2012 IEEE 53rd Annual Symposium on Foundations of Computer Science

Iterative rounding approximation algorithms for degree-bounded node-connectivity network design

Takuro Fukunaga
Graduate School of Informatics,
Kyoto University, Japan
Email: takuro@amp.i.kyoto-u.ac.jp

R. Ravi
Tepper School of Business,
Carnegie Mellon University, USA
Email: ravi@cmu.edu

***Abstract*—We consider the problem of finding a minimum edge cost subgraph of an undirected or a directed graph satisfying given connectivity requirements and degree bounds $b(\cdot)$ on nodes. We present an iterative rounding algorithm for this problem. When the graph is undirected and the connectivity requirements are on the element-connectivity with maximum value k, our algorithm computes a solution that is an $O(k)$-approximation for the edge cost in which the degree of each node v is at most $O(k) \cdot b(v)$. We also consider the no edge cost case where the objective is to find a subgraph satisfying connectivity requirements and degree bounds. Our algorithm for this case outputs a solution in which the degree of each node v is at most $6 \cdot b(v) + O(k^2)$. These algorithms can be extended to other well-studied undirected node-connectivity requirements such as uniform, subset and rooted connectivity. When the graph is directed and the connectivity requirement is k-out-connectivity from a root, our algorithm computes a solution that is a 2-approximation for the edge cost in which the degree of each node v is at most $2 \cdot b(v) + O(k)$.**

I. Introduction

A. Problem definition

The degree-bounded survivable network design for undirected graphs is the problem of constructing a subgraph of a given undirected graph that satisfies both degree-bounds on nodes and certain connectivity requirements between nodes with minimum edge cost. More formally, the problem is defined as follows.

***Degree-bounded Survivable Network Design* (*SND*)**: An undirected graph $G = (V, E)$ with edge costs $c : E \to \mathbb{R}_+$, connectivity requirements $r : V \times V \to \mathbb{Z}_+$, and degree-bounds $b : B \to \mathbb{Z}_+$ on a subset B of V are given. Find a minimum cost $F \subseteq E$ such that the degree of $v \in B$ is at most $b(v)$ and the connectivity between $u, v \in V$ is at least $r(u, v)$ in the subgraph (V, F) of G.

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

If $B = V$, $b(v) = 2$ for all $v \in B$, and a solution is required to be a connected spanning subgraph, then the degree-bounded SND is 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 constraints on degree-bounds. We say that an algorithm is $(\alpha, \beta(b(v)))$-*approximation* for $\alpha \in \mathbb{R}_+$ and a function $\beta : \mathbb{Z}_+ \to \mathbb{R}_+$ if it always outputs a solution such that its cost is at most α times the optimal value, and the degree of each $v \in B$ is at most $\beta(b(v))$ for each instance which admits a feasible solution.

We also discuss a special case of degree-bounded SND with no edge costs that we call the degree-bounded subgraph problem.

Degree-bounded subgraph problem: Given an undirected graph $G = (V, E)$, connectivity requirements $r : V \times V \to \mathbb{Z}_+$, a subset B of V, and degree-bounds $b : B \to \mathbb{Z}_+$, find $F \subseteq E$ such that the degree of $v \in B$ is at most $b(v)$ and the connectivity between $u, v \in V$ is at least $r(u, v)$ in the subgraph (V, F) of G.

We say that an algorithm is $\beta(b(v))$-*approximation* for some function $\beta : \mathbb{Z}_+ \to \mathbb{R}_+$ if it outputs a subgraph such that the connectivity between $u, v \in V$ is at least $r(u, v)$ and the degree of $v \in B$ is at most $\beta(b(v))$ for each instance which admits a feasible solution.

Notice that the degree-bounded subgraph problem contains the problem of finding a subgraph of required connectivity minimizing the maximum degree. 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.

In this paper, we are interested in element-connectivity and node-connectivity requirements. The definition of the element-connectivity supposes 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. On the other hand, the node-connectivity $\kappa(u, v)$ between two vertices $u, v \in V$ is defined as the maximum number of (u, v)-paths that are pair-wise openly (or node) disjoint. Special cases of the above problems are defined according to the given connectivity requirements as follows.

- The degree-bounded SND is called ***element-connectivity SND*** and the degree-bounded subgraph problem is called ***element-connectivity subgraph problem*** if they demand $\lambda_T(u,v) \geq r(u,v)$ for each $u,v \in T$ where $T \subseteq V$ is a given terminal set.

- The degree-bounded SND is called ***node-connectivity SND*** and the degree-bounded subgraph problem is called ***node-connectivity subgraph problem*** if they demand $\kappa(u,v) \geq r(u,v)$ for each $u,v \in V$.
- ***Rooted k-connectivity SND*** and ***rooted k-connectivity subgraph problem*** are respectively special cases of the node-connectivity SND and the node-connectivity subgraph problem in which $r(u,v) = k$ holds if $\{u,v\} \subseteq T$ and $\{u,v\}$ contains a specified vertex s called *root*, and $r(u,v) = 0$ otherwise.
- ***Subset k-connectivity SND*** and ***subset k-connectivity subgraph problem*** are respectively special cases of the node-connectivity SND and the node-connectivity subgraph problem such that $r(u,v) = k$ if $\{u,v\} \subseteq T$, and $r(u,v) = 0$ otherwise where $T \subseteq V$ is a given terminal set.
- ***k-connectivity SND*** and ***k-connectivity subgraph problem*** are respectively special cases of the subset-connectivity SND and the subset-connectivity subgraph problem in which $T = V$.

Whereas we defined the above problems for undirected graphs, they can be defined also for digraphs. In this paper, we investigate the following problems for digraphs.

k-out-connectivity SND and directed k-connectivity SND: We are given a digraph $G = (V,E)$ with arc costs $c : E \to \mathbb{Q}_+$, in-degree-bounds $b^- : B^- \to \mathbb{Z}_+$ on $B^- \subseteq V$, and out-degree-bounds $b^+ : B^+ \to \mathbb{Z}_+$ on $B^+ \subseteq V$. In k-out-connectivity SND, the connectivity requirements demand that $\kappa(s,v) \geq k$ for each $v \in V \setminus \{s\}$ with a given root $s \in V$. In directed k-connectivity SND, the connectivity requirements demand that $\kappa(u,v) \geq k$ for each $u,v \in V$. The task is to find a minimum cost $F \subseteq E$ such that the in-degree of $v \in B^-$ is at most $b^-(v)$, the out-degree of $v \in B^+$ is at most $b^+(v)$, and the connectivity requirements are satisfied in the subgraph (V,F).

For $\alpha \in \mathbb{R}_+$ and functions $\beta, \beta' : \mathbb{Z}_+ \to \mathbb{R}_+$, an algorithm for k-out-connectivity SND or directed k-connectivity SND is called $(\alpha, \beta(b^-(v)), \beta'(b^+(v)))$*-approximation* if it outputs a solution such that its cost is at most α times the optimal value, the in-degree of each $v \in B^-$ is at most $\beta(b^-(v))$, and the out-degree of each $v \in B^+$ is at most $\beta'(b^+(v))$ for each instance which admits a feasible solution.

B. Previous work

SND without degree-bounds is a typical combinatorial optimization problem, and a large number of studies on it have been presented so far especially for the case with edge-connectivity requirements. One of the most important achievement among them is iterative rounding, that was invented in the context of a 2-approximation algorithm by Jain [9]. He showed that every basic optimal solution to an LP relaxation for the edge-connectivity SND always has a variable of value at least $1/2$. The 2-approximation algorithm is obtained by repeatedly rounding up such variables and iterating the procedure until the rounded subgraph is feasible.

The degree-bounded SND 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 [13], [14] and Singh and Lau [21]. They gave a $(2, 2b(v)+3)$-approximation for the degree-bounded edge-connectivity SND, and a $(1, b(v)+1)$-approximation algorithm for the degree-bounded spanning tree problem. The former result was improved to a $(2, b(v)+6k+3)$-approximation by Lau and Singh [16] afterwards. After their work, many efficient algorithms have been proposed for various types of degree-bounded SND such as directed network design problem [3], matroid base and submodular flow problems [11], and matroid intersection and optimization over lattice polyhedra [2]. Almost all of them deal with edge-connectivity requirements and are based on iterative rounding. See [15] for a comprehensive survey on iterative rounding.

