ITERATIVE ROUNDING APPROXIMATION ALGORITHMS FOR DEGREE-BOUNDED NODE-CONNECTIVITY NETWORK DESIGN*

TAKURO FUKUNAGA[†], ZEEV NUTOV[‡], AND R. RAVI[§]

Abstract. We consider the problem of finding a minimum edge cost subgraph of a graph satisfying both given node-connectivity requirements and degree upper bounds on nodes. We present an iterative rounding algorithm of the biset LP relaxation for this problem. For directed graphs and k-out-connectivity requirements from a root, our algorithm computes a solution that is a 2-approximation on the cost, and the degree of each node v in the solution is at most $2b(v) + O(k)$ where $b(v)$ is the degree upper bound on v. For undirected graphs and element-connectivity requirements with maximum connectivity requirement k, our algorithm computes a solution that is a 4-approximation on the cost, and the degree of each node v in the solution is at most $4b(v) + O(k)$. These ratios improve the previous $O(\log k)$-approximation on the cost and $O(2^k b(v))$ approximation on the degrees. Our algorithms can be used to improve approximation ratios for other node-connectivity problems such as undirected k-out-connectivity, directed and undirected k-connectivity, and undirected rooted k-connectivity and subset k-connectivity.

1. Introduction.

1.1. Problem definition.

We consider the problem of finding a minimum edge cost subgraph that satisfies both degree-bounds on nodes and certain connectivity requirements between nodes. More formally, the problem is defined as follows.

> ***Degree-Bounded Survivable Network***
> A directed/undirected graph $G = (V, E)$ with edge costs $c\colon E \to \mathbb{R}_+$, connectivity requirements $r\colon V \times V \to \mathbb{Z}_+$, and degree-bounds $b\colon B \to \mathbb{Z}_+$ on a subset B of V are given. The goal is to find a minimum cost edge set $F \subseteq E$ such that in the subgraph (V, F) of G, the uv-connectivity is at least $r(u, v)$ for any $(u, v) \in V \times V$, and the out-degree/degree of each $v \in B$ is at most $b(v)$.

In the case of digraphs, our algorithms easily extend to the case when we also have *in-degree* bounds $b^-: B^- \to \mathbb{Z}_+$ on $B^- \subseteq V$, and require that the in-degree of each $v \in B^-$ is at most $b^-(v)$.

A node $v \in V$ is called *terminal* if there exists $u \in V \setminus \{v\}$ such that $r(u, v) > 0$ or $r(v, u) > 0$. We let T denote the set of terminals. We represent $\max_{u,v \in V} r(u, v)$ by k and $|V|$ by n throughout the paper.

If G is an undirected graph, $B = V$, $b(v) = 2$ for all $v \in B$, and a solution is required to be a connected spanning subgraph, then we get the Hamiltonian Path problem, and hence it is NP-hard even to find a feasible solution. Therefore we consider bi-criteria approximations by relaxing the degree-bounds. We say that an algorithm for Degree Bounded Survivable Network is $(\alpha, \beta(b(v)))$-approximation, or that it has ratio $(\alpha, \beta(b(v)))$, for $\alpha \in \mathbb{R}_+$ and a function $\beta\colon \mathbb{Z}_+ \to \mathbb{R}_+$, if it always outputs a solution such that its cost is at most α times the optimal value, and the degree (the out-degree, in the case of directed graphs) of each $v \in B$ is at most $\beta(b(v))$, for any instance which admits a feasible solution.

Notice that Degree Bounded Survivable Network includes the problem of finding a minimum cost subgraph of required connectivity minimizing the maximum degree.

---

*A preliminary version of this paper appeared in proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2012).

[†]National Institute of Informatics, Japan. JST, ERATO, Kawarabayashi Large Graph Project
[‡]Open University of Israel
[§]Carnegie Mellon University

This can be done by letting $B = V$, and defining $b(v)$ as the uniform bound on the optimal value for all $v \in B$, which can be computed by the binary search. For LP based algorithms, such as ours, when a lower bound is obtained by solving an LP relaxation, these algorithms also establish a "relaxed" integrality gap for the LP relaxation (relaxed since the solutions violate the exact degree requirements).

In this paper, we are interested in node-connectivity and element-connectivity requirements. The node-connectivity $\kappa(u, v)$ is the maximum number of (u, v)-paths that are pair-wise internally (node) disjoint. The definition of element-connectivity assumes that a terminal set T is given. The element-connectivity $\lambda^T(u, v)$ between two terminals $u, v \in T$ is the maximum number of (u, v)-paths that are pair-wise disjoint in edges and in non-terminal nodes. The main two problems we consider are as follows.

Degree Bounded k-Out-connected Subgraph
This is a particular case of Degree Bounded Survivable Network when it is required that (V, F) is k-outconnected from a given root s, namely, when $\kappa(s, v) \geq k$ for each $v \in V \setminus \{s\}$.

Degree Bounded Element-Connectivity Survivable Network
This is a particular case of Degree Bounded Survivable Network when the input graph is *undirected*, and it is required that $\lambda^T(u, v) \geq r(u, v)$ for each $u, v \in T$, where $T \subseteq V$ is a given terminal set.

Our main results for Degree Bounded k-Out-connected Subgraph is for directed graphs. We also present similar results for undirected graphs, but these are derived from the ones for the directed case. Another fundamental problem that we consider for both directed and undirected graphs is as follows.

Degree Bounded k-Connected Subgraph
This is a particular case of Degree Bounded Survivable Network where it is required that (V, F) is k-connected, namely, that $\kappa(u, v) \geq k$ for each $u, v \in V$.

Other special cases of Degree Bounded Survivable Network on undirected graphs are defined according to the connectivity requirements, as follows.

- ***Degree Bounded Node-Connectivity Survivable Network*** requires $\kappa(u, v) \geq r(u, v)$ for each $u, v \in V$.
- ***Degree Bounded Rooted Survivable Network*** is a special case of *Degree Bounded Node-Connectivity Survivable Network* where a *root* node s is specified, and $r(u, v) = 0$ holds if $u \neq s$ and $v \neq s$.
- ***Degree Bounded Subset k-Connected Subgraph*** is a special case of *Degree Bounded Node-Connectivity Survivable Network* where $r(u, v) = k$ if $\{u, v\} \subseteq T$, and $r(u, v) = 0$ otherwise, for a given terminal set $T \subseteq V$.

1.2. Previous work. Survivable Network without degree-bounds is a typical combinatorial optimization problem, that was studies extensively; c.f. [14] for a survey on exact algorithms, and [23] for a survey on approximation algorithms for various Survivable Network problems and their classification w.r.t. costs and requirements. One of the most important methods for these problems is iterative rounding, that was invented in the context of a 2-approximation algorithm by Jain [17]. He showed that every basic optimal solution to an LP relaxation for the undirected Edge-Connectivity Survivable Network always has a variable of value at least $1/2$; see [31] for a simplified proof. The 2-approximation algorithm is obtained by repeatedly rounding up such variables and iterating the procedure until the rounded subgraph is feasible.

Degree Bounded Survivable Network, even with edge-connectivity requirements, was regarded as a difficult problem for a long time because of the above-mentioned hardness on feasibility. A breakthrough was given by Lau, Naor, Salavatipour and Singh [26] and Singh and Lau [36]. They gave a $(2, 2b(v) + 3)$-approximation for the Degree-Bounded Edge-Connectivity Survivable Network problem, and a $(1, b(v) + 1)$-approximation algorithm for the Degree-Bounded Spanning Tree problem. The former result was improved (for large $b(v)$) to a $(2, b(v) + 6k + 3)$-approximation by Lau and Singh [28] afterwards. After their work, many efficient algorithms have been proposed for various types of Degree Bounded Edge-Connectivity Survivable Network problems, such as directed Degree Bounded k-Out-connected Subgraph problems [4], matroid base and submodular flow problems [22], and matroid intersection and optimization over lattice polyhedra [3]. All of them are based on iterative rounding. For applying iterative rounding to a problem with degree-bounds, we need to show that every basic optimal solution to an LP relaxation has a high fractional variable or the subgraph induced by its support has a low degree node on which a degree-bound is given. Once this property is proven, a bi-criteria approximation algorithm can be obtained by repeating rounding up the high fractional variable or dropping the degree-bound on the low degree node. See [27] for a survey on the iterative rounding method.

Despite the success of iterative rounding for edge-connectivity requirements, Degree Bounded Survivable Network with node- and element-connectivity requirements still remain difficult to address with this method. Both [26] and [28] mention that their algorithms for edge-connectivity extend to element-connectivity, but they assumed that degree-bounds were given on terminals only. In [26] it is also shown that undirected Degree Bounded Subset k-Connected Subgraph with $k = \Omega(n)$ admits no $2^{\log^{1-\epsilon} n} b(v)$ degree approximation unless $\mathrm{NP} \subseteq \mathrm{DTIME}(n^{\mathrm{polylog}(n)})$. For the Degree Bounded k-Connected Subgraph problem without costs, Feder, Motwani and Zhu [10] presented an $O(k \log n \cdot b(v))$-approximation algorithm, which runs in $n^{O(k)}$ time. Khandekar, Kortsarz and Nutov [21] proposed a $(4, 6b(v) + 6)$-approximation algorithm for Degree Bounded 2-Connected Subgraph, using iterative rounding. [35] extended the idea of [21] to obtain ratio $(O(\log k), O(2^k) \cdot b(v))$ for Degree Bounded k-Out-connected Subgraph and Degree Bounded Element-Connectivity Survivable Network.

As for Survivable Network without degree bounds, an iterative rounding algorithm computes optimal solutions for Degree Bounded k-Out-connected Subgraph on directed graphs. Fleischer, Jain and Williamson [11] showed that iterative rounding achieves ratio 2 for Element-Connectivity Survivable Network, and also for Node-Connectivity Survivable Network with $k = 2$. Aazami, Cheriyan and Laekhanukit [1] presented an instance of undirected k-Connected Subgraph (without degree-bounds) for which the basic optimal solution to the standard LP relaxation has all variables of value $O\left(\frac{1}{\sqrt{k}}\right)$. Their instance belongs to a special case called *augmentation version*, in which the given graph has a $(k-1)$-connected subgraph of cost zero. On the other hand, several works showed that this augmentation version can be decomposed into a small number p of problems similar to k-Out-connected Subgraph each. The bound on p was subsequently improved [5, 9, 20], culminating in the currently best known bound $O\left(\log \frac{n}{n-k}\right)$ [32], that applies for both directed and undirected graphs. When one applies this method for the general version, an additional factor of $O(\log k)$ is invoked, giving the approximation ratio $O\left(\log k \log \frac{n}{n-k}\right)$ [32]. Cheriyan and Vegh

[7] showed that for undirected graphs with $n \geq k^3(k-1)+k$ this $O(\log k)$ factor can be saved: After solving only two k-Out-connected Subgraph instances, iterative rounding gives a 2-approximation by the work of [11, 18]. This gives ratio 6 for undirected graphs with $n \geq k^3(k-1)+k$.

The decomposition approach was also used for other Survivable Network problems. In [33] it is shown that the augmentation version of Rooted Node-Connectivity Survivable Network is decomposed into $p = O(k)$ instances of a problem that is similar to Element-Connectivity Survivable Network, while in [8] it is shown that Node-Connectivity Survivable Network is decomposed into $p = O(k^3 \log n)$ instances of Element-Connectivity Survivable Network. In [25, 34], it is shown that the augmentation version of Subset k-Connected Subgraph with $k \leq (1-\epsilon)T$ and $0 < \epsilon < 1$ is decomposed into $\frac{1}{\epsilon}O(\log k)$ instances of Rooted Survivable Network.

Summarizing, many connectivity problems are decomposed into k-Out-connected Subgraph and Element-Connectivity Survivable Network problems. Hence it is sufficient to consider these two problems in the degree bounded setting. However, as is indicated in [35], compared to edge-connectivity problems, there is a substantial difficulty in proving that iterative rounding achieves a good result for these problems. We resolve this difficulty by introducing several novel ideas. Moreover, we believe that it is worthwhile investigating the iterative rounding approach for node-connectivity requirements. One reason is that iterative rounding seems to be a promising approach for Degree Bounded Survivable Network problems, as we demonstrate in this paper. A second reason is that it may give new insights for improving the approximability of Node-Connectivity Survivable Network problems (without degree-bounds) with rooted requirements, subset k-connectivity requirements, and general requirements.

1.3. Our results. In this paper, we show that iterative rounding works well for Degree Bounded k-Out-connected Subgraph and Degree Bounded Element-Connectivity Survivable Network problems. Our main results for these two problems are summarized in the following two theorems.

THEOREM 1.1. Degree Bounded Directed k-Out-connected Subgraph *admits approximation ratio* $\left(\alpha, \alpha b(v) + \left\lceil \frac{2(k-1)}{\alpha-1} \right\rceil + 1\right)$ *for any integer* $\alpha \geq 2$.

THEOREM 1.2. Degree Bounded Element-Connectivity Survivable Network *admits the following approximation ratios.*

(i) $\left(\alpha, \alpha b(v) + \left\lceil 4\frac{k+1}{\alpha-2} \right\rceil + 4\right)$ *for any integer* $\alpha \geq 4$.

(ii) $(\infty, 2b(v) + 1.5k^2 + 4.5k + 9)$.

Note that in Theorem 1.2, the degree approximation in part (ii) may be better than the one in part (i) if $b(v) > k^2$.

The ratios in Theorems 1.1 and 1.2 improve the ratio $(O(\log k), O(2^k b(v)))$ of [35].

In [35] it is shown that ratio $(\alpha, \beta(b(v)))$ for Degree Bounded Directed k-Out-connected Subgraph implies ratio $(2\alpha, \beta(b(v)) + k)$ for the undirected case. Thus Theorem 1.1 implies for the undirected case the ratio $(2\alpha, \alpha b(v) + O(k))$ for any integer $\alpha \geq 2$. In particular, for $\alpha = 2$ the ratio is $(4, 2b(v) + O(k))$.

Next, we consider the Degree Bounded k-Connected Subgraph problem. In [35] it is shown that if Degree Bounded k-Out-connected Subgraph admits approximation ratio $(\alpha, \beta(b(v)))$, then Degree Bounded k-Connected Subgraph admits approximation ratio $(\alpha + O(k), \beta(b(v)) + O(k^2))$. We improve this as follows.

