Theory of Computing Systems manuscript No.
(will be inserted by the editor)

Improved approximation algorithms for minimum cost node-connectivity augmentation problems

Zeev Nutov

Received: date / Accepted: date

Abstract Let $\kappa_G(s,t)$ denote the maximum number of pairwise internally disjoint st-paths in a graph $G=(V,E)$. For a set $T \subseteq V$ of terminals, G is k-T-connected if $\kappa_G(s,t) \geq k$ for all $s,t \in T$; if $T=V$ then G is k-connected. Given a root node s, G is k-(T,s)-connected if $\kappa_G(t,s) \geq k$ for all $t \in T$. We consider three well studied min-cost connectivity augmentation problems, where we are given a graph $G=(V,E)$ of connectivity k, and an additional edge set $\hat{E}$ on V with costs. The goal is to compute a minimum cost edge set $J \subseteq \hat{E}$ such that $G \cup J$ has connectivity $k+1$. In the k-T-Connectivity Augmentation problem G is k-T-connected and $G \cup J$ should be $(k+1)$-T-connected. In the k-Connectivity Augmentation problem G is k-connected and $G \cup J$ should be $(k+1)$-connected. In the k-(T,s)-Connectivity Augmentation problem G is k-(T,s)-connected and $G \cup J$ should be $(k+1)$-(T,s)-connected.

For the k-T-Connectivity Augmentation problem when $\hat{E}$ is an edge set on T we obtain ratio $O\left(\ln \frac{|T|}{|T|-k}\right)$, improving the ratio $O\left(\frac{|T|}{|T|-k} \cdot \ln \frac{|T|}{|T|-k}\right)$ of [29]. For the k-Connectivity Augmentation problem we obtain the following approximation ratios. For $n \geq 3k-5$, we obtain ratio 3 for directed graphs and 4 for undirected graphs, improving the previous ratio 5 of [29]. For directed graphs and $k=1$, or $k=2$ and n odd, we further improve to 2.5 the previous ratios 3 and 4, respectively. For the undirected 2-(T,s)-Connectivity Augmentation problem we achieve ratio $4\frac{2}{3}$, improving the previous best ratio 12 of [27]. For the special case when all the edges in $\hat{E}$ are incident to s, we give a polynomial time algorithm, improving the ratio $4\frac{17}{30}$ of [28,23] for this variant.

Keywords node-connectivity augmentation · approximation algorithm · crossing biset family

Z. Nutov
The Open University of Israel
Tel.: +972-9-778 1254
Fax: +972-9-778 0605
E-mail: nutov@openu.ac.il

1 Introduction

1.1 Problems and results

Let $\kappa_G(s,t)$ denote the maximum number of pairwise internally disjoint st-paths in a graph $G=(V,E)$. For a set $T \subseteq V$ of terminals, G is **k-T-connected** if $\kappa_G(s,t) \geq k$ for all $s,t \in T$; if $T=V$ then G is **k-connected**. Given a root node s, G is **k-(T,s)-connected** if $\kappa_G(t,s) \geq k$ for all $t \in T$. We consider three extensively studied minimum cost connectivity augmentation problems. In all problems we are given an integer $k \geq 0$, a graph $G=(V,E)$ of connectivity k, and an additional edge set $\hat{E}$ on V with costs. The goal is to compute a minimum cost edge set $J \subseteq \hat{E}$ such that $G \cup J$ has connectivity $k+1$. More formally, our problems are as follows.

> k-T-Connectivity Augmentation
> Here for a given set $T \subseteq V$ of terminals, G is k-T-connected and $G \cup J$ should be $(k+1)$-T-connected. We consider the version of the problem when $\hat{E}$ is an edge set on T, namely, every edge in $\hat{E}$ has both endnodes in T.

> k-Connectivity Augmentation
> Here G is k-connected and $G \cup J$ should be $(k+1)$-connected. This is a particular case of k-T-Connectivity Augmentation when $T=V$.

> k-(T,s)-Connectivity Augmentation
> Here we are also given a root node s and a set $T \subseteq V$ of terminals, G is k-(T,s)-connected, and $G \cup J$ should be $(k+1)$-(T,s)-connected.

One important particular case of k-(T,s)-Connectivity Augmentation is when all edges of positive cost are incident to s. This variant is closely related to Source Location problems, see [14,23].

These problems were studied extensively, see [1,4–6,9,12,14,19,21–23,27–29,32] for only a small sample of papers in the area. For $k=0$ and undirected graphs our problems include the Minimum Spanning Tree problem and the Steiner Tree problem; for directed graphs we get the Minimum Cost Strongly Connected Subgraph problem (that admits ratio 2 by taking a union of minimum cost in- and out-arborescences), and the Directed Steiner Tree problem.

The version of k-T-Connectivity Augmentation that we consider (when $\hat{E}$ is an edge set on T) admits ratio $O(\ln|T|)$; this was implicitly proved in [9], see also [3,29] for explicit proofs and generalizations. For $|T|>k$ the problem admits ratio $O\left(\frac{|T|}{|T|-k} \cdot \ln \frac{|T|}{|T|-k}\right)$ [29]. We improve the latter ratio as follows.

Theorem 1 *For both directed and undirected graphs, k-T-Connectivity Augmentation such that $\hat{E}$ is an edge set on T and $|T|>k$ admits ratio $O\left(\ln \frac{|T|}{|T|-k}\right)$.*

To state our result for the k-Connectivity Augmentation problem we need some definitions. Let q be the largest integer such that $2q-1 \leq n-k$, namely,

$q = \lfloor \frac{n-k+1}{2} \rfloor$. Let

$$\mu = \left\lfloor \frac{n}{q+1} \right\rfloor = \left\lfloor \frac{n}{\lfloor (n-k+3)/2 \rfloor} \right\rfloor = \begin{cases} \left\lfloor \frac{2n}{n-k+3} \right\rfloor & \text{if } n-k \text{ is odd} \\ \left\lfloor \frac{2n}{n-k+2} \right\rfloor & \text{if } n-k \text{ is even} \end{cases} .$$

It is not hard to see that:

- $\mu = 1$ if and only if $k = 0$, or $k = 1$, or $k = 2$ and n is odd.
- $\mu = 2$ if and only if one of the following holds: $k = 2$ and n is even, or $k \geq 3$ and one of the following holds: $n \geq 3k - 8$ and n, k have distinct parities, or $n \geq 3k - 5$ and n, k have the same parity.
- $\mu \leq 3$ if and only if one of the following holds: $n \geq 2k - 5$ and n, k have distinct parities, or $n \geq 2k - 3$ and n, k have the same parity.

Let $H(k)$ denote the kth harmonic number. For both directed and undirected graphs k-Connectivity Augmentation admits ratio $2H(\mu) + 2$ [29] (which is a constant unless $k = n - o(n)$), and also ratio $O(\ln(n-k))$ [29]. Specifically, for $n \geq 3k - 5$, the previous best ratio was 5, for both directed and undirected graphs. For small values of k better ratios are known: $k + 2$ for $k \leq 2$ in the case of directed graphs [22], and $\lceil k/2 \rceil + 1$ for $k \leq 6$ in the case of undirected graphs [2,8,22]. We prove the following (for comparison with previous ratios see Table 1):

Theorem 2 k-Connectivity Augmentation *admits the following approximation ratios:*

(i) *For directed graphs, ratio* $H(\mu) + \frac{3}{2}$. *In particular:*
  - *For* $k = 1$, *and for* $k = 2$ *and* n *odd,* $\mu = 1$, $H(\mu) = 1$, *so the ratio is* 2.5.
  - *For* $n \geq 3k - 5$, $\mu \leq 2$, $H(\mu) \leq 3/2$, *so the ratio is* 3.
  - *For* $n \geq 2k - 3$, $\mu \leq 3$, $H(\mu) \leq 11/6$, *so the ratio is* $3\frac{1}{3}$.

(ii) *For undirected graphs, ratio* $2H(\mu)+1$. *In particular, for* $n \geq 3k-5$, $\mu \leq 2$, $H(\mu) \leq 3/2$, *so the ratio is* 4.

For directed graphs our ratios improve over the previous ones for any $k \geq 1$. For undirected graphs our ratio matches the best known ratio 4 for $k = 6, 7$, and it improves over the previous ratios for any $k \geq 8$.

			directed		**undirected**	
range	μ	$H(\mu)$	previous	this paper	previous	this paper
$k = 0$	1	1	2		in P	
$k = 1, 2$	1	1	$3, 4$ [22]	2.5	2 [20,2]	
$3 \leq k \leq 6$	2	1.5	5 [29]	3	$\lceil k/2 \rceil + 1$ [8,22]	
$n \geq 3k - 5$	2	1.5	5 [29]	3	5 [29]	4
$n \geq 2k - 3$	3	$1\frac{5}{6}$	$5\frac{2}{3}$ [29]	$3\frac{1}{3}$	$5\frac{2}{3}$ [29]	$4\frac{2}{3}$
$n < 2k - 3$			$2H(\mu) + 2$ [29]	$H(\mu) + 1.5$	$2H(\mu) + 2$ [29]	$2H(\mu) + 1$

Table 1 Previous and our ratios for k-Connectivity Augmentation; for $k = 2$ our ratio 2.5 for directed graphs is valid when n is odd.

We now state our results for the k-(T,s)-Connectivity Augmentation problem. The best known ratio for k-(T,s)-Connectivity Augmentation is $O(k \log k)$, and it was 12 for $k=2$ [27]. For the version when all edges in $\hat{E}$ are incident to s the best ratio was $2H(2k+1)$ [23], which for $k=2$ is $2H(5)=4\frac{17}{30}>4.5$. We consider the case $k=2$, and significantly improve over the previous ratios.

Theorem 3 *Undirected* 2-(T,s)-Connectivity Augmentation *admits ratio* $4\frac{2}{3}$*; if all edges in* $\hat{E}$ *are incident to* s*, then the problem admits a polynomial time algorithm.*

The rest of the Introduction we survey some related work. In Section 2 we cast our problems as a problem of finding a minimum cost edge cover of a biset family, and state some properties of relevant biset families. In subsequent section 3, 4, and 5 we prove the corresponding theorems. In Section 6 we provide a short proof of a theorem from [29] that is used by our algorithms.

1.2 Some previous and related work

We consider *node-connectivity* problems for which classic techniques like the primal dual method [17] and iterative rounding [18] do not seem to be applicable directly. Ravi and Williamson [31] gave an example of a k-Connectivity Augmentation instance when the primal dual method has ratio $\Omega(k)$. Aazami, Cheriyan and Laekhanukit [1] presented a related instance for which the basic optimal solution to the LP-relaxation has all variables of value $O(1/\sqrt{k})$, ruling out the iterative rounding method. On the other hand, several works showed that node-connectivity problems can be decomposed into a small number p of "good" problems. The bound on p was subsequently improved, culminating in the currently best known bounds $O(\log \frac{n}{n-k})$ for directed/undirected k-Connectivity Augmentation [29], and $O(k)$ for undirected k-(T,s)-Connectivity Augmentation [27]. In fact, [25] shows that for $k=\Omega(n)$ the approximability of the directed and undirected variants of these problems is the same, up to a factor of 2. We refer the reader to [4,24] for various hardness results on k-(T,s)-Connectivity Augmentation. We note that the version of k-Connectivity Augmentation when any edge can be added by a cost of 1 can be solved in polynomial time for both directed [12] and undirected [32] graphs. But for general costs, determining whether k-Connectivity Augmentation admits a constant ratio for $k=n-o(n)$ is one of the most challenging problems.

