# Spanning Tree Congestion and Computation of Generalized Györi-Lovász Partition 

L. Sunil Chandran*1, Yun Kuen Cheung ${ }^{\dagger 2}$, and Davis Issac ${ }^{2}$<br>1 Department of Computer Science and Automation, Indian Institute of Science, India<br>sunil@csa.iisc.ernet.in<br>2 Max Planck Institute for Informatics, Saarland Informatics Campus, Germany ycheung@mpi-inf.mpg.de , dissac@mpi-inf.mpg.de


#### Abstract

We study a natural problem in graph sparsification, the Spanning Tree Congestion (STC) problem. Informally, the STC problem seeks a spanning tree with no tree-edge routing too many of the original edges. The root of this problem dates back to at least 30 years ago, motivated by applications in network design, parallel computing and circuit design. Variants of the problem have also seen algorithmic applications as a preprocessing step of several important graph algorithms.

For any general connected graph with $n$ vertices and $m$ edges, we show that its STC is at most $\mathcal{O}(\sqrt{m n})$, which is asymptotically optimal since we also demonstrate graphs with STC at least $\Omega(\sqrt{m n})$. We present a polynomial-time algorithm which computes a spanning tree with congestion $\mathcal{O}(\sqrt{m n} \cdot \log n)$. We also present another algorithm for computing a spanning tree with congestion $\mathcal{O}(\sqrt{m n})$; this algorithm runs in sub-exponential time when $m=\omega\left(n \log ^{2} n\right)$.

For achieving the above results, an important intermediate theorem is generalized GyöriLovász theorem, for which Chen et al. [14] gave a non-constructive proof. We give the first elementary and constructive proof by providing a local search algorithm with running time $\mathcal{O}^{*}\left(4^{n}\right)$, which is a key ingredient of the above-mentioned sub-exponential time algorithm. We discuss a few consequences of the theorem concerning graph partitioning, which might be of independent interest.

We also show that for any graph which satisfies certain expanding properties, its STC is at $\operatorname{most} \mathcal{O}(n)$, and a corresponding spanning tree can be computed in polynomial time. We then use this to show that a random graph has $\operatorname{STC} \Theta(n)$ with high probability.


1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Spanning Tree Congestion, Graph Sparsification, Graph Partitioning, Min-Max Graph Partitioning, $k$-Vertex-Connected Graphs, Győri-Lovász Theorem

## 1 Introduction

Graph Sparsification/Compression generally describes a transformation of a large input graph into a smaller/sparser graph that preserves certain feature (e.g., distance, cut, congestion, flow) either exactly or approximately. The algorithmic value is clear, since the smaller

[^0]
© L. Sunil Chandran, Yun Kuen Cheung and Davis Issac;
licensed under Creative Commons License CC-BY
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany
graph might be used as a preprocessed input to an algorithm, so as to reduce subsequent running time and memory requirement. In this paper, we study a natural problem in graph sparsification, the Spanning Tree Congestion (STC) problem. Informally, the STC problem seeks a spanning tree with no tree-edge routing too many of the original edges. The problem is well-motivated by network design applications, where designers aim to build sparse networks that meet traffic demands, while ensuring no connection (edge) is too congested. Indeed, the root of this problem dates back to at least 30 years ago under the name of "load factor" $[8,36]$, with natural motivations from parallel computing and circuit design applications. The STC problem was formally defined by Ostrovskii [30] in 2004, and since then a number of results have been presented. The probabilistic version of the STC problem, coined as probabilistic capacity mapping, also finds applications in several important graph algorithm problems, e.g., the Min-Bisection problem.

Two canonical goals for graph sparsification problems are to understand the trade-off between the sparsity of the output graph(s) and how well the feature is preserved, and to devise (efficient) algorithms for computing the sparser graph(s). These are also our goals for the STC problem. We focus on two scenarios: (A) general connected graphs with $n$ vertices and $m$ edges, and (B) graphs which exhibit certain expanding properties:

- For (A), we show that the spanning tree congestion (STC) is at most $\mathcal{O}(\sqrt{m n})$, which is a factor of $\Omega(\sqrt{m / n})$ better than the trivial bound of $m$. We present a polynomial-time algorithm which computes a spanning tree with congestion $\mathcal{O}(\sqrt{m n} \cdot \log n)$. We also present another algorithm for computing a spanning tree with congestion $\mathcal{O}(\sqrt{m n})$; this algorithm runs in sub-exponential time when $m=\omega\left(n \log ^{2} n\right)$. For almost all ranges of average degree $2 m / n$, we also demonstrate graphs with STC at least $\Omega(\sqrt{m n})$.
- For (B), we show that the expanding properties permit us to devise polynomial-time algorithm which computes a spanning tree with congestion $\mathcal{O}(n)$. Using this result, together with a separate lower-bound argument, we show that a random graph has $\Theta(n)$ STC with high probability.

For achieving the results for (A), an important intermediate theorem is generalized Györi-Lovász theorem, which was first proved by Chen et al. [14]. Their proof uses advanced techniques in topology and homology theory, and is non-constructive.

- Definition 1. In a graph $G=(V, E)$, a $k$-connected-partition is a $k$-partition of $V$ into $\cup_{j=1}^{k} V_{j}$, such that for each $j \in[k], G\left[V_{j}\right]$ is connected.
- Theorem 2 ([14, Theorems 25, 26]). Let $G=(V, E)$ be a $k$-connected ${ }^{1}$ graph. Let $w$ be a weight function $w: V \rightarrow \mathbb{R}^{+}$. For any $U \subset V$, let $w(U):=\sum_{v \in U} w(v)$. Given any $k$ distinct terminal vertices $t_{1}, \cdots, t_{k}$, and $k$ positive integers $T_{1}, \cdots, T_{k}$ such that for each $j \in[k]$, $T_{j} \geq w\left(t_{j}\right)$ and $\sum_{i=1}^{k} T_{i}=w(V)$, there exists a $k$-connected-partition of $V$ into $\cup_{j=1}^{k} V_{j}$, such that for each $j \in[k], t_{j} \in V_{j}$ and $w\left(V_{j}\right) \leq T_{j}+\max _{v \in V} w(v)-1$.

One of our main contributions is to give the first elementary and constructive proof by providing a local search algorithm with running time $\mathcal{O}^{*}\left(4^{n}\right):^{2}$

- Theorem 3. (a) There is an algorithm which given a $k$-connected graph, computes a $k$-connected-partition satisfying the conditions stated in Theorem 2 in time $\mathcal{O}^{*}\left(4^{n}\right)$.

[^1](b) If we need a $(\lfloor k / 2\rfloor+1)$-partition instead of $k$-partition (the input graph remains assumed to be $k$-connected), the algorithm's running time improves to $\mathcal{O}^{*}\left(2^{\mathcal{O}((n / k) \log k)}\right)$.

We make three remarks. First, the $\mathcal{O}^{*}\left(2^{\mathcal{O}((n / k) \log k)}\right)$-time algorithm is a key ingredient of our algorithm for computing a spanning tree with congestion $\mathcal{O}(\sqrt{m n})$. Second, since Theorem 2 guarantees the existence of such a partition, the problem of computing such a partition is not a decision problem but a search problem. Our local search algorithm shows that this problem is in the complexity class PLS [20]; we raise its completeness in PLS as an open problem. Third, the running times do not depend on the weights.

The STC Problem, Related Problems and Our Results. Given a connected graph $G=(V, E)$, let $T$ be a spanning tree. For an edge $e=(u, v) \in E$, its detour with respect to $T$ is the unique path from $u$ to $v$ in $T$; let $\mathrm{DT}(e, T)$ denote the set of edges in this detour. The stretch of $e$ with respect to $T$ is $|\mathrm{DT}(e, T)|$, the length of its detour. The dilation of $T$ is $\max _{e \in E}|\mathrm{DT}(e, T)|$. The edge-congestion of an edge $e \in T$ is ec $(e, T):=$ $|\{f \in E: e \in \mathrm{DT}(f, T)\}|$, i.e., the number of edges in $E$ whose detours contain $e$. The congestion of $T$ is cong $(T):=\max _{e \in T} \mathrm{ec}(e, T)$. The spanning tree congestion (STC) of the graph $G$ is $\operatorname{STC}(G):=\min _{T}$ cong $(T)$, where $T$ runs over all spanning trees of $G$.

We note that there is an equivalent cut-based definition for edge-congestion, which we will use in our proofs. For each tree-edge $e \in T$, removing $e$ from $T$ results in two connected components; let $U_{e}$ denote one of the components. Then ec $(e, T):=\left|E\left(U_{e}, V \backslash U_{e}\right)\right|$.

Various types of congestion, stretch and dilation problems are studied in computer science and discrete mathematics. In these problems, one typically seeks a spanning tree (or some other structure) with minimum congestion or dilation. We mention some of the well-known problems, where minimization is done over all the spanning trees of the given graph:

1. The Low Stretch Spanning Tree (LSST) problem is to find a spanning tree which minimizes the total stretch of all the edges of $G$. [3] It is easy to see that minimizing the total stretch is equivalent to minimizing the total edge-congestion of the selected spanning tree.
2. The STC problem is to find a spanning tree of minimum congestion. [30]
3. Tree Spanner Problem is to find a spanning tree of minimum dilation. [13] The more general Spanner problem is to find a sparser subgraph of minimum distortion. [4]

There are other congestion and dilation problems which do not seek a spanning tree, but some other structure. The most famous among them is the Bandwidth problem and the Cutwidth problem; see the survey [34] for more details.

Among the problems mentioned above, several strong results were published in connection with the LSST problem. Alon et al. [3] had shown a lower bound of $\Omega(\max \{n \log n, m\})$. Upper bounds have been derived and many efficient algorithms have been devised; the current best upper bound is $\tilde{\mathcal{O}}(m \log n)$. [3, 15, 1, 22, 2] Since total stretch is identical to total edge-congestion, the best upper bound for the LSST problem automatically implies an $\tilde{\mathcal{O}}\left(\frac{m}{n} \log n\right)$ upper bound on the average edge-congestion. But in the STC problem, we concern the maximum edge-congestion; as we shall see, for some graphs, the maximum edge-congestion has to be a factor of $\tilde{\Omega}\left(\sqrt{n^{3} / m}\right)$ larger than the average edge-congestion.

In comparison, there were not many strong and general results for the STC Problem, though it was studied extensively in the past 13 years. The problem was formally proposed by Ostrovskii [30] in 2004. Prior to this, Simonson [36] had studied the same parameter under a different name to approximate the cut width of outer-planar graph. A number of graph-theoretic results were presented on this topic [31, 25, 24, 23, 10]. Some complexity results were also presented recently [29, 9], but most of these results concern special classes
of graphs. The most general result regarding STC of general graphs is an $\mathcal{O}(n \sqrt{n})$ upper bound by Löwenstein, Rautenbach and Regen in 2009 [27], and a matching lower bound by Ostrovskii in 2004 [30]. Note that the above upper bound is not interesting when the graph is sparse, since there is also a trivial upper bound of $m$. In this paper we come up with a strong improvement to these bounds after 8 years:

Theorem (informal): For a connected graph $G$ with $n$ vertices and $m$ edges, its spanning tree congestion is at most $\mathcal{O}(\sqrt{m n})$. In terms of average degree $d_{\text {avg }}=2 m / n$, we can state this upper bound as $\mathcal{O}\left(n \sqrt{d_{\text {avg }}}\right)$. There is a matching lower bound.

Our proof for achieving the $\mathcal{O}(\sqrt{m n})$ upper bound is constructive. It runs in exponential time in general; for graphs with $m=\omega\left(n \log ^{2} n\right)$ edges, it runs in sub-exponential time. By using an algorithm of Chen et al. [14] for computing single-commodity confluent flow from single-commodity splittable flow, we improve the running time to polynomial, but with a slightly worse upper bound guarantee of $\mathcal{O}(\sqrt{m n} \cdot \log n)$.

Motivated by an open problem raised by Ostrovskii [32] concerning STC of random graphs, we formulate a set of expanding properties, and prove that for any graph satisfying these properties, its STC is at most $\mathcal{O}(n)$. We devise a polynomial time algorithm for computing a spanning tree with congestion $\mathcal{O}(n)$ for such graphs. This result, together with a separate lower-bound argument, permit us to show that for random graph $\mathcal{G}(n, p)$ with $1 \geq p \geq \frac{c \log n}{n}$ for some small constant $c>1,{ }^{3}$ its STC is $\Theta(n)$ with high probability, thus resolving the open problem raised by Ostrovskii completely.

Min-Max Graph Partitioning and the Generalized Győri-Lovász Theorem. It looks clear that the powerful Theorem 2 can make an impact on graph partitioning. We discuss a number of its consequences which might be of wider interest.