THEOREM 1.3. *If* Degree Bounded k-Out-connected Subgraph *admits approximation ratio* $(\alpha, \beta(b(v)))$*, then* Degree Bounded k-Connected Subgraph *admits approximation ratio* $(\mu\alpha + O(k), \beta(b(v)) + O(k\sqrt{k}))$*, where* $\mu = 1$ *for undirected graphs and*

$\mu = 2$ *for digraphs. Consequently,* Degree Bounded k-Connected Subgraph *admits approximation ratio* $(2\mu + O(k), 2b(v) + O(k\sqrt{k}))$.

Cheriyan and Vegh [7] showed that for undirected graphs, an instance of k-Connected Subgraph with $n > k^3(k-1) + k$, can be decomposed into two instances of k-Out-connected Subgraph and a problem similar to Element-Connectivity Survivable Network. We improve their bound $n > k^3(k-1) + k$ to $n \geq k(k-1)(k-1.5) + k$, thus obtaining the following result.

THEOREM 1.4. *For undirected graphs with* $n \geq k(k-1)(k-1.5)+k$, k-Connected Subgraph *admits a* 6*-approximation algorithm, and if* $n \geq 2k(k-1)(k-0.5)+k$ *then* Degree Bounded k-Connected Subgraph *admits approximation ratio* $(12, 8b(v)+O(k))$.

Finally, using known decompositions, we obtain results for other undirected variants of Survivable Network problems, as follows.

THEOREM 1.5. Survivable Network *problems on undirected graphs admits the following approximation ratios for any integer* $\alpha \geq 1$.

(i) $O(k^3 \log |T|) \cdot (\alpha, \alpha b(v) + k/\alpha)$ *for* Degree Bounded Node-Connectivity Survivable Network.
(ii) $O(k \log k) \cdot (\alpha, \alpha b(v) + k/\alpha)$ *for* Degree Bounded Rooted Survivable Network.
(iii) $\frac{1}{\epsilon} O(k \log^2 k) \cdot (\alpha, \alpha b(v) + k/\alpha)$ *for* Degree Bounded Subset k-Connected Subgraph *with* $k \leq (1-\epsilon)|T|$ *and* $0 < \epsilon < 1$.

The rest of this paper is organized as follows. In Section 2 we formulate Theorems 1.1 and 1.2 in terms of biset functions, see Theorems 2.1 and 2.2, respectively. In Section 3 we describe the iterative rounding algorithm that we use, and formulate the latter two theorems in terms of extreme points of appropriate polytopes; see Theorems 3.1 and 3.2, respectively. These two theorems are the key ingredients in proving Theorems 1.1 and 1.2, respectively; we prepare some tools for proving them in Section 4 and prove them formally in Section 5. In the subsequent three Sections 6, 7, and 8, we employ all the tools developed to complete the proofs of Theorems 1.3, 1.4, and 1.5, respectively.

2. Biset edge-covering formulation of Survivable Network problems (Theorems 1.1 and 1.2).

We use a standard setpair LP relaxation, due to Frank and Jordan [12], but we formulate it in equivalent but more convenient terms of *bisets*, as was suggested by Frank [13], and used in several other recent papers [32, 33, 34, 35].

A *biset* is an ordered pair $\hat{S} = (S, S^+)$ of subsets of V such that $S \subseteq S^+$. S is called the *inner-part* of $\hat{S}$ and S^+ is called the *outer-part* of $\hat{S}$. We call $S^+ \setminus S$ the *boundary* of $\hat{S}$, denoted by $\Gamma(\hat{S})$. In the case of undirected graphs, an edge e covers a biset $\hat{S}$ if it has one end-node in S and the other in $V \setminus S^+$, and we denote by $\delta_E(\hat{S})$ the set of edges in E covering $\hat{S}$. In the case of directed graphs, an edge e covers $\hat{S}$ if it enters $\hat{S}$, namely, if e has tail in $V \setminus S^+$ and head in S; we denote by $\delta_E^-(\hat{S})$ the set of edges in E that cover $\hat{S}$. We also denote by $\delta_E^+(\hat{S})$ the set of edges in E that leave $\hat{S}$, namely, edges in E with tail in S and head in $V \setminus S^+$.

Let $\mathcal{V}$ denote the set of all bisets of a groundset V. A graph (V, F) satisfies the connectivity requirements if $|\delta_F(\hat{S})| \geq f(\hat{S})$ for each $\hat{S} \in \mathcal{V}$ in undirected graphs, and $|\delta_F^-(\hat{S})| \geq f(\hat{S})$ for each $\hat{S} \in \mathcal{V}$ in directed graphs, where f is the biset function derived from the connectivity requirements; in this case we say that the graph (V, F) is f*-connected*.

Every set S can be considered as the biset (S, S). Hence degree constraints can be represented by using bisets. For a node v, define a biset $\hat{S}_v = (\{v\}, \{v\})$. Then the degree constraint on a node v in undirected graphs is represented by $|\delta_F(\hat{S}_v)| \leq b(v)$. In digraphs, the out-degree constraint on v is $|\delta^+_F(\hat{S}_v)| \leq b(v)$, and the in-degree constraint on v is $|\delta^-_F(\hat{S}_v)| \leq b^-(v)$. We sometimes abuse the notation to identify a node $v \in V$ as the biset $(\{v\},\{v\})$.

We consider the following generic problem.

> ***Degree Bounded f-Connected Subgraph***
> A graph $G = (V, E)$ with edge/arc costs $c\colon E \to \mathbb{R}_+$, a biset function f on $\mathcal{V}$, and degree-bounds $b\colon B \to \mathbb{Z}_+$ on a subset B of V are given. The goal is to find a minimum cost edge set $F \subseteq E$ such that (V, F) is f-connected, and the degree/out-degree of each $v \in B$ is at most $b(v)$.

Let $\hat{X} = (X, X^+)$ and $\hat{Y} = (Y, Y^+)$ be two bisets. Define

$$\hat{X} \cap \hat{Y} = (X \cap Y, X^+ \cap Y^+),\ \hat{X} \cup \hat{Y} = (X \cup Y, X^+ \cup Y^+),\ \hat{X} \setminus \hat{Y} = (X \setminus Y^+, X^+ \setminus Y)\ .$$

For a biset function f and bisets $\hat{X}, \hat{Y}$, the *supermodular inequality* and the *posimodular inequality*, respectively, are defined as follows

$$f(\hat{X}) + f(\hat{Y}) \leq f(\hat{X} \cap \hat{Y}) + f(\hat{X} \cup \hat{Y}) \tag{2.1}$$

$$f(\hat{X}) + f(\hat{Y}) \leq f(\hat{X} \setminus \hat{Y}) + f(\hat{Y} \setminus \hat{X}) \tag{2.2}$$

A biset function $f\colon \mathcal{V} \to \mathbb{Z}$ is *intersecting supermodular* if any $\hat{X}, \hat{Y}$ with $X \cap Y \neq \emptyset$ satisfy the supermodular inequality (2.1); f is *skew supermodular* if any $\hat{X}, \hat{Y}$ satisfy the supermodular inequality (2.1) or the posimodular inequality (2.2).

For $s \in V$ and an integer $k \geq 1$ define a biset function $g\colon \mathcal{V} \to \mathbb{Z}$ as

$$g(\hat{S}) = \begin{cases} k - |\Gamma(\hat{S})| & \text{if } S \neq \emptyset \text{ and } s \notin S^+ \\ 0 & \text{otherwise.} \end{cases}$$

Then a digraph (V, F) is k-outconnected from s (namely, satisfies $\kappa(s, v) \geq k$ for each $v \in V \setminus \{s\}$) if and only if $|\delta^-_F(\hat{S})| \geq g(\hat{S})$ for each $\hat{S} \in \mathcal{V}$, namely if and only if (V, F) is g-connected. Thus g represents the k-out-connectivity requirements. It is known that g is intersecting supermodular [13].

For $T \subseteq V$ and $r\colon T \times T \to \mathbb{Z}_+$ define a biset function $h\colon \mathcal{V} \to \mathbb{Z}$ as

$$h(\hat{S}) = \begin{cases} \max\limits_{u \in S \cap T, v \in T \setminus S^+} r(u, v) - |\Gamma(\hat{S})| & \text{if } S \cap T \neq \emptyset \neq T \setminus S^+ \text{ and } T \cap \Gamma(\hat{S}) = \emptyset \\ 0 & \text{otherwise.} \end{cases}$$

By a "mixed-connectivity" version of Menger's Theorem (see, e.g. [23]), an undirected graph (V, F) satisfies $\lambda_T(u, v) \geq r(u, v)$ for each $u, v \in T$ if and only if $|\delta_F(\hat{S})| \geq h(\hat{S})$ for each $\hat{S} \in \mathcal{V}$, namely, if and only if (V, F) is h-connected. Thus h represents element-connectivity requirements. By [11], h is skew supermodular.

Let $x(e) \in [0, 1]$ be a variable indicating whether an edge $e \in E$ is chosen or not. Given an edge-set F let $x(F) = \sum_{e \in F} x(e)$. Our LP relaxation for the degree-bounded f-connected subgraph problem is $\min\{c \cdot x \colon x \in P(f, b, E)\}$, where $P(f, b, E)$ is a polytope defined as follows. In the case of directed graphs $P(f, b, E)$ is defined by the constraints

$$\begin{array}{ll} x(\delta^-_E(\hat{S})) \geq f(\hat{S}) & \text{for each } \hat{S} \in \mathcal{V}, \\ x(\delta^+_E(v)) \leq b(v) & \text{for each } v \in B, \\ 0 \leq x(e) \leq 1 & \text{for each } e \in E. \end{array}$$

In the case of undirected graphs, $P(f,b,E)$ is defined by the same constraints with replacing both $\delta^-_E(\hat{S})$ and $\delta^+_E(\hat{S})$ by $\delta_E(\hat{S})$.

Given a biset function f we denote $\gamma=\gamma_f=\max_{f(\hat{S})>0}|\Gamma(\hat{S})|$. Observe that if f is g or h as above, then $\gamma \leq k-1$. Given a biset function f and an edge-set J the *residual biset function* of f is $f_J(\hat{S})=f(\hat{S})-|\delta^-_J(\hat{S})|$ in the case of directed graphs and $f_J(\hat{S})=f(\hat{S})-|\delta_J(\hat{S})|$ in the case of undirected graphs. Given a parameter $\alpha \geq 1$ the *residual degree bounds* are $b^\alpha_J(v)=b(v)-|\delta^+_J(v)|/\alpha$ in the case of directed graphs and $b^\alpha_J(v)=b(v)-|\delta_J(v)|/\alpha$ in the case of undirected graphs. We assume that for any $G=(V,E)$, J, and $\alpha \geq 1$, one can find an extreme point solution to $\min\{c\cdot x\colon x \in P(f_j,b^\alpha_J,E)\}$; this assumption holds for the functions g and h above, and we omit the somewhat standard implementation details. Under this assumption, the following two theorems imply Theorems 1.1 and 1.2.

THEOREM 2.1 (Implies Theorem 1.1). *If f is intersecting supermodular then directed* Degree Bounded f-Connected Subgraph *admits approximation ratio* $(\alpha, \alpha b(v)+\left\lceil\frac{2\gamma}{\alpha-1}\right\rceil+1)$ *for any integer* $\alpha \geq 2$.

THEOREM 2.2 (Implies Theorem 1.2). *If f is skew supermodular then undirected* Degree Bounded f-Connected Subgraph *admits the following approximation ratios.*

(i) $\left(\alpha, \alpha b(v)+\left\lceil 4\frac{\gamma+2}{\alpha-2}\right\rceil+4\right)$ *for any integer* $\alpha \geq 4$.

(ii) $(\infty, 2b(v)+1.5\gamma^2+7.5\gamma+15)$.

3. Iterative rounding algorithms (Theorems 2.1 and 2.2). Here we describe the version of the iterative rounding method we use, and formulate Theorems 2.1 and 2.2 in terms of extreme points of appropriate polytopes. To apply iterative rounding, we define $J \subseteq E$ as the set of edges that have already been chosen as a part of the current solution by the algorithm. We also denote by $\deg_E(v)$ the out-degree/degree of v w.r.t. E, so $\deg_E(v)=|\delta^+_E(v)|$ in the case of digraphs and $\deg_E(v)=|\delta_E(v)|$ in the case of undirected graphs.

Algorithm ITEROUNDING
Input: A graph $G=(V,E)$, $B \subseteq V$, degree-bounds $b\colon B \to \mathbb{Z}_+$, edge costs $c\colon E \to \mathbb{Z}_+$, a biset function $f\colon \mathcal{V} \to \mathbb{Z}$, and integers $\alpha \geq 1$, $\beta \geq 0$ and $\sigma \leq \alpha$.
Output: An f-connected subgraph (V,J) of G.
Step 1: $J:=\emptyset$.
Step 2: Compute a basic optimal solution x^* to $\min\{c\cdot x\colon x \in P(f_J,b^\alpha_J,E)\}$.
Step 3: If there is $e \in E$ such that $x^*(e)=0$ then remove e from E.
Step 4: If there is $e \in E$ such that $x^*(e) \geq 1/\alpha$ then move e from E to J.
Step 5: If there is $v \in B$ such that $\deg_E(v) \leq \sigma b^\alpha_J(v)+\beta$ then remove v from B.
Step 6: If $E \neq \emptyset$, return to Step 2. Otherwise, output (V,J).

The performance of various versions of this algorithm are analyzed in several papers, c.f. [4, 26, 27, 28, 35]. Assume that at each iteration, there exists an edge $e \in E$ as in Steps 3 or 4, or a node $v \in B$ as in Step 5. Then the algorithm ITEROUNDING computes an edge set J of cost at most α times the optimal, such that $\deg_J(v) \leq \alpha b(v)+\beta$ for every $v \in B$ and if $\sigma=0$ then $\deg_J(v) \leq \alpha b(v)+\max\{\beta-1,0\}$ for every $v \in B$ (since $b(v) \geq 1$ for $v \in B$; for details see e.g. [35]). Note that if we care only about the degree approximation, as in part (ii) of Theorems 1.2 and 2.2, then at Step 4 we can also move from E to J any edge e that has no tail/end-node in B, without changing the approximability of the degrees.

