Approximating Steiner Network Activation Problems

Zeev Nutov

The Open University of Israel
nutov@openu.ac.il

Abstract. In the Steiner Networks Activation problem we are given a graph $G = (V, E)$, $S \subseteq V$, a family $\{f^{uv}(x_u, x_v) : uv \in E\}$ of monotone non-decreasing *activating functions* from $\mathbb{R}^2_+$ to $\{0, 1\}$ each, and *connectivity requirements* $\{r(u, v) : u, v \in V\}$. The goal is to find a *weight assignment* $\mathbf{w} = \{w_v : v \in V\}$ of minimum total weight $\mathbf{w}(V) = \sum_{v \in V} w_v$, such that: for all $u, v \in V$, the *activated graph* $G_{\mathbf{w}} = (V, E_{\mathbf{w}})$, where $E_{\mathbf{w}} = \{uv : f^{uv}(w_u, w_v) = 1\}$, contains $r(u, v)$ pairwise edge-disjoint uv-paths such that no two of them have a node in $S \setminus \{u, v\}$ in common. This problem was suggested recently by Panigrahi [14], generalizing the Node-Weighted Steiner Network and the Minimum-Power Steiner Network problems, as well as several other problems with motivation in wireless networks. We give new approximation algorithms for this problem.
For undirected/directed graphs, our ratios are $O(k \log n)$ for k-Out/In-connected Subgraph Activation and k-Connected Subgraph Activation. For directed graphs this solves a question from [14] for $k = 1$, while for the min-power case and k arbitrary this solves an open question from [11]. For other versions on undirected graphs, our ratios match the best known ones for the Node-Weighted Steiner Network problem [10].

1 Introduction

In Network Design problems, we are given a graph $G = (V, E)$, a function $\omega : 2^E \to \mathbb{R}_+$, and a *monotone* property Π of subgraphs of G; monotonicity of Π means that $H \in \Pi$ implies $H' \in \Pi$ for any $H \subseteq H' \subseteq G$. The goal is to find $F \subseteq E$ with $\omega(F)$ minimum, such that $(V, F) \in \Pi$. In Edge-Costs Network Design problems $\omega(F) = \mathbf{c}(F) = \sum_{e \in F} c_e$ for given edge-costs $\mathbf{c} = \{c_e : e \in E\}$. For an edge-set F on V let $V(F)$ denote the set of endnodes of the edges in F. In Node-Weighted Network Design problems, instead of edge-costs we are given node-weights $\mathbf{w} = \{w_v : v \in V\}$, and seek a node subset $V' \subseteq V$ of minimum total weight $\mathbf{w}(V') = \sum_{v \in V'} w_v$ such that the subgraph (V', F) of G induced by V' satisfies Π; equivalently, we seek an edge subset $F \subseteq E$ such that the graph (V, F) satisfies Π and $\mathbf{w}(V(F))$ is minimum. Panigrahi [14] suggested the following generalization of Node-Weighted Network Design problems, that captures also several known problems in wireless network design. For further motivation, applications, and history of the problem, see the paper of Panigrahi [14].

Definition 1. *Let $G = (V, E)$ be a graph and let $\{f^{uv} : uv \in E\}$ be a family of* activating functions*, where each f^{uv} is from $D^{uv} \subseteq \mathbb{R}^2_+$ to $\{0,1\}$, and $f^{uv}(x_u, x_v) = f^{vu}(x_v, x_u)$ if G is undirected. Let $\mathbf{w} = \{w_v : v \in V\}$ be a* non-negative weight assignment *on V. An edge $uv \in E$ is* activated *by $\mathbf{w}$ if $f^{uv}(w_u, w_v) = 1$. Let $E_\mathbf{w} = \{uv \in E : f^{uv}(w_u, w_v) = 1\}$ be the set of edges activated by $\mathbf{w}$. For $V' \subseteq V$ let $\mathbf{w}(V') = \sum_{v \in V'} w_v$ be the weight of V'.*

We consider connectivity variants of the following problem.

Network Activation

Instance: A graph $G = (V, E)$, a family $\{f^{uv}(x_u, x_v) : uv \in E\}$ of activating functions from $D^{uv} \subseteq \mathbb{R}^2_+$ to $\{0,1\}$ each, and a graph property Π.

Objective: Find a weight assignment $\mathbf{w} = \{w_v : v \in V\}$ with $\mathbf{w}(V)$ minimum such that the graph $G_\mathbf{w} = (V, E_\mathbf{w})$ activated by $\mathbf{w}$ satisfies Π.

Unless stated otherwise, or is clear from the context, graphs can be undirected or directed. We will assume that each activating function f^{uv} admits a polynomial time evaluation oracle, and also use the following assumptions.

Assumption 1. *For every $uv \in E$, f^{uv} is monotone non-decreasing, namely, $f^{uv}(x_u, x_v) = 1$ implies $f^{uv}(y_u, y_v) = 1$ whenever $y_u, y_v \in D^{uv}$, $y_u \geq x_u$, and $y_v \geq x_v$.*

Assumption 2. *For every edge $e = uv \in E$, we can compute in polynomial time some optimal weight assignment $\mathbf{x}^e = \mathbf{x}^{uv}$ activating e; here $\mathbf{x}^e$ has values $x^e_u = x^{uv}_u$ on u and $x^e_v = x^{uv}_v$ on v (such that $f^{uv}(x^e_u, x^e_v) = 1$ and $x^e_u + x^e_v$ is minimal), and is zero otherwise.*

Assumption 3. *For every $uv \in E$, $D^{uv} = D^u \times D^v$ where $|D^u|, |D^v|$ are polynomial in $n = |V|$.*

We are not aware of any specific problems that do not satisfy Assumption 1 or Assumption 2. For justification of Assumption 3 see the paper of Panigrahi [14]. Note that Assumption 3 implies Assumption 2, since it enables to compute in polynomial time all weight assignments activating uv.

Network Activation generalizes Node-Weighted Network Design problems, by setting $f^{uv}(x_u, x_v) = 1$ if $x_u \geq w_u$, $x_v \geq w_v$, and $uv \in E$. Another famous example is the Minimum-Power Network Design problem, where instead of activating functions we are given edge-costs $\mathbf{c} = \{c_{uv} : uv \in E\}$. Here an edge uv is activated by a weight assignment $\mathbf{w}$ if $w_u, w_v \geq c_{uv}$ in the case of undirected graphs, or if $w_u \geq c_{uv}$ in the case of directed graphs. An equivalent formulation is as follows. For an undirected/directed edge-set F and a node v let $\delta_F(v)$ denote the set of edges in F incident to v. If F is directed, $\delta^{out}_F(v)$ is the set of edges in F leaving v. The $\mathbf{c}$-*power* of F is defined by

$$p_\mathbf{c}(F) = \sum_{\delta_F(v) \neq \emptyset} \max_{e \in \delta_F(v)} c(e) \qquad \text{if } F \text{ is undirected}$$

$$p_\mathbf{c}(F) = \sum_{\delta^{out}_F(v) \neq \emptyset} \max_{e \in \delta^{out}_F(v)} c(e) \qquad \text{if } F \text{ is directed}$$

Now consider the directed variant of the Network Activation problem when each activating function $f^{uv}(x_u, x_v) = g^{uv}(x_u)$ depends on the weight at u only, and does not depend on x_v; namely, $f^{uv}(x_u, a) = f^{uv}(x_u, b)$ for all x_u, a, b. Under Assumptions 1 and 2, this variant is equivalent to the directed Minimum-Power Network Design problem with edge-costs $c_{uv} = x_u^{uv} = \min\{x_u : g^{uv}(x_u) = 1\}$.

