

Approximating minimum-cost edge-covers of crossing biset-families*

Zeev Nutov

The Open University of Israel

nutov@openu.ac.il

Abstract

An ordered pair $\hat{S} = (S, S^+)$ of subsets of a groundset V is called a *biset* if $S \subseteq S^+$; $(V \setminus S^+, V \setminus S)$ is the *co-biset* of \hat{S} . Two bisets \hat{X}, \hat{Y} *intersect* if $X \cap Y \neq \emptyset$ and *cross* if both $X \cap Y \neq \emptyset$ and $X^+ \cup Y^+ \neq V$. The intersection and the union of two bisets \hat{X}, \hat{Y} are defined by $\hat{X} \cap \hat{Y} = (X \cap Y, X^+ \cap Y^+)$ and $\hat{X} \cup \hat{Y} = (X \cup Y, X^+ \cup Y^+)$. A biset-family \mathcal{F} is *crossing* (*intersecting*) 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 (intersect). A directed edge covers a biset \hat{S} if it goes from S to $V \setminus S^+$. We consider the problem of covering a crossing biset-family \mathcal{F} by a minimum-cost set of directed edges. While for intersecting \mathcal{F} , a standard primal-dual algorithm computes an optimal solution, the approximability of the case of crossing \mathcal{F} is not yet understood, as it includes several NP-hard problems, for which a poly-logarithmic approximation was discovered only recently or is not known. Let us say that a biset-family \mathcal{F} is *k-regular* if $\hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y} \in \mathcal{F}$ for any $\hat{X}, \hat{Y} \in \mathcal{F}$ with $|V \setminus (X \cup Y)| \geq k + 1$ that intersect. In this paper we obtain an $O(\log |V|)$ -approximation algorithm for arbitrary crossing \mathcal{F} ; if in addition both \mathcal{F} and the family of co-bisets of \mathcal{F} are *k-regular*, our ratios are: $O\left(\log \frac{|V|}{|V|-k}\right)$ if $|S^+ \setminus S| = k$ for all $\hat{S} \in \mathcal{F}$, and $O\left(\frac{|V|}{|V|-k} \log \frac{|V|}{|V|-k}\right)$ if $|S^+ \setminus S| \leq k$ for all $\hat{S} \in \mathcal{F}$. Using these generic algorithms, we derive for some network design problems the following approximation ratios: $O\left(\log k \cdot \log \frac{n}{n-k}\right)$ for *k-Connected Subgraph*, and $O(\log k) \cdot \min\left\{\frac{n}{n-k} \log \frac{n}{n-k}, \log k\right\}$ for *Subset k-Connected Subgraph* when all edges with positive cost have their endnodes in the subset.

1 Introduction

1.1 Problem definition and main results

Following [6], an ordered pair $\hat{S} = (S, S^+)$ of subsets of a groundset V is called a *biset* if $S \subseteq S^+$; S is the *inner part* and S^+ is the *outer part* of \hat{S} , and $\Gamma(\hat{S}) = S^+ \setminus S$ is the *boundary* of \hat{S} . The *co-biset* of \hat{S} is the biset $(V \setminus S^+, V \setminus S)$. Any set S can be considered as a biset $\hat{S} = (S, S)$ with $\Gamma(\hat{S}) = \emptyset$.

*Part of this paper appeared in the preliminary version [16].

Definition 1.1 Two bisets \hat{X}, \hat{Y} on V intersect if $X \cap Y \neq \emptyset$ and are disjoint otherwise; \hat{X}, \hat{Y} cross if $X \cap Y \neq \emptyset$ and $X^+ \cup Y^+ \neq V$. The intersection and the union of bisets \hat{X}, \hat{Y} are defined by $\hat{X} \cap \hat{Y} = (X \cap Y, X^+ \cap Y^+)$ and $\hat{X} \cup \hat{Y} = (X \cup Y, X^+ \cup Y^+)$. We say that a biset-family \mathcal{F} is:

- crossing (intersecting) 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 (intersect).
- k -regular if $\hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y} \in \mathcal{F}$ for any $\hat{X}, \hat{Y} \in \mathcal{F}$ with $|V \setminus (X \cup Y)| \geq k + 1$ that intersect.

A biset \hat{S} is *proper* if $S, V \setminus S^+$ are both nonempty. All biset-families in this paper are assumed to contain only proper bisets. A directed edge e leaves/covers a (proper) biset \hat{S} if it goes from S to $V \setminus S^+$. An edge-set/graph J is an *edge-cover* of \mathcal{F} if every biset in \mathcal{F} is covered by some edge in J . We consider the following generic algorithmic problem.

Biset-Family Edge-Cover

Instance: A directed graph $G = (V, E)$ with edge-costs $\{c_e : e \in E\}$ and a biset-family \mathcal{F} on V .

Objective: Find a minimum-cost edge-cover $J \subseteq E$ of \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 $X \subseteq Y$ and $X^+ \subseteq Y^+$; similarly, $\hat{X} \subset \hat{Y}$ and \hat{Y} properly contains \hat{X} if $X \subset Y$ and $X^+ \subseteq Y^+$, or if $X = Y$ and $X^+ \subset Y^+$.

Definition 1.2 A biset $\hat{C} \in \mathcal{F}$ is a *core* of a biset-family \mathcal{F} , or an \mathcal{F} -core for short, if \hat{C} contains no biset in $\mathcal{F} \setminus \{\hat{C}\}$, namely, if \hat{C} is an inclusion-minimal member of \mathcal{F} . Let $\mathcal{C}(\mathcal{F})$ denote the family of \mathcal{F} -cores and let $\nu(\mathcal{F}) = |\mathcal{C}(\mathcal{F})|$ denote the number of \mathcal{F} -cores.

In the Biset-Family Edge-Cover problem, \mathcal{F} may not be given explicitly, and a polynomial in $n = |V|$ implementation of our algorithms requires that certain queries related to \mathcal{F} can be answered in polynomial time. Given an edge set J on V , the *residual family* \mathcal{F}^J of \mathcal{F} consists of all members of \mathcal{F} that are uncovered by the edges of J . It is known that if \mathcal{F} is crossing or intersecting, so is \mathcal{F}^J , for any J ; this follows from the fact that if an edge e covers $\hat{X} \cap \hat{Y}$ or $\hat{X} \cup \hat{Y}$ then e covers \hat{X} or \hat{Y} (see Fact 3.3(i) in Section 3). The *co-family* of \mathcal{F} is the biset-family $\{(V \setminus S^+, V \setminus S) : (S, S^+) \in \mathcal{F}\}$ of co-bisets of the bisets in \mathcal{F} . It is easy to see that \mathcal{F} is crossing if, and only if, its co-family is crossing, and that J covers \mathcal{F} if, and only if, the reverse edge-set of J covers the co-family of \mathcal{F} . We assume that for any edge set J on V and any $u, v \in V$ we are able to compute in polynomial time the cores of the biset-family $\mathcal{F}(u, v) = \{\hat{S} \in \mathcal{F}^J : u \in S, v \in V \setminus S^+\}$ and also the cores of its co-family, or to determine that $\mathcal{F}(u, v)$ is empty. In specific graph problems we consider, this can be implemented in polynomial time using the Ford-Fulkerson Max-Flow Min-Cut Algorithm; we omit the somewhat standard implementation details.

For intersecting \mathcal{F} , Biset-Family Edge-Cover can be solved in polynomial time by a standard primal-dual algorithm; in fact, even a more general problem of covering an intersecting supermodular biset-function by a digraph can also be solved in polynomial time [6]. However, the case of crossing \mathcal{F} includes the min-cost k -Connectivity Augmentation problem which is NP-hard, and its

approximability is not yet understood. Given a biset-family \mathcal{F} let $\gamma(\mathcal{F}) = \max_{\hat{S} \in \mathcal{F}} |\Gamma(\hat{S})|$. It is known that any crossing set-family \mathcal{F} (namely, a crossing biset-family \mathcal{F} with $\gamma(\mathcal{F}) = 0$) is decomposable into two intersecting families as follows. Let $s \in V$ be arbitrary, and let $\mathcal{F}^{out} = \{\hat{S} \in \mathcal{F} : s \in V \setminus S^+\}$ and $\mathcal{F}^{in} = \{\hat{S} \in \mathcal{F} : s \in S\}$. The families \mathcal{F}^{out} and the co-family of \mathcal{F}^{in} are both intersecting, and an edge-set J covers \mathcal{F} if, and only if, J covers \mathcal{F}^{out} and the reverse edge-set of J covers the co-family of \mathcal{F}^{in} . This implies ratio 2 for **Biset-Family Edge-Cover** for crossing \mathcal{F} with $\gamma(\mathcal{F}) = 0$. In a similar way we can decompose any crossing \mathcal{F} into $2(\gamma(\mathcal{F}) + 1)$ intersecting biset-families. This implies ratio $2(\gamma(\mathcal{F}) + 1)$ for **Biset-Family Edge-Cover** with crossing \mathcal{F} . Using ideas from [18, 11, 5, 16], we give approximation algorithms with logarithmic or constant ratios.

