experimentally and mathematically—that on the complete graph, about  $n^{1.5}$  edges are necessary to preserve the determinant, this is one of the first graph sparsification phenomena that requires the number of edges to be between  $\gg n$ . The proof of correctness of this procedure also hinges upon combinatorial arguments based on the matrix-tree theorem in ways motivated by a result for Janson for complete graphs [Jan94], instead of the more common matrix-concentration bound based proofs [SS11, Tro12, CP15, Coh16]. Furthermore, this algorithm appears far more delicate than spectral sparsification: it requires global control on the number of samples, high quality estimates of resistances (which is the running time bottleneck in Theorem 4.1 below), and only holds with constant probability. Nonetheless, the use of this procedure into our determinant estimation and spanning tree generation algorithms still demonstrates that it can serve as a useful algorithmic tool.

## 1.1 Our Results

We will use  $G = (V, E, \mathbf{w})$  to denote weighted multigraphs, and  $\mathbf{d}_u \stackrel{\text{def}}{=} \sum_{e: e \ni u} \mathbf{w}_e$  to denote the weighted degree of vertex  $u$ . The weight of a spanning tree in a weighed undirected multigraph is:

$$\mathbf{w}(T) \stackrel{\text{def}}{=} \prod_{e \in T} \mathbf{w}_e.$$

We will use  $\mathcal{T}_G$  to denote the total weight of trees,  $\mathcal{T}_G \stackrel{\text{def}}{=} \sum_{T \in \mathcal{T}} \mathbf{w}(T)$ . Our key sparsification result can be described by the following theorem:

**Theorem 1.1.** *Given any graph  $G$  and any parameter  $\delta$ , we can compute in  $O(n^2\delta^{-2})$  time a graph  $H$  with  $O(n^{1.5}\delta^{-2})$  edges such that with constant probability we have*

$$(1 - \delta) \mathcal{T}_G \leq \mathcal{T}_H \leq (1 + \delta) \mathcal{T}_G.$$

This implies that graphs can be sparsified in a manner that preserves the determinant, albeit to a density that is not nearly-linear in  $n$ .

We show how to make our sparsification routine to errors in estimating leverage scores, and how our scheme can be adapted to implicitly sparsify dense objects that we do not have explicit access to. In particular, we utilize tools such as rejection sampling and high quality effective resistance estimation via projections to extend this routine to give determinant-preserving sparsification algorithms for Schur complements, which are intermediate states of Gaussian elimination on graphs, using ideas from the sparsification of random walk polynomials.

We use these extensions of our routine to obtain a variety of algorithms built around our graph sparsifiers. Our two main algorithmic applications are as follows. We achieve the first algorithm for estimating the determinant of an SDDM matrix that is faster than general purpose algorithms for the matrix determinant problem. Since the determinant of an SDDM matrix corresponds to the determinant of a graph Laplacian with one row/column removed.

**Theorem 1.2.** *Given an SDDM matrix  $\mathbf{M}$ , there is a routine DETAPPROX which in  $\tilde{O}(n^2\delta^{-2})$  time outputs  $D$  such that  $D = (1 \pm \delta) \det(\mathbf{M})$  with high probability*

A crucial thing to note which distinguishes the above guarantee from most other similar results is that we give a multiplicative approximation of  $\det(\mathbf{M})$ . This is much stronger than giving a multiplicative approximation of  $\log \det(\mathbf{M})$ , which is what other work typically tries to achieve.

The sparsifiers we construct will also approximately preserve the spanning tree distribution, which we leverage to yield a faster algorithm for sampling random spanning trees. Our new algorithm improves upon the current fastest algorithm for general weighted graphs when one wishes to achieve constant—or slightly sub-constant—total variation distance.

**Theorem 1.3.** *Given an undirected, weighted graph  $G = (V, E, \mathbf{w})$ , there is a routine APPROXTREE which in expected time  $\tilde{O}(n^2\delta^{-2})$  outputs a random spanning tree from a distribution that has total variation distance  $\leq \delta$  from the  $\mathbf{w}$ -uniform distribution on  $G$ .*

## 1.2 Prior Work

### 1.2.1 Graph Sparsification

In the most general sense, a graph sparsification procedure is a method for taking a potentially dense graph and returning a sparse graph called a *sparsifier* that approximately still has many of the same properties of the original graph. It was introduced in [EGIN97] for preserving properties related to minimum spanning trees, edge connectivity, and related problems. [BK96] defined the notion of *cut sparsification* in which one produces a graph whose cut sizes approximate those in the original graph. [ST11] defined the more general notion of *spectral sparsification* which requires that the two graphs' Laplacian matrices approximate each other as quadratic forms.<sup>1</sup> In particular, this spectral sparsification samples  $\tilde{O}(n/\epsilon^2)$  edges from the original graph, yielding a graph with  $\tilde{O}(n/\epsilon^2)$  whose quadratic forms—and hence, eigenvalues—approximate each other within a factor of  $(1 \pm \epsilon)$ . This implies that their determinants approximate each other within  $(1 \pm \epsilon)^n$ . This is not useful from the perspective of preserving the determinant: since one would need to sample  $\Omega(n^3)$  edges to get a constant factor approximation, one could instead exactly compute the determinant or sample spanning trees using exact algorithms with this runtime.

All of the above results on sparsification are for undirected graphs. Recently, [CKP<sup>+</sup>17] has defined a useful notion of sparsification for directed graphs along with a nearly linear time algorithm for constructing sparsifiers under this notion of sparsification.

### 1.2.2 Determinant Estimation

Exactly calculating the the determinant of an arbitrary matrix is known to be equivalent to matrix multiplication [BS83]. For approximately computing the log of the determinant, [IL11] uses the identity  $\log(\det(A)) = \text{tr}(\log(B)) + \text{tr}(\log(B^{-1}A))$  to do this whenever one can find a matrix  $B$  such that the  $\text{tr}(\log(B)) = \log(\det(B))$  and  $\text{tr}(\log(B^{-1}A)) = \log(\det(B^{-1}A))$  can both be quickly approximated.<sup>2</sup>

For the special case of approximating the log determinant of an SDD matrix, [HAB14b] applies this same identity recursively where the  $B$  matrices are a sequence of ultrasparsifiers that are inspired by the recursive preconditioning framework of [ST14]. They obtain a running time of  $O(m(n^{-1}\epsilon^{-2} + \epsilon^{-1})\text{polylog}(n\kappa/\epsilon))$  for estimating the log determinant to additive error  $\epsilon$ .

[BDKZ15b] estimates the log determinant of arbitrary positive definite matrices, but has runtime that depends linearly on the condition number of the matrix.

In contrast, our work is the first we know of that gives a multiplicative approximation of the determinant itself, rather than its log. Despite achieving a much stronger approximation guarantee, our algorithm has essentially the same runtime as that of [HAB14b] when the graph is dense. Note also that if one wishes to conduct an “apples to apples” comparison by setting their value of  $\epsilon$  small enough in order to match our approximation guarantee, their algorithm would only achieve a runtime bound of  $O(mn\delta^{-2}\text{polylog}(n\kappa/\epsilon))$ , which is never better than our runtime and can be as bad as a factor of  $n$  worse.<sup>3</sup>

---

<sup>1</sup>If two graphs Laplacian matrices approximate each other as quadratic forms then their cut sizes also approximate each other.

<sup>2</sup>Specifically, they take  $B$  as the diagonal of  $A$  and prove sufficient conditions for when the log determinant of  $B^{-1}A$  can be quickly approximated with this choice of  $B$ .

<sup>3</sup>This simplification of their runtime is using the substitution  $\epsilon = \delta/n$  which gives roughly  $(1 \pm \delta)$  multiplicative

### 1.2.3 Sampling Spanning Trees

Previous works on sampling random spanning trees are a combination of two ideas: that they could be generated using random walks, and that they could be mapped from a random integer via Kirchoff's matrix tree theorem. The former leads to running times of the form  $O(nm)$  [Bro89, Ald90], while the latter approach[Gue83, Kul90, CMN96, HX16] led to routines that run in  $O(n^\omega)$  time, where  $\omega \approx 2.373$  is the matrix multiplication exponent [Wil12].

These approaches have been combined in algorithms by Kelner and Madry [KM09] and Madry, Straszak and Tarnawski [MST15]. These algorithms are based on simulating the walk more efficiently on parts of the graphs, and combining this with graph decompositions to handle the more expensive portions of the walks globally. Due to the connection with random-walk based spanning tree sampling algorithms, these routines often have inherent dependencies on the edge weights. Furthermore, on dense graphs their running times are still worse than the matrix-multiplication time routines.

The previous best running time for generating a random spanning tree from a weighted graph was  $\tilde{O}(n^{5/3}m^{1/3}\log^2(1/\delta))$  achieved by [DKP<sup>+</sup>16]. It works by combining a recursive procedure similar to those used in the more recent  $O(n^\omega)$  time algorithms [HX16] with spectral sparsification ideas, achieving a runtime of  $\tilde{O}(n^{5/3}m^{13})$ . When  $m = \Theta(n^2)$ , the algorithm in [DKP<sup>+</sup>16] takes  $\tilde{O}(n^{7/3})$  time to produce a tree from a distribution that is  $o(1)$  away from the  $w$ -uniform distribution, which is slower by nearly a  $n^{1/3}$  factor than the algorithm given in this paper.

Our algorithm can be viewed as a natural extension of the sparsification-based approach from [DKP<sup>+</sup>16]: instead of preserving the probability of a single edge being chosen in a random spanning tree, we instead aim to preserve the entire distribution over spanning trees, with the sparsifier itself also considered as a random variable. This allows us to significantly reduce the sizes of intermediate graphs, but at the cost of a higher total variation distance in the spanning tree distributions. This characterization of a random spanning tree is not present in any of the previous works, and we believe it is an interesting direction to combine our sparsification procedure with the other algorithms.

## 1.3 Organization

Section 2 will introduce the necessary notation and some of the previously known fundamental results regarding the mathematical objects that we work with throughout the paper. Section 3 will give a high-level sketch of our primary results and concentration bounds for total tree weight under specific sampling schemes. Section 4 leverages these concentration bounds to give a quadratic time sparsification procedure (down to  $\Omega(n^{1.5})$  edges) for general graphs. Section 5 uses random walk connections to extend our sparsification procedure to the Schur complement of a graph. Section 6 utilizes the previous routines to achieve a quadratic time algorithm for computing the determinant of SDDM matrices. Section 7 combines our results and modifies previously known routines to give a quadratic time algorithm for sampling random spanning trees with low total variation distance. Section 8 extends our concentration bounds to random samplings where an arbitrary tree is fixed, and is necessary for the error accounting of our

---

error in estimating the determinant for their algorithm. This simplification is also assuming  $\delta \leq 1$ , which is the only regime we analyze our algorithm in and thus the only regime in which we can compare the two.

random spanning tree sampling algorithm. Section 9 proves the total variation distance bounds given for our random sampling tree algorithm.

## 2 Background

### 2.1 Graphs, Matrices, and Random Spanning Trees

The goal of generating a random spanning tree is to pick tree  $T$  with probability proportional to its weight, which we formalize in the following definition.

**Definition 2.1** ( $\mathbf{w}$ -uniform distribution on trees). Let  $\mathbf{Pr}_T^G(\cdot)$  be a probability distribution on  $\mathcal{T}_G$  such that

$$\mathbf{Pr}_T^G(T = T_0) = \frac{\prod_{e \in T_0} \mathbf{w}_e}{\mathcal{T}_G}.$$

We refer to  $\mathbf{Pr}_T^G(\cdot)$  as the  $\mathbf{w}$ -uniform distribution on the trees of  $G$ .

When the graph  $G$  is unweighted, this corresponds to the uniform distribution on  $\mathcal{T}_G$ .

We refer to  $\mathbf{Pr}_T^G(\cdot)$  as the  $\mathbf{w}$ -uniform distribution on  $\mathcal{T}_G$ . When the graph  $G$  is unweighted, this corresponds to the uniform distribution on  $\mathcal{T}_G$ . Furthermore, as we will manipulate the probability of a particular tree being chosen extensively, we will denote such probabilities with  $\mathbf{Pr}^G(\widehat{T})$ , aka:

$$\mathbf{Pr}^G(\widehat{T}) \stackrel{\text{def}}{=} \mathbf{Pr}_T^G(T = \widehat{T}).$$

The Laplacian of a graph  $G = (V, E, \mathbf{w})$  is an  $n \times n$  matrix specified by:

$$\mathbf{L}_{uv} \stackrel{\text{def}}{=} \begin{cases} \mathbf{d}_u & \text{if } u = v \\ -\mathbf{w}_{uv} & \text{if } u \neq v \end{cases}$$

We will write  $\mathbf{L}^G$  when we wish to indicate which graph  $G$  that the Laplacian corresponds to and  $\mathbf{L}$  when the context is clear. When the graph has multi-edges, we define  $\mathbf{w}_{uv}$  as the sum of weights of all the edges  $e$  that go between vertices  $u, v$ . Laplacians are natural objects to consider when dealing with random spanning trees due to the matrix tree theorem, which states that the determinant of  $\mathbf{L}$  with any row/column corresponding to some vertex removed is the total weight of spanning trees. We denote this removal of a vertex  $u$  as  $\mathbf{L}_{-u}$ . As the index of vertex removed does not affect the result, we will usually work with  $\mathbf{L}_{-n}$ . Furthermore, we will use  $\det(\mathbf{M})$  to denote the determinant of a matrix. As we will work mostly with graph Laplacians, it is also useful for us to define the ‘positive determinant’  $\det_+$ , where we remove the last row and column. Using this notation, the matrix tree theorem can be stated as:

$$\mathcal{T}_G = \det(\mathbf{L}_{-n}^G) = \det_+(\mathbf{L}^G).$$

We measure the distance between two probability distributions by total variation distance.

**Definition 2.2.** Given two probability distributions  $p$  and  $q$  on the same index set  $\Omega$ , the *total variation distance* between  $p$  and  $q$  is given by

$$d_{TV}(p, q) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{x \in \Omega} |p(x) - q(x)|.$$

Let  $G = (V, E, \mathbf{w})$  be a graph and  $e \in E$  an edge. We write  $G/e$  to denote the graph obtained by contracting the edge  $e$ , i.e., identifying the two endpoints of  $e$  and deleting any self loops formed in the resulting graph. We write  $G \setminus e$  to denote the graph obtained by deleting the edge  $e$  from  $G$ . We extend these definitions to  $G/F$  and  $G \setminus F$  for  $F \subseteq E$  to refer to the graph obtained by contracting all the edges in  $F$  and deleting all the edges in  $F$ , respectively.

Also, for a subset of vertices  $V_1$ , we use  $G[V_1]$  to denote the graph induced on the vertex of  $V_1$ , letting  $G(V_1)$  be the edges associated with  $\mathbf{L}_{[V_1, V_1]}$  in the Schur complement.

## 2.2 Effective Resistances and Leverage Scores

The matrix tree theorem also gives connections to another important algebraic quantity: the *effective resistance* between two vertices. This quantity is formally given as  $\mathcal{R}_{\text{eff}}(u, v) \stackrel{\text{def}}{=} \chi_{uv}^\top \mathbf{L}^{-1} \chi_{uv}$  where  $\chi_{uv}$  is the indicator vector with 1 at  $u$ , -1 at  $v$ , and 0 everywhere else. Via the adjugate matrix, it can be shown that the effective resistance of an edge is precisely the ratio of the number of spanning trees in  $G/e$  over the number in  $G$ :

$$\mathcal{R}_{\text{eff}}(u, v) = \frac{\mathcal{T}_{G/e}}{\mathcal{T}_G}.$$

As  $\mathbf{w}_e \cdot \mathcal{T}_{G/e}$  is the total weight of all trees in  $G$  that contain edge  $e$ , the fraction<sup>4</sup> of spanning trees that contain  $e = uv$  is given by  $\mathbf{w}_e \mathcal{R}_{\text{eff}}(u, v)$ . This quantity is called the *statistical leverage score* of an edge, and we denote it by  $\bar{\tau}_e$ . It is fundamental component of many randomized algorithms for sampling / sparsifying graphs and matrices [SS11, Vis12, Tro12].

The fact that  $\bar{\tau}_e$  is the fraction of trees containing  $e$  also gives one way of deriving the sum of these quantities:

**Fact 2.3.** (*Foster's Theorem*) *On any graph  $G$  we have*

$$\sum_e \bar{\tau}_e = n - 1.$$

The resistance  $\mathcal{R}_{\text{eff}}(u, v)$ , and in turn the statistical leverage scores  $\bar{\tau}_e$  can be estimated using linear system solves and random projections [SS11]. For simplicity, we follow the abstraction utilized by Madry, Straszak, and Tarnawski [MST15], except we also allow the intermediate linear system solves to utilize a sparsifier instead of the original graph.

**Lemma 2.4.** (*Theorem 2.1. of [MST15]*)

Let  $G = (V, E)$  be a graph with  $m$  edges. For every  $\epsilon > 0$  we can find in  $\tilde{O}(\min\{m\epsilon^{-2}, m + n\epsilon^{-4}\})$  time an embedding of the effective resistance metric into  $\mathfrak{R}^{O(\epsilon^{-2 \log m})}$  such that with high probability allows one to compute an estimate  $\tilde{\mathcal{R}}_{\text{eff}}(u, v)$  of any effective resistance satisfying

$$\forall u, v \in V \quad (1 - \epsilon) \tilde{\mathcal{R}}_{\text{eff}}(u, v) \leq \mathcal{R}_{\text{eff}}(u, v) \leq (1 + \epsilon) \tilde{\mathcal{R}}_{\text{eff}}(u, v).$$

Specifically, each vertex  $u$  in this embedding is associated with an (explicitly stored)  $\mathbf{z}_u \in \mathfrak{R}^{O(\epsilon^{-2 \log m})}$ , and for any pair of vertices, the estimate  $\tilde{\mathcal{R}}_{\text{eff}}(u, v)$  is given by:

$$\tilde{\mathcal{R}}_{\text{eff}}(u, v) = \|\mathbf{z}_u - \mathbf{z}_v\|_2^2,$$

which takes  $O(\epsilon^{-2} \log m)$  time to compute once we have the embedding.

<sup>4</sup>provided one thinks of an edge with weight  $w$  as representing  $w$  parallel edges, or equivalently, counts spanning trees with multiplicity according to their weight

## 2.3 Schur Complements

For our applications, we will utilize our determinant-preserving sparsification algorithms in recursions based on Schur complements. A partition of the vertices, which we will denote using

$$V = V_1 \sqcup V_2,$$

partitions the corresponding graph Laplacian into blocks which we will denote using indices in the subscripts:

$$\mathbf{L} = \begin{bmatrix} \mathbf{L}_{[V_1, V_1]} & \mathbf{L}_{[V_1, V_2]} \\ \mathbf{L}_{[V_2, V_1]} & \mathbf{L}_{[V_2, V_2]} \end{bmatrix}.$$

The Schur complement of  $G$ , or  $\mathbf{L}$ , onto  $V_1$  is then:

$$\text{Sc}(G, V_1) = \text{Sc}(\mathbf{L}^G, V_1) \stackrel{\text{def}}{=} \mathbf{L}_{[V_1, V_1]}^G - \mathbf{L}_{[V_1, V_2]}^G \left( \mathbf{L}_{[V_2, V_2]}^G \right)^{-1} \mathbf{L}_{[V_2, V_1]}^G,$$

and we will use  $\text{Sc}(G, V_1)$  and  $\text{Sc}(\mathbf{L}^G, V_1)$  interchangeably. We further note that we will always consider  $V_1$  to be the vertex set we Schur complement onto, and  $V_2$  to be the vertex set we eliminate, except for instances in which we need to consider both  $\text{Sc}(G, V_1)$  and  $\text{Sc}(G, V_2)$ .

Schur complements behave nicely with respect to determinants determinants, which suggests the general structure of the recursion we will use for estimating the determinant.

**Fact 2.5.** *For any matrix  $\mathbf{M}$  where  $\mathbf{M}_{[V_2, V_2]}$  is invertible,*

$$\det(\mathbf{M}_{-n}) = \det(\mathbf{M}_{[V_2, V_2]}) \cdot \det_+(\text{Sc}(\mathbf{M}, V_1)).$$

This relationship also suggests that there should exist a bijection between spanning tree distribution in  $G$  and the product distribution given by sampling spanning trees independently from  $\text{Sc}(\mathbf{L}, V_1)$  and the graph Laplacian formed by adding one row/column to  $\mathbf{L}_{[V_2, V_2]}$ .

Finally, our algorithms for approximating Schur complements rely on the fact that they preserve certain marginal probabilities. The algorithms of [CDN89, CMN96, HX16, DKP<sup>+</sup>16] also use variants of some of these facts, which are closely related to the preservation of the spanning tree distribution on  $\text{Sc}(\mathbf{L}, V_1)$ . (See Section 7 for details.)

**Fact 2.6.** *Let  $V_1$  be a subset of vertices of a graph  $G$ , then for any vertices  $u, v \in V_1$ , we have:*

$$\mathcal{R}_{\text{eff}}^G(u, v) = \mathcal{R}_{\text{eff}}^{\text{Sc}(G, V_1)}(u, v).$$

**Theorem 2.7** (Burton and Pemantle [BP93]). *For any set of edges  $F \subseteq E$  in a graph  $G = (V, E, \mathbf{w})$ , the probability  $F$  is contained in a  $\mathbf{w}$ -uniform random spanning tree is*

$$\mathbf{Pr}_T^G(F \subseteq T) = \det(\mathbf{M}_{(L, F)}),$$

where  $\mathbf{M}_{(L, F)}$  is a  $|F| \times |F|$  matrix whose  $(e, f)$ ’th entry, for  $e, f \in F$ , is given by  $\sqrt{\mathbf{w}(e)\mathbf{w}(f)}\chi_e^T \mathbf{L}^\dagger \chi_f$ .

By a standard property of Schur complements (see [HJ12]), we have

$$(\mathbf{L}^\dagger)_{[V_1, V_1]} = \text{Sc}(G, V_1)^\dagger.$$

Here  $(\mathbf{L}^\dagger)_{[V_1, V_1]}$  is the minor of  $\mathbf{L}^\dagger$  with row and column indices in  $V_1$ . This immediately implies that when  $F$  is incident only on vertices in  $V_1$ , we have  $\mathbf{M}_{(L, F)} = \mathbf{M}_{(\text{Sc}(G, V_1), F)}$ . Putting these together, we have

**Fact 2.8.** *Given a partition of the vertices  $V = V_1 \sqcup V_2$ . For any set of edges  $F$  contained in  $G[V_1]$ , we have*

$$\mathbf{Pr}_T^G(F \subseteq T) = \mathbf{Pr}_T^{\text{SC}(G, V_1)}(F \subseteq T).$$

### 3 Sketch of the Results

The starting point for us is the paper by Janson [Jan94] which gives (among other things) the limiting distribution of the number of spanning trees in the  $\mathcal{G}_{n,m}$  model of random graphs. Our concentration result for the number of spanning trees in the sparsified graph is inspired by this paper, and our algorithmic use of this sparsification routine is motivated by sparsification based algorithms for matrices related to graphs [PS14, CCL<sup>+</sup>15, KLP<sup>+</sup>16]. The key result we will prove is a concentration bound on the number of spanning trees when the graph is sparsified by sampling edges with probability approximately proportional to effective resistance.

#### 3.1 Concentration Bound

Let  $G$  be a weighted graph with  $n$  vertices and  $m$  edges, and  $H$  be a random subgraph obtained by choosing a subset of edges of size  $s$  uniformly randomly. The probability of a subset of edges, which could either be a single tree, or the union of several trees, being kept in  $H$  can be bounded precisely. Since we will eventually choose  $s > n^{1.5}$ , we will treat the quantity  $n^3/s^2$  as negligible. The probability of  $H$  containing a fixed tree was shown by Janson to be:

**Lemma 3.1.** *If  $m \geq \frac{s^2}{n}$ , then for any tree  $T$ , the probability of it being included in  $H$  is*

$$\Pr_H[T \in H] = \frac{(s)_{n-1}}{(m)_{n-1}} = p^{n-1} \cdot \exp\left(-\frac{n^2}{2s} - O\left(\frac{n^3}{s^2}\right)\right).$$

where  $(a)_b$  denotes the product  $a \cdot (a-1) \cdots (a-(b-1))$ .

By linearity of expectation, the expected total weight of spanning trees in  $H$  is:

$$\mathbb{E}_H[\mathcal{T}_H] = \mathcal{T}_G \cdot p^{n-1} \cdot \exp\left(-\frac{n^2}{2s} - O\left(\frac{n^3}{s^2}\right)\right). \quad (1)$$

As in [Jan94], the second moment,  $\mathbb{E}_H[\mathcal{T}_H^2] = \mathbb{E}_H\left[\sum_{(T_1, T_2)} \mathbf{w}(T_1)\mathbf{w}(T_2)\mathbf{Pr}(T_1, T_2 \in H)\right]$ , can be written as a sum over all pairs of trees  $(T_1, T_2)$ . Due to symmetry, the probability of a particular pair of trees  $T_1, T_2$  both being subgraphs of  $H$  depends only on the size of their intersection. The following bound is shown in Appendix A.

**Lemma 3.2.** *Let  $G$  be a graph with  $n$  vertices and  $m$  edges, and  $H$  be a uniformly random subset of  $s > 10n$  edges chosen from  $G$ , where  $m \geq \frac{s^2}{n}$ . Then for any two spanning trees  $T_1$  and  $T_2$  of  $G$  with  $|T_1 \cap T_2| = k$ , we have:*

$$\Pr_H[T_1, T_2 \in H] \leq p^{2n-2} \exp\left(-\frac{2n^2}{s}\right) \left(\frac{1}{p} \left(1 + \frac{2n}{s}\right)\right)^k,$$

where  $p = s/m$ .

The crux of the bound on the second moment in Janson's proof is getting a handle on the number of tree pairs  $(T_1, T_2)$  with  $|T_1 \cap T_2| = k$  in the complete graph where all edges are symmetric. An alternate way to obtain a bound on the number of spanning trees can also be obtained using leverage scores, which describe the fraction of spanning trees that utilize a single edge. A well known fact about random spanning tree distributions [BP93] is that the edges are negatively correlated:

**Fact 3.3** (Negative Correlation). *Suppose  $F$  is subset of edges in a graph  $G$ , then*

$$\mathbf{Pr}_T^G(F \subseteq T) \leq \prod_{e \in F} \mathbf{Pr}_T^G(e \in T).$$

An easy consequence of Fact 3.3 is

**Lemma 3.4.** *For any subset of edges  $F$  we have that the total weight of all spanning trees containing  $F$  is given by*

$$\sum_{\substack{T \text{ is a spanning tree of } G \\ F \subseteq T}} \mathbf{w}(T) \leq \mathcal{T}_G \prod_{e \in F} \bar{\tau}_e.$$

The combinatorial view of all edges being interchangeable in the complete graph can therefore be replaced with an algebraic view in terms of the leverage scores. Specifically, invoking Lemma 3.4 in the case where all edges have leverage score at most  $\frac{n}{m}$  gives the following lemma which is proven in Appendix A.

