Approximating subset k-connectivity problems

Zeev Nutov

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

Abstract. A subset $T \subseteq V$ of terminals is k-connected to a root s in a directed/undirected graph J if J has k internally-disjoint vs-paths for every $v \in T$; T is k-connected in J if T is k-connected to every $s \in T$. We consider the Subset k-Connectivity Augmentation problem: given a graph $G = (V,E)$ with edge/node-costs, node subset $T \subseteq V$, and a subgraph $J = (V, E_J)$ of G such that T is k-connected in J, find a minimum-cost augmenting edge-set $F \subseteq E \setminus E_J$ such that T is $(k+1)$-connected in $J \cup F$. The problem admits trivial ratio $O(|T|^2)$. We consider the case $|T| > k$ and prove that for directed/undirected graphs and edge/node-costs, a ρ-approximation for Rooted Subset k-Connectivity Augmentation implies the following ratios for Subset k-Connectivity Augmentation:
(i) $b(\rho + k) + \left(\frac{3|T|}{|T|-k}\right)^2 H\left(\frac{3|T|}{|T|-k}\right)$ and (ii) $\rho \cdot O\left(\frac{|T|}{|T|-k} \log k\right)$,
where $b = 1$ for undirected graphs and $b = 2$ for directed graphs, and $H(k)$ is the kth harmonic number. The best known values of ρ on undirected graphs are $\min\{|T|, O(k)\}$ for edge-costs and $\min\{|T|, O(k \log |T|)\}$ for node-costs; for directed graphs $\rho = |T|$ for both versions. Our results imply that unless $k = |T| - o(|T|)$, Subset k-Connectivity Augmentation admits the same ratios as the best known ones for the rooted version. This improves the ratios in [19, 14].

1 Introduction

In the Survivable Network problem we are given a graph $G = (V,E)$ with edge/node-costs and pairwise connectivity requirements $\{r(u,v) : u, v \in T \subseteq V\}$ on a set T of terminals. The goal is to find a minimum-cost subgraph of G that contains $r(u,v)$ internally-disjoint uv-paths for all $u, v \in T$. In Rooted Subset k-Connectivity problem there is $s \in T$ such that $r(s,t) = k$ for all $t \in T \setminus \{s\}$ and $r(u,v) = 0$ otherwise. In Subset k-Connectivity problem $r(u,v) = k$ for all $u, v \in T$ and $r(u,v) = 0$ otherwise. In the *augmentation versions*, G contains a subgraph J of cost zero with $r(u,v) - 1$ internally disjoint paths for all $u, v \in T$. A subset $T \subseteq V$ of terminals is k-connected to a root s in a directed/undirected graph J if J has k internally-disjoint vs-paths for every $v \in T$; T is k-connected in J if T is k-connected to every $s \in T$. Formally, the versions of Survivable Network we consider are as follows, where we revise our notation to $k \leftarrow k+1$.

Rooted Subset k-Connectivity Augmentation
Instance: A graph $G = (V,E)$ with edge/node-costs, a set $T \subseteq V$ of terminals, root $s \in T$, and a subgraph $J = (V, E_J)$ of G such that $T \setminus \{s\}$ is k-connected to s in J.

Objective: Find a minimum-cost augmenting edge-set $F \subseteq E \setminus E_J$ such that $T \setminus \{s\}$ is $(k+1)$-connected to s in $J \cup F$.

Subset k-Connectivity Augmentation

Instance: A graph $G = (V, E)$ with edge/node-costs, subset $T \subseteq V$, and a subgraph $J = (V, E_J)$ of G such that T is k-connected in J.

Objective: Find a minimum-cost augmenting edge-set $F \subseteq E \setminus E_J$ such that T is $(k+1)$-connected in $J \cup F$.

The Subset k-Connectivity Augmentation is Label-Cover hard to approximate [9]. It is known and easy to see that for both edge-costs and node-costs, if Subset k-Connectivity Augmentation admits approximation ratio $\rho(k)$ such that $\rho(k)$ is a monotone increasing function, then Subset k-Connectivity admits ratio $k \cdot \rho(k)$. Moreover, for edge costs, if in addition the approximation $\rho(k)$ is w.r.t. a standard setpair/biset LP-relaxation to the problem, then Subset k-Connectivity admits ratio $H(k) \cdot \rho(k)$, where $H(k)$ denotes the kth harmonic number. For edge-costs, a standard LP-relaxation for Survivable Network (due to Frank and Jordán [5]) is:

$$\min \left\{ \sum_{e \in E} c_e x_e : \sum_{e \in E(X,X^*)} x_e \geq r(X, X^*), X, X^* \subseteq V, X \cap X^* = \emptyset, 0 \leq x_e \leq 1 \right\}$$

where $r(X, X^*) = \max\{r(u,v) : u \in X, v \in X^*\}$ and $E(X, X^*)$ is the set of edges in E from X to X^*.

The Subset k-Connectivity problem admits trivial ratios $O(|T|^2)$ for both edge-costs and node-costs, by computing for every $u, v \in V$ an optimal edge-set of k internally-disjoint uv-paths (this is essentially a Min-Cost k-Flow problem, that can be solved in polynomial time), and taking the union of the computed edge-sets. We note that for metric edge-costs the problem admits an $O(1)$ ratio [2]. For $|T| \geq k+1$ the problem can also be decomposed into k instances of Rooted Subset k-Connectivity problems, c.f. [11] for the case $T = V$, where it is also shown that for $T = V$ the number of of Rooted Subset k-Connectivity Augmentation instances can be reduced to $O\left(\frac{|T|}{|T|-k} \log k\right)$, which is $O(\log k)$ unless $k = |T| - o(|T|)$.