For an edge-set or a graph J and a biset \hat{S} on V let $\delta_J(\hat{S})$ denote the set of edges in J covering \hat{S} . Let $\tau(\mathcal{S})$ denote the optimal value of an LP-relaxation for covering a biset-family \mathcal{S} , namely,

$$\begin{aligned}
 \text{(P)} \quad \tau(\mathcal{S}) = \quad & \min \quad \sum_{e \in E} c_e x_e \\
 \text{s.t.} \quad & \sum_{e \in \delta_E(\hat{S})} x_e \geq 1 \quad \forall \hat{S} \in \mathcal{S} \\
 & x_e \geq 0 \quad \forall e \in E
 \end{aligned}$$

Let $H(n) = \sum_{i=1}^n (1/i)$ denote the n th harmonic number. Our main result is the following.

Theorem 1.1 *Biset-Family Edge-Cover with crossing \mathcal{F} admits a polynomial time algorithm that computes a solution of cost at most $\rho \cdot \tau(\mathcal{F})$ where:*

- (i) $\rho = O(\log \nu(\mathcal{F})) = O(\log n)$ for arbitrary crossing \mathcal{F} .
- (ii) $\rho = O\left(\log \min\left\{\frac{n}{n-k}, n-k\right\}\right)$ if both \mathcal{F} and the co-family of \mathcal{F} are k -regular and if $|\Gamma(\hat{S})| = k$ for all $\hat{S} \in \mathcal{F}$.
- (iii) $\rho = O\left(\frac{n}{n-k} \log \frac{n}{n-k}\right)$ if both \mathcal{F} and the co-family of \mathcal{F} are k -regular and if $\gamma(\mathcal{F}) \leq k$.

We note that Theorem 1.1 easily extends to the undirected case, with a loss of a factor of 2 in the approximation ratio.

1.2 Related work and applications

A directed/undirected graph is k -connected if there are k internally-disjoint paths from any of its nodes to any other node. A fundamental problem in network design is the following:

k -Connected Subgraph

Instance: A graph $G = (V, E)$ with edge-costs $\{c_e : e \in E\}$ and an integer k .

Objective: Find a minimum cost k -connected spanning subgraph of G .

See [18, 4, 11, 5, 16, 12] for a history of the problem. Let ℓ -Connectivity Augmentation be the restriction of k -Connected Subgraph to instances in which G contains an ℓ -connected spanning subgraph G_0 of cost zero, and we seek to increase the connectivity of G_0 from $\ell = k - 1$ to $\ell + 1 = k$. Parts of the following statement were implicitly proved in [7] and [18], see also [4, 11, 12].

Proposition 1.2 *ℓ -Connectivity Augmentation is a particular case of Biset-Family Edge-Cover with crossing \mathcal{F} , such that $|\Gamma(\hat{S})| = \ell$ for all $\hat{S} \in \mathcal{F}$, and such that both \mathcal{F} and the co-family of \mathcal{F} are ℓ -regular. Furthermore, if the latter problem admits a polynomial time algorithm that computes a solution of cost $\leq \alpha(n, \ell) \cdot \tau(\mathcal{F})$, then k -Connected Subgraph admits a polynomial time algorithm that computes a solution of cost $\leq \text{opt}_k \cdot \sum_{\ell=1}^k \frac{\alpha(n, \ell)}{k-\ell+1}$, where $\text{opt}_k = \min\{\sum_{e \in E} c_e x_e : x(\delta_E(\hat{S})) \geq k - |\Gamma(\hat{S})| \forall \text{ proper biset } \hat{S} \text{ on } V\}$ is the optimal value of a natural LP-relaxation for the problem. In particular, if $\alpha(n, \ell)$ is increasing in ℓ then the cost of the solution computed $\leq \text{opt}_k \cdot \alpha(n, k) \cdot H(k)$.*

Fackaroenphol and Laekhamukit [5] gave an $O(\log n)$ -approximation algorithm for the ℓ -Connectivity Augmentation problem. Part (i) of Theorem 1.1 extends this to an arbitrary crossing biset-family \mathcal{F} , by a similar proof. Part (ii) improves the result of [5] using the k -regularity property.

Corollary 1.3 ([16]) *k -Connectivity Augmentation admits a polynomial time algorithm that computes a solution of cost $O\left(\log \frac{n}{n-k}\right) \cdot \text{opt}_k$; the approximation ratio is $O(1)$, unless $k = n - o(n)$. The problem also admits a polynomial time algorithm that computes a solution of cost $O(\log(n-k)) \cdot \text{opt}_k$. k -Connected Subgraph admits an $O\left(\log k \cdot \log \frac{n}{n-k}\right)$ -approximation algorithm; the ratio is $O(\log k)$, unless $k = n - o(n)$. The problem also admits a polynomial time algorithm that computes a solution of cost $O\left(\sum_{\ell=1}^k \frac{\log(n-\ell)}{k-\ell+1}\right) \cdot \text{opt}_k$.*

Now let us consider the following known generalization of the k -Connected Subgraph problem. Let us say that a subset T of nodes of a directed/undirected graph is k -connected if in the graph there are k internally-disjoint paths from every node in T to any other node in T .

Subset k -Connected Subgraph

Instance: A graph $G = (V, E)$ with edge-costs $\{c_e : e \in E\}$, $T \subseteq V$, and an integer k .

Objective: Find a minimum cost subgraph of G in which T is k -connected.

The k -Connected Subgraph problem is a particular case of Subset k -Connected Subgraph when $T = V$. Let Subset ℓ -Connectivity Augmentation be the restriction of Subset k -Connected Subgraph to instances in which G contains a subgraph G_0 of cost zero such that T is ℓ -connected in G_0 , and we seek to increase the connectivity of T from $\ell = k - 1$ to $\ell + 1 = k$. When the costs are arbitrary, Subset k -Connected Subgraph is unlikely to admit a polylogarithmic approximation [10] (see also [14] for a simpler proof). The currently best known ratio for this problem for $|T| > k$ is $b(k + \rho) + \frac{|T|^2}{(|T|-k)^2} O\left(\log \frac{|T|}{|T|-k}\right)$, where $b = 1$ for undirected graphs and $b = 2$ for directed graphs, and ρ is the ratio for the rooted version of the problem [13, 17]; currently, $\rho = \min\{\tilde{O}(k), |T|\}$ for undirected graphs [15], and $\rho = |T|$ for directed graphs. For $|T| \leq k$ the best ratio is $\frac{b}{2}|T|(|T| - 1)$. We consider the version when every edge with positive cost has both of its endnodes in T . The

relevant biset family \mathcal{F} on T is the projection on T of the family of “tight sets”, namely, $\mathcal{F} = \{(X \cap T, X^+ \cap T) : \hat{X} \text{ is tight}\}$, where a biset \hat{X} on V is tight if $X \cap T$ and $T \setminus X^+$ are both non-empty and $|\delta_{G_0}(\hat{X})| + |\Gamma(\hat{X})| = \ell$ (c.f. [12, 13, 17]). Then a similar statement to the one in Proposition 1.2 applies, except that \mathcal{F} is a biset-family on T and $|\Gamma(\hat{S})| \leq \ell$ for all $\hat{S} \in \mathcal{F}$. Furthermore, when $|T| > k$, then by applying b times an approximation algorithm for the rooted version of the problem, we can reduce the number of cores to $O(k^2)$; such a procedure is described in [13, 17]. The rooted version when every edge has its tail in T admits a polynomial time algorithm [6]. Thus parts (i) and (iii) of Theorem 1.1 imply the following.