**Lemma 3.5.** *In a graph  $G$  where all edges have leverage scores at most  $\frac{n}{m}$ , we have*

$$\sum_{\substack{T_1, T_2 \\ |T_1 \cap T_2|=k}} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2) \leq \mathcal{T}_G^2 \cdot \frac{1}{k!} \left( \frac{n^2}{m} \right)^k$$

With Lemma 3.5, we can finally prove the following bound on the second moment which gives our concentration result.

**Lemma 3.6.** *Let  $G$  be a graph on  $n$  vertices and  $m$  edges such that all edges have statistical leverage scores  $\leq \frac{n}{m}$ . For a random subset of  $s > 10n$  edges,  $H$ , where  $m \geq \frac{s^2}{n}$  we have:*

$$\mathbb{E}_H [\mathcal{T}_H^2] \leq \mathcal{T}_G^2 p^{2n-2} \exp \left( -\frac{n^2}{s} + O \left( \frac{n^3}{s^2} \right) \right) = \mathbb{E}_H [\mathcal{T}_H]^2 \exp \left( O \left( \frac{n^3}{s^2} \right) \right).$$

*Proof.* By definition of the second moment, we have:

$$\mathbb{E}_H [\mathcal{T}_H^2] = \sum_{T_1, T_2} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2) \cdot \Pr_H [T_1 \cup T_2 \subseteq H].$$

Re-writing the above sum in terms of the size of the intersection  $k$ , and invoking Lemma 3.2 gives:

$$\mathbb{E}_H [\mathcal{T}_H^2] \leq \sum_{k=0}^{n-1} \sum_{\substack{T_1, T_2 \\ |T_1 \cap T_2|=k}} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2) \cdot p^{2n-2} \exp \left( -\frac{2n^2}{s} \right) \left( \frac{1}{p} \left( 1 + \frac{2n}{s} \right) \right)^k.$$

Note that the trailing term only depends on  $k$  and can be pulled outside the summation of  $T_1, T_2$ , so we then use Lemma 3.5 to bound this by:

$$\mathbb{E}_H [\mathcal{T}_H^2] \leq \sum_{k=0}^{n-1} \mathcal{T}_G^2 \cdot \frac{1}{k!} \left( \frac{n^2}{m} \right)^k \cdot p^{2n-2} \exp \left( -\frac{2n^2}{s} \right) \left( \frac{1}{p} \left( 1 + \frac{2n}{s} \right) \right)^k.$$

Which upon pulling out the terms that are independent of  $k$ , and substituting in  $p = s/m$  gives:

$$\mathbb{E}_H [\mathcal{T}_H^2] \leq \mathcal{T}_G^2 \cdot p^{2n-2} \cdot \exp \left( -\frac{2n^2}{s} \right) \cdot \sum_{k=0}^{n-1} \frac{1}{k!} \cdot \left( \frac{n^2}{s} \left( 1 + \frac{2n}{s} \right) \right)^k.$$

From the Taylor expansion of  $\exp(\cdot)$ , we have:

$$\begin{aligned} \mathbb{E}_H [\mathcal{T}_H^2] &\leq \mathcal{T}_G^2 \cdot p^{2n-2} \cdot \exp \left( -\frac{2n^2}{s} \right) \cdot \exp \left( \frac{n^2}{s} \left( 1 + \frac{2n}{s} \right) \right) \\ &= \mathcal{T}_G^2 \cdot p^{2n-2} \cdot \exp \left( -\frac{n^2}{s} \right) \cdot \exp \left( O \left( \frac{n^3}{s^2} \right) \right). \end{aligned}$$

□

This bound implies that once we set  $s^2 > n^3$ , the variance becomes less than the square of the expectation. It forms the basis of our key concentration results, which we show in Section 4, and also leads to Theorem 1.1. In particular, we demonstrate that this sampling scheme extends to importance sampling, where edges are picked with probabilities proportional to (approximations of) of their leverage scores.

A somewhat surprising aspect of this concentration result is that there is a difference between models  $\mathcal{G}_{n,m}$  and the Erdos-Renyi model  $\mathcal{G}_{n,p}$  when the quantity of interest is the number of spanning trees. In particular, the number of spanning trees of a graph  $G \sim \mathcal{G}_{n,m}$  is approximately normally distributed when  $m = \omega(n^{1.5})$ , whereas it has approximate log-normal distribution when  $G \sim \mathcal{G}_{n,p}$  and  $p < 1$ .

An immediate consequence of this is that we can now approximate  $\det_+(\mathbf{L}^G)$  by computing  $\det_+(\mathbf{L}^H)$ . It also becomes natural to consider speedups of random spanning tree sampling algorithms that generate a spanning tree from a sparsifier. Note however that we cannot hope to preserve the distribution over all spanning trees via a single sparsifier, as some of the edges are no longer present.

To account for this change in support, we instead consider the randomness used in generating the sparsifier as also part of the randomness needed to produce spanning trees. In Section 9.1, we show that just bounds on the variance of  $\mathcal{T}_H$  suffices for a bound on the TV distances of the trees.

**Lemma 3.7.** *Suppose  $\mathcal{H}$  is a distribution over rescaled subgraphs of  $G$  such that for some parameter some  $0 < \delta < 1$  we have*

$$\frac{\mathbb{E}_{H \sim \mathcal{H}} [\mathcal{T}_H^2]}{\mathbb{E}_{H \sim \mathcal{H}} [\mathcal{T}_H]^2} \leq 1 + \delta,$$

and for any tree  $\hat{T}$  and any graph from the distribution that contain it,  $H$  we have:

$$\mathbf{w}^H(\hat{T}) = \mathbf{w}^G(\hat{T}) \cdot \Pr_{H' \sim \mathcal{H}}[\hat{T} \subseteq H']^{-1} \cdot \frac{\mathbb{E}_{H' \sim \mathcal{H}}[\mathcal{T}_{H'}]}{\mathcal{T}_G},$$

then the distribution given by  $\mathbf{Pr}^G(T)$ ,  $p$ , and the distribution induced by  $\mathbb{E}_{H \sim \mathcal{H}}[\mathbf{Pr}^H(T)]$ ,  $\tilde{p}$  satisfies

$$d_{TV}(p, \tilde{p}) \leq \sqrt{\delta}.$$

Note that uniform sampling meets the property about  $\mathbf{w}^H(T)$  because of linearity of expectation. We can also check that the importance sampling based routine that we will discuss in Section 4.2 also meets this criteria. Combining this with the running time bounds from Theorem 1.1, as well as the  $\tilde{O}(m^{1/3}n^{5/3})$  time random spanning tree sampling algorithm from [DKP<sup>+</sup>16] then leads to a faster algorithm.

**Corollary 3.8.** *For any graph  $G$  on  $n$  vertices and any  $\delta > 0$ , there is an algorithm that generates a tree from a distribution whose total variation is at most  $\delta$  from the random tree distribution of  $G$  in time  $\tilde{O}(n^{\frac{13}{6}=2.1666\ldots}\delta^{-2/3} + n^2\delta^{-2})$ .*

### 3.2 Integration Into Recursive Algorithms

As a one-step invocation of our concentration bound leads to speedups over previous routines, we investigate tighter integrations of the sparsification routine into algorithms. In particular, the sparsified Schur complement algorithms [KLP<sup>+</sup>16] provide a natural place to substitute spectral sparsifiers with determinant-preserving ones. In particular, the identity of

$$\det_+(\mathbf{L}) = \det(\mathbf{L}_{[V_2, V_2]}) \cdot \det_+(\text{SC}(\mathbf{L}, V_1)).$$

where  $\det_+$  is the determinant of the matrix minor, suggests that we can approximate  $\det(\mathbf{L}_{-n})$  by approximating  $\det(\mathbf{L}_{[V_2, V_2]})$  and  $\det_+(\text{SC}(\mathbf{L}, V_1))$  instead. Both of these subproblems are smaller by a constant factor, and we also have  $|V_1| + |V_2| = n$ . So this leads to a recursive scheme where the total number of vertices involved at all layers is  $O(n \log n)$ . This type of recursion underlies both our determinant estimation and spanning tree sampling algorithms.

The main difficulty remaining for the determinant estimation algorithm is then sparsifying  $\text{SC}(G, V_1)$  while preserving its determinant. For this, we note that some  $V_1$  are significantly easier than others: in particular, when  $V_2 = V \setminus V_1$  is an independent set, the Schur complement of each of the vertices in  $V_2$  can be computed independently. Furthermore, it is well understood how to sample these complements, which are weighted cliques, by a distribution that exceeds their true leverage scores.

**Lemma 3.9.** *There is a procedure that takes a graph  $G$  with  $n$  vertices, a parameter  $\delta$ , and produces in  $\tilde{O}(n^2\delta^{-1})$  time a subset of vertices  $V_1$  with  $|V_1| = \Theta(n)$ , along with a graph  $H^{V_1}$  such that*

$$\mathcal{T}_{\text{SC}(G, V_1)} \exp(-\delta) \leq \mathbb{E}_{H^{V_1}}[\mathcal{T}_{H^{V_1}}] \leq \mathcal{T}_{\text{SC}(G, V_1)} \exp(\delta),$$

and

$$\frac{\mathbb{E}_{H^{V_1}}[\mathcal{T}_{H^{V_1}}^2]}{\mathbb{E}_{H^{V_1}}[\mathcal{T}_{H^{V_1}}]^2} \leq \exp(\delta).$$

Lemma 2.4 holds w.h.p., and we condition on this event. In our algorithmic applications we will be able to add the polynomially small failure probability of Lemma 2.4 to the error bounds.

The bound on variance implies that the number of spanning trees is concentrated close to its expectation,  $\mathcal{T}_{\text{SC}(G, V_1)}$ , and that a random spanning tree drawn from the generated graph  $H^{V_1}$  is—over the randomness of the sparsification procedure—close in total variation distance to a random spanning tree of the true Schur complement.

As a result, we can design schemes that:

1. Finds an  $O(1)$ -DD subset  $V_2$ , and set  $V_1 \leftarrow V \setminus V_2$ .
2. Produce a determinant-preserving sparsifier  $H^{V_1}$  for  $\text{SC}(G, V_1)$ .
3. Recurse on both  $\mathbf{L}_{[V_2, V_2]}$  and  $H^{V_1}$ .

However, in this case, the accumulation of error is too rapid for yielding a good approximation of determinants. Instead, it becomes necessary to track the accumulation of variance during all recursive calls. Formally, the cost of sparsifying so that the variance is at most  $\delta$  is about  $n^2\delta^{-1}$ , where  $\delta$  is the size of the problem. This means that for a problem on  $G_i$  of size  $\beta_i n$  for  $0 \leq \beta_i \leq 1$ , we can afford an error of  $\beta_i \delta$  when working with it, since:

1. The sum of  $\beta_i$  on any layer is at most 2,<sup>5</sup> so the sum of variance per layer is  $O(\delta)$ .
2. The cost of each sparsification step is now  $\beta_i n^2\delta^{-1}$ , which sums to about  $n^2\delta^{-1}$  per layer.

Our random spanning tree sampling algorithm in Section 7 is similarly based on this careful accounting of variance. We first modify the recursive Schur complement algorithm introduced by Coulburn et al. [CDN89] to give a simpler algorithm that only branches two ways at each step in Section 7.1, leading to a high level scheme fairly similar to the recursive determinant algorithm. Despite these similarities, the accumulation of errors becomes far more involved here due to the choice of trees in earlier recursive calls affecting the graph in later steps. More specifically, the recursive structure of our determinant algorithm can be considered analogous to a breadth-first-search, which allows us to consider all subgraphs at each layer to be independent. In contrast, the recursive structure of our random spanning tree algorithm, which we show in Section 7.2 is more analogous to a depth-first traversal of the tree, where the output solution of one subproblem will affect the input of all subsequent subproblems.

These dependency issues will be the key difficulty in considering variance across levels. The total variation distance tracks the discrepancy over all trees of  $G$  between their probability of being returned by the overall recursive algorithm, and their probability in the  $\mathbf{w}$ -uniform distribution. Accounting for this over all trees leads us to bounding variances in the probabilities of individual trees being picked. As this is, in turn, is equivalent to the weight of the tree divided by the determinant of the graph, the inverse of the probability of a tree being picked can play a similar role to the determinant in the determinant sparsification algorithm described above. However, tracking this value requires analyzing extending our concentration bounds to the case where an arbitrary tree is fixed in the graph and we sample from the remaining edges. We study this in Section 8, prove bounds analogous to the concentration bounds from Section 4, and incorporate the guarantees back into the recursive algorithm in Section 7.2.

---

<sup>5</sup>each recursive call may introduce one new vertex

## 4 Determinant Preserving Sparsification

In this section we will ultimately prove Theorem 1.1, our primary result regarding determinant-preserving sparsification. However, most of this section will be devoted to proving the following general determinant-preserving sparsification routine that also forms the core of subsequent algorithms:

**Theorem 4.1.** *Given an undirected, weighted graph  $G = (V, E, \mathbf{w})$ , an error threshold  $\epsilon > 0$ , parameter  $\rho$  along with routines:*

1. *SAMPLEEDGE $_G()$  that samples an edge  $e$  from a probability distribution  $\mathbf{p}$  ( $\sum_e \mathbf{p}_e = 1$ ), as well as returning the corresponding value of  $\mathbf{p}_e$ . Here  $\mathbf{p}_e$  must satisfy:*

$$\frac{\bar{\tau}_e}{n-1} \leq \rho \cdot \mathbf{p}_e$$

*where  $\bar{\tau}_e$  is the true leverage score of  $e$  in  $G$ .*

2. *APPROXLEVERAGE $_G(u, v, \epsilon)$  that returns the leverage score of an edge  $u, v$  in  $G$  to an error of  $\epsilon$ . Specifically, given an edge  $e$ , it returns a value  $\tilde{\tau}_e$  such that:*

$$(1 - \epsilon) \bar{\tau}_e \leq \tilde{\tau}_e \leq (1 + \epsilon) \bar{\tau}_e.$$

*There is a routine DETSPARSIFY( $G, s, \epsilon$ ) that computes a graph  $H$  with  $s$  edges such that its tree count,  $\mathcal{T}_H$ , satisfies:*

$$\mathbb{E}_H [\mathcal{T}_H] = \mathcal{T}_G \left( 1 \pm O \left( \frac{n^3}{s^2} \right) \right),$$

*and:*

$$\frac{\mathbb{E}_H [\mathcal{T}_H^2]}{\mathbb{E}_H [\mathcal{T}_H]^2} \leq \exp \left( \frac{\epsilon^2 n^2}{s} + O \left( \frac{n^3}{s^2} \right) \right)$$

*Furthermore, the expected running time is bounded by:*

1.  *$O(s \cdot \rho)$  calls to SAMPLEEDGE $_G(e)$  and APPROXLEVERAGE $(e)$  with constant error,*
2.  *$O(s)$  calls to APPROXLEVERAGE $(e)$  with  $\epsilon$  error.*

We establish guarantees for this algorithm using the following steps:

1. Showing that the concentration bounds as sketched in Section 3 holds for approximate leverage scores in Section 4.1.
2. Show via taking the limit of probabilistic processes that the analog of this process works for sampling a general graph where edges can have varying leverage scores. This proof is in Section 4.2.
3. Show via rejection sampling that (high error) one sided bounds on statistical leverage scores, such as those that suffice for spectral sparsification, can also be used to do the initial round of sampling instead of two-sided approximations of leverage scores. This, as well as pseudocode and guarantees of the overall algorithm are given in Section 4.3.

## 4.1 Concentration Bound with Approximately Uniform Leverage Scores

Similar to the simplified proof as outlined in Section 3, our proofs relied on uniformly sampling  $s$  edges from a multi-graph with  $m \geq \frac{s^2}{n}$  edges, such that all edges have leverage score within multiplicative  $1 \pm \epsilon$  of  $\frac{n-1}{s}$ , aka. approximately uniform. The bound that we prove is an analog of Lemma 3.6

**Lemma 4.2.** *Given a weighted multi-graph  $G$  such that  $m \geq \frac{s^2}{n}$ ,  $s \geq n$ , and all edges  $e \in E$  have  $\frac{(1-\epsilon)(n-1)}{m} \leq \bar{\tau}_e \leq \frac{(1+\epsilon)(n-1)}{m}$ , with  $0 \leq \epsilon < 1$ , then*

$$\frac{\mathbb{E}_H [\mathcal{T}_H^2]}{\mathbb{E}_H [\mathcal{T}_H]^2} \leq \exp \left( \frac{n^2 \epsilon^2}{s} + O \left( \frac{n^3}{s^2} \right) \right)$$

Similar to the proof of Lemma 3.6 in Section 3, we can utilize the bounds on the probability of  $k$  edges being chosen using Lemma 3.2. The only assumption that changed was the bounds on  $\bar{\tau}_e$ , which does not affect  $\mathbb{E}_H [\mathcal{T}_H]^2$ . The only term that changes is our upper bound the total weight of trees that contain some subset of  $k$  edges that was the produce of  $k$  leverage scores. At a glance, this product can change by a factor of up to  $(1 + \epsilon)^k$ , which when substituted naively into the proof of Lemma 3.2 directly would yield an additional term of

$$\exp \left( \frac{n^2 \epsilon}{s} \right),$$

and in turn necessitating  $\epsilon < n^{-1/2}$  for a sample count of  $s \approx n^{1.5}$ .

However, note that this is the worst case distortion over a subset  $F$ . The upper bound that we use, Lemma 3.5 sums over these bounds over all subsets, and over all edges we still have

$$\sum e \in G \bar{\tau}_e = n - 1.$$

Incorporating this allows us to show a tighter bound that depends on  $\epsilon^2$ .

Similar to the proof of Lemma 3.5, we can regroup the summation over all  $\binom{m}{k}$  subsets of  $E(G)$ , and bound the fraction of trees containing each subset  $F$  via  $\sum_{T: F \subseteq T} \mathbf{w}(T) \leq \mathcal{T}_G \prod_{e \in F} \tau_e$  via Lemma 3.4.

$$\sum_{\substack{T_1, T_2 \\ |T_1 \cap T_2| = k}} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2) \leq \sum_{\substack{F \subseteq E \\ |F| = k}} \mathcal{T}_G^2 \prod_{e \in F} \bar{\tau}_e^2$$

The proof will heavily utilize the fact that  $\sum_{e \in E} \bar{\tau}_e = n - 1$ . We bound this in first two steps: first treat it as a symmetric product over  $\bar{\tau}_e^2$ , and bound the total as a function of

$$\sum_e \bar{\tau}_e^2,$$

then we bound this sum using the fact that  $\sum_e \bar{\tau}_e = n - 1$ .

The first step utilizes the concavity of the product function, and bound the total by the sum:

**Lemma 4.3.** *For any set of non-negative values  $\mathbf{x}_1 \dots \mathbf{x}_m$  with  $\sum_i \mathbf{x}_i \leq z$ , we have*

$$\sum_{\substack{F \subseteq [1 \dots m] \\ |F| = k}} \prod_{i \in F} \mathbf{x}_i \leq \binom{m}{k} \left( \frac{z}{m} \right)^k.$$

*Proof.* We claim that this sum is maximized when  $\mathbf{x}_i = (\frac{z}{m})$  for all  $i$ .

Consider fixing all variables other than some  $\mathbf{x}_i$  and  $\mathbf{x}_j$ , which we assume to be  $\mathbf{x}_1 \leq \mathbf{x}_2$  without loss of generality as the function is symmetric on all variables:

$$\sum_{\substack{F \subseteq [1 \dots m] \\ |F|=k}} \prod_{i \in F} \mathbf{x}_i = \mathbf{x}_1 \mathbf{x}_2 \left( \sum_{\substack{F \subseteq [3 \dots m] \\ |F|=k-2}} \prod_{i \in F} \mathbf{x}_i \right) + (\mathbf{x}_1 + \mathbf{x}_2) \cdot \left( \sum_{\substack{F \subseteq [3 \dots m] \\ |F|=k-1}} \prod_{i \in F} \mathbf{x}_i \right) + \sum_{\substack{F \subseteq [3 \dots m] \\ |F|=k}} \prod_{i \in F} \mathbf{x}_i.$$

Then if  $\mathbf{x}_1 < \mathbf{x}_2$ , locally changing their values to  $\mathbf{x}_1 + \epsilon$  and  $\mathbf{x}_2 - \epsilon$  keeps the second term the same. While the first term becomes

$$(\mathbf{x}_1 + \epsilon)(\mathbf{x}_2 - \epsilon) = \mathbf{x}_1 \mathbf{x}_2 + \epsilon(\mathbf{x}_2 - \mathbf{x}_1) - \epsilon^2,$$

which is greater than  $\mathbf{x}_1 \mathbf{x}_2$  when  $0 < \epsilon < (\mathbf{x}_2 - \mathbf{x}_1)$ .

This shows that the overall summation is maximized when all  $\mathbf{x}_i$  are equal, aka

$$\mathbf{x}_i = \frac{z}{m},$$

which upon substitution gives the result.  $\square$

The second step is in fact the  $k=1$  case of Lemma 3.5.

**Lemma 4.4.** *For any set of values  $\mathbf{y}_e$  such that*

$$\sum_e \mathbf{y}_e = n - 1,$$

and

$$\frac{(1 - \epsilon) n}{m} \leq \mathbf{y}_e \leq \frac{(1 + \epsilon) n}{m},$$

we have

$$\sum_e \mathbf{y}_e^2 \leq \frac{(1 + \epsilon^2)(n - 1)^2}{m}.$$

*Proof.* Note that for any  $a \leq b$ , and any  $\epsilon$ , we have

$$(a - \epsilon)^2 + (b + \epsilon)^2 = a^2 + b^2 + 2\epsilon^2 + 2\epsilon(b - a),$$

and this transformation must increase the sum for  $\epsilon > 0$ . This means the sum is maximized when half of the leverage scores are  $\frac{(1-\epsilon)(n-1)}{m}$  and the other half are  $\frac{(1+\epsilon)(n-1)}{m}$ . This then gives

$$\sum_{e \in E} \mathbf{y}_e^2 \leq \frac{m}{2} \left( \frac{(1 + \epsilon)(n - 1)}{m} \right)^2 + \frac{m}{2} \left( \frac{(1 - \epsilon)(n - 1)}{m} \right)^2 = \frac{(1 + \epsilon^2)(n - 1)^2}{m}.$$

$\square$

*Proof.* (of Lemma 4.2)

We first derive an analog of Lemma 3.5 for bounding the total weights of pairs of trees containing subsets of size  $k$ , where we again start with the bounds

$$\sum_{\substack{T_1, T_2 \\ |T_1 \cap T_2|=k}} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2) \leq \sum_{\substack{F \subseteq E \\ |F|=k}} \sum_{\substack{T_1, T_2 \\ F \subseteq T_1 \cap T_2}} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2) = \sum_{\substack{F \subseteq E \\ |F|=k}} \left( \sum_{T: F \subseteq T} \mathbf{w}(T) \right)^2$$

Applying Lemma 3.4 to the inner term of the summation then gives

$$\sum_{\substack{T_1, T_2 \\ |T_1 \cap T_2|=k}} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2) \leq \sum_{\substack{F \subseteq E \\ |F|=k}} \mathcal{T}_G^2 \cdot \prod_{e \in F} \bar{\tau}_e^2$$

The bounds on  $\bar{\tau}_e$  and  $\sum_e \bar{\tau}_e = n - 1$  gives, via Lemma 4.4

$$\sum_e \bar{\tau}_e^2 \leq \frac{(1 + \epsilon^2)(n - 1)^2}{m}.$$

Substituting this into Lemma 4.3 with  $\mathbf{x}_i = \bar{\tau}_e^2$  then gives

$$\sum_{\substack{F \subseteq E \\ |F|=k}} \prod_{e \in F} \bar{\tau}_e^2 \leq \binom{m}{k} \left( \frac{(1 + \epsilon^2)n^2}{m^2} \right)^k \leq \frac{m^k}{k!} \left( \frac{(1 + \epsilon^2)n^2}{m^2} \right)^k = \frac{1}{k!} \left( \frac{(1 + \epsilon^2)n^2}{m} \right)^k.$$

which implies our analog of Lemma 3.5

$$\sum_{\substack{T_1, T_2 \\ |T_1 \cap T_2|=k}} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2) \leq \mathcal{T}_G^2 \cdot \frac{1}{k!} \left( \frac{(1 + \epsilon^2)n^2}{m} \right)^k.$$

We can then duplicate the proof of Lemma 3.6. Similar to that proof, we can regroup the summation by  $k = |T_1 \cap T_2|$  and invoking Lemma 3.2 to get:

$$\mathbb{E}_H [\mathcal{T}_H^2] \leq \sum_{k=0}^{n-1} \sum_{\substack{T_1, T_2 \\ |T_1 \cap T_2|=k}} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2) \cdot p^{2n-2} \exp \left( -\frac{2n^2}{s} \right) \left( \frac{1}{p} \left( 1 + \frac{2n}{s} \right) \right)^k.$$

where  $p = s/m$ . When incorporated with our analog of Lemma 3.5 gives:

$$\begin{aligned} \mathbb{E}_H [\mathcal{T}_H^2] &\leq \sum_{k=0}^{n-1} p^{2n-2} \exp \left( -\frac{2n^2}{s} \right) \left( \frac{1}{p} \left( 1 + \frac{2n}{s} \right) \right)^k \cdot \mathcal{T}_G^2 \frac{1}{k!} \left( \frac{(1 + \epsilon^2)n^2}{m} \right)^k \\ &= \mathcal{T}_G^2 p^{2n-2} \cdot \exp \left( -\frac{2n^2}{s} \right) \cdot \sum_{k=0}^{n-1} \frac{1}{k!} \cdot \left( \frac{(1 + \epsilon^2)n^2}{s} \left( 1 + \frac{2n}{s} \right) \right)^k. \end{aligned}$$

Substituting in the Taylor expansion of  $\sum_k \frac{z^k}{k!} \leq \exp(z)$  then leaves us with:

$$\mathbb{E}_H [\mathcal{T}_H^2] \leq \mathcal{T}_G^2 \cdot p^{2n-2} \cdot \exp \left( -\frac{n^2}{s} + \frac{n^2 \epsilon^2}{s} + O \left( \frac{n^3}{s^2} \right) \right)$$

and finishes the proof.  $\square$

## 4.2 Generalization to Graphs with Arbitrary Leverage Score Distributions

The first condition of  $m \geq \frac{s^2}{n}$  will be easily achieved by splitting each edge a sufficient number of times, which does not need to be done explicitly in the sparsification algorithm. Furthermore, from the definition of statistical leverage score splitting an edge into  $k$  copies will give each copy a  $k$ th fraction of the edge's leverage score. Careful splitting can then ensure the second condition, but will require  $\epsilon$ -approximate leverage score estimates on the edges. The simple approach would compute this for all edges, then split each edge according to this estimate and draw from the resulting edge set. Instead, we only utilize this algorithm as a proof technique, and give a sampling scheme that's equivalent to this algorithm's limiting behavior as  $m \rightarrow \infty$ . Pseudocode of this routine is in Algorithm 1.