Despite the success of iterative rounding for edge-connectivity requirements, the degree-bounded SND with element- and node-connectivity requirements still remain difficult to address with this method. The $(2, 2b(v)+3)$-approximation algorithm due to Lau, Naor, Salavatipour and Singh [13], and the $(2, b(v)+6k+3)$-approximation algorithm due to Lau and Singh [16] for the edge-connectivity SND can be extended to the element-connectivity SND but they need to assume that degree-bounds were given on terminals only. Lau, Naor, Salavatipour and Singh also showed that the subset k-connectivity subgraph problem admits no $2^{\log^{1-\epsilon} n}$-approximation algorithm for some $\epsilon > 0$ unless NP $\subseteq$ DTIME$(n^{\mathrm{polylog}(n)})$ in [14]. For the k-connectivity subgraph problem, Feder, Motwani and Zhu [7] presented an $O(k \log n)$-approximation algorithm, which runs in $n^{O(k)}$ time. Khandekar, Kortsarz and Nutov [10] proposed a $(4, 6b(v)+6)$-approximation algorithm for the k-connectivity SND with $k = 2$. Very recently, Nutov [20] applied iterative rounding for the element- and node-connectivity SNDs both for undirected graphs and digraphs. The approximation guarantees achieved by his algorithms can be found in Table I. Notice that his approximation guarantees on the degree-bounds on non-terminals for the element-connectivity SND and on arbitrary nodes for other problems are exponential on k.

As we can observe from the above, degree-bounds on non-terminals in the element-connectivity SND and those (on arbitrary nodes) in the node-connectivity SND incur large violation in previous work. One reason for this is the difficulty in dealing with node-connectivity requirements by iterative rounding directly even if we have no degree-bounds. Most approximation algorithms proposed for node-connectivity SND without degree-bounds so far are based on a decomposition approach [6], [12], [17], [18], [19]. Roughly speaking, papers adopting this approach solve the problem by decomposing it into instances of problems with edge- or element-connectivity requirements, or their generalizations.

As for iterative rounding, Cheriyan, Vempala and Vetta [4] showed that it achieves $O(\sqrt{n/\epsilon})$-approximation for the k-connectivity SND where ϵ is a positive real with $k \leq (1-\epsilon)n$. Fleischer, Jain and Williamson [8] showed that iterative rounding achieves a 2-approximation for the node-connectivity SND with $k \leq 2$. Both of these results hold only for the case without degree-bounds, and use a standard LP relaxation in which connectivity requirements are formulated by bisets (or its equivalent notion set-pairs). Since the methods based on the decomposition approach in [6], [17] only achieves $O(\log k \log \frac{n}{n-k})$-approximation for the k-connectivity SND without degree-bounds, iterative rounding gives a much better result than the decomposition approaches but only for $k \leq 2$. Aazami, Cheriyan and Laekhanukit [1] presented an instance of the k-connectivity SND without degree-bounds for which the basic optimal solution to the LP relaxation has no variable of value $\Omega(\frac{1}{\sqrt{k}})$. Their instance shows that it is hard for iterative rounding to achieve an approximation factor better than $O(\sqrt{k})$ for the k-connectivity SND with general k.

C. Our results and techniques

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 $\Gamma(\hat{S})$. We represent the size of $\Gamma(\hat{S})$ by $\gamma(\hat{S})$. When E is a set of undirected edges, $\delta_E(\hat{S})$ denotes the set of edges in E joining nodes in S with those in $V \setminus S^+$. When E is a set of arcs, $\delta_E^-(\hat{S})$ denotes the set of arcs in E that have their heads in S and their tails in $V \setminus S^+$. As we will see in Section II, the element- and node-connectivity requirements can be represented by requiring

$$\sum_{e \in \delta_E(\hat{S})} x(e) \geq f(\hat{S}) \quad \text{for each biset } \hat{S} \tag{1}$$

in undirected graphs, and

$$\sum_{e \in \delta_E^-(\hat{S})} x(e) \geq f(\hat{S}) \quad \text{for each biset } \hat{S} \tag{2}$$

in digraphs where $x(e) \in \{0,1\}$ is a variable for representing whether an edge/arc $e \in E$ is chosen or not, and f is some biset function defined from the connectivity requirements.

1) Element-connectivity SND: We give an iterative rounding $(O(k), O(k) \cdot b(v))$-approximation algorithm for the element-connectivity SND. Our algorithm exploits the LP relaxation in which the connectivity requirements are formulated by (1). This is a natural extension of the LP relaxation usually used for SND with edge-connectivity requirements. In SND with edge-connectivity requirements, the analysis of iterative rounding depends on the laminarity of the tight cut constraints defining the basic solutions. On the other hand, it is known [20] that if the biset function f is defined from the element-connectivity, then the tight biset constraints defining the basic solutions to our LP has a certain type of laminarity, that will be introduced in Section IV. Despite this observation, we still have some difficulty in carrying out the standard token argument from edge-connectivity problem particularly in the case of bisets that have exactly one child in the laminar family. To tackle this, we use a careful structural lemma (Lemma 2) that bounds the number of such one-child bisets that need careful handling. This leads to our bounds for the element-connectivity SND (Corollary 1).

2) k-out-connectivity SND: We have the same laminarity of the independent tight biset constraints for the k-out-connectivity. Hence it is not difficult to see that our idea for the element-connectivity SND gives an $(O(k), O(k) \cdot b^-(v), O(k) \cdot b^+(v))$-approximation for the k-out-connectivity SND. However we find that it is possible to achieve a better result. Here the key interaction is between the size of the laminar family of independent tight constraints for the extreme point, and the size of the set of nodes whose degree-bounds are tight and are linearly independent at this extreme point. Depending upon the relative sizes of these two sets, we modify the earlier token argument to get our improved result. When the set of nodes with tight degree-bounds is smaller than the laminar family of independent tight constraints, we show that a known iterative rounding proof for the case without degree-bounds can be modified for the case with degree-bounds. When the set of nodes with tight degree-bounds is larger, we use that our structural lemma on the laminar family to give a better token distribution scheme work for the k-out-connectivity SND. By these observations, we give a $(2, k, 2b^+(v) + O(k))$-approximation algorithm for the k-out-connectivity SND.

3) Element-connectivity subgraph problem: The above token distribution idea for the k-out-connectivity SND does not work for undirected graphs. However, by introducing a more careful token redistribution idea based on more sophisticated structural lemma than the one used for the element-connectivity SND, we can see that there exists a fractional edge in the basic solution for the corresponding LP relaxation with value at least $1/6$ or a degree-bounded node of degree at most $16k^2 - 4k - 7$ in the support graph of the basic solution (Theorem 5). This achieves approximation factor $6b(v) + O(k^2)$ for the element-connectivity subgraph problem.

4) Other node-connectivity requirements: It has been shown in the previous work on the node-connectivity SND that the other connectivity requirements can be decomposed into element-connectivity or k-out-connectivity requirements. Hence out results for the element-connectivity and

Table I

APPROXIMATION GUARANTEES FOR THE DEGREE-BOUNDED SND AND THE DEGREE-BOUNDED SUBGRAPH PROBLEM

	edge cost (approx. factors)	degree	note
element-connectivity (undirected)	2	$(b(v)+6k+3,+\infty)$*	Lau, Singh [16]
	$O(\log k)$	$(O(\log k\cdot b(v)+k), O(2^k)\cdot b(v))$*	Nutov [20]
	$4k-1$	$(4k-1)\cdot b(v)+4k-2$	This paper
	$+\infty$	$6b(v)+O(k^2)$	This paper
node-connectivity (undirected)	$O(k^3\log k\log\|T\|)$	$O(2^k k^3\log\|T\|)\cdot b(v)$	Nutov [20]
	$O(k^4\log\|T\|)$	$O(k^4\log\|T\|)\cdot b(v)$	This paper
	$+\infty$	$O(k^3\log\|T\|\cdot b(v)+k^5\log\|T\|)$	This paper
rooted k-connectivity (undirected)	$O(\log k)$	$O(2^k)\cdot b(v)$	only for $T=V$, Nutov [20]
	$O(k^2\log k\log\|T\|)$	$O(2^k k^2\log\|T\|)\cdot b(v)$	Nutov [20]
	4	$2b(v)+5k-1$	only for $T=V$, This paper
	$O(k^2\log k)$	$O(k^2\log k)\cdot b(v)$	This paper
	$+\infty$	$O(k\log k\cdot b(v)+k^3\log k)$	This paper
subset k-connectivity (undirected)	$O(k^2\log k\log\|T\|)$	$O(2^k k^2\log\|T\|)\cdot b(v)$	Nutov [20]
	$O(k^2\log k)$	$O(k^2\log k)\cdot b(v)$	This paper
	$+\infty$	$O(k\log k\cdot b(v)+k^3\log k)$	This paper
k-connectivity (undirected)	$O(k)$	$O(2^k)\cdot b(v)$	Nutov [20]
	$O(k)$	$2b(v)+O(k^2)$	This paper
k-out-connectivity (directed)	$O(\log k)$	$(+\infty, O(2^k)\cdot b^+(v))$	Nutov [20]
	1	$(b^-(v),+\infty)$	Nutov [20]
	2	$(k, 2b^+(v)+4k-1)$	This paper
directed k-connectivity (directed)	$O(k)$	$(+\infty, O(2^k)\cdot b^+(v))$	Nutov [20]
	$O(k)$	$(O(k^2), 2b^+(v)+O(k^2))$	This paper

*The two-tuples denote the degrees of terminals and of non-terminals respectively.

k-out-connectivity give new approximation algorithms for these requirements. The approximation guarantees achieved by our algorithms are summarized in Table I. Approximation guarantee $\beta(b(v))$ for the degree-bounded subgraph problem is represented by $(+\infty, \beta(b(v)))$ in the table.