Corollary 1.4 *In the case when every edge with positive cost has both of its endnodes in T , Subset k -Connectivity Augmentation admits a polynomial time algorithm that computes a solution of cost $\rho \cdot \text{opt}_k$, where $\rho = O\left(\min\left\{\frac{|T|}{\max\{|T|-k, 1\}} \log \frac{|T|}{\max\{|T|-k, 1\}}, \log \min\{k, |T|\}\right\}\right)$ and $\text{opt}_k = \min\{x(E) : x(\delta_E(S)) \geq k - |\Gamma(\hat{S})| \forall \text{ biset } \hat{S} \text{ on } V \text{ with } S \cap T, T \setminus S^+ \neq \emptyset\}$. Subset k -Connected Subgraph admits an $O(\rho \log k)$ -approximation algorithm.*

Cheriyān and Laekhanukit [2] considered the following directed edge-connectivity problem, that generalizes the Subset k -Connected Subgraph problem. Given two disjoint subsets S, T in a graph G , we say that G is k -edge-outconnected from S to T , or that G is k - (S, T) -edge-connected, if G has k edge-disjoint st -paths for every $(s, t) \in S \times T$.

k - (S, T) -Edge-Connected Subgraph

Instance: A directed graph $G = (V, E)$ with edge-costs $\{c_e : e \in E\}$, disjoint subsets $S, T \subseteq V$, and an integer k .

Objective: Find a minimum cost k -edge-outconnected from S to T subgraph of G .

In the so called “standard version” of the problem we have $E \subseteq S \times T$. Let the ℓ - (S, T) -Edge-Connectivity Augmentation problem be the restriction of k - (S, T) -Edge-Connected Subgraph to instances in which G contains a subgraph G_0 of cost zero such that G_0 is ℓ - (S, T) -connected, and we seek to increase at minimum cost the (S, T) -connectivity from $\ell = k - 1$ to $\ell + 1 = k$.

Let us say that two sets X, Y (S, T) -cross if $X \cap Y \cap S, T \setminus (X \cup Y) \neq \emptyset$. A set-family \mathcal{F} is (S, T) -crossing if $X \cap Y, X \cup Y \in \mathcal{F}$ for any $X, Y \in \mathcal{F}$ that (S, T) -cross. Generalizing the algorithm of Fackaroenphol and Laekhanukit [5] for the k -Connected Subgraph problem, Cheriyān and Laekhanukit [2] gave an approximation algorithm with ratio $O(\log(|S||T|))$ for the standard version of the k - (S, T) -Edge-Connectivity Augmentation problem. They also implicitly proved that it is a particular case of the Set-Family Edge-Cover problem with $V = S \cup T$, $E \subseteq S \times T$, and (S, T) -crossing \mathcal{F} . Our algorithm in Theorem 1.1(i) easily extends to the problem of covering an (S, T) -crossing family by a minimum-cost edge-set. Here we preferred the biset-family setting for simplicity of exposition, and since the concept of k -regularity is not a natural one for (S, T) -crossing families. Furthermore, the case of an (S, T) -crossing set family \mathcal{F} is reduced to the case of a crossing biset-family \mathcal{F}' , where for every set $X \in \mathcal{F}$ there is a biset $\hat{X}' = (X \cap S, S \cup (X \cap T))$ in \mathcal{F}' ; it is not hard to verify that if \mathcal{F} is an (S, T) -crossing family then \mathcal{F}' is a crossing biset-family, and that

an edge from S to T covers a set X if, and only if, it covers \hat{X}' . Thus from Theorem 1.1(i) we have the following generalization of the result of [2].

Corollary 1.5 *Set-Family Edge-Cover with (S, T) -crossing \mathcal{F} and $E \subseteq S \times T$ admits a polynomial time algorithm that computes a solution of cost $O(\log |S \cup T|) \cdot \tau(\mathcal{F})$.*

2 Proof of Theorem 1.1

2.1 Proof of part (i)

Recall that we assume that for any edge set J on V and any $u, v \in V$ we are able to compute in polynomial time the cores of the biset-family $\mathcal{F}(u, v) = \{\hat{S} \in \mathcal{F}^J : u \in S, v \in V \setminus S^+\}$ and also the cores of its co-family, or to determine that $\mathcal{F}(u, v)$ is empty. Note that if \mathcal{F} is crossing, then $\mathcal{F}(u, v)$ has a unique core, and the co-family of $\mathcal{F}(u, v)$ also has a unique core.

Lemma 2.1 *A crossing biset-family \mathcal{F} has at most $n(n - 1)$ cores and they can be computed in polynomial time.*

Proof: For every ordered pair of nodes $(u, v) \in V \times V$ we compute the core \hat{C}_{uv} of the biset-family $\mathcal{F}(u, v) = \{\hat{S} \in \mathcal{F}^J : u \in S, v \in V \setminus S^+\}$, if $\mathcal{F}(u, v)$ is non-empty. Then $\mathcal{C}(\mathcal{F})$ consists of the inclusion-minimal members (cores) of the biset-family $\{\hat{C}_{uv} : u, v \in V\}$. \square

Definition 2.1 *Given a core $\hat{C} \in \mathcal{C}(\mathcal{F})$ of \mathcal{F} let $\mathcal{F}(\hat{C})$ denote the family of the bisets in \mathcal{F} that contain \hat{C} and contain no other core of \mathcal{F} .*

The following statement is easily verified; it is implicitly proved in [11, 2].

Claim 2.2 *Let \mathcal{F} be a biset-family and J a set of directed edges on V . If for some $\hat{C} \in \mathcal{C}(\mathcal{F})$, J covers $\mathcal{F}(\hat{C})$ and covers no core distinct from \hat{C} , then $\mathcal{C}(\mathcal{F}^J) = \mathcal{C}(\mathcal{F}) \setminus \{\hat{C}\}$ and $\nu(\mathcal{F}^J) = \nu(\mathcal{F}) - 1$.*

Lemma 2.3 *Let $\hat{C}_1, \hat{C}_2 \in \mathcal{C}(\mathcal{F})$ be distinct cores of a crossing biset-family \mathcal{F} and let $\hat{S}_1 \in \mathcal{F}(\hat{C}_1)$ and $\hat{S}_2 \in \mathcal{F}(\hat{C}_2)$. Then \hat{S}_1, \hat{S}_2 do not cross. Consequently, no edge covers both \hat{S}_1, \hat{S}_2 .*

Proof: Suppose to the contrary that \hat{S}_1 and \hat{S}_2 cross. Then $\hat{S}_1 \cap \hat{S}_2 \in \mathcal{F}$. Thus $\hat{S}_1 \cap \hat{S}_2$ contains some \mathcal{F} -core \hat{C} . We cannot have $\hat{C} \neq \hat{C}_1$ as $\hat{C} \subseteq \hat{C}_1$ and \hat{C}_1 is a core, and we cannot have $\hat{C} = \hat{C}_1$ as $\hat{C} \subseteq \hat{C}_2$, $\hat{C}_1 \neq \hat{C}_2$, and \hat{C}_2 is a core. This gives a contradiction. The second statement follows from the observation that an edge covers two bisets \hat{X}, \hat{Y} simultaneously if, and only if, it goes from $X \cap Y$ to $V \setminus (X^+ \cup Y^+)$, and hence \hat{X}, \hat{Y} cross (see Fact 3.3(iii) in Section 3). \square

Lemma 2.4 *Let $\hat{C} \in \mathcal{C}(\mathcal{F})$ be a core of a crossing biset-family \mathcal{F} , and let $\hat{X}, \hat{Y} \in \mathcal{F}(\hat{C})$. If \hat{X}, \hat{Y} cross then $\hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y} \in \mathcal{F}(\hat{C})$.*