Note that an arbitrary set F of directed edges satisfies $|\delta_F^-(\hat{X})| + |\delta_F^-(\hat{Y})| \geq |\delta_F^-(\hat{X} \cap \hat{Y})| + |\delta_F^-(\hat{X} \cup \hat{Y})|$ for any $\hat{X}, \hat{Y} \in \mathcal{V}$. Similarly, note that an arbitrary set F of undirected edges satisfies both $|\delta_F(\hat{X})| + |\delta_F(\hat{Y})| \geq |\delta_F(\hat{X} \cap \hat{Y})| + |\delta_F(\hat{X} \cup \hat{Y})|$ and $|\delta_F(\hat{X})| + |\delta_F(\hat{Y})| \geq |\delta_F(\hat{X} \setminus \hat{Y})| + |\delta_F(\hat{Y} \setminus \hat{X})|$ for any $\hat{X}, \hat{Y} \in \mathcal{V}$. Hence if f is intersecting supermodular, or if f is skew supermodular, so is its residual function f_J.

We will prove the following property of extreme points solutions.

Theorem 3.1 (Implies Theorem 2.1). *Let x^* be an extreme point of $P(f, b, E)$ where G is a directed graph and f is an intersecting supermodular biset function on V. Then for any integer $\alpha \geq 2$, there is $e \in E$ with $x^*(e) = 0$ or $x^*(e) \geq 1/\alpha$, or there is $v \in B$ with $|\delta_E^+(v)| \leq \alpha b(v) + \left\lceil \frac{2\gamma}{\alpha-1} \right\rceil + 1$.*

Theorem 3.2 (Implies Theorem 2.2). *Let x^* be an extreme point of $P(f, b, E)$ where $G = (V, E)$ is an undirected graph and f is a skew supermodular biset function on V. Then the following holds.*

- *For any integer $\alpha \geq 4$, there is $e \in E$ with $x^*(e) = 0$ or $x^*(e) \geq 1/\alpha$, or there is $v \in B$ with $|\delta_E(v)| \leq \left\lceil 4\frac{\gamma+2}{\alpha-2} \right\rceil + 5$.*
- *If every edge in E is incident to some node in B, then there is $e \in E$ with $x^*(e) = 0$ or $x^*(e) \geq 1/2$, or there is $v \in B$ with $|\delta_E(v)| \leq 1.5\gamma^2 + 7.5\gamma + 16$.*

In the next section we prepare some tools needed to prove Theorems 3.1 and 3.2, and prove the theorems in Section 5.

4. Laminar biset families.

A set family $\mathcal{F}$ is *laminar* if for any $X, Y \in \mathcal{F}$, either $X \subseteq Y$, $Y \subseteq X$, or $X \cap Y = \emptyset$. Note that $X \subseteq Y$ or $Y \subseteq X$ holds if and only if $\{X \cap Y, X \cup Y\} = \{X, Y\}$, and that $X \cap Y = \emptyset$ holds if and only if $\{X \setminus Y, Y \setminus X\} = \{X, Y\}$. Our aim is to establish that any extreme point of $P(f, b, E)$ with intersecting supermodular f or with skew supermodular f is defined by a "laminar" family of bisets. But how to define "laminarity" of bisets?

In the case of an intersecting supermodular f, it is natural to say that $\mathcal{F}$ is laminar if $\{\hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y}\} = \{\hat{X}, \hat{Y}\}$ for any $\hat{X}, \hat{Y} \in \mathcal{F}$ with $X \cap Y \neq \emptyset$. In the case of a skew supermodular f, it is natural to say that $\mathcal{F}$ is laminar if $\{\hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y}\} = \{\hat{X}, \hat{Y}\}$ or $\{\hat{X} \setminus \hat{Y}, \hat{Y} \setminus \hat{X}\} = \{\hat{X}, \hat{Y}\}$ for any $\hat{X}, \hat{Y} \in \mathcal{F}$; we refer to this property as "strong laminarity". Laminar biset families are used in [13], while strongly laminar biset families are defined in [11] in terms of setpairs. Following [35], we formulate both these concepts in terms of bisets, by establishing an inclusion order (namely, a partial order) on bisets.

Definition 4.1. *We say that a biset $\hat{Y}$ contains a biset $\hat{X}$ and write $\hat{X} \subseteq \hat{Y}$ if $X \subseteq Y$ and $X^+ \subseteq Y^+$; if also $\hat{X} \neq \hat{Y}$ then $\hat{X} \subset \hat{Y}$ and $\hat{Y}$ properly contains $\hat{X}$. A biset family $\mathcal{F}$ is laminar (resp., strongly laminar) if for any $\hat{X}, \hat{Y} \in \mathcal{L}$ either $\hat{X} \subseteq \hat{Y}$, $\hat{Y} \subseteq \hat{X}$, or $X \cap Y = \emptyset$ (resp., $X \cap Y^+ = \emptyset$ and $Y \cap X^+ = \emptyset$).*

Note that a strongly laminar biset family is laminar, and that the family of inner parts of a laminar biset family is a laminar set family. The above inclusion order on bisets defines a forest structure of a laminar biset family. In the rest of this paper, "minimal" and "maximal" are defined with respect to this inclusion order. For a laminar biset family $\mathcal{F}$ and $\hat{X}, \hat{Y} \in \mathcal{F}$, we say that $\hat{Y}$ is the *parent* of $\hat{X}$ and $\hat{X}$ is a *child* of $\hat{Y}$ if $\hat{Y}$ is the minimal biset with $\hat{X} \subset \hat{Y}$. A minimal biset in $\mathcal{F}$ is called a *leaf*.

Let $G = (V, E)$ be a graph, and let $\mathcal{F}$ be a biset family on V. For $E' \subseteq E$ the incidence vector of E' is an $|E|$-dimensional vector whose component corresponding to $e \in E$ is 1 if $e \in E'$, and 0 otherwise. If G is a directed graph, then let $\chi_E^+(\mathcal{F})$ denote the set of incidence vectors of the edge sets in $\{\delta_E^+(\hat{S})\colon \hat{S} \in \mathcal{F}\}$ and $\chi_E^-(\mathcal{F})$ is the set of incidence vectors of the edge sets in $\{\delta^-_E(\hat{S})\colon \hat{S} \in \mathcal{F}\}$. If G is an undirected graph, let $\chi_E(\mathcal{F})$ denote the set of incidence vectors of the edge sets in $\{\delta_E(\hat{S})\colon \hat{S} \in \mathcal{F}\}$. For $C \subseteq V$, $\chi_E(C) = \chi_E(\{(v,v)\colon v \in C\})$, and similarly $\chi^+_E(C)$ and $\chi^-_E(C)$ are defined. The following statement was proved for set-functions in [4, 26], and the proof for biset functions is similar (e.g., see [35] for the case of a skew supermodular f).

LEMMA 4.2. *Let x be an extreme point of $P(f,b,E)$ with $0 < x_e < 1$ for all $e \in E \neq \emptyset$. Then there exists a family $\mathcal{L}$ of bisets and $C \subseteq B$ such that $|\mathcal{L}| + |C| = |E|$ and such that the following holds:*

(i) *If G is a directed graph, then $x(\delta^-_E(\hat{S})) = f(\hat{S}) \geq 1$ for all $\hat{S} \in \mathcal{L}$, $x(\delta^+_E(v)) = b(v)$ for all $v \in C$, and the vectors in $\chi^-_E(\mathcal{L}) \cup \chi^+_E(C)$ are linearly independent. Furthermore, if f is intersecting supermodular, then there exists such $\mathcal{L}$ that is laminar.*

(ii) *If G is an undirected graph, then $x(\delta_E(\hat{S})) = f(\hat{S}) \geq 1$ for all $\hat{S} \in \mathcal{L}$, $x(\delta_E(v)) = b(v)$ for all $v \in C$, and the vectors in $\chi_E(\mathcal{L}) \cup \chi_E(C)$ are linearly independent. Furthermore, if f is skew supermodular, then there exists such $\mathcal{L}$ that is strongly laminar.*

The following parameter is defined in [35].

DEFINITION 4.3. *Let $\mathcal{L}$ be a laminar biset family on V, let $\hat{S} \in \mathcal{L}$ and let $v \in V$. We say that $\hat{S}$* owns *v if $\hat{S}$ is the minimal biset in $\mathcal{L}$ with $v \in S$. We say that $\hat{S}$* shares *v if $\hat{S}$ is a minimal biset in $\mathcal{L}$ with $v \in \Gamma(\hat{S})$. Let $\Delta_{\mathcal{L}}(v)$ denote the number of bisets in $\mathcal{L}$ that share v.*

From the definition it follows that every node v is owned by at most one biset, and that if $\hat{X}, \hat{Y}$ share the same node v then they are incomparable, namely, that none of them contains the other. From the definition of laminarity and strong laminarity we have the following.

LEMMA 4.4. *Let $\mathcal{L}$ be a laminar biset family, and let $v \in V$. Let $\mathcal{X}$ be a subfamily of $\mathcal{L}$ such that $v \in \Gamma(\hat{X})$ for each $\hat{X} \in \mathcal{X}$, and the bisets in $\mathcal{X}$ are pair-wise incomparable. For each $\hat{X} \in \mathcal{X}$, let $\hat{X}'$ be a biset in $\mathcal{L}$ such that $\hat{X} \subseteq \hat{X}'$ and $\hat{X}'$ contains no biset in $\mathcal{X} \setminus \{\hat{X}\}$. Then $v \in \Gamma(\hat{X}')$ or $v \in X'$ holds, and the latter holds for at most one biset in $\mathcal{X}' = \{\hat{X}' \mid \hat{X} \in \mathcal{X}\}$. Furthermore, if $\mathcal{L}$ is strongly laminar then the former holds for all bisets in $\mathcal{X}'$.*

Proof. Since $\Gamma(\hat{X})$ is contained by the outer-part of $\hat{X}'$, $v \in \Gamma(\hat{X})$ is either in $\Gamma(\hat{X}')$ or in X'. Any two bisets $\hat{X}'$ and $\hat{Y}'$ in $\mathcal{X}'$ are incomparable, and hence $X' \cap Y' = \emptyset$ by the laminarity of $\mathcal{L}$. It follows from this fact that v is contained by the inner-part of at most one biset in $\mathcal{X}'$. When $\mathcal{L}$ is strongly laminar, $X' \cap \Gamma(\hat{Y}') = \Gamma(\hat{X}') \cap Y' = \emptyset$ holds for any $\hat{X}', \hat{Y}' \in \mathcal{X}'$, and hence $v \in \Gamma(X')$ holds for all $\hat{X}' \in \mathcal{X}'$. ◻

LEMMA 4.5. *Let $\mathcal{L}$ be a biset family on V, let $C \subseteq V$, let $\mathcal{E}$ be the set of leaves of $\mathcal{L}$, and let $\gamma = \max_{\hat{S} \in \mathcal{L}} |\Gamma(\hat{S})|$. Then*

$$\sum_{v \in C} \max\{\Delta_{\mathcal{L}}(v), 1\} \leq 2\gamma(|\mathcal{E}| - 1) + |C| \quad \text{if } \mathcal{L} \text{ is laminar,} \tag{4.1}$$

$$\sum_{v \in C} \max\{\Delta_{\mathcal{L}}(v), 1\} \leq \gamma|\mathcal{E}| + |C| \quad \text{if } \mathcal{L} \text{ is strongly laminar.} \tag{4.2}$$

Proof. We may assume that $|\Delta_{\mathcal{L}}(v)| \geq 2$ for every $v \in C$, as if $\Delta_{\mathcal{L}}(v) \leq 1$ for some $v \in C$, then excluding v from C decreases both sides of each of (4.1) and (4.2) by exactly 1.

We prove (4.1). Let $\mathcal{L}'$ be the family of bisets in $\mathcal{L}$ whose parent has at least 2 children. It is known and easy to prove by induction that $|\mathcal{L}'| \leq 2|\mathcal{E}| - 2$. There are

$\Delta_{\mathcal{L}}(v)$ incomparable bisets $\hat{S}$ in the tree with $v \in \Gamma(\hat{S})$. By walking up to a biset that owns v until the first biset in $\mathcal{L}'$ is reached, we see that v belongs to the boundary of at least $\Delta_{\mathcal{L}}(v) - 1$ bisets in $\mathcal{L}'$. This implies

$$\sum_{v \in C} \max\{\Delta_{\mathcal{L}}(v), 1\} - |C| = \sum_{v \in C} (\max\{\Delta_{\mathcal{L}}(v), 1\} - 1) \leq \gamma|\mathcal{L}'| \leq 2\gamma(|\mathcal{E}| - 1) .$$

We prove (4.2) by induction on $|\mathcal{E}|$. Assume therefore that $|\mathcal{E}| \geq 2$, as otherwise $|\Delta_{\mathcal{L}}(v)| \leq 1$ for every $v \in C$. Then there exists $\hat{S} \in \mathcal{L}$ such that $\hat{S}$ has at least 2 children, but every proper descendant of $\hat{S}$ has at most one child. Let $\hat{R}$ be a child of $\hat{S}$, let $\hat{Z} \subseteq \hat{R}$ be a leaf of $\mathcal{L}$, and let $\mathcal{P} = \{\hat{Y} \in \mathcal{L} \colon \hat{Z} \subseteq \hat{Y} \subseteq \hat{R}\}$ be the "chain" from the child $\hat{R}$ of $\hat{S}$ to $\hat{Z}$ (possibly $\hat{R} = \hat{Z}$). Since we assume that $|\Delta_{\mathcal{L}}(v)| \geq 2$ for every $v \in C$, then by Lemma 4.4, every node that is shared by some biset in $\mathcal{P}$ belongs to $\Gamma(\hat{R})$, so there are at most $|\Gamma(\hat{R})| \leq \gamma$ such nodes. Hence excluding the bisets in $\mathcal{P}$ from $\mathcal{L}$ decreases the left hand side of (4.2) by at most γ, while $|\mathcal{E}|$ decreases by 1, hence the right hand side of (4.2) decreases by γ. ☐

5. Proof of Theorems 3.1 and 3.2. Let $x = x^*$ be an extreme point solution to the corresponding biset LP relaxation, and let $\mathcal{L}$ and C be as in Lemma 4.2. Let $\mathcal{E}$ be the set of leaf bisets in $\mathcal{L}$. For a biset $\hat{S} \in \mathcal{L}$ denote by $\mathcal{C}_S$ the set of children of $\hat{S}$, by E_S^+ the set of edges in E that cover $\hat{S}$ but not a child of $\hat{S}$, and by E_S^- the set of edges in E that cover a child of $\hat{S}$ but not $\hat{S}$. If $\hat{S} \in \mathcal{E}$, then $E_S^+ = \delta_E^-(\hat{S})$ when E is the set of arcs, and $E_S^+ = \delta_E(\hat{S})$ when E is the set of undirected edges. Hence $E_S^+ \neq \emptyset$ in this case. If $\hat{S} \in \mathcal{L} \setminus \mathcal{E}$, then $E_S^+ \cup E_S^- \neq \emptyset$ since otherwise the vectors in $\chi_E(\{\hat{S}\} \cup \mathcal{C}_S)$ are linearly dependent.