We mention some related work on the more general k-Connected Subgraph problem, where we seek a minimum cost k-connected spanning subgraph; k-Connectivity Augmentation is a particular case, when the target connectivity is $k+1$ and the edges of cost zero of the input graph form a k-connected spanning subgraph. Many papers that considered the k-Connected Subgraph problem built on the algorithm of Frank and Tardos [13] for a related problem of finding a minimum cost k-outconnected subgraph [20,2,8,6,21,9,5], but most papers that considered high values of k in fact designed algorithms for k-Connectivity

Augmentation [6,21,9,29]. These papers use the fact that ratio ρ w.r.t. the LP-relaxation for k-Connectivity Augmentation implies ratio $\rho H(k) = \rho \cdot O(\log k)$ for k-Connected Subgraph [30]. Recently, Cheriyan and Végh [5] showed that for undirected graphs with $n = \Omega(k^4)$ this $O(\log k)$ factor can be saved and ratio 6 can be achieved by a new decomposition of the problem. The bound $n = \Omega(k^4)$ of [5] was improved to $n = \Omega(k^3)$ in [16].

In the more general Survivable Network problem, we are given connectivity requirements $\{r_{uv} : u, v \in V\}$. The goal is to compute a minimum cost subgraph that has r_{uv} internally-disjoint uv-paths for all $u, v \in V$. For undirected graphs the problem admits ratio $O(k^3 \log n)$ due to Chuzhoy and Khanna [7]. For directed graphs, no non-trivial ratio is known even for 2-(T, s)-Connectivity Augmentation.

2 Preliminaries on biset families

While edge-cuts of a graph correspond to node subsets, a natural way to represent a node-cut of a graph is by a pair of sets called a "biset".

Definition 1 An ordered pair $\mathbb{A} = (A, A^+)$ of subsets of a groundset V is called a **biset** if $A \subseteq A^+$; A is the **inner part** and A^+ is the **outer part** of $\mathbb{A}$, and $\partial(\mathbb{A}) = \partial\mathbb{A} = A^+ \setminus A$ is the **boundary** of $\mathbb{A}$. The **co-set** of a biset $\mathbb{A} = (A, A^+)$ is $A^* = V \setminus A^+$; the **co-biset** of $\mathbb{A}$ is $\mathbb{A}^* = (A^*, V \setminus A)$.

Definition 2 A **biset family** is a family of bisets. The **co-family** of a biset family $\mathcal{F}$ is $\mathcal{F}^* = \{\mathbb{A}^* : \mathbb{A} \in \mathcal{F}\}$. $\mathcal{F}$ is **symmetric** if $\mathcal{F} = \mathcal{F}^*$.

Definition 3 An **edge covers a biset** $\mathbb{A}$ if it goes from A to A^*. Let $\delta_E(\mathbb{A})$ denote the set of edges in E that cover $\mathbb{A}$. The **residual family** of a biset family $\mathcal{F}$ w.r.t. an edge-set/graph J is denoted $\mathcal{F}^J$ and it consists of the members in $\mathcal{F}$ not covered by any $e \in J$, namely, $\mathcal{F}^J = \{\mathbb{A} \in \mathcal{F} : \delta_J(\mathbb{A}) = \emptyset\}$. We say that an **edge set/graph** J **covers** $\mathcal{F}$ or that J **is an** $\mathcal{F}$**-edge-cover** if every $\mathbb{A} \in \mathcal{F}$ is covered by some $e \in J$, namely, if $\mathcal{F}^J = \emptyset$.

We say that $\mathbb{A}$ is an st**-biset** if $s \in A$ and $t \in A^*$. Let $G = (V, E)$ is a (directed or undirected) graph and let $s, t \in V$ with $st \notin E$. In biset terms, Menger's Theorem says that $\kappa_G(s, t) \leq |\partial\mathbb{A}|$ for any st-biset $\mathbb{A}$ with $\delta_E(\mathbb{A}) = \emptyset$, and

$$\kappa_G(s, t) = \min\{|\partial\mathbb{A}| : \mathbb{A} \text{ is an } st\text{-biset}, \delta_E(\mathbb{A}) = \emptyset\} .$$

Given an instance of k-T-Connectivity Augmentation we will assume that G has no edge between two terminals by subdividing by a new node every such edge. Similarly, given an instance of k-(T, s)-Connectivity Augmentation we will assume that G has no edge from T to s. Then the biset families we need to cover in k-T-Connectivity Augmentation, k-Connectivity Augmentation, and k-(T, s)-Connectivity Augmentation, respectively, are:

$$\mathcal{F}_{k\text{-}T} = \{\mathbb{A} : |\partial\mathbb{A}| = k, \delta_E(\mathbb{A}) = \emptyset, A \cap T \neq \emptyset, A^* \cap T \neq \emptyset\} \quad (1)$$

$$\mathcal{F}_k = \{\mathbb{A} : |\partial\mathbb{A}| = k, \delta_E(\mathbb{A}) = \emptyset, A \neq \emptyset, A^* \neq \emptyset\} \quad (2)$$

$$\mathcal{F}_{k\text{-}(T,s)} = \{\mathbb{A} : |\partial\mathbb{A}| = k, \delta_E(\mathbb{A}) = \emptyset, A \cap T \neq \emptyset, s \in A^*\} \quad (3)$$

Recall that in the case of k-T-Connectivity Augmentation, we consider the version when only edges between nodes of T can be added. Then it is sufficient to cover the projection of $\mathcal{F}_{k\text{-}T}$ on T, namely the following biset family on T:

$$\mathcal{T} = \{(A \cap T, A^+ \cap T) : \mathbb{A} \in \mathcal{F}_{k\text{-}T}\} \tag{4}$$

Note that if $T = V$ then $\mathcal{F}_k = \mathcal{F}_{k\text{-}T} = \mathcal{T}$.

We thus consider the following generic algorithmic problem.

Biset-Family Edge-Cover
Input: A graph $(V, \hat{E})$ with edge-costs $\{c_e : e \in \hat{E}\}$ and a biset family $\mathcal{F}$.
Output: A minimum cost $\mathcal{F}$-edge-cover $J \subseteq F$.

Here the biset family $\mathcal{F}$ may not be given explicitly, and a polynomial time implementation in $n = |V|$ of our algorithms requires that the following query can be answered in time polynomial in n: Given an edge set/graph J on V and $s, t \in V$, find the inclusionwise minimal and the inclusionwise maximal members of the family $\{\mathbb{A} \in \mathcal{F}^J : s \in A, t \in V \setminus A^+\}$, if non-empty. For biset families arising from our problems, this query can be answered in polynomial time using max-flow min-cut computations (we omit the standard implementation details).

Definition 4 The **intersection** and the **union** of two bisets $\mathbb{A}, \mathbb{B}$ are defined by $\mathbb{A} \cap \mathbb{B} = (A \cap B, A^+ \cap B^+)$ and $\mathbb{A} \cup \mathbb{B} = (A \cup B, A^+ \cup B^+)$. The biset $\mathbb{A} \setminus \mathbb{B}$ is defined by $\mathbb{A} \setminus \mathbb{B} = (A \setminus B^+, A^+ \setminus B)$. We say that $\mathbb{B}$ **contains** $\mathbb{A}$ and write $\mathbb{A} \subseteq \mathbb{B}$ if $A \subseteq B$ and $A^+ \subseteq B^+$. We say that $\mathbb{A}, \mathbb{B}$ **intersect** if $A \cap B \neq \emptyset$ and $\mathbb{A}, \mathbb{B}$ **cross** if $A \cap B \neq \emptyset$ and $A^+ \cup B^+ \neq V$.

The following properties of bisets are known and easy to verify.

Fact 1 *For any bisets $\mathbb{A}, \mathbb{B}$ the following holds. If a directed/undirected edge e covers one of $\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cup \mathbb{B}$ then e covers one of $\mathbb{A}, \mathbb{B}$; if e is an undirected edge, then if e covers one of $\mathbb{A} \setminus \mathbb{B}, \mathbb{B} \setminus \mathbb{A}$, then e covers one of $\mathbb{A}, \mathbb{B}$. Furthermore*

$$|\partial\mathbb{A}| + |\partial\mathbb{B}| = |\partial(\mathbb{A} \cap \mathbb{B})| + |\partial(\mathbb{A} \cup \mathbb{B})| = |\partial(\mathbb{A} \setminus \mathbb{B})| + |\partial(\mathbb{B} \setminus \mathbb{A})| \ .$$

Definition 5 A biset family $\mathcal{F}$ is **intersecting/crossing** if $\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cup \mathbb{B} \in \mathcal{F}$ whenever $\mathbb{A}, \mathbb{B}$ intersect/cross. Let us say that a crossing biset family $\mathcal{F}$ is p**-crossing** if for any $\mathbb{A}, \mathbb{B} \in \mathcal{F}$ that intersect the following holds: $\mathbb{A} \cap \mathbb{B} \in \mathcal{F}$ if $|A \cup B| \leq n - p$, and $\mathbb{A} \cup \mathbb{B} \in \mathcal{F}$ if $|A \cup B| \leq n - p - 1$.

The following known lemma (c.f. [19,27]) can be deduced from Fact 1.

Lemma 2 *If G is k-T-connected then $\mathcal{T}$ and $\mathcal{T}^*$ are both k-crossing, where $\mathcal{T}$ is defined in (4). Furthermore, if $T = V$ then $|\partial\mathbb{A}| = k$ for all $\mathbb{A} \in \mathcal{T}$.*

Note that Fact 1 implies that if $\mathcal{F}$ is intersecting, crossing, or k-crossing, then so is the residual family $\mathcal{F}^J$ of $\mathcal{F}$, for any J.

Let $\tau(\mathcal{F})$ denote the optimal value of a standard **Biset-LP** for the problem of edge-covering a biset family $\mathcal{F}$, namely:

$$\textbf{Biset-LP} \quad \tau(\mathcal{F}) = \min\left\{\sum_{e\in\hat{E}} c_e x_e : \sum_{e\in\delta(\mathbb{A})} x_e \geq 1 \ \forall \mathbb{A} \in \mathcal{F}, x_e \geq 0 \ \forall e \in \hat{E}\right\} .$$

Directed Biset-Family Edge-Cover with intersecting $\mathcal{F}$ admits a polynomial time algorithm that computes an $\mathcal{F}$-edge-cover of cost $\tau(\mathcal{F})$ [11]; for undirected graphs the cost is $2\tau(\mathcal{F})$ for intersecting $\mathcal{F}$, by a standard "bidirection" reduction to the directed case.

In terms of bisets, we prove the following two theorems that imply Theorems 1 and 2. In these theorems $\mathcal{F}$ is a biset family on a groundset V of size $n = |V|$. Let us say that Biset-Family Edge-Cover admits **LP-ratio** ρ if there exists a polynomial time algorithm that computes an $\mathcal{F}$-cover of cost $\rho \cdot \tau(\mathcal{F})$.

Theorem 4 (Implies Theorem 1) Biset-Family Edge-Cover *such that $\mathcal{F}$ and $\mathcal{F}^*$ are k-crossing and $n \geq k+1$ admits LP-ratio $O(\ln \mu)$.*

Theorem 5 (Implies Theorem 2) Biset-Family Edge-Cover *such that $\mathcal{F}$ and $\mathcal{F}^*$ are k-crossing, $|\partial\mathbb{A}| \geq k$ for all $\mathbb{A} \in \mathcal{F}$, and $n \geq k+3$, admits the following LP-ratios:*

(i) *For directed graphs, ratio $H(\mu) + \frac{3}{2}$. In particular:*
   - *For $k = 1$, and for $k = 2$ and n odd, $\mu = 1$, $H(\mu) = 1$, so the ratio is* 2.5.
   - *For $n \geq 3k-5$, $\mu \leq 2$, $H(\mu) \leq 3/2$, so the ratio is* 3.
   - *for $n \geq 2k-3$, $\mu \leq 3$, $H(\mu) \leq 11/6$, so the ratio is* $3\frac{1}{3}$.