Graph partitioning/clustering is a prominent topic in graph theory/algorithms, and has a wide range of applications.A popular goal is to partition the vertices into sets such that the number of edges across different sets is small. While the min-sum objective, i.e., minimizing the total number of edges across different sets, is more widely studied, in various applications, the more natural objective is the min-max objective, i.e., minimizing the maximum number of edges leaving each set. The min-max objective is our focus here.

Depending on applications, there are additional constraints on the sets in the partition. Two natural constraints are (i) balancedness: the sets are (approximately) balanced in sizes, and (ii) induced-connectivity: each set induces a connected subgraph. The balancedness constraint appears in the application of domain decomposition in parallel computing, while the induced-connectivity constraint is motivated by divide-and-conquer algorithms for spanning tree construction. Imposing both constraints simultaneously is not feasible for every graph; for instance, consider the star graph with more than 6 vertices and one wants a 3 -partition. Thus, it is natural to ask, for which graphs do partitions satisfying both constraints exist. Theorem 2 implies a simple sufficient condition for existence of such partitions.

By setting the weight of each vertex in $G$ to be its degree, and using the elementary fact that the maximum degree $\Delta(G) \leq n \leq 2 m / k$ for any $k$-connected graph $G$ on $n$ vertices and $m$ edges, we have

- Proposition 4. If $G$ is a $k$-connected graph with $m$ edges, then there exists a $k$-connected-

[^2]partition, such that the total degree of vertices in each part is at most $4 \mathrm{~m} / k$. Consequently, the min-max objective is also at most $4 m / k$.

Due to expander graphs, this bound is optimal up to a small constant factor. This proposition (together with Lemma 9) implies the following crucial lemma for achieving some of our results.

- Lemma 5. Let $G$ be a $k$-connected graph with $m$ edges. Then $\operatorname{STC}(G) \leq 4 m / k$.

Proposition 4 can be generalized to include approximate balancedness in terms of number of vertices. By setting the weight of each vertex to be $\mathrm{cm} / n$ plus its degree in $G$, we have

- Proposition 6. Given any fixed $c>0$, if $G$ is a $k$-connected graph with $m$ edges and $n$ vertices, then there exists a $k$-connected-partition such that the total degree of vertices in each part is at most $(2 c+4) m / k$, and the number of vertices in each part is at most $\frac{2 c+4}{c} \cdot \frac{n}{k}$.

Further Related Work. Concerning STC problem, Okamoto et al. [29] gave an $\mathcal{O}^{*}\left(2^{n}\right)$ algorithm for computing the exact STC of a graph. The probabilistic version of the STC problem, coined as probabilistic capacity mapping, is an important tool for several graph algorithm problems, e.g., the Min-Bisection problem. Räcke [33] showed that in the probabilistic setting, distance and capacity are interchangeable, which briefly says a general upper bound for one objective implies the same general upper bound for the other. Thus, due to the above-mentioned results on LSST, there is an upper bound of $\tilde{\mathcal{O}}(\log n)$ on the maximum average congestion. Räcke's result also implies an $\mathcal{O}(\log n)$ approximation algorithm to the Min-Bisection problem, improving upon the $\mathcal{O}\left(\log ^{3 / 2} n\right)$ approximation algorithm of Feige and Krauthgamer [16]. However, in the deterministic setting, such interchanging phenomenon does not hold: there is a simple tight bound $\Theta(n)$ for dilation, but for congestion it can be as high as $\Theta(n \sqrt{n})$. For the precise definitions, more background and key results about the concepts we have just discussed, we recommend the writing of Andersen and Feige [5].

Graph partitioning/clustering is a prominent research topic with wide applications, so it comes no surprise that a lot of work has been done on various aspects of the topic; we refer readers to the two extensive surveys by Schaeffer [35] and by Teng [41]. Kiwi, Spielman and Teng [21] formulated the min-max $k$-partitioning problem and gave bounds for classes of graphs with small separators, which are then improved by Steurer [38]. On the algorithmic side, many of the related problems are NP-hard, so the focus is on devising approximation algorithms. Sparkled by the seminal work of Arora, Rao and Vazirani [6] on sparsest cut and of Spielman and Teng [37] on local clustering, graph partitioning/clustering algorithms with various constraints have attracted attention across theory and practice; we refer readers to [7] for a fairly recent account of the development. The min-sum objective has been extensively studied; the min-max objective, while striking as the more natural objective in some applications, has received much less attention. The only algorithmic work on this objective (and its variants) are Svitkina and Tardos [40] and Bansal et al. [7]. None of the above work addresses the induced-connectivity constraint.

The classical version of Győri-Lovász Theorem (i.e., the vertex weights are uniform) was proved independently by Győri [17] and Lovász [26]. Lovász's proof uses homology theory and is non-constructive. Győri's proof is elementary and is constructive implicitly, but he did not analyze the running time. Polynomial time algorithms for constructing the $k$-partition were devised for $k=2,3$ [39, 42], but no non-trivial finite-time algorithm was known for general graphs with $k \geq 4 .{ }^{4}$ Recently, Hoyer and Thomas [19] provided a clean presentation

[^3]of Győri's proof by introducing their own terminology, which we use for our constructive proof of Theorem 2.

Notation. Given a connected graph $G=(V, E)$, an edge set $F \subseteq E$ and two disjoint vertex subsets $V_{1}, V_{2} \subset V$, we let $F\left(V_{1}, V_{2}\right):=\left\{e=\left\{v_{1}, v_{2}\right\} \in F \mid v_{1} \in V_{1}\right.$ and $\left.v_{2} \in V_{2}\right\}$.

## 2 Technical Overview

To prove the generalized Győri-Lovász theorem constructively, we follow the same framework of Győri's proof [17], and we borrow terminology from the recent presentation by Hoyer and Thomas [19]. But it should be emphasized that proving our generalized theorem is not straight-forward, since in Győri's proof, at each stage a single vertex is moved from one set to other to make progress, while making sure that the former set remains connected. In our setting, in addition to this we also have to ensure that the weights in the partitions do not exceed the specified limit; and hence any vertex that can be moved from one set to another need not be candidate for being transferred. The proof is presented in Section 3.

As discussed, a crucial ingredient for our upper bound results is Lemma 5, which is a direct corollary of the generalized Győri-Lovász theorem. The lemma takes care of the highlyconnected cases; for other cases we provide a recursive way to construct a low congestion spanning tree; see Section 4 for details. For showing our lower bound for general graphs, the challenge is to maintain high congestion while keeping density small. To achieve this, we combine three expander graphs with little overlapping between them, and we further make those overlapped vertices of very high degree. This will force a tree-edge adjacent to the centroid of any spanning tree to have high congestion; see Section 5 for details.

We formulate a set of expanding properties which permit constructing a spanning tree of better congestion guarantee in polynomial time. The basic idea is simple: start with a vertex $v$ of high degree as the root. Now try to grow the tree by keep attaching new vertices to it, while keeping the invariant that the subtrees rooted at each of the neighbours of $v$ are roughly balanced in size; each such subtree is called a branch. But when trying to grow the tree in a balanced way, we will soon realize that as the tree grow, all the remaining vertices may be seen to be adjacent only to a few number of "heavy" branches. To help the balanced growth, the algorithm will identify a transferable vertex which is in a heavy branch, and it and its descendants in the tree can be transferred to a "lighter" branch. Another technique is to use multiple rounds of matching between vertices in the tree and the remaining vertices to attach new vertices to the tree. This will tend to make sure that all subtrees do not grow uncontrolled. By showing that random graph satisfies the expanding properties with appropriate parameters, we show that a random graph has STC of $\Theta(n)$ with high probability.

## 3 Generalized Győri-Lovász Theorem

We prove Theorem 3 in this section. Observe that the classical Győri-Lovász Theorem follows from Theorem 2 by taking $w(v)=1$ for all $v \in V$ and $T_{j}=n_{j}$ for all $j \in[k]$. We note that a
claimed a poly-time algorithm for all $k$. However, according to a recent study [18], Ma and Ma's algorithm can fall into an endless loop. Also, Győri said the algorithm should be wrong (see [28]).
perfect generalization where one requires that $w\left(V_{j}\right)=T_{j}$ is not possible - think when all vertex weights are even integers, while some $T_{j}$ is odd.

Let $G=(V, E)$ be a $k$-connected graph on $n$ vertices and $m$ edges, and $w: V \rightarrow \mathbb{R}^{+}$be a weight function. For any subset $U \subseteq V, w(U):=\sum_{u \in U} w(u)$. Let $w_{\max }:=\max _{v \in V} w(v)$.

### 3.1 Key Combinatorial Notions

We first highlight the key combinatorial notions used for proving Theorem 3; see Figures 1 and 2 for illustrations of some of these notions.

Fitted Partial Partition. First, we introduce the notion of fitted partial partition (FPP). An FPP $A$ is a tuple of $k$ subsets of $V,\left(A_{1}, \ldots, A_{k}\right)$, such that the $k$ subsets are pairwise disjoint, and for each $j \in[k]$ :

1. $t_{j} \in A_{j}$,
2. $G\left[A_{j}\right]$ is connected and
3. $w\left(A_{j}\right) \leq T_{j}+w_{\max }-1$ (we say the set is fitted for satisfying this inequality).

We say an FPP is a Strict Fitted Partial Partition (SFPP) if $A_{1} \cup \cdots \cup A_{k}$ is a proper subset of $V$. We say the set $A_{j}$ is light if $w\left(A_{j}\right)<T_{j}$, and we say it is heavy otherwise. Note that there exists at least one light set in any SFPP, for otherwise $w\left(A_{1} \cup \cdots \cup A_{k}\right) \geq \sum_{j=1}^{k} T_{j}=w(V)$, which means $A_{1} \cup \cdots \cup A_{k}=V$. Also note that by taking $A_{j}=\left\{t_{j}\right\}$, we have an FPP, and hence at least one FPP exists.

Configuration. For a set $A_{j}$ in an FPP $A$ and a vertex $v \in A_{j} \backslash\left\{t_{j}\right\}$, we define the reservoir of $v$ with respect to $A$, denoted by $R_{A}(v)$, as the vertices in the same connected component as $t_{j}$ in $G\left[A_{j}\right] \backslash\{v\}$. Note that $v \notin R_{A}(v)$.

For a heavy set $A_{j}$, a sequence of vertices $\left(z_{1}, \ldots, z_{p}\right)$ for some $p \geq 0$ is called a cascade of $A_{j}$ if $z_{1} \in A_{j} \backslash\left\{t_{j}\right\}$ and $z_{i+1} \in A_{j} \backslash R_{A}\left(z_{i}\right)$ for all $1 \leq i<p$. The cascade is called a null cascade if $p=0$, i.e., if the cascade is empty. Note that for light set, we do not need to define its cascade since we do not use it in the proof. (See Figure 1.)

A configuration $\mathcal{C}_{A}$ is defined as a pair $(A, D)$, where $A=\left(A_{1}, \cdots, A_{k}\right)$ is an FPP, and $D$ is a set of cascades, which consists of exactly one cascade (possibly, a null cascade) for each heavy set in $A$. A vertex that is in some cascade of the configuration is called a cascade vertex.

Given a configuration, we define rank and level inductively as follows. Any vertex in a light set is said to have level 0 . For $i \geq 0$, a cascade vertex is said to have rank $i+1$ if it has an edge to a level $i$ vertex but does not have an edge to any level- $i^{\prime}$ vertex for $i^{\prime}<i$. A vertex $u$ is said to have level $i$, for $i \geq 1$, if $u \in R_{A}(v)$ for some rank- $i$ cascade vertex $v$, but $u \notin R_{A}(w)$ for any cascade vertex $w$ such that rank of $w$ is less than $i$. A vertex that is not in $R_{A}(v)$ for any cascade vertex $v$ is said to have level $\infty$.