We are interested in Network Activation problems with graph property Π that for every node pair (u, v) ensures a certain number $r(u, v)$ of uv-paths, with the additional property that they cannot share edges and some nodes. For undirected graphs, generalizing the algorithm of Klein and Ravi [5] for Node-Weighted Steiner Forest, Panigrahi [14] gave an $O(\log n)$-approximation algorithm for Steiner Forest Activation and for 2-Connected Subgraph Activation. He asked whether similar results can be obtained for directed graphs, e.g. for the Arborescence Activation or the Strongly Connected Subgraph Activation problems. We answer this question, and moreover, generalize all this to high connectivity, by extending and significantly simplifying the generic approach developed in [11, 13, 10], as well as using some additional ideas.

Definition 2. *For a subset S of nodes in a graph G, let $\lambda_G^S(u, v)$ denote the maximum number of edge-disjoint uv-paths in G such that no two of them have a node in $S \setminus \{u, v\}$ in common. Given connectivity requirements $\mathbf{r} = \{r(u, v) : u, v \in U \subseteq V\}$, we say that G satisfies $\mathbf{r}$ if $\lambda_G^S(u, v) \geq r(u, v)$ for all $u, v \in U$.*

We consider variants of the following problem.

Steiner Network Activation
Instance: A graph $G = (V, E)$, $S \subseteq V$, a family $\{f^{uv}(x_u, x_v) : uv \in E\}$ of activating functions from $\mathbb{R}_+^2$ to $\{0, 1\}$ each, and connectivity requirements $\mathbf{r} = \{r(u, v) : u, v \in U \subseteq V\}$.
Objective: Find a weight assignment $\mathbf{w}$ on V with $\mathbf{w}(V)$ minimum such that the graph $G_{\mathbf{w}} = (V, E_{\mathbf{w}})$ activated by $\mathbf{w}$ satisfies $\mathbf{r}$.

Edge-connectivity is the case $S = \emptyset$, *node-connectivity* is the case $S = V$, and *element-connectivity* is the case $S \cap U = \emptyset$. Let $k = \max\{r(u, v) : u, v \in U\}$ denote the maximum requirement.

The simplest type of connectivity requirements is when $U = \{s, t\}$ and $r(s, t) = k$, namely, when we require k disjoint paths from a source s to the sink t. This gives the k Disjoint Paths Activation problem, which has several variants, depending whether the graph is undirected or directed, and on the choice of S: when $S = \emptyset$ the paths are edge-disjoint, and when $S = V$ the paths are internally-disjoint.

In Steiner Network Activation problems, the following types of requirements are often considered in the literature, c.f. [7, 11, 13, 1, 10].

- *Out-rooted requirements*: there is $s \in V$ such that $r(u, v) > 0$ implies $u = s$. *In-rooted requirements*: there is $s \in V$ such that $r(u, v) > 0$ implies $v = s$.
- *Subset uniform requirements*: $r(u, v) = k$ for all $u, v \in U \subseteq V$ and $r(u, v) = 0$ otherwise; *uniform requirements* is the case when $U = V$, namely, when $r(u, v) = k$ for all $u, v \in V$.

A graph is: *k-out-connected from s* if it contains k internally-disjoint paths from s to every $v \in V \setminus \{s\}$, and *k-in-connected to s* if it contains k internally-disjoint paths from every $v \in V \setminus \{s\}$ to s. A graph with at least $k+1$ nodes is *k-connected* if it contains k internally-disjoint paths from every node to the other. In the k-Out/In-connected Subgraph Activation problem $G_{\mathbf{w}}$ is required to be k-out/in-connected from/to a given root s; this is the case of *uniform out/in-rooted requirements* and $S = V$. In the k-Connected Subgraph Activation problem $G_{\mathbf{w}}$ is required to be k-connected; this is the case of *uniform requirements* and $S = V$.

In Steiner Network Activation Augmentation problems we are given a graph J such that $r(u,v) - \lambda_J^S(u,v) \leq 1$ for all $u, v \in V$, and seek a minimum weight assignment $\mathbf{w}$ such that the graph $(V, E_J \cup E_{\mathbf{w}})$ satisfies $\mathbf{r}$. Equivalently, given a set $\mathcal{T} = \{uv : r(u,v) - \lambda_J^S(u,v) = 1\}$ of *demand-edges* (the edges in $\mathcal{T}$ are undirected or directed, depending whether J is undirected or directed), we require that $\lambda^S_{J \cup E_{\mathbf{w}}}(u,v) \geq \lambda_J^S(u,v) + 1$ for all $uv \in \mathcal{T}$. It is known that a ρ-approximation for Steiner Network Activation Augmentation implies a $k\rho$-approximation for Steiner Network Activation. On the other hand, Steiner Network Activation Augmentation is a particular case of the Bifamily Edge-Cover Activation problem defined below (c.f. [7, 10]). We need some definitions to present this problem.

Definition 3. *An ordered pair $\hat{X} = (X, X^+)$ of subsets of a groundset V is called a* biset *if $X \subseteq X^+$; X is the* inner part*, X^+ is the* outer part*, and $\Gamma(\hat{X}) = X^+ \setminus X$ is the* boundary *of $\hat{X}$. A biset-family is called a* bifamily *if for any $\hat{X}, \hat{Y} \in \mathcal{F}$ the following holds: $X = Y$ implies $X^+ = Y^+$ (bijectiveness), and $X \subseteq Y$ implies $X^+ \subseteq Y^+$ (monotonicity).*

Definition 4. *The intersection and the union of two bisets $\hat{X}$ and $\hat{Y}$ is defined by $\hat{X} \cap \hat{Y} = (X \cap Y, X^+ \cap Y^+)$ and $\hat{X} \cup \hat{Y} = (X \cup Y, X^+ \cup Y^+)$. The biset $\hat{X} \setminus \hat{Y}$ is defined by $\hat{X} \setminus \hat{Y} = (X \setminus Y^+, X^+ \setminus Y)$ A bifamily $\mathcal{F}$ is:*

- uncrossable *if $\hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y} \in \mathcal{F}$ or $\hat{X} \setminus \hat{Y}, \hat{Y} \setminus \hat{X} \in \mathcal{F}$ for any $\hat{X}, \hat{Y} \in \mathcal{F}$.*
- intersecting *if $\hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y} \in \mathcal{F}$ for any $\hat{X}, \hat{Y} \in \mathcal{F}$ with $X \cap Y \neq \emptyset$.*
- *a* ring-bifamily *if $\mathcal{F}$ is an intersecting bifamily and if the intersection of the inner parts of all bisets in $\mathcal{F}$ is non-empty.*

A directed/undirected edge e covers a biset $\hat{X}$ if it goes from $V \setminus X^+$ to X. An edge set I covers a bifamily $\mathcal{F}$ if every $\hat{X} \in \mathcal{F}$ is covered by some edge $e \in I$. We consider the following generic problem.

Bifamily Edge-Cover Activation
Instance: A graph $G = (V, E)$, a family $\{f^{uv}(x_u, x_v) : uv \in E\}$ of activating functions, and a bifamily $\mathcal{F}$ on V.
Objective: Find a minimum-weight assignment $\mathbf{w}$ on V such that $E_{\mathbf{w}}$ covers $\mathcal{F}$.

Given an instance of Steiner Network Activation Augmentation, the corresponding Bifamily Edge-Cover Activation instance is obtained as follows. To avoid considering "mixed" cuts that contain both nodes and edges, we may assume that $st \notin E_J$ for all $st \in \mathcal{T}$. One way to achieve this is to subdivide every edge $st \in E_J$ with $st \in \mathcal{T}$ by a *dummy node*, and to add all these dummy nodes to S.