Although we do not give a proof due to the space limitation, if the biset function f is defined for the k-connectivity requirements and $n > 3k-3$, then the tight biset constraints defining a basic solution for our LP is laminar, which has been never observed before as far as we know. Therefore our results for the laminar bisets in undirected graphs show that applying iterative rounding directly gives an $(O(k), O(k)\cdot b(v))$-approximation for the k-connectivity SND and a $(6b(v)+O(k^2))$-approximation for the k-connectivity subgraph problem when $n > 3k-3$. This is not important in terms of approximation guarantees because we can achieve better guarantees by decomposing the k-connectivity requirement into the k-out-connectivity SND requirements. Nevertheless we believe that this fact is worth noting even for the case without degree-bounds because this $O(k)$ guarantee on the fractionality of the basic optimal solution to the LP relaxation is much better than the earlier bound $O(\sqrt{n/\epsilon})$ due to [4].

II. PRELIMINARIES ON BISETS

For a graph $G=(V,E)$, we denote the set of all bisets of V by $\mathcal{V}$. If $\hat{S}=(S,S^+)\in\mathcal{V}$ satisfies $S\neq\emptyset$ and $V\setminus S^+\neq\emptyset$, then $\hat{S}$ is called *proper*. For two bisets $\hat{X}=(X,X^+)$ and $\hat{Y}=(Y,Y^+)$, we define $\hat{X}\cap\hat{Y}$ as $(X\cap Y, X^+\cap Y^+)$, $\hat{X}\cup\hat{Y}$ as $(X\cup Y, X^+\cup Y^+)$, and $\hat{X}\setminus\hat{Y}$ as $(X\setminus Y^+, X^+\setminus Y)$.

Suppose that $G=(V,E)$ is undirected. For a biset $\hat{S}$, $\chi_E(\hat{S})$ denotes the incidence vector of $\delta_E(\hat{S})$ (i.e., $\chi_E(\hat{S})$ is the $|E|$-dimensional vector whose component corresponding to $e\in E$ is 1 if $e\in\delta_E(\hat{S})$, and 0 otherwise).

Let $g:\mathcal{V}\to\mathbb{Z}_+$. g is called *symmetric* when $g(\hat{S})=g(\hat{S}')$ holds for any $\hat{S}=(S,S^+)\in\mathcal{V}$ and $\hat{S}'=(V\setminus S^+, V\setminus S)$. g is called *intersecting supermodular* if for any bisets $\hat{X}$ and $\hat{Y}$ with $g(\hat{X})>0$, $g(\hat{Y})>0$, $X\cap Y\neq\emptyset$ and $V\setminus(X^+\cup Y^+)\neq\emptyset$, $g(\hat{X})+g(\hat{Y})\leq g(\hat{X}\cap\hat{Y})+g(\hat{X}\cup\hat{Y})$ holds, and is called *skew supermodular* if for any two bisets $\hat{X}$ and $\hat{Y}$ with $g(\hat{X})>0$ and $g(\hat{Y})>0$, we have (i) $g(\hat{X})+g(\hat{Y})\leq g(\hat{X}\cap\hat{Y})+g(\hat{X}\cup\hat{Y})$ or (ii) $g(\hat{X})+g(\hat{Y})\leq g(\hat{X}\setminus\hat{Y})+g(\hat{Y}\setminus\hat{X})$.

From $T\subseteq V$ and $r:T\times T\to\mathbb{Z}_+$, define a biset function $f_{\mathrm{elt}}:\mathcal{V}\to\mathbb{Z}$ by $f_{\mathrm{elt}}(\hat{S})=\max_{u\in S\cap T, v\in T\setminus S^+} r(u,v)-\gamma(\hat{S})$ if $S\cap T\neq\emptyset\neq T\setminus S^+$ and $T\cap\Gamma(\hat{S})=\emptyset$, and by $f_{\mathrm{elt}}(\hat{S})=0$ otherwise. By Menger's theorem, an undirected graph (V,E) satisfies $\lambda_T(u,v)\geq r(u,v)$ for each $u,v\in T$ if and only if $|\delta_E(\hat{S})|\geq f_{\mathrm{elt}}(\hat{S})$ for each $\hat{S}\in\mathcal{V}$. Thus f_{elt} represents element-connectivity requirements.

Theorem 1 ([8]): f_{elt} is symmetric and skew supermodular.

For $k\in\mathbb{Z}_+$ and $s\in V$, define a function $f_{\mathrm{out}}:\mathcal{V}\to\mathbb{Z}$ by $f_{\mathrm{out}}(\hat{S})=k-\gamma(\hat{S})$ if $S\neq\emptyset$ and $s\notin S^+$, and by $f_{\mathrm{out}}(\hat{S})=0$ otherwise. Then a digraph (V,E) satisfies $\kappa(s,v)\geq k$ for each $v\in V\setminus\{s\}$ if and only if $|\delta^-_E(\hat{S})|\geq$

$f_{\text{out}}(\hat{S})$ for each $\hat{S} \in \mathcal{V}$. In other words, f_{out} represents the k-out-connectivity requirements. It is not difficult to see the following property of f_{out}.

Theorem 2: f_{out} is intersecting supermodular.

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 $f_{\text{elt}} - |\delta_F(\cdot)|$ is skew supermodular. Similarly 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 $\hat{X}, \hat{Y} \in \mathcal{V}$. Hence $f_{\text{out}} - |\delta_F^-(\cdot)|$ is intersecting supermodular.

III. Iterative rounding algorithms

A. Degree-bounded SND with symmetric skew supermodular functions in undirected graphs

For applying iterative rounding, we define $F \subseteq E$ as the set of edges which have not been chosen by the algorithm yet. The edges in $E \setminus F$ have already been chosen as a part of the current solution by the algorithm. Let $x \in \mathbb{R}^F$ be a variable vector each component of which corresponds to an edge in F. For any $F' \subseteq F$, we let $x(F')$ denote $\sum_{e \in F'} x(e)$. Let $f : \mathcal{V} \to \mathbb{Z}_+$ be a symmetric skew supermodular function, where we define f by $f(\hat{S}) = f_{\text{elt}}(\hat{S}) - |\delta_{E \setminus F}(\hat{S})|$, $\hat{S} \in \mathcal{V}$ when we apply our algorithm to the element-connectivity SND. The LP relaxation we use is

$$
\begin{array}{ll}
\text{minimize} & \sum_{e \in F} c(e)x(e) \\
\text{subject to} & x(\delta_F(\hat{S})) \geq f(\hat{S}) \text{ for each } \hat{S} \in \mathcal{V}, \\
& x(\delta_F(v)) \leq b(v) \text{ for each } v \in B, \\
& 0 \leq x(e) \leq 1 \text{ for each } e \in F.
\end{array} \tag{3}
$$

Suppose that x^* is the basic optimal solution to (3). If the constraint corresponding to $\hat{S} \in \mathcal{V}$ (resp., $v \in B$) in (3) is tight with regard to x^*, then $\hat{S}$ (resp., v) is called *tight*. For iterative rounding to work, we need to show that the fractionality of x^* is low or there exists a node $v \in B$ of low degree in (V, F) [15].

Theorem 3: Let $G = (V, E)$ be an undirected graph G, and $f : \mathcal{V} \to \mathbb{Z}_+$ be a symmetric skew supermodular function such that $\max\{\gamma(\hat{S}) \mid f(\hat{S}) > 0\} < k$. Then there exists $e \in F$ such that $x^*(e) = 0$ or $x^*(e) \geq 1/(4k-1)$, for the basic optimal solution x^* to (3), or there exists $v \in B$ such that $|\delta_F(v)| \leq 4k - 1$.

We prove Theorem 3 in Section IV. If there exists an edge $e \in F$ such that $x^*(e) = 0$ or $x^*(e) \geq 1/\alpha$, or if there exists $v \in B$ such that $|\delta_F(v)| \leq \beta$, then a standard iterative rounding algorithm achieves $(\alpha, \alpha \cdot b(v) + \beta - 1)$-approximation for the degree-bounded SND (see [13], [14]). Hence Theorem 3 implies the following results.