(ii) *For undirected graphs, ratio $2H(\mu)+1$. In particular, for $n \geq 3k-5$, $\mu \leq 2$, $H(\mu) \leq 3/2$, so the ratio is* 4.

Theorem 3 relies on different "uncrossing" properties of the family $\mathcal{F}_{2\text{-}(T,s)}$, that will be given in Section 5

The following definition plays a key role in our algorithms.

Definition 6 The inclusionwise minimal members of a biset family $\mathcal{F}$ are called **$\mathcal{F}$-cores**, or simply **cores**, if $\mathcal{F}$ is clear from the context. 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. For $\mathbb{C} \in \mathcal{C}(\mathcal{F})$, the **halo-family** $\mathcal{F}(\mathbb{C})$ of $\mathbb{C}$ is the family of those members of $\mathcal{F}$ that contain $\mathbb{C}$ and contain no $\mathcal{F}$-core distinct from $\mathbb{C}$.

Let us say that two biset families $\mathcal{A}, \mathcal{B}$ are **independent** if no $\mathbb{A} \in \mathcal{A}$ and $\mathbb{B} \in \mathcal{B}$ cross. Note that if $\mathcal{F}_1, \ldots, \mathcal{F}_p$ is a collection of pairwise independent subfamilies of a biset family $\mathcal{F}$, then for $i \neq j$ no directed edge can cover $\mathbb{A}_i \in \mathcal{F}_i$ and $\mathbb{A}_j \in \mathcal{F}_j$, and thus $\sum_{i=1}^{p} \tau(\mathcal{F}_i) \leq \tau(\mathcal{F})$.

The following statement summarizes several relevant properties of halo families of crossing biset families, c.f. [22,9,3,29]. We provide a proof for completeness of exposition.

Lemma 3 *For any crossing biset family $\mathcal{F}$ the following holds.*

(i) *For any $\mathcal{F}$-core $\mathbb{C}$, $\mathcal{F}(\mathbb{C})$ is a crossing family and $\mathcal{F}(\mathbb{C})^* = \{\mathbb{A}^* : \mathbb{A} \in \mathcal{F}(\mathbb{C})\}$ (the co-family of $\mathcal{F}(\mathbb{C})$) is an intersecting family.*
(ii) *Halo families of distinct cores are independent.*
(iii) *For any $\mathcal{F}$-core $\mathbb{C}$, if J is an inclusion minimal edge set that covers $\mathcal{F}(\mathbb{C})$ then $\mathcal{C}(\mathcal{F}^J) = \mathcal{C}(\mathcal{F}) \setminus \{\mathbb{C}\}$.*

Proof We prove (i). Let $\mathbb{A}, \mathbb{B} \in \mathcal{F}(\mathbb{C})$ cross. Then $\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cup \mathbb{B} \in \mathcal{F}$. Since $\mathbb{A} \cap \mathbb{B} \subseteq \mathbb{A} \subseteq \mathbb{A} \cup \mathbb{B}$ and $\mathbb{A} \in \mathcal{F}(\mathbb{C})$, $\mathbb{A} \cap \mathbb{B} \in \mathcal{F}(\mathbb{C})$ and $\mathbb{C} \subseteq \mathbb{A} \cup \mathbb{B}$. We claim that $\mathbb{A} \cup \mathbb{B}$ contains no core $\mathbb{C}'$ distinct from $\mathbb{C}$. Otherwise, since none of $\mathbb{A}, \mathbb{B}$ can contain $\mathbb{C}'$, we must have that $\mathbb{C}', \mathbb{A}$ cross or $\mathbb{C}', \mathbb{B}$ cross, so $\mathbb{C}' \cap \mathbb{A} \in \mathcal{F}$ or $\mathbb{C}' \cap \mathbb{B} \in \mathcal{F}$; this contradicts that $\mathbb{C}'$ is a core. Thus $\mathcal{F}(\mathbb{C})$ is a crossing family. We prove that $\mathcal{F}(\mathbb{C})^*$ is an intersecting family. Let $\mathbb{A}, \mathbb{B} \in \mathcal{F}(\mathbb{C})^*$ intersect. Then $\mathbb{C} \subseteq \mathbb{A}^* \cap \mathbb{B}^*$ so $\mathbb{A}, \mathbb{B}$ cross. Thus since $\mathcal{F}(\mathbb{C})$ is a crossing family, we get that $\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cup \mathbb{B} \in \mathcal{F}(\mathbb{C})^*$.

We prove (ii). Let $\mathbb{A}_1 \in \mathcal{F}(\mathbb{C}_1)$ and $\mathbb{A}_2 \in \mathcal{F}(\mathbb{C}_2)$ cross, for $\mathbb{C}_1, \mathbb{C}_2 \in \mathcal{C}(\mathcal{F})$. Then $\mathbb{A}_1 \cap \mathbb{A}_2 \in \mathcal{F}$, so $\mathbb{A}_1 \cap \mathbb{A}_2$ contains some $\mathcal{F}$-core $\mathbb{C}$. We have $\mathbb{C} = \mathbb{C}_1$ since $\mathbb{C} \subseteq \mathbb{A}_1$ and $\mathbb{C} = \mathbb{C}_2$ since $\mathbb{C} \subseteq \mathbb{A}_2$, hence $\mathbb{C}_1 = \mathbb{C}_2$.

Part (iii) follows from part (ii), since every $e \in J$ covers some biset in $\mathcal{F}(\mathbb{C})$ (by the minimality of J) and thus by (ii) cannot cover a core distinct from $\mathbb{C}$. $\square$

The following statement was implicitly proved in [9] (see also [3]) and explicitly in [29]. We provide a proof for completeness of exposition.

Theorem 6 *Directed* Biset-Family Edge-Cover *with crossing $\mathcal{F}$ admits a polynomial time algorithm that given $\mathcal{C} \subseteq \mathcal{C}(\mathcal{F})$ and an integer $0 \leq t \leq |\mathcal{C}|$ computes an edge set $J \subseteq E$ such that the following holds:*

- *$\mathcal{C}(\mathcal{F}^J) = \mathcal{C}(\mathcal{F}) \setminus \mathcal{C}'$ for some $\mathcal{C}' \subseteq \mathcal{C}$ with $|\mathcal{C}'| = |\mathcal{C}| - t$.*
- *$c(J) \leq (H(|\mathcal{C}|) - H(t)) \cdot \tau(\mathcal{F}')$, where $\mathcal{F}'$ is the family of those members of $\mathcal{F}$ that contain no core in $\mathcal{C}(\mathcal{F}) \setminus \mathcal{C}$.*

Proof Consider the following algorithm. Start with a partial solution $J = \emptyset$. While $|\mathcal{C} \cap \mathcal{C}(\mathcal{F}^J)| \geq t+1$ continue with iterations. At iteration i, compute for each $\mathbb{C} \in \mathcal{C} \cap \mathcal{C}(\mathcal{F}^J)$ an optimal inclusion minimal edge cover $J_\mathbb{C}$ of the family $\mathcal{F}^J(\mathbb{C})$ (the halo family of $\mathbb{C}$ in $\mathcal{F}^J$); then add to J a minimum cost edge set J_i among the edge sets $\left\{J_\mathbb{C} : \mathbb{C} \in \mathcal{C} \cap \mathcal{C}\left(\mathcal{F}^J\right)\right\}$. By part (i) of Lemma 3, each $J_\mathbb{C}$ can be computed in polynomial time and $c(J_\mathbb{C}) = \tau(\mathcal{F}^J(\mathbb{C}))$. By part (ii) of Lemma 3, $\sum_{\mathbb{C} \in \mathcal{C} \cap \mathcal{C}(\mathcal{F}^J)} c(J_\mathbb{C}) \leq \tau(\mathcal{F}')$. Thus there is $\mathbb{C} \in \mathcal{C}$ such that $c(J_\mathbb{C}) \leq \tau(\mathcal{F}')/|\mathcal{C} \cap \mathcal{C}(\mathcal{F}^J)|$. By part (iii) of of Lemma 3, at iteration i we have $|\mathcal{C} \cap \mathcal{C}(\mathcal{F}^J)| \leq |\mathcal{C}| - i + 1$. Thus $c(J_i) \leq \tau(\mathcal{F}')/(|\mathcal{C}| - i + 1)$ at iteration i. The number of iterations is $|\mathcal{C}| - t$. Consequently,

$$c(J) \leq \sum_{i=1}^{|\mathcal{C}|-t} c(J_i) \leq \tau(\mathcal{F}') \sum_{i=1}^{|\mathcal{C}|-t} \frac{1}{|\mathcal{C}| - i + 1} = (H(|\mathcal{C}|) - H(t)) \cdot \tau(\mathcal{F}') ,$$

and the statement follows. $\square$

Note that for $t = 0$, the edge set J in Theorem 6 covers the family of those members of $\mathcal{F}$ that contain no core in $\mathcal{C}(\mathcal{F}) \setminus \mathcal{C}$ and $c(J) \leq H(|\mathcal{C}|)\tau(\mathcal{F}')$; if also $\mathcal{C} = \mathcal{C}(\mathcal{F})$ then J covers $\mathcal{F}$ and has cost $c(J) \leq H(\nu(\mathcal{F})) \cdot \tau(\mathcal{F})$.

3 Ratio $O(\ln \mu)$ for k-crossing families (Theorem 4)

Recall that q is a parameter eventually set to $q = \lfloor \frac{n-k+1}{2} \rfloor$, and $\mu = \left\lfloor \frac{n}{q+1} \right\rfloor$. Let us say that $\mathbb{A}$ is a **small biset/core** if $|A| \leq q$, and $\mathbb{A}$ is a **large biset/core** otherwise. We mention some definitions from [29] needed for the proof of Theorem 4.

Definition 7 A biset family $\mathcal{F}$ is **intersection-closed** if $\mathbb{A} \cap \mathbb{B} \in \mathcal{F}$ for any intersecting $\mathbb{A}, \mathbb{B} \in \mathcal{F}$. An intersection-closed $\mathcal{F}$ is **q-semi-intersecting** if $|A| \leq q$ for every $\mathbb{A} \in \mathcal{F}$ and if $\mathbb{A} \cup \mathbb{B} \in \mathcal{F}$ for any intersecting $\mathbb{A}, \mathbb{B} \in \mathcal{F}$ with $|A \cup B| \leq q$. The **q-truncated family** of $\mathcal{F}$ is $\mathcal{F}_{\leq q} := \{\mathbb{A} \in \mathcal{F} : |A| \leq q\}$. namely, $\mathcal{F}_{\leq q}$ is the family of the small bisets in $\mathcal{F}$.

We obtain a q-semi-intersecting family from a k-crossing family as follows.

Lemma 4 *Let $\mathcal{F}$ be a k-crossing biset family. If $2q-1 \leq n-k$ and $q \leq n-k-1$ (in particular if $q \leq \left\lfloor \frac{n-k+1}{2} \right\rfloor$ and $n \geq k+3$) then $\mathcal{F}_{\leq q}$ is q-semi-intersecting.*

Proof Let $\mathbb{A}, \mathbb{B} \in \mathcal{F}_{\leq q}$ intersect. Then $|A \cup B| \leq |A| + |B| - 1 \leq 2q - 1 \leq n - k$. Thus $\mathbb{A} \cap \mathbb{B} \in \mathcal{F}_{\leq q}$. If $|A \cup B| \leq q \leq n-k-1$ then $\mathbb{A} \cup \mathbb{B} \in \mathcal{F}_{\leq q}$. Hence if both $2q-1 \leq n-k$ and $q \leq n-k-1$, then $\mathcal{F}_{\leq q}$ is q-semi-intersecting. □

The following theorem is the main result of [29]. It says that if $\mathcal{F}$ is q-semi-intersecting, then we can find a "cheap" edge set J such that $\nu(\mathcal{F}^J)$ is "small". We will provide a relatively simple proof of this theorem in Section 6.