Recently, Laekhanukit [14] made an important observation that the method of [11] can be extended for the case of arbitrary $T \subseteq V$. Specifically, he proved that if $|T| \geq 2k$, then $O(\log k)$ instances of Rooted Subset k-Connectivity Augmentation will suffice. Thus for $|T| \geq 2k$, the $O(k)$-approximation algorithm of [19] for Rooted Subset k-Connectivity Augmentation leads to the ratio $O(k \log k)$ for Rooted Subset k-Connectivity Augmentation. By cleverly exploiting an additional property of the algorithm of [19] (see [14, Lemma 14]), he reduced the ratio to $O(k)$ in the case $|T| \geq k^2$.

However, using a different approach, we will show that all this is not necessary, as for both directed and undirected graphs and edge-costs and node-costs, Subset k-Connectivity Augmentation can be reduced to solving *one* instance (or two instances, in the case of directed graphs) of Rooted Subset k-Connectivity Augmentation and $O\left(\frac{3|T|}{|T|-k}\right)^2 H\left(\frac{3|T|}{|T|-k}\right)$ instances of Min-Cost k-Flow problem.

This leads to a much simpler algorithm, improves the result of Laekhanukit [14] for $|T| < k^2$, and applies also for node-costs and directed graphs. In addition, we give a more natural and much simpler extension of the algorithm of [11] for $T = V$, that also enables the same bound $O\left(\frac{|T|}{|T|-k} \log k\right)$ as in [11] for arbitrary T with $|T| \geq k+1$, and in addition applies also for directed graphs, for node-costs, and for an arbitrary type of edge-costs, e.g., metric costs, or uniform costs, or $0,1$-costs. When we say "$0,1$-edge-costs" we mean that the input graph G is complete, and the goal is to add to the subgraph J of G formed by the zero-cost edges a minimum size edge-set F (any edge is allowed) such that $J \cup F$ satisfies the connectivity requirements. Formally, our result is the following.

Theorem 1. *For both directed and undirected graphs, and edge-costs and node-costs the following holds. If* Rooted Subset k-Connectivity Augmentation *admits approximation ratio $\rho = \rho(k, |T|)$, then for $|T| \geq k+1$* Subset k-Connectivity Augmentation *admits the following approximation ratios:*

(i) $b(\rho + k) + \left(\frac{|T|}{|T|-k}\right)^2 O\left(\log \frac{|T|}{|T|-k}\right)$, *where $b = 1$ for undirected graphs and $b = 2$ for directed graphs.*
(ii) $\rho \cdot O\left(\frac{|T|}{|T|-k} \log \min\{k, |T|-k\}\right)$, *and this is so also for $0,1$-edge-costs.*

Furthermore, if for edge-costs the approximation ratio ρ is w.r.t. a standard LP-relaxation for the problem, then so are the ratios in (i) *and* (ii).

For $|T| > k$, the best known values of ρ on undirected graphs are $O(k)$ for edge-costs and $\min\{O(k \log |T|), |T|\}$ for node-costs [19]; for directed graphs $\rho = |T|$ for both versions. For $0,1$-edge-costs $\rho = O(\log k)$ [20] for undirected graphs and $\rho = O(\log |T|)$ [18] for directed graphs. For edge-costs, these ratios are w.r.t. a standard LP-relaxation. Thus Theorem 1 implies the following.

Corollary 1. *For $|T| \geq k+1$,* Subset k-Connectivity Augmentation *admits the following approximation ratios.*

- *For undirected graphs, the ratios are $O(k) + \left(\frac{|T|}{|T|-k}\right)^2 O\left(\log \frac{|T|}{|T|-k}\right)$ for edge-costs, $O(k \log |T|) + \left(\frac{|T|}{|T|-k}\right)^2 O\left(\log \frac{|T|}{|T|-k}\right)$ for node-costs, and $\frac{|T|}{|T|-k} \cdot O\left(\log^2 k\right)$ for $0,1$-edge-costs.*
- *For directed graphs, the ratio is $2(|T|+k) + \left(\frac{|T|}{|T|-k}\right)^2 O\left(\log \frac{|T|}{|T|-k}\right)$ for both edge-costs and node-costs, and $\frac{|T|}{|T|-k} \cdot O(\log |T| \log k)$ for $0,1$ edge-costs.*

For Subset k-Connecivity, *the ratios are larger by a factor of $H(k)$ for edge-costs, and by a factor k for node-costs.*

Note that except the case of $0,1$-edge-costs, Corollary 1 is deduced from part (i) of Theorem 1. However, part (ii) of Theorem 1 might become relevant if Rooted Subset k-Connectivity Augmentation admits ratio better than $O(k)$. In addition, part (ii) applies for *any type* of edge-costs, e.g. metric or $0,1$-edge-costs.

We conclude this section by mentioning some additional related work. The case $T = V$ of Rooted Subset k-Connectivity problem is the k-Outconnected Subgraph problem; this problem admits a polynomial time algorithm for directed graphs [6], which implies ratio 2 for undirected graphs. For arbitrary T, the problem harder than Directed Steiner Tree [15]. The case $T = V$ of Subset k-Connectivity problem is the k-Connected Subgraph problem. This problem is NP-hard, and the best known ratio for it is $O\left(\log k \log \frac{n}{n-k}\right)$ for both directed and undirected graphs [17]; for the augmentation version of increasing the connectivity by one the ratio in [17] is $O\left(\log \frac{n}{n-k}\right)$. For metric costs the problem admits ratios $2 + \frac{k-1}{n}$ for undirected graphs and $2 + \frac{k}{n}$ for directed graphs [10]. For $0,1$-edge-costs the problem is solvable for directed graphs [5], which implies ratio 2 for undirected graphs. The Survivable Network problem is Label-Cover hard [9], and the currently best known non-trivial ratios for it on undirected graphs are: $O(k^3 \log |T|)$ for arbitrary edge-costs by Chuzhoy and Khanna [3], $O(\log k)$ for metric costs due to Cheriyan and Vetta [2], $O(k) \cdot \min\left\{\log^2 k, \log |T|\right\}$ for $0,1$-edge-costs [20, 13], and $O(k^4 \log^2 |T|)$ for node-costs [19].

