Node-Connectivity Survivable Network Problems

Zeev Nutov

The Open University of Israel

E-mail: nutov@openu.ac.il

1.1 Introduction

We survey approximation algorithms and hardness of approximation results for the Survivable Network problem in which we seek a low edge-cost directed/undirected subgraph that satisfies prescribed connectivity demands w.r.t. a given "connectivity measure". These problems include the following well known problems: Minimum Spanning Tree, Traveling Salesman Problem, Steiner Tree, Steiner Forest, and their directed variants. These are examples of *low connectivity* Survivable Network problems. In this survey we will consider *high connectivity* Survivable Network problems; some examples are Min-Cost k-Flow, k-Inconnected Subgraph, k-Connected Subgraph, and Rooted Survivable Network. See previous surveys on such problem in [30, 26].

Many common connectivity measures can be defined by the following unified framework. Let $G = (V, E)$ be a (possibly directed) graph. For $Q \subseteq V$, the **Q-connectivity** $\lambda_G^Q(s,t)$ of a node pair (s,t) is the maximum number of st-paths such that no two of them have an edge or a node in $Q \setminus \{s,t\}$ in common. The case $Q = \emptyset$ is the case of **edge-connectivity**, and we use the notation $\lambda_G(s,t) := \lambda_G^{\emptyset}(s,t)$; the case $Q = V$ is the case of **node-connectivity**, and we use the notation $\kappa_G(s,t) := \lambda_G^V(s,t)$. Even more generally, given node capacities $\{q_v : v \in V\}$, the **q-connectivity** $\lambda_G^q(s,t)$ is the maximum number of pairwise edge disjoint st-paths such that for every $v \in V \setminus \{s,t\}$ at most q_v of the paths contain v; Q-connectivity is the particular case when $q_v \in \{1, \infty\}$ for all $v \in V$ and $Q = \{v \in V : q_v = 1\}$. We will consider mainly the node-connectivity case $Q = V$, when the paths are required to be pairwise internally node disjoint. However, most algorithms presented can be adjusted to the q-connectivity case with the same approximation ratio.

Given positive integral connectivity **demands** $\{r_{st} \geq 1 : st \in D\}$ over a set D of demand pairs on V and $Q \subseteq V$ we say that G **satisfies** r if $\lambda_G^Q(s,t) \geq r_{st}$ for all $st \in D$. Our problem is:

<u>Survivable Network</u>

Input: A (multi)graph $\hat{G} = (V, \hat{E})$ with edge costs $\{c_e : e \in \hat{E}\}$, $Q \subseteq V$, and connectivity demands $\{r_{st} > 0 : st \in D\}$ on a set $D \subseteq V \times V$ of demand pairs.

Output: A minimum cost subgraph of $\hat{G}$ that satisfies r.

Assume that the input numbers are integers bounded by a polynomial in $n = |V|$. A node is a **terminal** if it belongs to some demand pair. Let T denote the set of terminals. An important Q-connectivity measure is **element-connectivity**, when $Q \subseteq V \setminus T$. Let k-Survivable Network be the restriction of Survivable Network to instances with demands at most k, and unless stated otherwise $k = \max_{st \in D} r_{st}$ is the maximum demand.

We may assume that the input graph $\hat{G}$ is complete, by assigning infinite costs to "forbidden" edges. Under this assumption, we consider four types of edge costs:

- **$\{0,1\}$-costs:** Here we are given a graph G, and the goal is to find a minimum size augmenting edge set J of new edges (any edge is allowed, and parallel edges are allowed) such that $G \cup J$ satisfies r.
- **$\{1,\infty\}$-costs:** Here we seek a subgraph of $\hat{G}$ with minimum number of edges that satisfies r.
- **Metric Costs:** The costs satisfy the triangle inequality $c_{uv} \leq c_{uw} + c_{wv}$ for all $u, v, w \in V$.
- **General Costs:** The costs are arbitrary non-negative integers or ∞.

For each type of costs, we have the following types of demands:

- **Uniform demands:** $r_{st} = k$ for all $s, t \in V$; this is the k-Connected Subgraph problem.
- **Rooted demands:** All pairs in D share the same node s; this is the Rooted Survivable Network problem. If the demands are $r_{ts} = k$ for all $t \in T \setminus \{s\}$ then we get **rooted subset uniform demands** and the Subset k-Inconnected Subgraph problem.
- **Subset uniform demands:** $r_{st} = k$ for all $s, t \in T$; this is the Subset k-Connected Subgraph problem.
- **Arbitrary (non-negative) demands**.