Theorem 7 ([29]) *Directed* Biset-Family Edge-Cover *with q-semi-intersecting $\mathcal{F}$ admits a polynomial time algorithm that computes an edge-set $J \subseteq E$ such that $\nu(\mathcal{F}^J) \leq \lfloor n/(q+1) \rfloor$ and $c(J) \leq \tau(\mathcal{F})$.*

From Theorem 7 and Lemma 4 we have the following.

Corollary 1 *Directed* Biset-Family Edge-Cover *with k-crossing $\mathcal{F}$ and $n \geq k+3$ admits a polynomial time algorithm that for any $q \leq \lfloor \frac{n-k+1}{2} \rfloor$ computes $J \subseteq E$ such that $\nu(\mathcal{F}^J_{\leq q}) \leq \lfloor n/(q+1) \rfloor$ and $c(J) \leq \tau(\mathcal{F}_{\leq q})$.*

We note that each of the statements in Theorem 7, Corollary 1, and Theorem 6, applies also for undirected graphs and symmetric $\mathcal{F}$, but with an additional factor of 2 in the cost. In this case we have $c(J) \leq 2\tau(\mathcal{F})$ in Theorem 7 and Corollary 1, and $c(J) \leq 2H(|\mathcal{C}|) \cdot \tau(\mathcal{F}')$ in Theorem 6. This is achieved by the following standard reduction. In each of the cases, we bidirect the edges of G (namely, replace every undirected edge e with endnodes u, v by two opposite directed edges uv, vu of cost c_e each), compute a set of directed edges for the obtained directed problem, and return the corresponding set of undirected edges.

A weaker version of the following statement is implicitly proved in [26].

Lemma 5 *Let $\mathcal{F}$ be a biset family such that both $\mathcal{F}$ and $\mathcal{F}^*$ are k-crossing. Then $\nu(\mathcal{F}) \leq \nu(\mathcal{F}_{\leq q}) + \nu(\mathcal{F}^*_{\leq q}) + \mu^2 H(\mu)$.*

Proof Note that for any distinct $\mathbb{A}, \mathbb{B} \in \mathcal{C}(\mathcal{F})$ we have $\mathbb{A} \cap \mathbb{B} \notin \mathcal{F}$. We show that if an arbitrary $\mathcal{A} \subseteq \mathcal{F}$ has this property then $|\mathcal{A}| \leq \nu(\mathcal{F}_{\leq q}) + \nu(\mathcal{F}^*_{\leq q}) + \mu^2 H(\mu)$. Let $\mathcal{B} = \{\mathbb{A} \in \mathcal{A} : |A|, |A^*| \geq q+1\}$. Clearly, $|\mathcal{A}| \leq |\mathcal{A}_{\leq q}| + |\mathcal{A}^*_{\leq q}| + |\mathcal{B}|$. Note that $|\mathcal{A}_{\leq q}| \leq \nu(\mathcal{F}_{\leq q})$, since $\mathcal{F}_{\leq q}$ is intersection closed, by Lemma 4. Similarly, $|\mathcal{A}^*_{\leq q}| \leq \nu(\mathcal{F}^*_{\leq q})$. To see that $|\mathcal{B}| \leq \mu^2 H(\mu)$, consider the hypergraph $\mathcal{H}$ formed by the inner parts of the bisets in $\mathcal{B}$. Let Δ be the maximum degree in $\mathcal{H}$. Recall that a hitting set of a hypergraph/set family is a set U of nodes that intersects every hyperedge/set of $\mathcal{H}$. A fractional hitting set is a function $h : V \longrightarrow [0,1]$ such that $h(A) = \sum_{v \in A} h(v) \geq 1$ for every hyperedge A. It is known that if h is a fractional hitting set of $\mathcal{H}$ then $\mathcal{H}$ has a hitting set of size at most $H(\Delta) \cdot h(V)$. Note the following:

(i) $\Delta \leq \mu$. This is so since no two bisets in $\mathcal{B}$ cross, and thus for any $v \in V$ the sets in the family $\{A^* : \mathbb{A} \in \mathcal{B}, v \in A\}$ are pairwise disjoint; hence their number is at most $\nu(\mathcal{B}^*) \leq \left\lfloor \frac{n}{q+1} \right\rfloor = \mu$.

(ii) $\mathcal{H}$ has a hitting set U of size $|U| \leq \mu H(\Delta) \leq \mu H(\mu)$. This is so since $\mathcal{H}$ has a fractional hitting set h of value μ defined by $h(v) = \frac{1}{q+1}$ for all $v \in V$.

Since $\mathcal{H}$ has at most $|U| \cdot \Delta$ hyperedges, the bound $|\mathcal{B}| \leq \mu^2 H(\mu)$ follows. □

The algorithm as in Theorem 4 is as follows.

Algorithm 1: DIRECTED-COVER$(\mathcal{F}, \hat{G}, c)$ ($\mathcal{F}, \mathcal{F}^*$ are both k-crossing)

1 **compute** $J_1 \subseteq E$ with $\nu(\mathcal{F}^{J_1}_{\leq q}) \leq \mu$ and $c(J_1) \leq \tau(\mathcal{F}_{\leq q})$ using the algorithm from **Corollary 1**
compute a similar edge set $J^*_1 \subseteq E$ for the family $\mathcal{F}^*_{\leq q}$
2 **compute** $J_2 \subseteq E$ covering $\mathcal{F}^{J_1 \cup J^*_1}$ using the algorithm from **Lemma 3**
3 **return** $J = J_1 \cup J^*_1 \cup J_2$

By Lemma 5, $\left|\mathcal{C}\left(\mathcal{F}^{J_1 \cup J^*_1}\right)\right| = O(\mu^2 \ln \mu)$ and thus $c(J_2) = \tau(\mathcal{F}) O(\ln \mu)$. Consequently, the cost of the solution computed is bounded by

$$\tau(\mathcal{F})(c(J_1) + c(J^*_1) + c(J_2)) \leq \tau(\mathcal{F})(1 + 1 + O(\ln \mu)) = O(\ln \mu) \ .$$

4 Proof of Theorem 5

4.1 Directed graphs

Recall that biset families $\mathcal{A}, \mathcal{B}$ are independent if no $\mathbb{A} \in \mathcal{A}$ and $\mathbb{B} \in \mathcal{B}$ cross. In Lemma 3 and Theorem 6 we used the observation that if $\mathcal{A}, \mathcal{B}$ are two independent subfamilies of a biset family $\mathcal{F}$ then $\tau(\mathcal{A}) + \tau(\mathcal{B}) \leq \tau(\mathcal{F})$. Here we use a different novel setting, where $\mathcal{A}, \mathcal{B}$ may not be independent, but $\mathcal{A} \setminus \mathcal{B}$ and $\mathcal{B} \setminus \mathcal{A}$ are independent.

Lemma 6 *Let $\mathcal{A},\mathcal{B}$ be subfamilies of a biset family $\mathcal{F}$ such that $\mathcal{A}\cup\mathcal{B}=\mathcal{F}$ and the families $\mathcal{A}\setminus\mathcal{B}$ and $\mathcal{B}\setminus\mathcal{A}$ are independent. Suppose that for any $J\subseteq\hat{E}$* Biset-Family Edge-Cover *with $\mathcal{A}^J$ admits LP-ratio α and with $\mathcal{B}^J$ admits LP-ratio β. Then* Biset-Family Edge-Cover *with $\mathcal{F}$ admits LP-ratio $\alpha+\beta-\frac{\alpha\beta}{\alpha+\beta}$.*

Proof We claim that the following algorithm achieves LP-ratio $\alpha+\beta-\frac{\alpha\beta}{\alpha+\beta}$:

Algorithm 2: INDEPENDENT-COVER($\mathcal{A},\mathcal{B},G,c$)

1 $J_\mathcal{A}\leftarrow\alpha$-approximate $\mathcal{A}$-cover $J'_\mathcal{B}\leftarrow\beta$-approximate $\mathcal{B}^{J_\mathcal{A}}$-cover
2 $J_\mathcal{B}\leftarrow\beta$-approximate $\mathcal{B}$-cover $J'_\mathcal{A}\leftarrow\alpha$-approximate $\mathcal{A}^{J_\mathcal{B}}$-cover
3 **return** the cheaper edge set J among $J_\mathcal{A}\cup J'_\mathcal{B}, J_\mathcal{B}\cup J'_\mathcal{A}$.

Note that since $\mathcal{A}\setminus\mathcal{B}$ and $\mathcal{B}\setminus\mathcal{A}$ are independent, so are $\mathcal{B}^{J_\mathcal{A}}$ and $\mathcal{A}^{J_\mathcal{B}}$. Thus no $\mathbb{B}\in\mathcal{B}^{J_\mathcal{A}}$ and $\mathbb{A}\in\mathcal{A}^{J_\mathcal{B}}$ cross, so no directed edge can cover both $\mathbb{A}$ and $\mathbb{B}$. Therefore

$$\tau\left(\mathcal{B}^{J_\mathcal{A}}\right)+\tau\left(\mathcal{A}^{J_\mathcal{B}}\right)\leq\tau(\mathcal{F})\ .$$

Denoting $\tau=\tau(\mathcal{F})$ and $\tau'=\tau\left(\mathcal{B}^{J_\mathcal{A}}\right)$, we have $\tau\left(\mathcal{A}^{J_\mathcal{B}}\right)\leq\tau-\tau'$. We also have:

$$\begin{array}{ll} c(J_\mathcal{A})\leq\alpha\tau(\mathcal{A})\leq\alpha\tau & c(J'_\mathcal{B})\leq\beta\tau\left(\mathcal{B}^{J_\mathcal{A}}\right)=\beta\tau' \\ c(J_\mathcal{B})\leq\beta\tau(\mathcal{B})\leq\beta\tau & c(J'_\mathcal{A})\leq\alpha\tau\left(\mathcal{A}^{J_\mathcal{B}}\right)\leq\alpha(\tau-\tau') \end{array}$$

Thus the cost of the edge set produced by the algorithm is bounded by

$$c(J)=\min\{c(J_\mathcal{A})+c(J'_\mathcal{B}),c(J_\mathcal{B})+c(J'_\mathcal{A})\}\leq\min\{\alpha\tau+\beta\tau',\beta\tau+\alpha(\tau-\tau')\}\ .$$

The worst case is when $\alpha\tau+\beta\tau'=\beta\tau+\alpha(\tau-\tau')$, namely $\tau'=\frac{\beta}{\alpha+\beta}\tau$. Then

$$c(J)=\alpha\tau+\beta\tau'=\tau\left(\alpha+\frac{\beta^2}{\alpha+\beta}\right)=\tau\frac{\alpha^2+\alpha\beta+\beta^2}{\alpha+\beta}=\tau\left(\alpha+\beta-\frac{\alpha\beta}{\alpha+\beta}\right)\ .$$

This concludes the proof of the lemma. □

Recall that $\mathbb{A}$ is a small biset/core if $|A|\leq q$, and $\mathbb{A}$ is a large biset/core otherwise. Let $\mathcal{F}$ be as in Theorem 5, namely, $\mathcal{F}$ and $\mathcal{F}^*$ are k-crossing, $|\partial\mathbb{A}|\geq k$ for all $\mathbb{A}\in\mathcal{F}$, and $n\geq k+3$. We show that then the following two subfamilies $\mathcal{A},\mathcal{B}$ of $\mathcal{F}$ satisfy the assumptions of Lemma 6 with $\alpha=\beta=1$; note that then $\alpha+\beta-\frac{\alpha\beta}{\alpha+\beta}=3/2$.