Proof: Since \mathcal{F} is crossing, $\hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y} \in \mathcal{F}$. Since $\hat{X} \cap \hat{Y} \subseteq \hat{X} \subseteq \hat{X} \cup \hat{Y}$ and since $\hat{X} \in \mathcal{F}(\hat{C})$, it follows that $\hat{X} \cap \hat{Y} \in \mathcal{F}(\hat{C})$ and that $\hat{C} \subseteq \hat{X} \cup \hat{Y}$. It remains to prove that $\hat{X} \cup \hat{Y}$ contains no

core distinct from \hat{C} . Suppose to the contrary that $\hat{X} \cup \hat{Y}$ contains a core \hat{S} distinct from \hat{C} . Since $\hat{X}, \hat{Y} \in \mathcal{F}(\hat{C})$, none of \hat{X}, \hat{Y} contains \hat{S} . This implies that \hat{S}, \hat{X} cross or \hat{S}, \hat{Y} cross, so $\hat{S} \cap \hat{X} \in \mathcal{F}$ or $\hat{S} \cap \hat{Y} \in \mathcal{F}$. This contradicts that \hat{S} is a core. \square

Lemma 2.5 *Let \mathcal{F} be a crossing biset-family on V and let $\hat{C} \in \mathcal{C}(\mathcal{F})$. Then the co-family $\mathcal{R}(\hat{C}) = \{(V \setminus S^+, V \setminus S) : \hat{S} \in \mathcal{F}(\hat{C})\}$ of $\mathcal{F}(\hat{C})$ is intersecting, and its cores can be found in polynomial time.*

Proof: Let $\hat{X}_0, \hat{Y}_0 \in \mathcal{R}(\hat{C})$ be the co-bisets of $\hat{X}, \hat{Y} \in \mathcal{F}(\hat{C})$, respectively. Suppose that \hat{X}_0, \hat{Y}_0 intersect. Then \hat{X}, \hat{Y} cross, hence $\hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y} \in \mathcal{F}(\hat{C})$, by Lemma 2.4. The co-bisets of $\hat{X} \cap \hat{Y}$ and $\hat{X} \cup \hat{Y}$ are $\hat{X}_0 \cup \hat{Y}_0$ and $\hat{X}_0 \cap \hat{Y}_0$, hence $\hat{X}_0 \cup \hat{Y}_0, \hat{X}_0 \cap \hat{Y}_0 \in \mathcal{R}(\hat{C})$. This implies that $\mathcal{R}(\hat{C})$ is an intersecting biset family. Now we show how to find the cores of $\mathcal{R}(\hat{C})$ in polynomial time. For an \mathcal{F} -core $\hat{S} \neq \hat{C}$ let $K_{\hat{S}} = \{uv : u \in S, v \in V \setminus S^+\}$ be the set of all edges from S to $V \setminus S^+$. Let $K = \bigcup_{\hat{S} \in \mathcal{C}(\mathcal{F}) \setminus \{\hat{C}\}} K_{\hat{S}}$. We claim that $\mathcal{F}^K = \mathcal{F}(\hat{C})$. To see this, note that for any $\hat{S} \in \mathcal{C}(\mathcal{F}) \setminus \{\hat{C}\}$: (i) $K_{\hat{S}}$ covers all bisets in \mathcal{F} that contain \hat{S} ; (ii) $K_{\hat{S}}$ does not cover any biset in $\mathcal{F}(\hat{C})$, by Lemma 2.3. Now choose $u \in C$, and for every $v \in V$ compute the core \hat{C}_v of the co-family of $\mathcal{F}^K(u, v)$. The $\mathcal{R}(\hat{C})$ -cores are the inclusion-minimal members of the family $\{\hat{C}_v : v \in V\}$. \square

Corollary 2.6 *Biset-family Edge-Cover with crossing \mathcal{F} admits a polynomial time algorithm that given a core $\hat{C} \in \mathcal{C}(\mathcal{F})$ computes an $\mathcal{F}(\hat{C})$ -cover $J_{\hat{C}} \subseteq E$ of cost $c(J_{\hat{C}}) = \tau(\mathcal{F}(\hat{C}))$. Moreover, $\sum_{\hat{C} \in \mathcal{C}} \tau(\mathcal{F}(\hat{C})) \leq \tau(\mathcal{F})$, and thus there exists $\hat{C} \in \mathcal{C}(\mathcal{F})$ such that $c(J_{\hat{C}}) \leq \tau(\mathcal{F})/\nu(\mathcal{F})$.*

Proof: By Lemma 2.5, the co-family $\mathcal{R}(\hat{C})$ of $\mathcal{F}(\hat{C})$ is intersecting. Thus, after reversing the edges in E , we can apply a standard primal-dual algorithm to compute an edge-cover of $\mathcal{R}(\hat{C})$ of cost $\tau(\mathcal{R}(\hat{C})) = \tau(\mathcal{F}(\hat{C}))$; $J_{\hat{C}}$ is the reverse edge set of this cover. This primal-dual algorithm can be implemented in polynomial time if the cores of $\mathcal{R}(\hat{C})$ can be found in polynomial time, which is possible by Lemma 2.5. The second statement of the lemma follows from Lemma 2.3. \square

Now we finish the proof of part (i) of Theorem 1.1. The algorithm starts with $J = \emptyset$. At iteration i , it finds $\hat{C}_i \in \mathcal{C}(\mathcal{F}^J)$ and $J_i \subseteq E \setminus J$ with $c(J_i) \leq \tau(\mathcal{F}^J)/\nu(\mathcal{F}^J)$ and $\nu(\mathcal{F}^{J \cup J_i}) = \nu(\mathcal{F}^J) - 1$, and adds J_i to J ; such J_i exists and can be found in polynomial time by Lemma 2.1, Claim 2.2, Lemma 2.3, and Corollary 2.6. At each iteration $\nu(\mathcal{F}^J)$ decreases by 1, by Claim 2.2. Thus at the end of iteration i we have $\nu(\mathcal{F}^J) = \nu(\mathcal{F}) - i$. Consequently, $c(J_i) \leq \tau(\mathcal{F}^J)/\nu(\mathcal{F}^J) \leq \tau(\mathcal{F})/(\nu(\mathcal{F}) - i)$. Thus at the end of the algorithm, $c(J) \leq \sum_i c(J_i) \leq \tau(\mathcal{F}) \sum_i 1/(\nu(\mathcal{F}) - i) = \tau(\mathcal{F}) \cdot H(\nu(\mathcal{F}))$.

2.2 Proof of part (ii)

The following concepts play a central role in the proof of part (ii) of Theorem 1.1.

Definition 2.2 *A biset-family \mathcal{F} is intersection-closed if $\hat{X} \cap \hat{Y} \in \mathcal{F}$ for any $\hat{X}, \hat{Y} \in \mathcal{F}$ that intersect. An intersection-closed biset-family \mathcal{S} is q -semi-intersecting if $|S| \leq q$ for all $\hat{S} \in \mathcal{S}$ and if $\hat{X} \cup \hat{Y} \in \mathcal{S}$ for any intersecting $\hat{X}, \hat{Y} \in \mathcal{S}$ with $|X \cup Y| \leq q$.*

Note that if \mathcal{F} is intersection closed, so is \mathcal{F}^J , for any edge set J ; this follows from the fact that if an edge e covers $\hat{X} \cap \hat{Y}$ then e covers \hat{X} or \hat{Y} (see Fact 3.3(i) in Section 3).

The following statement is straightforward.

Lemma 2.7 *Let \mathcal{F} be an intersection-closed biset-family. If $\hat{C} \in \mathcal{C}(\mathcal{F})$ and $\hat{S} \in \mathcal{F}$ intersect, then $\hat{C} \subseteq \hat{S}$. Thus the \mathcal{F} -cores are pairwise disjoint.*

The following statement is implicitly proved in [11].

Claim 2.8 *Let \mathcal{F} be an intersection-closed biset-family, let $\mathcal{C} \subseteq \mathcal{C}(\mathcal{F})$, and let J a set of directed edges on V that covers $\mathcal{F}(\hat{C})$ for every $\hat{C} \in \mathcal{C}$. Then every \mathcal{F}^J -core that contains some $\hat{C} \in \mathcal{C}$ contains an \mathcal{F} -core distinct from \hat{C} ; thus $\nu(\mathcal{F}^J) \leq \nu(\mathcal{F}) - |\mathcal{C}|/2$.*

The following is implicitly proved in several papers, c.f. [9, 18].