For $X \subseteq V$, let X^+ be the union of X and the set of those nodes that have a neighbor in X. Let us say that a biset $\hat{X} = (X, X^+)$ is *tight* if $\Gamma(\hat{X}) \subseteq S$ and there exists $st \in \mathcal{T}$ that covers $\hat{X}$ such that $|\Gamma(\hat{X})| = \lambda_J^S(s,t)$.

By Menger's Theorem, $J \cup E_{\mathbf{w}}$ satisfies the connectivity requirements if, and only if, $E_{\mathbf{w}}$ covers the family $\mathcal{F}_{J,\mathcal{T}}$ of tight bisets, c.f. [7]. It is easy to see that $\mathcal{F}_{J,\mathcal{T}}$ is a bifamily. This bifamily is uncrossable in the case of element-connectivity requirements [3], and intersecting in the case of out-rooted uniform requirements [4]. In the case of undirected graphs and out/in-rooted requirements, it is sufficient to cover the bifamily $\mathcal{F}^s_{J,\mathcal{T}} = \{\hat{X} \in \mathcal{F}_{J,\mathcal{T}} : s \notin X^+\}$. This bifamily is intersecting for rooted uniform requirements, c.f. [4].

A polynomial time implementation of our algorithms requires that certain queries related to $\mathcal{F}$ can be answered in polynomial time. Given an edge set I on V, the *residual bifamily* $\mathcal{F}(I)$ *of* $\mathcal{F}$ (w.r.t. I) consists of all members of $\mathcal{F}$ that are uncovered by the edges of I. It is easy to verify that if $\mathcal{F}$ is uncrossable, then so is $\mathcal{F}(I)$, for any I, c.f. [10].

Definition 5. *A set $C \in \{X : \hat{X} \in \mathcal{F}\}$ is a* core *(or an $\mathcal{F}$-*core*) of a bifamily $\mathcal{F}$, if C does not contain as subsets two distinct inclusion-minimal members of the set-family $\{X : (X, X^+) \in \mathcal{F}\}$. An inclusion-minimal (inclusion-maximal) core is a* min-core *(*max-core*). Let $\mathcal{C}_{\mathcal{F}}$ ($\mathcal{M}_{\mathcal{F}}$) denote the set-family of min-cores (max-cores) of $\mathcal{F}$.*

Assumption A. Given the inner part X of a biset $\hat{X} \in \mathcal{F}$, the outer part X^+ of $\hat{X}$ can be computed in polynomial time.

Assumption B. For any edge set I on V, the families $\mathcal{C}_{\mathcal{F}(I)}$ of min-cores and $\mathcal{M}_{\mathcal{F}(I)}$ of max-cores of $\mathcal{F}(I)$ can be computed in polynomial time.

Using standard max-flow min-cut methods, it is easy to see that Assumptions A and B hold for the family of tight bisets. Summarizing, we have the following.

Corollary 1. *Given an instance of* Steiner Network Activation Augmentation *(with $st \notin E_J$ for all $\{s,t\} \in \mathcal{T}$), $J \cup E_{\mathbf{w}}$ satisfies the requirements if, and only if, $E_{\mathbf{w}}$ covers the bifamily $\mathcal{F}_{J,\mathcal{T}}$ of tight bisets. Furthermore, Assumptions A and B hold for $\mathcal{F}_{J,\mathcal{T}}$.* □

For a graph (V,F) let $\Delta_F = \max_{v \in V} |\delta_F(v)|$ denote the maximum number of edges in F incident to a node in (V,F). Our first result is the following simple relation between Network Activation and Edge-Costs Network Design problems.

Theorem 1. *Suppose that for some graph property Π the following holds.*

- *There exists an integer Δ such that $\Delta_F \leq \Delta$ holds for any inclusion minimal edge-set F with $(V,F) \in \Pi$.*
- Edge-Costs Network Design *with property Π admits a θ-approximation algorithm.*

Then Network Activation *with property Π admits an $\theta\Delta$-approximation algorithm, under Assumptions 1 and 2.*

Theorem 1 has the following consequence (to be proved formally in Section 4).

Corollary 2. *The* k Internally-Disjoint Paths Activation *problem admits a* 2*-approximation algorithm, if Assumption 1 holds and if* D^s, D^t *are polynomial in* n.

The main result of this paper is the following.

Theorem 2. *Under Assumptions 1,3,A,B,* Bifamily Edge-Cover Activation *admits the following approximatio ratios:* 2 *for ring bifamilies, and* $O(\log |\mathcal{C}_{\mathcal{F}}|)$ *for undirected graphs with uncrossable* $\mathcal{F}$*, or for directed graphs with intersecting* $\mathcal{F}$.

In $[6, 11, 13, 10, 8]$ it is shown how various Steiner Network problems can be decomposed into Bifamily Edge-Cover problems. Using this, we deduce from Theorem 2 the following result (to be proved formally in Section 4), that for the particular case of directed graphs and $k = 1$ answers a question from [14].

Theorem 3. Steiner Network Activation *problem admit the following approximation ratios under Assumptions 1 and 3. For both undirected/directed graphs,* k-Out/In-connected Subgraph Activation *and* k-Connected Subgraph Activation *admit ratio* $O(k \log n)$*, and* k Disjoint Paths Activation *admits ratio* $2k$*. For undirected graphs, the following ratios are also achievable:*

- $O(\log |U|)$ *for requirements in* $\{0, 1, 2\}$.
- $O(k \log |U|)$ *for element-connectivity requirements.*
- $O(k^2 \log |U|)$ *for rooted requirements and for subset uniform requirements.*
- $O(k^4 \log^2 |U|)$ *for general requirements.*

2 Proof of Theorem 1

Recall that by Assumption 2, for every $e = uv \in E$, we can compute in polynomial time some optimal weight function $\mathbf{x}^e$ activating e, with values $x^e_u = x^{uv}_u$ on u and $x^e_v = x^{uv}_v$ on v, and zero otherwise; hence $x^e_u + x^e_v = \min\{x_u + x_v : f^{uv}(x_u, x_v) = 1\}$. In the proof of Theorem 1, the key observation is the following statement, which applies for both directed and undirected graphs.

Lemma 1. *Let* $G = (V, E)$ *be a (directed or undirected) graph and let* $E' \subseteq E$. *Let* $\mathbf{w}'$ *be a weight function on* V *defined by* $w'_u = \max\limits_{e \in \delta_{E'}(u)} x^e_u$ *if* $u \in V(E')$ *and* $w'_u = 0$ *otherwise, and let* $\mathbf{c}$ *be a cost function on* E *defined by* $c_e = x^e_u + x^e_v$ *for all* $e = uv \in E$. *Then* $E' \subseteq E_{\mathbf{w}'}$*, and* $\mathbf{w}'(V) \leq \mathbf{c}(E') \leq \Delta_{E'} \cdot \mathbf{w}(V)$ *for any weight function* $\mathbf{w}$ *such that* $E' \subseteq E_{\mathbf{w}}$.

Proof. To see that $E' \subseteq E_{\mathbf{w}'}$, note that $w'_u \geq x^{uv}_u$ and $w'_v \geq x^{uv}_v$ for every $uv \in E'$, by the definition of $\mathbf{w}'$. Hence $uv \in E_{\mathbf{w}'}$, by Assumption 1.

We prove that $\mathbf{w}'(V) \leq \mathbf{c}(E')$. Let D' be a set of directed edges on $V(E')$ obtained from E' by choosing for every $u \in V(E')$ some maximum $\mathbf{c}$-cost edge $e \in \delta_{E'}(u)$ incident to u, and picking into D' the orientation of e with tail u.