- $\mathcal{A}$ is the family of bisets in $\mathcal{F}$ that contain some small $\mathcal{F}$-core.
- $\mathcal{B}$ is the family of bisets in $\mathcal{F}$ that contain some large $\mathcal{F}$-core.

Lemma 7 *The families $\mathcal{A},\mathcal{B}$ above satisfy the assumption properties of Lemma 6 with $\alpha=\nu(\mathcal{F}_{\leq q})$ (so $\alpha=1$ if $\nu(\mathcal{F}_{\leq q})=1$) and $\beta=1$.*

Proof Clearly, $\mathcal{A} \cup \mathcal{B} = \mathcal{F}$. We prove that $\mathcal{A} \setminus \mathcal{B}$ and $\mathcal{B} \setminus \mathcal{A}$ are independent. Let $\mathbb{A} \in \mathcal{A}$ and $\mathbb{B} \in \mathcal{B}$ cross. Then $\mathbb{A} \cap \mathbb{B} \in \mathcal{F}$, since $\mathcal{F}$ is a crossing family. Thus $\mathbb{A} \cup \mathbb{B}$ contains an $\mathcal{F}$-core $\mathbb{C}$. If $\mathbb{C}$ is small then $\mathbb{A}, \mathbb{B} \in \mathcal{A}$ and thus $\mathbb{B} \notin \mathcal{B} \setminus \mathcal{A}$. If $\mathbb{C}$ is large then $\mathbb{A}, \mathbb{B} \in \mathcal{B}$ and thus $\mathbb{A} \notin \mathcal{A} \setminus \mathcal{B}$. In both cases we cannot have $\mathbb{A} \in \mathcal{A} \setminus \mathcal{B}$ and $\mathbb{B} \in \mathcal{B} \setminus \mathcal{A}$, hence $\mathcal{A} \setminus \mathcal{B}$ and $\mathcal{B} \setminus \mathcal{A}$ are independent.

To prove the claimed approximability of covering $\mathcal{A}$ and $\mathcal{B}$, we show that $\mathcal{A}^*$ is a union of $\nu(\mathcal{F}_{\leq q})$ intersecting biset families, and that $\mathcal{B}^*$ is an intersecting biset family. For a core $\mathbb{C}$ denote $\mathcal{F}_{\mathbb{C}} = \{\mathbb{A} \in \mathcal{F} : \mathbb{C} \subseteq \mathbb{A}\}$. Note that $\mathcal{A} = \bigcup \mathcal{F}_{\mathbb{C}}$ and $\mathcal{A}^* = \bigcup \mathcal{F}^*_{\mathbb{C}}$, where the union is taken over all small cores $\mathbb{C}$ of $\mathcal{F}$. It is easy to see that since $\mathcal{F}$ is crossing, then each family $\mathcal{F}^*_{\mathbb{C}}$ is an intersecting family. Hence $\mathcal{A}^*$ is a union of $\nu(\mathcal{F}_{\leq q})$ intersecting families

We prove that $\mathcal{B}^*$ is an intersecting family. Consider the inclusionwise maximal members of $\mathcal{B}^*$; each maximal member of $\mathcal{B}^*$ is the co-biset $\mathbb{C}^*$ of some large $\mathcal{F}$-core $\mathbb{C}$. We claim that if $\mathbb{C}_i, \mathbb{C}_j$ are distinct large $\mathcal{F}$-cores then $C^*_i \cap C^*_j = \emptyset$. Note that $|C_i|, |C_j| \geq q+1$, hence $|C^*_i|, |C^*_j| \leq n-k-q-1$. If $C^*_i \cap C^*_j \neq \emptyset$ then for $q \geq \frac{n-k-2}{2}$, and in particular for $q = \lfloor \frac{n-k+1}{2} \rfloor$, we have

$$|C^*_i \cup C^*_j| \leq |C^*_i| + |C^*_j| - 1 \leq 2n - 2k - 2q - 3 \leq n - k - 1$$

Since $\mathcal{F}^*$ is k-crossing, we get that $\mathbb{C}^*_i \cup \mathbb{C}^*_j \in \mathcal{F}^*$, contradicting the maximality of $\mathbb{C}^*_i, \mathbb{C}^*_j$. This implies that if $\mathbb{A}, \mathbb{B} \in \mathcal{B}^*$ intersect, then $\mathbb{A}, \mathbb{B}$ are contained in the same inclusionwise maximal member of $\mathcal{B}^*$, namely, $\mathbb{A}, \mathbb{B} \subseteq \mathbb{C}^*$ for some large $\mathcal{F}$-core $\mathbb{C}^*$. Note that $\mathbb{C} \subseteq \mathbb{A}^* \cap \mathbb{B}^*$. Thus if $\mathbb{A}, \mathbb{B}$ cross, and since $\mathcal{F}^*$ is a crossing family, $\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cup \mathbb{B} \in \mathcal{F}^*$. Moreover, $\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cup \mathbb{B} \subseteq \mathbb{C}^*$, which implies $\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cup \mathbb{B} \in \mathcal{B}^*$. Consequently, $\mathcal{B}^*$ is an intersecting family. $\square$

From Lemmas 6 and 7 we have:

Corollary 2 *Suppose that $\mathcal{F}$ is crossing, $\mathcal{F}^*$ is k-crossing, that $|\partial \mathbb{A}| \geq k$ for all $\mathbb{A} \in \mathcal{F}$, and that $n \geq k+3$ and $q = \lfloor \frac{n-k+1}{2} \rfloor$. Then directed* Biset-Family Edge-Cover *admits a polynomial time algorithm if $\nu(\mathcal{F}_{\leq q}) = 0$ and approximation ratio $3/2$ if $\nu(\mathcal{F}_{\leq q}) = 1$.*

Now we use Corollaries 1 and 2, and Theorem 6, to prove the directed part of Theorem 5. Note that the following algorithm uses all the assumptions on $\mathcal{F}$ in Theorem 5: $\mathcal{F}$ is k-crossing in Corollary 1, crossing in Theorem 6, and in Corollary 2 $\mathcal{F}$ is crossing, $\mathcal{F}^*$ is k-crossing, and $|\partial \mathbb{A}| \geq k$ for all $\mathbb{A} \in \mathcal{F}$.

Algorithm 3: DIRECTED-COVER$(\mathcal{F}, G, c)$

Using the algorithm from **Corollary 1** compute $J_1 \subseteq E$ such that $\nu(\mathcal{F}^{J_1}_{\leq q}) \leq \mu$ and $c(J_1) \leq \tau(\mathcal{F}_{\leq q})$

Using the algorithm from **Theorem 6** with $\mathcal{C} = \mathcal{C}(\mathcal{F}^{J_1}_{\leq q})$ and $t = 1$, compute $J_2 \subseteq E \setminus J_1$ such that $\nu\left(\mathcal{F}^{J_1 \cup J_2}_{\leq q}\right) \leq 1$ and $c(J_2) \leq (H(\mu) - 1)\tau\left(\mathcal{F}^{J_1}\right)$

Using the algorithm from **Corollary 2** compute an $\mathcal{F}^{J_1 \cup J_2}$-cover $J_3 \subseteq E \setminus (J_1 \cup J_2)$ such that $c(J_3) \leq \frac{3}{2}\tau(\mathcal{F}^{J_1 \cup J_2})$

return $J = J_1 \cup J_2 \cup J_3$

Tp summarize, the algorithm sequentially computes three edge sets:

1. J_1 reduces the number of small cores to μ by cost τ (Corollary 1).
2. J_2 further reduces the number of small cores to 1 by cost $(H(\mu) - H(1))\tau$ (Theorem 6).
3. J_3 covers the remaining members of $\mathcal{F}$ by cost $\frac{3}{2}\tau$ (Corollary 2).

Clearly, the algorithm computes a feasible solution. The approximation ratio is bounded by $1 + (H(\mu) - 1) + 3/2 = H(\mu) + 3/2$.

The proof of the directed part of Theorem 5 is complete.

4.2 Undirected graphs

To prove the undirected part of Theorem 5 we prove the following lemma.

Lemma 8 *Suppose that $\mathcal{F}$ is symmetric, k-crossing, that $|\partial\mathbb{A}| \geq k$ for all $\mathbb{A} \in \mathcal{F}$, and that $n \geq k+3$. Let $q = \lfloor \frac{n-k+1}{2} \rfloor$. Then undirected* Biset-Family Edge-Cover *admits a polynomial time algorithm if $\nu(\mathcal{F}_{\leq q}) = 1$, and LP-ratio 2 if $\nu(\mathcal{F}_{\leq q}) = 2$.*

Proof We claim that if $\nu(\mathcal{F}_{\leq q}) \leq 2$ then there exist a pair $s, t \in V$ such that

$$\nu\left(\mathcal{F}^{\{st\}}_{\leq q}\right) \leq \nu(\mathcal{F}_{\leq q}) - 1 \ . \tag{5}$$

Namely, adding the edge st reduces the number of small cores by at least 1. Note that such a pair s, t can be found in polynomial time by computing $\nu(\mathcal{F}_{\leq q})$ and $\nu\left(\mathcal{F}^{\{st\}}_{\leq q}\right)$ for every $s, t \in V$. Once such pair s, t is found, we compute a minimum cost cover J_{st} of the family $\{\mathcal{F}_{st} = \mathbb{A} \in \mathcal{F} : s \in A, t \in A^*\}$. This family is intersecting and has a unique core; such a family is sometimes called a **ring family**. Thus we get that in the case $\nu(\mathcal{F}_{\leq q}) \leq 2$, the problem of edge covering $\mathcal{F}$ is reduced to edge covering $\nu(\mathcal{F}_{\leq q})$ ring families. It is known that Biset-Family Edge-Cover with a ring family admits a polynomial time algorithm that computes a solution of cost $\tau(\mathcal{F})$. Consequently, we get a polynomial time algorithm if $\nu(\mathcal{F}_{\leq q}) = 1$ and ratio 2 if $\nu(\mathcal{F}_{\leq q}) = 2$.

We now prove existence of a pair s, t as above. Let $\mathbb{C} \in \mathcal{C}(\mathcal{F}_{\leq q})$ and let $\mathcal{M}_{\mathbb{C}}$ be the family of inclusionwise maximal bisets in $\mathcal{F}_{\leq q}$ that contain $\mathbb{C}$. If $\mathcal{M}_{\mathbb{C}}$ has a unique biset $\mathbb{A}_{\mathbb{C}}$, then (5) holds for any $s \in C$ and $t \in A^*_{\mathbb{C}}$. Suppose that $|\mathcal{M}_{\mathbb{C}}| \geq 2$. Note that by Lemma 4 and by the symmetry of $\mathcal{F}$, if $\mathbb{A}, \mathbb{B} \in \mathcal{F}_{\leq q}$ intersect, then $\mathbb{A} \cup \mathbb{B} \in \mathcal{F}_{\leq q}$ or $(\mathbb{A} \cup \mathbb{B})^* \in \mathcal{F}_{\leq q}$. Thus for any distinct $\mathbb{A}, \mathbb{B} \in \mathcal{M}_{\mathbb{C}}$, $(\mathbb{A} \cup \mathbb{B})^* \in \mathcal{F}_{\leq q}$ holds, by the maximality of the bisets in $\mathcal{M}_{\mathbb{C}}$. Consequently, since $\nu(\mathcal{F}_{\leq q}) \leq 2$, there is a unique $\mathcal{F}_{\leq q}$-core $\mathbb{C}'$ distinct from $\mathbb{C}$, such that $\mathbb{C}' \subseteq (\mathbb{A} \cup \mathbb{B})^*$ for any distinct $\mathbb{A}, \mathbb{B} \in \mathcal{M}_{\mathbb{C}}$. This implies that (5) holds for any $s \in C$ and $t \in C'$. □

Let us now show that Lemma 8 implies the undirected part of Theorem 5. The algorithm is similar to the one for the directed case; it returns a solution $J = J_1 \cup J_2 \cup J_3$ where:

1. J_1 reduces the number of small cores to μ by cost 2τ (Corollary 1).
2. If $\mu \geq 3$ then J_2 further reduces the number of small cores to 2 by cost $2(H(\mu) - H(2))\tau$ (Theorem 6).
3. J_3 covers the remaining members of $\mathcal{F}$ by cost τ if $\mu = 1$ and by cost 2τ otherwise (Lemma 8).