A configuration is called a valid configuration if for each heavy set $A_{j}$, rank is defined for each of its cascade vertices and the rank is strictly increasing in the cascade, i.e., if $\left\{z_{1}, \ldots, z_{p}\right\}$ is the cascade, then $\operatorname{rank}\left(z_{1}\right)<\cdots<\operatorname{rank}\left(z_{p}\right)$. Note that by taking $A_{j}=\left\{t_{j}\right\}$ and taking the null cascade for each heavy set (in this case $A_{j}$ is heavy if $w\left(t_{j}\right)=T_{j}$ ), we get a valid configuration. (See Figure 2.)
Configuration Vectors and Their Total Ordering. For any vertex, we define its neighborhood level as the smallest level of any vertex adjacent to it. A vertex $v$ of level $\ell$ is said to satisfy maximality property if each vertex adjacent on it is either a rank- $(\ell+1)$


Figure 1 Given a configuration $(A, D)$ and a heavy set $A_{j}$ in $A$, the figure shows a cascade $\left(z_{1}, z_{2}, z_{3}\right)$ for the heavy set $A_{j}$ and several reservoirs of the cascade vertices.
For any $z_{\ell}, z_{\ell} \notin R_{A}\left(z_{\ell}\right)$. A cascade vertex $z_{\ell}$ is a cut-vertex of $G\left[A_{j}\right]$, i.e., $G\left[A_{j} \backslash\left\{z_{\ell}\right\}\right]$ is disconnected. The removal of $z_{\ell}$ from $A_{j}$ will lead to at least two connected components in $G\left[A_{j} \backslash\left\{z_{\ell}\right\}\right]$, and the connected component containing $t_{j}$ is the reservoir of $z_{\ell}$.
We identify $t_{j}=z_{0}$, but we clarify that a terminal vertex is never in a cascade. Each epoch between $z_{\ell}$ and $z_{\ell+1}$, and also the epoch above $z_{3}$, is a subset of vertices $B \subset A_{j}$, where $B \ni z_{\ell}$ and $G[B]$ is connected. Note that in general, it is possible that there is no vertex above the last cascade vertex.
cascade vertex, has a level of at most $\ell+1$, or is one of the terminals $t_{j}$ for some $j$. For any $\ell \geq 0$, a valid configuration is called an $\ell$-maximal configuration if all vertices having level at most $\ell-1$ satisfy the maximality property. Note that by definition, any valid configuration is a 0 -maximal configuration.

For a configuration $\mathcal{C}_{A}=\left(\left(A_{1}, \ldots, A_{k}\right), D\right)$, we define $S_{A}:=V \backslash\left(A_{1} \cup \cdots \cup A_{k}\right)$. An edge $u v$ is said to be a bridge in $\mathcal{C}_{A}$ if $u \in S_{A}, v \in A_{j}$ for some $j \in[k]$, and level $(v) \neq \infty$.

A valid configuration $\mathcal{C}_{A}$ is said to be $\ell$-good if the highest rank of a cascade vertex in $\mathcal{C}_{A}$ is exactly $\ell$ (if there are no cascade vertices, then we take the highest rank as 0 ), $\mathcal{C}_{A}$ is $\ell$-maximal, and all bridges $u v$ in $\mathcal{C}_{A}$ (if any) are such that $u \in S_{A}$ and $\operatorname{level}(v)=\ell$. Note that taking $A_{j}=\left\{t_{j}\right\}$ and taking the null cascade for each heavy set gives a 0 -good configuration.

For each configuration $\mathcal{C}_{A}=(A, D)$, we define a configuration vector as below:

$$
\left(L_{A}, N_{A}^{0}, N_{A}^{1}, N_{A}^{2}, \ldots, N_{A}^{n}\right)
$$

where $L_{A}$ is the number of light sets in $A$, and $N_{A}^{\ell}$ is the total number of all level- $\ell$ vertices in $\mathcal{C}_{A}$.

Next, we define ordering on configuration vectors. Let $\mathcal{C}_{A}$ and $\mathcal{C}_{B}$ be configurations. We say $\mathcal{C}_{A}>{ }_{0} \mathcal{C}_{B}$ if

- $L_{A}<L_{B}$, or
- $L_{A}=L_{B}$, and $N_{A}^{0}>N_{B}^{0}$.


Figure 2 An instance of a valid configuration. Every blue segment/curve represent an edge from a cascade vertex to a vertex in some reservoir or light set.
Every cascade vertex connected to a light set has rank 1, and all vertices in the epoch immediately below a rank 1 cascade vertex are of level 1. Inductively, every cascade vertex connected to a vertex of level $i$ has rank $i+1$, and all vertices in the epoch immediately below a rank $i$ cascade vertex are of level $i$. All vertices above the last cascade vertex of each cascade has level $\infty$.

We say $\mathcal{C}_{A}={ }_{0} \mathcal{C}_{B}$ if $L_{A}=L_{B}$ and $N_{A}^{0}=N_{B}^{0}$. We say $\mathcal{C}_{A} \geq{ }_{0} \mathcal{C}_{B}$ if $\mathcal{C}_{A}={ }_{0} \mathcal{C}_{B}$ or $\mathcal{C}_{A}>{ }_{0} \mathcal{C}_{B}$. We say $\mathcal{C}_{A}={ }_{\ell} \mathcal{C}_{B}$ if $L_{A}=L_{B}$, and $N_{A}^{\ell^{\prime}}=N_{B}^{\ell^{\prime}}$ for all $\ell^{\prime} \leq \ell$.

For $1 \leq \ell \leq n$, we say $\mathcal{C}_{A}>_{\ell} \mathcal{C}_{B}$ if

- $\mathcal{C}_{A}>_{\ell-1} \mathcal{C}_{B}$, or
- $\mathcal{C}_{A}={ }_{\ell-1} \mathcal{C}_{B}$, and $N_{A}^{\ell}>N_{B}^{\ell}$.

We say $\mathcal{C}_{A} \geq_{\ell} \mathcal{C}_{B}$ if $\mathcal{C}_{A}={ }_{\ell} \mathcal{C}_{B}$ or $\mathcal{C}_{A}>_{\ell} \mathcal{C}_{B}$. We say $\mathcal{C}_{A}>\mathcal{C}_{B}\left(\mathcal{C}_{A}\right.$ is strictly better than $\left.\mathcal{C}_{B}\right)$ if $\mathcal{C}_{A}>{ }_{n} \mathcal{C}_{B}$.

### 3.2 Proof of Theorem 3

We use two technical lemmas about configuration vectors and their orderings to prove Theorem 3(a). The proof of Theorem 3(b) follows closely with the proof of Theorem 3(a), but makes use of an observation about the rank of a vertex in the local search algorithm, to give an improved bound on the number of configuration vectors navigated by the algorithm.

- Lemma 7. Given any $\ell$-good configuration $\left.\mathcal{C}_{A}=\left(A=\left(A_{1}, \ldots, A_{k}\right), D_{A}\right)\right)$ that does not have a bridge, we can find an $(\ell+1)$-good configuration $\mathcal{C}_{B}=\left(B=\left(B_{1}, B_{2}, \ldots, B_{k}\right), D_{B}\right)$ in polynomial time such that $\mathcal{C}_{B}>\mathcal{C}_{A}$.

Proof. Since $\mathcal{C}_{A}$ is $\ell$-maximal, any vertex that is at level $\ell^{\prime}<\ell$ satisfies maximality property. So, for satisfying $(\ell+1)$-maximality, we only need to worry about the vertices that are at level $\ell$. Let $X_{j}$ be the set of all vertices $x \in A_{j}$ such that $x$ is adjacent to a level- $\ell$ vertex, $\operatorname{level}(x) \geq \ell+1$ (i.e., level $(x)=\infty$ as the highest rank of any cascade vertex is $\ell$ ), $x \neq t_{j}$, and $x$ is not a cascade vertex of rank $\ell$.

We claim that there exists at least one $j$ for which $X_{j}$ is not empty. If that is not the case, then we exhibit a cut set of size at most $k-1$. For each $j$ such that $A_{j}$ is a heavy set with a non-null cascade, let $y_{j}$ be the highest ranked cascade vertex in $A_{j}$. For each $j$ such that $A_{j}$ is a heavy set with a null cascade, let $y_{j}$ be $t_{j}$. Let $Y$ be the set of all $y_{j}$ such that $A_{j}$ is a heavy set. Note that $|Y| \leq k-1$ as $A$ is an SFPP and hence has at least one light set. Let $Z_{\infty}$ be the set of all vertices in $V \backslash Y$ that have level $\infty$ and $Z$ be the remaining vertices in $V \backslash Y$. Since $A$ is an SFPP, $S_{A} \neq \emptyset$, and since all vertices in $S_{A}$ have level $\infty$, we have that $Z_{\infty} \neq \emptyset . Z$ is not empty because there exists at least one light set in $A$ and the vertices in a light set have level 0 . We show that there is no edge between $Z_{\infty}$ and $Z$ in $G$. Suppose there exists an edge $u v$ such that $u \in Z_{\infty}$ and $v \in Z$. If $u \in S_{A}$, then $u v$ is a bridge which is a contradiction by our assumption that $\mathcal{C}_{A}$ does not have a bridge. Hence $u \in A_{j}$ for some $j \in[k]$. Note that $A_{j}$ has to be a heavy set, otherwise $u$ has level 0 . We have that $u$ is not a cascade vertex (as all cascade vertices with level $\infty$ are in $Y$ ) and $u \neq t_{j}$ (as all $t_{j}$ such that level $\left(t_{j}\right)=\infty$ are in $Y$ ). Also, $v$ is not of level $\ell$ as otherwise, $u \in X_{j}$ but we assumed $X_{j}$ is empty. But then, $v$ has level at most $\ell-1, u$ has level $\infty$, and there is an edge $u v$. This means that $\mathcal{C}_{A}$ was not $\ell$-maximal, which is a contradiction. Thus, there exists at least one $j$ for which $X_{j}$ is not empty.

For any $j$ such that $X_{j} \neq \emptyset$, there is at least one vertex $x_{j}$ such that $X_{j} \backslash\left\{x_{j}\right\} \subseteq R_{A}\left(x_{j}\right)$. Now we give the configuration $\mathcal{C}_{B}$ as follows. We set $B_{j}=A_{j}$ for all $j \in[k]$. For each heavy set $A_{j}$ such that $X_{j} \neq \emptyset$, we take the cascade of $B_{j}$ as the cascade of $A_{j}$ appended with $x_{j}$. For each heavy set $A_{j}$ such that $X_{j}=\emptyset$, we take the cascade of $B_{j}$ as the cascade of $A_{j}$. It is easy to see that $\mathcal{C}_{B}$ is $(\ell+1)$-maximal as each vertex that had an edge to level- $\ell$ vertices in $\mathcal{C}_{A}$ is now either a rank $\ell+1$ cascade vertex or a level- $(\ell+1)$ vertex or is $t_{j}$ for some $j$. Also, notice that all the new cascade vertices that we introduce (i.e., the $x_{j}$ 's) have their rank as $\ell+1$ and there is at least one rank $\ell+1$ cascade vertex as $X_{j}$ is not empty for some $j$. Since there were no bridges in $\mathcal{C}_{A}$, all bridges in $\mathcal{C}_{B}$ has to be from $S_{B}$ to a vertex having level $\ell+1$. Hence, $\mathcal{C}_{B}$ is $(\ell+1)$-good. All vertices that had level at most $\ell$ in $\mathcal{C}_{A}$ retained their levels in $\mathcal{C}_{B}$. And, at least one level- $\infty$ vertex of $\mathcal{C}_{A}$ became a level- $(\ell+1)$ vertex in $\mathcal{C}_{B}$ because the cascade vertex that was at rank $\ell$ becomes level- $(\ell+1)$ vertex now in at least one set. Since $\mathcal{C}_{A}$ had no level $(\ell+1)$ vertices, this means that $\mathcal{C}_{B}>\mathcal{C}_{A}$.

- Lemma 8. Given an $\ell$-good configuration $\mathcal{C}_{A}=\left(A=\left(A_{1}, \ldots, A_{k}\right), D_{A}\right)$ having a bridge, we can find in polynomial time a valid configuration $\mathcal{C}_{B}=\left(B=\left(B_{1}, \ldots, B_{k}\right), D_{B}\right)$ such that one of the following holds:
- $\mathcal{C}_{B}>_{\ell} \mathcal{C}_{A}$, and $\mathcal{C}_{B}$ is an $\ell$-good configuration, or
- $\mathcal{C}_{B} \geq{ }_{\ell-1} \mathcal{C}_{A}$, there is a bridge $u^{\prime} v^{\prime}$ in $\mathcal{C}_{B}$ such that $u^{\prime} \in S_{B}$ and level $\left(v^{\prime}\right) \leq \ell-1$, and $\mathcal{C}_{B}$ is an $(\ell-1)$-good configuration.

Proof. Let $u v$ be a bridge where $u \in S_{A}$. Let $A_{j^{*}}$ be the set containing $v$. Note that $\operatorname{level}(v)=\ell$ because $\mathcal{C}_{A}$ is $\ell$-good. We keep $B_{j}=A_{j}$ for all $j \neq j^{*}$. But we modify $A_{j^{*}}$ to get $B_{j^{*}}$ as described below. We maintain that if $A_{j}$ is a heavy set then $B_{j}$ is also a heavy set for all $j$, and hence maintain that $L_{B} \leq L_{A}$.