Corollary 1: The degree-bounded SND admits the following approximation guarantees:

(i) $(4k-1, (4k-1) \cdot b(v) + 4k - 2)$ for the element-connectivity SND;
(ii) $(O(k^4 \log |T|), O(k^4 \log |T|) \cdot b(v))$ for the node-connectivity SND;
(iii) $(O(k^2 \log k), O(k^2 \log k) \cdot b(v))$ for the rooted k-connectivity SND;
(iv) $(O(k^2 \log k), O(k^2 \log k) \cdot b(v))$ for the subset k-connectivity SND.

Proof: (i) is immediate from the simmetry and the skew supermodularity of $f_{\text{elt}} - |\delta_{E \setminus F}(\cdot)|$. (ii) is obtained from (i) and the decomposition of the node-connectivity SND into $O(k^3 \log |T|)$ instances of the element-connectivity SND due to Chuzhoy and Khanna [5]. (iii) and (iv) also can be derived by a decomposition presented in [18], [19]. ∎

B. Degree-bounded SND with intersecting supermodular functions in digraphs

We let $f : \mathcal{V} \to \mathbb{Z}_+$ be an intersecting supermodular function. When we apply our algorithm to the k-out-connectivity SND, we define f by $f(\hat{S}) = f_{\text{out}}(\hat{S}) - |\delta^-_{E \setminus F}(\hat{S})|$ $S \in \mathcal{V}$. Our LP relaxation for this case is

$$
\begin{array}{ll}
\text{minimize} & \sum_{e \in F} c(e)x(e) \\
\text{subject to} & x(\delta_F^-(\hat{S})) \geq f(\hat{S}) \text{ for each } \hat{S} \in \mathcal{V}, \\
& x(\delta_F^+(v)) \leq b^+(v) \text{ for each } v \in B^+, \\
& 0 \leq x(e) \leq 1 \text{ for each } e \in F.
\end{array} \tag{4}
$$

Theorem 4: Let G be a digraph, $f : \mathcal{V} \to \mathbb{Z}_+$ be a intersecting supermodular function such that $\max\{\gamma(S) \mid f(S) > 0\} < k$. Then there exists an arc $e \in F$ such that $x^*(e) = 0$ or $x^*(e) \geq 1/2$, for the basic optimal solution x^* to (4), or there exists a node $v \in B^+$ such that $|\delta_F^+(v)| \leq 4k$.

We prove Theorem 4 in Section V.

Corollary 2: The degree-bounded SND admits the following approximation guarantees:

(i) $(2, k, 2b^+(v) + 4k - 1)$-approximation for the k-out-connectivity SND;
(ii) $(4, 2b(v) + 5k - 1)$-approximation for the rooted k-connectivity SND with $T = V$;
(iii) $(O(k), O(k^2), 2b^+(v) + O(k^2))$-approximation for the directed k-connectivity SND;
(iv) $(O(k), 2b(v) + O(k^2))$-approximation for the k-connectivity SND.

Proof: The guarantees on the arc costs and the out-degree in (i) is immediate from Theorem 4 and the intersecting supermodularity of $f_{\text{out}} - |\delta^-_{E \setminus F}(\cdot)|$. Nutov [20] showed that the in-degree of each node is at most k in every minimal solution for the degree-bounded SND in digraphs when the connectivity requirements are represented by intersecting supermodular functions. The guarantee on the in-degree in (i) follows from this fact.

(ii) is proven by applying (i) to the digraph obtained by replacing each undirected edge uv by arcs uv and vu with degree-bounds defined by $B^- = B^+ = B$ and $b^-(v) = b^+(v) = b(v)$ for each $v \in B$.

(iii) follows from (i) and the theorem due to Nutov [20] which shows that if the k-out-connectivity SND problem admits an $(\alpha, \beta(b^-(v)), \beta'(b^+(v)))$-approximation, then the directed spanning SND problem admits an $(\alpha +$

$O(k), \beta(b^-(v))+O(k^2), \beta'(b^+(v))+O(k^2))$-approximation. (iv) is obtained from (iii) as (ii) is obtained from (i). ■

C. Degree-bounded subgraph problem with symmetric skew supermodular functions in undirected graphs

For the degree-bounded subgraph problem, we use the following LP relaxation.

$$\begin{array}{ll} \text{maximize} & x(F) \\ \text{subject to} & x(\delta_F(\hat{S})) \geq f(\hat{S}) \text{ for each } \hat{S} \in \mathcal{V}, \\ & x(\delta_F(v)) \leq b(v) \text{ for each } v \in B, \\ & 0 \leq x(e) \leq 1 \text{ for each } e \in F. \end{array} \quad (5)$$

In this case, Theorem 3 can be improved to obtain the next theorem.

Theorem 5: Let $G = (V, E)$ be an undirected graph, and $f : \mathcal{V} \to \mathbb{Z}_+$ be a symmetric skew supermodular function such that $\max\{\gamma(\hat{S}) \mid f(\hat{S}) > 0\} < k$. Then there exists $e \in F$ such that $x^*(e) = 0$, or $x^*(e) \geq 1/6$ for the basic optimal solution x^* to (5), or there exists $v \in B$ such that $|\delta_F(v)| \leq 16k^2 - 4k - 7$.

We prove Theorem 5 in Section VI.

Corollary 3: The degree-bounded subgraph problem admits the following approximation guarantees:

(i) $6b(v) + O(k^2)$ for the element-connectivity subgraph problem;
(ii) $O(k^3 \log |T| \cdot b(v) + k^5 \log |T|)$ for the node-connectivity subgraph problem;
(iii) $O(k \log k \cdot b(v) + k^3 \log k)$ for the rooted k-connectivity subgraph problem;
(iv) $O(k \log k \cdot b(v) + k^3 \log k)$ for the subset k-connectivity subgraph problem.

Proof: The claims can be proven by applying Theorem 5 as in the proof of Corollary 1. ■

IV. Proof of Theorem 3

Here we discuss the structure of tight bisets defining x^*. Our aim is to prove "laminarity" of the tight bisets, but what is the laminarity of bisets? We can find an answer to this question in previous works [8], [20].

Definition 1: A family $\mathcal{F}$ of bisets is called *laminar* if it satisfies all of the following conditions: (i) Set family $\{S \mid \hat{S} \in \mathcal{F}\}$ is laminar; (ii) If $\hat{X}, \hat{Y} \in \mathcal{F}$ satisfy $X \subset Y$, then $X^+ \subseteq Y^+$; (iii) If $\hat{X}, \hat{Y} \in \mathcal{F}$ satisfy $X = Y$, then $Y^+ \subseteq X^+$ or $X^+ \subseteq Y^+$.

We have the following property of tight constraints which define the basic optimal solution x^* to (3). This can be proven by a standard uncrossing technique. See other papers such as [4], [8], [20] for the proof.

Lemma 1: Let G be an undirected graph, and $f : \mathcal{V} \to \mathbb{Z}_+$ be a skew supermodular function such that $\max\{\gamma(S) \mid f(\hat{S}) > 0\} < k$. Let x^* be the basic optimal solution to (3) such that $0 < x^*(e) < 1$ for each $e \in F$, and C be a maximal subset of $\{v \in B \mid x^*(\delta_F(v)) = b(v)\}$ such that $\{\chi_F(v) \mid v \in C\}$ is linearly independent. Then there exists a laminar family $\mathcal{L}$ of proper tight bisets such that $|\mathcal{L}|+|C| = |F|$, the vectors in $\{\chi_F(\hat{S}) \mid \hat{S} \in \mathcal{L}\} \cup \{\chi_F(v) \mid v \in C\}$ are linearly independent, and x^* is the unique solution to $\{x(\delta_F(\hat{S})) = f(\hat{S}) \mid \hat{S} \in \mathcal{L}\} \cup \{x(\delta_F(v)) = b(v) \mid v \in C\}$.

We now prove Theorem 3. Suppose for a contradiction that $0 < x^*(e) < 1/(4k-1)$ for each $e \in F$ and $|\delta_F(v)| \geq 4k$ for each $v \in B$. Let $\mathcal{L}$ be the laminar family of tight bisets and C be the set of tight degree-bounded nodes in Lemma 1.