Clearly, the algorithm computes a feasible solution. In the case $\mu \geq 2$ the approximation ratio is $2 + 2(H(\mu) - H(2)) + 2 = 2H(\mu) + 1$. This is so also in the case $\mu = 1$, since then the ratio is $2 + 1 = 3 = 2H(1) + 1$.

This concludes the proof of the undirected part of Theorem 5.

5 Algorithm for 2-(T, s)-Connectivity Augmentation (Theorem 3)

Here we prove Theorem 3. We need some definitions.

Definition 8 Let us say that bisets $\mathbb{A}, \mathbb{B}$ **T-intersect** if $A \cap B \cap T \neq \emptyset$ and **T-co-cross** if both $A \cap B^* \cap T$ and $B \cap A^* \cap T$ are nonempty. A biset family $\mathcal{F}$ is **T-uncrossable** if for any $\mathbb{A}, \mathbb{B} \in \mathcal{F}$ the following holds: $\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cup \mathbb{B} \in \mathcal{F}$ if $\mathbb{A}, \mathbb{B}$ T-intersect, and $\mathbb{A} \setminus \mathbb{B}, \mathbb{B} \setminus \mathbb{A} \in \mathcal{F}$ if $\mathbb{A}, \mathbb{B}$ T-co-cross.

The following known lemma (c.f. [27]) can be easily deduced from Fact 1.

Lemma 9 ([27]) *Let G be an undirected k-(T, s)-connected graph. Then the biset family $\mathcal{F}_{k\text{-}(T,s)} = \{\mathbb{A} : |\partial \mathbb{A}| = k, \delta_E(\mathbb{A}) = \emptyset, A \cap T \neq \emptyset, s \in A^*\}$ defined in (3) is T-uncrossable.*

Halo families of a T-uncrossable family have the following property.

Lemma 10 ([27]) *Let $\mathcal{F}$ be an arbitrary T-uncrossable biset family and let $\mathbb{A}_i \in \mathcal{F}(\mathbb{C}_i)$ and $\mathbb{A}_j \in \mathcal{F}(\mathbb{C}_j)$, where $\mathbb{C}_i, \mathbb{C}_j \in \mathcal{C}(\mathcal{F})$.*

(i) *If $i = j$ (so $\mathbb{A}_i, \mathbb{A}_j$ contain the same $\mathcal{F}$-core) then $\mathbb{A}_i \cap \mathbb{A}_j, \mathbb{A}_i \cup \mathbb{A}_j \in \mathcal{F}(\mathbb{C}_i)$.*
(ii) *If $i \neq j$ and $\mathbb{A}_i, A_j$ T-co-cross then $\mathbb{A}_i \setminus \mathbb{A}_j \in \mathcal{F}(\mathbb{C}_i)$ and $\mathbb{A}_j \setminus \mathbb{A}_i \in \mathcal{F}(\mathbb{C}_j)$.*

A **simple biset family** $\mathcal{F}$ has no biset that contains 2 distinct cores, namely, $\mathcal{F}$ is the union of its halo families. The best known ratio for edge-covering an uncrossable biset family $\mathcal{F}$ is 2. Fukunaga [15] showed that for simple uncrossable biset families one can achieve ratio 4/3.

Definition 9 For $\mathcal{A} \subseteq \mathcal{F}$ and $U \subseteq V$ the **U-mesh graph** $\mathcal{G} = \mathcal{G}(\mathcal{A}, U)$ of $\mathcal{A}$ has node set $\mathcal{A}$ and edge set $\{\mathbb{A}_i\mathbb{A}_j : \partial \mathbb{A}_i \cap A_j \cap U \neq \emptyset$ or $\partial \mathbb{A}_j \cap A_i \cap U \neq \emptyset\}$.

Lemma 10(i) implies that if $\mathcal{F}$ is T-uncrossable, then for every $\mathbb{C}_i \in \mathcal{C}(\mathcal{F})$, the halo family of $\mathbb{C}_i$ has a unique maximal member (the union of the bisets in $\mathcal{F}(\mathbb{C}_i)$). The following statement easily follows from Lemma 10.

Corollary 3 ([27]) *Let $\mathcal{F}$ be an arbitrary T-uncrossable biset family and let $\mathcal{A}$ be the family of the maximal members of the halo families of the $\mathcal{F}$-cores. Let $\mathcal{A}'$ be an independent set in the T-mesh graph of $\mathcal{A}$. Then the union of the halo families of the bisets in $\mathcal{A}'$ is a simple uncrossable biset family.*

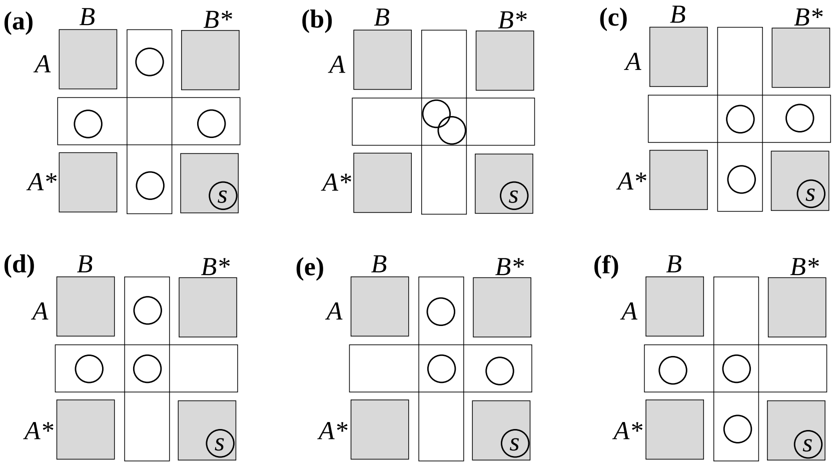

Fig. 1 Illustration to the proof of Lemma 11.

In the rest of this section let $G = (V, E)$ be a 2-(T, s)-connected graph and unless stated otherwise let $\mathcal{F} = \mathcal{F}_{2\text{-}(T,s)}$ be the biset family we want to cover. In the following lemma we summarize additional "uncrossing" properties of the bisets in $\mathcal{F}$ that we need.

Lemma 11 *Let $\mathbb{A}, \mathbb{B} \in \mathcal{F}$ such that $A \cap B \cap T = \emptyset$. Then either $\partial\mathbb{A} \cap B, \partial\mathbb{B} \cap A$ are both empty, or the following holds (see Fig. 1(a)):*

(i) *Each one of the sets $\partial\mathbb{A} \cap B, \partial\mathbb{A} \cap B^*, \partial\mathbb{B} \cap A, \partial\mathbb{B} \cap A^*$ is a singleton.*
(ii) *If $B \cap A^* \cap T \neq \emptyset$ then $\mathbb{B} \setminus \mathbb{A} \in \mathcal{F}$; if $A \cap B^* \cap T \neq \emptyset$ then $\mathbb{A} \setminus \mathbb{B} \in \mathcal{F}$.*
(iii) *If $|A \cap T| \geq 2$ and $|B \cap T| \geq 2$ then $\mathbb{A}, \mathbb{B}$ T-co-cross.*

Proof Fig. 1 depicts all possible cases of two bisets $\mathbb{A}, \mathbb{B}$ with $|\partial\mathbb{A}| = |\partial\mathbb{B}| = 2$. For part (i), we claim that if $A \cap B \cap T = \emptyset$ and if one of $\partial\mathbb{A} \cap B, \partial\mathbb{B} \cap A$ is non-empty, then the only possible case is the one depicted in (a). In cases (b,c) the sets $\partial\mathbb{A} \cap B, \partial\mathbb{B} \cap A$ are both empty. In the other cases (d,e,f) there is a biset $\mathbb{C}$ such that $C \cap T \neq \emptyset$ and $|\partial\mathbb{C}| = 1$, contradicting that G is 2-(T, s)-connected: $\mathbb{C} = \mathbb{A} \cup \mathbb{B}$ in case (d), $\mathbb{C} = \mathbb{B} \setminus \mathbb{A}$ in case (e), and $\mathbb{C} = \mathbb{A} \setminus \mathbb{B}$ in case (f).

For part (ii), assume that $B \cap A^* \cap T \neq \emptyset$; the proof of the case $A \cap B^* \cap T \neq \emptyset$ is similar. In the possible cases (a,b,c) we have $|\partial(\mathbb{B} \setminus \mathbb{A})| = 2$, so $\mathbb{B} \setminus \mathbb{A} \in \mathcal{F}$ in these cases, while the other case (d,e,f) are not possible.

Part (iii) is immediate from part (i). $\square$

Corollary 4 *Let $\mathbb{A}, \mathbb{B} \in \mathcal{C}(\mathcal{F})$. Then either $\mathbb{A} \subseteq \mathbb{B}^*$ and $\mathbb{B} \subseteq \mathbb{A}^*$, or each one of the sets $A \cap T, B \cap T$ is a singleton, and $\partial\mathbb{B} \cap A = A \cap T$ and $\partial\mathbb{A} \cap B = B \cap T$.*

Lemma 12 *Let $\mathcal{A} \subseteq \mathcal{F}$. If $A_i \cap A_j \cap T = \emptyset$ for any distinct $\mathbb{A}_i, \mathbb{A}_j \in \mathcal{A}$ then the V-mesh graph $\mathcal{G}$ of $\mathcal{A}$ is a forest.*

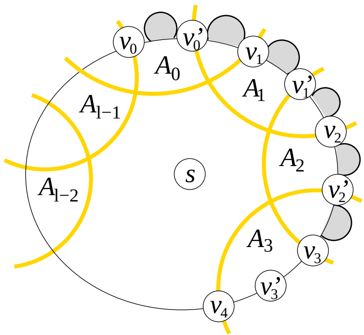

Fig. 2 Illustration to the proof of Lemma 12.

Proof Suppose to the contrary that $\mathcal{G}$ is not a forest. Let $(\mathbb{A}_0, \mathbb{A}_1, \ldots, \mathbb{A}_{l-1}, \mathbb{A}_0)$ be a cycle in $\mathcal{G}$. Assume that the indices are modulo l. By Lemma 11(i), for every i we have (see Fig. 2, and note that for any i, $v_i = v'_i$ may hold):

- $A_i \cap \partial\mathbb{A}_{i-1}$ is a singleton which we denote by v_i.
- $A_i \cap \partial\mathbb{A}_{i+1}$ is a singleton which we denote by v'_i.