Case 1: $A_{j^{*}}$ is a light set (i.e., when $\ell=0$ ). We take $B_{j^{*}}=A_{j^{*}} \cup\{u\}$. For all $j$ such that $B_{j}$ is a heavy set, cascade of $B_{j}$ is taken as the null cascade. We have $w\left(A_{j^{*}}\right) \leq T_{j}-1$ because $A_{j^{*}}$ is a light set. So, $w\left(B_{j^{*}}\right)=w\left(A_{j^{*}}\right)+w(u) \leq\left(T_{j}-1\right)+w_{\max }$, and hence $B_{j^{*}}$ is fitted. Also, $G\left[B_{j^{*}}\right]$ is connected and hence $\left(B_{1}, \ldots, B_{k}\right)$ is an FPP. We have $\mathcal{C}_{B}>_{0} \mathcal{C}_{A}$ because either $B_{j^{*}}$ became a heavy set in which case $L_{B}<L_{A}$, or it is a light set in which case $L_{B}=L_{A}$ and $N_{B}^{0}>N_{A}^{0}$. It is easy to see that $\mathcal{C}_{B}$ is 0 -good.

Case 2: $A_{j^{*}}$ is a heavy set i.e., when $\ell \geq 1$.
Case 2.1: $w\left(A_{j^{*}} \cup\{u\}\right) \leq T_{j}+w_{\max }-1$. We take $B_{j^{*}}=A_{j^{*}} \cup\{u\}$. For each $j$ such that $B_{j}$ is a heavy set ( $A_{j}$ is also heavy set for such $j$ ), the cascade of $B_{j}$ is taken as the cascade of $A_{j}$. G[ $\left.B_{j^{*}}\right]$ is clearly connected and $B_{j^{*}}$ is fitted by assumption of the case that we are in. Hence $B$ is indeed an FPP. Observe that all vertices that had level $\ell^{\prime} \leq \ell$ in $\mathcal{C}_{A}$ still has level $\ell^{\prime}$ in $\mathcal{C}_{B}$. Since level $(v)$ was $\ell$ in $\mathcal{C}_{A}$ by $\ell$-goodness of $\mathcal{C}_{A}, u$ also has level $\ell$ in $\mathcal{C}_{B}$; and $u$ had level $\infty$ in $\mathcal{C}_{A}$. Hence, $\mathcal{C}_{B}>{ }_{\ell} \mathcal{C}_{A}$. It is also easy to see that $\mathcal{C}_{B}$ remains $\ell$-good.

Case 2.2: $w\left(A_{j^{*}} \cup\{u\}\right) \geq T_{j}+w_{\max }$. Let $z$ be the cascade vertex of rank $\ell$ in $A_{j^{*}}$. Note that $A_{j^{*}}$ should have such a cascade vertex as $v \in A_{j^{*}}$ has level $\ell$. Let $\bar{R}$ be $A_{j^{*}} \backslash\left(R_{A}(z) \cup z\right)$, i.e., $\bar{R}$ is the set of all vertices in $A_{j^{*}} \backslash\{z\}$ with level $\infty$. We initialize $B_{j^{*}}:=A_{j^{*}} \cup\{u\}$. Now, we delete vertices one by one from $B_{j^{*}}$ in a specific order until $B_{j^{*}}$ becomes fitted. We choose the order of deleting vertices such that $G\left[B_{j^{*}}\right]$ remains connected. Consider a spanning tree $\tau$ of $G[\bar{R} \cup\{z\}]$. $\tau$ has at least one leaf, which is not $z$. We delete this leaf from $B_{j^{*}}$ and $\tau$. We repeat this process until $\tau$ is just the single vertex $z$ or $B_{j^{*}}$ becomes fitted. If $B_{j^{*}}$ is not fitted even when $\tau$ is the single vertex $z$, then delete $z$ from $B_{j^{*}}$. If $B_{j^{*}}$ is still not fitted then delete $u$ from $B_{j^{*}}$. Note that at this point $B_{j^{*}} \subset A_{j^{*}}$ and hence is fitted. Also, note that $G\left[B_{j^{*}}\right]$ remains connected. Hence $\left(B_{1}, \ldots, B_{k}\right)$ is an FPP. $B_{j^{*}}$ does not become a light set because $B_{j}$ became fitted when the last vertex was deleted from it. Before this vertex was deleted, it was not fitted and hence had weight at least $T_{j^{*}}+w_{\max }$ before this deletion. Since the last vertex deleted has weight at most $w_{\max }, B_{j^{*}}$ has weight at least $T_{j^{*}}$ and hence is a heavy set. Now we branch into two subcases for defining the cascades.

Case 2.2.1: $z \in B_{j *}$ (i.e, $z$ was not deleted from $B_{j^{*}}$ in the process above). For each $j$ such that $B_{j}$ is a heavy set, the cascade of $B_{j}$ is taken as the cascade of $A_{j}$. Since a new $\ell$ level vertex $u$ is added and all vertices that had level at most $\ell$ retain their level, we have that $\mathcal{C}_{B}>_{\ell} \mathcal{C}_{A}$. It is also easy to see that $\mathcal{C}_{B}$ remains $\ell$-good.

Case 2.2.2: $z \notin B_{j^{*}}$ (i.e, $z$ was deleted from $B_{j^{*}}$ ). For each $j$ such that $B_{j}$ is a heavy set, the cascade of $B_{j}$ is taken as the cascade of $A_{j}$ but with the rank $\ell$ cascade vertex (if it has any) deleted from it. $\mathcal{C}_{B} \geq_{\ell-1} \mathcal{C}_{A}$ because all vertices that were at a level of $\ell^{\prime}=\ell-1$ or smaller, retain their levels. Observe that there are no bridges in $\mathcal{C}_{B}$ to vertices that are at a level at most $\ell-2$, all vertices at a level at most $\ell-2$ still maintain the maximality property, and we did not introduce any cascade vertices. Hence, $\mathcal{C}_{B}$ is $(\ell-1)$-good. It only remains to prove that there is a bridge $u^{\prime} v^{\prime}$ in $\mathcal{C}_{B}$ such that level $\left(v^{\prime}\right) \leq \ell-1$. We know $z \in S_{B}$. Since $z$ was a rank $\ell$ cascade vertex in $\mathcal{C}_{A}, z$ had an edge to $z^{\prime}$ such that $z^{\prime}$ had level $\ell-1$ in $\mathcal{C}_{A}$. Observe that level of $z^{\prime}$ is at most $\ell-1$ in $\mathcal{C}_{B}$ as well. Hence, taking $u^{\prime} v^{\prime}=z z^{\prime}$ completes the proof.

Proof of Theorem 3(a): We always maintain a configuration $\mathcal{C}_{A}=\left(A, D_{A}\right)$ that is $\ell$-good for some $\ell \geq 0$. If the FPP $A$ is not an SFPP at any point, then we are done. So assume $A$ is an SFPP.

We start with the 0 -good configuration where $A_{j}=\left\{t_{j}\right\}$ and the cascades of all heavy sets are null cascades. If our current configuration $\mathcal{C}_{A}$ is an $\ell$-good configuration that has
no bridge, then we use Lemma 7 to get a configuration $\mathcal{C}_{B}$ such that $\mathcal{C}_{B}>\mathcal{C}_{A}$ and $B$ is $(\ell+1)$-good. We take $\mathcal{C}_{B}$ as the new current configuration $\mathcal{C}_{A}$. If our current configuration $\mathcal{C}_{A}$ is an $\ell$-good configuration with a bridge, then we get an $\ell^{\prime}$-good configuration $\mathcal{C}_{B}$ for some $\ell^{\prime} \geq 0$ such that $\mathcal{C}_{B}>\mathcal{C}_{A}$ by repeatedly applying Lemma 8 at most $\ell$ times. So in either case, we get a strictly better configuration that is $\ell^{\prime}$-good for some $\ell^{\prime} \geq 0$ in polynomial time. We call this an iteration of our algorithm.

Notice that the number of iterations possible is at most the number of distinct configuration vectors possible. It is easy to see that the number of distinct configuration vectors with highest rank at most $r$ is at most $\binom{n+r-1}{n}$. Since rank of any point is at most $n$, the number of iterations of our algorithm is at most $(k+1) \cdot\binom{2 n}{n}$, which is at most $n \cdot 4^{n}$. Since each iteration runs in polynomial time as guaranteed by the two lemmas, the required running time is $\mathcal{O}^{*}\left(4^{n}\right)$.

When the algorithm terminates, the FPP given by the current configuration is not an SFPP and this gives the required partition.

Proof of Theorem 3(b): Since any $k$-connected graph is also $(\lfloor k / 2\rfloor+1)$-vertex connected, the algorithm will give the required partition due to Theorem $3(\mathrm{a})$. We only need to prove the better running time claimed by Theorem 3(b). For this, we show that the highest rank attained by any vertex during the algorithm is at most $2 n /(k-2)$. Since the number of distinct configuration vectors with highest rank $r$ is at most $\binom{n+r-1}{n}$, we then have that the running time is $\mathcal{O}^{*}\binom{n+\frac{2 n}{k-2}-1}{n}$, which is $\mathcal{O}^{*}\left(2^{\mathcal{O}}((n / k) \log k)\right.$, as claimed. Hence, it only remains to prove that the highest rank is at most $2 n /(k-2)$.

For this, observe that in an $\ell$-good configuration, for each $0 \leq i<\ell$, the union of all vertices having level $i$ and the set of $(\lfloor k / 2\rfloor+1)$ terminals together forms a cutset. Since the graph is $k$-connected, this means that for each $0 \leq i<\ell$, the number of vertices having level $i$ is at least $k / 2-1$. The required bound on the rank easily follows.

## 4 Upper Bounds for Spanning Tree Congestion

We first state the following easy lemma, which together with Proposition 4, implies Lemma 5.

- Lemma 9. In a graph $G=(V, E)$, let $t_{1}$ be a vertex, and let $t_{2}, \cdots, t_{\ell}$ be any $(\ell-1)$ neighbours of $t_{1}$. Suppose that there exists a $\ell$-connected-partition $\cup_{j=1}^{\ell} V_{\ell}$ such that for all $j \in \ell, t_{j} \in V_{j}$, and the sum of degree of vertices in each $V_{j}$ is at most $D$. Let $\tau_{j}$ be an arbitrary spanning tree of $G\left[V_{j}\right]$. Let $e_{j}$ denote the edge $\left\{t_{1}, t_{j}\right\}$. Let $\tau$ be the spanning tree of $G$ defined as $\tau:=\left(\cup_{j=1}^{\ell} \tau_{j}\right) \cup\left(\cup_{j=2}^{\ell} e_{j}\right)$. Then $\tau$ has congestion at most $D$.
- Theorem 10. For any connected graph $G=(V, E)$, there is an algorithm which computes a spanning tree with congestion at most $8 \sqrt{m n}$ in time $\mathcal{O}^{*}\left(2^{\mathcal{O}(n \log n / \sqrt{m / n})}\right)$.
- Theorem 11. For any connected graph $G=(V, E)$, there is a polynomial time algorithm which computes a spanning tree with congestion at most $16 \sqrt{m n} \cdot \log n$.

The two algorithms follow the same framework, depicted in Algorithm 1. It is a recursive algorithm; the parameter $\hat{m}$ is a global parameter, which is the number of edges in the input graph $G$ in the first level of the recursion; let $\hat{n}$ denote the number of vertices in this graph.