5.1. Intersecting supermodular f and directed graphs (Theorem 3.1). Assume to the contrary that $0 < x(e) < 1/\alpha$ for every $e \in E$. Assign 1 token to every edge $e = uv \in E$, putting $1 - \alpha x(e) > 0$ "tail-tokens" at u and $\alpha x(e) > 0$ "head-tokens" at v. We will show that these tokens can be distributed such that every member of $\mathcal{L} \cup C$ gets 1 token, and some tokens are not assigned. This gives the contradiction $|E| > |\mathcal{L}| + |C|$.

For every $\hat{S} \in \mathcal{L}$, the amount of head-tokens of edges in E_S^+ and tail-tokens of edges in E_S^- is $\alpha x(E_S^+) + |E_S^-| - \alpha x(E_S^-)$. Note that this is an integer, since α is an integer, and since $x(E_S^+) - x(E_S^-) = f(\hat{R}) - \sum_{\hat{R} \in \mathcal{C}_S} f(\hat{R})$. It is a positive integer since $x(E_S^+) > 0$ if $E_S^+ \neq \emptyset$, and $|E_S^-| - \alpha x(E_S^-) > 0$ if $E_S^- \neq \emptyset$. Thus if we assign to every $\hat{S} \in \mathcal{L}$ the head-tokens of edges in E_S^+ and the tail-tokens of edges in E_S^-, then every member of $\mathcal{L}$ will get at least one token, and the tail-tokens entering the maximal members of $\mathcal{L}$ are not assigned. It is easy to see that no token is assigned twice. We note that if $C = \emptyset$, and in particular if there are no degree bounds, then this implies that the extreme points of the polytope $P(f, b, E)$ are all integral. Every leaf $\hat{S}$ gets $\alpha x(\delta_E^-(\hat{S})) = \alpha f(\hat{S}) \geq \alpha$ head-tokens from edges in E_S^+. Hence we have $\alpha|\mathcal{E}|$ tokens at leaves.

LEMMA 5.1. *If $(\alpha - 1)|\mathcal{E}| \geq |C|$ then there is $e \in E$ with $x(e) \geq 1/\alpha$.*

Proof. By the assumption $(\alpha - 1)|\mathcal{E}| \geq |C|$ we have $|\mathcal{E}| + |C| \leq \alpha|\mathcal{E}|$, hence the $\alpha|\mathcal{E}|$ tokens at leaves suffice to give 1 token to each member of $\mathcal{E} \cup C$. Every non-leaf biset $\hat{S} \in \mathcal{L}$ gets the head-tokens from edges in E_S^+ and the tail-tokens from edges in E_S^-, so at least 1 token. Consequently, every members of $\mathcal{L} \cup C$ gets 1 token, and the tail-tokens of the edges entering the maximal members of $\mathcal{L}$ are not assigned. This gives the contradiction $|E| > |\mathcal{L}| + |C|$. ☐

LEMMA 5.2. *If $|C| > (\alpha-1)|\mathcal{E}|$, then there is $e \in E$ with $x(e) \geq 1/\alpha$ or there is $v \in C$ with $|\delta_E^+(v)| \leq \alpha b(v) + \beta$, where $\beta = \left\lceil \frac{2\gamma}{\alpha-1} \right\rceil + 1$.*

Proof. Assume that $|\delta_E^+(v)| \geq \alpha b(v) + \beta + 1$ for every $v \in C$. Then the amount of tail-tokens at each $v \in C$ is at least $\alpha b(v) + \beta + 1 - \alpha x(\delta_E^+(v)) = \beta + 1$. Hence we have at least $\alpha|\mathcal{E}| + (\beta+1)|C|$ tokens at leaves and nodes in C. From these tokens, we give 1 token to every leaf and $\Delta_{\mathcal{L}}(v) + 2$ tokens to every $v \in C$, and spare tokens remain. This is possible, since by Lemma 4.5 and the assumption $|C| > (\alpha-1)|\mathcal{E}|$

$$\begin{aligned}
|\mathcal{E}| + \sum_{v \in C} (\Delta_{\mathcal{L}}(v) + 2) &\leq |\mathcal{E}| + 3|C| + 2\gamma(|\mathcal{E}| - 1) \\
&= \alpha|\mathcal{E}| + (2\gamma + 1 - \alpha)|\mathcal{E}| + 3|C| - 2\gamma \\
&= \alpha|\mathcal{E}| + |C| \left((2\gamma + 1 - \alpha)\frac{|\mathcal{E}|}{|C|} + 3 \right) - 2\gamma \\
&< \alpha|\mathcal{E}| + |C| \left(\frac{2\gamma}{\alpha - 1} + 2 \right) - 2\gamma \\
&\leq \alpha|\mathcal{E}| + |C|(\beta + 1) \ .
\end{aligned}$$

Every $v \in C$ will keep one token, and from the remaining at least $\Delta_{\mathcal{L}}(v) + 1$ tokens v will give 1 token to every biset that owns or shares v. Now let $\hat{S} \in \mathcal{L}$ be a biset that is not a leaf and that does not own or share a node in C. Then $\hat{S}$ gets the head-tokens from edges in E_S^+ and the tail tokens from edges in E_S^-, so at least one token as argued above. Consequently, every members of $\mathcal{L} \cup C$ gets 1 token. This gives the contradiction $|E| > |\mathcal{L}| + |C|$. ◻

Theorem 3.1 follows by combining Lemmas 5.1 and 5.2.

5.2. Skew supermodular f and undirected graphs (Part (i) of Theorem 3.2). We deduce part (i) of Theorem 3.2 from the following two lemmas.

LEMMA 5.3. *If $(\theta-1)|\mathcal{E}| \geq |C|$ for an integer $\theta \geq 2$, then there is $e \in E$ with $x(e) \geq \frac{1}{2\theta}$.*

Proof. We generalize the approach from [31]. Assume to the contrary that $0 < x(e) < \frac{1}{2\theta}$ for every $e \in E$. Assign 1 token to every $e = uv \in E$, putting $\theta x(e) > 0$ "end-tokens" at each of u, v and $1 - 2\theta x(e) > 0$ "middle-tokens" at e. We will show that these tokens can be distributed such that every member of $\mathcal{L} \cup C$ gets 1 token, and the middle-tokens of the edges entering the maximal members of $\mathcal{L}$ are not assigned. This gives the contradiction $|E| > |\mathcal{L}| + |C|$.

Every leaf $\hat{S}$ gets $\theta x(\delta_E(\hat{S})) = \theta f(\hat{S}) \geq \theta$ end-tokens from edges in E_S^+. Hence we have $\theta|\mathcal{E}|$ tokens at leaves. By the assumption $(\theta-1)|\mathcal{E}| \geq |C|$ we have $|\mathcal{E}| + |C| \leq \theta|\mathcal{E}|$, so these tokens suffice to give 1 token to each member of $\mathcal{E} \cup C$.

Now let $\hat{S} \in \mathcal{L}$ be a non-leaf biset. Denote by $t(\hat{S})$ the amount of end-tokens of edges in E_S^+ at nodes owned by $\hat{S}$ and middle-tokens of edges in E_S^-. Note that $t(\hat{S}) > 0$, by the linear independence. We claim that $t(\hat{S})$ is an integer, hence $t(\hat{S}) \geq 1$. Let E_1 be the set of edges in E_S^- that cover exactly one child of $\hat{S}$ and let $E_2 = E_S^- \setminus E_1$ be the set of edges in E_S^- that cover two distinct children of $\hat{S}$. Let E' be the set of edges in E that cover both $\hat{S}$ and some child of $\hat{S}$. Then

$$\begin{aligned} t(\hat{S}) &= \theta x(E_S^+) + (|E_1| - \theta x(E_1)) + (|E_2| - 2\theta x(E_2)) \\ &= \theta[x(E_S^+) + x(E')] - \theta[x(E_1) + 2x(E_2) + x(E')] + |E_1| + |E_2| \\ &= \theta x(\delta_E(\hat{S})) - \theta \sum_{\hat{R} \in \mathcal{C}_S} x(\delta_E(\hat{R})) + |E_S^-| = \theta \left(f(\hat{S}) - \sum_{\hat{R} \in \mathcal{C}_S} f(\hat{R}) \right) + |E_S^-| \ . \end{aligned}$$

Consequently, every members of $\mathcal{L} \cup C$ gets 1 token, and the middle-tokens of the edges entering the maximal members of $\mathcal{L}$ are not assigned. This gives the contradiction $|E| > |\mathcal{L}| + |C|$. ◻

We note that if $C = \emptyset$, i.e., if there are no degree bounds, then the same proof applies for $\theta = 1$ to show that any extreme point of $P(f, b, E)$ has an edge $e \in E$ with $x(e) \geq 1/2$. This gives a simple proof of the result of [11], using the idea of [31].

LEMMA 5.4. *If $|C| > (\alpha/2 - 1)|\mathcal{E}|$, then there is $e \in E$ with $x(e) \geq 1/\alpha$ or there is $v \in C$ with $|\delta_E(v)| \leq \beta$, where $\beta = \left\lceil 4\frac{\gamma+2}{\alpha-2} \right\rceil + 5$.*

Proof. Assume to the contrary that $0 < x(e) < 1/\alpha$ for every $e \in E$ and that $|\delta_E(v)| \geq \beta + 1$ for every $v \in C$. We give 1 token to each end-node of every $e \in E$. We will show that these tokens can be distributed such that every member of $\mathcal{L} \cup C$ gets 2 tokens, and each maximal member of $\mathcal{L}$ gets 4 tokens leading to the contradiction that $|E| > |\mathcal{L}| + |C|$.

The amount of tokens at each $v \in C$ is at least $\beta + 1$. Hence we have at least $(\beta + 1)|C|$ tokens at the nodes in C. From these tokens, we give 4 tokens to every leaf and $2(\Delta_{\mathcal{L}}(v) + 2)$ tokens to every $v \in C$. This is possible since by Lemma 4.5 and the assumption $|C| > (\alpha/2 - 1)|\mathcal{E}|$, as we verify below.

$$\begin{aligned} 4|\mathcal{E}| + 2\sum_{v \in C} (\Delta_{\mathcal{L}}(v) + 2) &\leq 4|\mathcal{E}| + 2(\gamma|\mathcal{E}| + 3|C|) = 2|\mathcal{E}|(\gamma + 2) + 6|C| \\ &= |C| \left(6 + 2\frac{|\mathcal{E}|}{|C|}(\gamma + 2) \right) \leq |C| \left(6 + 2(\gamma + 2)\frac{2}{\alpha - 2} \right) \leq |C|(1 + \beta). \end{aligned}$$

Every $v \in C$ will keep 2 tokens, and from the remaining $2(\Delta_{\mathcal{L}}(v) + 1)$ tokens v will give 2 tokens to every biset that owns or shares v. Now let $\hat{S} \in \mathcal{L}$ be a biset that is not a leaf. Then $\hat{S}$ gets 2 tokens from each child. If $\hat{S}$ has at least 2 children then we are done. If $\hat{S}$ has 1 child and owns or shares a node $v \in C$, then $\hat{S}$ gets 2 tokens from its child and 2 tokens from v. Otherwise, $\hat{S}$ gets 2 tokens from its child and 2 tokens from edges in $E_S^+ \cup E_S^-$; note that $|E_S^+ \cup E_S^-| \geq 2$, by linear independence and the integrality of f. Consequently, every member of $\mathcal{L} \cup C$ gets 2 tokens, and each maximal member of $\mathcal{L}$ gets 4 tokens. This gives the contradiction $|E| > |\mathcal{L}| + |C|$. ◻

Applying Lemma 5.3 with $\theta = \lfloor \alpha/2 \rfloor$, we get that if $(\lfloor \alpha/2 \rfloor - 1)|\mathcal{E}| \geq |C|$, and in particular if $(\alpha/2 - 1)|\mathcal{E}| \geq |C|$, then there is $e \in E$ with $x(e) \geq \frac{1}{2\lfloor \alpha/2 \rfloor} \geq \frac{1}{\alpha}$. Together with Lemma 5.4 this implies Theorem 3.2.

5.3. Degree approximation only (Part (ii) of Theorem 3.2). We call a biset in $\mathcal{L}$ *strictly black* if it owns a node in C, *black* if one of its descendents is strictly black (i.e., its inner-part contains a node in C), and *white* otherwise (i.e., its inner-part contains no node in C). Let $\mathcal{E}_{\mathrm{b}}$ and $\mathcal{E}_{\mathrm{w}}$ denote the family of strictly black bisets and white bisets in $\mathcal{E}$, respectively.

LEMMA 5.5. *If $|\mathcal{E}| \leq (\gamma + 4)|C|$, then there is $e \in E$ with $x(e) \geq 1/2$, or there is $v \in C$ with $|\delta_E(v)| \leq 1.5\gamma^2 + 7.5\gamma + 16$.*

Proof. Assume to the contrary that $0 < x(e) < 1/2$ for every $e \in E$ and $|\delta_E(v)| \geq 1.5\gamma^2 + 7.5\gamma + 17$ for every $v \in C$. Identifying a node $v \in C$ as a biset $(\{v\}, \{v\})$, we regard $\mathcal{L} \cup C$ as a biset family. $\mathcal{L} \cup C$ may not be strongly laminar, but it is laminar. Therefore we can define the inclusion order on $\mathcal{L} \cup C$.

We assign two tokens to every edge in E, putting one end-token at each of its end-nodes. We will show that these tokens can be distributed such that every members of $\mathcal{L} \cup C$ gets two tokens, and an extra token remains. This gives the contradiction that $|E| > |\mathcal{L}| + |C|$.

Let $e = uv \in E$. Note that there always exists a biset $\hat{X} \in \mathcal{L} \cup C$ such that $e \in \delta_E(\hat{X})$. Suppose that $\hat{X}$ is a minimal one among such bisets. Without loss of generality, let $u \in X$. We give the end-token of e at u to $\hat{X}$. If there also exists a biset $\hat{Y} \in \mathcal{L} \cup C$ such that $e \in \delta_E(\hat{Y})$ and $v \in Y$, then we give the end-token of e at v to the minimal such biset $\hat{Y}$. Otherwise, the end-token of e at v is given to the minimal biset $\hat{X}'$ such that $\hat{X} \subset \hat{X}'$ and $e \notin \delta_E(\hat{X}')$.