Assign cost c'_{uv} to every edge $uv \in D'$ as follows; $c'_{uv} = c_{uv}$ if uv does not belong to a cycle of length 2 of D' and $c'_{uv} = x'_u$ otherwise. It is easy to see that $\mathbf{c}(E') \geq \mathbf{c}'(D')$ and that $\mathbf{c}'(D') \geq \mathbf{w}'(V)$. The statement follows.

Let now $\mathbf{w}$ be any weight function such that $E' \subseteq E_{\mathbf{w}}$. We prove that $\mathbf{c}(E') \leq \Delta_{E'} \cdot \mathbf{w}(V)$. Note that $c_{uv} \leq w_u + w_v$ for every $uv \in E'$, by the definition of $\mathbf{c}$ and since $E' \subseteq E_{\mathbf{w}}$. This implies:

$$\mathbf{c}(E') = \sum_{uv \in E'} c_{uv} \leq \sum_{uv \in E'} (w_u + w_v) = \sum_{u \in V} |\delta_{E'}(u)| w_u \leq \Delta_{E'} \cdot \mathbf{w}(V) \ .$$

This concludes the proof of the lemma. □

We now finish the proof of Theorem 1. The algorithm is as follows. With edge-cost function $\mathbf{c}$ as in Lemma 1, compute an α-approximate $\mathbf{c}$-cost solution E' satisfying the property Π, and return the weight function $\mathbf{w}'$ as in Lemma 1. This can be done in polynomial time, by Assumption 2. By Lemma 1, $E' \subseteq E_{\mathbf{w}'}$, hence the solution $\mathbf{w}'$ returned is feasible, namely, $E_{\mathbf{w}'}$ satisfies Π, by the monotonicity of Π.

Let $\mathbf{w}$ be an optimal solution to Network Activation, and let $F \subseteq E_{\mathbf{w}}$ be an inclusion minimal edge set that satisfies Π. By the assumption, $\Delta_F \leq \Delta$. Using Lemma 1 and the fact that E' is an θ-approximate $\mathbf{c}$-cost solution, while $(V, F) \in \Pi$, we get:

$$\mathbf{x}'(V) \leq \mathbf{c}(E') \leq \theta \cdot \mathbf{c}(F) \leq \theta \cdot \Delta_F \cdot \mathbf{w}(V) \leq \theta \cdot \Delta \cdot \mathbf{w}(V) \ .$$

The proof of Theorem 1 is now complete.

3 Proof of Theorem 2

We need the concept "spider-cover" introduced in [11, 13]. For a bifamily $\mathcal{F}$ on V, a min-core $C \in \mathcal{C}_{\mathcal{F}}$, and $s \in V$ let

$$\mathcal{F}(C) = \{\hat{X} \in \mathcal{F} : X \supseteq C, X \text{ is an } \mathcal{F}\text{-core}\}$$
$$\mathcal{F}(s, C) = \{\hat{X} \in \mathcal{F}(C) : s \notin X^+\}$$

Definition 6. *Let $\mathcal{F}$ be a bifamily on V and let $\mathcal{C} \subseteq \mathcal{C}_{\mathcal{F}}$. We say that an undirected/directed edge-set S on V is an* $\mathcal{F}(s, \mathcal{C})$-spider-cover *if $s \in V(S)$ and if S can be partitioned into $\mathcal{F}(s, C)$-covers $\{P_C : C \in \mathcal{C}\}$ such that the node sets $\{V(P_C) \setminus \{s\} : C \in \mathcal{C}\}$ are pairwise disjoint. We say that S is an* $\mathcal{F}(\mathcal{C})$-spider-cover*, or a* spider-cover *if $\mathcal{C}$ is clear from the context, if the following holds:*

- *If $|\mathcal{C}| \geq 2$ then there exists $s \in V$ (a* center *of the spider-cover) such that S is an $\mathcal{F}(s, \mathcal{C})$-spider-cover.*
- *If $|\mathcal{C}| = 1$, say $\mathcal{C} = \{C\}$, then S covers $\mathcal{F}(C)$.*

Equivalently, for $|\mathcal{C}| \geq 2$, an $\mathcal{F}(\mathcal{C})$-spider-cover S with a chosen center s is a union of $\mathcal{F}(s, C)$-covers $\{P_C : C \in \mathcal{C}\}$ so that only s can be a common end-node of two of them.

Definition 7. *Let $\mathcal{F}$ be a bifamily on V and let $\mathcal{C} \subseteq \mathcal{C}_{\mathcal{F}}$. We say that a collection $\mathcal{S} = \{S_1, \ldots, S_h\}$ of edge-sets* spider-covers *$\mathcal{C}$ if the following holds:*

- *The node-sets $V(S_1), \ldots, V(S_h)$ are pairwise disjoint.*
- *$\mathcal{C}$ admits a partition $\{\mathcal{C}_1, \ldots, \mathcal{C}_h\}$ such that each S_t is an $\mathcal{F}(\mathcal{C}_t)$-spider-cover.*

In [11] directed covers of intersecting *set-families* are considered, For this case, [11, Theorem 2.3] states that any cover I of $\mathcal{F}$ admits a "tail-disjoint" subpartition that spider-covers a subfamily $\mathcal{C} \subseteq \mathcal{C}_{\mathcal{F}}$ of size at least $|\mathcal{C}| \geq \lceil 2|\mathcal{C}_{\mathcal{F}}|/3 \rceil$; in the setting of [11] this bound is the best possible. [13, Theorem 2.3] states that any (undirected) cover I of an uncrossable *set-family* $\mathcal{F}$ admits a subpartition that spider-covers the entire family $\mathcal{C}_{\mathcal{F}}$ of $\mathcal{F}$-cores. In the case of bifamilies, the following is proved in [10].

Theorem 4 ([10]). *Any undirected cover I of an uncrossable bifamily $\mathcal{F}$ admits a subpartition that spider-covers $\mathcal{C}_{\mathcal{F}}$.*

For the case of *directed covers* of intersecting *bifamilies*, we use a novel method to prove the following. Let us say that a bifamily $\mathcal{F}$ is *simple* if the inner part of every member of $\mathcal{F}$ is a core.

Theorem 5. *Let $\mathcal{F}$ be a simple bifamily such that the $\mathcal{F}$-cores are pairwise disjoint and such that $\mathcal{F}(C)$ is a ring-bifamily for every $C \in \mathcal{C}_{\mathcal{F}}$. Then any directed cover I of $\mathcal{F}$ admits a subpartition that spider-covers a subfamily $\mathcal{C} \subseteq \mathcal{C}_{\mathcal{F}}$ of $\mathcal{F}$-cores of size at least $|\mathcal{C}| \geq \lceil 2|\mathcal{C}_{\mathcal{F}}|/3 \rceil$.*

The following statement is well known, c.f. [10].

Lemma 2. *If a bifamily $\mathcal{F}$ is uncrossable or intersecting, then so is the bifamily $\{\hat{X} \in \mathcal{F} : X$ is an $\mathcal{F}$-core$\}$, the min-cores of $\mathcal{F}$ are pairwise disjoint, and $\mathcal{F}(C)$ is a ring-bifamily for every min-core $C \in \mathcal{C}_{\mathcal{F}}$. In particular, for every min-core C there is a unique max-core containing C.*

Note that Definitions 6 and 7 consider covers only of bisets in $\mathcal{F}$ for which the inner parts are cores, namely, the relevant bifamily is $\{\hat{X} \in \mathcal{F} : X$ is an $\mathcal{F}$-core$\}$; this bifamily is uncrossable if $\mathcal{F}$ is, by Lemma 2. Any uncrossable simple bifamily satisfies the assumptions of Theorem 5, by Lemma 2. Thus Theorem 5 implies the following.