2 Proof of Theorem 1

We start by proving the following essentially known statement.

Proposition 1. *Suppose that* Rooted Subset k-Connectivity Augmentation *admits an approximation ratio ρ. If for an instance of* Subset k-Connectivity Augmentation *we are given a set of q edges (when any edge is allowed) and p stars (directed to or from the root) on T whose addition to G makes T $(k+1)$-connected, then we can compute a $(\rho p+q)$-approximate solution F to this instance in polynomial time. Furthermore, for edge-costs, if the ρ-approximation is w.r.t. a standard LP-relaxation, then $c(F) \leq (\rho p + q)\tau^*$, where τ^* is an optimal standard LP-relaxation value for* Subset k-Connectivity Augmentation.

Proof. For every edge uv among the q edges compute a minimum-cost edge-set $F_{uv} \subseteq E \setminus E_J$ such that $J \cup F_{uv}$ contains k internally-disjoint uv-paths. This can be done in polynomial time for both edge and node costs, using a Min-Cost k-Flow algorithm. For edge-costs, it is known that $c(F_{uv}) \leq \tau^*$. Then replace uv by F_{uv}, and note that T remains k-connected. Similarly, for every star S with center s and leaf-set T', compute an α-approximate augmenting edge-set $F_S \subseteq E \setminus E_J$ such that $J \cup F_S$ contains k internally-disjoint sv-paths (or vs-paths, in the case of directed graphs and S being directed twords the root) for every $v \in T'$. Then replace S by F_S, and note that T remains k-connected. For edge-costs, it is known that if the ρ-approximation for the rooted version is w.r.t. a standard LP-relaxation, then $c(F_S) \leq (\alpha p + q)\tau^*$. The statement follows. □

Motivated by Proposition 1, we consider the following question:
Given a k-connected subset T in a graph J, how many edges and/or stars on T

one needs to add to J such that T will become $(k+1)$-connected?

We emphasize that we are interested in obtaining *absolute bounds* on the number of edges in the question, expressed in certain parameters of the graph; namely we consider the *extremal graph theory* question and not the *algorithmic problem*. Indeed, the algorithmic problem of adding the minimum number of edges on T such that T will become $(k+1)$-connected can be shown to admit a polynomial-time algorithm for directed graphs using the result of Frank and Jordán [5]; this also implies a 2-approximation algorithm for undirected graphs. However, in terms of the parameters $|T|, k$, the result in [5] implies only the trivial bound $O(|T|^2)$ on the the number of edges one needs to add to J such that T will become $(k+1)$-connected.

Our bounds will be derived in terms of the family of the "deficient" sets of the graph J. We need some definitions to state our results.

Definition 1. *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 *and X^+ is the* outer part *of $\hat{X}$, $\Gamma(\hat{X}) = X^+ \setminus X$ is the boundary of $\hat{X}$, and $X^* = V \setminus X^+$ is the* complementary set *of $\hat{X}$.*

Given an instance of Subset k-Connectivity Augmentation we may assume that T is an independent set in J. Otherwise, we obtain an equivalent instance by subdividing every edge $uv \in J$ with $u, v \in T$ by a new node.

Definition 2. *Given a k-connected independent set T in a graph $J = (V, E_J)$ let us say that a biset $\hat{X}$ on V is* (T,k)-tight *in J if $X \cap T, X^* \cap T \neq \emptyset$, X^+ is the union of X and the set of neighbors of X in J, and $|\Gamma(\hat{X})| = k$.*

An edge covers a biset $\hat{X}$ if it goes from X to X^*. By Menger's Theorem, F is a feasible solution to Subset k-Connectivity Augmentation if, and only if, F covers the biset-family $\mathcal{F}$ of tight bisets; see [12, 20]. Thus our question can be reformulated as follows:
Given a k-connected independent set T in a graph J, how many edges and/or stars on T are needed to cover the family $\mathcal{F}$ of (T,k)-tight bisets?

Definition 3. *The intersection and the union of two bisets $\hat{X}, \hat{Y}$ is defined by $\hat{X} \cap \hat{Y} = (X \cap Y, X^+ \cap Y^+)$ and $\hat{X} \cup \hat{Y} = (X \cup Y, X^+ \cup Y^+)$. Two bisets $\hat{X}, \hat{Y}$* intersect *if $X \cap Y \neq \emptyset$; if in addition $X^* \cap Y^* \neq \emptyset$ then $\hat{X}, \hat{Y}$* cross. *We say that a biset-family $\mathcal{F}$ is:*

- crossing *if $\hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y} \in \mathcal{F}$ for any $\hat{X}, \hat{Y} \in \mathcal{F}$ that cross.*
- k-regular *if $|\Gamma(\hat{X})| \leq k$ for every $\hat{X} \in \mathcal{F}$, and if $\hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y} \in \mathcal{F}$ for any intersecting $\hat{X}, \hat{Y} \in \mathcal{F}$ with $|X \cup Y| \leq |T| - k - 1$.*

The following statement is essentially known.