---

**Algorithm 1:** IDEALSPARSIFY( $G, \tilde{\tau}, s$ ): Sample  $s$  (multi) edges of  $G$  to produce  $H$  such that  $\mathcal{T}_G \approx \mathcal{T}_H$ .

---

**Input:** Graph  $G$ , approximate leverage scores  $\tilde{\tau}$ , sample count  $s$

- 1 Initialize  $H$  as the empty graph,  $H \leftarrow \emptyset$ ;
  - 2 **for**  $i = 1 \dots s$  **do**
  - 3     Pick edge  $e$  with probability proportional to  $\tilde{\tau}_e$ ;
  - 4     Add  $e$  to  $H$  with new weight:
  - 5
- $$\frac{w_e(n-1)}{\tilde{\tau}_e s} \exp \left( \frac{n^2}{2(n-1)s} \right).$$
- 5 Output  $H$
- 

Note that this sampling scheme is with replacement: the probability of a ‘collision’ as the number of copies tend to  $\infty$  is sufficiently small that it can be covered by the proof as well.

The guarantee that we will show for Algorithm 1 is:

**Lemma 4.5.** *For any graph  $G$  and any set of approximate leverage scores  $\tilde{\tau}$  such that*

$$(1 - \epsilon) \tau_e \leq \tilde{\tau}_e \leq (1 + \epsilon) \tau_e$$

*for all edges  $e$ . The graph  $H = \text{IDEALSPARSIFY}(G, \tilde{\tau}, s)$  satisfies:*

$$\left( 1 - O \left( \frac{n^3}{s^2} \right) \right) \mathcal{T}_G \leq \mathbb{E}_H [\mathcal{T}_H] \leq \mathcal{T}_G,$$

and

$$\frac{\mathbb{E}_H [\mathcal{T}_H^2]}{\mathbb{E}_H [\mathcal{T}_H]^2} \leq \exp \left( O \left( \frac{\epsilon^2 n^2}{s} + \frac{n^3}{s^2} \right) \right).$$

Our proof strategy is simple: claim that this algorithm is statistically close to simulating splitting each edge into a very large number of copies. Note that these proofs are purely for showing the convergence of statistical processes, so all that's needed is for the numbers that arise in this proof (in particular,  $m$ ) to be finite.

We first show that  $G$  and  $\tilde{\tau}$  can be perturbed to become rational numbers.

**Lemma 4.6.** *For any graph  $G$  and any set of  $\tilde{\tau}$  such that  $(1 - \epsilon)\bar{\tau}_e^{(G)} \leq \tilde{\tau}_e \leq (1 + \epsilon)\bar{\tau}_e^{(G)}$  for all edges  $e$  for some constant  $\epsilon > 0$ , and any perturbation threshold  $\delta$ , we can find graph  $G'$  with all edge weights rationals, and  $\tilde{\tau}'$  with all entries rational numbers such that:*

1.  $\mathcal{T}_G \leq \mathcal{T}_{G'} \leq (1 + \delta)\mathcal{T}_G$ , and
2.  $(1 - 2\epsilon)\bar{\tau}_e^{(G')} \leq \tilde{\tau}'_e \leq (1 + 2\epsilon)\bar{\tau}_e^{(G')}$  for all edges  $e$ .

*Proof.* This is a direct consequence of the rational numbers being everywhere dense, and that perturbing edge weights by a factor of  $1 \pm \alpha$  perturbs leverage scores by a factor of up to  $1 \pm O(\alpha)$ , and total weights of trees by a factor of  $(1 \pm \alpha)^{n-1}$ .  $\square$

Having all leverage scores as integers means that we can do an exact splitting by setting  $m$ , the total number of split edges, to a multiple of the common denominator of all the  $\tilde{\tau}'_e$  values times  $n - 1$ . Specifically, an edge with approximate leverage score  $\tilde{\tau}'_e$  becomes

$$\tilde{\tau}'_e \cdot \frac{m}{n - 1}$$

copies, each with weight

$$\frac{w_e(n-1)}{\tilde{\tau}'_e m},$$

and ‘true’ leverage score

$$\frac{\bar{\tau}_e(n-1)}{\tilde{\tau}'_e m}.$$

In particular, since

$$(1 - 2\epsilon) \leq \frac{\bar{\tau}_e}{\tilde{\tau}'_e} \leq (1 + 2\epsilon),$$

this splitted graph satisfies the condition of Lemma 4.2. This then enables us to obtain the guarantees of Lemma 4.5 by once again letting  $m$  tend to  $\infty$ .

*Proof.* (of Lemma 4.5) We first show that Algorithm 1 works for the graph with rational weights and approximate leverage scores as generated by Lemma 4.6.

The condition established above means that we can apply Lemma 4.2 to the output of picking  $s$  random edges among these  $m$  split copies. This graph  $H'$  satisfies

$$\mathbb{E}_{H'}[\mathcal{T}_{H'}] = \mathcal{T}_{G'} \left( \frac{s}{m} \right)^{n-1} \exp \left( -\frac{n^2}{2s} - O \left( \frac{n^3}{s^2} \right) \right),$$

and

$$\frac{\mathbb{E}_{H'}[\mathcal{T}_{H'}^2]}{\mathbb{E}_{H'}[\mathcal{T}_{H'}]^2} \leq \exp \left( \frac{n^2 \epsilon^2}{s} + O \left( \frac{n^3}{s^2} \right) \right).$$

The ratio of the second moment is not affected by rescaling, so the graph

$$H'' \leftarrow \frac{m}{s} \exp \left( \frac{n^2}{2s(n-1)} \right)$$

meets the requirements on both the expectation and variances. Furthermore, the rescaled weight of a single edge being picked is:

$$\frac{\mathbf{w}_e(n-1)}{\tilde{\tau}'_e m} \cdot \frac{m}{s} \exp \left( \frac{n^2}{2s(n-1)} \right) = \frac{\mathbf{w}_e(n-1)}{\tilde{\tau}'_e s} \exp \left( \frac{n^2}{2s(n-1)} \right),$$

which is exactly what Algorithm 1 assigns.

It remains to resolve the discrepancy between sampling with and without replacement: the probability of the same edge being picked twice in two different steps is at most  $1/m$ , so the total probability of a duplicate sample is bounded by  $s^2/m$ . We then give a finite bound on the size of  $m$  for which this probability becomes negligible in our routine. The rescaling factor of a single edge is (very crudely) bounded by

$$\frac{(n-1)}{\tilde{\tau}'_e s} \exp \left( \frac{n^2}{2s(n-1)} \right) \leq \exp(n^3) \frac{1}{\min_e \tilde{\tau}'_e},$$

which means that any of the  $H''$  returned must satisfy

$$\mathcal{T}_{H''} \leq \exp(n^4) \left( \frac{1}{\min_e \tilde{\tau}'_e} \right)^n \mathcal{T}_{G'},$$

which is finite. As a result, as  $m \rightarrow \infty$ , the difference that this causes to both the first and second moments become negligible.

The result for  $H \leftarrow \text{IDEALSPARSIFY}(G, \tilde{\tau}, s)$  then follows from the infinitesimal perturbation made to  $G$ , as the rational numbers are dense everywhere.  $\square$

### 4.3 Incorporating Crude Edge Sampler Using Rejection Sampling

Under Lemma 4.5 we assumed access to  $\epsilon$ -approximate leverage scores, which could be computed with  $m$  calls to our assumed subroutine  $\text{APPROXLEVERAGE}_G$ , where  $m$  here is the number of edges of  $G$ . However, we roughly associate  $\text{APPROXLEVERAGE}_G$  with Lemma 2.4 that requires  $\tilde{O}(\epsilon^{-2})$  time per call (and we deal with the w.h.p. aspect in the proof of Theorem 1.1), and to achieve our desired sparsification of  $O(n^{1.5})$  edges, we will need  $\epsilon = n^{-1/4}$  for the necessary concentration bounds. Instead, we will show that we can use rejection sampling to take  $s$  edges drawn from approximate leverage scores using a cruder distribution  $\mathbf{p}_e$ , which will only require application of  $\text{APPROXLEVERAGE}_G$  with error  $\epsilon$  for an expected  $O(s)$  number of edges.

Rejection sampling is a known technique that allows us to sample from some distribution  $f$  by instead sampling from a distribution  $g$  that approximates  $f$  and accept the sample with a specific probability based on the probability of drawing that sample from  $f$  and  $g$ .

More specifically, suppose we are given two probability distributions  $f$  and  $g$  over the same state space  $X$ , such that for all  $x \in X$  we have  $Cg(x) \geq f(x)$  for some constant  $C$ . Then we can draw from  $f$  by instead drawing  $x \sim g$ , and accepting the draw with probability  $\frac{f(x)}{Cg(x)}$ .

This procedure only requires a lower bound on  $g$  with respect to  $f$ , but in order to accept a draw with constant probability, there need to be weaker upper bound guarantees. Our guarantees on  $\tilde{\tau}_e$  will fulfill these requirements, and the rejection sampling will accept a constant fraction of the draws. By splitting into a sufficient number of edges, we ensure that drawing the same multi-edge from any split edge will occur with at most constant probability.

Specifically, each sample is drawn via the following steps:

1. Draw a sample according the distribution  $g, e$ .
2. Evaluate the values of  $f(e)$  and  $g(e)$ .
3. Keep the sample with probability  $f(e)/g(e)$ .

As the running time of  $\text{APPROXLEVERAGE}_G(e, \epsilon)$  will ultimately depend on the value of  $\epsilon$  apply this algorithmic framework, we also need to perform rejection sampling twice, once with constant error, and once with leverage scores extracted from the true approximate distribution. Pseudocode of this routine is shown in Algorithm 2.

---

**Algorithm 2:**  $\text{DETSPARSIFY}(G, s, \text{SAMPLEEDGE}_G(), \rho, \text{APPROXLEVERAGE}_G(u, v, \epsilon))$ :  
Sample  $s$  (multi) edges of  $G$  to produce  $H$  such that  $\mathcal{T}_G \approx \mathcal{T}_H$ .

---

**Input:** Graph  $G$ .

Sample count  $s$ , leverage score approximation error  $0 < \epsilon < 1/2$ ,

$\text{SAMPLEEDGE}_G()$  that samples an edge  $e$  from a probability distribution  $\mathbf{p}$  ( $\sum_e \mathbf{p}_e = 1$ ), and returning the corresponding value of  $\mathbf{p}_e$ .

$\rho$  that bounds the under-sampling rate of  $\text{SAMPLEEDGE}_G()$ .

$\text{APPROXLEVERAGE}_G(u, v, \epsilon)$  that returns the approximate leverage score of an edge  $u, v$  in  $G$  to an error of  $\epsilon$ .

1 Initialize  $H$  as the empty graph,  $H \leftarrow \emptyset$ ;

2 **while**  $H$  has fewer than  $s$  edges **do**

3    $e, \mathbf{p}_e \leftarrow \text{SAMPLEEDGE}_G()$ .

4   Let  $\mathbf{p}'_e \leftarrow \frac{2}{n-1} \text{APPROXLEVERAGE}_G(u, v, 0.1)$

5   Reject  $e$  with probability  $1 - \mathbf{p}'_e / (4\rho \cdot \mathbf{p}_e)$ .

6   Let  $\mathbf{p}''_e \leftarrow \frac{1}{n-1} \text{APPROXLEVERAGE}_G(u, v, \epsilon)$

7   Reject  $e$  with probability  $1 - \mathbf{p}''_e / \mathbf{p}'_e$ .

8   Add  $e$  to  $H$  with new weight

$$\frac{\mathbf{w}_e}{\mathbf{p}''_e s} \exp\left(\frac{n^2}{2(n-1)s}\right).$$

9 Output  $H$

---

We first show that this routine will in fact sample edges according to  $\epsilon$ -approximate leverage scores, as was assumed in  $\text{IDEALSPARSIFY}$

**Lemma 4.7.** *The edges are being sampled with probability proportional to  $\tilde{\tau}^{(G, \epsilon)}$ , the leverage score estimates given by  $\text{APPROXLEVERAGE}_G(\cdot, \epsilon)$ .*

Note that this algorithm does not, at any time, have access to the full distribution  $\tilde{\tau}^{(G, \epsilon)}$ .

*Proof.* Our proof will assume the known guarantees of rejection sampling, which is to say that the following are true:

1. Given distributions  $\mathbf{p}$  and  $\mathbf{p}'$ , sampling an edge  $e$  from  $\mathbf{p}$  and accepting with probability  $\mathbf{p}'_e/(4\rho \cdot \mathbf{p}_e)$  is equivalent to drawing an edge from  $\mathbf{p}'$  as long as  $\mathbf{p}'_e/(4\rho \cdot \mathbf{p}_e) \in [0, 1]$  for all  $e$ .
2. Given distributions  $\mathbf{p}'$  and  $\mathbf{p}''$ , sampling an edge  $e$  from  $\mathbf{p}'$  and accepting with probability  $\mathbf{p}''_e/\mathbf{p}'_e$  is equivalent to drawing an edge from  $\mathbf{p}''$  as long as  $\mathbf{p}''_e/\mathbf{p}'_e \in [0, 1]$  for all  $e$ .

As a result, we only need to check that  $\mathbf{p}'_e/(4\rho \mathbf{p}_e)$  and  $\mathbf{p}''_e/\mathbf{p}'_e$  are at most 1.

The guarantees of `SAMPLEEDGEG()` gives

$$\frac{\bar{\tau}_e}{n-1} \leq \rho \mathbf{p}_e.$$

As  $\mathbf{p}'_e$  was generated with error 1.1, we have

$$\mathbf{p}'_e \leq \frac{2.2\bar{\tau}_e}{(n-1)} \leq 2.2\rho \mathbf{p}_e,$$

so  $\mathbf{p}'_e/(4\rho \mathbf{p}_e) \leq 1$ . To show  $\mathbf{p}''_e/\mathbf{p}'_e \leq 1$ , once again the guarantees of `SAMPLEEDGEG()` gives:

$$\mathbf{p}''_e \leq (1+\epsilon) \frac{\bar{\tau}_e}{n-1} \leq 2 \cdot 0.9 \frac{\bar{\tau}_e}{n-1} \leq \mathbf{p}'_e.$$

□

It remains to show that this rejection sampling process still makes sufficiently progress, yet also does not call `APPROXLEVERAGEG(e,  $\epsilon$ )` (the more accurate leverage score estimator) too many times.

**Lemma 4.8.** *At each step, the probability of `DETPARSIFY` calling `APPROXLEVERAGEG(e,  $\epsilon$ )` is at most  $\frac{1}{\rho}$ , while the probability of it adding an edge to  $H$  is at least  $\frac{1}{8\rho}$ .*

*Proof.* The proof utilizes the fact  $\sum_e \bar{\tau}_e = n-1$  (Fact 2.3) extensively.

If the edge  $e$  is picked, `APPROXLEVERAGEG(e,  $\epsilon$ )` is called with probability

$$\frac{\mathbf{p}'_e}{4\rho \cdot \mathbf{p}_e} \leq \frac{2.2\bar{\tau}_e}{4\rho \cdot \mathbf{p}_e \cdot (n-1)}$$

Summing over this over all edge  $e$  by the probability of picking them gives:

$$\sum_e \mathbf{p}_e \frac{2.2\bar{\tau}_e}{4\rho \cdot \mathbf{p}_e \cdot (n-1)} = \frac{2.2 \sum_e \bar{\tau}_e}{4\rho \cdot (n-1)} \leq \frac{1}{\rho}.$$

On the other hand, the probability of picking edge  $e$ , and not rejecting it is:

$$\mathbf{p}_e \cdot \frac{\mathbf{p}'_e}{4\rho \cdot \mathbf{p}_e} \cdot \frac{\mathbf{p}''_e}{\mathbf{p}'_e} = \frac{\tilde{\tau}^{(G, \epsilon)}}{4\rho(n-1)}.$$

where this follows by cancellation and how we set  $\mathbf{p}_e''$  in our algorithm. Summing over all edges then gives the probability of not rejecting an edge to be

$$\sum_e \frac{\tilde{\tau}^{(G, \epsilon)}}{4\rho(n-1)} \geq \sum_e \frac{(1-\epsilon)\bar{\tau}_e}{4\rho(n-1)} = \frac{(1-\epsilon)\sum_e \bar{\tau}_e}{4\rho(n-1)} \geq \frac{1}{8\rho}$$

□

*Proof.* (of Theorem 4.1) Lemma 4.7 implies that edges are sampled in DETSPARSIFY with probability proportional to  $\epsilon$ -approximate leverage scores guaranteed by APPROXLEVERAGE $_G(\cdot, \epsilon)$ . Therefore, we can apply Lemma 4.5 to achieve the desired expectation and concentration bounds. Finally, Lemma 4.8 implies that we expect to sample at most  $O(s \cdot \rho)$  edges, each of which require a call to SAMPLEEDGE $_G(e)$  and APPROXLEVERAGE $_G$  with constant error. It additionally implies that we expect to make  $O(s)$  calls to APPROXLEVERAGE $_G$  with  $\epsilon$  error.

□

Directly invoking this theorem leads to the sparsification algorithm.

*Proof.* (of Theorem 1.1) Consider invoking Theorem 4.1 with parameters

$$\begin{aligned} s &\leftarrow O(n^{1.5}\delta^{-2}), \\ \epsilon &\leftarrow n^{-1/4}. \end{aligned}$$

This gives:

$$\frac{\epsilon^2 n^2}{s}, \frac{n^3}{s^2} \leq \delta,$$

which then implies

$$(1 - O(\delta^2)) \mathcal{T}_G \leq \mathbb{E}_H[\mathcal{T}_H] \leq (1 + O(\delta^2)) \mathcal{T}_G,$$

and

$$\mathbb{E}_H[\mathcal{T}_H^2] \leq (1 + O(\delta^2)) \mathbb{E}_H[\mathcal{T}_H]^2.$$

The second condition is equivalent to  $\mathbf{Var}_H[\mathcal{T}_H] \leq \delta^2 \mathbb{E}_H[\mathcal{T}_H]$ , which by Chebyshev inequality gives that with constant probability we have

$$(1 - O(\delta)) \mathcal{T}_G \leq \mathcal{T}_H \leq (1 + O(\delta)) \mathcal{T}_G.$$

Combining this with the bounds on  $\mathbb{E}_H[\mathcal{T}_H]$ , and adjusting constants gives the overall bound.

Constructing the probability distribution  $\mathbf{p}$  for sampling edges only requires computing constant approximate leverage scores for all edges, and then sampling proportionally for each edge, giving a constant value for  $\rho$ . By Lemma 2.4, this requires  $\tilde{O}(m)$  time. The running time then is dominated by the  $O(s)$  calls made to the effective resistance oracle with error  $\epsilon = n^{-1/4}$ . Invoking Lemma 2.4 gives that this cost is bounded by

$$O(n\epsilon^{-4} + s\epsilon^{-2}) = O(n^2\delta^{-2}).$$

Furthermore, because Lemma 2.4 holds w.h.p. we can absorb the probability of failure into our constant probability bound

□

Another immediate consequence of our sparsification routine in Theorem 4.1, along with bounds on total variation distance that we prove in Section 9, is that we can give a faster spanning tree sampling algorithm for dense graphs by plugging the sparsified graph into previous algorithms for generating random spanning trees.

*Proof.* (of Corollary 3.8) As in the proof of Theorem 1.1, we invoke Theorem 4.1 with parameters

$$s \leftarrow O(n^{1.5}\delta^{-2}), \\ \epsilon \leftarrow n^{-1/4}.$$

giving

$$\frac{\mathbb{E}_H [\mathcal{T}_H^2]}{\mathbb{E}_H [\mathcal{T}_H]^2} \leq 1 + \delta^2.$$

Applying Lemma 3.7, which is proven in Section 9.1, we then have that drawing a tree from  $H$  according to the  $\mathbf{w}$ -uniform distribution gives a total variation distance of  $\delta$  from drawing a tree according to the  $\mathbf{w}$ -uniform distribution of  $G$ . The running time of drawing  $H$  is dominated by the  $O(s)$  calls made to the effective resistance oracle with error  $\epsilon = n^{-1/4}$ . Invoking Lemma 2.4 gives that this cost is bounded by

$$O(n\epsilon^{-4} + s\epsilon^{-2}) = O(n^2\delta^{-2}).$$

Furthermore, because Lemma 2.4 holds w.h.p. we can absorb the probability of failure into our total variation distance bound (where we implicitly assume that  $\delta$  is at most polynomially small).

We then use the  $\tilde{O}(m^{1/3}n^{5/3})$  time algorithm in [DKP<sup>+</sup>16] with  $m = O(n^{1.5}\delta^{-2})$  to draw a tree from  $H$ . This then achieves our desired running time and total variation distance bound.  $\square$

## 5 Implicit Sparsification of the Schur Complement

Note that the determinant sparsification routine in Theorem 4.1 only requires an oracle that samples edges by an approximate distribution to resistance, as well as access to approximate leverage scores on the graph. This suggests that a variety of naturally dense objects, such as random walk matrices [CCL<sup>+</sup>15, JKPS17] and Schur complements [KLP<sup>+</sup>16, DKP<sup>+</sup>16] can also be sparsified in ways that preserve the determinant (of the minor with one vertex removed) or the spanning tree distributions. The latter objects, Schur complements, have already been shown to lead to speedups in random spanning tree generation algorithms recently [DKP<sup>+</sup>16].

Furthermore the fact that Schur complements preserve effective resistances exactly (2.6) means that we can directly invoke the effective resistances data structure as constructed in Lemma 2.4 to produce effective resistance estimates on any of its Schur complements. As a result, the main focus of this section is an efficient way of producing samples from a distribution that approximates drawing a multi-edge from the Schur complement with probabilities proportional to its leverage score. Here we follow the template introduced in [KLP<sup>+</sup>16] of only eliminating  $(1 + \alpha)$ -diagonally-dominant subsets of vertices, as it in turn allows the use of walk sampling based implicit sparsification similar to those in [CCL<sup>+</sup>15, JKPS17].

$(1 + \alpha)$ -diagonally-dominant subsets have been used in Schur complement based linear system solvers to facilitate the convergence of iterative methods in the  $\mathbf{L}_{[V_2, V_2]}$  block [KLP<sup>+</sup>16]. Formally, the condition that we require is:

**Definition 5.1.** In a weighted graph  $G = (V, E, \mathbf{w})$ , a subset of vertices  $V_2 \subseteq V$  is  $(1 + \alpha)$ -diagonally-dominant, or  $(1 + \alpha)$ -DD if for every  $u \in V_2$  with weighted degree  $\mathbf{d}_u$  we have:

$$\sum_{v \sim u, v \notin V_2} \mathbf{w}_{uv} \geq \frac{1}{1 + \alpha} \mathbf{d}_u = \frac{1}{1 + \alpha} \sum_{v \sim u} \mathbf{w}_{uv}.$$

It was shown in [KLP<sup>+</sup>16] that large sets of such vertices can be found by trimming a uniformly random sample.

**Lemma 5.2.** (Lemma 3.5. of [KLP<sup>+</sup>16] instantiated on graphs)

There is a routine `ALMOSTINDEPENDENT`( $G, \alpha$ ) that for a graph  $G$  with  $n$  vertices, and a parameter  $\alpha \geq 0$ , returns in  $O(m)$  expected time a subset  $V_2$  with  $|V_2| \geq n/(8(1 + \alpha))$  such that  $\mathbf{L}_{G, [V_2, V_2]}$  is  $(1 + \alpha)$ -DD.

Given such a subset  $V_2$ , we then proceed to sample edges in  $\text{SC}(G, V_1)$  via the following simple random walk sampling algorithm:

1. Pick a random edge in  $G$ .
2. Extend both of its endpoints in random walks until they first reach somewhere in  $V_1$ .

Incorporating this scheme into the determinant preserving sparsification schemes then leads these guarantees:

**Theorem 5.3.** Conditioned on Lemma 2.4 holding, there is a procedure `SCHURSPARSE` that takes a graph  $G$ , and an 1.1-DD subset of vertices  $V_2$ , returns a graph  $H^{V_1}$  in  $\tilde{O}(n^2 \delta^{-1})$  expected time such that the distribution over  $H^{V_1}$  satisfies:

$$\mathcal{T}_{\text{SC}(G, V_1)} \exp(-\delta) \leq \mathbb{E}_{H^{V_1}} [\mathcal{T}_{H^{V_1}}] \leq \mathcal{T}_{\text{SC}(G, V_1)} \exp(\delta),$$

and

$$\frac{\mathbb{E}_{H^{V_1}} [\mathcal{T}_{H^{V_1}}^2]}{\mathbb{E}_{H^{V_1}} [\mathcal{T}_{H^{V_1}}]^2} \leq \exp(\delta).$$

Furthermore, the number of edges of  $H^{V_1}$  can be set to anywhere between  $O(n^{1.5} \delta^{-1})$  and  $O(n^2 \delta^{-1})$  without affecting the final bound.

We let this subset of vertices produced to be  $V_2$ , and let its complement be  $V_1$ . Our key idea is to view  $\text{SC}(G, V_1)$  as a multi-graph where each multi-edge corresponds to a walk in  $G$  that starts and ends in  $V_1$ , but has all intermediate vertices in  $V_2$ . Specifically a length  $k$  walk

$$u_0, u_1, \dots, u_k,$$

with  $u_0, u_k \in V_1$  and  $u_i \in V_2$  for all  $0 < i < k$ , corresponds to a multi-edge between  $u_0$  and  $u_k$  in  $\text{SC}(G, V_1)$  with weight given by

$$\mathbf{w}_{u_0, u_1, \dots, u_k}^{\text{SC}(G, V_1)} \stackrel{\text{def}}{=} \frac{\prod_{0 \leq i < k} \mathbf{w}_{u_i u_{i+1}}^G}{\prod_{0 < i < k} \mathbf{d}_{u_i}^G}. \quad (2)$$

We check formally that this multi-graph defined on  $V_1$  is exactly the same as  $\text{SC}(G, V_1)$  via the Taylor expansion of  $\mathbf{L}_{[V_2, V_2]}^{-1}$  based Jacobi iteration.