We mention two important cases of Rooted Survivable Network. In the Min-Cost k-Flow problem $|D| = 1$. This problem admits a polynomial time algorithm for both directed and undirected graphs, c.f. [15, 44]. Another case is the k-Inconnected Subgraph problem where the demands are $r_{ts} = k$ for every $t \in V \setminus \{s\}$; this is a particular case of the Subset k-Inconnected Subgraph problem, when $T = V$. For directed graphs, k-Inconnected Subgraph can be solved in polynomial time [18] (see also [14]).

Survivable Network admits a trivial approximation ratio $\min\{|D|, O(n)\}$. Ratio $|D|$ is obtained by applying the above Min-Cost k-Flow algorithm for each pair in D and taking the union of the $|D|$ computed edge sets. Ratio $O(n)$ is obtained by randomized rounding of an LP-relaxation, see Section 1.2.

Problem	**Demands**	**Approximability**
directed/undirected Shortest Path (Min-Cost 1-Flow)	rooted, $\|D\| = 1$	in P
Minimum Spanning Tree (undirected 1-Connected Subgraph)	uniform/rooted, $T = V$	in P
Steiner Tree (undirected Subset 1-Connected Subgraph)	rooted/subset uniform	$\ln 4 + \varepsilon$ [3] ([21])
Steiner Forest (undirected 1-Survivable Network)	general	2 [1] ([22])
Minimum Arborescence (directed 1-Inconnected Subgraph)	rooted, $T = V$	in P
Strongly Connected Subgraph (directed 1-Connected Subgraph)	uniform	2 for general costs 3/2 for $\{1, \infty\}$-costs [45]
Directed Steiner Tree (directed Rooted 1-Survivable Network)	rooted uniform	$O\left(\ell^3 \|T\|^{2/\ell}\right)$ in $O\left(\|T\|^{2\ell} n^\ell\right)$ time [47, 33, 4, 24], $\Omega(\ln^2 n)$ [23]
Directed Steiner Forest (directed 1-Survivable Network)	general	$n^\varepsilon \cdot \min\left\{\|D\|^{1/2}, n^{2/3}\right\}$ [5, 2] ([11]) Φ [10]

Table 1.1: Known approximability status of 1-Survivable Network problems. References in parenthesis give a simplified proof, or a slight improvement needed to achieve the approximation ratio/threshold stated.

In Table 1 we list famous particular cases of 1-Survivable Network. For a survey on the k-Connected Subgraph problem (the case of uniform demands) see [36]. For the metric costs case see the original paper of Cheriyan and Vetta [8]. Here we consider all the other variants of the problem. We focus on *node-connectivity* and k *arbitrary*, although there are many interesting results for edge-connectivity and element-connectivity, as well as for small requirements. We survey only *approximation algorithms* (for polynomially solvable cases see [15, 17]), with the currently best known approximation ratios, as summarized in Table 2, where, for comparison, we also include results for edge-connectivity.

c*,*r	**Node-Connectivity**		**Edge-Connectivity**	
	Undirected	*Directed*	*Undirected*	*Directed*
$\{0,1\}$,R,RU	$O(\min\{\ln^2 k, \ln n\})$ [31] $\Omega(\ln n)$ [37]	$O(\ln n)$ [31, 37] $\Omega(\ln n)$ [13]	in P [13]	$O(\ln n)$ [31] $\Omega(\ln n)$ [13]
$\{0,1\}$,SU	$\frac{\lvert T\rvert \ln k}{\lvert T\rvert - k} O(\min\{\ln^2 k, \ln n\})$ [42] Φ [29, 38]	$\frac{\lvert T\rvert \ln k}{\lvert T\rvert - k} \cdot O(\ln n)$ [42] Φ [29, 38]	in P [13]	$O(\ln n)$ [31] $\Omega(\ln n)$ [13]
$\{0,1\}$,G	$k \cdot O(\min\{\ln^2 k, \ln n\})$ [41, 31] Φ [29, 38]	$O(k \ln n)$ [31] Φ [29, 38]	in P [13]	$O(\ln n)$ [31] $\Omega(\ln n)$ [13]
$\{1,\infty\}$,R	$O(k^2)$ [40], Φ [7]	Φ [7]	2 [25]	Φ [7]
$\{1,\infty\}$,RU	$O(k \ln k)$ [40], Φ [7]	Φ [7]	2 [25]	Φ [7]
$\{1,\infty\}$,SU	$\frac{\lvert T\rvert}{\lvert T\rvert - k} O(k \ln k)$ [42], Φ [29]	Φ [10]	2 [25]	Φ [10]
$\{1,\infty\}$,G	$O(k^3 \ln \lvert T\rvert)$ [9], Φ [29]	Φ [10]	2 [25]	Φ [10]
MC,R	$O(\ln k)$ [8]	Φ [7]	2 [25]	Φ [7]
MC,RU,SU	24 [8]	Φ [7]	2 [25]	Φ [7]
MC,G	$O(\ln k)$ [8]	Φ [10]	2 [25]	Φ [10]
GC,R	$O(k^2)$ [40], Φ [7]	Φ [7]	2 [25]	Φ [7]
GC,RU	$O(k \ln k)$ [40], Φ [7]	Φ [7]	2 [25]	Φ [7]
GC,SU	$\frac{\lvert T\rvert}{\lvert T\rvert - k} O(k \ln k)$ [42], Φ [29]	Φ [10]	2 [25]	Φ [10]
GC,G	$O(k^3 \ln \lvert T\rvert)$ [9], Φ [29]	Φ [10]	2 [25]	Φ [10]