By Definition 1, $\{S \mid \hat{S} \in \mathcal{L}\}$ is laminar, and if $\hat{X}, \hat{Y} \in \mathcal{L}$ satisfy $X = Y$, then $X^+ \subset Y^+$ or $Y^+ \subset X^+$. Identifying each $v \in C$ as a biset $(\{v\}, \{v\})$, we use $\mathcal{L}'$ to denote the family $\mathcal{L} \cup C$ of bisets. $\mathcal{L}'$ is also laminar. We define a partial order $\prec$ on $\mathcal{L}'$ so that $\hat{X} \prec \hat{Y}$ holds when (i) $X \subset Y$, or (ii) $X = Y$ and $X^+ \subset Y^+$. This order defines a forest structure on $\mathcal{L}'$. In the rest of this paper, "minimal" and "maximal" are defined with respect to this order. For $\hat{X}, \hat{Y} \in \mathcal{L}'$, 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 $\hat{X} \prec \hat{Y}$.

Let $e = uv \in F$; we distribute two tokens for this edge e to some bisets in $\mathcal{L}'$ using the following rules.

(i) Note that there is a biset $\hat{X} \in \mathcal{L}'$ such that $u \in X$ and $v \in V \setminus X^+$ since e is included in F. We denote the minimal such biset $\hat{X}$ by $\hat{S}_{(e,u)}$, and give it the first token of e.
(ii) Suppose there also exists a biset $\hat{Y} \in \mathcal{L}'$ such that $v \in Y$ and $u \in V \setminus Y^+$ then we denote the minimal such biset $\hat{Y}$ by $\hat{S}_{(e,v)}$ and give it the second token of e. Otherwise, e gives the second token to the biset $\hat{Z}$ which is minimal in $\{\hat{Z} \in \mathcal{L}' \mid \hat{S}_{(e,u)} \prec \hat{Z}, v \in Z^+\}$. In the former case when both $\hat{S}_{(e,u)}$ and $\hat{S}_{(e,v)}$ exist, we say that e is *missed* by the biset $\hat{Z}$ which is minimal in $\{\hat{Z} \in \mathcal{L}' \mid \hat{S}_{(e,u)} \prec \hat{Z}, v \in Z^+\}$.

Induction proof

Observe that each edge in F distributes at most two tokens. In what follows, we show that it is possible to rearrange the tokens so that each biset in $\mathcal{L}'$ obtains at least two tokens, each maximal biset in $\mathcal{L}'$ obtains at least $2k+2$ tokens, and obtains $4k$ tokens unless it has exactly one child. This proves Theorem 3 because it means that $2|\mathcal{L}'| = 2(|\mathcal{L}|+|C|) = 2|F| \geq$ (the number of distributed tokens) $> 2|\mathcal{L}'|$, which is a contradiction.

Our proof is by induction on the height of the forest defined from the partial order on $\mathcal{L}'$. The base case of our induction is when the height of the forest is 1. That is to say, each biset in $\mathcal{L}'$ has no child. In this case, each biset $\hat{S} \in \mathcal{L}'$ obtains one token from each edge in $\delta_F(\hat{S})$. If $\hat{S} \in \mathcal{L}$, then $|\delta_F(\hat{S})| \geq 4k$ follows from $x^*(\delta_F(\hat{S})) = f(\hat{S}) \geq 1$ and $x^*(e) < 1/(4k-1)$ for $e \in \delta_F(\hat{S})$. If $S = (\{v\}, \{v\})$ for some $v \in C$, then the assumption indicates $|\delta_F(\hat{S})| \geq 4k$. In either case, $\hat{S}$ obtains at least $4k$ tokens.

Now let us discuss the case where the forest contains a tree of height more than one. Let $\hat{R}$ be the biset maximal in

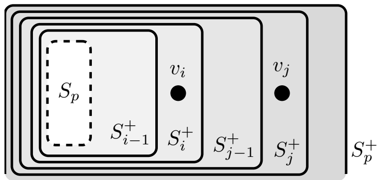

Figure 1. $\hat{S}_i$ and $\hat{S}_j$ with $i, j \in I$ and $i \neq j$ in the proof of Lemma 2

the tree. Suppose that $\hat{R}$ has at least two children. Then by induction, we can arrange the tokens in the subtree rooted at each child so that the claim holds, i.e., each biset has at least two tokens and each child of $\hat{R}$ has at least $2k+2$ tokens. From each child of $\hat{R}$, move $2k$ tokens to $\hat{R}$. Then $\hat{R}$ obtains $4k$ tokens and each child keeps 2 tokens. Thus the claim holds in this case.

The remaining case is when $\hat{R}$ has exactly one child. Let $\hat{S}_1$ be the maximal descendent of $\hat{R}$ which has more than one child or has no child. Let $\hat{S}_2, \hat{S}_3, \ldots, \hat{S}_{p-1}$ be the bisets on the path from $\hat{S}_1$ to $\hat{R}$ in the tree, and assume that $\hat{S}_1 \prec \hat{S}_2 \prec \cdots \prec \hat{S}_{p-1} \prec \hat{S}_p = \hat{R}$. Note that the only child of each $\hat{S}_i$, $i \in \{2, 3, \ldots, p\}$ is $\hat{S}_{i-1}$. We apply the induction hypothesis for the subtree rooted at $\hat{S}_1$. That is to say, we arrange the tokens in the subtree so that each biset below $\hat{S}_1$ has two tokens, and $\hat{S}_1$ has $4k$ tokens. Then we show that the tokens owned by $\hat{S}_1, \hat{S}_2, \ldots, \hat{S}_p$ can be distributed so that each $\hat{S}_i$, $i \in \{1, 2, \ldots, p-1\}$ has at least two tokens and $\hat{S}_p$ has $2k+2$ tokens.

Let $i \in \{2, 3, \ldots, p\}$. The linear independence between $\chi_F(\hat{S}_i)$ and $\chi_F(\hat{S}_{i-1})$ implies that $|\delta_F(\hat{S}_i) \setminus \delta_F(\hat{S}_{i-1})| + |\delta_F(\hat{S}_{i-1}) \setminus \delta_F(\hat{S}_i)| > 0$. In particular $|\delta_F(\hat{S}_i) \setminus \delta_F(\hat{S}_{i-1})| + |\delta_F(\hat{S}_{i-1}) \setminus \delta_F(\hat{S}_i)| \geq 2$ by the fact that $x^*(\delta_F(\hat{S}_i))$ and $x^*(\delta_F(\hat{S}_{i-1}))$ are integers and $x^*(e) < 1$ for each $e \in F$. $\hat{S}_i$ obtains one token from each edge in $\delta_F(\hat{S}_i) \setminus \delta_F(\hat{S}_{i-1})$. Hence we are done if $|\delta_F(\hat{S}_i) \setminus \delta_F(\hat{S}_{i-1})| \geq 2$.

Consider the case where $|\delta_F(\hat{S}_i) \setminus \delta_F(\hat{S}_{i-1})| \leq 1$. Then $|\delta_F(\hat{S}_{i-1}) \setminus \delta_F(\hat{S}_i)| \geq 1$. Let $e = uv \in \delta_F(\hat{S}_{i-1}) \setminus \delta_F(\hat{S}_i)$ with $u \in S_{i-1} \subseteq S_i$ and $v \in S^+_i \setminus S^+_{i-1}$. If $\hat{S}_{(e,v)}$ does not exist, then $\hat{S}_i$ gets the second token of e. Otherwise, e is missed by $\hat{S}_i$. Thus $\hat{S}_i$ obtains two tokens unless $\hat{S}_i$ misses an edge.

After applying the induction hypothesis, $\hat{S}_1$ has $4k$ tokens. As we have observed, $\hat{S}_i$, $i \geq 2$ has two tokens unless $\hat{S}_i$ misses an edge. In the next lemma, we show that at most $k-1$ bisets in $\hat{S}_2, \hat{S}_3, \ldots, \hat{S}_p$ miss some edges. Then we can prove the claim by making $\hat{S}_1$ give two tokens to each of such bisets (potentially including $\hat{S}_p$), and $2k$ tokens to $\hat{S}_p$. This completes the proof of Theorem 3.

Lemma 2: At most $k-1$ bisets in $\hat{S}_2, \hat{S}_3, \ldots, \hat{S}_p$ miss some edges.

Proof: Let $I \subseteq \{2, \ldots, p\}$ be the set of indices such that each $\hat{S}_i$, $i \in I$ misses an edge $e_i = u_i v_i \in F$. Suppose that $u_i \in S_{i-1} \subseteq S_i$ and $v_i \in S^+_i \setminus S^+_{i-1}$ for each $i \in I$.