**Lemma 5.4.** *Given a graph  $G$  and a partition of its vertices into  $V_1$  and  $V_2$ , the graph  $G^{V_1}$  formed by all the multi-edges corresponding to walks starting and ending at  $V_1$ , but stays entirely within  $V_2$  with weights given by Equation 2 is exactly  $\text{SC}(G, V_1)$ .*

*Proof.* Consider the Schur complement:

$$\text{SC}(G, V_1) = \mathbf{L}_{[V_1, V_1]} - \mathbf{L}_{[V_2, V_1]} \mathbf{L}_{[V_2, V_2]}^\dagger \mathbf{L}_{[V_1, V_2]}.$$

If there are no edges leaving  $V_2$ , then the result holds trivially. Otherwise,  $\mathbf{L}_{[V_2, V_2]}$  is a strictly diagonally dominant matrix, and is therefore full rank. We can write it as

$$\mathbf{L}_{[V_2, V_2]} = \mathbf{D} - \mathbf{A}$$

where  $\mathbf{D}$  is the diagonal of  $\mathbf{L}_{[V_2, V_2]}$  and  $\mathbf{A}$  is the negation of the off-diagonal entries, and then expand  $\mathbf{L}_{[V_2, V_2]}^{-1}$  via the Jacobi series:

$$\begin{aligned} \mathbf{L}_{[V_2, V_2]}^{-1} &= (\mathbf{D} - \mathbf{A})^{-1} = \mathbf{D}^{-1/2} \left( \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2} \right)^{-1} \mathbf{D}^{-1/2} \\ &= \mathbf{D}^{-1/2} \left[ \sum_{k=0}^{\infty} \left( \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2} \right)^k \right] \mathbf{D}^{-1/2} = \sum_{k=0}^{\infty} (\mathbf{D}^{-1} \mathbf{A})^k \mathbf{D}^{-1}. \quad (3) \end{aligned}$$

Note that this series converges because the strict diagonal dominance of  $\mathbf{L}_{[V_2, V_2]}$  implies  $(\mathbf{A} \mathbf{D}^{-1})^k$  tends to zero as  $k \rightarrow \infty$ . Substituting this in place of  $\mathbf{L}_{[V_2, V_2]}^{-1}$  gives:

$$\text{SC}(G, V_1) = \mathbf{L}_{[V_1, V_1]} - \sum_{k=0}^{\infty} \mathbf{L}_{[V_1, V_2]} (\mathbf{D}^{-1} \mathbf{A})^k \mathbf{D}^{-1} \mathbf{L}_{[V_2, V_1]}.$$

As all the off-diagonal entries in  $\mathbf{L}$  are non-positive, we can replace  $\mathbf{L}_{[V_1, V_2]}$  with  $-\mathbf{L}_{[V_1, V_2]}$  to make all the terms in the trailing summation positive. As these are the only ways to form new off-diagonal entries, the identity based on matrix multiplication of

$$\left[ (-\mathbf{L}_{[V_1, V_2]}) (\mathbf{D}^{-1} \mathbf{A})^k \mathbf{D}^{-1} (-\mathbf{L}_{[V_2, V_1]}) \right]_{u_0, u_k} = \sum_{u_1 \dots u_{k-1}} \frac{\prod_{0 \leq i < k} \mathbf{w}_{u_i u_{i+1}}^G}{\prod_{0 < i < k} \mathbf{d}_{u_i}^G}$$

gives the required identity.  $\square$

This characterization of  $\text{SC}(G, V_1)$ , coupled with the  $(1 + \alpha)$ -diagonal-dominance of  $V_2$ , allows us to sample the multi-edges in  $\text{SC}(G, V_1)$  in the same way as the (short) random walk sparsification algorithms from [CCL<sup>+</sup>15, JKPS17].

**Lemma 5.5.** *Given any graph  $G = (V, E, \mathbf{w})$ , an  $(1 + \alpha)$ -DD subset  $V_2$ , and access to 2-approximations of statistical leverage scores on  $G$ ,  $\tilde{\tau}^G$ , `SAMPLEEDGESCHUR` returns edges in  $G$  according to the distribution  $\mathbf{p}_e$  in  $O(\alpha)$  expected time per sample. Furthermore, the distribution that it samples edges in  $\text{SC}(G, V_1)$  from,  $\mathbf{p}$ , satisfies*

$$O(1) \cdot \mathbf{p}_{u_0, \dots, u_k} \geq \frac{\bar{\tau}_{u_0, \dots, u_k}^{\text{SC}(G, V_1)}}{n-1}.$$

for every edge in  $\text{SC}(G, V_1)$  corresponding to the walk  $u_0, \dots, u_k$ .

---

**Algorithm 3:** SAMPLEEDGESCHUR( $G = (V, E, \mathbf{w})$ ,  $V_1$ ): samples an edge from  $\text{SC}(G, V_1)$ 


---

**Input:** Graph  $G$ , vertices  $V_1$  to complement onto, and (implicit) access to a 2-approximation of the leverage scores of  $G$ ,  $\tilde{\tau}^G$ .

**Output:** A multi-edge  $e$  in  $\text{SC}(G, V_1)$  corresponding to a walk  $u_0, u_1, \dots, u_k$ , and the probability of it being picked in this distribution  $\mathbf{p}_{u_0, u_1, \dots, u_k}$

- 1 Sample an edge  $e$  from  $G$  randomly with probability drawn from  $\tilde{\tau}_e^G$ ;
- 2 Perform two independent random walks from the endpoints of  $e$  until they both reach some vertex in  $V_1$ , let the walk be  $u_0 \dots u_k$ ;
- 3 Output edge  $u_0 u_k$  (corresponding to the path  $u_0, u_1, \dots, u_k$ ) with

$$\mathbf{w}_{u_0 \dots u_k} \leftarrow \frac{\prod_{0 \leq i < k} \mathbf{w}_{u_i u_{i+1}}^G}{\prod_{0 < i < k} d_{u_i}^G}, \quad (\text{same as Equation 2})$$

$$\mathbf{p}_{u_0 \dots u_k} \leftarrow \frac{1}{\sum_{e'} \tilde{\tau}_{e'}^G} \sum_{0 \leq i < k} \tilde{\tau}_{u_i u_{i+1}}^G \cdot \left( \prod_{0 \leq j < i} \frac{\mathbf{w}_{u_j u_{j+1}}^G}{d_{u_{j+1}}} \cdot \prod_{i+1 \leq j < k} \frac{\mathbf{w}_{u_j u_{j+1}}^G}{d_{u_j}} \right).$$


---

The guarantees of this procedure are analogous to the random walk sampling sparsification scheme from [CCL<sup>+</sup>15, JKPS17], with the main difference being the terminating condition for the walks leads to the removal of an overhead related to the number of steps in the walk. The modification of the initial step to picking the initial edge from  $G$  by resistance is necessary to get  $\rho$  to a constant, as the about  $n^{1.5}$  samples limits the amount of overhead that we can have per sample.

*Proof.* We first verify that  $\mathbf{p}$  is indeed a probability on the multi-edges of  $\text{SC}(G, V_1)$ , partitioned by the walks that they correspond to in  $G$ , or formally

$$\sum_{\substack{u_0, u_1, \dots, u_k: \\ u_0, u_k \in V_1, \\ u_i \in V_2 \ \forall 1 \leq i < k}} \mathbf{p}_{u_0, u_1, \dots, u_k} = 1.$$

To obtain this equality, note that for any random walk starting at vertex  $i$ , the total probabilities of walks starting at  $i$  and ending in  $V_1$  is upper bounded by 1. Algebraically this becomes:

$$\sum_{u_1, u_2, \dots, u_k} \prod_{0 \leq i < k} \frac{\mathbf{w}_{u_i u_{i+1}}}{d_{u_i}} = 1,$$

so applying this to both terms of each edge  $e$  gives that the total probability mass over any starting edge is  $\frac{\tilde{\tau}_e^G}{\sum_{e'} \tilde{\tau}_{e'}^G}$ , and in turn the total.

For the running time, since  $V_2$  is  $(1 + \alpha)$ -almost independent, each step of the walk takes expected time  $O(\alpha)$ . Also, the value of  $\mathbf{p}_{u_0, u_1, \dots, u_k}$  can be computed in  $O(k)$  time by computing prefix/suffix products of the transition probabilities along the path (instead of evaluating each summand in  $O(k)$  time for a total of  $O(k^2)$ ).

Finally, we need to bound the approximation of  $\mathbf{p}$  compared to the true leverage scores  $\bar{\tau}$ . As  $\tilde{\tau}_e^G$  is a 2-approximation of the true leverage scores,  $\sum_e \tilde{\tau}_e^G$  is within a constant factor of  $n$ .

So it suffices to show

$$O(1) \cdot \sum_{0 \leq i < k} \tilde{\tau}_{u_i u_{i+1}}^G \left( \prod_{0 \leq j < i} \frac{\mathbf{w}_{u_j u_{j+1}}^G}{\mathbf{d}_{u_{j+1}}} \cdot \prod_{i+1 \leq j < k} \frac{\mathbf{w}_{u_j u_{j+1}}^G}{\mathbf{d}_{u_j}} \right) \geq \mathcal{R}_{eff}^{Sc(G, V_1)}(u_0, u_k) \cdot \mathbf{w}_{u_0, u_1, \dots, u_k}.$$

Here we invoke the equivalence of effective resistances in  $G$  and  $Sc(G, V_1)$  given by Fact 2.6 in the reverse direction. Then by Rayleigh's monotonicity principle, we have

$$\mathcal{R}_{eff}^{Sc(G, V_1)}(u_0, u_k) = \mathcal{R}_{eff}^G(u_0, u_k) \leq \sum_{0 \leq i < k} \frac{2\tilde{\tau}_{u_i u_{i+1}}^G}{\mathbf{w}_{u_i u_{i+1}}},$$

which when substituted into the expression for  $\mathbf{w}_{u_0, u_1, \dots, u_k}$  from Equation 2 gives

$$\left( \sum_{0 \leq i < k} \frac{2\tilde{\tau}_{u_i u_{i+1}}^G}{\mathbf{w}_{u_i u_{i+1}}} \right) \mathbf{w}_{u_0, u_1, \dots, u_k} = \sum_{0 \leq i < k} 2\tilde{\tau}_{u_i u_{i+1}}^G \left( \prod_{0 \leq j < i} \frac{\mathbf{w}_{u_j u_{j+1}}^G}{\mathbf{d}_{u_{j+1}}} \cdot \prod_{i+1 \leq j < k} \frac{\mathbf{w}_{u_j u_{j+1}}^G}{\mathbf{d}_{u_j}} \right).$$

□

This sampling procedure can be immediately combined with Theorem 4.1 to give algorithms for generating approximate Schur complements. Pseudocode of this routine is in Algorithm 4.

---

**Algorithm 4:** SCHURSPARSE( $G, V_1, \delta$ )

---

**Input:** Graph  $G$ , 1.1-DD subset of vertices  $V_2$  and error parameter  $\delta$

**Output:** Sparse Schur complement of  $Sc(G, V_1)$

- 1 Set  $\epsilon \leftarrow 0.1$ ;
  - 2 Set  $s \leftarrow n^2\delta^{-1}$ ;
  - 3 Build leverage score data structure on  $G$  with errors 0.1 (via Lemma 2.4);
  - 4 Let  $H^{V_1} \leftarrow$   
DETSPARSIFY( $Sc(G, V_1), s, \text{SAMPLEEDGECHUR}(G, V_1), \text{LEVERAGEAPPROX}_G, \epsilon$ );
  - 5 Output  $H^{V_1}$ ;
- 

*Proof.* (Of Theorem 5.3) Note that the choices of  $\epsilon$  and  $s$  must ensure that

$$\begin{aligned} \frac{n^2\epsilon^2}{s} &= \delta \\ \frac{n^3}{s^2} &\leq \delta \end{aligned}$$

This is then equivalent to  $s \geq n^{1.5}\delta^{-1}$  and  $\frac{s}{\epsilon^2} = n^2\delta^{-1}$ . This further implies that  $\epsilon \geq n^{1/4}$ . Our  $\epsilon$  and  $s$  in SCHURSPARSE meet these conditions (and the ones specifically chosen in the algorithm will also be necessary for one of our applications). The guarantees then follow from putting the quality of the sampler from Lemma 5.5 into the requirements of the determinant preserving sampling procedure from Theorem 4.1. Additionally, Lemma 5.5 only requires access to 2-approximate leverage scores, which can be computed by Lemma 2.4 in  $\tilde{O}(m)$  time. Furthermore, Lemma 5.5 gives that our  $\rho$  value is constant, and our assumption in Theorem 5.3 that we

are given an 1.1-DD subset  $V_2$  implies that our expected  $O(s \cdot \rho)$  calls to SAMPLEEDGESCHUR will require  $O(1)$  time. The only other overheads are the computation and invocations of the various copies of approximate resistance data structures. Since  $m \leq n^2$  and  $\epsilon \geq n^{1/4}$ , Lemma 2.4 gives that this cost is bounded by  $\tilde{O}(m + n^2 + \frac{s}{\epsilon^2}) = \tilde{O}(n^2\delta^{-1})$ .  $\square$

## 6 Approximate Determinant of SDDM Matrices

In this section, we provide an algorithm for computing an approximate determinant of SDDM matrices, which are minors of graph Laplacians formed by removing one row/column.

Theorem 1.1 allows us to sparsify a dense graph while still approximately preserving the determinant of the graph minor. If there were some existing algorithm for computing the determinant that had good dependence on sparsity, we could achieve an improved runtime for determinant computation by simply invoking such an algorithm on a minor of the sparsified graph.<sup>6</sup> Unfortunately, current determinant computation algorithms (that achieve high-accuracy) are only dependent on  $n$ , so simply reducing the edge count does not directly improve the runtime for determinant computation. Instead the algorithm we give will utilize Fact 2.5

$$\det_+(\mathbf{L}) = \det(\mathbf{L}_{[V_2, V_2]}) \cdot \det_+(\text{SC}(\mathbf{L}, V_1)).$$

(where we recall that  $\det_+$  is the determinant of the matrix minor) to recursively split the matrix. Specifically, we partition the vertex set based upon the routine ALMOSTINDEPENDENT from Lemma 5.2, then compute Schur complements according to SCHURSPARSE in Theorem 5.3. Our algorithm will take as input a Laplacian matrix. However, this recursion naturally produces two matrices, the second of which is a Laplacian and the first of which is a submatrix of a Laplacian. Therefore, we need to convert  $\mathbf{L}_{[V_2, V_2]}$  into a Laplacian. We do this by adding one vertex with appropriate edge weights such that each row and column sums to 0. Pseudocode of this routine is in Algorithm 5, and we call it with the parameters  $\mathbf{L}^{V_2} \leftarrow \text{ADDROWCOLUMN}(\mathbf{L}_{[V_2, V_2]})$ .

---

**Algorithm 5:** ADDROWCOLUMN( $\mathbf{M}$ ) : complete  $\mathbf{M}$  into a graph Laplacian by adding one more row/column

---

**Input:** SDDM Matrix  $\mathbf{M}$

**Output:** Laplacian matrix  $\mathbf{L}$  with one extra row / column than  $\mathbf{M}$

- 1 Let  $n$  be the dimension of  $\mathbf{M}$ ;
  - 2 **for**  $i = 1$  to  $n$  **do**
  - 3     Sum non-zero entries of row  $i$ , call  $\mathbf{s}_i$ ;
  - 4     Set  $\mathbf{L}(n+1, i), \mathbf{L}(i, n+1) \leftarrow -\mathbf{s}_i$ ;
  - 5 Let  $\mathbf{L}(n+1, n+1) \leftarrow \sum_{i=1}^n \mathbf{s}_i$ ;
  - 6 Output  $\mathbf{L}$ ;
- 

The procedure ADDROWCOLUMN outputs a Laplacian  $\mathbf{L}^{V_2}$  such that  $\mathbf{L}_{[V_2, V_2]}$  can be obtained if one removes this added row/column. This immediately gives  $\det_+(\mathbf{L}^{V_2}) = \det(\mathbf{L}_{[V_2, V_2]})$  by definition, and we can now give our determinant computation algorithm of the minor of a graph Laplacian.

---

<sup>6</sup>To get with high probability one could use standard boosting tricks involving taking the median of several estimates of the determinant obtained in this fashion.

---

**Algorithm 6:** DETAPPROX( $\mathbf{L}, \delta, \bar{n}$ ) : Compute  $\det_+(\mathbf{L})$  with error parameter  $\delta$ 


---

**Input:** Laplacian matrix  $\mathbf{L}$ , top level error threshold  $\delta$ , and top level graph size  $\bar{n}$

**Output:** Approximate  $\det_+(\mathbf{L})$

```

1 if this is the top-level invocation of this function in the recursion tree then
2    $\delta' \leftarrow \Theta(\delta^2 / \log^3 n)$ 
3 else
4    $\delta' \leftarrow \delta$ 
5 if  $\mathbf{L}$  is  $2 \times 2$  then
6   return the weight on the (unique) edge in the graph
7  $V_2 \leftarrow \text{ALMOSTINDEPENDENT}(\mathbf{L}, \frac{1}{10})$                                      {Via Lemma 5.2}
8  $V_1 \leftarrow V \setminus V_2$  ;
9  $\mathbf{L}^{V_1} \leftarrow \text{SCHURSPARSE}(\mathbf{L}, V_1, \delta');$                                  $\{|V_1| / \bar{n}$  is the value of  $\beta$  in Lemma 6.1.}
10  $\mathbf{L}^{V_2} \leftarrow \text{ADDROWCOLUMN}(\mathbf{L}_{[V_2, V_2]})$ ;
11 Output DETAPPROX( $\mathbf{L}^{V_1}, \delta' |V_1| / \bar{n}, \bar{n}$ )  $\cdot$  DETAPPROX( $\mathbf{L}^{V_2}, \delta' |V_2| / \bar{n}, \bar{n}$ );

```

---

Our analysis of this recursive routine consists of bounding the distortions incurred at each level of the recursion tree. This in turn uses the fact that the number of vertices across all calls within a level and the total “amount” of  $\delta$  across all calls within a level both remain unchanged from one level to the next. This can be summarized by the following Lemma which bounds the error accumulated within one level of recursion in our algorithm.

**Lemma 6.1.** *Suppose we are given some small  $\delta \geq 0$  and non-negative  $\beta_1, \dots, \beta_k$  such that  $\sum_{i=1}^k \beta_i = O(1)$ , along with Laplacian matrices  $\mathbf{L}(1), \dots, \mathbf{L}(k)$  and each having a corresponding vertex partition  $V_1(i), V_2(i)$ , where*

$$\mathbf{L}(i) = \begin{bmatrix} \mathbf{L}(i)_{[V_1(i), V_1(i)]} & \mathbf{L}(i)_{[V_1(i), V_2(i)]} \\ \mathbf{L}(i)_{[V_2(i), V_1(i)]} & \mathbf{L}(i)_{[V_2(i), V_2(i)]} \end{bmatrix}.$$

Let  $\mathbf{L}^{V_1(i)}$  denote the result of running SCHURSPARSE to remove the  $V_2(i)$  block in each of these matrices:<sup>7</sup>

$$\mathbf{L}^{V_1(i)} \stackrel{\text{def}}{=} \text{SCHURSPARSE}(\mathbf{L}(i), V_1(i), \beta_i \delta).$$

Then conditioning upon a with high probability event<sup>8</sup> in each of these calls to SCHURSPARSE, for any  $p$  we have with probability at least  $1 - p$ :

$$\prod_{i=1}^k \det_+(\mathbf{L}(i)) = \left(1 \pm O\left(\sqrt{\delta/p}\right)\right) \prod_{i=1}^k \det(\mathbf{L}_{[V_2(i), V_2(i)]}(i)) \cdot \det_+(\mathbf{L}^{V_1(i)}).$$

<sup>7</sup>This Lemma only applies when the matrices are fixed with respect to the randomness used in the invocations of SCHURSPARSE mentioned in the Lemma. In other words, it only applies when the result of running SCHURSPARSE on each of these  $\mathbf{L}(i)$  matrices is independent of the result of running it on the other matrices. This is why the Lemma only immediately bounds error within a level of the recursion—where this independence holds—rather than for the entire algorithm.

<sup>8</sup>namely, the event that all the leverage score estimation calls to Lemma 2.4 from SCHURSPARSE succeed

Here the  $\beta_i$  corresponds to the  $|V_1|/n$  and  $|V_2|/n$  values that  $\delta$  is multiplied against in each call parameter to SCHURSPARSE. An example of the main steps in this determinant approximation algorithm, as well as the graphs corresponding to applying Lemma 6.1 to one of the layers is in Figure 1.

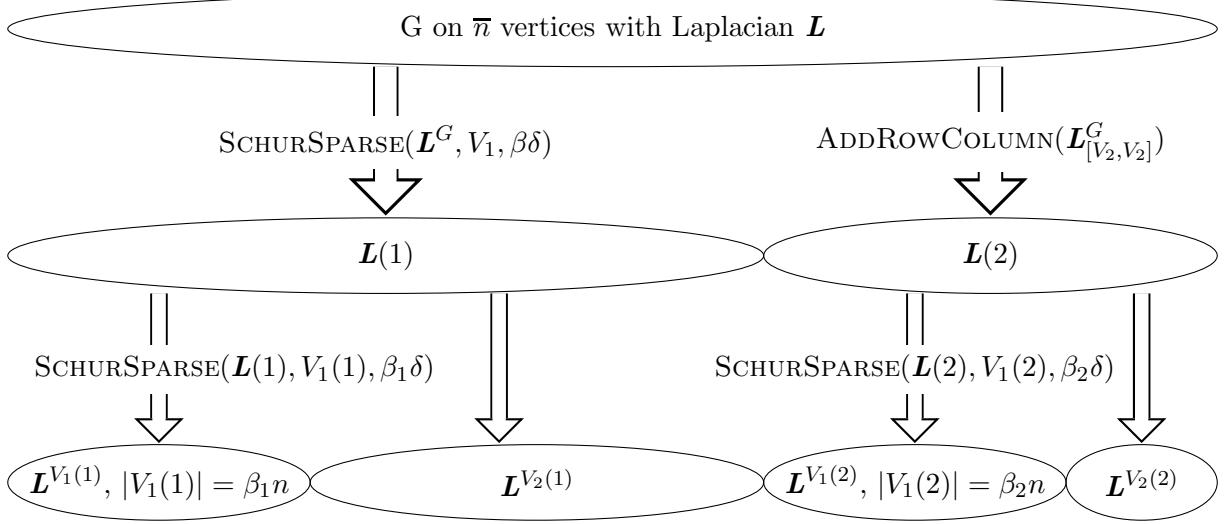


Figure 1: Two layers of the call Structure of the determinant approximation algorithm DETAPPROX (algorithm 6), with the transition from the first to the second layer labeled as in Lemma 6.1.

Applying Lemma 6.1 to all the layers of the recursion tree gives the overall guarantees.

*Proof of Theorem 1.2.*

**Running Time:** Let the number of vertices and edges in the current graph corresponding to  $L$  be  $n$  and  $m$  respectively. Calling ALMOSTINDEPENDENT takes expected time  $O(m)$  and guarantees

$$\frac{n}{16} \leq |V_2| \leq \frac{n}{8},$$

which means the total recursion terminates in  $O(\log n)$  steps.

For the running time, note that as there are at most  $O(n)$  recursive calls, the total number of vertices per level of the recursion is  $O(n)$ . The running time on each level are also dominated by the calls to SCHURSPARSE, which comes out to

$$\tilde{O}\left(|V_1(i)|^2 \frac{n}{\delta' |V_1(i)|}\right) = \tilde{O}\left(|V_1(i)| n \delta^{-2}\right),$$

and once again sums to  $\tilde{O}(n^2 \delta^{-2})$ . We note that this running time can also be obtained from more standard analyses of recursive algorithms, specifically applying guess-and-check to a running time recurrence of the form of:

$$T(n, \delta) = T(\theta n, \theta \delta) + T((1 - \theta)n + 1, (1 - \theta)\delta) + \tilde{O}(n^2 \delta^{-1}).$$

**Correctness.** As shown in the running time analysis, our recursion tree has depth at most  $O(\log n)$ , and there are at most  $O(n)$  total vertices at any given level. We associate each level of the recursion in our algorithm with the list of matrices which are given as input to the calls making up that level of recursion. For any level in our recursion, consider the product of  $\det_+$  applied to each of these matrices. We refer to this quantity for level  $j$  as  $q_j$ . Notice that  $q_0$  is the determinant we wish to compute and  $q_{\#\text{levels}-1}$  is what our algorithm actually outputs. As such, it suffices to prove that for any  $j$ ,  $q_j = (1 \pm \frac{\delta}{\#\text{levels}})q_{j-1}$  with probability of failure at most  $\frac{1}{10 \cdot \#\text{levels}}$ . However, by the fact that we set  $\delta' = \Theta(\delta^2 / \log^3 n)$  in the top level of recursion with sufficiently small constants, this immediately follows from Lemma 6.1.

A minor technical issue is that Lemma 6.1 only gives guarantees conditioned on a WHP event. However, we only need to invoke this Lemma a logarithmic number of times, so we can absorb this polynomially small failure probability into the our total failure probability without issue.

Standard boosting techniques—such as running  $O(\log n)$  independent instances and taking the medians of the estimates—give our desired with high probability statement.  $\square$

It remains to bound the variances per level of the recursion.

*Proof.* (Of Lemma 6.1) As a result of Fact 2.5

$$\prod_{i=1}^k \det_+(\mathbf{L}(i)) = \prod_{i=1}^k \det\left(\mathbf{L}(i)_{[V_2(i), V_2(i)]}\right) \det_+(\text{SC}(\mathbf{L}(i), V_1(i))).$$

Consequently, it suffices to show that with probability at least  $1 - p$

$$\prod_{i=1}^k \det_+(\text{SC}(\mathbf{L}(i), V_1(i))) = \left(1 \pm O\left(\sqrt{\delta/p}\right)\right) \prod_{i=1}^k \det_+(\mathbf{L}^{V_1(i)}).$$

Recall that  $\mathbf{L}^{V_1(i)}$  denotes the random variable that is the approximate Schur complement generated through the call to  $\text{SCHURSPARSE}(\mathbf{L}(i), V_1(i), \beta_i \delta)$ .