Lemma 1. *Let T be a k-connected independent set in a graph $J = (V, E_J)$, and let $\hat{X}, \hat{Y}$ be (T,k)-tight bisets. If $(X \cap T, X^+ \cap T), (Y \cap T, Y^+ \cap T)$ cross or if $|(X \cup Y) \cap T| \leq |T| - k - 1$ then $\hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y}$ are both (T,k)-tight.*

Proof. The case $(X \cap T, X^+ \cap T), (Y \cap T, Y^+ \cap T)$ was proved in [20] and [14]. The proof of the case $|(X \cup Y) \cap T| \le |T| - k - 1$ is identical to the proof of [7, Lemma 1.2] where the case $T = V$ is considered. □

Corollary 2. *The biset-family*

$$\mathcal{F} = \{(X \cap T, X^+ \cap T) : (X, X^+) \text{ is a } (T,k)\text{-tight biset in } J\}$$

is crossing and k-regular, and the reverse family $\bar{\mathcal{F}} = \{(T \setminus X^+, T \setminus X) : \hat{X} \in \mathcal{F}\}$ *of $\mathcal{F}$ is also crossing and k-regular. Furthermore, if J is undirected then $\mathcal{F}$ is symmetric, namely, $\mathcal{F} = \bar{\mathcal{F}}$.*

Given two bisets $\hat{X}, \hat{Y}$ we write $\hat{X} \subseteq \hat{Y}$ and say that $\hat{Y}$ contains $\hat{X}$ if $X \subseteq Y$ or if $X = Y$ and $X^+ \subseteq Y^+$; $\hat{X} \subset \hat{Y}$ and $\hat{Y}$ properly contains $\hat{X}$ if $X \subset Y$ or if $X = Y$ and $X^+ \subset Y^+$.

Definition 4. *A biset $\hat{C}$ is a* core *of a biset-family $\mathcal{F}$ if $\hat{C} \in \mathcal{F}$ and $\hat{C}$ contains no biset in $\mathcal{F} \setminus \{\hat{C}\}$; namely, a core is an inclusion-minimal biset in $\mathcal{F}$. Let $\mathcal{C}(\mathcal{F})$ be the family of cores of $\mathcal{F}$ and let $\nu(\mathcal{F}) = |\mathcal{C}(\mathcal{F})|$ denote their number.*

Given a biset-family $\mathcal{F}$ and an edge-set I on T, the residual biset-family $\mathcal{F}_I$ of $\mathcal{F}$ consists of the members of $\mathcal{F}$ uncovered by I. We will assume that for any I, the cores of $\mathcal{F}_I$ and of $\bar{\mathcal{F}}_I$ can be computed in polynomial time. For $\mathcal{F}$ being the family of (T,k)-tight bisets this can be implemented in polynomial time using the Ford-Fulkerson Max-Flow Min-Cut algorithm, c.f. [20]. It is known and easy to see that if $\mathcal{F}$ is crossing and/or k-regular, so is $\mathcal{F}_I$, for any edge-set I.

Definition 5. *For a biset-family $\mathcal{F}$ on T let $\nu(\mathcal{F})$ be the maximum number of bisets in $\mathcal{F}$ which inner parts are pairwise-disjoint. For an integer k let $\mathcal{F}^k = \{\hat{X} \in \mathcal{F} : |X| \le (|T| - k)/2\}$.*

Lemma 2. *Let $\mathcal{F}$ be a k-regular biset-family on T and let $\hat{X}, \hat{Y} \in \mathcal{F}^k$ intersect. Then $\hat{X} \cap \hat{Y} \in \mathcal{F}^k$ and $\hat{X} \cup \hat{Y} \in \mathcal{F}$.*

Proof. Since $|X|, |Y| \le \frac{|T|-k}{2}$, we have $|X \cup Y| = |X| + |Y| - |X \cap Y| \le |T| - k - 1$. Thus $\hat{X} \cap \hat{Y}, \hat{X} \cap \hat{Y} \in \mathcal{F}$, by the k-regularity of $\mathcal{F}$. Moreover, $\hat{X} \cap \hat{Y} \in \mathcal{F}^k$, since $|X \cap Y| \le |X| \le \frac{|T|-k}{2}$. □

We will prove the following two theorems that imply Theorem 1.

Theorem 2. *Let $\mathcal{F}$ be a biset-family on T such that both $\mathcal{F}, \bar{\mathcal{F}}$ are crossing and k-regular. Then there exists a polynomial-time algorithm that computes an edge-cover I of $\mathcal{F}$ of size $|I| = \nu\left(\mathcal{F}^k\right) + \nu\left(\bar{\mathcal{F}}^k\right) + \left(\frac{3|T|}{|T|-k}\right)^2 H\left(\frac{3|T|}{|T|-k}\right)$. Furthermore, if $\mathcal{F}$ is symmetric then $|I| = \nu\left(\mathcal{F}^k\right) + \left(\frac{3|T|}{|T|-k}\right)^2 H\left(\frac{3|T|}{|T|-k}\right)$.*

Theorem 3. *Let $\mathcal{F}$ be a biset-family on T such that both $\mathcal{F}$ and $\bar{\mathcal{F}}$ are k-regular. Then there exists a collection of $O\left(\frac{|T|}{|T|-k}\lg\min\{\nu,|T|-k\}\right)$ stars on T which union covers $\mathcal{F}$, and such a collection can be computed in polynomial time. Furthermore, the total number of edges in the stars is at most $\nu\left(\mathcal{F}^k\right)+\nu\left(\bar{\mathcal{F}}^k\right)+\left(\frac{|T|}{|T|-k}\right)^2\cdot O\left(\log\frac{|T|}{|T|-k}\right)$.*

Note that the second statement in Theorem 3 implies (up to constants) the bound in Theorem 2. However, the proof of Theorem 2 is much simpler than the proof of Theorem 3, and the proof of Theorem 2 is a part of the proof of the second statement in Theorem 3.