Table 2.2: Known approximability status of Survivable Network problems. MC and GC stand for metric and general costs, R, RU, SU, and G stand for rooted, rooted uniform, subset uniform, and general demands, respectively. All problems admit ratio $\min\{|D|, O(n)\}$. Here $\Phi = 2^{\ln^{1-\epsilon} n}$, meaning that the problem cannot be approximated within $2^{\ln^{1-\varepsilon} n}$ for any fixed $\epsilon > 0$ (unless NP has quasi-polynomial time algorithms).

The approximation lower bounds results in the tables are given in terms of n. In terms of the parameters k and $|D|$, the following lower bounds are established in [34] for rooted demands: $k^{1/2-\epsilon}$ for directed graphs, $k^{1/10-\epsilon}$ for undirected graphs, and $|D|^{1/4-\epsilon}$ for both directed and undirected graphs. For undirected graphs and arbitrary demands [34] shows a better bound $k^{1/6-\epsilon}$.

The reader may observe that most of the results for the edge-connectivity case follow from 4 basic algorithms. In the undirected case we have ratio 2 for almost all variants by the seminal work of Jain [25], except that in the case of $\{0,1\}$-costs we have a polynomial time algorithm by the work of Frank [13]. In the directed case, we have the trivial ratio $\min\{|D|, O(n)\}$ for almost all variants, except that in the case of $\{0,1\}$-costs ratio $O(\ln n)$ is achievable using Set-Cover techniques [31]. Unfortunately, node-connectivity Survivable Network problems are more complicated, and often one has to decompose the problem into a "small" number of subproblems on which traditional methods can be applied.

We mention some additional results not appearing in the tables. Survivable Network admits ratio 2 for $k = 2$ [12], but no constant ratio is known for $k = 3$. Element-connectivity Survivable Network admits ratio 2 [12]; in the case of $\{0,1\}$-costs the problem is NP-hard for $k = 2$ [27], and admits ratios 7/4 for arbitrary demands and 3/2 for $\{0,1,2\}$ or $\{0,k\}$ demands [38]. When the sum of the demands is a constant, directed Rooted Survivable Network admits a polynomial time algorithm [7]. In the (undirected) Multiroot Connectivity problem the demands are $\max\{r_s, r_t\}$ for given non-negative integers $\{r_v : v \in V\}$. The best known ratios for this problem are: $2(k-1)$ for general costs, 3 for metric costs, and 2 for $\{1,\infty\}$-costs [6].

Here is some notation used. An edge from u to v is denoted by uv. A uv-**path** is a path from u to v. For arbitrary sets A, B of nodes and edges (or graphs) $A \setminus B$ is the set (or graph) obtained by deleting B from A, where deletion of a node implies also deletion of all the edges incident to it; similarly, $A \cup B$ is the set (graph) obtained by adding B to A. For real values $\{x_u : u \in U\}$ let $x(U) = \sum_{u \in U} x_u$ and $\max(U) = \max_{u \in U} x_u$.

Organization. In the next section 1.2 we give some reductions between various versions of our problem. In Section 1.3 we describe an $O(k^3 \ln n)$-approximation algorithm for undirected Survivable Network. In Section 1.4 we give a Biset-LP formulation of Survivable Network problems, and some properties of relevant biset families. In Section 1.5 we consider rooted and subset uniform demands, while in Section 1.6 we consider the $\{0,1\}$-costs case. We conclude in Section 1.7 with some open problems.

1.2 Some simple reductions

Note that for directed graphs, the node-connectivity and the edge-connectivity cases have often the same approximability in Table 2. The following statement shows that this is not a coincidence.