Using the fact that our calls to  $\text{SCHURSPARSE}$  are independent along with the assumption of  $\sum_{i=1}^k \beta_i = O(1)$ , we can apply the guarantees of Theorem 5.3 to obtain

$$\begin{aligned} \mathbb{E}_{\mathbf{L}^{V_1(1)} \dots \mathbf{L}^{V_1(k)}} \left[ \prod_{i=1}^k \det_+(\mathbf{L}^{V_1(i)}) \right] &= \prod_{i=1}^k \mathbb{E}_{\mathbf{L}^{V_1(i)}} \left[ \det_+ \mathbf{L}^{V_1(i)} \right] \\ &= (1 \pm O(\delta)) \prod_{i=1}^k \det_+(\text{SC}(\mathbf{L}(i), V_1(i))), \end{aligned}$$

and

$$\begin{aligned} \frac{\mathbb{E}_{\mathbf{L}^{V_1(1)} \dots \mathbf{L}^{V_1(k)}} \left[ \prod_{i=1}^k \det_+ \left( \mathbf{L}^{V_1(i)} \right)^2 \right]}{\mathbb{E}_{\mathbf{L}^{V_1(1)} \dots \mathbf{L}^{V_1(k)}} \left[ \prod_{i=1}^k \det_+ \left( \mathbf{L}^{V_1(i)} \right)^2 \right]} &= \prod_{i=1}^k \frac{\mathbb{E}_{\mathbf{L}^{V_1(i)}} \left[ \det_+ \left( \mathbf{L}^{V_1(i)} \right)^2 \right]}{\mathbb{E}_{\mathbf{L}^{V_1(i)}} \left[ \det_+ \left( \mathbf{L}^{V_1(i)} \right) \right]^2} \\ &\leq \prod_{i=1}^k \exp(O(\beta_i \delta)) \leq \exp(O(\delta)). \end{aligned}$$

By assumption  $\delta$  is small, so we can approximate  $\exp(O(\delta))$  with  $1+O(\delta)$ , which with bound above gives

$$\mathbf{Var}_{\mathbf{L}^{V_1(1)} \dots \mathbf{L}^{V_1(k)}} \left[ \prod_{i=1}^k \det_+ \left( \mathbf{L}^{V_1(i)} \right) \right] \leq O(\delta) \mathbb{E}_{\mathbf{L}^{V_1(1)} \dots \mathbf{L}^{V_1(k)}} \left[ \prod_{i=1}^k \det_+ \left( \mathbf{L}^{V_1(i)} \right) \right]^2,$$

Then applying the approximation on  $\mathbb{E} \left[ \prod_{i=1}^k \det_+ (\text{SCHURSPARSE}(\mathbf{L}(i), V_1(i), \beta_i \delta)) \right]$  gives

$$\mathbf{Var}_{\mathbf{L}^{V_1(1)} \dots \mathbf{L}^{V_1(k)}} \left[ \prod_{i=1}^k \det_+ \left( \mathbf{L}^{V_1(i)} \right) \right] \leq O(\delta) \left( \prod_{i=1}^k \det_+ (\text{SC}(\mathbf{L}(i), V_1(i))) \right)^2.$$

At which point we can apply Chebyshev's inequality to obtain our desired result.  $\square$

## 7 Random Spanning Tree Sampling

In this section we will give an algorithm for generating a random spanning tree from a weighted graph, that uses `SCHURSPARSE` as a subroutine, and ultimately prove Theorem 1.3.

In order to do so, we will first give an  $O(n^\omega)$  time recursive algorithm using Schur complement that exactly generates a random tree from the  $\mathbf{w}$ -uniform distribution. The given algorithm is inspired by the one introduced in [CDN89], and its variants utilized in [CMN96, HX16, DKP<sup>+</sup>16]. However, we will (out of necessity for our further extensions) reduce the number of branches in the recursion to two, by giving an efficient algorithmic implementation of a bijective mapping between spanning trees in  $G$  and spanning trees in  $\text{SC}(G, V_2)$  when  $V_1$ , the set of vertices removed, is an independent set. We note that this also yields an alternative algorithm for generating random spanning trees from the  $\mathbf{w}$ -uniform distribution in  $O(n^\omega)$  time.

The runtime of this recursion will then be achieved similar to our determinant algorithm. We reduce  $\delta$  proportional to the decrease in the number of vertices for every successive recursive call in exactly the same was as the determinant approximation algorithm from Section 6. As has been previously stated and which is proven in Section 9.1, drawing a random spanning tree from a graph after running our sparsification routine which takes  $\tilde{O}(n^2\delta^{-1})$ , will have total variation distance  $\sqrt{\delta}$  from the  $\mathbf{w}$ -uniform distribution.

Similar to our analysis of the determinant algorithm, we cannot directly apply this bound to each tree because the lower levels of the recursion will contribute far too much error when  $\delta$  is not decreasing at a proportional rate to the total variation distance. Thus we will again need to give better bounds on the variance across each level, allowing stronger bounds on the contribution to total variation distance of the entire level.

This accounting for total variance is more difficult here due to the stronger dependence between the recursive calls. Specifically, the input to the graph on  $V_2$  depends on the set of edges chosen in the first recursive call on  $V_1$ , specifically  $\text{SC}(G, V_1)$ , or a sparsified version of it.

Accounting for this dependency will require proving additional concentration bounds shown in Section 8, which we specifically achieve by sampling  $s = O(n^2\delta^{-1})$  edges in each call to `SCHURSPARSE`. While this might seem contradictory to the notion of “sampling”, we instead consider this to be sampling from the graph in which all the edges generated from the Schur complement are kept separate and could be far more than  $n^2$  edges.

## 7.1 Exact $O(n^\omega)$ Time Recursive Algorithm

We start by showing an algorithm that samples trees from the exact  $\mathbf{w}$ -uniform distribution via the computation of Schur complements. Its pseudocode is in Algorithm 7, and it forms the basis of our approximate algorithm: the faster routine in Section 7.2 is essentially the same as inserting sparsification steps between recursive calls.

---

**Algorithm 7:** EXACTTREE( $G$ ) : Take a graph and output a tree randomly from the  $\mathbf{w}$ -uniform distribution

---

**Input:** Graph  $G$

**Output:** A tree randomly generated from the  $\mathbf{w}$ -uniform distribution of  $G$

- 1 If there is only one edge  $e$  in  $G$ , return  $G$  ;
  - 2 Partition  $V$  evenly into  $V_1$  and  $V_2$ ;
  - 3  $T_1 = \text{EXACTTREE}(\text{SC}(G, V_1))$ ;
  - 4 **for** each  $e \in T_1$  **do**
  - 5    with probability  $\frac{w_e(G)}{w_e(\text{SC}(G, V_1))}$ ,  $G \leftarrow G/e$ ,  $T \leftarrow T \cup e$  ;
  - 6    Delete the remaining edges, i.e.,  $G \leftarrow G \setminus E(V_1)$ ;
  - 7  $T_2 = \text{EXACTTREE}(\text{SC}(G, V_2))$ ;
  - 8  $T \leftarrow T \cup \text{PROLONGATETREE}(G, V_1 \sqcup V_2, T_2)$ ;
  - 9 Output  $T$ ;
- 

The procedure PROLONGATETREE is invoked when  $V_1 = V \setminus V_2$  maps a tree  $T_2$  from the Schur complement  $\text{SC}(G, V_2)$  to a tree back in  $G$ . It crucially uses the property that  $V_1$  is an independent set, and its pseudocode is given in Algorithm 8.

**Lemma 7.1.** *The procedure EXACTTREE( $G$ ) will generate a random tree of  $G$  from the  $\mathbf{w}$ -uniform distribution in  $O(n^\omega)$  time.*

The algorithm we give is similar to the divide and conquer approaches of [CDN89, CMN96, HX16, DKP<sup>+</sup>16]. The two main facts used by these approaches can be summarized as follows:

1. Schur complements preserves the leverage score of original edges, and
2. The operation of taking Schur complements, and the operation of deleting or contracting an edge are associative.

We too will make use of these two facts. But unlike all previous approaches, at every stage we need to recurse on only two sub-problems. All previous approaches have a branching factor of at least four.

We can do this by exploiting the structure of the Schur complement when one eliminates an independent set of vertices. We formalize this in Lemma 7.5.

Before we can prove the lemma, we need to state an important property of Schur complements that follows from Fact 2.8. Recall the notation from Section 2 that for a weighted graph  $G = (V, E, \mathbf{w})$ ,  $\mathbf{Pr}_T^G(\cdot)$  denotes the probability of  $\cdot$  over trees  $T$  picked from the  $\mathbf{w}$ -uniform distribution on spanning trees of  $G$ .

---

**Algorithm 8:** PROLONGATETREE( $G, V_1 \sqcup V_2, T_2$ ): prolongating a tree on  $\text{SC}(G, V_2)$  to a tree on  $G$ .

---

**Input:** A graph  $G$ , a splitting of vertices  $V_1 \sqcup V_2$  such that  $V_1$  is an independent set, tree  $T_2$  of  $\text{SC}(G, V_2)$ .

**Output:** A tree in  $G$

```

1  $T \leftarrow \emptyset;$ 
2 for each  $e = xy \in T_2$  do
3   Create distribution  $\lambda_e$ , set  $\lambda_e(\emptyset) = \mathbf{w}_e(G)$ ;
4   for each  $v \in V_1$  such that  $(v, x), (v, y) \in E(G)$  do
5      $\lambda_e(v) = \mathbf{w}_{(v,x)}(G) \mathbf{w}_{(v,y)}(G) \mathbf{d}_v(G)^{-1}$ ;
6   Randomly assign  $f(e)$  to  $\{\emptyset \cup V_1\}$  with probability proportional to  $\lambda$ ;
7 for each  $v \in V_1$  do
8   for each  $e = (x, y) \in T_2$  such that  $(v, x), (v, y) \in E(G)$  do
9     if  $f(e) \neq v$  then
10      Contract  $x$  and  $y$  ;
11   for each contracted vertex  $X$  in the neighborhood of  $v$  do
12     Connect  $X$  to  $v$  with edge  $(v, u) \in G$  with probability proportional to  $w_G((v, u))$  ;
13    $T \leftarrow T \cup (v, u)$ ;
14 Output  $T$ ;

```

---

**Lemma 7.2.** Let  $G$  be a graph with a partition of vertices  $V = V_1 \sqcup V_2$ . Then for any set of edges  $F$  contained in  $G[V_1]$ , the induced subgraph on  $V_1$ , we have:

$$\mathbf{Pr}_T^G(T \cap E(G[V_1]) = F) = \mathbf{Pr}_T^{\text{SC}(G, V_1)}(T \cap E(G[V_1]) = F),$$

where the edges in  $\text{SC}(G, V_1)$  are treated as the sum of  $G[V_1]$  and  $G_{sc}[V_1]$ , the new edges added to the Schur complement.

*Proof.* If  $F$  contains a cycle, then  $\mathbf{Pr}_T^G(T \cap E(G[V_1]) = F) = 0 = \mathbf{Pr}_T^{\text{SC}(G, V_1)}(T \cap E(G[V_1]) = F)$ . Therefore, we will assume  $F$  does not contain any cycle, and we will prove by induction on the size of  $F$ . If  $|F| > |V_1| - 1$ , then  $F$  will have to contain a cycle. When  $|F| = |V_1| - 1$ , then  $F$  will have to be the edge set of a tree in  $\text{SC}(G, V_1)$ . Then by Fact 2.8, the corollary holds. Now suppose that the corollary holds for all  $F$  with  $|F| = |V_1| - 1 - k$ . Now consider some  $F$  with  $|F| = |V_1| - 1 - (k + 1)$ . We know

$$\mathbf{Pr}_T^G(F \subseteq T) = \mathbf{Pr}_T^G(F = (T \cap E(G[V_1]))) + \sum_{F' \supset F} \mathbf{Pr}_T^G(F' = (T \cap E(G[V_1]))) .$$

Since  $|F'| > |F|$ , by assumption

$$\sum_{F' \supset F} \mathbf{Pr}_T^G(F' = (T \cap E(G[V_1]))) = \sum_{F' \supset F} \mathbf{Pr}_T^{\text{SC}(G, V_1)}(F' = (T \cap E(G[V_1]))) ,$$

then by Fact 2.8 we have  $\mathbf{Pr}_T^G(F \subseteq T) = \mathbf{Pr}_T^{\text{SC}(G, V_1)}(F \subseteq T)$ , which implies

$$\mathbf{Pr}_T^G(F = (T \cap E(G[V_1]))) = \mathbf{Pr}_T^{\text{SC}(G, V_1)}(F = (T \cap E(G[V_1]))) .$$

□

The tracking of edges from various layers of the Schur complement leads to another layer of overhead in recursive algorithms. They can be circumvented by merging the edges, generating a random spanning tree, and the ‘unsplit’ the edge by random spanning. The following is a direct consequence of the definition of  $\mathbf{w}(T)$ :

**Lemma 7.3.** *Let  $\widehat{G}$  be a multi-graph, and  $G$  be the simple graph formed by summing the weights of overlapping edges. Then the procedure of:*

1. *Sampling a random spanning tree from  $G$ ,  $T$ .*
2. *For each edge  $e \in T$ , assign it to an original edge from  $\widehat{G}$ ,  $\widehat{e}$  with probability*

$$\frac{\mathbf{w}_{\widehat{e}}(\widehat{G})}{\mathbf{w}_e(G)}.$$

*Produces a  $\mathbf{w}$ -uniform spanning tree from  $\widehat{G}$ , the original multi-graph.*

This then leads to the following proto-algorithm:

1. Partition the vertices (roughly evenly) into

$$V = V_1 \sqcup V_2.$$

2. Generate a  $\mathbf{w}$ -uniform tree of  $\text{Sc}(G, V_1)$ , and create  $F_1 = T \cap E(G[V_1])$  by re-sampling edges in  $G[V_1]$  using Lemma 7.3. By Lemma 7.2, this subset is precisely the intersection of a random spanning tree with  $G[V_1]$ .
3. This means we have ‘decided’ on all edges in  $G[V_1]$ . So we can proceed by contracting all the edges of  $F_1$ , and deleting all the edges corresponding to  $E(G[V_1]) \setminus F_1$ . Let the resulting graph be  $G'$  and let  $V'_1$  be the remaining vertices in  $V_1$  after this contraction.
4. Observe that  $V'_1$  is an independent set, and its complement is  $V_2$ . We can use another recursive call to generate a  $\mathbf{w}$ -uniform tree in  $\text{Sc}(G', V_2)$ . Then we utilize the fact that  $V'_1$  is an independent set to lift this to a tree in  $G'$  efficiently via Lemma 7.5.

Our key idea for reducing the number of recursive calls of the algorithm, that when  $V_1$  (from the partition of vertices  $V = V_1 \sqcup V_2$ ) is an independent set, we can directly lift a tree from  $\text{Sc}(G, V_2)$  to a tree in  $G$ . This will require viewing  $G_{\text{SC}}[V_2]$  as a sum of cliques, one per vertex of  $V_1$ , plus the original edges in  $G[V_2]$ .

**Fact 7.4.** *Given a graph  $G$  and a vertex  $v$ , the graph  $\text{Sc}(G, V \setminus v)$  is the induced graph  $G[V \setminus \{v\}]$  plus a weighted complete graph  $K(v)$  on the neighbors of  $v$ . This graph  $K(v)$  is formed by adding one edge  $xy$  for every pair of  $x$  and  $y$  incident to  $v$  with weight*

$$\frac{\mathbf{w}_{(v,x)} \mathbf{w}_{(v,y)}}{\deg_v},$$

where  $\mathbf{d}_v \stackrel{\text{def}}{=} \sum_x \mathbf{w}_{(v,x)}$  is the weighted degree of  $v$  in  $G$ .

**Lemma 7.5.** *Let  $G$  be a graph on  $n$  vertices and  $V_1$  an independent set. If  $T$  is drawn from the  $\mathbf{w}$ -uniform distribution of  $\text{Sc}(G, V_2)$ , then in  $O(n^2)$  time  $\text{PROLONGATETREE}(G, V_1 \sqcup V_2, T_2)$  returns a tree from the  $\mathbf{w}$ -uniform distribution of  $G$ .*

*Proof.* The running time of  $\text{PROLONGATETREE}$  is  $O(n^2)$  as  $T_2$  has  $\leq n - 1$  edges and  $|V_1| \leq n$ .

Now we will show the correctness. Let  $V_1 = \{v_1, \dots, v_k\}$ . We will represent  $\text{Sc}(G, V_2)$  as a multi-graph arising by Schur complementing out the vertices in  $V_1$  one by one and keeping the new edges created in the process separate from each other as a multi-graph. We represent this multi-graph as

$$\text{Sc}(G, V_2) = G[V_2] + K(v_1) + \dots + K(v_k),$$

where  $G[V_2]$  is the induced subgraph on  $V_2$  and  $K(v_i)$  is the weighted complete graph on the neighbors of  $v_i$ . Then

- By the unsplitting procedure from Lemma 7.3, the function  $f$  maps  $T_2$  to a tree in the multi-graph  $G[V_2] + K(v_1) + \dots + K(v_k)$ , and
- the rest of the sampling steps maps this tree to one in  $G$ .

We will now prove correctness by induction on the size of the independent set  $V_1$ . The case of  $|V_1| = 0$  follows from  $\text{Sc}(G, V_2) = G$ . If  $|V_1| = 1$ , i.e.,  $V_1 = \{v\}$  for some vertex  $v$ , then  $\text{Sc}(G, V_2)$  is  $G[V_2] + K(v)$ . Given a tree  $T_2$  of  $\text{Sc}(G, V_2)$ , the creation of  $f$  will first map  $T_2$  to a tree in the multigraph  $G[V_2] + K(v)$  by randomly deciding for each edge  $e \in T$  to be in  $G(V_1)$  or  $K(v)$  depending on its weight. If we let  $T'(V_2) = T' \cap G[V_2]$ , then by Lemma 7.2,

$$\Pr_T^G(T \cap E(G[V_2]) = T'(V_2)) = \Pr_T^{G[V_2] + K(v)}(T \cap E(G[V_2]) = T'(V_2)).$$

Therefore, we can contract all the edges of  $T'(V_2) \cap G[V_2]$  and delete all other edges of  $G[V_2]$ . This results in a multi-graph star with  $v$  at the center. Now,  $\text{PROLONGATETREE}$  does the following to decide on the remaining edges. For every multi-edge of the star graph obtained by contracting or deleting edges in  $G[V_2]$ , we choose exactly one edge, randomly according to its weight. This process generates a random tree of multi-graph star.

Now we assume that the lemma holds for all  $V'_1$  with  $|V'_1| < k$ . Let  $V_1 = \{v_1, \dots, v_k\}$ . The key thing to note is that when  $V_1$  is an independent set, we can write

$$\text{Sc}(G, V_2) = G[V_2] + K(v_1) + \dots + K(v_k),$$

and

$$\text{Sc}(G, V_2 \cup v_k) = G[V_2 \cup v_k] + K(v_1) + \dots + K(v_{k-1}).$$

Therefore, by the same reasoning as above, we can take a random tree  $T'$  of the multi-graph  $G[V_2] + K(v_1) + \dots + K(v_k)$  and map it to a tree on  $G[V_2 \cup v_k] + K(v_1) + \dots + K(v_{k-1}) = \text{Sc}(G, V_2 \cup v_k)$  by our procedure  $\text{PROLONGATETREE}$ . We then apply our inductive hypothesis on the set  $V_1 \setminus \{v_k\}$  to map  $\text{Sc}(G, V_2 \cup v_k)$  to a tree of  $G$  by  $\text{PROLONGATETREE}$ , which implies the lemma.  $\square$

We also remark that the running time of  $\text{PROLONGATETREE}$  can be reduced to  $O(m \log n)$  using dynamic trees, which can be abstracted as a data structure supporting operations on rooted forests [ST85, AHLT05]. We omit the details here as this does not bottleneck the running time.

With this procedure fixed, we can now show the overall guarantees of the exact algorithm.

*Proof.* of Lemma 7.1 Correctness follows immediately from Lemmas 7.2 and 7.5. The running time of PROLONGATETREE is  $O(n^2)$  and contracting or deleting all edges contained in  $G[V_1]$  takes  $O(m)$  time. Note that in this new contracted graph, the vertex set containing  $V_1$  is an independent set. Furthermore, computing the Schur complement takes  $O(n^\omega)$  time, giving the running time recurrence

$$T(n) = 2T(n/2) + O(n^\omega) = O(n^\omega).$$

□

## 7.2 Fast Random Spanning Tree Sampling using Determinant Sparsification of Schur complement

Next, we note that the most expensive operation from the exact sampling algorithm from Section 7.1 was the Schur complement procedure. Accordingly, we will substitute in our sparse Schur complement procedure to speed up the running time.

However, this will add some complication in applying Line 5 of EXACTTREE. To address this, we need the observation that the SCHURSPARSE procedure can be extended to distinguish edges from the original graph, and the Schur complement in the multi-graph that it produces.

**Lemma 7.6.** *The procedure SCHURSPARSE( $G, V_1, \delta$ ) given in Algorithm 4 can be modified to record whether an edge in its output,  $H^{V_1}$  is a rescaled copy of an edge from the original induced subgraph on  $V_1$ ,  $G[V_1]$ , or one of the new edges generated from the Schur complement,  $G_{SC}(V_1)$ .*

*Proof.* The edges for  $H^{V_1}$  are generated by the random walks via SAMPLEEDGESCHUR( $G, V_1$ ), whose pseudocode is given in Algorithm 3. Each of these produces a walk between two vertices in  $V_1$ , and such a walk belongs to  $G[V_1]$  if it is length 1, and  $G_{SC}(V_1)$  otherwise. □

We can now give our algorithm for generating random spanning trees and prove the guarantees that lead to the main result from Theorem 1.3.

Note that the splitting on Line 7 is mapping  $T_1$  first back to a tree on a the sparsified multi-graph of  $SC(G, V_1)$ : where the rescaled edges that originated from  $G[V_1]$  are tracked separately from the edges that arise from new edges involving random walks that go through vertices in  $V_2$ .

The desired runtime will follow equivalently to the analysis of the determinant algorithm in Section 6 as we are decreasing  $\delta$  proportionally to the number of vertices. It remains to bound the distortion to the spanning tree distribution caused by the calls to SCHURSPARSE.

Bounds on this distortion will not follow equivalently to that of the determinant algorithm, which also substitutes SCHURSPARSE for exact Schur complements, due to the dependencies in our recursive structure. In particular, while the calls to SCHURSPARSE are independent, the graphs that they are called upon depend on the randomness in Line 6 and PROLONGATETREE, which more specifically, are simply the resulting edge contractions/deletions in previously visited vertex partitions within the recursion. Each subgraph SCHURSPARSE is called upon is additionally dependent on the vertex partitioning from ALMOSTINDEPENDENT.

The key idea to our proof will then be a layer-by-layer analysis of distortion incurred by SCHURSPARSE at each layer to the probability of sampling a *fixed* tree. By considering an alternate procedure where we consider exactly sampling a random spanning tree after some layer, along with the fact that our consideration is restricted to a *fixed* tree, this will allow

---

**Algorithm 9:** APPROXTREE( $G, \delta, \bar{n}$ ) Take a graph and output a tree randomly from a distribution  $\delta$ -close to the  $\mathbf{w}$ -uniform distribution

---

**Input:** Graph  $G$ , error parameter  $\delta$ , and initial number of vertices  $\bar{n}$   
**Output:** A tree randomly generated from a distribution  $\delta$ -close to the  $\mathbf{w}$ -uniform distribution of  $G$

```

1  $V_2 \leftarrow \text{ALMOSTINDEPENDENT}(G, \frac{1}{10})$ ;                                {Via Lemma 5.2}
2  $H_1 \leftarrow \text{SCHURSPARSE}(G, V_1, \delta \cdot |V_1|/\bar{n})$ , while tracking whether the edge is from  $G[V_1]$  via the modifications from Lemma 7.6 ;
3  $T_1 = \text{APPROXTREE}(H_1, \delta, \bar{n})$ ;
4  $G' \leftarrow G$  ;
5 for each  $e \in T_1$  do
6   if  $\text{RAND}[0, 1] \leq \mathbf{w}_e^{\text{ori}}(G_1)/\mathbf{w}_e(G_1)$  then
7      $G' \leftarrow G' / \{e\}$  ;
8      $T \leftarrow T \cup \{e\}$ ;
9 Delete all edges between (remaining) vertices in  $V_1$  in  $G'$ ,  $G' \leftarrow G' \setminus E(G'[V_1])$  ;
10  $H_2 \leftarrow \text{SCHURSPARSE}(G', V_2, \delta \cdot |V_2|/\bar{n})$  ;
11  $T_2 = \text{APPROXTREE}(H_2, \delta, n)$ ;
12  $T \leftarrow T \cup \text{PROLONGATETREE}(G, V_1 \sqcup V_2, T_2)$  ;
13 Output  $T$ ;
```

---

us to separate the randomness incurred by calls to SCHURSPARSE from the other sources of randomness mentioned above. Accordingly, we will provide the following definition.

**Definition 7.7.** For any  $L \geq 0$ , the **level- $L$  truncated algorithm** is the algorithm given by modifying APPROXTREE( $G, \delta, \bar{n}$ ) so that all computations of sparsified Schur complements are replaced by exact calls to Schur complements (aka.  $\text{SC}(G, V_1)$  or  $\text{SC}(G', V_2)$ )) after level  $l$ .

The tree distribution  $\mathcal{T}^{(L)}$  is defined as the output of the level- $L$  truncated algorithm.

Note that in particular,  $\mathcal{T}^{(0)}$  is the tree distribution produced by EXACTTREE( $G$ ), or the  $\mathbf{w}$ -uniform distribution; while  $\mathcal{T}^{(O(\log n))}$  is the distribution outputted by APPROXTREE( $G, \delta$ ).

The primary motivation of this definition is that we can separate the randomness between  $\mathcal{T}^{(l)}$  and  $\mathcal{T}^{(l+1)}$  by only the calls to SCHURSPARSE at level  $l+1$ , which will ultimately give the following lemma that we prove at the end of this section

**Lemma 7.8.** For an invocation of APPROXTREE on a graph  $G$  with variance bound  $\delta$ , for any layer  $L > 0$ , we have

$$d_{TV} \left( \mathcal{T}^{(L-1)}, \mathcal{T}^{(L)} \right) \leq O(\sqrt{\delta}).$$

To begin, we consider the differences between  $\mathcal{T}^{(0)}$  and  $\mathcal{T}^{(1)}$  and the probability of sampling a fixed tree  $\hat{T}$  on a recursive call on  $G$ . The most crucial observation is that the two recursive calls to APPROXTREE( $G_1, \delta, \bar{n}$ ) and APPROXTREE( $G_2, \delta, \bar{n}$ ) can be viewed as independent:

**Claim 7.9.** For a call to APPROXTREE( $G, \delta, \bar{n}$ ) (Algorithm 9) to return  $\hat{T}$ , there is only one possible choice of  $G'$  as generated via Lines 4 to 9.

*Proof.* Note that the edges removed from Line 7 are precisely the edges in  $T$  with both endpoints contained in  $V_1$ ,  $E(T[V_1])$ . For a fixed  $\widehat{T}$ , this set is unique, so  $G'$  is unique as well.  $\square$

This allows us to analyze a truncated algorithm by splitting the probabilities into those that occur at level  $l$  or above. Specifically, at the first level, this can be viewed as pairs of graphs  $\text{SC}(G, V_1)$  and  $\text{SC}(G, V_2)$  along with the ‘intended’ trees from them:

**Definition 7.10.** We define the level-one probabilities of returning a pair of trees  $T_1$  and  $T_2$  that belong a pair of graphs  $G_1, G_2$ ,

$$p^{(\leq 1)} \left( (G, G_1, G_2), (T_1, T_2), \widehat{T} \right).$$

as the product of:

1. The probability (from running `ALMOSTINDEPENDENT`) that  $G$  is partitioned into  $V_1 \sqcup V_2$  so that  $\text{SC}(G, V_1) = G_1$  and  $\text{SC}(G', V_2) = G_2$ , where  $G'$  is  $G$  with the edges  $T \cap G[V_1]$  contracted and all other edges in  $G[V_1]$  are deleted.
2. The probability that  $T_1$  is mapped to  $\widehat{T}[V_1]$  in Line 6.
3. The probability that  $T_2$  is mapped to  $\widehat{T}/\widehat{T}[V_1]$  by the call to `PROLONGATETREE` on Line 12.