Let us show that Theorems 2 and 3 imply Theorem 1. For that, all we need is to show that by applying one time the α-approximation algorithm for the Rooted Subset k-Connectivity Augmentation, we obtain an instance with $\nu\left(\mathcal{F}^k\right),\nu\left(\bar{\mathcal{F}}^k\right)\le k+1$. This is achieved by the following procedure due to Khuller and Raghavachari [8] that originally considered the case $T=V$, see also [1, 4, 10]; the same procedure is also used by Laekhanukit in [14].

Choose an arbitrary subset $T'\subseteq T$ of $k+1$ nodes, add a new node s (the root) and all edges between s and T' of cost zero each, both to G and to J. Then, using the α-approximation algorithm for the Rooted Subset k-Connectivity Augmentation, compute an augmenting edge set F such that $J\cup F$ contains k internally disjoint vs-paths and sv-paths for every $v\in T'$. Now, add F to J and remove s from J. It is a routine to show that $c(F)\le b\mathsf{opt}$, and that for edge-costs $c(F)\le b\tau^*$. It is also known that if $\hat{X}$ is a tight biset of the obtained graph J, then $X\cap T',X^*\cap T'\neq\emptyset$, c.f. [1, 14]. Combined with Lemma 2 we obtain that $\nu\left(\mathcal{F}^k\right),\nu\left(\bar{\mathcal{F}}^k\right)\le|T'|\le k+1$ for the obtained instance, as claimed.

3 Proof of Theorem 2

Definition 6. *Given a biset-family $\mathcal{F}$ on T, let $\Delta(\mathcal{F})$ denote the maximum degree in the hypergraph $\mathcal{F}^{in}=\{X:\hat{X}\in\mathcal{F}\}$ of the inner parts of the bisets in $\mathcal{F}$. We say that $T'\subseteq T$ is a* transversal *of $\mathcal{F}$ if $T'\cap X\neq\emptyset$ for every $X\in\mathcal{F}^{in}$; a function $t:T\to[0,1]$ is a* fractional transversal *of $\mathcal{F}$ if $\sum_{v\in X}t(v)\ge 1$ for every $X\in\mathcal{F}^{in}$.*

Lemma 3. *Let $\mathcal{F}$ be a crossing biset-family. Then $\Delta(\mathcal{C}(\mathcal{F}))\le\nu\left(\bar{\mathcal{F}}\right)$.*

Proof. Since $\mathcal{F}$ is crossing, the members of $\mathcal{C}(\mathcal{F})$ are pairwise non-crossing. Thus if $\mathcal{H}$ is a subfamily of $\mathcal{C}(\mathcal{F})$ such that the intersection of the inner parts of the bisets in $\mathcal{H}$ is non-empty, then $\bar{\mathcal{H}}$ is a subfamily of $\bar{\mathcal{F}}$ such that the inner parts of the bisets in $\bar{\mathcal{H}}$ are pairwise disjoint, so $|\bar{\mathcal{H}}|\le\nu\left(\bar{\mathcal{F}}\right)$. The statement follows. □

Lemma 4. *Let T' be a transversal of a biset-family $\mathcal{F}'$ on T and let I' be an edge-set on T obtained by picking for every $s\in T'$ an edge from s to every inclusion member of the set-family $\{X^*:\hat{X}\in\mathcal{F}',s\in X\}$. Then I' covers $\mathcal{F}'$. Moreover, if $\mathcal{F}'$ is crossing then $|I'|\le|T'|\cdot\nu(\bar{\mathcal{F}}')$.*

Proof. The statement that I' covers $\mathcal{F}'$ is obvious. If $\mathcal{F}'$ is crossing, then for every $s \in T$ the inclusion-minimal members of $\{X^* : \hat{X} \in \mathcal{F}', s \in X\}$ are pairwise-disjoint, hence their number is at most $\nu(\bar{\mathcal{F}}')$. The statement follows. ☐

Lemma 5. *Let $\mathcal{F}$ be a k-regular biset-family on T. Then the following holds.*

(i) $\nu(\mathcal{F}) \leq \nu\left(\mathcal{F}^k\right) + \frac{2|T|}{|T|-k}$.

(ii) *If $\nu\left(\mathcal{F}^k_{\{e\}}\right) = \nu\left(\mathcal{F}^k\right)$ holds for every edge e on T then $\nu\left(\mathcal{F}^k\right) \leq \frac{|T|}{|T|-k}$.*

(iii) *There exists a polynomial time algorithm that finds a transversal T' of $\mathcal{C}(\mathcal{F})$ of size at most $|T'| \leq \left(\nu\left(\mathcal{F}^k\right) + \frac{2|T|}{|T|-k}\right) \cdot H(\Delta(\mathcal{C}(\mathcal{F})))$.*

Proof. Part (i) is immediate.

We prove (ii). Let $\hat{C} \in \mathcal{C}\left(\mathcal{F}^k\right)$ and let $\hat{U}_C$ be the union of the bisets in $\mathcal{F}^k$ that contain $\hat{C}$ and contain no other member of $\mathcal{C}\left(\mathcal{F}^k\right)$. If $|U_C| \leq |T| - k - 1$ then $\hat{U}_C \in \mathcal{F}$, by the k-regularity of $\mathcal{F}$. In this case $\nu\left(\mathcal{F}^k_{\{e\}}\right) \leq \nu\left(\mathcal{F}^k\right) - 1$ for any edge from C to U^*_C. Hence $|U_C| \geq |T| - k$ must hold for every $\hat{C} \in \mathcal{C}(\mathcal{F})$. By Lemma 2, the sets in the set family $\{U_C : \hat{C} \in \mathcal{C}(\mathcal{F})\}$ are pairwise disjoint. The statement follows.