The only difference between the two algorithms is in Line 15 on how this step is executed, with trade-off between the running time of the step $T\left(\hat{m}, n_{H}, m_{H}\right)$, and the guarantee
$D\left(\hat{m}, n_{H}, m_{H}\right)$. For proving Theorem 10, we use Theorem 3(b), Proposition 4 and Lemma 9, yielding $D\left(\hat{m}, n_{H}, m_{H}\right) \leq 8 m_{H} \sqrt{n_{H} / \hat{m}}$ and $T\left(\hat{m}, n_{H}, m_{H}\right)=\mathcal{O}^{*}\left(2^{\mathcal{O}\left(n_{H} \log n_{H} / \sqrt{\hat{m} / n_{H}}\right)}\right)$. For proving Theorem 11, we make use of an algorithm in Chen et al. [14], which yields $D\left(\hat{m}, n_{H}, m_{H}\right) \leq 16 m_{H} \sqrt{n_{H} / \hat{m}} \cdot \log n_{H}$ and $T\left(\hat{m}, n_{H}, m_{H}\right)=\operatorname{poly}\left(n_{H}, m_{H}\right)$.

```
Algorithm 1: FindLCST( \(H, \hat{m}\) )
    Input : A connected graph \(H=\left(V_{H}, E_{H}\right)\) on \(n_{H}\) vertices and \(m_{H}\) edges
    Output: A spanning tree \(\tau\) of \(H\)
    if \(m_{H} \leq 8 \sqrt{\hat{m} n_{H}}\) then
        return an arbitrary spanning tree of \(H\)
    end
    \(k \leftarrow\left\lceil\sqrt{\hat{m} / n_{H}}\right\rceil\)
    \(Y \leftarrow\) a global minimum vertex cut of \(H\)
    if \(|Y|<k\) then
        \(X \leftarrow\) the smallest connected component in \(H\left[V_{H} \backslash Y\right]\) (See Figure 3)
        \(Z \leftarrow V_{H} \backslash(X \cup Y)\)
        \(\tau_{1} \leftarrow \operatorname{FindLCST}(H[X], \hat{m})\)
        \(\tau_{2} \leftarrow \operatorname{FindLCST}(H[Y \cup Z], \hat{m}) ; \quad(H[Y \cup Z]\) is connected as \(Y\) is a global min cut \()\)
        return \(\tau_{1} \cup \tau_{2} \cup\) (an arbitrary edge between \(X\) and \(Y\) )
    else
        \(t_{1} \leftarrow\) an arbitrary vertex in \(V_{H}\)
        Pick \(\lfloor k / 2\rfloor\) neighbours of \(t_{1}\) in the graph \(H\); denote them by \(t_{2}, t_{3}, \cdots, t_{\lfloor k / 2\rfloor+1}\).
        Let \(e_{j}\) denote edge \(t_{1} t_{j}\) for \(2 \leq j \leq\lfloor k / 2\rfloor+1\). (See Figure 4)
        Compute a \((\lfloor k / 2\rfloor+1)\)-connected-partition of \(H\), denoted by \(\cup_{j=1}^{\lfloor k / 2\rfloor+1} V_{j}\), such that
        for each \(j \in[\lfloor k / 2\rfloor+1], t_{j} \in V_{j}\), and the total degree (w.r.t. graph \(H\) ) of vertices in
        each \(V_{j}\) is at most \(D\left(\hat{m}, n_{H}, m_{H}\right)\). Let the time needed be \(T\left(\hat{m}, n_{H}, m_{H}\right)\).
        For each \(j \in[\lfloor k / 2\rfloor+1], \tau_{j} \leftarrow\) an arbitrary spanning tree of \(G\left[V_{j}\right]\)
        return \(\left(\cup_{j=1}^{\lfloor k / 2\rfloor+1} \tau_{j}\right) \cup\left(\cup_{j=2}^{\lfloor k / 2\rfloor+1} e_{j}\right)\)
    end
```



Figure 3 The scenario in Algorithm 1 when the graph has low connectivity. The vertex set $Y$ is a global minimum vertex cut of the graph. The vertex set $X$ is the smallest connected component after the removal of $Y$, and $Z$ is the union of all the other connected components.


Figure 4 The scenario in Algorithm 1 when the graph has high connectivity.

In the rest of this section, we first discuss the algorithm in Chen et al., then we prove Theorem 11. The proof of Theorem 10 is almost identical, and is deferred to Appendix A.2.

Single-Commodity Confluent Flow and The Algorithm of Chen et al. In a singlecommodity confluent flow problem, the input includes a graph $G=(V, E)$, a demand function $w: V \rightarrow \mathbb{R}^{+}$and $\ell$ sinks $t_{1}, \cdots, t_{\ell} \in V$. For each $v \in V$, a flow of amount $w(v)$ is routed from $v$ to one of the sinks. But there is a restriction: at every vertex $u \in V$, the outgoing flow must leave $u$ on at most 1 edge, i.e., the outgoing flow from $u$ is unsplittable. The problem is to seek a flow satisfying the demands which minimizes the node congestion, i.e., the maximum incoming flow among all vertices. Since the incoming flow is maximum at one of the sinks, it is equivalent to minimize the maximum flow received among all sinks. (Here, we assume that no flow entering a sink will leave.)

Single-commodity splittable flow problem is almost identical to single-commodity confluent flow problem, except that the above restriction is dropped, i.e., now the outgoing flow at $u$ can split along multiple edges. Note that here, the maximum incoming flow might not be at a sink. It is known that single-commodity splittable flow can be solved in polynomial time. For brevity, we drop the phrase "single-commodity" from now on.

- Theorem 12 ([14, Section 4]). Suppose that given graph $G$, demand $w$ and $\ell$ sinks, there is a splittable flow with node congestion $q$. Then there exists a polynomial time algorithm which computes a confluent flow with node congestion at most $(1+\ln \ell) q$ for the same input.
- Corollary 13. Let $G$ be a $k$-connected graph with $m$ edges. Then for any $\ell \leq k$ and for any $\ell$ vertices $t_{1}, \cdots, t_{\ell} \in V$, there exists a polynomial time algorithm which computes an $\ell$-connected-partition $\cup_{j=1}^{\ell} V_{\ell}$ such that for all $j \in \ell, t_{j} \in V_{j}$, and the total degrees of vertices in each $V_{j}$ is at most $4(1+\ln \ell) m / \ell$.

Corollary 13 follows from Theorem 12 and Proposition 4. See Appendix A. 1 for details.
Congestion Analysis. We view the whole recursion process as a recursion tree. There is no endless loop, since down every path in the recursion tree, the number of vertices in the input graphs are strictly decreasing. On the other hand, note that the leaf of the recursion tree is resulted by either (i) when the input graph $H$ to that call satisfies $m_{H} \leq 8 \sqrt{\hat{m} n_{H}}$, or (ii) when Lines 13-17 are executed. An internal node appears only when the vertex-connectivity of the input graph $H$ is low, and it makes two recursion calls.

We prove the following statement by induction from bottom-up: for each graph which is the input to some call in the recursion tree, the returned spanning tree of that call has congestion at most $16 \sqrt{\hat{m} n_{H}} \log n_{H}$.

We first handle the two basis cases (i) and (ii). In case (i), FindLCST returns an arbitrary spanning tree, and the congestion is bounded by $m_{H} \leq 8 \sqrt{\hat{m} n_{H}}$. In case (ii), by Corollary 13 and Lemma 9, FindLCST returns a tree with congestion at most $16 m_{H} \sqrt{n_{H} / \hat{m}} \cdot \log n_{H} \leq$ $16 \sqrt{\hat{m} n_{H}} \cdot \log n_{H}$.

Next, let $H$ be the input graph to a call which is represented by an internal node of the recursion tree. Recall the definitions of $X, Y, Z, \tau_{1}, \tau_{2}$ in the algorithm.

Let $|X|=x$. Note that $1 \leq x \leq n_{H} / 2$. Then by induction hypothesis, the congestion of the returned spanning tree is at most

$$
\begin{align*}
& \quad \max \left\{\text { congestion of } \tau_{1} \text { in } H[X], \text { congestion of } \tau_{2} \text { in } H[Y \cup Z]\right\}+|X| \cdot|Y| \\
& \leq 16 \sqrt{\hat{m}\left(n_{H}-x\right)} \log \left(n_{H}-x\right)+\left(\sqrt{\hat{m} / n_{H}}+1\right) \cdot x \tag{1}
\end{align*}
$$

Viewing $x$ as a real variable, by taking derivative, it is easy to see that the above expression is maximized at $x=1$. Thus the congestion is at most
$16 \sqrt{\hat{m}\left(n_{H}-1\right)} \log \left(n_{H}-1\right)+\sqrt{\hat{m} / n_{H}}+1 \leq 16 \sqrt{\hat{m} n_{H}} \log n_{H}$, as desired by Theorem 11.
Runtime Analysis. At every internal node of the recursion tree, the algorithm makes two recursive calls with two vertex-disjoint and strictly smaller (w.r.t. vertex size) inputs. The dominating knitting cost is in Line 5 for computing a global minimum vertex cut, which is well-known that it can be done in polynomial time. Since at every leaf of the recursion tree the running time is polynomial, by standard analysis on divide-and-conquer algorithms, the running time of the whole algorithm is polynomial, which completes the proof of Theorem 11.

## 5 Lower Bound for Spanning Tree Congestion

Here, we give a lower bound on spanning tree congestion which matches our upper bound.

- Theorem 14. For any sufficiently large $n$, and for any $m$ satisfying $n^{2} / 2 \geq m \geq$ $\max \{16 n \log n, 100 n\}$, there exists a connected graph with $N=(3-o(1)) n$ vertices and $M \in[m, 7 m]$ edges, for which the spanning tree congestion is at least $\Omega(\sqrt{m n})$.

We start with the following lemma, which states that for a random graph $\mathcal{G}(n, p)$, when $p$ is sufficiently large, its edge expansion is $\Theta(n p)$ with high probability. The proof of the lemma uses only fairly standard arguments and is deferred to Appendix A.3.

- Lemma 15. For any integer $n \geq 4$ and $1 \geq p \geq 32 \cdot \frac{\log n}{n}$, let $\mathcal{G}(n, p)$ denote the random graph with $n$ vertices, in which each edge occurs independently with probability $p$. Then with probability at least $1-\mathcal{O}(1 / n)$, (i) the random graph is connected, (ii) the number of edges in the random graph is between $p n^{2} / 4$ and $p n^{2}$, and (iii) for each subset of vertices $S$ with $|S| \leq n / 2$, the number of edges leaving $S$ is at least $\frac{p}{2} \cdot|S| \cdot(n-|S|)$.

In particular, for any sufficiently large integer $n$, when $n^{2} / 2 \geq m \geq 16 n \log n$, by setting $p=2 m / n^{2}$, there exists a connected graph with $n$ vertices and $[m / 2,2 m]$ edges, such that for each subset of vertices $S$ with $|S| \leq n / 2$, the number of edges leaving $S$ is at least $\frac{m}{2 n} \cdot|S|=\Theta(m / n) \cdot|S|$. We denote such a graph by $H(n, m)$.


Figure 5 Our lower-bound construction for spanning tree congestion. $V_{1}, V_{2}, V_{3}$ are three vertex subsets of the same size. In each of the subsets, we embed expander $H(n, m)$. There is a small overlap between $V_{2}$ and $V_{1}, V_{3}$, while $V_{1}, V_{3}$ are disjoint. For any vertex $v_{1} \in V_{1} \cap V_{2}$, we add edges between it and any other vertex in $V_{1} \cup V_{2}$; similarly, for any vertex $v_{3} \in V_{3} \cap V_{2}$, we add edges (not shown in figure) between it and any other vertex in $V_{3} \cup V_{2}$.

We discuss our construction here (see Figure 5) before delving into the proof. The vertex set $V$ is the union of three vertex subsets $V_{1}, V_{2}, V_{3}$, such that $\left|V_{1}\right|=\left|V_{2}\right|=\left|V_{3}\right|=n$, $\left|V_{1} \cap V_{2}\right|=\left|V_{2} \cap V_{3}\right|=\sqrt{m / n}$, and $V_{1}, V_{3}$ are disjoint. In each of $V_{1}, V_{2}$ and $V_{3}$, we embed $H(n, m)$. The edge sets are denoted $E_{1}, E_{2}, E_{3}$ respectively. Up to this point, the construction is similar to that of Ostrovskii [30], except that we use $H(n, m)$ instead of a complete graph.

The new component in our construction is adding the following edges. For each vertex $v \in V_{1} \cap V_{2}$, add an edge between $v$ and every vertex in $\left(V_{1} \cup V_{2}\right) \backslash\{v\}$. The set of these edges are denoted $F_{1}$. Similarly, for each vertex $v \in V_{3} \cap V_{2}$, add an edge between $v$ and every vertex in $\left(V_{3} \cup V_{2}\right) \backslash\{v\}$. The set of these edges are denoted $F_{3}$. This new component is crucial: without it, we could only prove a lower bound of $\Omega(m / \sqrt{n})=\Omega\left(\sqrt{m n} \cdot \frac{\sqrt{m}}{n}\right)$.

Proof of Theorem 14: Let $G=(V, E)$ be the graph constructed as above. The whole graph has $3 n-2 \sqrt{m / n}$ vertices. The number of edges is at least $m$ (due to edges in $E_{1}$ and $E_{3}$ ), and is at most $6 m+2 \sqrt{m / n} \cdot 2 n=6 m+4 \sqrt{m n}$, which is at most $7 m$ for all sufficiently large $n$.