Corollary 3. *Any directed cover I of an intersecting bifamily $\mathcal{F}$ admits a subpartition that spider-covers a subfamily $\mathcal{C} \subseteq \mathcal{C}_{\mathcal{F}}$ of $\mathcal{F}$-cores of size at least $|\mathcal{C}| \geq \lceil 2|\mathcal{C}_{\mathcal{F}}|/3 \rceil$.*

We now prove Theorem 5, and at the end of this section describe how Theorem 4 and Corollary 3 imply Theorem 2.

For an edge-set I and a biset $\hat{X}$ on a node set V let $\zeta_I(\hat{X})$ denote the set of edges in I covering $\hat{X}$. We need the following (known) statement.

Lemma 3. *Let I be an inclusion-minimal directed cover of a ring bifamily $\mathcal{F}$ and let C be the min-core of $\mathcal{F}$. Then $|\zeta_I(\hat{C})| = 1$.*

Proof. Clearly, $|\zeta_I(\hat{C})| \geq 1$ since I covers $\mathcal{F}$ and since $\hat{C} \in \mathcal{F}$. Suppose to the contrary that there are distinct $e, f \in \zeta_I(\hat{C})$. By the minimality of I, there are $\hat{W}_e, \hat{W}_f \in \mathcal{F}$ such that $\zeta_I(\hat{W}_e) = \{e\}$ and $\zeta_I(\hat{W}_f) = \{f\}$. There is an edge in I covering $\hat{W}_e \cup \hat{W}_f$, because $\hat{W}_e \cup \hat{W}_f \in \mathcal{F}$. This edge must be one of e, f, because if for arbitrary bisets $\hat{X}, \hat{Y}$ an edge covers $\hat{X} \cup \hat{Y}$ then it also covers one of $\hat{X}, \hat{Y}$. Each of e, f covers $\hat{W}_e \cap \hat{W}_f$, because each of e, f has an endnode in C and $C \subseteq W_e \cap W_f$. Consequently, one of e, f covers both $\hat{W}_e \cap \hat{W}_f$ and $\hat{W}_e \cup \hat{W}_f$. However, if for arbitrary bisets $\hat{X}, \hat{Y}$ an edge covers both $\hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y}$ then it cover both $\hat{X}$ and $\hat{Y}$. Hence one of e, f covers both $\hat{W}_e, \hat{W}_f$. This is a contradiction, since $\zeta_I(\hat{W}_e) = \{e\}$, $\zeta_I(\hat{W}_f) = \{f\}$, and $e \neq f$. $\square$

The proof of the following key statement is similar to the proof of [11, Lemma 2.6] where directed covers of ring *set-families* are considered.

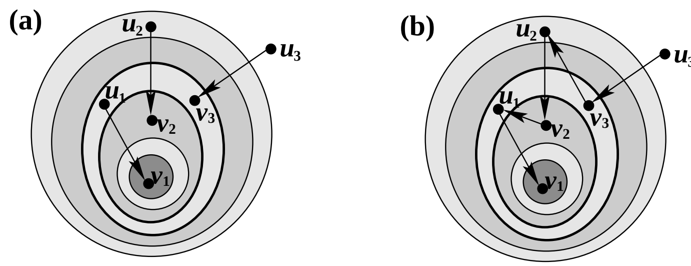

Fig. 1. (a) Illustration to Lemma 4; inner parts of the bisets are shown by darker ellipses. (b) Construction of the path P_C.

Lemma 4. *Let I be an inclusion-minimal directed cover of a ring bifamily $\mathcal{F}$. There exists an ordering $e_1, \ldots, e_q$ of I and a nested family $C_1 \subset \cdots \subset C_q$ of sets in $\{X : \hat{X} \in \mathcal{F}\}$ such that for every $j = 1, \ldots, q$ the following holds (see Fig. 1(a)).*

(i) *C_j is the min-core of $\mathcal{F}_{I_{j-1}}$, where $I_j = \{e_1, \ldots, e_j\}$ and $I_0 = \emptyset$, and e_j is the unique edge in I covering $\hat{C}_j$.*
(ii) *If $e_j = u_j v_j$ where $v_j \in C_j$, then I_j is an $\mathcal{F}(u_j, C)$-cover and I_{j-1} is an $\mathcal{F}(v_j, C)$-cover, where C is the min-core of $\mathcal{F}$.*

Proof. Let $C_1 = C$. By Lemma 3 there is a unique edge $e_1 \in I$ covering $\hat{C}_1$. If $e_1 = u_1 v_1$ where $v_1 \in C_1$, then clearly $I_0 = \emptyset$ is an $\mathcal{F}(v_1, C)$-cover and $I_1 = \{e_1\}$ is an $\mathcal{F}(u_1, C)$-cover. Thus if e_1 covers $\mathcal{F}$ we are done. Otherwise, let C_2 be the min-core of $\mathcal{F}_{I_1}$. Then $C_1 \subset C_2$. Let $e_2 = u_2 v_2$ be the edge in I covering $\hat{C}_2$, where $v_2 \in C_2$. As C_2 is the min-core of $\mathcal{F}_{I_1}$ and $v_2 \in C_2$, it follows that I_1 is an $\mathcal{F}(v_2, C)$-cover and $I_2 = I_1 \cup \{e_2\}$ is an $\mathcal{F}(u_2, C)$-cover. We can continue this process until some edge e_q covers $\mathcal{F}_{I_{q-1}}$. Namely, given the edge set $I_{j-1} = \{e_1, \ldots, e_{j-1}\}$ that still does not cover $\mathcal{F}$, C_j is the min-core of $\mathcal{F}_{I_{j-1}}$, and $e_j = u_j v_j$ is the edge in I covering $\hat{C}_j$, where $v_j \in C_j$. Then $C_{j-1} \subset C_j$. As C_j is a min-core of $F_{I_{j-1}}$ and $v_j \in C_j$, it follows that I_{j-1} is an $\mathcal{F}(v_j, C)$-cover and I_j is an $\mathcal{F}(u_j, C)$-cover. The lemma follows. $\square$

Recall that a *directed spider* is an arborescence (directed tree) with at most one node (the root) of outdegree ≥ 2. The following statement is an immediate consequence from [1, Theorem 4].

Lemma 5 (Chuzhoy and Khanna [1]). *Let $\mathcal{Q}$ be a set of directed simple paths ending at a set $A = \{a_P : P \in \mathcal{P}\}$ of distinct nodes. There exists $\mathcal{P} \subseteq \mathcal{Q}$ with $\mathcal{P} \geq \lceil 2|\mathcal{Q}|/3 \rceil$ such that the following holds. Every $P \in \mathcal{P}$ has a subpath P' (possibly of length zero) that ends at a_P and has no internal node in A, such that in the (simple) graph J induced by the subpaths $\{P' : P \in \mathcal{P}\}$, every connected component is either a directed spider with at least 2 nodes in A, or is a path in $\mathcal{P}$.*