This definition then allows us to formalize the splitting of probabilities above and below level 1. More importantly, we note that if we instead call `SCHURSPARSE` to generate  $G_1$  and  $G_2$ , this will not affect the level-one probability because (1) both the calls to `ALMOSTINDEPENDENT` and `PROLONGATETREE` do not depend on  $G_1$  and  $G_2$ , and (2) we can consider  $T_1$  to be drawn from the multi-graph of  $G_1$  where we track which edges are from the original graph and which were generated by the Schur complement.

Consequently, the only difference between the distributions  $\mathcal{T}^{(0)}$  and  $\mathcal{T}^{(1)}$  will be the distortion of drawing  $T_1$  and  $T_2$  from  $G_1$  and  $G_2$  vs the sparsified version of  $G_1$  and  $G_2$ . This handling of sparsifiers of the Schur complements is further simplified with by the following observation:

**Claim 7.11.** *The output of `SCHURSPARSE`( $G, V', \delta$ ) is identical to the output of*

$$\text{IDEALSPARSIFY} \left( \text{SC}(G, V'), \tilde{\tau}, n^2 \delta^{-1} \right),$$

for some set of 1.1-approximate statistical leverage scores of  $\text{SC}(G, V')$ ,  $\tilde{\tau}$ .

This can be seen by revisiting the Schur complement sparsification and rejection sampling algorithms from Section 5 and 4.3 which show that this statement also extends to the approximate Schur complements produced on lines 2 and 10 of Algorithm 9.

This means we can let  $\mathcal{H}_1$  and  $\mathcal{H}_2$  denote the distribution produced by `IDEALSPARSIFY` on  $G_1$  and  $G_2$  respectively.

**Lemma 7.12.** *There exists a collection of graphs and tree pairs  $(\vec{\mathcal{G}}, \vec{\mathcal{T}})^{\leq 1}$  such that for any tree  $\widehat{T}$ , with the probabilities given above in Definition 7.10 we have:*

$$\mathbf{Pr}^{\mathcal{T}^{(0)}} \left( \widehat{T} \right) = \sum_{((G, G_1, G_2), (T_1, T_2)) \in (\mathcal{G}, \mathcal{T})^{(\leq 1)}} p^{(\leq 1)} \left( (G, G_1, G_2), (T_1, T_2), \widehat{T} \right) \cdot \mathbf{Pr}^{G_1} (T_1) \cdot \mathbf{Pr}^{G_2} (T_2).$$

and

$$\mathbf{Pr}^{\mathcal{T}^{(1)}}(\widehat{T}) = \sum_{((G, G_1, G_2), (T_1, T_2)) \in (\mathcal{G}, \mathcal{T})^{(\leq 1)}} p^{(\leq 1)}((G, G_1, G_2), (T_1, T_2), \widehat{T}) \cdot \mathbb{E}_{H_1 \in \mathcal{H}_1} [\mathbf{Pr}^{H_1}(T_1)] \cdot \mathbb{E}_{H_2 \in \mathcal{H}_2} [\mathbf{Pr}^{G_2}(T_2)].$$

We can then in turn extend this via induction to multiple levels. It is important to note that in comparing the distributions  $\mathcal{T}^{(L-1)}$  and  $\mathcal{T}^{(L)}$  for  $L \geq 1$  both will make calls to IDEALSPARSIFY through level  $L$ . We will then need to additionally consider the possible graphs generated by sparsification through level  $L$ , then restrict to the corresponding exact graphs at level  $L+1$ .

**Definition 7.13.** We will use  $\vec{\mathcal{G}}^{(\leq L)}, \vec{\mathcal{T}}^{(L)}$  to denote a sequence of graphs on levels up to  $L-1$ , plus the peripheral exact Schur complements on level  $L$ , along with the spanning trees generated on these peripheral graphs.

As these graphs and trees can exist on different vertex sets, we will use  $(\vec{\mathcal{G}}, \vec{\mathcal{T}})^{(\leq L)}$  to denote the set of graph/tree pairs that are on the same set of vertices. For a sequence of graphs  $\vec{G}^{\leq L}$  and a sequence of trees on their peripherals,  $\vec{T}^L$ , we will use

$$p^{(\leq L)}(\vec{G}^{(\leq L)}, \vec{T}^{(L)}, \widehat{T})$$

to denote the product of the probabilities of the level-by-level vertex split and resulting trees mapping back correctly as defined in Definition 7.10, times the probabilities that the subsequent graphs are generated as sparsifiers of the ones above

Furthermore, we will use  $\vec{G}^{(L)}$  to denote just the peripheral graphs, and  $\vec{\mathcal{H}}(\vec{G}^{(L)})$  to denote the product distribution over sparsifiers of these graphs, and  $\vec{H}^{(L)}$  to denote one particular sequence of such sparsifiers on this level. We can also define the probabilities of trees being picked in a vector-wise sense:

$$\mathbf{Pr}^{\vec{G}^{(L)}}(\vec{T}^{(L)}) \stackrel{\text{def}}{=} \prod_j \mathbf{Pr}^{\vec{G}_j^{(L)}}(\vec{T}_j^{(L)}), \quad \mathbf{Pr}^{\vec{H}^{(L)}}(\vec{T}^{(L)}) \stackrel{\text{def}}{=} \prod_j \mathbf{Pr}^{\vec{H}_j^{(L)}}(\vec{T}_j^{(L)}).$$

Applying Lemma 7.12 inductively then allows us to extend this to multiple levels.

**Corollary 7.14.** *There exists a collection of graphs and tree pairs  $(\vec{\mathcal{G}}, \vec{\mathcal{T}})^{(\leq L)}$  such that for any tree  $\widehat{T}$  we have:*

$$\mathbf{Pr}^{\mathcal{T}^{(L-1)}}(\widehat{T}) = \sum_{(\vec{G}^{(\leq L)}, \vec{T}^{(L)}) \in (\mathcal{G}, \mathcal{T})^{(\leq L)}} p^{(\leq L)}(\vec{G}^{(\leq L)}, \vec{T}^{(L)}, \widehat{T}) \cdot \mathbf{Pr}^{\vec{G}^{(L)}}(\vec{T}^{(L)}),$$

and

$$\mathbf{Pr}^{\mathcal{T}^{(L)}}(\widehat{T}) = \sum_{(\vec{G}^{(\leq L)}, \vec{T}^{(L)}) \in (\mathcal{G}, \mathcal{T})^{(\leq L)}} p^{(\leq L)}(\vec{G}^{(\leq L)}, \vec{T}^{(L)}, \widehat{T}) \cdot \mathbb{E}_{\vec{H}^{(L)} \sim \vec{\mathcal{H}}(\vec{G}^{(L)})} [\mathbf{Pr}^{\vec{H}^{(L)}}(\vec{T}^{(L)})].$$

This reduces our necessary proof of bounding the total variation distance between  $\mathcal{T}^{(L-1)}$  and  $\mathcal{T}^{(L)}$  to examining the difference between

$$\mathbf{Pr}^{\vec{G}^{(L)}}(\vec{T}^{(L)}) \quad \text{and} \quad \mathbb{E}_{\vec{H}^{(L)} \sim \vec{\mathcal{H}}(\vec{G}^{(L)})} [\mathbf{Pr}^{\vec{H}^{(L)}}(\vec{T}^{(L)})].$$

Recalling the definition of  $\mathbf{Pr}^{\vec{H}^{(L)}}(\vec{T}^{(L)})$ : we have that the inverse of each probability in the expectation is

$$\mathbf{Pr}^{\vec{H}_j^{(L)}}(\vec{T}_j^{(L)})^{-1} = \frac{\mathcal{T}_{\vec{H}_j^{(L)}}}{\mathbf{w}^{\vec{H}_j^{(L)}}(\vec{T}_j^{(L)})},$$

and we have concentration bounds for the total trees in  $\vec{H}_j^{(L)}$ . However, it is critical to note that this probability is 0 (and cannot be inverted) when  $\vec{T}_j^{(L)}$  is not contained in  $\vec{H}_j^{(L)}$  for some  $j$ .

This necessitates extending our concentration bounds to random graphs where we condition upon a certain tree remaining in the graph. This will be done in the following Lemma, proven in Section 8, and we recall that we set  $s$  such that  $\delta = O(\frac{n^2}{s})$  in SCHURSPARSE.

**Lemma 7.15.** *Let  $G$  be a graph on  $n$  vertices and  $m$  edges,  $\tilde{\tau}$  be an 1.1-approximate estimates of leverage scores,  $s$  be a sample count such that  $s \geq 4n^2$  and  $m \geq \frac{s^2}{n}$ . Let  $\mathcal{H}$  denote the distribution over the outputs of IDEALSPARSIFY( $G, \tilde{\tau}, s$ ), and for any fixed spanning  $\hat{T}$ , let  $\mathcal{H}|_{\hat{T}}$  denote the distribution formed by conditioning on the graph containing  $\hat{T}$ . Then we have:*

$$\Pr_{H \sim \mathcal{H}}[\hat{T} \subseteq H]^{-1} \cdot \mathbb{E}_{H|_{\hat{T}} \sim \mathcal{H}|_{\hat{T}}} \left[ \mathbf{Pr}^{H|_{\hat{T}}}(\hat{T})^{-1} \right] = \left(1 \pm O\left(\frac{n^2}{s}\right)\right) \mathbf{Pr}^G(\hat{T})^{-1},$$

and

$$\Pr_{H \sim \mathcal{H}}[\hat{T} \subseteq H]^{-2} \cdot \mathbf{Var}_{H|_{\hat{T}} \sim \mathcal{H}|_{\hat{T}}} \left[ \mathbf{Pr}^{H|_{\hat{T}}}(\hat{T})^{-1} \right] \leq O\left(\frac{n^2}{s}\right) \mathbf{Pr}^G(\hat{T})^{-2}.$$

Due to the independence of each call to IDEALSPARSIFY, we can apply these concentration bounds across the product

$$\mathbf{Pr}^{\vec{H}^{(L)}}(\vec{T}^{(L)}) = \prod_j \mathbf{Pr}^{\vec{H}_j^{(L)}}(\vec{T}_j^{(L)})$$

and use the fact that  $\delta$  decreases proportionally to vertex size in our algorithm:

**Corollary 7.16.** *For any sequence of peripheral graphs  $\vec{G}^{(l)}$ , with associated sparsifier distribution  $\mathcal{H}^S$ , and any sequence of trees  $\vec{T}^{(L)}$  as defined in Definition 7.13 such that  $\mathbf{Pr}^{\vec{G}^{(L)}}(\vec{T}^{(L)}) > 0$ , we have*

$$\begin{aligned} \Pr_{\vec{H}^{(L)} \sim \vec{\mathcal{H}}(G^{(L)})} \left[ \mathbf{Pr}^{\vec{H}^{(L)}}(\vec{T}^{(L)}) > 0 \right]^{-1} \cdot \mathbb{E}_{\vec{H}^{(L)} \sim \vec{\mathcal{H}}(G^{(L)})} \left[ \mathbf{Pr}^{\vec{H}^{(L)}}(\vec{T}^{(L)})^{-1} \right] \\ = (1 \pm \delta) \mathbf{Pr}^{\vec{G}^{(L)}}(\vec{T}^{(L)})^{-1}, \end{aligned}$$

and

$$\begin{aligned} \Pr_{\vec{H}^{(L)} \sim \vec{\mathcal{H}}(G^{(L)})} \left[ \mathbf{Pr}^{\vec{H}^{(L)}}(\vec{T}^{(L)}) > 0 \right]^{-2} \cdot \mathbb{E}_{\vec{H}^{(L)} \sim \vec{\mathcal{H}}(G^{(L)})} \left[ \mathbf{Pr}^{\vec{H}^{(L)}}(\vec{T}^{(L)})^{-2} \right] \\ \leq (1 + \delta) \mathbf{Pr}^{\vec{G}^{(L)}}(\vec{T}^{(L)})^{-2}. \end{aligned}$$

*Proof.* The independence of the calls to IDEALSPARSIFY, and the definition of

$$\mathbf{Pr}^{\vec{G}^{(L)}}\left(\vec{T}^{(L)}\right) \stackrel{\text{def}}{=} \prod_j \mathbf{Pr}^{\vec{G}_j^{(L)}}\left(\vec{T}_j^{(L)}\right), \quad \mathbf{Pr}^{\vec{H}^{(L)}}\left(\vec{T}^{(L)}\right) \stackrel{\text{def}}{=} \prod_j \mathbf{Pr}^{\vec{H}_j^{(L)}}\left(\vec{T}_j^{(L)}\right).$$

Applying Lemma 7.15 to each call of IDEALSPARSIFY, where  $s$  was set such that  $\delta/\bar{n} = \frac{n^2}{s}$  gives that the total error bounded by

$$\exp\left(\sum_j \frac{|V(G^{(L)})|}{\bar{n}}\right),$$

and the bound then follows from the total size of each level of the recursion being  $O(\bar{n})$ .  $\square$

It then remains to use concentration bounds on the inverse of the desired probability to bound the total variation distance, which can be done by the following lemma which can be viewed as an extension of Lemma 3.7, and is also proven in Section 9.

**Lemma 7.17.** *Let  $\mathcal{U}$  be a distribution over a universe of elements,  $u$ , each associated with random variable  $P_u$  such that*

$$\mathbb{E}_{u \sim \mathcal{U}} [\mathbb{E}[P_u]] = 1,$$

*and for each  $P_u$  we have*

1.  $P_u \geq 0$ , and
2.  $\Pr[P_u > 0]^{-1} \cdot \mathbb{E}_{p \sim P_u | p > 0} [p^{-1}] = 1 \pm \delta$ , and
3.  $\Pr[P_u > 0]^{-2} \mathbb{E}_{p \sim P_u | p > 0} [p^{-2}] \leq 1 + \delta$ ,

*then*

$$\mathbb{E}_{u \sim \mathcal{U}} [|1 - \mathbb{E}[P_u]|] \leq O(\sqrt{\delta}).$$

To utilize this lemma, we observe that the values

$$p^{(\leq L)}\left(\vec{G}^{(\leq L)}, \vec{T}^{(L)}, \hat{T}\right) \cdot \mathbf{Pr}^{\vec{G}^{(L)}}\left(\vec{T}^{(L)}\right)$$

forms a probability distribution over tuples  $\vec{G}^{(\leq L)}, \vec{T}^{(L)}, \hat{T}$ , while the distribution  $\mathcal{H}(\vec{G}^{(L)})$ , once rescaled, can play the role of  $P_u$ . Decoupling the total variation distance per tree into the corresponding terms on pairs of  $\vec{G}^{(\leq L)}, \vec{T}^{(L)}$  then allows us to bound the overall total variation distance between  $\mathcal{T}^{(L-1)}$  and  $\mathcal{T}^{(L)}$ .

*Proof of Lemma 7.8.* By the definition of total variation distance

$$d_{TV}\left(\mathcal{T}^{(L-1)}, \mathcal{T}^{(L)}\right) = \sum_{\hat{T}} \left| \mathbf{Pr}^{\mathcal{T}^{(L-1)}}\left(\hat{T}\right) - \mathbf{Pr}^{\mathcal{T}^{(L)}}\left(\hat{T}\right) \right|.$$

By Corollary 7.14 and triangle inequality we can then upper bound this probability by

$$d_{TV} \left( \mathcal{T}^{(L-1)}, \mathcal{T}^{(L)} \right) \leq \sum_{\hat{T}} \sum_{(\vec{G}^{(\leq L)}, \vec{T}^{(L)}) \in (\mathcal{G}, \mathcal{T})^{(\leq L)}} p^{(\leq L)} \left( \vec{G}^{(\leq L)}, \vec{T}^{(L)}, \hat{T} \right) \\ \cdot \left| \mathbf{Pr}^{\vec{G}^{(L)}} \left( \vec{T}^{(L)} \right) - \mathbb{E}_{\vec{H}^{(L)} \sim \vec{\mathcal{H}}(\vec{G}^{(L)})} \left[ \mathbf{Pr}^{\vec{H}^{(L)}} \left( \vec{T}^{(L)} \right) \right] \right|.$$

The scalar  $p^{(\leq L)}(\vec{G}^{(\leq L)}, \vec{T}^{(L)}, \hat{T})$  is crucially the same for each, and the inner term in the summation is equivalent to

$$\left| p^{(\leq L)} \left( \vec{G}^{(\leq L)}, \vec{T}^{(L)}, \hat{T} \right) \cdot \mathbf{Pr}^{\vec{G}^{(L)}} \left( \vec{T}^{(L)} \right) - p^{(\leq L)} \left( \vec{G}^{(\leq L)}, \vec{T}^{(L)}, \hat{T} \right) \cdot \mathbb{E}_{\vec{H}^{(L)} \sim \vec{\mathcal{H}}(\vec{G}^{(L)})} \left[ \mathbf{Pr}^{\vec{H}^{(L)}} \left( \vec{T}^{(L)} \right) \right] \right|$$

Our goal is to use Lemma 7.17 where  $\mathcal{U}$  here is the distribution over tuples  $(\vec{G}^{(L)}, \vec{T}^{(L)}, \hat{T})$  with density equaling:

$$p^{(\leq L)} \left( \vec{G}^{(\leq L)}, \vec{T}^{(L)}, \hat{T} \right) \cdot \mathbf{Pr}^{\vec{G}^{(L)}} \left( \vec{T}^{(L)} \right),$$

and  $P_u$  is the distribution over the corresponding value of  $\mathcal{H}(\vec{G}^{(L)})$ , with the same density, and values equaling to:

$$\mathbf{Pr}^{\vec{G}^{(L)}} \left( \vec{T}^{(L)} \right)^{-1} \mathbf{Pr}^{\vec{H}^{(L)}} \left( \vec{T}^{(L)} \right).$$

Note that the fact that each  $\vec{T}^L$  maps back to some tree  $\hat{T}$  imply that  $\mathcal{U}$  is a distribution, as well as  $\mathbb{E}_{u \sim \mathcal{U}} [\mathbb{E}[P_u]] = 1$ . A rescaled version of Corollary 7.16 then gives the required conditions for Lemma 7.17, which in turn gives the overall bound.  $\square$

*Proof of Theorem 1.3.* The running time follows the same way as the analysis of the determinant estimation algorithm in the Proof of Theorem 1.2 at the end of Section 6.

For correctness, the total variation distance bound is implied by appropriately setting  $\delta$ , and then invoking the per-layer bound from Lemma 7.8. Note that factors of  $\log n$  are absorbed by the  $\tilde{O}$  notation.

Finally, note that for simplicity our analysis of total variation distance does not account for the failure probability of Lemma 2.4. To account for these, we can simply use the fact that only  $O(n \log n)$  calls to SCHURSPARSE are made. Hence, the probability of any call failing is polynomially small, which can be absorbed into the total variation distance.  $\square$

## 8 Conditional Concentration Bounds

In this section, we extend our concentration bounds to conditioning on a certain tree being in the sampled graph, specifically with the goal of proving Lemma 7.15. By edge splitting arguments similar to those in Section 4.2, it suffices to analyze the case where all edges have about the same leverage score.

**Lemma 8.1.** *Let  $G$  be a graph on  $n$  vertices and  $m$  edges such that all edges have statistical leverage scores  $\bar{\tau}_e \leq \frac{2n}{m}$ , and  $s$  be a sample count such that  $s \geq 4n^2$  and  $m \geq \frac{s^2}{n}$ . Let  $H$  be a subgraph containing  $s$  edges picked at random without replacement, and let  $\mathcal{H}$  denote this*

distribution over subgraphs on  $s$  edges. Furthermore for any fixed spanning tree,  $\hat{T}$ , let  $\mathcal{H}|_{\hat{T}}$  denote the distribution induced by those in  $\mathcal{H}$  that contain  $\hat{T}$ , and use  $H|_{\hat{T}}$  to denote such a graph, then

$$\Pr_{H \sim \mathcal{H}} [\hat{T} \subseteq H]^{-1} \cdot \mathbb{E}_{H|_{\hat{T}} \sim \mathcal{H}|_{\hat{T}}} [\mathbf{Pr}^{H|_{\hat{T}}} (\hat{T})^{-1}] = \left(1 \pm O\left(\frac{n^2}{s}\right)\right) \mathbf{Pr}^G (\hat{T})^{-1},$$

and

$$\Pr_{H \sim \mathcal{H}} [\hat{T} \subseteq H]^{-2} \cdot \mathbf{Var}_{H|_{\hat{T}} \sim \mathcal{H}|_{\hat{T}}} [\mathbf{Pr}^{H|_{\hat{T}}} (\hat{T})^{-1}] \leq O\left(\frac{n^2}{s}\right) \mathbf{Pr}^G (\hat{T})^{-2}.$$

Note that the ‘uniform leverage score’ requirement here is not as strict as the analysis from Lemma 4.2. This is because we’re eventually aiming for a bound of  $s \approx n^2$  samples. This also means that constant factor leverage score approximations suffices for this routine.

The starting point of this proof is the observation that because we’re doing uniform sampling, the only term in

$$\mathbf{Pr}^{H|_{\hat{T}}} (\hat{T}) = \frac{\mathbf{w}^{H|_{\hat{T}}} (\hat{T})}{\mathcal{T}_{H|_{\hat{T}}}} = \frac{\mathbf{w}^G (\hat{T})}{\mathcal{T}_{H|_{\hat{T}}}}$$

that is dependent on  $H|_{\hat{T}}$  is  $\mathcal{T}_{H|_{\hat{T}}}$ . The proof will then follow by showing concentration of this variable which will be done similarly to the concentration of  $\mathcal{T}_H$  that was done in Section 3 and 4.

The primary difficulty of extending the proof will come from the fact that trees will have different probabilities of being in the sampled graph depending on how many edges they share with  $\hat{T}$ . Much of this will be dealt with by the assumption that  $s \geq 4n^2$ , which makes the exponential terms in the probabilities associated with a tree being in a sampled graph negligible. Additionally, this assumption implies that for any fixed tree  $\hat{T}$  the expected number of edges it shares with a random tree is close to 0. As a result, trees that intersect with  $\hat{T}$  will have negligible contributions, and our analysis can follow similarly to that in Section 3 and 4.

We further note that due to the larger sample count of  $s \geq 4n^2$ , the concentration bounds in this section will also hold, and would in fact be slightly simpler to prove, if the edges were sampled independently with probability  $s/m$ . We keep our assumption of sampling  $s$  edges globally without replacement though in order to avoid changing our algorithm, and the analysis will not require much additional work.

The section will be organized as follows: In Section 8.1 we give upper and lower bounds on the expectation of  $\mathcal{T}_{H|_{\hat{T}}}$ . In Section 8.2 we give an upper bound on the variance of  $\mathcal{T}_{H|_{\hat{T}}}$ . In Section 8.3 we combine the bounds from the previous two sections to prove Lemma 8.1.

## 8.1 Upper and Lower Bounds on Conditional Expectation

In order to prove upper and lower bounds on  $\mathbb{E}_{H|_{\hat{T}}} [\mathcal{T}_{H|_{\hat{T}}}]$ , we will first give several helpful definitions, corollaries, and lemmas to assist in the proof. Our examination of  $\mathbb{E}_{H|_{\hat{T}}} [\mathcal{T}_{H|_{\hat{T}}}]$  will require approximations of  $\Pr_{H|_{\hat{T}}} [T \subseteq H|_{\hat{T}}]$ , and, as we are now fixing  $n - 1$  edges and drawing  $s - n + 1$  edges from the remaining  $m - n + 1$  edges, each edge will now have probability  $\frac{s-n+1}{m-n+1}$  of being in the sampled graph. We will denote this probability with

$$\hat{p} \stackrel{\text{def}}{=} \frac{s - n + 1}{m - n + 1}.$$

It will often be easier to exchange  $\hat{p}$  for

$$p \stackrel{\text{def}}{=} \frac{s}{m},$$

the probability of a single edge being picked without the conditioning on  $\hat{T}$ . The errors of doing so is governed by:

$$\left(1 - \frac{n}{s}\right)p = \frac{s-n}{m} \leq \frac{s-n+1}{m-n+1} = \hat{p} \leq \frac{s}{m} = p. \quad (4)$$

We remark that these errors turn out to be acceptable even when  $\hat{p}$  is raised to the  $O(n)$  power.

Furthermore, our assumption of  $s \geq 4n^2$  implies that we expect a randomly chosen tree not to intersect with  $\hat{T}$ . This will often implicitly show up in the form of the geometric series below, for which a bound is immediately implied by our assumption.

**Lemma 8.2.** *If  $s \geq 4n^2$ , then*

$$\sum_{k=1}^{\infty} \left(\frac{2n^2}{s}\right)^k = O\left(\frac{n^2}{s}\right).$$

The change in our sampling procedure will alter the formulation of  $\Pr_{H|\hat{T}}[T \subseteq H|\hat{T}]$ , so we first want to write  $\mathbb{E}_{H|\hat{T}}[\mathcal{T}_{H|\hat{T}}]$  in terms of values that we are familiar with while only losing small errors. Additionally, many of the exponential terms in the previous analysis will immediately be absorbed into approximation error by our assumption that  $s \geq 4n^2$ .

**Lemma 8.3.** *Let  $G$  be a graph on  $n$  vertices and  $m$  edges and  $s$  a value such that  $m \geq \frac{s^2}{n}$ . Fix some tree  $\hat{T} \in G$ . For a random subset of  $s \geq 4n^2$  edges containing  $\hat{T}$ ,  $H|\hat{T} \supseteq \hat{T}$ , we have*

$$\mathbb{E}_{H|\hat{T}}[\mathcal{T}_{H|\hat{T}}] = \left(1 - O\left(\frac{n^2}{s}\right)\right) \sum_{k=0}^{n-1} p^{n-1-k} \sum_{T: |T \cap \hat{T}|=k} \mathbf{w}(T),$$

where  $p = s/m$  is the probability of each edge being picked in the sample.