We prove (iii). Let T^k be an inclusion-minimal transversal of $\mathcal{F}^k$. By Lemma 2, $\left|T^k\right| = \nu\left(\mathcal{F}^k\right)$. Setting $t(v) = 1$ if $v \in T^k$ and $t(v) = \frac{2}{|T|-k}$ otherwise, we obtain a fractional transversal of $\mathcal{C}(\mathcal{F})$ of value at most $\nu\left(\mathcal{F}^k\right) + \frac{2|T|}{|T|-k}$. Consequently, the greedy algorithm of Lovász [16] finds a transversal T' as claimed. ☐

The algorithm for computing I as in Theorem 2 starts with $I = \emptyset$ and then continues as follows.

Phase 1
While there exists an edge e on T such that $\nu\left(\mathcal{F}^k_{I\cup\{e\}}\right) \leq \nu\left(\mathcal{F}^k_I\right) - 1$, or such that $\nu\left(\bar{\mathcal{F}}^k_{I\cup\{e\}}\right) \leq \nu\left(\bar{\mathcal{F}}^k_I\right) - 1$, add e to I.

Phase 2
Find a transversal T' of $\mathcal{C}(\mathcal{F}')$ as in Lemma 5(iii), where $\mathcal{F}' = \mathcal{F}_I$. Then find an edge-cover I' of $\mathcal{F}'$ as in Lemma 4 and add I' to I.

The edge-set I computed covers $\mathcal{F}$ by Lemma 4. Clearly, the number of edges in I at the end of Phase 1 is at most $\nu\left(\mathcal{F}^k\right) + \nu\left(\bar{\mathcal{F}}^k\right)$, and is at most $\nu\left(\mathcal{F}^k\right)$ if $\mathcal{F}$ is symmetric. Now we bound the size of I'. Note that at the end of Phase 1 we have $\nu\left(\mathcal{F}^k_I\right), \nu\left(\bar{\mathcal{F}}^k_I\right) \leq \frac{|T|}{|T|-k}$ (by Lemma 5(ii)) and thus $\nu\left(\bar{\mathcal{F}}_I\right) \leq \frac{3|T|}{|T|-k}$ (by Lemma 5(i)) and $\Delta(\mathcal{C}(\mathcal{F}_I)) \leq \nu\left(\bar{\mathcal{F}}_I\right) \leq \nu\left(\bar{\mathcal{F}}^k_I\right) + \frac{2|T|}{|T|-k} \leq \frac{3|T|}{|T|-k}$ (by Lemma 3). Consequently, $|T'| \leq \left(\nu\left(\mathcal{F}^k_I\right) + \frac{2|T|}{|T|-k}\right) \cdot H(\Delta(\mathcal{C}(\mathcal{F}_I))) \leq \frac{3|T|}{|T|-k} \cdot H\left(\frac{3|T|}{|T|-k}\right)$. From this we get $|I'| \leq |T'| \cdot \nu\left(\bar{\mathcal{F}}_I\right) \leq \left(\frac{3|T|}{|T|-k}\right)^2 \cdot H\left(\frac{3|T|}{|T|-k}\right)$.

The proof of Theorem 2 is now complete.

4 Proof of Theorem 3

We start by analyzing the performance of a natural Greedy Algorithm for covering $\nu\left(\mathcal{F}^k\right)$, that starts with $I=\emptyset$ and while $\nu(\mathcal{F}^k_I) \geq 1$ adds to I a star S for which $\nu(\mathcal{F}^k_{I\cup S})$ is minimal. It is easy to see that the algorithm terminates since any star with center s in the inner part of some core of $\mathcal{F}^k_I$ and edge set $\{vs : v \in T \setminus \{s\}\}$ reduces the number of cores by one. The proof of the following statement is similar to the proof of the main result of [11].

Lemma 6. *Let $\mathcal{F}$ be a k-regular biset-family and let $\mathcal{S}$ be the collection of stars computed by the* Greedy Algorithm*. Then*

$$|\mathcal{S}| = O\left(\frac{|T|}{|T|-k}\ln\min\left\{\nu\left(\mathcal{F}^k\right), |T|-k\right\}\right) .$$

Recall that given $\hat{C} \in \mathcal{C}\left(\mathcal{F}^k_I\right)$ we denote by $\hat{U}_C$ the union of the bisets in $\mathcal{F}^k_I$ that contain $\hat{C}$ and contain no other member of $\mathcal{C}\left(\mathcal{F}^k_I\right)$, and that by Lemma 2, the sets in the set-family $\{U_C : \hat{C} \in \mathcal{C}(\mathcal{F})\}$ are pairwise disjoint.

Definition 7 ([11]). *Let us say that $s \in V$* out-covers *$\hat{C} \in \mathcal{C}\left(\mathcal{F}^k\right)$ if $s \in U^*_C$.*

Lemma 7. *Let $\mathcal{F}$ be k-regular biset-family and let $\nu = \nu\left(\mathcal{F}^k\right)$.*

(i) *There is $s \in T$ that out-covers at least $\nu\left(1-\frac{k}{|T|}\right)-1$ members of $\mathcal{C}\left(\mathcal{F}^k\right)$.*
(ii) *Let s out-cover the members of $\mathcal{C} \subseteq \mathcal{C}\left(\mathcal{F}^k\right)$ and let S be a star with one edge from s to the inner part of each member of $\mathcal{C}$. Then $\nu(\mathcal{F}^k) \leq \nu(\mathcal{F}^k_S) - |\mathcal{C}|/2$.*

Consequently, there exists a star S on T such that

$$\nu(\mathcal{F}^k_S) \leq \frac{1}{2}\left(1+\frac{k}{|T|}\right)\cdot\nu + \frac{1}{2} = \alpha\cdot\nu+\beta \ . \tag{1}$$