Fact 2.9 *If a biset-family \mathcal{F} is k -regular, then the subfamily $\mathcal{S} = \{\hat{S} \in \mathcal{F} : |\mathcal{S}| \leq q\}$ of \mathcal{F} is q -semi-intersecting (and in particular, is intersection closed) for any $q \leq (n - k)/2$.*

Let $q = \lfloor (n - k)/2 \rfloor$ and let $\mathcal{S} = \{\hat{S} \in \mathcal{F} : |\mathcal{S}| \leq q\}$. In Lemmas 2.10 and 2.12 to follow we show that if $\gamma(\mathcal{F}) \leq k$, then there exists a polynomial time algorithm that computes an \mathcal{S} -cover J of cost $c(J) \leq \tau(\mathcal{F}) \cdot \min\{1 + H(\lfloor 2n/(n - k + 2) \rfloor), \lfloor \log_2 \lfloor (n - k + 2)/2 \rfloor \rfloor\}$.

Lemma 2.10 *Biset-family Edge-Cover with crossing k -regular \mathcal{F} with $\gamma(\mathcal{F}) \leq k$ admits a polynomial time algorithm that returns an \mathcal{S} -cover J of costs $c(J) \leq \tau(\mathcal{F}) \cdot \lfloor \log_2 \lfloor (n - k + 2)/2 \rfloor \rfloor$.*

Proof: The algorithm that computes an \mathcal{S} -cover of cost at most $\tau(\mathcal{F}) \cdot \lfloor \log_2 \lfloor (n - k + 2)/2 \rfloor \rfloor$ is as follows. Start with $J = \emptyset$, and iteratively, until $\nu(\mathcal{S}^J) = 0$, find and add to J a cover $\bigcup_{\hat{C} \in \mathcal{C}(\mathcal{F})} J_{\hat{C}}$ of cost $\leq \tau(\mathcal{F}^J)$ of all families $\mathcal{F}(\hat{C})$ of \mathcal{S}^J -cores, as in Corollary 2.6. By Lemma 2.7 and Claim 2.2, after step i we have $2^i \leq |\mathcal{C}| \leq \frac{n-k}{2}$ for every \mathcal{S}^J -core \hat{C} . Hence the number of iterations is at most $\lfloor \log_2 \lfloor (n - k + 2)/2 \rfloor \rfloor$. As at every iteration we add to J an edge set of cost $\leq \tau(\mathcal{F})$, the total cost of the \mathcal{S} -cover computed is at most $\tau(\mathcal{F}) \cdot \lfloor \log_2 \lfloor (n - k + 2)/2 \rfloor \rfloor$. \square

The algorithm that computes an \mathcal{S} -cover of cost at most $\tau(\mathcal{F}) (1 + H(\lfloor 2n/(n - k + 2) \rfloor))$ is based on the following theorem (to be proved in Section 3), which is our main technical result, and which we believe is of independent interest.

Theorem 2.11 *Biset-Family Edge-Cover with q -semi-intersecting biset-family \mathcal{S} admits a polynomial time algorithm that computes an edge-set $J \subseteq E$ such that $\nu(\mathcal{S}^J) \leq \lfloor n/(q + 1) \rfloor$ and $c(J) \leq \tau(\mathcal{S})$.*

Note that for $q = n$ any intersecting biset-family is q -semi-intersecting, hence q -semi-intersecting biset-families generalize intersecting biset-families, and then the algorithm in Theorem 2.11 computes an optimal \mathcal{S} -cover of cost $\tau(\mathcal{S})$.

Lemma 2.12 *Biset-family Edge-Cover with crossing k -regular \mathcal{F} with $\gamma(\mathcal{F}) \leq k$ admits a polynomial time algorithm that returns an \mathcal{S} -cover cover J of costs $c(J) \leq \tau(\mathcal{F}) \cdot (1 + H(\lfloor 2n/(n-k+2) \rfloor))$.*

Proof: The proof combines Theorem 2.11 with a modification of the algorithm from part (i). Compute an edge set J as in Theorem 2.11 with $q = \lfloor (n-k)/2 \rfloor$ and $\mathcal{S} = \{\hat{S} \in \mathcal{F} : |S| \leq q\}$, so $\nu(\mathcal{S}^J) \leq \lfloor \frac{n}{q+1} \rfloor = \lfloor \frac{n}{\lfloor (n-k)/2 \rfloor + 1} \rfloor \leq \lfloor \frac{2n}{n-k+2} \rfloor$. Then apply the algorithm from part (i) on \mathcal{F}^J using the cores from \mathcal{S}^J only: iteratively, until $\nu(\mathcal{S}^J) = 0$, at iteration i compute for each $\hat{C} \in \mathcal{C}(\mathcal{S}^J)$ an $\mathcal{F}(\hat{C})$ -cover $J_{\hat{C}} \subseteq E$ as in Corollary 2.6, and add to J the cheapest set J_i among the edge sets $\{J_{\hat{C}} : \hat{C} \in \mathcal{C}(\mathcal{S}^J)\}$. By Claim 2.8 $\nu(\mathcal{S}^{J \cup J_i}) = \nu(\mathcal{S}^J) - 1$, thus at iteration i we have

$$c(J_i) \leq \frac{\tau(\mathcal{F}^J)}{\nu(\mathcal{S}^J)} \leq \frac{\tau(\mathcal{F}^J)}{\lfloor 2n/(n-k+2) \rfloor - i}.$$

Consequently, the cost of the \mathcal{S} -cover J computed is bounded by

$$c(J) \leq \tau(\mathcal{S}) + H(\lfloor 2n/(n-k+2) \rfloor) \cdot \tau(\mathcal{F}) \leq \tau(\mathcal{F}) \cdot (1 + H(\lfloor 2n/(n-k+2) \rfloor)).$$

□

Corollary 2.13 *Biset-family Edge-Cover with crossing k -regular \mathcal{F} with $\gamma(\mathcal{F}) \leq k$ admits a polynomial time algorithm that returns an edge set J of cost $c(J) = O\left(\log \min\left\{\frac{n}{n-k}, n-k\right\}\right) \cdot \tau(\mathcal{F})$ with the following property: the size of the inner part of every biset in \mathcal{F}^J or in the co-family of \mathcal{F}^J is larger than $q = \frac{n-k}{2}$.*

Proof: Apply each of the algorithms from Lemmas 2.10 and 2.12 twice: once on \mathcal{F}, E and once on the "reversed" instance with the biset-family being the co-family $\{(V - S^+, V - S) : (S, S^+) \in \mathcal{F}\}$ of \mathcal{F} and with the reverse edge-set $\{uv : vu \in E\}$ of E ; after a solution J' to the reversed instance is computed, we return the reversed edge-set of J' . The union J of the two computed edge sets has cost $c(J) = \tau(\mathcal{F}) \cdot O(\log(n-k))$ if we use the algorithm from Lemma 2.10, and $c(J) = \tau(\mathcal{F}) \cdot O\left(\log \frac{n}{n-k}\right)$ if we use the algorithm from Lemma 2.12. □

Let J be as in Corollary 2.13. In the case $|\Gamma(\hat{S})| = k$ for all $\hat{S} \in \mathcal{F}$, we must have $\mathcal{F}^J = \emptyset$, hence J covers \mathcal{F} and has cost as stated in part (ii) of Theorem 1.1.

2.3 Proof of part (iii)

Corollary 2.13 implies that to prove part (iii) of Theorem 1.1 it is sufficient to prove the following.

Lemma 2.14 *Biset-Family Edge-Cover with crossing \mathcal{F} admits a polynomial time algorithm that computes an edge-cover of \mathcal{F} of cost at most $\tau(\mathcal{F}) \cdot (n/q) \cdot H(\lfloor n/q \rfloor)$, provided that $|S| \geq q$ holds for every biset \hat{S} that belongs to \mathcal{F} or to the co-family of \mathcal{F} .*

In the rest of this section we prove Lemma 2.14. The key observation is the following.

Lemma 2.15 *Let \mathcal{F} be a crossing biset-family on V with $|S| \geq q$ for every biset \hat{S} in \mathcal{F} or in the co-family of \mathcal{F} . Then there exists $T \subseteq V$ of size $|T| \leq (n/q) \cdot H(\lfloor n/q \rfloor)$ such that $T \cap S \neq \emptyset$ for every $\hat{S} \in \mathcal{F}$.*

Proof: Consider the hypergraph $\mathcal{H} = \{C : \hat{C} \in \mathcal{C}(\mathcal{F})\}$ of the inner parts of the cores of \mathcal{F} . We combine the following two observations.