*Proof.* Given that all edges of  $\hat{T}$  are in  $H|\hat{T}$ , the remaining  $s-n+1$  edges are chosen uniformly from all  $m-n+1$  edges not in  $\hat{T}$ . Accordingly, for any tree  $T \in G$ , the probability  $\Pr_{H|\hat{T}}[T \subseteq H|\hat{T}]$  is obtained by dividing the number of subsets of  $s-n+1$  edges that contain all edges in  $T \setminus \hat{T}$ , against the number of subsets of  $s-n+1$  edges from  $m-n+1$ :

$$\Pr_{H|\hat{T}}[T \subseteq H|\hat{T}] = \binom{m-n+1 - |T \setminus \hat{T}|}{s-n+1 - |T \setminus \hat{T}|} / \binom{m-n+1}{s-n+1} = \frac{(s-n+1)^{|T \setminus \hat{T}|}}{(m-n+1)^{|T \setminus \hat{T}|}}.$$

Following the proof Lemma 3.1, this reduces to

$$\Pr_{H|\hat{T}}[T \subseteq H|\hat{T}] = \hat{p}^{|T \setminus \hat{T}|} \exp\left(-\frac{|T \setminus \hat{T}|^2}{2s} - O\left(\frac{n^3}{s^2}\right)\right),$$

which we can further reduce using the assumption of  $s \geq 4n^2$  to:

$$\Pr_{H|\hat{T}} [T \subseteq H|\hat{T}] = \left(1 - O\left(\frac{n^2}{s}\right)\right) \hat{p}^{|T \setminus \hat{T}|},$$

and in turn obtain via linearity of expectation:

$$\mathbb{E}_{H|\hat{T}} [\mathcal{T}_{H|\hat{T}}] = \left(1 - O\left(\frac{n^2}{s}\right)\right) \sum_T \mathbf{w}(T) \hat{p}^{|T \setminus \hat{T}|}.$$

We then subdivide the summation based on the amount of edges in the intersection of  $T$  and  $\hat{T}$  and move our  $\hat{p}$  term inside the summation

$$\mathbb{E}_{H|\hat{T}} [\mathcal{T}_{H|\hat{T}}] = \left(1 - O\left(\frac{n^2}{s}\right)\right) \sum_{k=0}^{n-1} \hat{p}^{n-1-k} \sum_{T: |T \cap \hat{T}|=k} \mathbf{w}(T).$$

Finally, we can use Equation 4 to replace  $\hat{p}$  by  $p$  because

$$1 \geq \left(1 - \frac{n}{s}\right)^n \geq \left(1 - \frac{2n^2}{s}\right)$$

where  $n^2s < 0.1$ .

□

We will also require a strong lower bound of the expectation. The following lemma shows that most of the trees do not intersect with  $\hat{T}$ . Restricting our consideration to such trees will be much easier to work in obtaining the lower bound on  $\mathbb{E}_{H|\hat{T}} [\mathcal{T}_{H|\hat{T}}]$ .

**Lemma 8.4.** *Let  $G$  be a graph on  $n$  vertices and  $m$  edges such that  $m \geq 4n^2$  and all edges have statistical leverage scores  $\leq \frac{2n}{m}$ . For any tree  $\hat{T} \in G$ .*

$$\sum_{T: |T \cap \hat{T}|=0} \mathbf{w}(T) \geq \left(1 - O\left(\frac{n^2}{s}\right)\right) \mathcal{T}_G.$$

*Proof.* By definition, we can classify the trees by their intersection with  $\hat{T}$ :

$$\mathcal{T}_G = \sum_{k=0}^{n-1} \sum_{T: |T \cap \hat{T}|=k} \mathbf{w}(T).$$

Consider each inner summation and further separating into each possible forest of  $\hat{T}$  with  $k$  edges gives:

$$\sum_{T: |T \cap \hat{T}|=k} \mathbf{w}(T) = \sum_{\substack{F \subseteq \hat{T} \\ |F|=k}} \sum_{\substack{T \\ F=T \cap \hat{T}}} \mathbf{w}(T) \leq \sum_{\substack{F \subseteq \hat{T} \\ |F|=k}} \sum_{\substack{T: F \subseteq T \\ T \subseteq G}} \mathbf{w}(T).$$

Invoking Lemma 3.4 on the inner summation and the fact that there are  $\binom{n-1}{k}$  forests of  $\widehat{T}$  with  $k$  edges, gives an upper bound of

$$\sum_{T: |T \cap \widehat{T}|=k} \mathbf{w}(T) \leq \binom{n-1}{k} \mathcal{T}_G \left( \frac{2n}{m} \right)^k \leq \mathcal{T}_G \left( \frac{2n^2}{m} \right)^k.$$

We will utilize this upper bound for all  $k > 0$  and achieve a lower bound from rearranging our initial summation

$$\sum_{T: |T \cap \widehat{T}|=0} \mathbf{w}(T) = \mathcal{T}_G - \sum_{k=1}^{n-1} \sum_{T: |T \cap \widehat{T}|=k} \mathbf{w}(T) \geq \mathcal{T}_G \left( 1 - \sum_{k=1}^{n-1} \left( \frac{2n^2}{m} \right)^k \right).$$

Applying the assumption of  $m \geq 4n^2$  and Lemma 8.2 gives our desired result.  $\square$

With the necessary tools in place, we will now give upper and lower bounds on the expectation in terms of  $\mathcal{T}_G p^{n-1}$ , which we note is also a close approximation of  $\mathbb{E}_H [\mathcal{T}_H]$  by our assumption that  $s \geq 4n^2$ .

**Lemma 8.5.** *Let  $G$  be a graph on  $n$  vertices and  $m$  edges such that all edges have statistical leverage scores  $\leq \frac{2n}{m}$ , and let  $s$  be such that  $m \geq \frac{s^2}{n}$ . Fix some tree  $\widehat{T} \in G$ . For a random subset of  $s \geq 4n^2$  edges that contain  $\widehat{T}$ ,  $H|_{\widehat{T}} \subseteq \widehat{T}$  we have:*

$$\mathbb{E}_{H|_{\widehat{T}}} [\mathcal{T}_{H|_{\widehat{T}}}] = \left( 1 \pm O \left( \frac{n^2}{s} \right) \right) \mathcal{T}_G p^{n-1}.$$

*Proof.* We will first prove the upper bound. From Lemma 8.3 we have

$$\mathbb{E}_{H|_{\widehat{T}}} [\mathcal{T}_{H|_{\widehat{T}}}] \leq \sum_{k=0}^{n-1} p^{n-1-k} \sum_{T: |T \cap \widehat{T}|=k} \mathbf{w}(T),$$

while a proof similar to Lemma 8.4 gives

$$\sum_{T: |T \cap \widehat{T}|=k} \mathbf{w}(T) \leq \mathcal{T}_G \left( \frac{2n^2}{m} \right)^k.$$

Moving  $p^{n-1}$  outside the summation and substituting  $\frac{s}{m}$  for  $p$  gives

$$\mathbb{E}_{H|_{\widehat{T}}} [\mathcal{T}_{H|_{\widehat{T}}}] \leq \mathcal{T}_G p^{n-1} \sum_{k=0}^{n-1} \left( \frac{2n^2}{s} \right)^k,$$

and applying Corollary 8.2 to upper bound the summation gives

$$\mathbb{E}_{H|_{\widehat{T}}} [\mathcal{T}_{H|_{\widehat{T}}}] \leq \left( 1 + O \left( \frac{n^2}{s} \right) \right) \mathcal{T}_G p^{n-1}.$$

For the lower bound, we again first using Lemma 8.3 and then restrict to trees that do not intersect  $\widehat{T}$  using Lemma 8.4. Formally we have:

$$\begin{aligned}\mathbb{E}_{H|\widehat{T}} [\mathcal{T}_{H|\widehat{T}}] &= \left(1 - O\left(\frac{n^2}{s}\right)\right) \sum_{k=0}^{n-1} p^{n-1-k} \sum_{T: |T \cap \widehat{T}|=k} \mathbf{w}(T) \\ &\geq \left(1 - O\left(\frac{n^2}{s}\right)\right) p^{n-1} \sum_{T: |T \cap \widehat{T}|=0} \mathbf{w}(T) \geq \left(1 - O\left(\frac{n^2}{s}\right)\right) p^{n-1} \mathcal{T}_G.\end{aligned}$$

□

## 8.2 Upper Bound on Conditional Variance

The bound on variance is by upper bounding  $\mathbb{E}_{H|\widehat{T}} [\mathcal{T}_{H|\widehat{T}}^2]$  in a way similar to Lemma 3.6. Once again, the assumption of  $s > 4n^2$  means the situation is simpler because the exponential term is negligible.

As with the proof of Lemma 3.6, we will often separate summations of pairs of trees based upon the number of edges in their intersection, then frequently invoke Lemma 3.4. However there will be more moving pieces in each summation due to intersections with  $\widehat{T}$ , so Lemma 8.7 proven later in this section, which is analogous to Lemma 3.5, will be much more involved.

**Lemma 8.6.** *Let  $G$  be a graph on  $n$  vertices and  $m$  edges such that all edges have statistical leverage scores  $\leq \frac{2n}{m}$ , and  $s$  a sample count such that  $m \geq \frac{s^2}{n}$ . For some tree  $\widehat{T} \in G$ , let  $H|\widehat{T}$  denote a random subset of  $s$  edges such that  $\widehat{T} \subseteq H|\widehat{T}$ , then:*

$$\frac{\mathbb{E}_{H|\widehat{T}} [\mathcal{T}_{H|\widehat{T}}^2]}{\mathbb{E}_{H|\widehat{T}} [\mathcal{T}_{H|\widehat{T}}]^2} \leq \left(1 + O\left(\frac{n^2}{s}\right)\right).$$

*Proof.* By analogous reasoning to the proof in Lemma 8.3, for any pair of trees  $T_1, T_2 \in G$  we have

$$\Pr_{H|\widehat{T}} [T_1, T_2 \subseteq H|\widehat{T}] = \binom{m-n+1 - |(T_1 \cup T_2) \setminus \widehat{T}|}{s-n+1 - |(T_1 \cup T_2) \setminus \widehat{T}|} / \binom{m-n+1}{s-n+1} = \frac{(s-n+1)_{|(T_1 \cup T_2) \setminus \widehat{T}|}}{(m-n+1)_{|(T_1 \cup T_2) \setminus \widehat{T}|}}.$$

As a consequence of Equation 4, specifically the bound  $\frac{s-k}{m-k} \leq \frac{s}{m}$  when  $k \geq 0$ , we can obtain the upper bound

$$\Pr_{H|\widehat{T}} [T_1, T_2 \subseteq H|\widehat{T}] \leq p^{|(T_1 \cup T_2) \setminus \widehat{T}|},$$

and in turn summing over all pairs of trees:

$$\mathbb{E}_{H|\widehat{T}} [\mathcal{T}_{H|\widehat{T}}^2] \leq \sum_{T_1, T_2} \mathbf{w}(T_1) \mathbf{w}(T_2) p^{|(T_1 \cup T_2) \setminus \widehat{T}|}.$$

We note that  $|(T_1 \cup T_2) \setminus \widehat{T}| = |T_1 \setminus \widehat{T}| + |T_2 \setminus \widehat{T}| - |(T_1 \cap T_2) \setminus \widehat{T}|$ . Furthermore,  $|T_1 \setminus \widehat{T}| = n - 1 - |T_1 \cap \widehat{T}|$ , so we separate the summation as per usual by each possible size of  $|T_1 \cap \widehat{T}|$  and  $|T_2 \cap \widehat{T}|$ , and bring the terms outside of the summation that only depend on these values.

$$\mathbb{E}_{H|\widehat{T}} \left[ \mathcal{T}_{H|\widehat{T}}^2 \right] \leq p^{2n-2} \sum_{k_1, k_2} p^{-k_1-k_2} \sum_{\substack{T_1, T_2 \\ |T_1 \cap \widehat{T}|=k_1 \\ |T_2 \cap \widehat{T}|=k_2}} \mathbf{w}(T_1) \mathbf{w}(T_2) p^{-(T_1 \cap T_2) \setminus \widehat{T}}.$$

In order to deal with the inner most summation we will need to again separate based on the size of  $|(T_1 \cup T_2) \setminus \widehat{T}|$ , and we further note that  $|(T_1 \cap T_2) \setminus \widehat{T}| = |(T_1 \setminus \widehat{T}) \cap (T_2 \setminus \widehat{T})|$ :

$$\mathbb{E}_{H|\widehat{T}} \left[ \mathcal{T}_{H|\widehat{T}}^2 \right] \leq p^{2n-2} \sum_{k_1, k_2} p^{-k_1-k_2} \sum_{k=0}^{n-1} p^{-k} \sum_{\substack{T_1, T_2 \\ |T_1 \cap \widehat{T}|=k_1 \\ |T_2 \cap \widehat{T}|=k_2 \\ |(T_1 \setminus \widehat{T}) \cap (T_2 \setminus \widehat{T})|=k}} \mathbf{w}(T_1) \mathbf{w}(T_2).$$

The last term is bounded in Lemma 8.7, which is stated and proven immediately after this. Incorporating the resulting bound, and grouping the terms by the summations over  $k_1$ ,  $k_2$ , and  $k$  respectively gives:

$$\begin{aligned} \mathbb{E}_{H|\widehat{T}} \left[ \mathcal{T}_{H|\widehat{T}}^2 \right] &\leq p^{2n-2} \sum_{k_1, k_2} p^{-k_1-k_2} \sum_{k=0}^{n-1} p^{-k} \binom{m}{k} \binom{n}{k_1} \binom{n}{k_2} \left( \frac{2n}{m} \right)^{2k+k_1+k_2} \mathcal{T}_G^2 \\ &= \mathcal{T}_G^2 p^{2n-2} \left( \sum_{k_1=0}^{n-1} p^{-k_1} \binom{n}{k_1} \left( \frac{2n}{m} \right)^{k_1} \right) \left( \sum_{k_2=0}^{n-1} p^{-k_2} \binom{n}{k_2} \left( \frac{2n}{m} \right)^{k_2} \right) \left( \sum_{k=0}^{n-1} p^{-k} \binom{m}{k} \left( \frac{2n}{m} \right)^{2k} \right). \end{aligned}$$

We then plug in  $\frac{s}{m}$  for  $p$  in each summation and use the very crude upper bound  $\binom{a}{b} \leq a^b$ :

$$\mathbb{E}_{H|\widehat{T}} \left[ \mathcal{T}_{H|\widehat{T}}^2 \right] \leq \mathcal{T}_G^2 p^{2n-2} \left( \sum_{k_1=0}^{n-1} \left( \frac{2n^2}{s} \right)^{k_1} \right) \left( \sum_{k_2=0}^{n-1} \left( \frac{2n^2}{s} \right)^{k_2} \right) \left( \sum_{k=0}^{n-1} \left( \frac{2n^2}{s} \right)^k \right).$$

Lemma 8.2 then upper bounds each summation by  $1 + O(n^2/s)$ , giving

$$\mathbb{E}_{H|\widehat{T}} \left[ \mathcal{T}_{H|\widehat{T}}^2 \right] \leq \left( 1 + O \left( \frac{n^2}{s} \right) \right) \mathcal{T}_G^2 p^{2n-2}.$$

□

It remains to prove the following bound on the number of pairs of trees with a certain intersection size with  $\widehat{T}$ , and each other. The following Lemma is a generalization to Lemma 3.5, and is proven analogously using the negative correlation of edges in spanning trees from Fact 3.3 and Lemma 3.4.

**Lemma 8.7.** Let  $G$  be graph with  $m$  edges and  $n$  vertices such that every edges has leverage score  $\leq \frac{2n}{m}$ . For any tree  $\widehat{T} \in G$  and any integers  $k, k_1, k_2 \in [0, n-1]$ ,

$$\sum_{\substack{T_1, T_2 \\ |T_1 \cap \widehat{T}| = k_1 \\ |T_2 \cap \widehat{T}| = k_2 \\ |(T_1 \setminus \widehat{T}) \cap (T_2 \setminus \widehat{T})| = k}} \mathbf{w}(T_1) \mathbf{w}(T_2) \leq \binom{m}{k} \binom{n}{k_1} \binom{n}{k_2} \left(\frac{2n}{m}\right)^{2k+k_1+k_2} \mathcal{T}_G^2.$$

*Proof.* We will first separate the summation over all possible forests  $F$  of size  $k$  that could be the intersection of  $T_1 \setminus \widehat{T}$  and  $T_2 \setminus \widehat{T}$ :

$$\sum_{\substack{T_1, T_2 \\ |T_1 \cap \widehat{T}| = k_1 \\ |T_2 \cap \widehat{T}| = k_2 \\ |(T_1 \setminus \widehat{T}) \cap (T_2 \setminus \widehat{T})| = k}} \mathbf{w}(T_1) \mathbf{w}(T_2) = \sum_{\substack{F \subseteq E \\ |F| = k}} \sum_{\substack{T_1, T_2 \\ |T_1 \cap \widehat{T}| = k_1 \\ |T_2 \cap \widehat{T}| = k_2 \\ F = (T_1 \setminus \widehat{T}) \cap (T_2 \setminus \widehat{T})}} \mathbf{w}(T_1) \mathbf{w}(T_2).$$

We first consider the inner summation, and will relax the requirement to only needing

$$F \subseteq (T_1 \setminus \widehat{T}) \cap (T_2 \setminus \widehat{T}),$$

which we note is equivalent to  $F \subseteq (T_1 \setminus \widehat{T})$  and  $F \subseteq (T_2 \setminus \widehat{T})$ . This then allows us to separate the summation again for a particular  $F$  into terms involving just  $T_1$  and  $T_2$ :

$$\sum_{\substack{T_1, T_2 \\ |T_1 \cap \widehat{T}| = k_1 \\ |T_2 \cap \widehat{T}| = k_2 \\ F = (T_1 \setminus \widehat{T}) \cap (T_2 \setminus \widehat{T})}} \mathbf{w}(T_1) \mathbf{w}(T_2) \leq \left( \sum_{\substack{T_1: |T_1 \cap \widehat{T}| = k_1 \\ F \subseteq (T_1 \setminus \widehat{T})}} \mathbf{w}(T_1) \right) \left( \sum_{\substack{T_2: |T_2 \cap \widehat{T}| = k_2 \\ F \subseteq (T_2 \setminus \widehat{T})}} \mathbf{w}(T_2) \right).$$

We further examine the first term in the product, and the second will follow equivalently. Once again, we will split the summation by all possible forests  $\widehat{F}$  of  $\widehat{T}$  with size  $k_1$  that  $T_1 \setminus \widehat{T}$  could intersect in, and further relax to them only having to contain  $\widehat{F}$ .

$$\sum_{\substack{T_1: |T_1 \cap \widehat{T}| = k_1 \\ F \subseteq (T_1 \setminus \widehat{T})}} \mathbf{w}(T_1) \leq \sum_{\substack{\widehat{F} \subseteq \widehat{T} \\ |\widehat{F}| = k_1}} \sum_{\substack{T_1 \\ \widehat{F} \subseteq (T_1 \cap \widehat{T}) \\ F \subseteq (T_1 \setminus \widehat{T})}} \mathbf{w}(T_1).$$

Since  $T_1 \cap \widehat{T}$  and  $T_1 \setminus \widehat{T}$  are disjoint, we can restrict to  $\widehat{F}$  that are disjoint from  $F$ , as well as relaxing to requiring  $(\widehat{F} \cup F) \subseteq T_1$  (instead of  $\widehat{F} \subseteq (T_1 \cap \widehat{T})$  and  $F \subseteq (T_1 \setminus \widehat{T})$ ):

$$\sum_{\substack{T_1: |T_1 \cap \widehat{T}| = k_1 \\ F \subseteq (T_1 \setminus \widehat{T})}} \mathbf{w}(T_1) \leq \sum_{\substack{\widehat{F} \subseteq \widehat{T} \\ |\widehat{F}| = k_1 \\ (\widehat{F} \cap F) = \emptyset}} \sum_{\substack{(F \cup \widehat{F}) \subseteq T}} \mathbf{w}(T).$$

The assumption of  $\widehat{F}$  and  $F$  being disjoint means their union must have exactly  $k + k_1$  edges. We can then apply Lemma 3.4 to the inner summation and use the fact that there are at most  $\binom{n-1}{k_1}$  sets  $\widehat{F}$  to achieve the upper bound

$$\sum_{\substack{T_1: |T_1 \cap \widehat{T}| = k_1 \\ F \subseteq (T_1 \setminus \widehat{T})}} \mathbf{w}(T_1) \leq \binom{n}{k_1} \left(\frac{2n}{m}\right)^{k+k_1} \mathcal{T}_G.$$

Similarly, we can also obtain

$$\sum_{\substack{T_2: |T_2 \cap \widehat{T}| = k_2 \\ F \subseteq (T_2 \setminus \widehat{T})}} \mathbf{w}(T_2) \leq \binom{n}{k_2} \left(\frac{2n}{m}\right)^{k+k_2} \mathcal{T}_G,$$

which, along with the fact that there are  $\binom{m}{k}$  edge sets  $F$  of size  $k$ , gives our desired bound.  $\square$

### 8.3 Concentration of Inverse Probabilities

We now complete a proof of Lemma 8.1 using the concentration results on the number of trees in a sampled graph, conditioned upon a certain tree being contained in the graph.

*Proof of Lemma 8.1.* The definition of

$$\mathbf{Pr}^{H|\widehat{T}}(\widehat{T})^{-1} = \frac{\mathcal{T}_{H|\widehat{T}}}{\mathbf{w}(\widehat{T})}$$

and Lemma 3.1 give

$$\Pr_{H \sim \mathcal{H}} [\widehat{T} \subseteq H]^{-1} \cdot \mathbb{E}_{H|\widehat{T} \sim \mathcal{H}_{|\widehat{T}}} [\mathbf{Pr}^{H|\widehat{T}}(\widehat{T})^{-1}] = \left(\frac{1}{p}\right)^{n-1} \exp\left(\frac{n^2}{2s} + O\left(\frac{n^3}{s^2}\right)\right) \frac{\mathbb{E}_{H|\widehat{T} \sim \mathcal{H}_{|\widehat{T}}} [\mathcal{T}_{H|\widehat{T}}]}{\mathbf{w}(\widehat{T})}.$$

Our condition of  $s \geq 4n^2$  allows us to bound the term  $\exp(n^2/(2s) + O(n^3/s^2))$  by  $(1 + O(n^2/s))$ , and incorporating our approximation of  $\mathbb{E}_{H|\widehat{T} \sim \mathcal{H}_{|\widehat{T}}} [\mathcal{T}_{H|\widehat{T}}]$  from Lemma 8.5 gives

$$\Pr_{H \sim \mathcal{H}} [\widehat{T} \subseteq H]^{-1} \cdot \mathbb{E}_{H|\widehat{T} \sim \mathcal{H}_{|\widehat{T}}} [\mathbf{Pr}^{H|\widehat{T}}(\widehat{T})^{-1}] = \left(1 \pm O\left(\frac{n^2}{s}\right)\right) \cdot \frac{\mathcal{T}_G}{\mathbf{w}(\widehat{T})},$$

and the definition of  $\mathbf{Pr}^G(\widehat{T})^{-1}$  implies the bounds on expectation.

For the variance bound, we use the identity

$$\mathbf{Var}_{H|\widehat{T} \sim \mathcal{H}_{|\widehat{T}}} [\mathbf{Pr}^{H|\widehat{T}}(\widehat{T})^{-1}] = \mathbb{E}_{H|\widehat{T} \sim \mathcal{H}_{|\widehat{T}}} [\mathbf{Pr}^{H|\widehat{T}}(\widehat{T})^{-2}] - \mathbb{E}_{H|\widehat{T} \sim \mathcal{H}_{|\widehat{T}}} [\mathbf{Pr}^{H|\widehat{T}}(\widehat{T})^{-1}]^2,$$

which by the definition

$$\mathbf{Pr}^{H|\widehat{T}}(\widehat{T})^{-1} = \frac{\mathcal{T}_{H|\widehat{T}}}{\mathbf{w}(\widehat{T})}$$

reduces to

$$\mathbf{Var}_{H|\hat{T} \sim \mathcal{H}|\hat{T}} \left[ \mathbf{Pr}^{H|\hat{T}} \left( \hat{T} \right)^{-1} \right] = \frac{\mathbb{E}_{H|\hat{T} \sim \mathcal{H}|\hat{T}} \left[ \mathcal{T}_{H|\hat{T}}^2 \right] - \mathbb{E}_{H|\hat{T} \sim \mathcal{H}|\hat{T}} \left[ \mathcal{T}_{H|\hat{T}} \right]^2}{\mathbf{w} \left( \hat{T} \right)^2} \leq O \left( \frac{n^2}{s} \right) \cdot \frac{\mathcal{T}_G^2 p^{2n-2}}{\mathbf{w}(\hat{T})^2},$$

where the last inequality is from incorporating Lemmas 8.5 and 8.6. Applying Lemma 3.1, and once again using the condition of  $s \geq 4n^2$  to bound

$$\exp \left( \frac{n^2}{2s} + O \left( \frac{n^3}{s^2} \right) \right) \leq \left( 1 + O \left( \frac{n^2}{s} \right) \right) \leq O(1)$$

gives:

$$\Pr_{H \sim \mathcal{H}} \left[ \hat{T} \subseteq H \right]^{-2} \cdot \mathbf{Var}_{H|\hat{T} \sim \mathcal{H}|\hat{T}} \left[ \mathbf{Pr}^{H|\hat{T}} \left( \hat{T} \right)^{-1} \right] \leq O \left( \frac{n^2}{s} \right) \cdot \frac{\mathcal{T}_G^2}{\mathbf{w}(\hat{T})^2},$$

and the variance bound follows from the definition of  $\mathbf{Pr}^G \left( \hat{T} \right)^{-1}$ .  $\square$

## 9 Bounding Total Variation Distance

In this section we will first bound the total variation distance between drawing a tree from the  $\mathbf{w}$ -uniform distribution of  $G$ , and uniformly sampling  $s$  edges,  $H$ , from  $G$ , then drawing a tree from the  $\mathbf{w}$ -uniform distribution of  $H$ . The first bound will only be based on a concentration for the number of trees in  $H$ , and will give the  $\tilde{O}(n^{13/6})$  time algorithm for sampling spanning trees from Corollary 3.8.

Next we will give a more general bound on the total variation distance between two distributions based on concentration of inverse probabilities. The resulting Lemma 7.17 is used for proving the bound on total variation distance in the recursive algorithm given in Section 7. However, as this bound requires a higher sample count of about  $n^2$ , the direct derivation of TV distances from concentration bounds is still necessary for uses of the  $\tilde{O}(n^{1.5})$  edge sparsifier in Corollary 3.8.

### 9.1 Simple Total Variation Distance Bound from Concentration Bounds

We give here a proof of total variation distance being bounded based on the concentration of spanning trees in the sampled graph.