Let $i \in I$. Since $S^+_i \subseteq S^+_p$, $v_i \in S^+_p$. If $v_i \in S_p$, then $S_{(e_i, v_i)} \subseteq S_p$ by the laminarity of $\{S \mid \hat{S} \in \mathcal{L}\}$, but this contradicts the fact that each $\hat{S}_i$, $i \geq 2$ has only one child. Thus $v_i \in \Gamma(\hat{S}_p)$. $v_i \neq v_j$ holds for any $i, j \in I$ with $i \neq j$ because $v_i \in S^+_i \setminus S^+_{i-1}$ for each $i \in I$. Figure 1 illustrates these facts. Since $\gamma(\hat{S}_p) \leq k-1$, $|I| \leq k-1$. In other words, at most $k-1$ bisets in $\hat{S}_2, \hat{S}_3, \ldots, \hat{S}_p$ miss some edges. ■

V. Proof of Theorem 4

The following lemma shows that a family of tight bisets characterizing the basic optimal solution to (4) is laminar.

Lemma 3 ([20]): Let $G = (V, E)$ be a digraph, $f : \mathcal{V} \to \mathbb{Z}_+$ be an intersecting supermodular function, and x^* be the basic optimal solution to (4) such that $0 < x^*(e) < 1$ for each $e \in F$. Then there exists a laminar family $\mathcal{L}$ of proper tight bisets and a set $C^+ \subseteq B^+$ of tight degree-bounded nodes such that $|\mathcal{L}| + |C^+| = |F|$, the vectors in $\{\chi^-_F(\hat{S}) \mid \hat{S} \in \mathcal{L}\} \cup \{\chi^+_F(v) \mid v \in C^+\}$ are linearly independent, and x^* is the unique solution to $\{x(\delta^-_F(\hat{S})) = f(\hat{S}) \mid \hat{S} \in \mathcal{L}\} \cup \{x(\delta^+_F(v)) = b^+(v) \mid v \in C^+\}$.

For the sake of arriving at a contradiction, suppose that each arc $e \in F$ satisfies $0 < x^*(e) < 1/2$, and each $v \in B^+$ satisfies $|\delta^+_F(v)| \geq 4k+1$. Define $\mathcal{L}$ and C^+ as in Lemma 3. Let $\mathcal{E}$ denote the set of leaf bisets in $\mathcal{L}$. We use one of two arguments according to whether $|\mathcal{E}| \geq |C^+|$ holds or not.

First, let us consider the case where $|\mathcal{E}| \geq |C^+|$. We modify the proof for the existence of an integer-valued variable when degree-bounds are not given [15]. For $e \in F$, we let $\hat{S}_e$ denote the minimal biset $\hat{S} \in \mathcal{L}$ such that $e \in \delta^-_F(\hat{S})$. $\hat{S}_e$ always exists since otherwise $x^*(e) = 0$. We make each $e \in F$ distribute one token to $\hat{S}_e$. Then each $\hat{S} \in \mathcal{E}$ has $|\delta^-_F(\hat{S})| \geq 2f(\hat{S})+1$ tokens. We make each biset in $\mathcal{E}$ give one token to a node in C^+. Since $|\mathcal{E}| \geq |C^+|$, each node in C^+ obtains one token and each $\hat{S} \in \mathcal{E}$ owns $2f(\hat{S}) \geq f(\hat{S})+1$ tokens after this. As in the proof of [15] for the no degree-bounds case, we can show via induction that the tokens owned by bisets in $\mathcal{L}$ can be redistributed so that each biset $\hat{S} \in \mathcal{L}$ has at least one token, and at least $f(\hat{S})+1$ tokens if $\hat{S}$ is maximal. This implies a contradiction that $|F| < |\mathcal{L}| + |C^+|$.

Now we consider the case where $|\mathcal{E}| < |C^+|$. The argument for this case is similar to the one for Theorem 3. We make each $e = uv \in F$ distribute two tokens according to the following rules:

(i) e gives one token to $\hat{S}_e$;
(ii) If $u \in C^+$, then e gives one token to u;
(iii) If $u \notin C^+$, then e gives one token to the minimal biset $\hat{X} \in \mathcal{L}$ such that $\hat{S}_e \prec \hat{X}$ and $e \notin \delta^-_F(\hat{X})$.

The total number of tokens distributed by arcs is at most $2|F|$. We show how to redistribute these tokens such that each biset in $\mathcal{L}$ and each node in C^+ can obtain at least two tokens and one extra token remains, giving a contradiction.

Let $v \in C^+$. Since v obtains one token from each arc in $\delta_F^+(v)$ and $|\delta_F^+(v)| \geq 4k+1$, v obtains at least $4k+1$ tokens. These tokens are redistributed as follows.

(iv) v keeps 2 tokens;

(v) If the minimal biset $\hat{X}$ in $\{\hat{S} \in \mathcal{L} \mid v \in S\}$ has a unique minimal descendent $\hat{Y}$ such that $v \in Y^+$, then 2 tokens are given to $\hat{Y}$;

(vi) $4k-3$ tokens are given to a biset in $\mathcal{E}$.

Since $|\mathcal{E}| < |C^+|$, each biset in $\mathcal{E}$ receives $4k-3$ tokens from a node in C^+ by rule (vi). By rule (iv), each node in C^+ has already owned 2 tokens. In what follows, we discuss the tokens for $\mathcal{L}$. We claim that it is possible to arrange the tokens given to them so that

- each biset in $\mathcal{L}$ has at least 2 tokens;
- each maximal biset in $\mathcal{L}$ has at least $2k+2$ tokens, and has $4k$ tokens unless it has exactly one child.

We prove the claim by induction on the height of the forest representing $\mathcal{L}$. In the base case when the height of the forest is one, it is not difficult to prove the claim. Hence let us consider the case where the height is at least two. Let $\hat{R}$ be a maximal biset of $\mathcal{L}$. If $\hat{R}$ has at least two children, then the proof is immediate; Since each child has at least $2k+2$ tokens by the induction hypothesis, let the child keep 2 tokens and give $2k$ tokens to $\hat{R}$. Thus suppose that $\hat{R}$ has only one child.

Let $\hat{S}_1$ be the maximal descendent of $\hat{R}$ such that it has more than one child or has no child. Let $\hat{S}_2, \hat{S}_3, \ldots, \hat{S}_{p-1}$ be the bisets on the path from $\hat{S}_1$ to $\hat{R}$ in the forest, and assume that $\hat{S}_1 \prec \hat{S}_2 \prec \cdots \prec \hat{S}_{p-1} \prec \hat{S}_p = \hat{R}$. We apply the induction hypothesis for the subtree rooted at $\hat{S}_1$ to arrange the tokens in the subtree so that each biset below $\hat{S}_1$ has 2 tokens, and $\hat{S}_1$ has $4k$ tokens. As in Section IV, we can observe that $\hat{S}_i$ has two tokens from the edges unless the tail of each $e \in \delta_F^-(\hat{S}_{i-1}) \setminus \delta_F^-(\hat{S}_i)$ is in $(\Gamma(\hat{S}_i) \cap \Gamma(\hat{S}_p)) \setminus S_{i-1}^+$ for each $i \in \{2, 3, \ldots .p\}$, and there are at most $k-1$ such bisets in $\hat{S}_2, \hat{S}_3, \ldots, \hat{S}_p$. By moving two tokens from $\hat{S}_1$ give to each of such bisets, and $2k$ tokens to $\hat{S}_p$, we can prove the claim.

VI. Proof of Theorem 5

For proving Theorem 5, we suppose for the sake of arriving at a contradiction that each $e \in F$ satisfies $0 < x^*(e) < 1/6$, and $|\delta_F(v)| \geq 16k^2 - 4k - 6$ holds for each $v \in B$. Even if we replace (3) by (5), we have a laminar family $\mathcal{L}$ of tight bisets and a set $C \subseteq B$ of tight degree-bounded nodes in Lemma 1 because (3) and (5) are different only in their objective functions. In fact $\mathcal{L}$ is not only laminar, but it is *non-overlapping*, which is a stronger property than laminarity. Refer to [4] for the definition of non-overlapping families of bisets. Since the property of non-overlapping is closed under replacing a biset (S, S^+) by $(V \setminus S^+, V \setminus S)$, we assume without loss of generality that each $\hat{S} \in \mathcal{L}$ satisfies $|S| \leq |V \setminus S^+|$.

Moreover, if x^* is the basic optimal solution to (5), we have an extra property that every edge in F is incident to a node in C; Each $e \in F$ is incident to a tight degree-bounded node in B since otherwise $x^*(e) = 1$; Since C is a maximal set of tight degree-bounded nodes such that $\{\chi_F(v) \mid v \in C\}$ is linearly independent, at least one of the end nodes of every $e \in F$ is included in C.

Below we derive a contradiction that $|F|$ is larger than $|\mathcal{L}| + |C|$ by new token arguments. Recall that $\mathcal{E}$ is the set of leaves of $\mathcal{L}$. We use one of two arguments according to whether $|\mathcal{E}| < (4k+3)|C|$ holds or not.