Let $U = \bigcup_{i=1}^{l}\{v_i, v'_i\}$. For every $v \in U$, $v = v_i$ or $v = v'_i$ for some i, and we have: $v \in A_i$, $s \in A_i^*$, $\partial\mathbb{A}_i \subseteq U$, and $G \setminus \partial\mathbb{A}_i$ has no sv-path. We claim that there exists $u \in U$ such that G has no su-path. To see this, consider the shortest path P from s to U and the endnode u of P in U. Then P is an su-path that has no internal node in U. Since $\partial\mathbb{A}_u \subseteq U$, P is an su-path in $G \setminus \partial\mathbb{A}_u$. This contradicts the assumption that $G \setminus \partial\mathbb{A}_u$ has no st-path. On the other hand, G has an sv-path for every node v that belongs to the boundary of some tight biset, and thus G has an sv-path for every $v \in U$. This is a contradiction. □

Corollary 5 *Let $\mathcal{A}$ be obtained by picking for each core $\mathbb{C}_i \in \mathcal{C}(\mathcal{F})$ a biset $\mathbb{A}_i$ in the halo-family $\mathcal{F}(\mathbb{C}_i)$ of $\mathbb{C}_i$ (possibly $\mathbb{A}_i = \mathbb{C}_i$). Then the V-mesh graph of $\mathcal{A}$ is a forest. Furthermore, if $\mathbb{A}_i = \mathbb{C}_i$ for each i then the T-mesh graph $\mathcal{G}$ of $\mathcal{A}$ is a collection of node disjoint paths.*

Proof Since $\mathcal{F}$ is T-uncrossable, bisets from distinct halo families cannot T-intersect. Thus $A_i \cap A_j \cap T = \emptyset$ for distinct $\mathbb{A}_i, \mathbb{A}_j \in \mathcal{A}$, and the V-mesh graph of $\mathcal{A}$ is a forest by Lemma 12.

We prove that if $\mathbb{A}_i = \mathbb{C}_i$ for each i then $\mathcal{G}$ has no node of degree ≥ 3. Suppose to the contrary that $\mathcal{G}$ has a node $\mathbb{C}_0$ with 3 distinct neighbors $\mathbb{C}_1, \mathbb{C}_2, \mathbb{C}_3$. Then $C_i \cap C_j \cap T = \emptyset$ for distinct $0 \leq i, j \leq 3$. By Corollary 4 $C_i \cap T \subseteq \partial\mathbb{C}_0$ for $i = 1, 2, 3$, and we get the contradiction $|\partial\mathbb{C}_0| \geq 3$. □

Corollary 6 *Let $\mathcal{C}$ be the set family of the inner parts of the bisets in $\mathcal{C}(\mathcal{F})$. Then the maximum degree of a node in the hypergraph $(V, \mathcal{C})$ is at most 2.*

Proof Let $v \in V$ and let $\mathcal{C}_v = \{\mathbb{C} \in \mathcal{C}(\mathcal{F}) : v \in C\}$ be the family of cores whose inner part contains v. Consider the the V-mesh graph $\mathcal{G}_v$ of $\mathcal{C}_v$. By

Corollary 4 $\mathcal{G}_v$ is a clique, while by Corollary 5 $\mathcal{G}_v$ is a path. Thus $\mathcal{G}_v$ has at most 2 nodes. □

Now we prove the following.

Lemma 13 *If* Biset-Family Edge-Cover *admits approximation ratio α for simple uncrossable families and approximation ratio β for uncrossable families, then* 2-(T,s)-Connectivity Augmentation *admits approximation ratio $2\alpha+\beta$.*

Proof Let $\mathcal{F}=\mathcal{F}_{2\text{-}(T,s)}$. Let $\mathcal{A}$ be the family of the maximal members of the halo families of the $\mathcal{F}$-cores. Let $\mathcal{G}$ be the T-mesh graph of $\mathcal{A}$. By Lemma 12 $\mathcal{G}$ is a forest. Thus $\mathcal{G}$ is 2-colorable, so its nodes can be partitioned into 2 independent sets $\mathcal{A}'$ and $\mathcal{A}''$. The rest of the analysis coincides with [27]. Let $\mathcal{C}'$ and $\mathcal{C}''$ the set of $\mathcal{F}$-cores that correspond to $\mathcal{A}'$ and $\mathcal{A}''$, respectively. By Corollary 3, each one of the families $\mathcal{F}'=\bigcup_{\mathbb{C}\in\mathcal{C}'}\mathcal{F}(\mathbb{C})$ and $\mathcal{F}''=\bigcup_{\mathbb{C}\in\mathcal{C}''}\mathcal{F}(\mathbb{C})$ is uncrossable and simple. Thus the problem of covering $\mathcal{F}'\cup\mathcal{F}''$ admits ratio 2β. After the family $\mathcal{F}'\cup\mathcal{F}''$ is covered, the inner part of every core of the residual family contains at least 2 terminals. Hence by Lemma 11(iii), the residual family is uncrossable, and thus the problem of covering it admits ratio β. Consequently, the overall ratio is $2\alpha+\beta$, as claimed. □

As was mentioned, the currently best known values of α and β are $\alpha=4/3$ [15] and $\beta=2$ [10], so we get ratio $2\cdot 4/3+2=4\frac{2}{3}$.

Now let us consider the case when all edges in $\hat{E}$ are incident to s. Let $\mathcal{C}$ be the set family of the inner parts of the $\mathcal{F}$-cores. Recall that a hitting set of a hypergraph/set family is a set of nodes that intersects every hyperedge/set. Note that $J\subseteq\hat{E}$ is a feasible solution for our problem if and only if the set $\{v\in V : sv\in J\}$ is a hitting set of $\mathcal{C}$. Thus by assigning for every node v weight $w(v)=c(sv)$ (or a sufficiently large weight, if $sv\notin F$) we get that our problem is equivalent to finding a minimum-weight hitting set of the hypergraph $(V,\mathcal{C})$. By Corollary 6, the maximum degree in this hypergraph is ≤ 2. Finding a minimum-weight hitting set in hypergraph with maximum degree ≤ 2 can be done in polynomial time, as this is essentially an Edge-Cover problem. Consequently, we get a polynomial time algorithm for the case when all edges in $\hat{E}$ are incident to s, and the proof of Theorem 3 is complete.

6 A short proof of Theorem 7

Let $\mathcal{F}$ be a q-semi-intersecting biset family. Consider the dual program of the Biset-LP for covering $\mathcal{F}$

$$\max\left\{\sum_{\mathbb{A}\in\mathcal{F}} y_{\mathbb{A}} : \sum_{\delta_{\hat{E}}(\mathbb{A})\ni e} y_{\mathbb{A}}\leq c_e\ \forall e\in\hat{E},\ y_{\mathbb{A}}\geq 0\ \forall\mathbb{A}\in\mathcal{F}\right\}\ .$$

Given a dual solution y let us say that the dual constraint of an edge e is **tight**, or that e is a **tight edge** if $\sum_{\delta_{\hat{E}}(\mathbb{A})\ni e} y_{\mathbb{A}}=c_e$. Now consider the following primal-dual algorithm for covering $\mathcal{F}$.

Algorithm 4: q-Semi-Intersecting Family Edge-Cover$(\mathcal{F}, G, c)$

1 $J \leftarrow \emptyset$, $y \leftarrow 0$, $\mathcal{L} \leftarrow \emptyset$.
2 **while** $\nu(\mathcal{F}^J) \geq 1$ **do**
3   add some $\mathbb{C} \in \mathcal{C}(\mathcal{F}^J)$ to $\mathcal{L}$
4   raise $y_{\mathbb{C}}$ until the dual constraint of some $e \in \delta_{\hat{E} \setminus J}(\mathbb{C})$ becomes tight and add e to J
5 Let $e_1, \ldots, e_j$ be the order in which the edges were added to J
6 **for** $i = j$ **downto** 1 **do**
7   if $J \setminus \{e_i\}$ covers the family $\mathcal{F}' = \{\mathbb{A} \in \mathcal{F} : \mathbb{A} \subseteq \mathbb{B}$ for some $\mathbb{B} \in \mathcal{L}\}$ then do $J \leftarrow J \setminus \{e_i\}$
8 return J

Let I denote the set of edges in J right before the reverse-delete phase (steps 5,6,7). Note that I covers $\mathcal{F}$, but in the reverse-delete phase we care to cover just the subfamily $\mathcal{F}'$ of $\mathcal{F}$. In fact, the algorithm coincides with a standard primal-dual algorithm for covering the biset family $\mathcal{F}'$. We will show that $\mathcal{F}'$ is an intersecting biset family and conclude that $c(J) = \tau(\mathcal{F}') \leq \tau(\mathcal{F})$. In what follows, let $\mathcal{M}$ denote the family of inclusionwise maximal members of $\mathcal{L}$, and for an $\mathcal{F}^J$-core $\mathbb{C}_i$ let $\mathcal{M}_i$ denote the family of bisets in $\mathcal{M}$ that intersect with $\mathbb{C}_i$, and $\mathbb{B}_i$ the union of $\mathbb{C}_i$ and the bisets in $\mathcal{M}_i$.

Note that each family $\mathcal{M}_i$ is non-empty, since $\mathbb{C}_i$ is covered by some edge $e \in I \setminus J$, and since any edge $e \in I$ covers some $\mathbb{A} \in \mathcal{L}$. Let us say that a biset family $\mathcal{L}$ is **laminar** if for any $\mathbb{A}, \mathbb{B} \in \mathcal{L}$ that intersect $\mathbb{A} \subseteq \mathbb{B}$ or $\mathbb{B} \subseteq \mathbb{A}$ holds. In the following lemma we establish some properties of the families $\mathcal{L}$ and $\mathcal{F}'$.

Lemma 14 *At the end of the algorithm the following holds:*

(i) *$\mathcal{L}$ is a laminar biset family and $\mathcal{F}'$ is an intersecting biset family.*
(ii) *For any $\mathbb{A} \in \mathcal{M}$ there is a unique edge $e_{\mathbb{A}}$ in I that covers $\mathbb{A}$, and $e_{\mathbb{A}} \in J$. Furthermore, if $\mathbb{A}$ and an $\mathcal{F}^J$-core $\mathbb{C}$ intersect, then $\delta_J(\mathbb{A} \cap \mathbb{C}) = \{e_{\mathbb{A}}\}$.*

Proof We prove (i). Let $\mathbb{A}_1, \mathbb{A}_2 \in \mathcal{L}$ intersect where $\mathbb{A}_1$ was added to $\mathcal{L}$ before $\mathbb{A}_2$. When $\mathbb{A}_1$ was added to $\mathcal{L}$, we had $\mathbb{A}_1 \in \mathcal{C}(\mathcal{F}^J)$ and $\mathbb{A}_2 \in \mathcal{F}^J$. Thus $\mathbb{A}_1 \cap \mathbb{A}_2 = \mathbb{A}_1$ (namely, $\mathbb{A}_1 \subseteq \mathbb{A}_2$) by the minimality of $\mathbb{A}_1$ and since $\mathcal{F}$ (and thus also $\mathcal{F}^J$) is intersection closed. This implies that $\mathcal{L}$ is laminar. We show that $\mathcal{F}'$ is an intersecting biset family. Let $\mathbb{A}_1, \mathbb{A}_2 \in \mathcal{F}'$ intersect. Then, since $\mathcal{L}$ is laminar, $\mathbb{A}_1 \cup \mathbb{A}_2 \subseteq \mathbb{B}$ for some $\mathbb{B} \in \mathcal{L}$. Thus $\mathbb{A}_1 \cup \mathbb{A}_2 \in \mathcal{F}$, since $|A_1 \cup A_2| \leq |B| \leq q$ and since $\mathcal{F}$ is q-semi-intersecting. This implies $\mathbb{A}_1 \cup \mathbb{A}_2 \in \mathcal{F}'$, and clearly $\mathbb{A}_1 \cap \mathbb{A}_2 \in \mathcal{F}'$ since $\mathbb{A}_1 \cap \mathbb{A}_2 \subseteq \mathbb{B}$ and since $\mathcal{F}$ is intersection closed.