Proof of Theorem 5. For every $C \in \mathcal{C}_{\mathcal{F}}$ fix some inclusion-minimal cover $I_C \subseteq I$ of $\mathcal{F}(C)$. Let $e_1, \ldots, e_q$ be an ordering of I_C as in Lemma 4, where $e_j = u_j v_j$ is as in the lemma. Obtain a directed path P_C adding for every $j = q, \ldots, 2$ the directed edge $v_j u_{j-1}$, if $v_j \neq u_{j-1}$; hence if $v_j \neq u_{j-1}$ for all j, then the node sequence of P_C is $(u_q, v_q, u_{q-1}, v_{q-1}, \ldots, u_1, v_1)$. Let $a_C = v_1$ and note that $a_C \in C$. Let $\mathcal{Q} = \{P_C : C \in \mathcal{C}_{\mathcal{F}}\}$. As the min-cores of $\mathcal{F}$ are pairwise disjoint, the path in $\mathcal{Q}$ end at distinct nodes. Hence Lemma 5 applies, and thus there exists a subfamily $\mathcal{C} \subseteq \mathcal{C}_{\mathcal{F}}$ of $\mathcal{F}$-cores of size at least $|\mathcal{C}| \geq \lceil 2|\mathcal{C}_{\mathcal{F}}|/3 \rceil$, such that the following holds. Every P_C with $C \in \mathcal{C}$ has a subpath P'_C that ends at a_C, such that if $J_1, \ldots, J_h$ are the connected components of the (simple) graph J induced by the subpaths $\{P'_C : C \in \mathcal{C}\}$, every J_t is either a directed spider with at least 2 nodes in $\{a_C : C \in \mathcal{C}\}$, or is a path in $\mathcal{P}$. For every $t = 1, \ldots, h$ let $\mathcal{C}_t = \{C : v_C \in J_t\}$ and let $S_t = J \cap I$ be the set of those edges $e \in I$ that in J_t. From the construction and Lemma 4 it follows that S_t is an $\mathcal{F}(\mathcal{C}_t)$-spider-cover. Thus the collection $\mathcal{S} = \{S_1, \ldots, S_h\}$ of edge-sets spider-covers $\mathcal{C}$. Since $|\mathcal{C}| \geq \lceil 2|\mathcal{C}_{\mathcal{F}}|/3 \rceil$, Theorem 5 follows. $\square$

We now describe how Theorem 4 and Corollary 3 imply Theorems 2. We use a *Greedy Algorithm* for the following type of problems:

Covering Problem
Instance: A ground-set E and integral functions ν, ω on 2^E, where $\nu(E) = 0$.
Objective: Find $I \subseteq E$ with $\nu(I) = 0$ and with $\omega(I)$ minimized.

In the Covering Problem, the instance functions ν, ω may be given by an evaluation oracle; ν is the *deficiency function* that measures how far is I from being a feasible solution, and ω is the *weight function*. Given a partial solution I, the *density* of a set $S \subseteq E \setminus I$ is $\omega(S)/(\nu(I) - \nu(I \cup S))$. The ρ-Greedy Algorithm starts with $I = \emptyset$, and as long as $\nu(I) \geq 1$, it finds and adds to I an edge-set $S \subseteq E \setminus I$ of density at most $\rho \cdot \mathsf{opt}/\nu(I)$, where opt denotes the optimal solution value. The following statement is known, c.f. [13].

Theorem 6. *For any* Covering Problem *such that ν is decreasing, the ρ-Greedy Algorithm computes a collection $\mathcal{S}$ of subsets of E such that $I = \bigcup_{S \in \mathcal{S}} S$ is a feasible solution and such that $\sum_{S \in \mathcal{S}} \omega(S) \leq \rho \cdot (\ln(\nu(\emptyset)) + 1) \cdot \mathsf{opt}$. Furthermore, if ω is subadditive then $\omega(I) \leq \rho \cdot (\ln(\nu(\emptyset)) + 1) \cdot \mathsf{opt}$.*

In our setting, for $I \subseteq E$, let $\nu(I) = |\mathcal{C}(\mathcal{F}(I))|$ denote the number of min-cores of the residual bifamily $\mathcal{F}(I)$, and let $\omega(I) = \min\{\mathbf{w}(V) : I \subseteq E_{\mathbf{w}}\}$ be an optimal weight assignment that activates I. Clearly, ν is decreasing, and ω is sub-additive.

Unfortunately, we do not have a polynomial time evaluation oracle for the function ω, namely, we do not have a method to compute $\omega(S)$ in polynomial time for a given edge set S. However, we can show a 2-approximate polynomial time evaluation oracle for $\omega(S)$ if S is a spider. Note that if every node in the graph (V, S) has degree at most Δ, then Theorem 1 gives a Δ-approximation for $\omega(S)$ in polynomial time. In particular, we have a 2-approximation if S is a path. If S is a spider, then S has at most one node s of degree ≥ 2, and then with the help of Assumption 3, we can still obtain a 2-approximation for $\omega(S)$ as follows. We "guess" the weight $w_s \in D^s$ of s in some optimal weight assignment inducing S, and update each activating function $f^{sv}(x_s, x_v)$ to $f^{sv}(w_s, x_v)$. Then we apply the algorithm as in Theorem 1 on the obtained instance. For a "correct" guess of w_s our estimation for $\omega(S)$ will be between $\omega(S)$ and $2\omega(S) - w_s$.

Recall that in the Bifamily Edge-Cover problem we eventually need to compute a weight-assignment $\mathbf{w}$ and $I \subseteq E_{\mathbf{w}}$ such that I covers $\mathcal{F}$. To apply the Greedy Algorithm, we will show how to find a weight assignment $\mathbf{w} = \mathbf{w}^S$ and $S \subseteq E_{\mathbf{w}^S}$ (S may not be a spider-cover), such that for some constant ρ the following holds:

$$\frac{\mathbf{w}^S(V)}{\nu(I) - \nu(I \cup S)} \leq \rho \cdot \frac{\mathsf{opt}}{\nu(I)} \ .$$

Note that $\omega(S) \leq \mathbf{w}^S(V)$, hence such S has density at most $\rho \cdot \mathsf{opt}/\nu(I)$. Consequently, we can apply the ρ-Greedy Algorithm to compute a collection $\mathcal{S}$ of subsets of E such that $I = \bigcup_{S \in \mathcal{S}} S$ is a feasible solution (namely, $\nu(I) = 0$) and such that $\sum_{S \in \mathcal{S}} \mathbf{w}^S(V) \leq \rho \cdot (\ln(\nu(\emptyset)) + 1) \cdot \mathsf{opt}$. Setting $w(v) = \max_{S \in \mathcal{S}} w_v^S$ (or even $w_v = \sum_{S \in \mathcal{S}} w_v^S$) for every $v \in V$ and $I = \bigcup_{S \in \mathcal{S}} S$ gives a weight assignment $\mathbf{w}$ and a feasible solution $I \subseteq E_{\mathbf{w}}$ as required.

For simplicity of exposition, it is sufficient to consider the case $I = \emptyset$. We assume that E is a feasible solution, thus $\nu(E) = 0$. Let $\nu = \nu(\emptyset)$. Theorem 2 will be proved if we prove the following.

Lemma 6. *There exists an algorithm that given an instance of* Bifamily Edge-Cover Activation *with either undirected E and uncrossable $\mathcal{F}$, or with directed E and intersecting $\mathcal{F}$, finds under Assumptions 1,3,A,B in polynomial time a weight-assignment $\mathbf{w}$ and $S \subseteq E_{\mathbf{w}}$ (S may not be a spider-cover) such that*

$$\frac{\mathbf{w}(V)}{\nu - \nu(S)} \leq 9 \cdot \frac{\mathsf{opt}}{\nu} \ .$$

In the rest of this section we prove Lemma 6.

The following statement follows from Theorem 4 and Corollary 3 by a standard averaging argument.

Proposition 1. *Any optimal $\mathcal{F}$-cover contains an $\mathcal{F}(\mathcal{C}')$-spider-cover S' for some $\mathcal{C}' \subseteq \mathcal{C}_{\mathcal{F}}$, such that $\omega(S')/|\mathcal{C}'| \leq 3/2 \cdot \mathsf{opt}/\nu$.* $\square$

Using Proposition 1, we show under Assumptions 1,3,A,B how to find in polynomial time an edge-set S (which may not be a spider cover) and a weight-assignment $\mathbf{w}^S$ as in Lemma 6; note that Proposition 1 only establishes existence of a spider-cover S' of low density, but does not implies an algorithm for finding such S'.