Since bisets in $\mathcal{E}_{\mathrm{w}}$ and nodes in C are leaves of $\mathcal{L} \cup C$, they obtain one token from each edge incident to them after this distribution. Hence each biset $\hat{S} \in \mathcal{E}_{\mathrm{w}}$ has 3 tokens and each node $v \in C$ has $|\delta_E(v)|$ tokens. We make each $v \in C$ keep only two tokens, return back 1/3 tokens to each edge in $\delta_E(v)$, and release the other tokens. Then the total number of released tokens is

$$\sum_{v \in C} \left(\frac{2}{3} |\delta_E(v)| - 2 \right) > |C|(\gamma^2 + 5\gamma + 9) \geq (1 + \gamma)|\mathcal{E}| + 5|C|,$$

where the first inequality follows from $|\delta_E(v)| > 1.5\gamma^2 + 7.5\gamma + 16.5$, $v \in C$ and the last one follows from $|\mathcal{E}| \leq (\gamma + 4)|C|$. We redistribute these tokens so that one token is given to each biset in $\mathcal{E}_{\mathrm{w}}$, 4 tokens to each in $\mathcal{E}_{\mathrm{b}}$, one token to each biset that shares or owns a node in C. Note that if a biset owns or shares more than one node in C, it obtains tokens from each of those nodes in C. This is possible because the number of tokens we need is

$$|\mathcal{E}_{\mathrm{w}}| + 4|\mathcal{E}_{\mathrm{b}}| + \sum_{v \in C} (1 + \Delta_{\mathcal{L}}(v)) \leq (|\mathcal{E}| - |C|) + 4|C| + |C| + \gamma|\mathcal{E}| + |C| = (1 + \gamma)|\mathcal{E}| + 5|C|,$$

where the above inequality follows from $|\mathcal{E}_{\mathrm{w}}| + |\mathcal{E}_{\mathrm{b}}| = |\mathcal{E}|$, $|\mathcal{E}_{\mathrm{b}}| \leq |C|$ and (4.2).

Now all tokens given to bisets in $\mathcal{E}$ and nodes in C have been redistributed so that each node in C has 2 tokens, each biset in $\mathcal{E}$ has 4 tokens, each pair of $v \in C$ and $e \in \delta_E(v)$ has 1/3 tokens, and each biset $\mathcal{L}$ has one token from each owning or sharing node in C. In what follows, we prove by induction that each biset in $\mathcal{L}$ can have at least two tokens, and each maximal biset in $\mathcal{L}$ can have 4 tokens.

It is obvious to see that the claim holds in the base case where the height of the forest defined from $\mathcal{L}$ is one. Let us consider the case where the height of the forest is more than one, and let $\hat{S}$ be a biset that is not a leaf. By the induction hypothesis, each descendent has at least two tokens, and each child of $\hat{S}$ has 4 tokens. $\hat{S}$ can obtain 2 tokens from each of its child. Hence if $\hat{S}$ has more than one child, if $\hat{S}$ owns a node in C, or if $\hat{S}$ shares at least two nodes in C, then $\hat{S}$ can collect 4 tokens.

By the linear independence, $|E_S^+ \cup E_S^-| \geq 2$ always holds, and $|E_S^+ \cup E_S^-| \geq 3$ holds when either E_S^+ or E_S^- is empty by the assumption that $x(e) < 1/2$, $e \in E$. Let $e = uv \in E_S^+ \cup E_S^-$. If $e \in E_S^+$, we let $u \in S$ and $v \in V \setminus S^+$. If $e \in E_S^-$, we let $u \in S$ be an end-node which is within the inner-part of a child of S, and $v \in S^+$. Notice that $v \in S^+$ implies that no biset $\hat{X} \in \mathcal{L}$ with $e \in \delta_E(\hat{X})$ contains v in its inner-part since $\mathcal{L}$ is strongly laminar.

$\hat{S}$ obtains the end-token of e at u if $e \in E_S^+$ and $u \notin C$. If $e \in E_S^+$ and $u \in C$, then $\hat{S}$ owns u, and hence $\hat{S}$ obtains one token corresponding to u. On the other hand, when $e \in E_S^-$, $\hat{S}$ obtains the end-token of e at v if $v \notin C$, and obtains one token corresponding to v if $v \in C$ since $\hat{S}$ owns or shares v. Therefore $\hat{S}$ can collect 4 tokens unless $E_S^- = \emptyset$ and all nodes in E_S^+ are incident to the same node in $C \cap S$, or unless $E_S^+ = \emptyset$ and all nodes in E_S^- are incident to the same node in $C \cap S^+$. In both of these two cases, $|E_S^+ \cup E_S^-| \geq 3$, and each of the edges in $E_S^+ \cup E_S^-$ has 1/3 tokens corresponding to the end-node to which all the edges are incident. Therefore $\hat{S}$ collects 4 tokens even in these cases. □

We next discuss the case where $|\mathcal{E}| \geq (\gamma + 4)|C|$, and prove the following lemma.

LEMMA 5.6. *If every edge in E is incident to some node in B and $|\mathcal{E}| \geq (\gamma+4)|C|$, then there is $e \in E$ with $x(e) \geq 1/2$.*

Under the assumption in Lemma 5.6, each edge $e \in E$ is incident to a node in $\{v \in B \mid x^*(\delta_E(v)) = b(v)\}$ since otherwise we can decrease $x^*(e)$. Recall that C is a maximal subset of $\{v \in B \mid x^*(\delta_E(v)) = b(v)\}$ such that the vectors in $\chi_E(C)$ are linearly independent. If C contains no end-node of $e \in E$, then the incidence vector of $\delta_E(v)$ defined from an end-node v of e is linearly independent from those in $\chi_E(C)$. Since this contradicts the maximality of C, we can observe that each edge $e \in E$ is incident to at least one node in C.

We again count bisets in $\mathcal{L}$ and nodes in C for proving Lemma 5.6, but the way of distributing tokens is different here. Let $e = uv \in E$. By the assumption, at least one of the end-nodes of e is in C. If C contains both end-nodes of e, then we assign no token to e. If C contains exactly one end-node, say v, of e, then we assign one token. This token will be given to a biset in $\mathcal{L}$ as follows. If $\mathcal{L}$ contains a biset $\hat{S}$ such that $e \in \delta_E(\hat{S})$ and $u \in S$, then the token is given to such a minimal biset. If there exists no such bisets and $\mathcal{L}$ contains a biset $\hat{X}$ such that $e \in \delta_E(\hat{X})$ and $v \in X$, then the token is given to the minimal biset in $\{\hat{Y} \in \mathcal{L} \mid \hat{X} \subset \hat{Y}, u \in Y^+\}$. Since the total number of tokens is at most $|E|$, it suffices to show that an extra token remains after redistributing tokens so that each biset in $\mathcal{L}$ and each node in C owns one token.

Let $\hat{S} \in \mathcal{E}_{\mathrm{w}}$. Since $x^*(e) < 1/2$ for each $e \in E$, $|\delta_E(\hat{S})| \geq 2f(\hat{S}) + 1$. Since S contains no nodes in C, each edge in $\delta_E(\hat{S})$ gives a token to $\hat{S}$. Thus $\hat{S}$ has $2f(\hat{S}) + 1 \geq f(\hat{S}) + 2$ tokens. We make each $\hat{S} \in \mathcal{E}_{\mathrm{w}}$ release one token. Then the number of released tokens is at least $|\mathcal{E}_{\mathrm{w}}| \geq |\mathcal{E}| - |C| \geq (\gamma + 3)|C|$. Recall that the number of strictly black bisets is at most $|C|$. We redistribute the released token to the nodes in C and the strictly black bisets so that each $v \in C$ has 1 token, and each strictly black biset has $\gamma + 2$ tokens. Note that each $\hat{S} \in \mathcal{E}_{\mathrm{w}}$ still has at least $f(\hat{S}) + 1$ tokens after this redistribution.

We first count tokens in a tree which consists of only white bisets.

LEMMA 5.7. *Let $\hat{R} \in \mathcal{L}$ be a white biset, and $\mathcal{L}' = \{\hat{S} \in \mathcal{L} \mid \hat{S} \subseteq \hat{R}\}$. We can distribute tokens owned by bisets in $\mathcal{L}'$ so that each biset in $\mathcal{L}'$ has at least one token, and $\hat{R}$ has at least $1 + f(\hat{R})$ tokens when $0 < x(e) < 1/2$ for each $e \in E$.*

Proof. We prove by the induction on the height of the tree representing $\mathcal{L}'$. If the height is one, then $\mathcal{L}' = \{\hat{R}\}$ and $\hat{R} \in \mathcal{E}_{\mathrm{w}}$. Thus the lemma follows in this case.

Assume that the height is at least 2. Applying the induction hypothesis to the trees rooted at the children of $\hat{R}$, we can allocate tokens so that each biset below the children of $\hat{R}$ has one token, and each child $\hat{S}$ of $\hat{R}$ has $1 + f(\hat{S})$ tokens. We can move $\sum_{\hat{S} \in \mathcal{C}_R} f(\hat{S})$ tokens from the children to $\hat{R}$.

If $\sum_{\hat{S} \in \mathcal{C}_R} f(\hat{S}) > f(\hat{R})$, then we are done. Hence consider the other case. When $\sum_{\hat{S} \in \mathcal{C}_R} f(\hat{S}) = f(\hat{R})$, $|E_R^+| \geq 1$ holds by the linear independence. When $\sum_{\hat{S} \in \mathcal{C}_R} f(\hat{S}) <$

$f(\hat{R})$, $|E_R^+| \geq 1 + 2(f(\hat{R}) - \sum_{\hat{S} \in \mathcal{C}_R} f(\hat{S}))$ holds by $x^*(e) < 1/2$, $e \in E$. In either case, $|E_R^+| \geq 1 + f(\hat{R}) - \sum_{\hat{S} \in \mathcal{C}_R} f(\hat{S})$. $\hat{R}$ is given a token from each $e \in E_R^+$ because e has an end-node $v \in R$ such that $\hat{R}$ is a minimal biset with $e \in \delta_E(\hat{R})$ and $v \in R$, and $v \notin C$ by $R \cap C = \emptyset$. Thus $\hat{R}$ has already owned $1 + f(\hat{R}) - \sum_{\hat{S} \in \mathcal{C}_R} f(\hat{S})$. With the tokens from the children, $\hat{R}$ obtains $1 + f(\hat{R})$ tokens. □

We next give a token distribution scheme for trees in which the maximal bisets are black. Together with Lemma 5.7, this finishes the proof of Lemma 5.6.

LEMMA 5.8. *Let $\hat{R} \in \mathcal{L}$ be a black biset, and $\mathcal{L}' = \{\hat{S} \in \mathcal{L} \mid \hat{S} \subseteq \hat{R}\}$. We can distribute tokens owned by bisets in $\mathcal{L}'$ so that each biset in $\mathcal{L}'$ has at least one token, and $\hat{R}$ has at least 2 tokens when $0 < x(e) < 1/2$ for each $e \in E$.*

Proof. We show how to rearrange the tokens so that each biset in $\mathcal{L}'$ obtains at least one token, and $\hat{R}$ obtains at least $2 + \gamma - |\Gamma(\hat{R}) \cap C| \geq 2$ tokens. Our proof is by the induction on the height of the tree. If the height is one, then the claim holds because it consists of a strictly black biset. Hence suppose that the height is at least two.

Let $\mathcal{B}$ be the set of black children of $\hat{R}$, and $\mathcal{W}$ be the set of white children of $\hat{R}$. Apply the induction hypothesis to the subtrees rooted at the black children, and Lemma 5.7 to the subtrees rooted at the white children. Then each biset below the children has one token, each $\hat{X} \in \mathcal{B}$ has $2 + \gamma - |\Gamma(\hat{X}) \cap C|$ tokens, and each $\hat{Y} \in \mathcal{W}$ has $1 + f(\hat{Y}) \geq 2$ tokens. If $\mathcal{B} = \emptyset$, then $\hat{R}$ is strictly black, and it has already given $\gamma + 2$ tokens. Since this finishes the claim, suppose that $\mathcal{B} \neq \emptyset$. Since each child of $\hat{R}$ needs only one token, we can move extra tokens from the children to $\hat{R}$. The number of tokens $\hat{R}$ obtains is at least

$$|\mathcal{W}| + \sum_{\hat{X} \in \mathcal{B}} (1 + \gamma - |\Gamma(\hat{X}) \cap C|). \tag{5.1}$$

Let $\hat{S}$ be an arbitrary biset in $\mathcal{B}$. A node $v \in \Gamma(\hat{S}) \cap C$ is either in R or $\Gamma(\hat{R})$. If v is in R, then we are done because R is a strictly black biset that owns v. Therefore assume that each $v \in \Gamma(\hat{S}) \cap C$ is in $\Gamma(\hat{R})$. This means that $\Gamma(\hat{S}) \cap C \subseteq \Gamma(\hat{R}) \cap C$, and hence $|\Gamma(\hat{S}) \cap C| \leq |\Gamma(\hat{R}) \cap C|$. Hence (5.1) is at least the required number of tokens if $|\mathcal{W}| \geq 1$, if $|\mathcal{B}| \geq 2$, or if $\Gamma(\hat{S}) \cap C \subset \Gamma(\hat{R}) \cap C$.

Let $|\mathcal{B}| = 1$, $|\mathcal{W}| = 0$ and $\Gamma(\hat{S}) \cap C = \Gamma(\hat{R}) \cap C$. Notice that $\hat{S}$ is the only child of $\hat{R}$ in this case. It suffices to find one more token for $\hat{R}$. The linear independence between $\chi_E(\hat{R})$ and $\chi_E(\hat{S})$ implies that at least one of E_R^+ and E_R^- is not empty.

Let $e \in E_R^+$. Then e has an end-node v in $R \setminus S$. If $v \in C$, then $\hat{R}$ is a strictly black biset that owns v, and hence $\hat{R}$ has the required number of tokens in this case as mentioned above. If $v \notin C$, then e gives a token to $\hat{R}$. Thus we are done when $E_R^+ \neq \emptyset$.