We prove (ii). Let $e_{\mathbb{A}}$ be the edge that was added to J at step 4 of the algorithm after $\mathbb{A}$ was added to $\mathcal{L}$ at step 3 (the first edge that covered $\mathbb{A}$). After $\mathbb{A}$ was added to $\mathcal{L}$, no biset that intersects with $\mathbb{A}$ was added to $\mathcal{L}$, since $\mathbb{A} \in \mathcal{M}$ and since $\mathcal{L}$ is laminar. Thus edges added to J after $e_{\mathbb{A}}$ do not cover $\mathbb{A}$, since their tails are in $V \setminus A$. Consequently, $e_{\mathbb{A}}$ is the unique edge in I that covers $\mathbb{A}$, and thus $e_{\mathbb{A}} \in J$. Now suppose that $\mathbb{A}$ and an $\mathcal{F}^J$-core $\mathbb{C}$ intersect. Then $\mathbb{A} \cap \mathbb{C} \in \mathcal{F}'$, since $\mathcal{F}$ is intersection closed and since $\mathbb{A} \cap \mathbb{C} \subseteq \mathbb{A}$. Thus

$\delta_J(\mathbb{A} \cap \mathbb{C}) \neq \emptyset$. Let $e \in \delta_J(\mathbb{A} \cap \mathbb{C})$. Then e covers $\mathbb{A}$, since e covers $\mathbb{A}$ or $\mathbb{C}$ by Fact 1, but e cannot cover C since $e \in J$ and J does not cover $\mathbb{C}$. Thus $e = e_{\mathbb{A}}$ for any $e \in \delta_J(\mathbb{A} \cap \mathbb{C})$, namely, $\delta_J(\mathbb{A} \cap \mathbb{C}) = \{e_{\mathbb{A}}\}$.

Lemma 15 *If $\mathcal{F}$ is q-semi-intersecting then at the end of the algorithm the following holds:*

(i) $|\delta_J(\mathbb{A})| = 1$ *for any* $\mathbb{A} \in \mathcal{L}$.
(ii) *The sets B_i are pairwise disjoint and each of them has size $\geq q + 1$.*

Proof For part (i), let $\mathbb{A} \in \mathcal{L}$ and suppose to the contrary that there are $e_1, e_2 \in \delta_J(\mathbb{A})$ with $e_1 \neq e_2$. For $i = 1, 2$ let $\mathbb{A}_i$ be some biset in $\mathcal{F}'$ that became uncovered when e_i was considered for deletion at step 7. Note that $\delta_J(\mathbb{A}_i) = \{e_i\}$ and that $\mathbb{A} \subseteq \mathbb{A}_i$, since the edges in J were considered for deletion in the reverse order. Thus $\mathbb{A} \subseteq \mathbb{A}_1 \cap \mathbb{A}_2$, and by Lemma 14(i) $\mathbb{A}_1 \cup \mathbb{A}_2 \in \mathcal{F}'$. Consequently, there is $e \in \delta_J(\mathbb{A}_1 \cup \mathbb{A}_2)$, hence $e \in \delta_J(\mathbb{A}_1)$ or $e \in \delta_J(\mathbb{A}_2)$, by Fact 1. Thus $e = e_1$ or $e = e_2$. Since the tail of each of e_1, e_2 is in $A \subseteq A_1 \cap A_2$, so is the tail of e. The head of e is in $A_1^* \cap A_2^*$. This gives the contradiction $e \in \delta_J(\mathbb{A}_1) \cap \delta_J(\mathbb{A}_2)$.

We prove part (ii). Let $\mathbb{C}_i, \mathbb{C}_j$ be distinct $\mathcal{F}^J$-cores. Note that no two bisets in $\mathcal{M}$ intersect (since $\mathcal{L}$ is laminar) and that $C_i \cap C_j = \emptyset$ (since $\mathcal{F}$ is intersection closed). Thus to prove that $B_i \cap B_j = \emptyset$ it is sufficient to prove that $\mathcal{M}_i \cap \mathcal{M}_j = \emptyset$. Suppose to the contrary that there is $\mathbb{A} \in \mathcal{M}_i \cap \mathcal{M}_j$. By Lemma 14(ii), the tail of $e_{\mathbb{A}}$ is both in $A \cap C_i$ and $A \cap C_j$. This contradicts $C_i \cap C_j = \emptyset$. We prove that $|B_i| \geq q + 1$. Note that $|B_i| \leq q$ implies $\mathbb{B}_i \in \mathcal{F}$, since $\mathcal{F}$ is q-semi-intersecting. Thus to prove that $|B_i| \geq q + 1$ it is sufficient to prove that $\delta_I(\mathbb{B}_i) = \emptyset$, since this implies $\mathbb{B}_i \notin \mathcal{F}$ (as I covers $\mathcal{F}$). Suppose to the contrary that there is $e \in \delta_I(\mathbb{B}_i)$. Then there is a biset $\mathbb{A} \in \mathcal{M}$ whose inner part contains the tail of e, and we must have $\mathbb{A} \in \mathcal{M}_i$, by the definition of $\mathbb{B}_i$ and since no two bisets in $\mathcal{M}$ intersect. As e covers the biset $\mathbb{B}_i$ that contains $\mathbb{A}$, e covers $\mathbb{A}$, and thus $e = e_{\mathbb{A}}$ and $\delta_J(\mathbb{A} \cap \mathbb{C}_i) = \{e_{\mathbb{A}}\}$, by Lemma 14(ii). The edge $e_{\mathbb{A}}$ has its tail in C_i and covers the biset $\mathbb{B}_i$ that contains $\mathbb{C}_i$. Consequently, $e_{\mathbb{A}}$ covers $\mathbb{C}_i$, contradicting that $\mathbb{C}_i \in \mathcal{F}^J$.

Lemma 15(ii) implies $\nu(\mathcal{F}^J) \leq \lfloor n/(q+1) \rfloor$. To see that $c(J) = \tau(\mathcal{F}')$ let $x \in \{0,1\}^F$ be the characteristic vector of J and y the dual solution produced by the algorithm. It is easy to see that x and y are feasible solutions for the primal and dual LPs, respectively, and that the Primal Complementary Slackness Conditions hold for x and y. The Dual Complementary Slackness Conditions are: $y_{\mathbb{A}} > 0$ implies $|\delta_J(\mathbb{A})| = 1$, and they hold by Lemma 15(i), since $\{\mathbb{A} : y_{\mathbb{A}} > 0\} \subseteq \mathcal{L}$.

This concludes the proof of Theorem 7.

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. V. Auletta, Y. Dinitz, Z. Nutov, and D. Parente. A 2-approximation algorithm for finding an optimum 3-vertex-connected spanning subgraph. *J. of Algorithms*, 32(1):21–30, 1999.
3. J. Cheriyan and B. Laekhanukit. Approximation algorithms for minimum-cost k-(S,T) connected digraphs. Manuscript, 2010.
4. J. Cheriyan, B. Laekhanukit, G. Naves, and A. Vetta. Approximating rooted steiner networks. *ACM Transactions on Algorithms*, 11(2):8:1–8:22, 2014.
5. J. Cheriyan and L. Végh. Approximating minimum-cost k-node connected subgraphs via independence-free graphs. *SIAM J. Computing*, 43(4):1342–1362, 2014.
6. J. Cheriyan, S. Vempala, and A. Vetta. An approximation algorithm for the min-cost k-vertex connected subgraph. *SIAM J. Comput.*, 32(4):1050–1055, 2003.
7. J. Chuzhoy and S. Khanna. An $O(k^3 \log n)$-approximation algorithm for vertex-connectivity survivable network design. *Theory of Computing*, 8(1):401–413, 2012.
8. Y. Dinitz and Z. Nutov. A 3-approximation algorithm for finding optimum $4,5$-vertex-connected spanning subgraphs. *J. of Algorithms*, 32(1):31–40, 1999.
9. J. Fackharoenphol and B. Laekhanukit. An $O(\log^2 k)$-approximation algorithm for the k-vertex connected subgraph problem. *SIAM J. Comput.*, 41:1095–1109, 2012.
10. L. Fleischer, K. Jain, and D. Williamson. Iterative rounding 2-approximation algorithms for minimum-cost vertex connectivity problems. *J. Computer and System Sciences*, 72(5):838–867, 2006.
11. A. Frank. Rooted k-connections in digraphs. *Discrete Applied Math.*, 157(6):1242–1254, 2009.
12. A. Frank and T. Jordán. Minimal edge-coverings of pairs of sets. *J. of Comb. Theory B*, 65:73–110, 1995.
13. A. Frank and E. Tardos. An application of submodular flows. *Linear Algebra and its Applications*, 114/115:329–348, 1989.
14. T. Fukunaga. Approximating minimum cost source location problems with local vertex-connectivity demands. *J. Discrete Algorithms*, 19:30–38, 2013.
15. T. Fukunaga. Approximating the generalized terminal backup problem via half-integral multiflow relaxation. In *STACS*, pages 316–328, 2015.
16. T. Fukunaga, Z. Nutov, and R. Ravi. Iterative rounding approximation algorithms for degree-bounded node-connectivity network design. *SIAM J. Computing*, 44(5):1202–1229, 2015.
17. M. Goemans, A. Goldberg, S. Plotkin, D. Shmoys, E. Tardos, and D. Williamson. Improved approximation algorithms for network design problems. In *SODA*, pages 223–232, 1994.
18. K. Jain. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica*, 21(1):39–60, 2001.
19. T. Jordán. On the optimal vertex-connectivity augmentation. *J. on Comb. Theory B*, 63:8–20, 1995.
20. S. Khuller and B. Raghavachari. Improved approximation algorithms for uniform connectivity problems. *J. of Algorithms*, 21:434–450, 1996.
21. G. Kortsarz and Z. Nutov. Approximating node connectivity problems via set covers. *Algorithmica*, 37:75–92, 2003.
22. G. Kortsarz and Z. Nutov. Approximating k-node connected subgraphs via critical graphs. *SIAM J. on Computing*, 35(1):247–257, 2005.
23. G. Kortsarza and Z. Nutov. Approximating source location and star survivable network problems. Manuscript, 2015.
24. B. Laekhanukit. Parameters of two-prover-one-round game and the hardness of connectivity problems. In *SODA*, pages 1626–1643, 2014.
25. Y. Lando and Z. Nutov. Inapproximability of survivable networks. *Theortical Computer Science*, 410(21-23):2122–2125, 2009.
26. Z. Nutov. Approximating subset k-connectivity problems. In *WAOA*, pages 9–20, 2011.
27. Z. Nutov. Approximating minimum cost connectivity problems via uncrossable bifamilies. *ACM Transactions on Algorithms*, 9(1):1, 2012.
28. Z. Nutov. Approximating node-connectivity augmentation problems. *Algorithmica*, 63(1-2):398–410, 2012.
29. Z. Nutov. Approximating minimum-cost edge-covers of crossing biset families. *Combinatorica*, 34(1):95–114, 2014.

30. R. Ravi and D. P. Williamson. An approximation algorithm for minimum-cost vertex-connectivity problems. *Algorithmica*, 18:21–43, 1997.
31. R. Ravi and D. P. Williamson. Erratum: an approximation algorithm for minimum-cost vertex-connectivity problems. *Algorithmica*, 34(1):98–107, 2002.
32. L. Végh. Augmenting undirected node-connectivity by one. *SIAM J. Discrete Math.*, 25(2):695–718, 2011.