Proof. We prove (i). Consider the hypergraph $\mathcal{H} = \left\{T \setminus \Gamma\left(\hat{U}_C\right) : \hat{C} \in \mathcal{C}\left(\mathcal{F}^k\right)\right\}$. Note that the number of members of $\mathcal{C}\left(\mathcal{F}^k\right)$ out-covered by any $v \in T$ is at least the degree of s in $\mathcal{H}$ minus 1. Thus all we need to prove is that there is a node $s \in T$ whose degree in $\mathcal{H}$ is at least $\nu\left(1-\frac{k}{|T|}\right)$. For every $C \in \mathcal{C}(\mathcal{F})$ we have $\left|T \setminus \Gamma\left(\hat{U}_C\right)\right| \geq |T|-k$, by the k-regularity of $\mathcal{F}$. Hence the bipartite incidence graph of $\mathcal{H}$ has at least $\nu(|T|-k)$ edges, and thus has a node $s \in T$ of degree at least $\nu\left(1-\frac{k}{|T|}\right)$, which equals the degree of s in $\mathcal{H}$. Part (i) follows.

We prove (ii). It is sufficient to show that every $\hat{C} \in \mathcal{C}\left(\mathcal{F}^k_S\right)$ contains some $\hat{C}' \in \mathcal{C}\left(\mathcal{F}^k\right) \setminus \mathcal{C}$ or contains at least two members in $\mathcal{C}$. Clearly, $\hat{C}$ contains some $\hat{C}' \in \mathcal{C}\left(\mathcal{F}^k\right)$. We claim that if $\hat{C}' \in \mathcal{C}$ then $\hat{C}$ must contain some $\hat{C}'' \in \mathcal{C}\left(\mathcal{F}^k\right)$ distinct from $\hat{C}'$. Otherwise, $\hat{C} \in \mathcal{F}^k(C)$. But as S covers all members of $\mathcal{F}^k(C)$, $\hat{C} \notin \mathcal{F}^k_S$. This is a contradiction. □

Let us use parameters $\alpha, \beta, \gamma, \delta$ and j set to

$$\alpha = \frac{1}{2}\left(1 + \frac{k}{|T|}\right) \qquad \beta = \frac{1}{2} \qquad \gamma = 1 - \frac{k}{|T|} = 2(1-\alpha) \qquad \delta = 1,$$

and j is the minimum integer such that $\alpha^j \left(\nu - \frac{\beta}{1-\alpha}\right) \leq \frac{2}{1-\alpha}$ (note that $\alpha < 1$), namely,

$$j = \left\lfloor \frac{\ln \frac{1}{2}(\nu(1-\alpha) - \beta)}{\ln(1/\alpha)} \right\rfloor \leq \left\lfloor \frac{\ln \frac{1}{2}\nu(1-\alpha)}{\ln(1/\alpha)} \right\rfloor . \qquad (2)$$

We assume that $\nu \geq \frac{2+\beta}{1-\alpha}$ to have $j \geq 0$ (otherwise Lemma 6 follows). Note that $\frac{\beta}{1-\alpha} = \frac{|T|}{|T|-k}$.

Lemma 8. *Let $0 \leq \alpha < 1$, $\beta \geq 0$, $\nu_0 = \nu$, and for $i \geq 1$ let*

$$\nu_{i+1} \leq \alpha\nu_i + \beta \qquad s_i = \gamma\nu_{i-1} - \delta .$$

Then $\nu_i \leq \alpha^i \left(\nu - \frac{\beta}{1-\alpha}\right) + \frac{\beta}{1-\alpha}$ and $\sum_{i=1}^{j} s_i \leq \frac{1-\alpha^j}{1-\alpha} \cdot \gamma \left(\nu - \frac{\beta}{1-\alpha}\right) + j\left(\frac{\gamma\beta}{1-\alpha} - \delta\right)$. Moreover, if j is given by (2) then $\nu_j \leq \frac{2+\beta}{1-\alpha} = \frac{5|T|}{|T|-k}$ and $\sum_{i=1}^{j} s_i \leq 2\left(\nu - \frac{|T|}{|T|-k}\right)$.

Proof. Unraveling the recursive inequality $\nu_{i+1} \leq \alpha\nu_i + \beta$ in the lemma we get:

$$\nu_i \leq \alpha^i\nu + \beta\left(1 + \alpha + \cdots + \alpha^{i-1}\right) = \alpha^i\nu + \beta\frac{1-\alpha^i}{1-\alpha} = \alpha^i\left(\nu - \frac{\beta}{1-\alpha}\right) + \frac{\beta}{1-\alpha} .$$

This implies $s_i \leq \gamma\left(\nu - \frac{\beta}{1-\alpha}\right)\alpha^{i-1} + \frac{\gamma\beta}{1-\alpha} - \delta$, and thus

$$\begin{aligned}\sum_{i=1}^{j} s_i &\leq \gamma\left(\nu - \frac{\beta}{1-\alpha}\right)\sum_{i=1}^{j}\alpha^{i-1} + j\left(\frac{\gamma\beta}{1-\alpha} - \delta\right)\\ &= \gamma\left(\nu - \frac{\beta}{1-\alpha}\right)\cdot\frac{1-\alpha^j}{1-\alpha} + j\left(\frac{\gamma\beta}{1-\alpha} - \delta\right)\end{aligned}$$

If j is given by (2) then $\nu_j \leq \alpha^i\left(\nu - \frac{\beta}{1-\alpha}\right) + \frac{\beta}{1-\alpha} \leq \frac{2}{1-\alpha} + \frac{\beta}{1-\alpha} = \frac{2+\beta}{1-\alpha}$, and

$$\begin{aligned}\sum_{i=1}^{j} s_i &\leq \frac{1-\alpha^j}{1-\alpha}\cdot\gamma\left(\nu - \frac{\beta}{1-\alpha}\right) + j\left(\frac{\gamma\beta}{1-\alpha} - \delta\right)\\ &\leq 2\left(\nu - \frac{\beta}{1-\alpha}\right) = 2\left(\nu - \frac{|T|}{|T|-k}\right) .\end{aligned}$$

□

We now finish the proof of Lemma 6. At each one of the first j iterations we out-cover at least $\nu\left(\mathcal{F}_I^k\right)\left(1 - \frac{k}{|T|}\right) - 1$ members of $\mathcal{C}\left(\mathcal{F}_I^k\right)$, by Lemmas 7.