- (i) The function t on V defined by $t(v) = \frac{1}{q}$ for all $v \in V$ is a fractional hitting-set of \mathcal{H} (namely $\sum_{v \in S} t(v) \geq 1$ for all $C \in \mathcal{H}$) of value $\sum_{v \in V} t(v) = n/q$.
- (ii) The maximum degree in the hypergraph \mathcal{H} is at most $\lfloor n/q \rfloor$.

Given observations (i) and (ii), the greedy algorithm computes a subset $T \subseteq V$ as stated. Observation (i) follows from the assumption that $|S| \geq q$ for all $\hat{S} \in \mathcal{F}$. We prove (ii). Since \mathcal{F} is crossing, the members of $\mathcal{C}(\mathcal{F})$ are pairwise non-crossing. Thus if $\mathcal{C} \subseteq \mathcal{C}(\mathcal{F})$ is a set of cores which inner parts contain the same element $v \in V$, then the sets in $\{V \setminus C^+ : \hat{C} \in \mathcal{C}\}$ are pairwise disjoint. As each of these sets is an inner part of a biset in the co-family of \mathcal{F} , the number of such sets is at most $\lfloor n/q \rfloor$. Observation (ii) follows. \square

Lemma 2.14 easily follows from Lemma 2.15. Note that if \mathcal{F} is crossing, then for every $s \in V$ the co-family of the biset-family $\mathcal{F}_s = \{\hat{S} \in \mathcal{F} : v \in S\}$ is intersecting. Thus given an instance of Biset-Family Edge-Cover and $s \in V$, we can compute in polynomial time an edge-cover J_s of \mathcal{F}_s of cost $c(J_s) \leq \tau(\mathcal{F}_s) \leq \tau(\mathcal{F})$. Now let T be as in Lemma 2.15. For every $s \in T$ we compute an edge-cover J_s of \mathcal{F}_s as above, and return $J = \cup_{s \in T} J_s$. This concludes the proof of Lemma 2.14 and thus also the proof of part (iii) of Theorem 1.1 is now complete.

3 Proof of Theorem 2.11

Definition 3.1 *For biset-families \mathcal{S} and \mathcal{U} on V let*

$$\mathcal{S}[\mathcal{U}] = \{\hat{S} \in \mathcal{S} : \hat{S} \subseteq \hat{U} \text{ for some } \hat{U} \in \mathcal{U}\} .$$

A biset-family \mathcal{S} is weakly-intersecting if $\mathcal{S}[\{\hat{U}\}]$ is an intersecting biset-family for every $\hat{U} \in \mathcal{S}$.

Clearly, any q -semi-intersecting biset-family is weakly-intersecting. Note that if $\mathcal{U} \subseteq \mathcal{S}$ and if the members of \mathcal{U} are pairwise disjoint, then $\mathcal{S}[\mathcal{U}]$ is an intersecting biset-family, if \mathcal{S} is weakly-intersecting. We will prove the following refinement of Theorem 2.11.

Lemma 3.1 *Biset-Family Edge-Cover with weakly-intersecting biset-family \mathcal{S} admits a polynomial time algorithm that computes a sub-family $\mathcal{U} \subseteq \mathcal{S}$ of pairwise disjoint bisets and an edge-set $J \subseteq E$ such that the following two properties hold.*

Property 1: For any \mathcal{S}^J -core \hat{C} the following holds:

- (i) If $\hat{U} \in \mathcal{U}$ and \hat{C} intersect then \hat{U} intersects no \mathcal{S}^J -core distinct from \hat{C} .
- (ii) The union \hat{B}_C of \hat{C} and the bisets in \mathcal{U} intersecting with \hat{C} is not in \mathcal{S} .

Property 2: J is an optimal cover of $\mathcal{S}[\mathcal{U}]$, thus $c(J) = \tau(\mathcal{S}[\mathcal{U}]) \leq \tau(\mathcal{S})$.

To see that Lemma 3.1 implies Theorem 2.11, it is sufficient to show that if \mathcal{S} is q -semi-intersecting, then Property 1 implies $\nu(\mathcal{S}^J) \leq \lfloor n/(q+1) \rfloor$. Let B_C be the inner part of \hat{B}_C . Note that the sets B_C are pairwise disjoint; this is since the members of \mathcal{U} are pairwise disjoint, the \mathcal{S}^J -cores are pairwise disjoint, and since by Property 1(i) every $\hat{U} \in \mathcal{U}$ intersects at most one \mathcal{S}^J -core. Now observe that if \mathcal{S} is q -semi-intersecting then $|B_C| \geq q+1$, since $\hat{B}_C \notin \mathcal{S}$, by Property 1(ii). Consequently, $\nu(\mathcal{S}^J) \leq \lfloor n/(q+1) \rfloor$, and Lemma 3.1 implies Theorem 2.11.

In the rest of this section we prove Lemma 3.1. The algorithm is a variation of a standard primal-dual algorithm for covering an *intersecting* biset-family, and it only needs that the \mathcal{S}^J -cores can be computed in polynomial time. The dual LP of the LP-relaxation **(P)** from the Introduction is:

$$\begin{aligned}
 \text{(D)} \quad & \max \quad \sum_{\hat{S} \in \mathcal{S}} y_{\hat{S}} \\
 & \text{s.t.} \quad \sum_{\hat{S} \in \mathcal{S}: e \in \delta(\hat{S})} y_{\hat{S}} \leq c_e \quad \forall e \in E \\
 & \quad \quad y_{\hat{S}} \geq 0 \quad \quad \quad \forall \hat{S} \in \mathcal{S}
 \end{aligned}$$

Given a partial solution y to **(D)**, an edge $e \in E$ is *tight* if the inequality in **(D)** that corresponds to e holds with equality. The algorithm produces an edge set $J \subseteq E$, a sub-family $\mathcal{U} \subseteq \mathcal{S}$ of \mathcal{S} , and a dual solution y to **(D)**, such that the following holds: J covers $\mathcal{S}[\mathcal{U}]$, y is a feasible solution to **(D)**, and (the characteristic vector of) J and y satisfy the complementary slackness conditions:

Primal Complementary Slackness Conditions: $e \in J \implies e$ is tight;

Dual Complementary Slackness Conditions: $y_{\hat{S}} > 0 \implies |\delta_J(\hat{S})| = 1$.

Phase 1 starts with $J = \emptyset$ and $\mathcal{U} = \emptyset$ and applies a sequence of iterations. At each iteration we choose some \mathcal{S}^J -core \hat{C} and do the following:

1. Add \hat{C} to \mathcal{U} and exclude from \mathcal{U} the bisets contained in \hat{C} .
2. Raise (possibly by zero) the dual variable corresponding to \hat{C} , until some edge $e \in E \setminus J$ covering \hat{C} becomes tight, and add e to J .

Phase 1 terminates when $\nu(\mathcal{S}^J) = 0$, namely, when J covers \mathcal{S} .

Phase 2 applies on J “reverse delete” like the family $\mathcal{S}[\mathcal{U}]$ is the one we want to cover, which means the following. Let $J = \{e_1, \dots, e_j\}$, where e_{i+1} was added after e_i . For $i = j$ downto 1, we

delete e_i from J if $J - \{e_i\}$ still covers the family $\mathcal{S}[\mathcal{U}]$. This can be implemented in polynomial time as follows. When an edge $e \in J$ is checked for deletion, $\mathcal{S}[\mathcal{U}]$ is covered by $J \setminus \{e\}$ if, and only if, no $\hat{U} \in \mathcal{U}$ contains an $\mathcal{S}^{J \setminus \{e\}}$ -core. At the end of the algorithm, J is output.

Summarizing, the algorithm is a variation of a standard primal-dual algorithm for intersecting biset-families, with the following changes.

1. Unlike a standard primal-dual algorithm in which *all* the dual variables corresponding to cores are raised uniformly, we raise the dual variable of only *one* core.
2. The algorithm maintains a biset-family $\mathcal{U} \subseteq \mathcal{S}$. In each iteration, we add to \mathcal{U} the corresponding tight \mathcal{F}^J -core \hat{C} , and exclude from \mathcal{U} all the bisets contained in \hat{C} , if any.
3. While at Phase 1 the algorithm intends to cover the entire biset-family \mathcal{S} , Phase 2 (reverse-delete) is applied like the family $\mathcal{S}[\mathcal{U}]$ is the one that we want to cover.