Let $e' = u'v' \in E_R^-$. Then e' has an end-node, say u', in $R^+ \setminus S^+$. If $u' \in C$, then $\Gamma(\hat{S}) \cap C = \Gamma(\hat{R}) \cap C$ implies that $u' \in R$, and hence $\hat{R}$ is a strictly black biset. Let $u' \notin C$. Then $v' \in C$. $\mathcal{L}$ has no biset $\hat{Z}$ with $e' \in \delta_E(\hat{Z})$ and $u' \in Z$ by the strong laminarity of $\mathcal{L}$. $e' \in \delta_E(\hat{S})$, $v' \in S \cap C$, and $\hat{R}$ is the minimal biset such that $\hat{S} \subset \hat{R}$ and $u' \in R^+$. Hence e' gives 1 token to $\hat{R}$ in this case, which completes the proof. □

6. Proof of Theorem 1.3. We need to describe the algorithm of [35] for Degree Bounded k-Connected Subgraph. The algorithm uses the following procedure due to Khuller and Raghavachari [24], that is also used in the next section.

Procedure External k-Out-connectivity
Input: A graph $G = (V, E)$, an integer k, and $R \subseteq V$ with $|R| = k$.
Output: A subgraph J of G.
Step 1: Let G' be obtained from G by adding a new node s and all edges between s and R, of cost zero each.
Step 2: Compute a k-outconnected from s spanning subgraph J' of G'.
Step 3: Return $J = (J' \setminus \{s\})$.

Assume that Degree Bounded k-Out-connected Subgraph admits an $(\alpha, \beta(b(v)))$-approximation algorithm. For undirected graphs, the algorithm of [35] is as follows.

Algorithm Degree Bounded k-Connectivity
Step 1: Apply Procedure External k-Out-connectivity, where J' is computed using the $(\alpha, \beta(b(v)))$-approximation algorithm for Degree Bounded k-Out-connected Subgraph with degree bounds $b'(v) = b(v) + 1$ if $v \in R$ and $b'(v) = b(v)$ otherwise.
Step 2: Let F be a set of edges on V such that $J \cup F$ is k-connected.
Step 3: For every $ut \in F$ compute a minimum-cost inclusion-minimal edge-set $I_{ut} \subseteq E \setminus J$ such that $J \cup I_{ut}$ contains k internally disjoint ut-paths.
Step 4: Return $J \cup I$, where $I = \cup_{ut \in F} I_{ut}$.

In the case of directed graphs, Procedure External k-Out-connectivity computes a subgraph $J' = J^- \cup J^+$, where: J^+ is k-outconnected from s and is computed by the $(\alpha, \beta(b(v)))$-approximation algorithm, while J^- is a minimum cost subgraph which is k-inconnected to s, namely, $\kappa_{J^-}(v, s) \geq k$ for each $v \in V \setminus \{s\}$.

Lemma 6.1 ([35]). *Algorithm* Degree Bounded k-Connectivity *has ratio* $(\alpha+|F|, \beta(b(v))+2|F|+kd/2)$ *for undirected graphs, and* $(\alpha+1+|F|, \beta(b(v))+k+|F|+kd/2)$ *for digraphs, where F is the edge set computed at Step 2,* $d = \max_{v \in V} |\delta_F(v)|$ *for undirected graphs, and* $d = \max_{v \in V} |\delta_F^+(v)|$ *for digraphs.*

In [19, 35] it is shown that at Step 2 there exists an edge set F on R such that $(J' \setminus \{s\}) \cup F$ is k-connected, and if F is an inclusion minimal such edge set, then: in the undirected case, F is a forest, and in the directed case, F is a forest in the corresponding bipartite graph. The latter property is a known consequence from the undirected and directed Critical Cycle Theorems of Mader [29, 30]. Hence we can find in polynomial time F with $|F| \leq |R| - 1 = k - 1$ in the undirected case, and $|F| \leq 2|R| - 1 = 2k - 1$ in the directed case. We improve this by showing that F as above can be converted in polynomial time into an edge set F' such that $|F'| = |F|$, $\max_{v \in V} |\delta_{F'}(v)| = O(\sqrt{k})$ for undirected graphs, and $\max_{v \in V} |\delta_{F'}^+(v)| = O(\sqrt{k})$ for digraphs.

6.1. Undirected graphs. We start by proving the following.

Lemma 6.2. *Let $G = (V, E \cup F)$ be a simple k-connected undirected graph such that $|\delta_E(v)| \geq p$ for all $v \in V$ and such that every edge in F is critical for k-connectivity of G. Let $d = \max\limits_{v \in V} |\delta_F(v)| \geq 4$, let $u \in V$ with $|\delta_F(u)| = d$, and let $e = ut \in \delta_F(u)$. Suppose that for every $v \in V$ with $|\delta_F(v)| \leq d - 2$ the graph $G \setminus \{ut\} \cup \{vt\}$ is not k-connected. Then $d(d + p - k - 2) \leq 3p - k + 1$.*

Proof. Since every edge in F is critical for k-connectivity of G, F is a forest, by Mader's Critical Cycle Theorem for undirected graphs. Consider the graph $G \setminus \{e\}$ and the biset family

$$\mathcal{F} = \{\hat{S} \in \mathcal{V} \colon u \in S, t \in V \setminus S^+, |\Gamma(\hat{S})| = k - 1, \delta_{E \cup F}(\hat{S}) = \{e\}\} \ .$$

It is known that $\mathcal{F}$ is a ring biset family, namely, that the intersection of the inner parts of the members of $\mathcal{F}$ is non-empty, and $\hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y} \in \mathcal{F}$ for any $\hat{X}, \hat{Y} \in \mathcal{F}$. This implies that $\mathcal{F}$ has a unique inclusion minimal member $\hat{S}$, and that for every $v \in S$ the graph $G \setminus \{ut\} \cup \{vt\}$ is k-connected. Thus $|\delta_F(v)| \geq d-1$ for every $v \in S$, implying that $|\delta_{E \cup F}(v)| \geq p+d-1$.

Let $K = \Gamma(\hat{S})$, so $|K| = k-1$. Let I be the set of edges in F with at least one end-node in S. Every edge in $I \setminus \{e\}$ has its both end-nodes in $S \cup K$. In G, every $v \in S \setminus \{u\}$ has at least $p+d-1$ neighbors in $(S \setminus \{v\}) \cup K$, implying $(|S|-1)+(k-1) \geq p+d-1$, so $p+d-k+1 \leq |S|$. Since I is a forest on a set $S \cup K \cup \{t\}$ of $|S|+k$ nodes, $|I| \leq |S|+k-1$. Let $\zeta_I(S)$ be the set of edges in I with both end-nodes in S, so I is a disjoint union of $\delta_I(S)$ and $\zeta_I(S)$. Hence $|\delta_I(S)| + |\zeta_I(S)| = |I| \leq |S| + k - 1$. On the other hand, $(d-1)|S| \leq \sum_{v \in S} |\delta_I(v)| = |\delta_I(S)| + 2|\zeta_I(S)|$. Summarizing, we have the following:

$$p + d - k + 1 \leq |S|, \tag{6.1}$$

$$|\delta_I(S)| + |\zeta_I(S)| \leq |S| + k - 1, \tag{6.2}$$

$$|\delta_I(S)| + 2|\zeta_I(S)| \geq (d-1)|S|. \tag{6.3}$$

Subtracting (6.2) from (6.3) gives $|\zeta_I(S)| \geq (d-2)|S| - k + 1$ and thus $|\delta_I(S)| \leq 2k - 2 - (d-3)|S|$. Since $|\delta_I(S)| \geq 0$ we get $(d-3)|S| \leq 2k-2$. Combining with (6.1) we get

$$p + d - k + 1 \leq |S| \leq \frac{2k-2}{d-3}.$$

Multiplying by $d-3$ and rearranging terms we obtain $d(d+p-k-2) \leq 3p-k+1$, as claimed. □

COROLLARY 6.3. *Let $G = (V, E \cup F)$ be a simple k-connected undirected graph such that $|\delta_E(v)| \geq k-1$ for all $v \in V$. Then there exists a polynomial time algorithm that finds a set F' of edges on V with $|F'| \leq |F|$ such that $G' = (V, E \cup F')$ is k-connected and such that $|\delta_{F'}(v)| \leq \max\left\{3, \frac{3}{2} + \sqrt{2k + \frac{1}{4}}\right\}$ for all $v \in V$.*

Proof. Let $u \in V$ be a node that maximizes $|\delta_F(u)|$. Lemma 6.2 with $p = k-1$ implies that if $|\delta_F(u)|$ is larger than the required value, then we can replace an edge $ut \in \delta_F(u)$ by another edge vt such that $|\delta_F(v)|$ is at most $|\delta_F(u)| - 2$, keeping the graph being k-connected. By repeating this replacement, we can obtain a required edge set F'. □

The undirected part of Theorem 1.3 follows from Lemma 6.1, Corollary 6.3, and our ability to find in polynomial time an edge set F with $|F| \leq k-1$ at Step 3 of the algorithm.

6.2. Digraphs. We start by proving the directed counterpart of Lemma 6.2.

LEMMA 6.4. *Let $G = (V, E \cup F)$ be a simple k-connected digraph such that $|\delta^+_E(v)| \geq p$ for all $v \in V$, and suppose that every arc in F is critical for k-connectivity of G. Let $d = \max_{v \in V} |\delta^+_F(v)| \geq 4$, let $u \in V$ with $|\delta^+_F(u)| = d$ and let $e = ut \in \delta^+_F(u)$. Suppose that for every $v \in V$ with $|\delta^+_F(v)| \leq d-2$ the graph $G \setminus \{e\} \cup \{vt\}$ is not k-connected. Then $d(d+p-k-2) \leq 3p-k+2$.*

Proof. Since every arc in F is critical for k-connectivity of G, F has no alternating cycle, and thus F is a forest in the corresponding bipartite graph, by Mader's Critical Cycle Theorem for digraphs. Consider the graph $G \setminus \{e\}$ and the biset family

$$\mathcal{F} = \{\hat{S} \in \mathcal{V} \colon u \in S, t \in V \setminus S^+, |\Gamma(\hat{S})| = k-1, \delta^+_{E \cup F}(\hat{S}) = \{e\}\}\ .$$

It is known that $\mathcal{F}$ is a ring biset family, so $\mathcal{F}$ has a unique inclusion minimal member $\hat{S}$, and that for every $v \in S$ the graph $G \setminus \{ut\} \cup \{vt\}$ is k-connected. Thus $|\delta_F^+(v)| \geq d-1$ for every $v \in S$, implying that $|\delta_{E\cup F}^+(v)| \geq p+d-1$ for every $v \in S$.

Let $K = \Gamma(\hat{S})$, so $|K| = k-1$. Let I be the set of arcs in F with tail in S. Every edge in $I \setminus \{e\}$ has its head in $S \cup K$. In G, every $v \in S \setminus \{u\}$ has at least $p+d-1$ neighbors in $(S \setminus \{v\}) \cup K$, implying $(|S|-1)+(k-1) \geq p+d-1$, so $d-k+p+1 \leq |S|$. Since I is an arc set without alternating cycle on a set $S \cup K \cup \{t\}$ of $|S|+k$ nodes, $|I| \leq 2(|S|+k)-1$. On the other hand, $(d-1)|S| \leq \sum_{v\in S} |\delta_I^+(v)| = |I|$ Summarizing, we have the following:

$$d-k+p+1 \leq |S|, \tag{6.4}$$

$$|I| \leq 2(|S|+k)-1, \tag{6.5}$$

$$|I| \geq (d-1)|S|. \tag{6.6}$$

From (6.5) and (6.6) we get $(d-1)|S| \leq 2(|S|+k)-1$ so $|S|(d-3) \leq 2k-1$. Combining with (6.4) we get

$$d-k+p+1 \leq |S| \leq \frac{2k-1}{d-3}.$$

Multiplying by $d-3$ and rearranging terms we obtain $d(d+p-k-2) \leq 3p-k+2$, as claimed. □

COROLLARY 6.5. *Let $G = (V, E \cup F)$ be a simple k-connected digraph such that $\min_{v\in V} |\delta_E^-(v)| \geq k-1$ and $\min_{v\in V} |\delta_E^+(v)| \geq k-1$. Then there exists a polynomial time algorithm that finds a set F' of arcs on V with $|F'| \leq |F|$ such that $G' = (V, E \cup F')$ is k-connected and such that $|\delta_{F'}^+(v)|$ and $|\delta_{F'}^-(v)|$ are both at most $\max\left\{3, 1.5+\sqrt{2k+1.25}\right\}$ for all $v \in V$.*

Proof. Let $u \in V$ be a node that maximizes $|\delta_F^+(u)|$. Lemma 6.4 with $p = k-1$ implies that if $|\delta_F^+(u)|$ is larger than the required value, then we can replace an edge $ut \in \delta_F^+(u)$ by another edge vt such that $|\delta_F^+(v)|$ is at most $|\delta_F^+(u)|-2$, keeping the graph being k-connected. By repeating this replacement, we can obtain F'' that satisfies the conditions on connectivity and out-degree. Notice that $|\delta_{F''}^-(v)| = |\delta_F^-(v)|$ for all $v \in V$. Similarly we can decrease the in-degree of a node in V if it is larger than the required value, by applying Lemma 6.4 to the graph obtained by reversing the directions of all arcs. This gives the required edge set F'. □

The directed part of Theorem 1.3 follows from Lemma 6.1, Corollary 6.5, and our ability to find in polynomial time an edge set F with $|F| \leq 2k-1$ at Step 3 of the algorithm.

7. Proof of Theorem 1.4. Define a biset function $f^k \colon \mathcal{V} \to \mathbb{Z}$ as

$$f^k(\hat{S}) = \begin{cases} k - |\Gamma(\hat{S})| & \text{if } S \neq \emptyset \text{ and } S^+ \neq V \\ 0 & \text{otherwise.} \end{cases}$$

By the node-connectivity version of Menger's Theorem, an undirected graph (V, F) is k-connected if and only if $|\delta_F(\hat{S})| \geq f^k(\hat{S})$ for each $\hat{S} \in \mathcal{V}$. Now suppose that our goal is to augment a given graph J by a minimum-cost edge set F such that $J \cup F$ is k-connected. A natural LP relaxation for this problem is as follows (see [12]).

$$\tau^* = \min\left\{\sum_{e\in E} c(e)x(e) \mid x(\delta_E(\hat{S})) \geq f_J(\hat{S}), 0 \leq x(e) \leq 1\right\} \tag{7.1}$$

where $f_J(\hat{S}) = f^k(\hat{S}) - |\delta_J(\hat{S})|$ is the residual biset function of f^k. We will denote $\mathcal{S}_J = \{\hat{S} \mid f_J(\hat{S}) > 0\}$. Recall also that we denote $\gamma = \max_{f(\hat{S})>0} |\Gamma(\hat{S})|$, and note that $\gamma \leq k-1$ for $f = f_J$. In what follows, we assume that $k \geq 2$.