In each one of the consequent iterations, we can reduce $\nu\left(\mathcal{F}_I^k\right)$ by at least one, if we choose the center of the star in C for some $\hat{C} \in \mathcal{C}\left(\mathcal{F}_I^k\right)$. Thus using Lemma 8, performing the necessary computations, and substituting the values of the parameters, we obtain that the number of stars in $\mathcal{S}$ is bounded by

$$j + \nu_j \leq \left\lfloor \frac{\ln \frac{1}{2}\nu(1-\alpha)}{\ln(1/\alpha)} \right\rfloor + \frac{5|T|}{|T|-k} = O\left(\frac{|T|}{|T|-k} \ln \min\{\nu, |T|-k\}\right) .$$

Now we discuss a variation of this algorithm that produces $\mathcal{S}$ with a small number of leaves. Here at each one of the first j iterations we out-cover *exactly* $\nu\left(1 - \frac{k}{|T|}\right) - 1$ min-cores. For that, we need be able to compute the bisets $\hat{U}_C$, and such a procedure can be found in [14]. The number of edges in the stars at the end of this phase is at most $2\left(\nu - \frac{|T|}{|T|-k}\right)$ and $\nu_j \leq \frac{5|T|}{|T|-k}$. In the case of non-symmetric $\mathcal{F}$ and/or directed edges, we apply the same algorithm on $\bar{\mathcal{F}}^k$. At this point, we apply Phase 2 of the algorithm from the previous section. Since the number of cores of each one of $\mathcal{F}_I^k, \bar{\mathcal{F}}_I^k$ is now $O\left(\frac{|T|}{|T|-k}\right)$, the size of the transversal T' computed is bounded by $|T'| = O\left(\frac{|T|}{|T|-k} \cdot \log \frac{|T|}{|T|-k}\right)$. The number of stars is at most the size $|T'|$, while the number of edges in the stars is at most $|T'| \cdot \nu\left(\bar{\mathcal{F}}_I\right) = \left(\frac{|T|}{|T|-k}\right)^2 \cdot O\left(\log \frac{|T|}{|T|-k}\right)$.

This concludes the proof of Theorem 3.

References

1. V. Auletta, Y. Dinitz, Z. Nutov, and D. Parente. A 2-approximation algorithm for finding an optimum 3-vertex-connected spanning subgraph. *J. Algorithms*, 32(1):21–30, 1999.
2. J. Cheriyan and A. Vetta. Approximation algorithms for network design with metric costs. *SIAM J. Discrete Mathematics*, 21(3):612–636, 2007.
3. J. Chuzhoy and S. Khanna. An $O(k^3 \log n)$-approximation algorithm for vertex-connectivity survivable network design. In *FOCS*, pages 437–441, 2009.
4. Y. Dinitz and Z. Nutov. A 3-approximation algorithm for finding optimum 4,5-vertex-connected spanning subgraphs. *J. Algorithms*, 32(1):31–40, 1999.
5. A. Frank and T. Jordán. Minimal edge-coverings of pairs of sets. *J. Combinatorial Theory, Ser. B*, 65(1):73–110, 1995.
6. A. Frank and E. Tardos. An application of submodular flows. *Linear Algebra and its Applications*, 114/115:329–348, 1989.
7. T. Jordán. On the optimal vertex-connectivity augmentation. *J. Combinatorial Theory, Ser. B*, 63(1):8–20, 1995.
8. S. Khuller and B. Raghavachari. Improved approximation algorithms for uniform connectivity problems. *J. Algorithms*, 21(2):434–450, 1996.
9. G. Kortsarz, R. Krauthgamer, and J. Lee. Hardness of approximation for vertex-connectivity network design problems. *SIAM J. Computing*, 33(3):704–720, 2004.
10. G. Kortsarz and Z. Nutov. Approximating node-connectivity problems via set covers. *Algorithmica*, 37:75–92, 2003.

11. G. Kortsarz and Z. Nutov. Approximating k-node connected subgraphs via critical graphs. *SIAM J. on Computing*, 35(1):247–257, 2005.
12. 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.
13. G. Kortsarz and Z. Nutov. Tight approximation algorithm for connectivity augmentation problems. *J. Computer and System Sciences*, 74(5):662–670, 2008.
14. B. Laekhamukit. An improved approximation algorithm for minimum-cost subset k-connectivity. In *ICALP*, 2011. To appear.
15. Y. Lando and Z. Nutov. Inapproximability of survivable networks. *Theoretical Computer Science*, 410(21-23):2122–2125, 2009.
16. L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13:383–390, 1975.
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, SODA 2009, 912–921.
18. Z. Nutov. Approximating rooted connectivity augmentation problems. *Algorithmica*, 44:213–231, 2006.
19. Z. Nutov. Approximating minimum cost connectivity problems via uncrossable bifamilies and spider-cover decompositions. In *FOCS*, pages 417–426, 2009.
20. Z. Nutov. Approximating node-connectivity augmentation problems. In *APPROX*, pages 286–297, 2009.