It is well known that for any tree on $n$ vertices, there exists a vertex $x$ called a centroid of the tree such that, removing $x$ decomposes the tree into connected components, each of size at most $n / 2$. Now, consider any spanning tree of the given graph, let $u$ be a centroid of the tree. Without loss of generality, we can assume that $u \notin V_{1}$; otherwise we swap the roles of $V_{1}$ and $V_{3}$. The removal of $u$ (and its adjacent edges) from the tree decomposes the tree into a number of connected components. For any of these components which intersects $V_{1}$, it must contain at least one vertex of $V_{1} \cap V_{2}$, thus the number of such components is at most $\sqrt{m / n}$, and hence there exists one of them, denoted by $U_{j}$, such that

$$
b_{1}:=\left|U_{j} \cap V_{1}\right| \geq n /(\sqrt{m / n})=n \sqrt{n / m}
$$

Let $e_{j}$ denote the tree-edge that connects $u$ to $U_{j}$. Then there are three cases:
Case 1: $n \sqrt{n / m} \leq b_{1} \leq n-n \sqrt{n / m}$. Due to the property of $H(n, m)$, the congestion of $e_{j}$ is at least $\Theta(m / n) \cdot \min \left\{b_{1}, n-b_{1}\right\} \geq \Theta(\sqrt{m n})$.

Case 2: $b_{1}>n-n \sqrt{n / m}$ and $\left|U_{j} \cap V_{1} \cap V_{2}\right| \leq \frac{1}{2} \cdot \sqrt{m / n}$. Let $W:=\left(V_{1} \cap V_{2}\right) \backslash U_{j}$. Note that by this case's assumption, $\left|W_{1}\right| \geq \frac{1}{2} \cdot \sqrt{m / n}$. Due to the edge subset $F_{1}$, the congestion of $e_{j}$ is at least

$$
\left|F_{1}\left(W, V_{1} \backslash W\right)\right| \geq\left(\frac{1}{2} \cdot \sqrt{m / n}\right) \cdot \frac{n}{2}=\Theta(\sqrt{m n})
$$

Case 3: $b_{1}>n-n \sqrt{n / m}$ and $\left|U_{j} \cap V_{1} \cap V_{2}\right|>\frac{1}{2} \cdot \sqrt{m / n}$. Let $W^{\prime}:=U_{j} \cap V_{1} \cap V_{2}$, and let $Z:=\left(V_{2} \backslash V_{1}\right) \cap U_{j}$.

Note that $b_{1}>n-n \sqrt{n / m} \geq 9 n / 10$. Suppose $|Z| \geq 6 n / 10$, then $\left|U_{j}\right|>9 n / 10+$ $6 n / 10>|V| / 2$, a contradiction to the assumption that $u$ is a centroid. Thus, $|Z|<6 n / 10$. Due to the edge subset $F_{2}$, the congestion of $e_{j}$ is at least

$$
\begin{aligned}
\left|F_{2}\left(W^{\prime} \cup Z, V_{2} \backslash\left(W^{\prime} \cup Z\right)\right)\right| & \geq\left|W^{\prime}\right| \cdot\left(n-\left|W^{\prime}\right|-|Z|\right) \\
& \geq\left(\frac{1}{2} \cdot \sqrt{m / n}\right) \cdot\left(n-\sqrt{m / n}-\frac{6 n}{10}\right)=\Theta(\sqrt{m n})
\end{aligned}
$$

## 6 Graphs with Expanding Properties

For any vertex subset $U, W \subset V$, let $N_{W}(U)$ denote the set of vertices in $W$ which are adjacent to a vertex in $U$. Let $N(U):=N_{V \backslash U}(U)$.

- Definition 16. A graph $G=(V, E)$ on $n$ vertices is an $\left(n, s, d_{1}, d_{2}, d_{3}, t\right)$-expanding graph if the following four conditions are satisfied:
(1) for each vertex subset $S$ with $|S|=s,|N(S)| \geq d_{1} n$;
(2) for each vertex subset $S$ with $|S| \leq s,|N(S)| \geq d_{2}|S|$;
(3) for each vertex subset $S$ with $|S| \leq n / 2$ and for any subset $S^{\prime} \subset S,\left|N_{V \backslash S}\left(S^{\prime}\right)\right| \geq\left|S^{\prime}\right|-t$.
(4) For each vertex subset $S,|E(S, V \backslash S)| \leq d_{3}|S|$.
- Theorem 17. For any connected graph $G$ which is an ( $\left.n, s, d_{1}, d_{2}, d_{3}, t\right)$-expanding graph, there is a polynomial time algorithm which computes a spanning tree with congestion at most

$$
d_{3} \cdot\left[4 \cdot \max \left\{s+1,\left\lceil\frac{3 d_{1} n}{d_{2}}\right\rceil\right\} \cdot\left(\frac{1}{2 d_{1}}\right)^{\log _{(2-\delta)} 2}+t\right], \quad \text { where } \delta=\frac{t}{d_{1} n}
$$

Next, we present the polynomial time algorithm in Theorem 17 and its analysis.
Algorithm. Let $G$ be an $\left(n, s, d_{1}, d_{2}, d_{3}, t\right)$-expanding graph. By Condition (2), every vertex has degree at least $d_{2}$. Let $v_{0}$ be a vertex of degree $d \geq d_{2}$, and let $v_{1}, \cdots, v_{d}$ be its $d$ neighbours. We maintain a tree $T$ rooted at $v_{0}$ such that $T=T_{1} \cup T_{2} \cup \cdots \cup T_{d} \cup$ $\left\{v_{0} v_{1}, v_{0} v_{2}, \ldots, v_{0} v_{d}\right\}$ where $T_{1}, T_{2}, \cdots, T_{d}$ are trees rooted at $v_{1}, v_{2}, \ldots, v_{d}$ respectively. We call the $T_{i}^{\prime} s$ as branches. (See Figure 6). We start with each branch $T_{i}=v_{i}$. In order to minimize congestion, we grow $T$ in a balanced way, i.e., we maintain that the $T_{i}$ 's are roughly of the same size. A branch is saturated if it contains at least max $\left\{s+1, \frac{3 d_{1} n}{d_{2}}\right\}$ vertices.

At any point of time, let $V_{T}$ be the set of vertices in $T$ and $\overline{V_{T}}$ be the vertices not in $T$. Often, we will move a subtree of a saturated branch $T_{i}$ to an unsaturated branch $T_{j}$ to ensure balance. For any $x \in V_{T}$, let $T_{x}$ denote the subtree of $T$ rooted at $x$. A vertex $x$ of a


Figure 6 The tree $T$ and its branches


Figure 7 Transfer of a subtree from a saturated branch to an unsaturated branch
saturated branch $T_{i}$ is called transferable (to branch $T_{j}$ ) if $x$ has a neighbour $y$ in $T_{j}$ and the tree $T_{j} \cup\{x y\} \cup T_{x}$ is unsaturated. (See Figure 7.)

The algorithm is divided into two phases which are described below. Throughout the algorithm, whenever a branch $T_{i}$ gets modified, $T$ gets modified accordingly, and whenever $T$ gets modified $V_{T}$ and $\overline{V_{T}}$ gets modified accordingly.

Phase 1: Repeatedly do one of the following two actions, until $\left|V_{T}\right| \geq d_{1} n$ :
(We will prove that the precondition of at least one of the actions is satisfied if $\left|V_{T}\right|<d_{1} n$ )

1. If there exists a $b \in \overline{V_{T}}$ such that $b$ has a neighbour $a$ in some unsaturated branch $T_{i}$ : Add the vertex $b$ and the edge $a b$ to branch $T_{i}$.
2. If there exists at least one transferable vertex: (see Figure 7)

Find the transferable vertex $x$ such that $T_{x}$ is the smallest. Let $T_{i^{*}}$ be the branch currently containing $x, T_{j}$ be a branch to which it is transferable, and $y$ be an arbitrarily chosen neighbour of $x$ in $T_{j}$.
a. Remove the subtree $T_{x}$ from $T_{i^{*}}$ and add it to $T_{j}$ with $x$ as a child of $y$.
b. Pick a $b \in \overline{V_{T}}$ that has a neighbour $a$ (arbitrarily chosen, if many) either in $T_{i^{*}}$ or in $T_{j}$. (We will show in the analysis that such $b$ exists). We add vertex $b$ and edge $a b$ to the branch containing $a$ (i.e. to $T_{i^{*}}$ or $T_{j}$ ).

Phase 2: While $\overline{V_{T}} \neq \emptyset$, repeat:
Find a maximum matching of $G\left[V_{T}, \overline{V_{T}}\right]$, the bipartite graph formed by edges of $G$ between $V_{T}$ and $\overline{V_{T}}$. Let $M$ be the matching. Add all edges of $M$ to $T$.

In the analysis below, we say that a tree is saturated if it contains at least $\mathcal{A}$ vertices; we will determine its appropriate value by the end of the analysis.

Analysis of Phase 1. We claim that during Phase 1, i.e. if $\left|V_{T}\right|<d_{1} n$, the precondition of either step 1 or step 2 is satisfied. We also show the existence of a vertex $b$ as specified in step 2 b , whenever step 2 b is reached. Given these and the fact that a vertex in $\overline{V_{T}}$ is moved to $V_{T}$ (either in step 1 or in step 2 b ) during each round of Phase 1, we have that Phase 1 runs correctly and terminates after a linear number of rounds.

During Phase 1, we will also maintain the invariant that each branch has at most $\mathcal{A}$ vertices; thus, each saturated branch has exactly $\mathcal{A}$ vertices. We call this invariant the balancedness. Note that balancedness is not violated due to step 1, as the new vertex is added to an unsaturated branch. It is not violated during step 2 as the branches $T_{i^{*}}$ and $T_{j}$ (as defined in step 2) become unsaturated at the end of the step.

We define the hidden vertices of $T$ (denoted by $H \equiv H_{T}$ ) as follows: they are the vertices which are not adjacent to any vertices outside the tree, i.e., to any vertex in $\overline{V_{T}}$. If there is an unsaturated branch with a non-hidden vertex, clearly the precondition of step 1 is satisfied. So, let us assume that all the vertices in all unsaturated branches are hidden. In such a case, we show that the precondition of step 2 is satisfied if $\left|V_{T}\right|<d_{1} n$.

We argue that in this case $|H| \leq s$ : otherwise, take a subset $H^{\prime} \subset H$ of cardinality $s$, then by condition $(2), N\left(H^{\prime}\right)$, which is contained in $V_{T}$, has cardinality at least $d_{1} n$, a contradiction.

Since $\left|V_{T}\right|<d_{1} n$, the number of saturated branches is at most $d_{1} n / \mathcal{A}$. To ensure that at least one unsaturated branch exists, we set $\mathcal{A}$ such that $d_{1} n / \mathcal{A}<d_{2}$. Let $U$ denote the set of vertices in all unsaturated branches. Since all vertices in $U$ are hidden vertices, $|U| \leq s$. Then by condition (2), $|N(U)| \geq d_{2}|U|$. Note that the vertices in $N(U)$ are all in the saturated branches. By the pigeon-hole principle, there exists a saturated branch containing at least

$$
N(U) /\left(d_{1} n / \mathcal{A}\right) \geq \frac{\mathcal{A} d_{2}|U|}{d_{1} n}
$$

vertices of $N(U)$. By setting $\mathcal{A} \geq \frac{3 d_{1} n}{d_{2}}$, the above calculation guarantees the existence of a saturated branch containing at least $3|U| \geq|U|+2$ vertices of $N(U)$; let $T_{i}$ be such a branch.

In $T_{i}$, pick a vertex $x \in T_{i} \cap N(U)$ such that $T_{x}$ does not contain any vertex in $N(U)$, except $x$. Then the size of $T_{x}$ is at most $\mathcal{A}-\left|N(U) \cap T_{i}\right|+1 \geq \mathcal{A}-(|U|+1)$. Let $y \in U$ be a vertex which is adjacent to $x$ and $T_{j}$ be the branch containing $y$. Since $T_{j}$ has at most $|U|$ vertices, $x$ is a transferable vertex (to $T_{j}$ ). Thus precondition of step 2 is satisfied.

We further set $\mathcal{A}>s$ so that in each saturated branch, there is at least one unhidden vertex. In particular, $T_{i}$ has an unhidden vertex, which is adjacent to some $b \in \overline{V_{T}}$. The vertex $b$ is either adjacent to a vertex in $T_{x}$, or a vertex in $T_{i} \backslash T_{x}$ as required in step 2b.

Analysis of Phase 2. Since $G$ is connected, $M$ is non-empty in each iteration of Phase 2, and hence Phase 2 terminates in linear number of rounds. At the end of Phase 2, since $\overline{V_{T}}$ is empty, $T$ is clearly a spanning tree. It only remains to estimate the congestion of this spanning tree. Towards this, we state the following modified Hall's theorem, which is an easy corollary of the standard Hall's theorem.

- Lemma 18. In a bipartite graph $(L, R)$ with $|L| \leq|R|$, for any vertex $w \in L$, let $R(w)$ denote the neighbours of $w$ in $R$; then for any $W \subset L$, let $R(W):=\cup_{w \in W} R(w)$. Suppose that there exist $t \geq 0$ such that for any $W \subset L$, we have $|R(W)| \geq|W|-t$. Then the bipartite graph admits a matching of size at least $|L|-t$.