Two bisets $\hat{X}, \hat{Y}$ *cross* if $X \cap Y \neq \emptyset$ and $X^+ \cup Y^+ \neq V$, and *posi-cross* if $X \setminus Y^+ \neq \emptyset$ and $Y \setminus X^+ \neq \emptyset$. A biset function f is *crossing supermodular* if any $\hat{X}, \hat{Y} \in \mathcal{V}$ that cross satisfy the supermodular inequality (2.1). f is *symmetric* if $f(S, S^+) = f(V \setminus S^+, V \setminus S)$ for any biset $\hat{S} = (S, S^+) \in \mathcal{V}$. It is known that the function f_J is crossing supermodular and symmetric for any edge set J.

A biset function f is *positively skew supermodular* if any $\hat{X}, \hat{Y}$ with $f(\hat{X}), f(\hat{Y}) > 0$ satisfy the supermodular inequality (2.1) or the posimodular inequality (2.2). Note that we add the prefix "positively" if we enforce the requirement only when the two participating bisets obey $f(\hat{X}), f(\hat{Y}) > 0$. However, it is straightforward to verify that the 2-approximation algorithm of Fleischer, Jain and Williamson [11] for f-Connected Subgraph, as well as our algorithm for Degree Bounded f-Connected Subgraph in Theorem 2.2, apply for positively skew supermodular biset functions as well.

A biset family $\mathcal{F}$ is *independence-free* if any $\hat{X}, \hat{Y} \in \mathcal{F}$ cross or posi-cross. We say that a biset function f is independence-free if the family $\{\hat{S} \in \mathcal{V} \mid f(\hat{S}) > 0\}$ is independence free, and that a graph J is independence-free if the biset function f_J is independence free (namely, if the family $\mathcal{S}_J$ is independence free). It was observed by Jackson and Jordán [18] that if f is crossing supermodular, symmetric, and independence-free, then the biset function defined by $f'(\hat{S}) = \max\{f(\hat{S}), 0\}$ is positively skew supermodular. Clearly, a graph (V, F) is f-connected if and only if it is f'-connected. Hence in this case we can apply the iterative rounding algorithm of [11] to compute an edge set $F \subseteq E$ with $c(F) \leq 2\tau^*$ such that $J \cup F$ is k-connected. The idea of the algorithm of Cheriyan and Vegh [7] is to find a "cheap" graph J such that f_J is independence free. More precisely, the algorithm is based on the following two statements.

LEMMA 7.1 ([24]). *Let J' be an undirected graph such that J' is k-outconnected from some node s, let R be the set of neighbors of s in J', and let $J = J' \setminus \{s\}$. Then $S \cap R \neq \emptyset$ for any $\hat{S} \in \mathcal{S}_J$.*

LEMMA 7.2 ([7]). *Let J' be an undirected graph such that J' is k-outconnected from some node s, let R be the set of neighbors of s, and let $J = J' \setminus \{s\}$. Let $U = \bigcup\{S \mid \hat{S} \in \mathcal{S}_J, |S| \leq k-1\}$. Then $|U| \leq |R|k^2(k-1)$. Furthermore, if $|V| \geq |U| + k$, then there exists a polynomial time algorithm that given an edge set E on V with costs returns one of the following:*

(i) *An edge set $F \subseteq E$ with $c(F) \leq 2\tau^*$ such that $J \cup F$ is k-connected.*

(ii) *The set U.*

Frank and Tardos [15] gave a polynomial-time algorithm for computing a subgraph that is spanning k-outconnected from a root node in a directed graph. This implies a 2-approximation algorithm for the same problem in undirected graphs [24]. Furthermore, if Procedure EXTERNAL k-OUT-CONNECTIVITY from the beginning of Section 6 uses this 2-approximation algorithm, then it computes a subgraph of cost at most $2\tau^*$.

The algorithm of Cheriyan and Vegh has four steps. At every step, a certain edge set of cost at most $2\tau^*$ is computed. If the algorithm terminates at Step 2, then it returns the union of the edge-sets computed at Steps 1 and 2, of overall cost at most $4\tau^*$. Else, the algorithm returns the union of the edge-sets computed at Steps 1,3, and 4, of overall cost at most $6\tau^*$.

Algorithm of Cheriyan and Vegh

Step 1: Compute a subgraph J_{CV} of G by applying Procedure EXTERNAL k-OUT-CONNECTIVITY for some $R \subseteq V$ with $|R| = k$.

Step 2: Apply the algorithm from Lemma 7.2. If the algorithm return an edge set F as in Lemma 7.2(i) then return $J_{CV} \cup F$ and STOP.

Step 3: If the algorithm from Lemma 7.2 return U_{CV}, then apply Procedure EXTERNAL k-OUT-CONNECTIVITY for some $R \subseteq V \setminus U_{CV}$ and add the computed edge set to J_{CV}.

Step 4: The graph J_{CV} is independence free, by Lemma 7.1. Apply the iterative rounding algorithm of [11] to compute an edge set $F \subseteq E$ such that $J_{CV} \cup F$ is k-connected.

We prove the following.

LEMMA 7.3. *Let J' be an undirected graph such that J' is k-outconnected from some node s with $|\delta_{J'}(s)| = k$ and let $J = J' \setminus \{s\}$. Let $U' = \bigcup\{S : \hat{S} \in \mathcal{S}_{J'}, |S| \leq p\}$ and $U = \bigcup\{S : \hat{S} \in \mathcal{S}_J, |S| \leq p\}$. Then $|U'| \leq pk(k-1.5)$ and $|U| \leq 2pk(k-0.5)$. In particular, for $p = k-1$,*

$$|U'| + k \leq k(k-1)(k-1.5) + k \quad \textit{and} \quad |U| + k \leq 2k(k-1)(k-0.5) + k \ .$$

Lemma 7.3 will be proved later. Now we use it to finish the proof of Theorem 1.4.

Let us consider the k-Connected Subgraph problem, without degree bounds. In the case of undirected graphs, Auletta et al. [2] gave a procedure for computing a spanning subgraph J' of G such that J' is k-outconnected from some node r, $|\delta_{J'}(r)| = k$, and $c(J') \leq 2\tau^*$ (this procedure does not apply in the degree bounded setting, if we care about the cost). We apply this procedure at Step 1, instead of the EXTERNAL k-OUT-CONNECTIVITY procedure to finally obtain a subgraph J and a set U (instead of J_{CV} and U_{CV} respectively). By Lemma 7.3, $|U'| + k \leq k(k-1)(k-1.5) + k$ for $p = k-1$. We apply Steps 2,3,4 with J' and U' instead of J_{CV} and U_{CV}, achieving the same performance for $|V| \geq k(k-1)(k-1.5) + k$.

Now let us consider the degree bounded setting. At Steps 1,3 we apply Procedure EXTERNAL k-OUT-CONNECTIVITY with degree bounds $b'(v) = b(v) + 1$ if $v \in R$ and $b'(v) = b(v)$ otherwise, using our algorithm for undirected Degree Bounded k-Out-connected Subgraph. At Steps 2,4 we use our algorithm for Degree Bounded f-Connected Subgraph with positively skew supermodular f. If $|V| \geq |U| + k$, then following [7], we can design a polynomial time algorithm that either returns the set U, or an edge set $F \subseteq E$ such that $J \cup F$ is k-connected, within the same ratio as our algorithm for Degree Bounded f-Connected Subgraph with skew supermodular f (with $\gamma = k-1$). Summarizing, we have the following corollary, that combined with Theorem 1.1 and Theorem 1.2 (the version with positively skew supermodular function f), concludes the proof of Theorem 1.4.

COROLLARY 7.4. *Suppose that for undirected graphs* Degree Bounded k-Out-connected Subgraph *admits ratio $(\alpha, \beta(b(v)))$ and* Degree Bounded f-Connected Subgraph *with positively skew supermodular f admits ratio $(\alpha', \beta'(b(v)))$. Then for $|V| \geq 2k(k-1)(k-0.5) + k$* Degree Bounded k-Connected Subgraph *admits ratio $(2\alpha + \alpha', 2\beta(b(v)) + \beta'(b(v)))$, with the correspondence $\gamma = k-1$.*

In the rest of this section we prove Lemma 7.3. Two bisets $\hat{X}, \hat{Y}$ are *strongly disjoint* if $\hat{X} \setminus \hat{Y} = \hat{X}$ or $\hat{Y} \setminus \hat{X} = \hat{Y}$ (note that this is equivalent to $\hat{X} \setminus \hat{Y} = \hat{X}$ and $\hat{Y} \setminus \hat{X} = \hat{Y}$; in particular, $X \subseteq V \setminus Y^+$ and $Y \subseteq V \setminus X^+$). Given a biset family $\mathcal{F}$ let $\nu_{\mathcal{F}}$ denote the maximum number of pairwise strongly disjoint bisets in $\mathcal{F}$.

Let us say that a biset family $\mathcal{F}$ is *weakly posi-uncrossable* if for any $\hat{X}, \hat{Y} \in \mathcal{F}$ with $X \setminus Y^+, Y \setminus X^+ \neq \emptyset$, one of the bisets $\hat{X} \setminus \hat{Y}, \hat{Y} \setminus \hat{X}$ is in $\mathcal{F}$. If f is crossing supermodular and symmetric then the biset family $\mathcal{F} = \{\hat{S} \mid f(\hat{S}) > 0\}$ is weakly posi-uncrossable. This is because any $\hat{X}, \hat{Y} \in \mathcal{F}$ with $X \setminus Y^+, Y \setminus X^+ \neq \emptyset$ satisfy $f(\hat{X} \setminus \hat{Y}) + f(\hat{Y} \setminus \hat{X}) \geq f(V \setminus X^+, V \setminus X) + f(Y) = f(\hat{X}) + f(\hat{Y}) > 0$, and hence $f(\hat{X} \setminus \hat{Y}) > 0$ or $f(\hat{Y} \setminus \hat{X}) > 0$ holds.

Lemma 7.5. *Let $\mathcal{F}$ be a weakly posi-uncrossable biset family. Denote $p = \max_{\hat{S} \in \mathcal{F}} |S|$ and $\gamma = \max_{\hat{S} \in \mathcal{F}} |\Gamma(\hat{S})|$. Then $\left|\bigcup_{\hat{S} \in \mathcal{F}} S\right| \leq p(2\gamma + 1)\nu_{\mathcal{F}}$.*

Proof. Let $\mathcal{F}' = \{\hat{S}_1, \hat{S}_2, \ldots, \hat{S}_\ell\}$ be a minimum size sub-family of $\mathcal{F}$ such that $\bigcup_{i=1}^{\ell} S_i = \bigcup_{\hat{S} \in \mathcal{F}} S$. We prove that $|\mathcal{F}'| \leq (2\gamma + 1)\nu_{\mathcal{F}}$. For every $\hat{S}_i \in \mathcal{F}'$ there is $v_i \in S_i$ such that $v_i \notin S_j$ for every $j \neq i$. Among all bisets in $\mathcal{F}$ contained in $\hat{S}_i$ for which the inner part contains v_i, let $\hat{C}_i$ be an inclusion minimal one. Since $\mathcal{F}$ is weakly posi-uncrossable, the minimality of $\hat{C}_i$ implies that one of the following must hold for any distinct $\hat{C}_i, \hat{C}_j$:

- $v_i \in \Gamma(\hat{C}_j)$ or $v_j \in \Gamma(\hat{C}_i)$;
- $\hat{C}_i = \hat{C}_i \setminus \hat{C}_j$ or $\hat{C}_j = \hat{C}_j \setminus \hat{C}_i$, namely, $\hat{C}_i, \hat{C}_j$ are strongly disjoint.

Construct an auxiliary directed graph $\mathcal{J}$ on node set $\mathcal{C} = \{\hat{C}_1, \hat{C}_2, \ldots, \hat{C}_\ell\}$. Add an arc $\hat{C}_i\hat{C}_j$ if $v_i \in \Gamma(\hat{C}_j)$. The in-degree in $\mathcal{J}$ of a node $\hat{C}_i$ is at most $|\Gamma(\hat{C}_i)| \leq \gamma$. This implies that every subgraph of the underlying graph of $\mathcal{J}$ has a node of degree $\leq 2\gamma$. A graph is d-degenerate if every subgraph of it has a node of degree $\leq d$. It is known that any d-degenerate graph is $(d+1)$-colorable. Hence $\mathcal{J}$ is $(2\gamma+1)$-colorable, so its node set can be partitioned into $2\gamma + 1$ independent sets. The members of each independent set are pairwise strongly disjoint, hence their number is at most $\nu_{\mathcal{C}}$. Consequently, $\ell \leq (2\gamma + 1)\nu_{\mathcal{C}} \leq (2\gamma + 1)\nu_{\mathcal{F}}$, as claimed. ▯

The following version of Lemma 7.1 is proved in [2]; we provide a proof-sketch for completeness of exposition.