Lemma 7 ([11, 10]). *Let $\mathcal{F}$ be a biset-family on V, let $\mathcal{C} \subseteq \mathcal{C}_{\mathcal{F}}$, and let S be a directed/undirected edge set on V such that the following holds.*

- *If $|\mathcal{C}| \geq 2$ then there is $s \in V$ such that S is a $\mathcal{F}(s, C)$-cover for every $C \in \mathcal{C}$.*
- *If $|\mathcal{C}| = 1$, say $\mathcal{C} = \{C\}$, then S covers all $\mathcal{F}$-cores containing C.*

If the min-cores of $\mathcal{F}$ and of $\mathcal{F}_S$ are pairwise disjoint then $\nu - \nu(S) \geq |\mathcal{C}|/3$.

It is known that under Assumptions A and B, a standard primal-dual algorithm computes a minimum-cost undirected/directed edge-cover of a ring bifamily $\mathcal{F}$. Furthermore, if P is an inclusion-minimal cover of $\mathcal{F}$ then $\Delta_P \leq 2$; for undirected P this is proved in the full version of [10], while for directed P easily follows from Lemma 4. Thus $\mathsf{Bifamily\ Edge\text{-}Cover\ Activation}$ with ring-bifamily $\mathcal{F}$ admits a 2-approximation algorithm. Combined with Theorem 1 and Lemma 2 we obtain the following.

Corollary 4. *Given an instance of* Bifamily Edge-Cover Activation *with either undirected E and uncrossable $\mathcal{F}$ or with directed E and intersecting $\mathcal{F}$, $s \in V$, and $C \in \mathcal{C}_{\mathcal{F}}$, the problem of finding an optimal weight assignment $\mathbf{w}$ on V such that $E_{\mathbf{w}}$ covers $\mathcal{F}(s, C)$ admits a 2-approximation algorithm, under Assumptions A and B.*

Now we describe the algorithm that finds S and $\mathbf{w}$ as in Lemma 6. The algorithm is essentially the same as the one in [13, 10] for the node-weighted case, except that we "guess" not only the center s of an optimal density spider but also the corresponding weight assignment $w_s \in D^s$ of s (for min-power problems, the approach of "guessing" both s and the power of s was used in [11]). For every fixed "guess" $s \in V$ and $w_s \in D^s$ compute a weight assignment $\mathbf{w} = \mathbf{w}^{(s,w_s)}$, a set of min-cores $\mathcal{C} = \mathcal{C}^{(s,w_s)}$, and an $\mathcal{F}(s, \mathcal{C})$-cover $S = S^{(s,w_s)} \subseteq E_{\mathbf{w}}$ as follows. Set temporarily the weight of s to w_s, and update each activating function $f^{sv}(x_s, x_v)$ to $g^{sv}(x_v) = f^{sv}(w_s, x_v)$ by setting $x_s = w_s$. For every $C \in \mathcal{C}_{\mathcal{F}}$ let $\mathbf{w}^C$ be the weight assignment and P^C the $\mathcal{F}(s, C)$-cover computed by the 2-approximation algorithm as in Corollary 4, with weight of s fixed to w_s. Let $W_C = \mathbf{w}^C(V) - w_s$, where $W_C = \infty$ if P^C does not exist. Sort the members of $\mathcal{C}_{\mathcal{F}}$ by increasing weight, say $W_{C_1} \leq W_{C_2} \leq \ldots \leq W_{C_q}$. Let M_i be the (unique, by Lemma 2) max-core containing C_i. Let σ_j be defined as follows:

- $\sigma_1 = w_s + \min\{W_{C_i} : s \in V \setminus M_i^+\}$ if $\{C_i : s \in V \setminus M_i^+\} \neq \emptyset$ and $\sigma_1 = \infty$ otherwise.
- $\sigma_j = W_j/j$ where $W_j = w_s + \sum_{i=1}^{j} W_{C_i}$, $j = 2, \ldots, q$.

Note that $\sigma_j \leq 2 \cdot \frac{\mathbf{w}'(V)}{j}$ for any weight assignment $\mathbf{w}'$ with $w'_s = w_s$ such that $E_{\mathbf{w}'}$ contains an $\mathcal{F}(s, \mathcal{C}')$-spider-cover S' with $|\mathcal{C}'| = j$.

Next we find an index j for which σ_j is minimum, which determines the corresponding weight assignment $\mathbf{w} = \mathbf{w}^{(s,w_s)}$, the set of min-cores $\mathcal{C} = \mathcal{C}^{(s,w_s)}$, and the $\mathcal{F}(s, \mathcal{C})$-cover $S = S^{(s,w_s)} \subseteq E_{\mathbf{w}}$. If $j = 1$ then $\mathbf{w} = \mathbf{w}^{C_i}$, $S = P^{C_i}$, and $\mathcal{C} = \{C_i\}$, where i is the index for which the minimum is attained in the definition of σ_1. If $j \geq 2$ then $w_v = \max\limits_{i \leq j} w_v^{C_i}$ for all $v \in V$, $S = \bigcup\limits_{i=1}^{j} P^{C_i}$ and $\mathcal{C} = \{C_1, \ldots, C_j\}$.

We compute such triple $\mathbf{w}^{(s,w_s)}$, $\mathcal{C}^{(s,w_s)}$, and $S^{(s,w_s)}$, for every $s \in V$ and $w_s \in D^s$. Then, among all triples computed we choose one tripple $\mathbf{w}, \mathcal{C}, S$ with $\frac{\mathbf{w}^{(S_s,w_s)}(V)}{|\mathcal{C}^{(s,w_s)}|}$ minimum. For this choice we have $\frac{\mathbf{w}(V)}{|\mathcal{C}|} \leq 2 \cdot \frac{\omega(S')}{|\mathcal{C}'|}$ for any $\mathcal{F}(\mathcal{C}')$-spider-cover S'. In particular, if S' is as in Proposition 1, then $\frac{\mathbf{w}(V)}{|\mathcal{C}|} \leq 2 \cdot \frac{\omega(S')}{|\mathcal{C}'|} \leq 3 \cdot \frac{\mathsf{opt}}{\nu}$. On the other hand, $\frac{\mathbf{w}(V)}{\nu - \nu(S)} \leq 3 \cdot \frac{\mathbf{w}(V)}{|\mathcal{C}|}$, by Lemma 7. Consequently, $\frac{\mathbf{w}(V)}{\nu - \nu(S)} \leq 9 \cdot \frac{\mathsf{opt}}{\nu}$, as required.

Time complexity. The time complexity to compute S as in Lemma 6 is the number $|V| \cdot \max\limits_{s \in V} |D^s|$ of "guesses" of the pair (s, w_s) (polynomial by Assumption 3) multiplied by the following: the time required to compute the family $\mathcal{C}_{\mathcal{F}}$ (polynomial by Assumption 5), plus $n|\mathcal{C}_{\mathcal{F}}|$ times the time required to check whether a given node v belongs to the max-$\mathcal{F}$-core M containing a give min-core C (polynomial by Assumptions A and B) plus $n|\mathcal{C}_{\mathcal{F}}|$ times the time required to apply the (polynomial time) algorithm in Corollary 4.

The proof of Lemma 6, and thus also of Theorem 2 is complete.