Recall that Phase 2 consists of multiple rounds of finding a matching between $V_{T}$ and $\overline{V_{T}}$. As long as $\left|V_{T}\right| \leq n / 2$, condition (3) (with $S=V_{T}$ ) plus the modified Hall's theorem (with $L=V_{T}$ and $R=\overline{V_{T}}$ ) guarantees that in each round, at least

$$
\left|V_{T}\right|-t \geq\left(1-\frac{t}{d_{1} n}\right) \cdot\left|V_{T}\right|=:(1-\delta)\left|V_{T}\right|
$$

number of vertices in $\overline{V_{T}}$ are matched. Thus, after at most $\left\lceil\log _{(2-\delta)} \frac{1}{2 d_{1}}\right\rceil$ rounds of matching, $\left|V_{T}\right| \geq n / 2$. After reaching $\left|V_{T}\right| \geq n / 2$, condition (3) (with $S=\overline{V_{T}}$ ) plus the modified Hall's theorem (with $L=\overline{V_{T}}$ and $R=V_{T}$ ) guarantees that after one more round of matching, all but $t$ vertices are left in $\overline{V_{T}}$.

By the end of Phase 1, each branch had at most $\mathcal{A}$ vertices. After each round of matching, the cardinality of each branch is doubled at most. Thus, the maximum possible number of vertices in each branch after running the whole algorithm is at most

$$
\mathcal{A} \cdot 2^{\left\lceil\log _{(2-\delta)} \frac{1}{2 d_{1}}\right\rceil+1}+t \leq 4 \mathcal{A} \cdot\left(\frac{1}{2 d_{1}}\right)^{\log _{(2-\delta)} 2}+t
$$

and hence the STC is at most

$$
d_{3} \cdot\left[4 \mathcal{A} \cdot\left(\frac{1}{2 d_{1}}\right)^{\log _{(2-\delta)} 2}+t\right]
$$

Recall that we need $\mathcal{A}$ to satisfy $d_{1} n / \mathcal{A}<d_{2}, \mathcal{A} \geq \frac{3 d_{1} n}{d_{2}}$ and $\mathcal{A}>s$. Thus we set $\mathcal{A}:=\max \left\{s+1,\left\lceil\frac{3 d_{1} n}{d_{2}}\right\rceil\right\}$.

### 6.1 Random Graph

Let $G \in \mathcal{G}(n, p)$ where $p \geq c_{0} \log n / n$, and $c_{0}=64$. The following lemmas show that with high probability $G$ is an $\left(n, s, d_{1}, d_{2}, d_{3}, t\right)$-expanding graph with $s=\Theta(1 / p), d_{1}=\Theta(1)$, $d_{2}=\Theta(n p), d_{3}=\Theta(n p), t=\Theta(1 / p)$ (and hence $\left.\delta=o(1)\right)$. The proof of the lemmas are deferred to Appendix B.2.

- Lemma 19. For any $S \subseteq V(G)$ such that $|S|=\lceil 1 / p\rceil$, we have $|N(S)| \geq c_{2} n$ with probability at least $1-e^{-n / 16}$, where $c_{2}=1 / 25$.
- Lemma 20. For any $S \subseteq V(G)$ such that $|S| \leq 1 / p$, we have $|N(S)| \geq c_{3} n p|S|$ with probability at least $1-\mathcal{O}\left(1 / n^{2}\right)$, where $c_{3}=1 / 16$.
- Lemma 21. For all $A \subseteq V(G)$ such that $|A| \leq n / 2$, and for all $S \subseteq A$, with probability at least $1-e^{-n}, S$ has at least $|S|-c_{4} / p$ neighbors in $V(G) \backslash A$, where $c_{4}=12$.
- Lemma 22. For all $S \subseteq V(G)$, the cut size $|E(S, V(G) \backslash S)|$ is at most $n p|S|$ with probability at least $1-n^{-c_{0} / 4}$.

Plugging the bounds from above lemmas into Theorem 17, together with a separate lower bound argument (Theorem 25 in Appendix B.1), we have the following theorem; in Appendix B.1, we also present a non-algorithmic proof of this theorem.

- Theorem 23. If $G \in \mathcal{G}(n, p)$ where $p \geq 64 \log n / n$, then with probability at least $1-\mathcal{O}(1 / n)$, its STC is $\Theta(n)$.


## 7 Discussion and Open Problems

In this paper, we provide thorough understanding, both combinatorially and algorithmically, on the spanning tree congestion of general graphs and random graphs. On course of doing so, we also provide the first constructive proof for the generalized Győri-Lovász theorem, which might be of independent interest. Following are some natural open problems:

- Finding the spanning tree with minimum congestion is NP-hard; indeed, Bodlaender et al. [9] showed a $(9 / 8-\epsilon)$-approximation NP-hardness for the STC problem. Does a constant or a poly-logarithmic factor approximation polynomial time algorithm exist?
- We present an algorithm for computing a spanning tree achieving congestion at most $\mathcal{O}(\sqrt{m n})$. The algorithm runs in sub-exponential time when $m=\omega\left(n \log ^{2} n\right)$. Is there a polynomial time algorithm for constructing such a spanning tree?
- For a $k$-connected graph, a connected $k$-partition where all parts are of size at most $\mathcal{O}((n / k) \log k)$ can be found in polynomial time due to an algorithm of Chen et al. [14]. Can we improve the sizes of parts to $\mathcal{O}(n / k)$ ?
- Is finding Győri-Lovász partition PLS-complete? If not, is it polynomial time solvable?


## References

1 Ittai Abraham, Yair Bartal, and Ofer Neiman. Nearly tight low stretch spanning trees. In FOCS 2008, pages 781-790, 2008.
2 Ittai Abraham and Ofer Neiman. Using petal-decompositions to build a low stretch spanning tree. In STOC 2012, pages 395-406, 2012.
3 Noga Alon, Richard M. Karp, David Peleg, and Douglas B. West. A graph-theoretic game and its application to the k-server problem. SIAM J. Comput., 24(1):78-100, 1995.
4 Ingo Althöfer, Gautam Das, David P. Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. Discrete \& Computational Geometry, 9:81-100, 1993.
5 Reid Andersen and Uriel Feige. Interchanging distance and capacity in probabilistic mappings. CoRR, abs/0907.3631, 2009.
6 Sanjeev Arora, Satish Rao, and Umesh V. Vazirani. Expander flows, geometric embeddings and graph partitioning. J. ACM, 56(2):5:1-5:37, 2009.
7 Nikhil Bansal, Uriel Feige, Robert Krauthgamer, Konstantin Makarychev, Viswanath Nagarajan, Joseph Naor, and Roy Schwartz. Min-max graph partitioning and small set expansion. SIAM J. Comput., 43(2):872-904, 2014.
8 Sandeep N. Bhatt, Fan R. K. Chung, Frank Thomson Leighton, and Arnold L. Rosenberg. Optimal simulations of tree machines (preliminary version). In FOCS 1986, pages 274-282, 1986.

9 Hans L. Bodlaender, Fedor V. Fomin, Petr A. Golovach, Yota Otachi, and Erik Jan van Leeuwen. Parameterized complexity of the spanning tree congestion problem. Algorithmica, 64(1):85-111, 2012.
10 Hans L. Bodlaender, Kyohei Kozawa, Takayoshi Matsushima, and Yota Otachi. Spanning tree congestion of k-outerplanar graphs. Discrete Mathematics, 311(12):1040-1045, 2011.
11 Béla Bollobás. Random Graphs. Cambridge University Press, 2001.
12 Béla Bollobás and Andrew Thomason. Random graphs of small order. North-Holland Mathematics Studies, 118:47-97, 1985.

13 Leizhen Cai and Derek G. Corneil. Tree spanners. SIAM J. Discrete Math., 8(3):359-387, 1995.

14 Jiangzhuo Chen, Robert D. Kleinberg, László Lovász, Rajmohan Rajaraman, Ravi Sundaram, and Adrian Vetta. (almost) tight bounds and existence theorems for singlecommodity confluent flows. J. ACM, 54(4):16, 2007.
15 Michael Elkin, Yuval Emek, Daniel A. Spielman, and Shang-Hua Teng. Lower-stretch spanning trees. SIAM J. Comput., 38(2):608-628, 2008.
16 Uriel Feige and Robert Krauthgamer. A polylogarithmic approximation of the minimum bisection. SIAM J. Comput., 31(4):1090-1118, 2002.
17 E. Győri. On division of graphs to connected subgraphs. Colloq. Math. Soc. Janos Bolyai, 18:485-494, 1976.
18 Ludovic Hofer and Thibaud Lambert. Study of the article: An $O\left(k^{2} n^{2}\right)$ algorithm to find a $k$-partition in a $k$-connected graph. 2014.
19 Alexander Hoyer and Robin Thomas. The Győri-Lovász theorem. arXiv, abs/1605.01474, 2016. URL: http://arxiv.org/abs/1706.08115.

20 David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? J. Comput. Syst. Sci., 37(1):79-100, 1988.
21 Marcos A. Kiwi, Daniel A. Spielman, and Shang-Hua Teng. Min-max-boundary domain decomposition. Theor. Comput. Sci., 261(2):253-266, 2001.
22 Ioannis Koutis, Gary L. Miller, and Richard Peng. A nearly $O(m \log n)$ time solver for SDD linear systems. In FOCS 2011, pages 590-598, 2011.
23 Kyohei Kozawa and Yota Otachi. Spanning tree congestion of rook's graphs. Discussiones Mathematicae Graph Theory, 31(4):753-761, 2011.
24 Kyohei Kozawa, Yota Otachi, and Koichi Yamazaki. On spanning tree congestion of graphs. Discrete Mathematics, 309(13):4215-4224, 2009.
25 Hiu Fai Law, Siu Lam Leung, and Mikhail I. Ostrovskii. Spanning tree congestions of planar graphs. Involve, 7(2):205-226, 2014.
26 László Lovász. A homology theory for spanning trees of a graph. Acta Math. Acad. Sci. Hungaricae, 30(3-4):241-251, 1977.
27 Christian Löwenstein, Dieter Rautenbach, and Friedrich Regen. On spanning tree congestion. Discrete Math., 309(13):4653-4655, 2009.
28 Shin-Ichi Nakano, Md. Saidur Rahman, and Takao Nishizeki. A linear-time algorithm for four-partitioning four-connected planar graphs. Inf. Process. Lett., 62(6):315-322, 1997.
29 Yoshio Okamoto, Yota Otachi, Ryuhei Uehara, and Takeaki Uno. Hardness results and an exact exponential algorithm for the spanning tree congestion problem. J. Graph Algorithms Appl., 15(6):727-751, 2011.
30 M. I. Ostrovskii. Minimal congestion trees. Discrete Math., 285:219-226, 2004.
31 M. I. Ostrovskii. Minimum congestion spanning trees in planar graphs. Discrete Math., 310:1204-1209, 2010.
32 M. I. Ostrovskii. Minimum congestion spanning trees in bipartite and random graphs. Acta Mathematica Scientia, 31(2):634-640, 2011.
33 Harald Räcke. Optimal hierarchical decompositions for congestion minimization in networks. In STOC, pages 255-264, 2008.
34 André Raspaud, Ondrej Sýkora, and Imrich Vrto. Congestion and dilation, similarities and differences: A survey. In SIROCCO 2000, pages 269-280, 2000.
35 Satu Elisa Schaeffer. Graph clustering. Computer Science Review, 1(1):27-64, 2007.
36 Shai Simonson. A variation on the min cut linear arrangement problem. Mathematical Systems Theory, 20(4):235-252, 1987.

37 Daniel A. Spielman and Shang-Hua Teng. A local clustering algorithm for massive graphs and its application to nearly linear time graph partitioning. SIAM J. Comput., 42(1):1-26, 2013.

38 David Steurer. Tight bounds for the min-max boundary decomposition cost of weighted graphs. In SPAA 2006, pages 197-206, 2006.
39 Hitoshi Suzuki, Naomi Takahashi, and Takao Nishizeki. A linear algorithm for bipartition of biconnected graphs. Inf. Process. Lett., 33(5):227-231, 1990.
40 Zoya Svitkina and Éva Tardos. Min-max multiway cut. In APPROX-RANDOM 2004, pages 207-218, 2004.
41 Shang-Hua Teng. Scalable algorithms for data and network analysis. Foundations and Trends in Theoretical Computer Science, 12(1-2):1-274, 2016.
42 Koichi Wada and Kimio Kawaguchi. Efficient algorithms for tripartitioning triconnected graphs and 3-edge-connected graphs. In Graph-Theoretic Concepts in Computer Science, 19th International Workshop, WG '93, Utrecht, The Netherlands, June 16-18, 1993, Proceedings, pages 132-143, 1993.