A. Token argument when $|\mathcal{E}| < (4k+3)|C|$

Each edge $e = vu \in F$ distributes two tokens. In particular, for each pair of $e = vu \in F$ and an end node v of e, one token is distributed according to the following rules applied in this order:

(i) If $v \in C$, then e gives one token to v;

(ii) If $v \notin C$ and $\hat{S}_{(e,v)}$ exists, then e gives one token to $\hat{S}_{(e,v)}$;

(iii) If $v \notin C$, $\hat{S}_{(e,v)}$ does not exist and $\hat{S}_{(e,u)}$ exists, then e gives one token to the biset $\hat{X} \in \mathcal{L}$ which is minimal in $\{\hat{X} \in \mathcal{L} \mid \hat{S}_{(e,u)} \prec \hat{X}, v \in X^+\}$.

Let $v \in C$. Since $|\delta_F(v)| \geq 16k^2 - 4k - 6$ holds, v obtains at least $16k^2 - 4k - 6$ tokens. These tokens are redistributed as follows.

(iv) v keeps 2 tokens;

(v) 4 tokens are given to the minimal biset $\hat{S} \in \mathcal{L}$ such that $v \in S$;

(vi) $16k^2 - 4k - 12 = (4k-4)(4k+3)$ tokens are given to $4k+3$ bisets in $\mathcal{E}$ so that each of those $4k+3$ bisets obtains $4k-4$ tokens.

Since $|\mathcal{E}| < (4k+3)|C|$, each biset in $\mathcal{E}$ receives $4k-4$ tokens by rule (vi). By rule (iv), each $v \in C$ has already owned 2 tokens. We claim that it is possible to arrange tokens so that

- each biset in $\mathcal{L}$ has at least 2 tokens;
- each maximal biset has at least $2k+2$ tokens, and has $4k$ tokens unless it has exactly one child.

We do not give a proof for this claim since it is almost same as the proof of Theorem 3. The main thing to note here is that the tokens given to nodes in C have already been distributed now.

B. Token argument when $|\mathcal{E}| \geq (4k+3)|C|$

We again make each $e = uv \in F$ distribute 2 tokens. For each pair of $e = uv \in F$ and an end node v of e, we make e give a token as follows.

(i) If $v \notin C$ and there exists $\hat{S}_{(e,v)}$, then e gives one token to $\hat{S}_{(e,v)}$;

(ii) If $v \in C$ or $\hat{S}_{(e,v)}$ does not exist, and if $\hat{S}_{(e,u)}$ exists, then e gives one token to the minimal biset $\hat{X}$ in $\{\hat{X} \in \mathcal{L} \mid \hat{S}_{(e,u)} \prec \hat{X}, v \in X^+\}$.

We call $\hat{S} \in \mathcal{L}$ *white* if $S \cap C = \emptyset$, and *black* otherwise. Note that by definition we cannot have a white biset $\hat{X}$ and a black biset $\hat{Y}$ such that $\hat{Y} \preceq \hat{X}$. If a black biset is minimal in $\{\hat{S} \in \mathcal{L} \mid v \in S\}$ for some $v \in C$, then it is called *strictly black*. The number of strictly black bisets is at most $|C|$.

Now each node in C has no token. A black biset may have tokens, but it received no tokens from the edges whose end nodes in its inner-part are in C. On the other hand, each white minimal biset $\hat{S}$ owns 7 tokens because $|\delta_F(\hat{S})| \geq 7$ by $x^*(\delta_F(\hat{S})) = f(\hat{S}) \geq 1$ and $x^*(e) < 1/6$ for $e \in \delta_F(\hat{S})$, and the end node of $e \in \delta_F(\hat{S})$ in S is not in C. We make each white minimal biset give one token. The total number of tokens given by the white minimal bisets is at least $|\mathcal{E}| - |C| \geq (4k+3)|C| - |C| = (4k+2)|C|$. We allocate these tokens to nodes in C and strictly black bisets so that

- Each strictly black biset has $4k$ tokens;
- Each node in C has 2 tokens.

We then prove by induction that we can rearrange the tokens so that each biset in $\mathcal{L}$ and each node in C obtains at least 2 tokens, and each maximal biset obtains more than 2 tokens.

White bisets: We begin with a tree which consists of only white bisets.

Lemma 4: Let $\hat{R} \in \mathcal{L}$ be a white biset, and $\mathcal{L}' = \{\hat{S} \in \mathcal{L} \mid \hat{S} \preceq \hat{R}\}$. If $|\mathcal{L}'| > 1$, then we can rearrange tokens so that each biset in $\mathcal{L}'$ has at least 2 tokens, each minimal biset in $\mathcal{L}'$ has 4 tokens, and $\hat{R}$ has 4 tokens. If $|\mathcal{L}'| = 1$, then $\hat{R}$ has at least 6 tokens.

Proof: We prove by the induction on the height of $\mathcal{L}'$. If the height is one, then the lemma is obvious. Hence suppose that the height is more than one.

If $\hat{R}$ has at least two children, then the lemma is proven by applying the induction hypothesis to the trees rooted on the children of $\hat{R}$ and by making each child give 2 tokens to $\hat{R}$. Hence assume that $\hat{R}$ has only one child $\hat{Q}$. Since $\hat{Q}$ can give 2 tokens to $\hat{R}$, it suffices to find two more tokens for $\hat{R}$. We show that each edge in $(\delta_F(\hat{R}) \setminus \delta_F(\hat{Q})) \cup (\delta_F(\hat{Q}) \setminus \delta_F(\hat{R}))$ gives one token to $\hat{R}$, which proves the lemma because there are at least two such edges.

Let $e \in \delta_F(\hat{R}) \setminus \delta_F(\hat{Q})$. Then e has an end node $v \in R \setminus Q$ and $\hat{R} = \hat{S}_{(e,v)}$. $v \notin C$ follows from the fact that $\hat{R}$ is white. Hence e gives one token to $\hat{R}$. Next, let $e \in \delta_F(\hat{Q}) \setminus \delta_F(\hat{R})$. The end node v of e in $V \setminus Q^+$ is in C because the other end node is not in C by the fact that $\hat{Q}$ is white. Notice that $v \in R^+$ because $e \notin \delta_F(\hat{R})$. Hence e gives this token to $\hat{R}$ by rule (ii). Therefore the lemma is proven. ■

Black bisets: We next present a proof for a tree whose maximal biset is black. In this case, we show how to rearrange the tokens so that each biset obtains at least two tokens, the maximal biset obtains at least $2k+2$ tokens, and obtains $4k$ tokens unless the number of its black children is exactly one. Our proof is by the induction on the height of the tree again. 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.

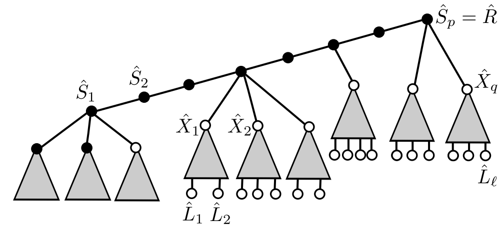

Figure 2. Structure of the tree rooted at a black biset $\hat{R}$ (black bisets are represented by filled circles and white bisets are represented by the void circles)

Let $\hat{R}$ be the maximal biset. If $\hat{R}$ has at least two black children, apply the induction hypothesis for the subtrees rooted at the black children of $\hat{R}$, and Lemma 4 for those rooted at the white children of $\hat{R}$. Then by making each of the black children give $2k$ tokens to $\hat{R}$, we can prove the claim. If $\hat{R}$ has no black child, then it already has $4k$ tokens because $\hat{R}$ is strictly black. In what follows, we discuss the case where $\hat{R}$ has exactly one black child.

Let $\hat{S}_1$ be the maximal black biset which has more than one black child or has no black children in the subtree rooted at $\hat{R}$. Let $\hat{S}_2, \hat{S}_3, \ldots, \hat{S}_{p-1}$ be the bisets on the path from $\hat{S}_1$ to $\hat{R}$, and assume that $\hat{S}_1 \prec \hat{S}_2 \prec \cdots \prec \hat{S}_{p-1} \prec \hat{S}_p = \hat{R}$. Note that each $\hat{S}_i$, $i \in \{1, 2, \ldots, p\}$ is black. Let $\hat{X}_1, \hat{X}_2, \ldots, \hat{X}_q$ be the white children of $\hat{S}_2, \hat{S}_3, \ldots, \hat{S}_p$, and $\hat{L}_1, \hat{L}_2, \ldots, \hat{L}_\ell$ be the minimal bisets in the subtrees rooted at $\hat{X}_1, \hat{X}_2, \ldots, \hat{X}_q$. Figure 2 illustrates these definitions.