4 Proof of Theorem 3

We start by proving Corollary 2. Note that in a graph that consists of k internally-disjoint st-paths, the degree of every node distinct from s, t is at most 2. Thus the following algorithm computes a 2-approximate solution to the k Internally-Disjoint Paths Activation problem. We "guess" the weights w_s of s and w_t of t in some optimal weight-assignment, update the activating functions of edges incident to s and to t accordingly, and apply the algorithm from Theorem 1.

We also need the following statements that follows from Theorem 2 by elementary constructions.

Corollary 5. *Under Assumptions 1,3,A,B, the directed problem of finding a minimum-weight assignment* $\mathbf{w}$ *such that the reverse edge set of* $E_{\mathbf{w}}$ *covers an intersecting bifamily* $\mathcal{F}$*, admits an* $O(\log |\mathcal{C}_{\mathcal{F}}|)$*-approximation algorithm.*

Proof. It is easy to see that a weight-assignment $\mathbf{w}$ is a feasible solution to an instance of the problem in the corollary if, and only if, $\mathbf{w}$ is a feasible solution to an instance of Bifamily Edge-Cover Activation obtained by replacing every edge $uv \in E$ by the edge vu with activating function $g^{vu}(x_v, x_u) = f^{uv}(x_u, x_v)$. Thus the statement follows from Theorem 2. $\square$

As was mentioned in the Introduction, for all the types of requirements considered, a ρ-approximation for Steiner Network Activation Augmentation implies a $k\rho$-approximation for Steiner Network Activation. Recall also that by Corollary 1, Steiner Network Activation Augmentation is reducible to the problem of covering the bifamily $\mathcal{F}_{J,\mathcal{T}}$ of tight bisets, or the bifamily $\mathcal{F}^s_{J,\mathcal{T}} = \{\hat{X} \in \mathcal{F}_{J,\mathcal{T}} : s \notin X\}$ in the case of out-rooted requirements.

Consider the directed/undirected k-Outconnected Subgraph Activation problem. In the augmentation version, the initial graph J is $(k-1)$-outconnected from s and $J \cup E_{\mathbf{w}}$ should be k-outconnected from s. The corresponding bifamily $\mathcal{F}^s_{J,\mathcal{T}}$ is intersecting in the directed case [4] and uncrossable in the undirected case [3]. Hence the augmentation version admits ratio $O(\log n)$, by Theorem 2. This implies an $O(k \log n)$-approximation for directed/undirected k-Outconnected Subgraph Activation. By Corollary 5, we have the same ratio for k-Inconnected Subgraph Activation. The ratio for directed/undirected k-Out/In-connected Subgraph Activation in Theorem 3 follows.

Consider the undirected/directed k-Connected Subgraph Activation problem. Our algorithm is a modification of the $O(k)$-approximation algorithm of [6] for the Minimum-Cost k-Connected Subgraph problem. For undirected graphs the algorithm is as follows.

1. Let $R \subseteq V$ be a subset of k nodes, so $|R| = k$. Construct a graph G' by adding to G a new node s and new edges $\{sv : v \in R\}$ with $f^{sv}(x_s, x_v) = 1$ for all $v \in R$.
2. Compute an $O(k \log n)$-approximate weight-assignment $\mathbf{w}$ such that $(V, E_{\mathbf{w}})$ is k-outconnected from s.
3. Let F be an inclusion minimal set of edges on R such that $(V, E_{\mathbf{w}}) \cup F$ is k-connected. Using the 2-approximation algorithm from Corollary 2, compute for every $uv \in F$ a 2-approximate weight-assignment $\mathbf{w}^{uv}$ such that $E_{\mathbf{w}^{uv}}$ contains k internally-disjoint uv-paths.
4. Output $\mathbf{w} + \sum\limits_{uv \in F} \mathbf{w}^{uv}$.

It is known and is shown in [6] that F as at Step 3 exists and is a forest on R. Hence $|F| \leq k - 1 = O(k)$. Consequently, the approximation ratio is $O(k \log n) + 2|F| = O(k \log n) + O(k) = O(k \log n)$.

The algorithm for directed graphs is similar, except that at step 2 we require that $(V, E_{\mathbf{w}})$ is both k-out-connected and k-inconnected to R. In this case we have $|F| \leq 2k - 1$, hence the approximation ratio is still $O(k \log n)$.

Other ratios in Theorem 3 are identical to the best known ones for the undirected Node-Weighted Steiner Network [13, 10, 8], and they are derived from Theorem 2 in the same way as the ratios in [13, 10, 8] are derived.

5 Conclusions

This paper generalizes a line of research of the author initiated in the conference version of [11] in 2006 on min-power and node-weighted connectivity problems. For the more general Steiner Network Activation problem, we now have ratio $O(k \log n)$ for k-Out/In-connected Subgraph Activation and k-Connected Subgraph Activation, for both undirected and directed graphs. For directed graphs, this solves a question from [14] for $k = 1$, and for the min-power case and k arbitrary this solves an open question from [11]. Except the undirected k-Outconnected Subgraph Activation and k-Connected Subgraph Activation problems, whose min-power variants admit an $O(\log k)$-approximation algorithm [2], our ratios match the best known ones for the easier min-power or the node-weighted problems. Our results rely on Theorem 1, Theorem 4 from [10], and Theorem 5 proved in this paper. In fact, the new unifying and simple approach (modulo the non-trivial technical Lemma 5 by [1]) in the proof of Theorem 5, can be used to prove a slightly weaker variant of Theorem 4 from [10] (undirected graphs and uncrossable bifamilies), as well as all the other previous "Spider-Cover Decomposition Theorems" from [11, 13].

There are still several open problems in the field, which is appropriate to state for the easier min-power and node-weighted problems, see [2, 10, 13, 9, 12].

References

1. J. Chuzhoy and S. Khanna. Algorithms for single-source vertex connectivity. In *FOCS*, pages 105–114, 2008.
2. N. Cohen and Z. Nutov. Approximating minimum-power edge-multicovers. Manuscript, 2011.
3. L. Fleischer, K. Jain, and D. Williamson. Iterative rounding 2-approximation algorithms for minimum-cost vertex connectivity problems. *J. of Computing and System Sciences*, 72(5):838–867, 2006.
4. A. Frank. Rooted k-connections in digraphs. *Discrete Applied Math.*, 157(6):1242–1254, 2009.
5. P. Klein and R. Ravi. A nearly best-possible approximation algorithm for node-weighted Steiner trees. *J. of Algorithms*, 19:104115, 1995.
6. G. Kortsarz and Z. Nutov. Approximating node-connectivity problems via set covers. *Algorithmica*, 37:75–92, 2003.
7. G. Kortsarz and Z. Nutov. *Approximating minimum-cost connectivity problems,* Ch. 58 in Approximation algorithms and Metaheuristics, Editor T. F. Gonzalez. Chapman & Hall/CRC, 2007.
8. G. Kortsarz and Z. Nutov. Subset and rooted connectivity problems revisited. Manuscript, 2011.
9. Y. Lando and Z. Nutov. On minimum power connectivity problems. *J. Discrete Algorithms*, 8(2):164–173, 2010.
10. Z. Nutov. Approximating minimum cost connectivity problems via uncrossable bifamilies and spider-cover decompositions. In *FOCS*, pages 417–426, 2009.
11. Z. Nutov. Approximating minimum power covers of intersecting families and directed edge-connectivity problems. *Theoretical Computer Science*, 411(26-28):2502–2512, 2010.

12. Z. Nutov. Approximating minimum-power k-connectivity. *Ad Hoc & Sensor Wireless Networks*, 9:129–137, 2010.
13. Z. Nutov. Approximating steiner networks with node-weights. *SIAM J. on Computing*, 37(7):3001–3022, 2010.
14. D. Panigrahi. Survivable network design problems in wireless networks. In *SODA*, 2011.