## A Missing Proofs in Sections 4 and 5

## A. 1 Proof of Corollary 13

First of all, we set the demand of each vertex in the flow problem to be the the degree of the vertex in $G$, and $t_{1}, \cdots, t_{\ell}$ as the sinks in the flow problem.

By Proposition 4, there exists an $\ell$-connected-partition $\cup_{j=1}^{\ell} U_{\ell}$ such that for all $j \in[\ell]$, $t_{j} \in U_{j}$, and the total degrees of vertices in each $U_{j}$ is at most $4 m / \ell$. With this, by routing the demand of a vertex in $U_{j}$ to $t_{j}$ via an arbitrary path in $G\left[U_{j}\right]$ only, we construct a splittable flow with node congestion at most $4 \mathrm{~m} / \ell$. By Theorem 12 , one can construct a confluent flow with node congestion at most $4(1+\ln \ell) m / \ell$ in polynomial time.

Obviously, in the confluent flow, all the flow originating from one vertex goes completely into one sink. Set $V_{j}$ to be the set of vertices such that the flows originating from these vertices go into $t_{j}$. It is then routine to check that $\cup_{j=1}^{\ell} V_{\ell}$ is our desired $\ell$-connected-partition.

## A. 2 Proof of Theorem 10

Instead of giving the full proof, we point out the differences from the proof of Theorem 11.
First, in handling the basis case (ii), by Theorem 3(b), Proposition 4 and Lemma 9, we have an improved upper bound on the congestion of the returned tree, which is $8 m_{H} / \sqrt{\hat{m} / n_{H}} \leq$ $8 \sqrt{\hat{m} n_{H}}$. Thus, (1) can be improved to

$$
8 \sqrt{\hat{m}\left(n_{H}-x\right)}+\sqrt{\frac{\hat{m}}{n_{H}}} \cdot x .
$$

Again, by viewing $x$ as a real variable and taking derivative, it is easy to see that the above expression is maximized at $x=1$. So the above bound is at most

$$
8 \sqrt{\hat{m}\left(n_{H}-1\right)}+\sqrt{\frac{\hat{m}}{n_{H}}} \leq 8 \sqrt{\hat{m} n_{H}}, \text { as desired. }
$$

Concerning the running time, it is clear that in the worst case, it is dominated by some calls to the algorithm in Theorem 3(b). Note that the number of such calls is at most $\hat{n}$, since each call to the algorithm is on a disjoint set of vertices.

There remains one concern, which is the connectedness of $H[Y \cup Z]$. Suppose the contrary that $H[Y \cup Z]$ is not connected. Let $C$ be one of its connected components, so that it contains the least number of vertices from $Y$. Then $C$ contains at most $\lfloor|Y| / 2\rfloor$ vertices from $Y$, i.e., $|C \cap Y|<|Y|$. Note that $C \cap Y$ is a vertex cut set of the graph $H$, thus contradicting that $Y$ is a global minimum vertex cut set.

## A. 3 Proof of Lemma 15

It is well known that the requirements (i) and (ii) are satisfied with probability $1-o(1 / n)$. [11] For each subset $S$ with $|S| \leq n / 2$, by the Chernoff bound,

$$
\mathbb{P}\left[|E(S, V \backslash S)| \leq \frac{p}{2} \cdot|S| \cdot(n-|S|)\right] \leq e^{-p|S|(n-|S|) / 8} \leq e^{-p n|S| / 16}
$$

Since $p \geq 32 \cdot \frac{\log n}{n}$, the above probability is at most $n^{-2|S|}$. Then by a union bound, the probability that (iii) is not satisfied is at most

$$
\sum_{s=1}^{\lfloor n / 2\rfloor}\binom{n}{s} \cdot n^{-2 s} \leq \sum_{s=1}^{\lfloor n / 2\rfloor} n^{s} \cdot n^{-2 s} \leq \sum_{s=1}^{\lfloor n / 2\rfloor} n^{-s} \leq \frac{2}{n} .
$$

## B Spanning Tree Congestion of Random Graphs

## B. 1 Non-Algorithmic Proof of Theorem 23

We first present a simple non-algorithmic proof that random graph has STC $\Theta(n)$ with high probability. Theorem 24 gives the upper bound and Theorem 25 gives the lower bound. The proof of Theorem 24 uses Lemma 5 and the fact that for random graphs, vertex-connectivity and minimum degree are equal with high probability. Theorem 24 does not give an efficient algorithm.

- Theorem 24. If $G \in \mathcal{G}(n, p)$ where $p \geq 8 \log n / n$, then the spanning tree congestion of $G$ is at most $16 n$ with probability at least $1-o(1 / n)$.

Proof. It is known that the threshold probability for a random graph being $k$-connected is same as the threshold probability for it having minimum degree at least $k$ [12]. Since $p \geq 8 \log n / n$, using Chernoff bound and taking union bound over all vertices gives that $G$ has minimum degree at least $n p / 2$ with probability at least $1-o(1 / n)$. Hence $G$ is $(n p / 2)$-connected with probability at least $1-o(1 / n)$. We also have that the number of edges in $G$ is at most $2 n^{2} p$ with probability at least $1-o(1 / n)$. Now, by using Lemma 5 , we have that with probability at least $1-o(1 / n)$, the spanning tree congestion is at most $16 n$.

- Theorem 25. If $G \in \mathcal{G}(n, p)$ where $p \geq 32 \log n / n$, then the spanning tree congestion of $G$ is $\Omega(n)$ with probability $1-\mathcal{O}(1 / n)$.

Proof. By using Chernoff Bounds and applying union bound, it is easy to show that with probability $1-o(1 / n)$, every vertex of $G$ has degree at most $c_{1} n p$ for a sufficiently large constant $c_{1}$. Also, by Lemma 15 , with probability $1-\mathcal{O}(1 / n)$, properties (i) and (iii) of that lemma holds. In the proof below, we conditioned on the above mentioned highly probable events.

Take a spanning tree $T$ of $G$ which gives the minimum congestion. Let $u$ be a centroid of the tree $T$, i.e., each connected component of $T \backslash\{u\}$ has at most $n / 2$ vertices. If there is a connected component with number of vertices at least $n / 4$, then define this connected component as $T^{\prime}$. Else, all connected components have at most $n / 4$ vertices. In this case, let $T^{\prime}$ be the forest formed by the union of a minimum number of connected components of $T \backslash\{u\}$ such that $\left|T^{\prime}\right| \geq n / 4$. It is easy to see that $\left|T^{\prime}\right| \leq n / 2$. Also, the number of edges in $T$ from $V\left(T^{\prime}\right)$ to $V(T) \backslash V\left(T^{\prime}\right)$ is at $\operatorname{most}^{\operatorname{deg}}{ }_{G}(u)$, which is at most $c_{1} n p$.

By property (iii) of Lemma 15 , the number of edges between $V\left(T^{\prime}\right)$ and $V(G) \backslash V\left(T^{\prime}\right)$ is $\Omega\left(n^{2} p\right)$. Each of these edges in $G$ between $V\left(T^{\prime}\right)$ and $V(G) \backslash V\left(T^{\prime}\right)$ have to contribute to the congestion of at least one of the edges in $T$ between $V\left(T^{\prime}\right)$ and $V(G) \backslash V\left(T^{\prime}\right)$. Now since $T^{\prime}$ sends at most $c_{1} n p$ tree edges to other parts of $T$, it follows that there exists one edge in $T$ with congestion at least $\Omega\left(n^{2} p\right) /\left(c_{1} n p\right)=\Omega(n)$, as claimed.

## B. 2 Random Graph Satisfies Expanding Properties

Constants. For easy reference, we list out the constants used.

$$
c_{0}=64, c_{2}=1 / 25, c_{3}=1 / 16, c_{4}=12
$$

Proof of Lemma 19: Let $\bar{S}=V(G) \backslash S$. The probability that a fixed vertex in $\bar{S}$ does not have edge to $S$ is at most $(1-p)^{|S|} \leq(1-p)^{1 / p} \leq e^{-1}$. Since $|\bar{S}| \geq n-2 / p \geq$ $n-2 n /\left(c_{0} \log n\right) \geq 31 n / 32$, the expected value of $|N(S)|$ is at least $(31 / 32) n\left(1-e^{-1}\right) \geq n / 2$. Hence, using Chernoff bound, the probability that $|N(S)|<c_{2} n=n / 25$ is at most $e^{-n / 8}$. Since the number of such $S$ is at most $n^{2 / p}=2^{2 n / c_{0}} \leq 2^{n / 32}$, we have the lemma by applying union bound.

Proof of Lemma 20: Let $\bar{S}=V(G) \backslash S$. Since $|S| \leq 1 / p \leq n / \log n$, we have $|\bar{S}| \geq n / 2$ for sufficiently large $n$. Divide $\bar{S}$ into groups of size $\lceil 1 /(p|S|)\rceil$. The probability that such a group does not have edge to $S$ is at most $(1-p)^{|S|(1 /(p|S|))} \leq 1 / e$. The expected number of groups having edge to $S$ is at least $(n p|S| / 2)(1-1 / e) \geq n p|S| / 4$. Thus, by Chernoff bound, the probability that $|N(S)| \leq n p|S| / 16$ is at most $e^{-n p|S| / 16} \leq 2^{-c_{0}|S| \log n / 16} \leq 2^{-4|S| \log n}$. The number of sets of size $|S|$ is at most $2^{|S| \log n}$. Hence, taking union bound over all $S$ with $|S| \leq 1 / p$, we get the required lemma.

Proof of Lemma 21: First, we prove that for all $C, D \subseteq V(G)$ such that $|C| \geq n / 4,|D| \geq$ $c_{4} / p$, and $C \cap D=\emptyset$, there exist at least one edge between $C$ and $D$ with high probability. The probability that there is no edge between such a fixed $C$ and $D$ is at most $(1-p)^{(n / 4)\left(c_{4} / p\right)} \leq$ $e^{-c_{4} n / 4}$. The number of pairs of such $C$ and $D$ is at most $2^{2 n}$. Hence, by taking union bound, the probability that for all $C$ and $D$, the claim holds is at least $1-e^{2 n-\left(c_{4} n / 4\right)} \geq 1-e^{-n}$.

Using the above claim, we prove that for all $S \subseteq A, S$ has at least $|S|-c_{4} / p$ neighbors in $\bar{A}:=V(G) \backslash A$ with high probability. Suppose there is an $S$ which violates the claim. Note that we can assume $|S| \geq c_{4} / p$, because otherwise the claim is vacuously true. Let $B:=\bar{A} \backslash N(S)$. There cannot be any edges between $S$ and $B$. Also, $|B| \geq(n / 2)-\left(|S|-\left(c_{4} / p\right)\right)$. So, $|B|$ is at least $c_{4} / p$ and when $|B|<n / 4,|S|$ is at least $n / 4$. Hence, using the previous claim, there is an edge between $S$ and $B$ with probability at least $1-e^{-n}$. Hence, we get a contradiction, and hence our claim is true with probability at least $1-e^{-n}$.

Proof of Lemma 22: Let $\mathcal{C}(S)$ denote $|E(S, V(G) \backslash S)|$. For a fixed vertex subset $S$, the expected value of $\mathcal{C}(S)$ is at most $n p|S|$. Therefore, probability that $\mathcal{C}(S)>n p|S| \geq$ $c_{0}|S| \log n$ is at most $n^{-c_{0}|S| / 2}$ using Chernoff bounds. The probability that $\mathcal{C}(S) \leq n p|S|$ for all sets $S$ of size $k$ is at least $1-n^{-c_{0} k / 2+k} \geq 1-n^{-c_{0} / 2+1}$ using union bound and using $k \geq 1$. The probability that $\mathcal{C}(S) \leq n p|S|$ for all vertex subsets $S$ is at least $1-n^{-c_{0} / 2+2} \geq$ using union bound over all $k \in[n]$.


[^0]:    * This work was done while this author was visiting Max Planck Institute for Informatics, Saarbrücken, Germany, supported by Alexander von Humboldt Fellowship.
    $\dagger$ Part of the work done while this author was a visitor at the Courant Institute, NYU. The visit was funded in part by New York University.

[^1]:    ${ }^{1}$ For brevity, we say " $k$-connected" for " $k$-vertex-connected" henceforth.
    ${ }^{2} \mathcal{O}^{*}$ notation hides all polynomial factors in input size.

[^2]:    ${ }^{3}$ Note that the STC problem is relevant only for connected graphs. Since the threshold function for graph connectivity is $\frac{\log n}{n}$, this result applies for almost all of the relevant range of values of $p$.

[^3]:    ${ }^{4}$ In 1994, there was a paper by Ma and Ma in Journal of Computer Science and Technology, which