Using Lemma 2.1, it is easy to see that the algorithm can be implemented in polynomial time, under the oracles assumed. Using Lemma 2.7 it is also easy to see the following.

Claim 3.2 *During the algorithm the following holds: the members of \mathcal{U} are pairwise disjoint, every edge in J has its tail in the inner part of some $\hat{U} \in \mathcal{U}$, and $|\delta_J(\hat{U})| = 1$ for every $\hat{U} \in \mathcal{U}$.*

Retrospectively, it turns out that our algorithm coincides with some run of an almost standard primal-dual algorithm that intends to cover the intersecting biset-family $\mathcal{S}[\mathcal{U}]$. The difference is in item 1 above, as we raise the dual variable of only one core; it is also possible to raise all the dual variables corresponding to cores, but this makes the analysis more complicated. We will show that this modified algorithm still computes an optimal solution. This will ensure Property 2 in Lemma 3.1; we give a formal proof of Property 2 after proving that Property 1 holds for J, \mathcal{U} at the end of the algorithm.

The following statement is easily verified, see Figure 1.

Fact 3.3 *Let \hat{X}, \hat{Y} be bisets and let e be an edge.*

- (i) *If e covers $\hat{X} \cap \hat{Y}$ or $\hat{X} \cup \hat{Y}$ then e covers \hat{X} or \hat{Y} .*
- (ii) *If e covers $\hat{X} \cup \hat{Y}$ and has tail in X then e covers \hat{X} .*
- (iii) *e covers both $\hat{X} \cap \hat{Y}$ and $\hat{X} \cup \hat{Y}$ if and only if e covers both \hat{X} and \hat{Y} ; furthermore, \hat{X}, \hat{Y} cross in this case.*

Let F be the set of edges stored in J at the end of Phase 1. Then F covers \mathcal{S} . Note that every edge in F has its tail in the inner part of some $\hat{U} \in \mathcal{U}$, and that $|\delta_F(\hat{U})| = 1$ for every $\hat{U} \in \mathcal{U}$. Thus the following statement implies that Property 1 holds for J after Phase 2, even if \mathcal{S} is only intersection-closed.

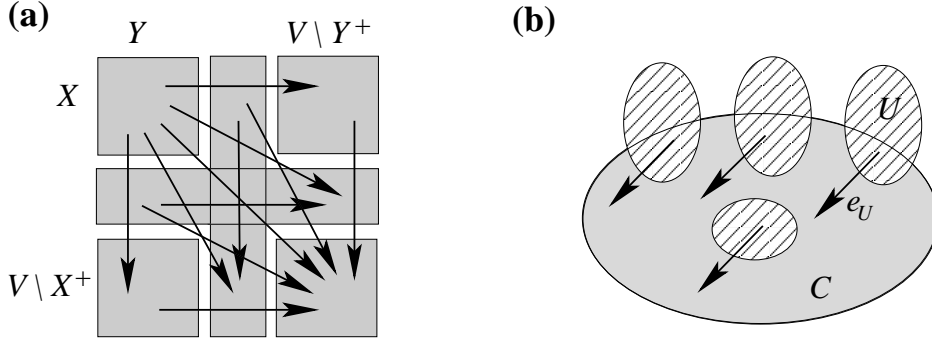


Figure 1: (a) All types of edges that can cover one of $\hat{X}, \hat{Y}, \hat{X} \cap \hat{Y}, \hat{X} \cup \hat{Y}$. (b) An \mathcal{S}^J -core C and the members of \mathcal{U} intersecting C .

Lemma 3.4 *Let F be an edge-cover of an intersection-closed biset-family \mathcal{S} and let $\mathcal{U} \subseteq \mathcal{S}$ be a sub-family of \mathcal{S} of pairwise disjoint bisets such that every edge in F has its tail in the inner part of some $\hat{U} \in \mathcal{U}$ and such that $\delta_F(\hat{U}) = \{e_U\}$ for every $\hat{U} \in \mathcal{U}$. Let $J \subseteq F$ be a cover of $\mathcal{S}[\mathcal{U}]$ and let \hat{C} be an \mathcal{S}^J -core. Then the following holds (see Figure 1(b)):*

- (i) *If $\hat{U} \in \mathcal{U}$ and \hat{C} intersect then $\delta_J(\hat{C} \cap \hat{U}) = \{e_U\}$; thus \hat{U} intersects no \mathcal{S}^J -core distinct from \hat{C} .*
- (ii) *$\delta_F(\hat{B}_C) = \emptyset$, where \hat{B}_C is the union of \hat{C} and the bisets in \mathcal{U} intersecting with \hat{C} ; thus $\hat{B}_C \notin \mathcal{S}$.*

Proof: We prove (i). Let $\hat{U} \in \mathcal{U}$ and \hat{C} intersect. Note that $\hat{C} \cap \hat{U} \in \mathcal{S}[\mathcal{U}]$, since \mathcal{S} is intersection-closed. Hence $\delta_J(\hat{C} \cap \hat{U}) \neq \emptyset$. Let $e \in \delta_J(\hat{C} \cap \hat{U})$. Then e covers \hat{C} or \hat{U} , by Fact 3.3(i). But e does not cover \hat{C} , since $e \in J$, and since J does not cover \hat{C} . Hence e covers \hat{U} . Consequently, $e = e_U$ for any $e \in \delta_J(\hat{C} \cap \hat{U})$, hence $\delta_J(\hat{C} \cap \hat{U}) = \{e_U\}$. This implies that the tail of e_U is in C , and it cannot be in the inner part of any other \mathcal{S}^J -core, since the \mathcal{S}^J -cores are pairwise disjoint. Thus U intersects no \mathcal{S}^J -core distinct from C .

We prove (ii). Suppose to the contrary that there is $e \in \delta_F(\hat{B}_C)$. Let $\hat{U} \in \mathcal{U}$ be the biset whose inner part contains the tail of e . Note that $U \cap C \neq \emptyset$. Let $\hat{X} = \hat{C} \cup \hat{U}$ and let \hat{Y} be the union of \hat{C} and the bisets in $\mathcal{U} \setminus \{\hat{U}\}$ that intersect \hat{C} . Note that $\hat{B}_C = \hat{X} \cup \hat{Y}$. Hence the tail of e lies in X and e covers $\hat{X} \cup \hat{Y}$. Thus e covers $\hat{C} \cup \hat{U}$, by Fact 3.3(ii). Applying the same argument on the bisets \hat{U}, \hat{Y} we obtain that e covers \hat{U} . Hence $e = e_U$, as e_U is the only edge in F that covers \hat{U} . By (i), e_U covers $\hat{C} \cap \hat{U}$. Consequently, e_U covers both $\hat{C} \cup \hat{U}$ and $\hat{C} \cap \hat{U}$, and hence e_U covers both \hat{C} and \hat{U} , by Fact 3.3(iii). However $e_U \in J$, contradicting that $\hat{C} \in \mathcal{S}^J$. \square

For Property 2 it is sufficient to prove the following.

Lemma 3.5 *If \mathcal{S} is weakly-intersecting, then at the end of the algorithm the following holds: J covers the biset-family $\mathcal{S}[\mathcal{U}]$, y is a feasible solution to (D) for $\mathcal{S}[\mathcal{U}]$, and J, y satisfy the complementary slackness conditions; hence J and y are optimal solutions.*

Proof: It is clear that J covers $\mathcal{S}[\mathcal{U}]$, and that during the algorithm y remains a feasible solution to (D). Since only tight edges enter J , after Phase 1 the *Primal Complementary Slackness Conditions* hold for J . So, the only part that requires proof is that after Phase 2, the *Dual Complementary Slackness Conditions* hold for J and y .

Claim: Consider an arbitrary $\hat{S} \in \mathcal{S}$ with $y_{\hat{S}} > 0$ and an edge $e \in \delta_J(\hat{S})$. Then there exists $\hat{W} \in \mathcal{S}[\mathcal{U}]$ such that $\delta_J(\hat{W}) = \{e\}$ and $\hat{S} \subseteq \hat{W} \subseteq \hat{U}$ for some $\hat{U} \in \mathcal{U}$.