*Proof.* (of Lemma 3.7) Substituting the definition of  $p$  and  $\tilde{p}$  into the definition of total variation distance gives:

$$d_{TV} (p, \tilde{p}) = \sum_{\hat{T}} \left| \mathbf{Pr}^G \left( \hat{T} \right) - \mathbb{E}_{H \sim \mathcal{H}} \left[ \mathbf{Pr}^H \left( \hat{T} \right) \right] \right|.$$

Substituting in the conditions of:

$$\begin{aligned} \mathbf{Pr}^H \left( \hat{T} \right) &= \frac{\mathbf{w}^H(\hat{T})}{\mathcal{T}_H}, & (\text{by definition of } \mathbf{Pr}^H(\hat{T})) \\ \mathbf{w}^H \left( \hat{T} \right) &= \mathbf{w}^G \left( \hat{T} \right) \cdot \Pr_{H' \sim \mathcal{H}} \left[ \hat{T} \subseteq H' \right]^{-1} \cdot \frac{\mathbb{E}_{H' \sim \mathcal{H}} [\mathcal{T}_{H'}]}{\mathcal{T}_G}, & (\text{by given condition}) \end{aligned}$$

Using the fact that

$$\mathbb{E}_{H \sim \mathcal{H}} \left[ \mathbf{1} \left( \widehat{T} \subseteq H \right) \right] = \Pr_{H' \sim \mathcal{H}} \left[ \widehat{T} \subseteq H' \right],$$

we can distribute the first term into:

$$d_{TV} (p, \tilde{p}) = \sum_{\widehat{T}} \left| \mathbb{E}_{H \sim \mathcal{H}} \left[ \mathbf{1} \left( \widehat{T} \subseteq H \right) \cdot \Pr_{H' \sim \mathcal{H}} \left[ \widehat{T} \subseteq H' \right]^{-1} \cdot \mathbf{Pr}^G \left( \widehat{T} \right) - \mathbf{Pr}^H \left( \widehat{T} \right) \right] \right|,$$

which by the condition on  $\mathbf{w}^H(\widehat{T})$  simplifies to:

$$d_{TV} (p, \tilde{p}) = \sum_{\widehat{T}} \left| \mathbb{E}_{H \sim \mathcal{H}} \left[ \mathbf{1} \left( \widehat{T} \subseteq H \right) \cdot \frac{\mathbf{w}^H \left( \widehat{T} \right)}{\mathbb{E}_{H' \sim \mathcal{H}} \left[ \mathcal{T}_{H'} \right]} - \mathbf{Pr}^H \left( \widehat{T} \right) \right] \right|.$$

As  $\mathbf{1}(\widehat{T} \subseteq H) = 1$  iff  $\mathbf{Pr}^H(\widehat{T}) > 0$ , this further simplifies into

$$d_{TV} (p, \tilde{p}) = \sum_{\widehat{T}} \Pr_{H' \sim \mathcal{H}'} \left[ \widehat{T} \subseteq H' \right] \left| \mathbb{E}_{H \sim \mathcal{H}|_T} \left[ \frac{\mathbf{w}^H \left( \widehat{T} \right)}{\mathbb{E}_{H' \sim \mathcal{H}} \left[ \mathcal{T}_{H'} \right]} - \mathbf{Pr}^H \left( \widehat{T} \right) \right] \right|,$$

which by triangle inequality gives:

$$d_{TV} (p, \tilde{p}) = \sum_{\widehat{T}} \Pr_{H' \sim \mathcal{H}'} \left[ \widehat{T} \subseteq H' \right] \cdot \mathbb{E}_{H \sim \mathcal{H}|_T} \left[ \left| \frac{\mathbf{w}^H \left( \widehat{T} \right)}{\mathbb{E}_{H' \sim \mathcal{H}} \left[ \mathcal{T}_{H'} \right]} - \mathbf{Pr}^H \left( \widehat{T} \right) \right| \right],$$

at which point we can rearrange the summation to obtain:

$$d_{TV} (p, \tilde{p}) \leq \mathbb{E}_H \left[ \sum_{\widehat{T} \subseteq H} \left| \mathbf{Pr}^H \left( \widehat{T} \right) - \frac{\mathbf{w}^H \left( \widehat{T} \right)}{\mathbb{E}_{H'} \left[ \mathcal{T}_{H'} \right]} \right| \right] = \mathbb{E}_H \left[ \sum_{\widehat{T} \subseteq H} \mathbf{w}^H \left( \widehat{T} \right) \cdot \left| \frac{1}{\mathcal{T}_H} - \frac{1}{\mathbb{E}_{H'} \left[ \mathcal{T}_{H'} \right]} \right| \right].$$

which by definition of  $\mathcal{T}_H$  simplifies to:

$$d_{TV} (p, \tilde{p}) \leq \mathbb{E}_H \left[ \left| 1 - \frac{\mathcal{T}_H}{\mathbb{E}_{H'} \left[ \mathcal{T}_{H'} \right]} \right| \right].$$

By the Cauchy-Schwarz inequality, which for distributions can be instantiated as  $\mathbb{E}_X [f(X)] \leq \sqrt{\mathbb{E}_X [f(X)^2]}$  for any random variable  $X$  and function  $f(X)$ , we then get:

$$d_{TV} (p, \tilde{p}) \leq \sqrt{\mathbb{E}_H \left[ \left( 1 - \frac{\mathcal{T}_H}{\mathbb{E}_{H'} \left[ \mathcal{T}_{H'} \right]} \right)^2 \right]} = \sqrt{\mathbb{E}_H \left[ \left( \frac{\mathcal{T}_H}{\mathbb{E}_{H'} \left[ \mathcal{T}_{H'} \right]} \right)^2 \right]} - 1 = \sqrt{\delta}.$$

□

## 9.2 Total Variation Distance Bound from Inverse Probability Concentration

We give here our proof of Lemma 7.17, that is a more general bound on total variation distance based upon concentration results of the inverse probabilities.

**Lemma 9.1.** *Let  $X$  be a random variable such that  $X > 0$  over its entire support, and given some  $\delta \geq 0$ , such that  $\mathbb{E}[X] = (1 \pm \delta)\mu$  and  $\mathbf{Var}[X] \leq \delta\mu^2$ , then*

$$\Pr\left[|X^{-1} - \mu^{-1}| > 4k\sqrt{\delta}\mu^{-1}\right] \leq \frac{1}{k^2}$$

if  $1 < k < \delta^{-1/2}/4$

*Proof.* Chebyshev's inequality gives

$$\Pr\left[|X - (1 \pm \delta)\mu| > k\sqrt{\delta}\mu\right] \leq \frac{1}{k^2}.$$

Furthermore, if we assume  $X$  such that

$$|X - (1 \pm \delta)\mu| \leq k\sqrt{\delta}\mu$$

which reduces to

$$(1 - 2k\sqrt{\delta})\mu \leq X \leq (1 + 2k\sqrt{\delta})\mu.$$

Inverting and reversing the inequalities gives

$$\frac{\mu^{-1}}{1 + 2k\sqrt{\delta}} \leq X^{-1} \leq \frac{\mu^{-1}}{1 - 2k\sqrt{\delta}}.$$

Using the fact that  $\frac{1}{1+\epsilon} = 1 - \frac{\epsilon}{1+\epsilon} \leq 1 - \epsilon$  for  $\epsilon > 0$ , and  $\frac{1}{1-\epsilon} = 1 + \frac{\epsilon}{1-\epsilon} \leq 1 + 2\epsilon$  for  $\epsilon \leq 1/2$ , we can then conclude,

$$(1 - 4k\sqrt{\delta})\mu^{-1} \leq X^{-1} \leq (1 + 4k\sqrt{\delta})\mu^{-1},$$

which implies

$$\Pr\left[|X^{-1} - \mu^{-1}| > 4k\sqrt{\delta}\mu^{-1}\right] \leq \Pr\left[|X - (1 \pm \delta)\mu| > k\sqrt{\delta}\mu\right]$$

and proves the lemma.  $\square$

This bound does not allow us to bound  $\mathbb{E}_X[|X - \mu|]$  because when  $X$  close to 0, the value of  $X^{-1}$  can be arbitrarily large, while this bound only bounds the probability of such events by  $O(\delta^{-1})$ . We handle this by treating the case of  $X$  small separately, and account for the total probability of such cases via summations over  $\mathcal{I}$  and  $\widehat{x}$ . First we show that once these distributions are truncated to avoid the small  $X$  case, its variance is bounded.

**Lemma 9.2.** *Let  $Y$  be a random variable such that for parameters  $\delta, \mu_Y > 0$  we have  $0 < Y \leq 2\mu_Y$  over its entire support, and that  $\mathbb{E}[Y^{-1}] = (1 \pm \delta)\mu_Y^{-1}$ ,  $\mathbf{Var}[Y^{-1}] \leq \delta\mu_Y^{-2}$ , then*

$$\mathbb{E}[|Y - \mu_Y|] \leq O\left(\sqrt{\delta}\right)\mu_Y.$$

*Proof.* Since  $|Y - \mu_Y| \leq \mu_Y$ , we can decompose this expected value into buckets of 2 via:

$$\mathbb{E} [|Y - \mu_Y|] \leq \sum_{i=0}^{\log(\delta^{-1/2}/4)} \Pr_Y [|Y - \mu_Y| \geq 2^i \sqrt{\delta} \mu_Y] \cdot (2^i \sqrt{\delta} \mu_Y),$$

where the last term is from the guarantee of  $Y \leq 1$ . Lemma 9.1 gives that each of the intermediate probability terms is bounded by  $O(2^{-2i})$ , while the last one is bounded by  $\frac{1}{\delta}$ , so this gives a total of

$$\mathbb{E} [|Y - \mu|] \leq \sum_{i=0}^{\log(\delta^{-1/2})} (2^i \sqrt{\delta} \mu_Y) O(2^{-2i}) \leq \sqrt{\delta} \mu_Y$$

□

We can now complete the proof via an argument similar to the proof of Lemma 3.7 in Section 9.1. The only additional step is the definition of  $BAD_u$ , which represents the portion of the random variable  $P_u$  with high deviation.

*Proof of Lemma 7.17.* For each  $u$ , we define a scaling factor corresponding to the probability that  $P_u$  is non-zero:

$$p_{u+} \stackrel{\text{def}}{=} \Pr_{p \sim P_u} [p > 0].$$

By triangle inequality, we have for each  $P_u$

$$|1 - \mathbb{E} [P_u]| \leq p_{u+} \cdot \mathbb{E}_{p \sim P_u | p > 0} [|p_{u+}^{-1} - p|].$$

We will handle the case where  $p$  is close and far from  $p_{u+}^{-1}$  separately. This requires defining the portion of  $P_u$  with non-zero values, but large variance as

$$BAD_u \stackrel{\text{def}}{=} \left\{ p \in \text{supp}(P_u) : |p_{u+}^{-1} - p| > \frac{1}{2} p_{u+}^{-1} \right\}.$$

Lemma 9.1 gives that for each  $u$ ,

$$\Pr_{p \sim P_u | p > 0} [p \in BAD_u] \leq O(\sqrt{\delta}),$$

which with the outer distribution and factoring the value of  $p_{u+}^{-1}$  gives gives:

$$\mathbb{E}_{u \sim \mathcal{U}} [p_{u+} \cdot \mathbb{E}_{p \sim P_u | p > 0} [\mathbf{1}(p \in BAD_u) \cdot p_{u+}^{-1}]] \leq O(\sqrt{\delta}), \quad (5)$$

$$\mathbb{E}_{u \sim \mathcal{U}} [p_{u+} \cdot \mathbb{E}_{p \sim P_u | p > 0} [\mathbf{1}(p \notin BAD_u) \cdot p_{u+}^{-1}]] \geq 1 - O(\sqrt{\delta}). \quad (6)$$

We then define the ‘fixed’ distributions  $\tilde{P}_u$  with the same distribution over  $p$  as  $P_u$ , but whose values are set to  $p_{u+}^{-1}$  whenever  $p \in BAD_u$ . Lemma 9.2 then gives:

$$\mathbb{E}_{p \sim \tilde{P}_u | p > 0} [|p_{u+}^{-1} - p|] \leq O(\sqrt{\delta} p_{u+}^{-1}),$$

or taken over the support of  $\mathcal{U}$ , and written with indicator variables:

$$\mathbb{E}_{u \sim \mathcal{U}} [p_{u+} \cdot \mathbb{E}_{p \sim P_u | p > 0} [\mathbf{1}(p \notin BAD_u) \cdot |p_{u+}^{-1} - p|]] \leq O(\sqrt{\delta}).$$

Combining this with the lower bound on the mass of  $p_{u+}^{-1}$  on the complements of the bad sets from Equation 6 via the triangle inequality  $p \geq p_{u+}^{-1} - |p_{u+}^{-1} - p|$  gives:

$$\mathbb{E}_{u \sim \mathcal{U}} [p_{u+} \cdot \mathbb{E}_{p \sim P_u | p > 0} [\mathbf{1}(p \notin BAD_u) \cdot p]] \geq 1 - O(\sqrt{\delta}),$$

or upon taking complement again:

$$\mathbb{E}_{u \sim \mathcal{U}} [p_{u+} \cdot \mathbb{E}_{p \sim P_u | p > 0} [\mathbf{1}(p \in BAD_u) \cdot p]] \leq O(\sqrt{\delta}),$$

which together with Equation 5 and the non-negativity of  $p_{u+}^{-1}$  and  $p$  gives

$$\mathbb{E}_{u \sim \mathcal{U}} [p_{u+} \cdot \mathbb{E}_{p \sim P_u | p > 0} [\mathbf{1}(p \in BAD_u) \cdot |p_{u+}^{-1} - p|]] \leq O(\sqrt{\delta}).$$

Combining these two summations, and invoking the triangle inequality at the start then gives the bound.  $\square$

## References

- [AGM<sup>+</sup>10] Arash Asadpour, Michel X. Goemans, Aleksander Madry, Shayan Oveis Gharan, and Amin Saberi. An  $o(\log n / \log \log n)$ -approximation algorithm for the asymmetric traveling salesman problem. In *Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '10, pages 379–389, Philadelphia, PA, USA, 2010. Society for Industrial and Applied Mathematics.
- [AHLT05] Stephen Alstrup, Jacob Holm, Kristian De Lichtenberg, and Mikkel Thorup. Maintaining information in fully dynamic trees with top trees. *Acm Transactions on Algorithms (talg)*, 1(2):243–264, 2005.
- [Ald90] David Aldous. The random walk construction of uniform spanning trees and uniform labelled trees. In *SIAM Journal on Discrete Mathematics*, pages 450–465, 1990.
- [BDKZ15a] Christos Boutsidis, Petros Drineas, Prabhanjan Kambadur, and Anastasios Zouzias. A randomized algorithm for approximating the log determinant of a symmetric positive definite matrix. *CoRR*, abs/1503.00374, 2015. Available at: <http://arxiv.org/abs/1503.00374>.
- [BDKZ15b] Christos Boutsidis, Petros Drineas, Prabhanjan Kambadur, and Anastasios Zouzias. A randomized algorithm for approximating the log determinant of a symmetric positive definite matrix. *CoRR*, abs/1503.00374, 2015.
- [BK96] András A. Benczúr and David R. Karger. Approximating s-t minimum cuts in  $\tilde{O}(n^2)$  time. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, STOC '96, pages 47–55, New York, NY, USA, 1996. ACM.

- [BP93] Robert Burton and Robin Pemantle. Local characteristics, entropy and limit theorems for spanning trees and domino tilings via transfer-impedances. *The Annals of Probability*, pages 1329–1371, 1993.
- [Bro89] Andrei Broder. Generating random spanning trees. In *Proceedings of the 30th annual Symposium on Foundations of Computer Science, FOCS 1989*, pages 442–447, 1989.
- [BS83] Walter Baur and Volker Strassen. The complexity of partial derivatives. *Theoretical Computer Science*, 22(3):317 – 330, 1983.
- [CCL<sup>+</sup>15] Dehua Cheng, Yu Cheng, Yan Liu, Richard Peng, and Shang-Hua Teng. Efficient sampling for Gaussian graphical models via spectral sparsification. *Proceedings of The 28th Conference on Learning Theory*, pages 364–390, 2015. Available at <http://jmlr.org/proceedings/papers/v40/Cheng15.pdf>.
- [CDN89] Charles J Colbourn, Robert PJ Day, and Louis D Nel. Unranking and ranking spanning trees of a graph. *Journal of Algorithms*, 10(2):271–286, 1989.
- [CKP<sup>+</sup>17] Michael B. Cohen, Jonathan A. Kelner, John Peebles, Richard Peng, Anup Rao, Aaron Sidford, and Adrian Vladu. Almost-linear-time algorithms for markov chains and new spectral primitives for directed graphs. Accepted to *STOC 2017*. Preprint available at <https://arxiv.org/abs/1611.00755>., 2017.
- [CMN96] Charles J Colbourn, Wendy J Myrvold, and Eugene Neufeld. Two algorithms for unranking arborescences. *Journal of Algorithms*, 20(2):268–281, 1996.
- [Coh16] Michael B Cohen. Nearly tight oblivious subspace embeddings by trace inequalities. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 278–287. SIAM, 2016.
- [CP15] Michael B. Cohen and Richard Peng.  $\ell_p$  row sampling by Lewis weights. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, STOC '15, pages 183–192, New York, NY, USA, 2015. ACM. Available at <http://arxiv.org/abs/1412.0588>.
- [DKP<sup>+</sup>16] David Durfee, Rasmus Kyng, John Peebles, Anup B. Rao, and Sushant Sachdeva. Sampling random spanning trees faster than matrix multiplication. *CoRR*, abs/1611.07451, 2016.
- [EGIN97] David Eppstein, Zvi Galil, Giuseppe F. Italiano, and Amnon Nissenzweig. Sparsification—a technique for speeding up dynamic graph algorithms. *J. ACM*, 44(5):669–696, September 1997.
- [FHHP11] Wai Shing Fung, Ramesh Hariharan, Nicholas JA Harvey, and Debmalya Panigrahi. A general framework for graph sparsification. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 71–80. ACM, 2011. <https://arxiv.org/abs/1004.4080>.

- [GRV09] Navin Goyal, Luis Rademacher, and Santosh Vempala. Expanders via random spanning trees. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '09, pages 576–585, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics.
- [Gue83] Alain Guenoche. Random spanning tree. *Journal of Algorithms*, 4(3):214–220, 1983.
- [HAB14a] Timothy Hunter, Ahmed El Alaoui, and Alexandre M. Bayen. Computing the log-determinant of symmetric, diagonally dominant matrices in near-linear time. *CoRR*, abs/1408.1693, 2014. Available at: <http://arxiv.org/abs/1408.1693>.
- [HAB14b] Timothy Hunter, Ahmed El Alaoui, and Alexandre M. Bayen. Computing the log-determinant of symmetric, diagonally dominant matrices in near-linear time. *CoRR*, abs/1408.1693, 2014.
- [HJ12] Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 2012.
- [HMS15] Insu Han, Dmitry Malioutov, and Jinwoo Shin. Large-scale log-determinant computation through stochastic chebyshev expansions. In *ICML*, pages 908–917, 2015. Available at: <https://arxiv.org/abs/1606.00942>.
- [HX16] Nicholas J. A. Harvey and Keyulu Xu. Generating random spanning trees via fast matrix multiplication. In *LATIN 2016: Theoretical Informatics*, volume 9644, pages 522–535, 2016.
- [IL11] Ilse C. F. Ipsen and Dean J. Lee. Determinant approximations, 2011.
- [Jan94] Svante Janson. The numbers of spanning trees, hamilton cycles and perfect matchings in a random graph. *Combinatorics, Probability and Computing*, 3(01):97–126, 1994.
- [JKPS17] Gorav Jindal, Pavel Kolev, Richard Peng, and Saurabh Sawlani. Density independent algorithms for sparsifying k-step random walks. *CoRR*, abs/1702.06110, 2017.
- [Kir47] Gustav Kirchhoff. Über die auflösung der gliechungen, auf welche man bei der untersuchung der linearen vertheilung galvanischer ströme geführt wird. In *Poggendorfs Ann. Phys. Chem.*, pages 497–508, 1847.
- [KLP<sup>+</sup>16] Rasmus Kyng, Yin Tat Lee, Richard Peng, Sushant Sachdeva, and Daniel A Spielman. Sparsified cholesky and multigrid solvers for connection laplacians. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 842–850. ACM, 2016. Available at <http://arxiv.org/abs/1512.01892>.
- [KM09] Jonathan Kelner and Aleksander Madry. Faster generation of random spanning trees. In *Proceedings of the 50th annual Symposium on Foundations of Computer Science, FOCS 2009*, pages 13–21, 2009. Available at <https://arxiv.org/abs/0908.1448>.

- [Kul90] Vidyadhar G. Kulkarni. Generating random combinatorial objects. *Journal of Algorithms*, 11(2):185–207, 1990.
- [MST15] Aleksander Madry, Damian Straszak, and Jakub Tarnawski. Fast generation of random spanning trees and the effective resistance metric. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015*, pages 2019–2036, 2015. Available at <http://arxiv.org/pdf/1501.00267v1.pdf>.
- [PS14] Richard Peng and Daniel A. Spielman. An efficient parallel solver for SDD linear systems. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC ’14, pages 333–342, New York, NY, USA, 2014. ACM. Available at <http://arxiv.org/abs/1311.3286>.
- [SS11] Daniel A. Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011.
- [ST85] Daniel Dominic Sleator and Robert Endre Tarjan. Self-adjusting binary search trees. *Journal of the ACM (JACM)*, 32(3):652–686, 1985.
- [ST11] Daniel A. Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM J. Comput.*, 40(4):981–1025, July 2011.
- [ST14] Daniel A. Spielman and Shang-Hua Teng. Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *SIAM Journal on Matrix Analysis and Applications*, 35(3):835–885, 2014. Available at <http://arxiv.org/abs/cs/0607105>.
- [Tro12] Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Found. Comput. Math.*, 12(4):389–434, August 2012. Available at <http://arxiv.org/abs/1004.4389>.
- [Vis12] N. K. Vishnoi.  $Lx = b$  laplacian solvers and their algorithmic applications, 2012.
- [Wil12] Virginia Vassilevska Williams. Multiplying matrices faster than coppersmith-winograd. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, STOC ’12, pages 887–898, New York, NY, USA, 2012. ACM. Available at: <https://pdfs.semanticscholar.org/3030/fa2aecda339d593b86a260bfab9988b42df7.pdf>.

## A Deferred Proofs

We now provide detailed proofs of the combinatorial facts about random subsets of edges that are discussed briefly in Section 3.

*Proof.* (of Lemma 3.1)

This probability is obtained by dividing the number of subsets of  $s$  edges that contain the  $n-1$  edges in  $T$ , against the number of subsets of  $s$  edges from  $m$ , which using  $\binom{a}{b} = \frac{(a)_b}{(b)_b}$ , gives:

$$\binom{m-n+1}{s-n+1} / \binom{m}{s} = \frac{(m-n+1)_{s-n+1} (s)_s}{(m)_s (s-n+1)_{s-n+1}}, \quad (7)$$

and the two terms can be simplified by the rule  $(a)_b / (a-k)_{b-k} = (a)_k$ .

Furthermore,

$$(a)_b = a^b \left(1 - \frac{1}{a}\right) \cdots \left(1 - \frac{b-1}{a}\right) = a^b \exp \left( \sum_{i=1}^{b-1} \ln \left(1 - \frac{i}{a}\right) \right)$$

We then use the Taylor expansion of  $\ln(1-x) = -\sum_{i=1}^{\infty} \frac{x^i}{i}$  to obtain

$$= a^b \exp \left( -\frac{\sum_{i=1}^{b-1} i}{a} - \frac{\sum_{i=1}^{b-1} i^2}{2a^2} - \frac{\sum_{i=1}^{b-1} i^3}{3a^3} - \dots \right) = a^b \exp \left( -\frac{b^2}{2a} - O\left(\frac{b^3}{a^2}\right) \right)$$

Substituting into  $\frac{(s)_{n-1}}{(m)_{n-1}}$  gives

$$p^{n-1} \exp \left( -\frac{n^2}{2s} + \frac{n^2}{2m} - O\left(\frac{n^3}{s^2}\right) + O\left(\frac{n^3}{m^2}\right) \right) = p^{n-1} \exp \left( -\frac{n^2}{2s} - O\left(\frac{n^3}{s^2}\right) \right)$$

where  $\frac{n^2}{2m}$  is absorbed by  $O\left(\frac{n^3}{s^2}\right)$  because  $m \geq \frac{s^2}{n}$  was assumed.

□

*Proof.* (Of Lemma 3.2)

As before, we have

$$\Pr_H [T_1, T_2 \in H] = p^{|T_1 \cup T_2|} \exp \left( -\frac{|T_1 \cup T_2|^2}{2s} - O\left(\frac{n^3}{s^2}\right) \right)$$

Invoking the identity:

$$|T_1 \cup T_2| = 2n - 2 - |T_1 \cap T_2|$$

gives

$$\Pr_H [T_1, T_2 \in H] = p^{2n-2} p^{-k} \exp \left( -\frac{(2n-2-k)^2}{2s} - O\left(\frac{n^3}{s^2}\right) \right).$$

Using the algebraic identity

$$(2n-2-k)^2 \geq 4n^2 + 4nk$$

and dropping the trailing (negative) lower order term gives:

$$\Pr_H [T_1, T_2 \in H] \leq p^{2n-2} \cdot p^{-k} \exp \left( -\frac{4n^2}{2s} + \frac{4nk}{2s} \right),$$

upon which we can pull out the  $\frac{4n^2}{2s}$  term in the exponential to get a term that only depends  $k$ . Grouping the  $p^{-k}$  term together with the  $\exp(\frac{2n}{s})^k$  term, and using the fact that  $\exp(t) \leq 1 + 2t$  when  $t \leq 0.1$  then gives the result. □

*Proof.* (of Lemma 3.5) We first separate the summation in terms of all possible forests  $F$  of size  $k$  that any pair of trees could intersect in

$$\sum_{\substack{T_1, T_2 \\ |T_1 \cap T_2|=k}} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2) = \sum_{\substack{F \subseteq E \\ |F|=k}} \sum_{\substack{T_1, T_2 \\ F=T_1 \cap T_2}} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2)$$

We then consider the inner summation, the number of pairs of trees  $T_1, T_2$  with  $T_1 \cap T_2 = F$  for some particular set  $F$  of size  $k$ . This is upper bounded by the square of the number of trees containing  $F$ :

$$\sum_{\substack{T_1, T_2 \\ F=T_1 \cap T_2}} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2) \leq \sum_{\substack{T_1, T_2 \\ F \subseteq T_1 \cap T_2}} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2) = \left( \sum_{T: F \subseteq T} \mathbf{w}(T) \right)^2$$

This allow us to directly incorporate the bounds from Lemma 3.4, and in turn the assumption of  $\tau_e \leq \frac{n}{m}$  to obtain the bound:

$$\sum_{\substack{T_1, T_2 \\ F=T_1 \cap T_2}} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2) \leq \left( \mathcal{T}_G \left( \frac{n}{m} \right)^k \right)^2.$$

Furthermore, the number of possible subsets of  $F$  is bounded by  $\binom{m}{k}$ , which can be bounded even more crudely by  $\frac{m^k}{k!}$ . Incorporating this then gives:

$$\sum_{\substack{T_1, T_2 \\ |T_1 \cap T_2|=k}} \mathbf{w}(T_1) \cdot \mathbf{w}(T_2) \leq \frac{m^k}{k!} \cdot \left( \mathcal{T}_G \left( \frac{n}{m} \right)^k \right)^2 = \mathcal{T}_G^2 \cdot \frac{1}{k!} \left( \frac{n^2}{m} \right)^k.$$

□