Lemma 7.6. *Let J' be an undirected graph such that J' is k-outconnected from some node s and let R be the set of neighbors of s in J'. Then $s \in \Gamma(\hat{S})$ and $|S \cap R| \geq 2$ for any $\hat{S} \in \mathcal{S}_{J'}$. Hence $\nu_{\mathcal{S}_{J'}} \leq \lfloor |R|/2 \rfloor$.*

Proof. Let $\hat{S} \in \mathcal{S}_{J'}$. If $s \notin \Gamma(\hat{S})$ then $s \in S$ or $s \in V \setminus S^+$. Since $S \neq \emptyset$ and $S^+ \neq V$, and since J' is k-outconnected from s, we easily obtain a contradiction to Menger's Theorem. We prove that $|S \cap R| \geq 2$. Let $v \in S$ and let $\ell = |\delta_{J'}(\hat{S})| + |\Gamma(\hat{S})| \leq k - 1$. Consider a set of k internally disjoint paths from s to v and the set $C = \delta_{J'}(\hat{S}) \cup \Gamma(\hat{S})$. At most $|S \cap R|$ of these paths may not contain a member in C. This implies that every one of the other at least $k - |S \cap R|$ of these paths must contain each at least one element from $C \setminus \{s\}$. Hence $\ell - 1 \geq k - |S \cap R|$. This implies $|S \cap R| \geq k - (\ell - 1) \geq 2$. ▯

Now let us get back to the proof of Lemma 7.3. Let $\mathcal{S}_J^p = \{\hat{S} \colon \hat{S} \in \mathcal{S}_J, |S| \leq p\}$. Since $\mathcal{S}_{J'}, \mathcal{S}_J$ are weakly posi-uncrossable, so are $\mathcal{S}_{J'}^p, \mathcal{S}_J^p$. By Lemmas 7.6 and 7.1, $\nu_{\mathcal{S}_{J'}^p} \leq \lfloor \frac{k}{2} \rfloor$ and $\nu_{\mathcal{S}_J^p} \leq k$. Furthermore, in the setting of Lemma 7.5 we may assume that $\gamma \leq k - 2$ for $\mathcal{S}_{J'}^p$, since by Lemma 7.6 $s \in \Gamma(\hat{S})$ for any $\hat{S} \in \mathcal{S}_{J'}$, so we can apply Lemma 7.5 after removing s from the boundary of every biset in $\mathcal{S}_{J'}^p$. Hence by Lemma 7.5 we have

$$
\begin{aligned}
|U'| &\leq p(2(k-2)+1)\nu_{\mathcal{S}_{J'}^p} \leq p(2k-3)\lfloor k/2 \rfloor \leq pk(k-1.5), \\
|U| &\leq p(2(k-1)+1)\nu_{\mathcal{S}_J^p} \leq p(2k-1)k = 2pk(k-0.5).
\end{aligned}
$$

This concludes the proof of Lemma 7.3, and thus also the proof of Theorem 1.4 is complete.

8. Proof of Theorem 1.5. Here we prove Theorem 1.5, stating that Survivable Network on undirected graphs admits the following approximation ratios for any integer $\alpha \geq 1$.

(i) $O(k^3 \log |T|) \cdot (\alpha, \alpha b(v) + k/\alpha)$ for Degree Bounded Node-Connectivity Survivable Network.
(ii) $O(k \log k) \cdot (\alpha, \alpha b(v) + k/\alpha)$ for Degree Bounded Rooted Survivable Network.
(iii) $\frac{1}{\epsilon} O(k \log^2 k) \cdot (\alpha, \alpha b(v) + k/\alpha)$ for Degree Bounded Subset k-Connected Subgraph with $k \leq (1-\epsilon)|T|$ and $0 < \epsilon < 1$.

Part (i) follows from Theorem 1.2 and the decomposition of Node-Connectivity Survivable Network into $O(k^3 \log |T|)$ instances of Element-Connectivity Survivable Network due to Chuzhoy and Khanna [8].

For proving (ii), we need to explain the algorithm of [33] for Rooted Survivable Network without degree-bounds. Augmenting version denotes instances of the problem in which G contains a subgraph J of zero edge cost such that $\kappa(s,v) \geq r(s,v) - 1$ for every $v \in T$. In [33] it is shown that the augmentation version can be decomposed into $O(k)$ instances of Degree Bounded f-Connected Subgraph with skew supermodular f. The algorithm for the general version has k iterations. At iteration ℓ, one adds to J an edge set that increases the connectivity by one for each node v such that $\kappa(s,v) = r(s,v) - k + \ell - 1$. After iteration ℓ we have $\kappa(s,v) \geq r(s,v) - k + \ell$, hence after k iterations the solution becomes feasible. In [33] it is shown that if the augmentation version admits an algorithm that computes a solution of cost at most α times the optimal value of the corresponding biset LP relaxation, then the general version admits ratio $O(\alpha \log k)$. This is because if x is a feasible solution to LP relaxation derived from an instance of Rooted Survivable Network, then $\frac{x}{k-\ell+1}$ is feasible to the LP relaxation derived from the augmentation version. For the case with degree-bounds, we act in the same way. When we solve an augmentation version instance, the degree-bounds b' is defined by $b'(v) = \lceil \frac{b(v)}{k-\ell+1} \rceil$ for $v \in B$. Then we can claim that if there exists an $(\alpha, \beta(b(v)))$-approximation algorithm for Degree Bounded f-Connected Subgraph with skew supermodular f, then Degree Bounded Rooted Survivable Network admits ratio $O(k \log k) \cdot (\alpha, \beta(b(v)))$. This and Theorem 2.2 prove (ii).

We prove (iii). In the augmentation version of Subset k-Connected Subgraph, the goal is to increase the connectivity between the terminals from $k-1$ to k, namely, G contains a subgraph J of zero edge cost such that $\kappa(u,v) \geq k-1$ for all $u, v \in T$. We use a result of [34] that the augmentation version of Subset k-Connected Subgraph with $k \leq (1-\epsilon)|T|$ is decomposed into $\frac{1}{\epsilon} O(\log k)$ instances of augmentation versions of Rooted Survivable Network with $r(s,v) = k$ for all $v \in T$. To solve the general version of Subset k-Connected Subgraph, we repeatedly solve k augmentation versions, at iteration ℓ increasing the connectivity between the nodes in T from $\ell - 1$ to ℓ. As in the rooted case, if x is a feasible solution to LP relaxation derived from an instance of Subset k-Connected Subgraph, then $\frac{x}{k-\ell+1}$ is feasible to the LP relaxation derived from the augmentation version. Hence if the augmentation version admits an algorithm that computes a solution of cost at most α times the optimal value of the corresponding biset LP relaxation, then the general version admits ratio $O(\alpha \log k)$. This extends to the degree bounded setting, if at iteration ℓ we scale the degree bounds to $b'(v) = \lceil \frac{b(v)}{k-\ell+1} \rceil$ for each $v \in B$. Then we can claim that if the augmentation version of Degree Bounded Rooted Survivable Network admits ratio $(\alpha, \beta(b(v)))$ then Degree Bounded Subset k-Connected Subgraph admits ratio $\frac{1}{\epsilon} O(\log^2 k) \cdot (\alpha, \beta(b(v)))$. By [33] the augmentation version of Rooted Survivable Network can be decomposed into $O(k)$ instances of f-Connected Subgraph with skew supermodular f, and this also extends to the degree bounded setting. Overall, we obtain that Degree Bounded f-Connected Subgraph with skew supermodular f admits ratio $(\alpha, \beta(b(v)))$ then Degree Bounded Subset k-Connected Subgraph admits ratio $\frac{1}{\epsilon}O(k \log^2 k) \cdot (\alpha, \beta(b(v)))$. This and Theorem 2.2 prove (iii).

Part (iii) does not mention the case $k > (1-\epsilon)|T|$. In this case, compute a minimum cost set of k internally disjoint (u, v)-paths for each pair of $u, v \in T$, and define a solution as the union of these paths. Note that the k internally disjoint (u, v)-paths can be computed by a minimum cost flow algorithms. This solution has the edge cost at most $O(k^2)$ times the optimal, and the degree of each node is at most $O(k^2)$ because $|T| = O(k)$.

9. Conclusion. We have presented iterative rounding algorithms and decomposition results for various Degree Bounded Survivable Network problems. We introduced several novel ideas in the field, which may be applicable also to Node-Connectivity Survivable Network problems without degree bounds. We believe that this is an important future work.

Acknowledgments. A part of this work was done when the third author was a visiting professor at RIMS, Kyoto University in Spring 2011 and when the first author was visiting Carnegie Mellon University in 2011-12, supported by Kyoto University Foundation. The first author was also supported by Japan Society for the Promotion of Science (JSPS), Grants-in-Aid for Young Scientists (B) 25730008. The third author was supported in part by NSF grants CCF 1143998 and CCF 1218382.

REFERENCES

[1] A. Aazami, J. Cheriyan, B. Laekhanukit. A bad example for the iterative rounding method for mincost k-connected spanning subgraphs. Discrete Optimization 10(1):25–41 (2013).

[2] V. Auletta, Y. Dinitz, Z. Nutov, D. Parente. A 2-approximation algorithm for finding an optimum 3-vertex-connected spanning subgraph. J. Algorithms 32(1):21–30 (1999).

[3] N. Bansal, R. Khandekar, J. Könemann, V. Nagarajan, B. Peis. On generalizations of network design problems with degree bounds. Math. Prog. 141(1-2):479–506 (2013).

[4] N. Bansal, R. Khandekar, V. Nagarajan. Additive guarantees for degree-bounded directed network design. SIAM J. Comput. 39 (2009) 1413–1431.

[5] J. Cheriyan, S. Vempala, A. Vetta. An approximation algorithm for the minimum-cost k-vertex connected subgraph. SIAM J. Comput. 32(4):1050–1055 (2003).

[6] J. Cheriyan, S. Vempala, A. Vetta. Network design via iterative rounding of setpair relaxations. Combinatorica 26(3):255–275 (2006).

[7] J. Cheriyan and L. Végh. Approximating minimum-cost k-node connected subgraphs via independence-free graphs. Proc. IEEE 54th Annual Symposium on Foundations of Computer Science, 2013, pp. 30–39.

[8] J. Chuzhoy and S. Khanna. An $O(k^3 \log n)$-approximation algorithms for vertex-connectivity survivable network design. Theory of Computing 8(1): 401–413 (2012).

[9] J. Fakcharoenphol and B. Laekhanukit. An $o(\log^2 k)$-approximation algorithm for the k-vertex connected spanning subgraph problem. SIAM J. Comput. 41(5):1095–1109 (2012).

[10] T. Feder, R. Motwani, A. Zhu. k-connected spanning subgraphs of low degree. Electronic Colloquium on Computational Complexity (ECCC) (2006).

[11] L. Fleischer, K. Jain, D.P. Williamson. Iterative rounding 2-approximation algorithms for minimum-cost vertex connectivity problems. J. Comput. Sys. Sci. 72(5):838–867 (2006).

[12] A. Frank and T. Jordán. Minimal edge-coverings of pairs of sets. J. Comb. Theory Ser. B, 65(1):73–110, 1995.

[13] A. Frank. Rooted k-connections in digraphs. Discrete Applied Math. 157(6):1242–1254, 2009.

[14] A. Frank. Connections in Combinatorial Optimization. Oxford Lecture Series in Mathematics and Its Applications 38, Oxford University Press, 2011.

[15] A. Frank and É. Tardos. An application of submodular flows. Linear Algebra and its Applications, 114:329–348, 1989.

[16] T. Fukunaga and R. Ravi. Iterative rounding approximation algorithms for degree-bounded node-connectivity network design. Proc. 53rd Annual IEEE Symposium on Foundations of Computer Science, 2012, pp. 263–272 (Longer version: CoRR abs/1203.3578 (2012)).

[17] K. Jain. A factor 2 approximation algorithm for the generalized Steiner network problem. Combinatorica 21(1):39–60 (2001).

[18] B. Jackson and T. Jordán. Independence free graphs and vertex connectivity augmentation. J. Comb. Theory, Ser. B 94(1):31–77 (2005).

[19] G. Kortsarz and Z. Nutov. Approximating node connectivity problems via set covers. Algorithmica 37(2):75–92 (2003).

[20] G. Kortsarz and Z. Nutov. Approximating k-node connected subgraphs via critical graphs. SIAM J. Comput. 35(1):247–257 (2005).

[21] R. Khandekar, G. Kortsarz, Z. Nutov. Network-design with degree constraints. Proc. 14th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, and 15th. International Workshop on Randomization and Computation, 2011, Lecture Notes in Computer Science 6845, pp. 289–301.

[22] T. Király, L.C. Lau, M. Singh. Degree bounded matroids and submodular flows. Combinatorica 32(6):703–720 (2012).

[23] G. Kortsarz and Z. Nutov. Approximating Minimum Cost Connectivity Problems. Ch. 58 In Handbook on Approximation Algorithms and Metaheuristics. Ed. T. F. Gonzalez, Chapman & Hall/CRC.

[24] S. Khuller and B. Raghavachari. Improved approximation algorithms for uniform connectivity problems. J. Algorithms 21:434–450 (1996).

[25] B. Laekhanukit. An improved approximation algorithm for minimum-cost subset k-connectivity. Proc. 38th International Colloquium on Automata, Languages and Programming, Part I, 2011, Lecture Notes in Computer Science 755, pp. 13–24.

[26] L.C. Lau, J. Naor, M.R. Salavatipour, M. Singh. Survivable network design with degree or order constraints. SIAM J. Comput. 39(3):1062–1087 (2009).

[27] L.C. Lau, R. Ravi, M. Singh. Iterative Method in Combinatorial Optimization. Cambridge University Press, 2011.

[28] L.C. Lau and M. Singh. Additive approximation for bounded degree survivable network design. Proc. 40th Annual ACM Symposium on Theory of Computing, 2008, pp. 759–768.

[29] W. Mader. Ecken vom grad n in minimalen n-fach zusammenhängenden graphen. Archive der Mathematik 23:219–224 (1972).

[30] W. Mader. Minimal n-fach in minimalen n-fach zusammenhängenden digraphen. J. Comb. Theory B 38(2):102–117 (1985).

[31] V. Nagarajan, R. Ravi, M. Singh. Simpler analysis of LP extreme points for traveling salesman and survivable network design problems. Oper. Res. Lett. 38(3):156–160 (2010).

[32] Z. Nutov. Approximating minimum-cost edge-covers of crossing biset-families. CoRR abs/1207.4366 (2012), to appear in Combinatorica. (Preliminary version: An almost $O(\log k)$-approximation for k-connected subgraphs. 20th Annual ACM-SIAM Symposium on Discrete Algorithms, 2009, pp. 912–921).

[33] Z. Nutov. Approximating minimum cost connectivity problems via uncrossable bifamilies. ACM Transactions on Algorithms 9(1):1 (2012)

[34] Z. Nutov. Approximating subset k-connectivity problems. J. Discrete Algorithms 17:51–59 (2012).

[35] Z. Nutov. Degree-constrained node-connectivity problems. Proc. 10th Latin American Theoretical Informatics Symposium, 2012, Lecture Notes in Computer Science 7256, pp. 582–593. To appear in Algorithmica.

[36] M. Singh and L.C. Lau. Approximating minimum bounded degree spanning trees to within one of optimal. Proc. 39th Annual ACM Symposium on Theory of Computing, 2007, pp. 661–670.

[37] D.P. Williamson and D.B. Shmoys. The Design of Approximation Algorithms, Cambridge University Press, 2011.