Proof: Such \hat{W} can be chosen as any member of $\mathcal{S}[\mathcal{U}]$ which becomes uncovered if we delete (instead of keeping) e at the reverse delete step when e was considered for deletion; note that the algorithm decided to keep e , hence such \hat{W} exists. Moreover, since the edges were deleted in the reverse order, $\hat{W} \in \mathcal{S}[\mathcal{U}]^{J \setminus \{e\}}$. Obviously, $\delta_J(\hat{W}) = \{e\}$ and \hat{S} and \hat{W} intersect. Finally, to see that $\hat{S} \subseteq \hat{W}$ note that: (i) at any iteration before e was added, \hat{W} was uncovered; (ii) since $y_{\hat{S}} > 0$, there was an iteration before e was added at which \hat{S} was an \mathcal{S}^J -core. Hence $\hat{S} \subseteq \hat{W}$, by Lemma 2.7. \square

Now assume to the contrary that there is $\hat{S} \in \mathcal{S}[\mathcal{U}]$ with $y_{\hat{S}} > 0$ such that there are $e', e'' \in \delta_J(\hat{S})$, $e' \neq e''$. Let $\hat{U} \in \mathcal{U}$ such that $\hat{S} \subseteq \hat{U}$. Let \hat{W}', \hat{W}'' be bisets for e', e'' as in the Claim above. so $\delta_J(\hat{W}') = \{e'\}$, $\delta_J(\hat{W}'') = \{e''\}$ and $\hat{S} \subseteq \hat{W}' \cap \hat{W}'' \subseteq \hat{U}$. In particular, \hat{W}', \hat{W}'' intersect and $\hat{W}', \hat{W}'' \subseteq \hat{U}$. Thus $\hat{W}' \cup \hat{W}'' \in \mathcal{S}[\mathcal{U}]$, since $\hat{U} \in \mathcal{S}$ and since \mathcal{S} is weakly intersecting. Consequently, there is an edge $e \in \delta_J(\hat{W}' \cup \hat{W}'')$. This implies that $e \in \delta_J(\hat{W}')$ or $e \in \delta_J(\hat{W}'')$, by Fact 3.3(i). Consequently, $e = e'$ or $e = e''$. Since the tail of each one of e', e'' is in $S \subseteq W' \cap W''$, so is the tail of e . The head of e is in $V \setminus (W'^+ \cup W''^+)$, since e covers $\hat{W}' \cup \hat{W}''$. We conclude that $e \in \delta_J(\hat{W}') \cap \delta_J(\hat{W}'')$. This is a contradiction since $\delta_J(\hat{W}') = \{e'\}$, $\delta_J(\hat{W}'') = \{e''\}$, and $e' \neq e''$. \square

The proof of Lemma 3.1, and thus also of Theorem 2.11 is complete.

4 Concluding remarks and open problems

The main open question is whether the k -Connectivity Augmentation problem admits a constant ratio approximation algorithm also for large values of k , say $k = n - \sqrt{n}$. In this context we mention four papers. In [14] it is shown that for values of k close to n , the approximability of the k -Connectivity Augmentation problem is the same for directed and undirected graphs, up to a factor of 2. Therefore, one should not expect to obtain a constant ratio for undirected graphs only. On the positive side, Frank and Jordan [7] showed that for directed graphs, k -Connected Subgraph can be solved in polynomial time when the input graph is complete and the costs are in $\{0, 1\}$. For arbitrary costs however, there are two negative results. In [19] Ravi and Williamson gave an example showing that the approximation ratio of a standard primal-dual algorithm that intends to edge-cover the biset-family \mathcal{S} as in Fact 2.9 has approximation ratio $\Omega(k)$. This does not exclude that some other variation of the primal-dual algorithm, that relies on concepts from [7] has a constant ratio for crossing biset-families. However recently, Aazami, Cheriyan, and Laekhanukit

[1] showed that the standard iterative rounding method that is based on a standard LP-relaxation for k -Connectivity Augmentation (thus intends to edge-cover a crossing k -regular biset-family) has approximation ratio $\Omega(\sqrt{k})$.

Another open question is whether for the k -Connected Subgraph problem, the factor $O(\log k)$ invoked by considering the augmentation version can be avoided. In particular, can we achieve ratio $O\left(\log \frac{n}{n-k}\right)$? In this context we mention a recent paper of Cheriyan and Véggh [3]. They show that in the case of *undirected graphs* with $n > k^3(k-1) + k$ the problem can be decomposed into 2 problems of covering an intersecting supermodular biset function and an additional problem of covering a skew-supermodular biset function. Each of these problems admits ratio 2, which results in an overall ratio of 6. Their bound $n > k^3(k-1) + k$ was recently improved to $n > k(k-1)^2$ in [8], where also the degree bounded version of the problem is considered. Their method does not apply to digraphs. Furthermore, for the k -Connectivity Augmentation on undirected graphs with $n > k(k-1)^2$, it can be shown that our algorithm has a slightly better approximation ratio of 5.

References

- [1] A. Aazami, J. Cheriyan, and B. Laekhanukit. A bad example for the iterative rounding method for mincost k -connected spanning subgraphs. *Discrete Optimization*, 10(1):25–41, 2013.
- [2] J. Cheriyan and B. Laekhanukit. Approximation algorithms for minimum-cost k - (S, T) connected digraphs. Manuscript, 2010.
- [3] J. Cheriyan and L. Véggh. Approximating minimum-cost k -node connected subgraphs via independence-free graphs. Manuscript, 2013.
- [4] J. Cheriyan, S. Vempala, and A. Vetta. An approximation algorithm for the minimum-cost k -vertex connected subgraph. *SIAM J. Computing*, (4):1050–1055, 2003.
- [5] J. Fackharoenphol and B. Laekhanukit. An $O(\log^2 k)$ -approximation algorithm for the k -vertex connected subgraph problem. In *STOC*, pages 153–158, 2008.
- [6] A. Frank. Rooted k -connections in digraphs. *Discrete Applied Math.*, 157(6):1242–1254, 2009.
- [7] A. Frank and T. Jordán. Minimal edge-coverings of pairs of sets. *J. of Comb. Theory B*, 65:73–110, 1995.
- [8] T. Fukunaga, Z. Nutov, and R. Ravi. Iterative rounding approximation algorithms for degree bounded node-connectivity problems. Manuscript, 2013.
- [9] T. Jordán. On the optimal vertex-connectivity augmentation. *J. on Comb. Theory B*, 63:8–20, 1995.

- [10] G. Kortsarz, R. Krauthgamer, and J. R. Lee. Hardness of approximation for vertex-connectivity network design problems. *SIAM Journal on Computing*, 33(3):704–720, 2004.
- [11] G. Kortsarz and Z. Nutov. Approximating k -node connected subgraphs via critical graphs. *SIAM Journal on Computing*, 35(1):247–257, 2005.
- [12] G. Kortsarz and Z. Nutov. Approximating minimum-cost connectivity problems. In T. F. Gonzalez, editor, Chapter 58 in *Approximation Algorithms and Metaheuristics*. Chapman & Hall/CRC, 2007.
- [13] B. Laekhanukit. An improved approximation algorithm for minimum-cost subset k -connectivity. In *ICALP*, pages 13–24, 2011.
- [14] Y. Lando and Z. Nutov. Inapproximability of survivable networks. *Theoretical Computer Science*, 410(21-23):2122–2125, 2009.
- [15] Z. Nutov. Approximating minimum cost connectivity problems via uncrossable bifamilies. Manuscript 2010. Preliminary version in FOCS 2009, pages 417-426.
- [16] Z. Nutov. An almost $O(\log k)$ -approximation for k -connected subgraphs. In *SODA*, pages 922–931, 2009.
- [17] Z. Nutov. Approximating subset k -connectivity problems. In *WAOA*, pages 9–20, 2011.
- [18] R. Ravi and D. P. Williamson. An approximation algorithm for minimum-cost vertex-connectivity problems. *Algorithmica*, 18:21–43, 1997.
- [19] R. Ravi and D. P. Williamson. Erratum: an approximation algorithm for minimum-cost vertex-connectivity problems. *Algorithmica*, 34(1):98–107, 2002.