Applying the induction hypothesis for the subtree rooted at $\hat{S}_1$, and Lemma 4 for the subtrees rooted at $\hat{X}_1, \hat{X}_2, \ldots, \hat{X}_q$, we allocate tokens so that

- $\hat{S}_1$ has $4k$ tokens,
- each of $\hat{X}_1, \hat{X}_2, \ldots, \hat{X}_q$ has at least 4 tokens,
- each of $\hat{L}_1, \hat{L}_2, \ldots, \hat{L}_\ell$ has at least 4 tokens,
- if $\hat{X}_i = \hat{L}_j$, then it has 6 tokens,
- each of the other bisets in the subtrees rooted at $\hat{S}_1$ and $\hat{X}_1, \hat{X}_2, \ldots, \hat{X}_q$ has 2 tokens.

Let $i \in \{2, 3, \ldots, p\}$. If $\hat{S}_i$ has a white child, then the child can give two tokens to $\hat{S}_i$. If $\hat{S}_i$ is strictly black, then it has $4k$ tokens. Thus assume that $\hat{S}_i$ has no white child (i.e., $\hat{S}_{i-1}$ is its only child) and it is not strictly black. $|\delta_F(\hat{S}_i) \setminus \delta_F(\hat{S}_{i-1})| + |\delta_F(\hat{S}_{i-1}) \setminus \delta_F(\hat{S}_i)| \geq 2$ holds.

Let $e = uv \in \delta_F(\hat{S}_i) \setminus \delta_F(\hat{S}_{i-1})$, and assume without loss of generality that $u \in S_i \setminus S_{i-1}$ and $v \in V \setminus S_i^+ \subseteq V \setminus S_{i-1}^+$. Then e gives one token to $\hat{S}_i$ because $\hat{S}_i = \hat{S}_{(e,u)}$ and $u \notin C$ since $\hat{S}_i$ is not strictly black. Next, let $e = uv \in \delta_F(\hat{S}_{i-1}) \setminus \delta_F(\hat{S}_i)$, and assume without loss of generality that $u \in S_{i-1} \subseteq S_i$ and $v \in (V \setminus S_{i-1}^+) \cap S_i^+$. If $v \in C$, then e gives one token to $\hat{S}_i$ by rule (ii). Then e gives one token to $\hat{S}_i$ unless $v \notin C$ and $\hat{S}_{(e,v)}$ exists. Summarizing, if $\hat{S}_i$ does not obtain two tokens, then there exists an edge

$e = uv \in \delta_F(\hat{S}_{i-1}) \setminus \delta_F(\hat{S}_i)$ such that $u \in S_{i-1} \subseteq S_i$, $v \in (S_i^+ \setminus S_{i-1}^+) \setminus C$, and $\hat{S}_{(e,v)}$ exists. The next lemma shows that there exist at most $k-1+\ell$ such bisets in $\hat{S}_2, \hat{S}_3, \ldots, \hat{S}_p$ (We do not present the proof due to the space limitation).

Lemma 5: Let $I \subseteq \{2, 3, \ldots, p\}$ be the set of indices i such that there exists an edge $e_i = u_i v_i \in \delta_F(\hat{S}_{i-1}) \setminus \delta_F(\hat{S}_i)$ with $u_i \in S_{i-1} \subseteq S_i$ and $v_i \in (S_i^+ \setminus S_{i-1}^+) \setminus C$, and $\hat{S}_{(e_i,v_i)}$ exists. If $\mathcal{L}$ is a non-overlapping family such that $|S| \leq |V \setminus S^+|$ for each $\hat{S} \in \mathcal{L}$, then $|I| \leq k - 1 + \ell$.

Recall that $\hat{S}_1$ has $4k-2$ extra tokens, and each $\hat{L}_i$, $i = 1, 2, \ldots, \ell$ has two extra tokens. From these tokens, give two tokens to $\hat{S}_i$ for each $i \in I$. Then each biset in the tree rooted at $\hat{S}_p$ (including $\hat{S}_p$) obtains two tokens, and $4k-2+2\ell-2|I| \geq 4k-2+2\ell-2(k-1+\ell) = 2k$ tokens still remain. By moving these tokens to $\hat{S}_p$, we can arrange tokens as required.

VII. Conclusion

We have presented iterative rounding algorithms for the degree-bounded SND and for the degree-bounded subgraph problem. Our result for the degree-bounded SND with requirements represented by skew supermodular functions are based on the laminarity of a family of tight bisets. However, the family has a stronger property, which is called strong laminarity in [20]. After announcing our results, Nutov pointed out to us that our analysis can be improved to $(O(1), O(1) \cdot b(v) + O(k))$-approximation for the element-connectivity SND using strong laminarity. We will include this in a full version of this paper.

Acknowledgments

This work was begun when the second author was a visiting professor at RIMS, Kyoto University, in Spring 2011 and completed when the first author was visiting Carnegie Mellon University in 2011-12, supported by Kyoto University Foundation. The second author was supported in part by NSF grant CCF 1143998. The authors thank Zeev Nutov for his suggestion on investigating k-out-connectivity SND.

References

[1] A. Aazami, J. Cheriyan, and B. Laekhanukit, "A bad example for the iterative rounding method for mincost k-connected spanning subgraphs," Manuscript, 2010.

[2] N. Bansal, R. Khandekar, J. Könemann, V. Nagarajan, and B. Peis, "On generalizations of network design problems with degree bounds," in *IPCO*, LNCS 6080, 2010, pp. 110–123.

[3] N. Bansal, R. Khandekar, and V. Nagarajan, "Additive guarantees for degree-bounded directed network design," *SIAM J. Comput.*, vol. 39, pp. 1413–1431, 2009.

[4] J. Cheriyan, S. Vempala, and A. Vetta, "Network design via iterative rounding of setpair relaxations," *Combinatorica*, vol. 26, pp. 255–275, 2006.

[5] J. Chuzhoy and S. Khanna, "An $O(k^3 \log n)$-approximation algorithms for vertex-connectivity survivable network design," in *50th FOCS*, 2009, pp. 437–441.

[6] J. Fakcharoenphol and B. Laekhanukit, "An $o(\log^2 k)$-approximation algorithm for the k-vertex connected spanning subgraph problem," in *40th STOC*, 2008, pp. 153–158.

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

[8] L. Fleischer, K. Jain, and D.P. Williamson, "Iterative rounding 2-approximation algorithms for minimum-cost vertex connectivity problems," *J. Comput. and Sys. Sci.*, vol. 72, pp. 838–867, 2006.

[9] K. Jain, "A factor 2 approximation algorithm for the generalized Steiner network problem," *Combinatorica*, vol. 21, pp. 39–60, 2001.

[10] R. Khandekar, G. Kortsarz, and Z. Nutov, "Network-design with degree constraints," in *APPROX-RANDOM 2011*, LNCS 6845, 2011, pp. 289–301.

[11] T. Király, L.C. Lau, and M. Singh, "Degree bounded matroids and submodular flows," in *13th IPCO*, LNCS 5035, 2008, pp. 259–272.

[12] B. Laekhanukit, "An improved approximation algorithm for minimum-cost subset k-connectivity," in *38th ICALP*, LNCS 6755, 2011, pp. 13–24.

[13] L.C. Lau, J. Naor, M.R. Salavatipour, and M. Singh, "Survivable network design with degree or order constraints," in *39th STOC*, 2007, pp. 651–660.

[14] L.C. Lau, J. Naor, M.R. Salavatipour, and M. Singh, "Survivable network design with degree or order constraints," *SIAM J. Comput.*, vol. 39, pp. 1062–1087, 2009.

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

[16] L.C. Lau and M. Singh, "Additive approximation for bounded degree survivable network design," in *40th STOC*, 2008, pp. 759–768.

[17] Z. Nutov, "Approximating minimum-cost edge-covers of crossing biset-families," Manuscript (Preliminary version: "An almost $O(\log k)$-approximation for k-connected subgraphs," in *20th SODA*, 2009, pp. 912–921).

[18] Z. Nutov, "Approximating minimum cost connectivity problems via uncrossable bifamilies and spider-cover decompositions," in *50th FOCS*, 2009, pp. 417–426.

[19] Z. Nutov, "Approximating subset k-connectivity problems," CoRR abs/1105.4250, 2011.

[20] Z. Nutov, "Degree-constrained node-connectivity problems," in *10th LATIN*, LNCS 7256, 2012, pp. 582–593.

[21] M. Singh and L.C. Lau, "Approximating minimum bounded degree spanning trees to within one of optimal," in *39th STOC*, 2007, pp. 661–670.