Lemma 1.1 *For directed graphs, if edge-connectivity* Survivable Network *admits ratio $\rho(n)$ then q-connectivity* Survivable Network *admits ratio $\rho(2n)$.*

Proof: Given a q-connectivity instance, obtain an edge-connectivity instance by a standard conversion of node capacities to edge-capacities: Replace every $v \in V$ by two nodes v^{out}, v^{in} connected by q_v edges from v^{in} to v^{out} of cost 0 each, and redirect the heads of the edges in $\hat{E}$ entering v to v^{in} and the tails of the edges in $\hat{E}$ leaving v to v^{out}, keeping costs. The demands are $r'(s^{out}, t^{in}) = r(s,t)$. Then J is a feasible solution to the q-connectivity instance iff $J' = \{u^{out}v^{in} : uv \in E\} \cup \{v^{in}v^{out} : v \in V\}$ is a feasible solution to the edge-connectivity instance, and clearly both solutions have the same cost. □

As an application of Lemma 1.1 we show ratio $O(n)$ for edge-connectivity Survivable Network, and via Lemma 1.1 obtain the same ratio for q-connectivity. For $A \subseteq V$ let $f_r(A) = \max\{r_{st} : s \in A, t \in V \setminus A\}$, where the maximum or a sum taken over the empty set is defined to be 0. Let $\delta(A)$ denote the set of edges in $\hat{E}$ going from A to $V \setminus A$. Consider the **Cut-LP** for directed edge-connectivity Survivable Network

$$\min\left\{\sum_{e\in\hat{E}} c_e x_e : \sum_{e\in\delta(A)} x_e \geq f_r(A)\ \forall A \subseteq V,\ 0 \leq x_e \leq 1\ \forall e \in E\right\}$$

This LP can be solved in polynomial time by the ellipsoid method. As the number of constraints in the LP is 2^n, we can get (expected) ratio $O(\ln 2^n) = O(n)$ via randomized rounding; the computed solution will be feasible with high probability.

Note that in Table 2 the $\{1,\infty\}$-costs case has the same approximability as the general costs case. This is also not a coincidence, since for problems we consider, the general costs case can be reduced to the $\{1,\infty\}$-costs case with a small loss in the approximation ratio, as follows. Assume that the set of 0 cost edges is not a feasible solution. Multiply the costs by $M = n^2/\varepsilon$, then assign cost 1 to 0 cost edges, and replace every edge $e = uv$ by a uv-path of length c_e of 1 cost edges. Any feasible solution of cost C is transformed into a solution of cost between MC and $MC + n^2 = M(C + \varepsilon)$, which causes arbitrarily small loss in the ratio.

We now show that for high values of k, the directed and undirected variants of Survivable Network have similar approximability. Let $\rho_{\mathsf{dir}}(k,n)$ and $\rho_{\mathsf{und}}(k,n)$ denote the best possible approximation ratio for the directed and undirected k-Survivable Network problem on graphs on n nodes, respectively.

Lemma 1.2 $\rho_{\mathsf{und}}(k,n) \le 2\rho_{\mathsf{dir}}(k,n)$.

Proof: Given an undirected Survivable Network instance $\mathcal{I} = (\hat{G}, Q, c, r)$, obtain a "bidirected" instance $\mathcal{I}' = (\hat{G}', Q, c', r)$ of directed Survivable Network by replacing every edge uv of $\hat{G}$ by the two opposite directed edges uv, vu each of the same cost as uv. Then compute a ρ approximate solution J' to $\mathcal{I}'$ and output its underlying graph J. It is easy to see that J is a feasible solution for $\mathcal{I}$. Furthermore, if G is an arbitrary subgraph of $\hat{G}$, and G' is the corresponding bidected subgraph of $\hat{G}'$, then $c'(G') = 2c(G)$, and G is a feasible solution for $\mathcal{I}$ if G' is a feasible solution for $\mathcal{I}'$. Thus $\mathsf{opt}' \le 2\mathsf{opt}$, where opt and opt' denote the optimal solution value of $\mathcal{I}$ and $\mathcal{I}'$, respectively. Consequently, $c(J) \le c'(J') \le \rho \cdot \mathsf{opt}' \le 2\rho \cdot \mathsf{opt}$. □

In the reduction in the proof of Lemma 1.2 it is also possible to "bidirect" the demands.

A pair (A, B) of disjoint subsets of V is called a **setpair**; (A, B) is an st**-setpair** if $s \in A$ and $t \in B$. For a graph $G = (V, E)$ let $d_G(A, B) = d_E(A, B)$ denote the number of edges in G going from A to B. Recall that $\kappa_G(s,t)$ denotes the maximum number of pairwise internally disjoint st-paths in G. If (A, B) is an st-setpair then $\kappa_G(s,t) \le d_G(A,B) + |V \setminus (A \cup B)|$, and we say that (A, B) is st**-tight** in G if $\kappa_G(s,t) = d_G(A,B) + |V \setminus (A \cup B)|$. The node-connectivity version of Menger's Theorem states that for any $s, t \in V$ an st-tight setpair exists, namely

$$\kappa_G(s,t) = \min\{d_G(A,B) + n - |A \cup B| : (A,B) \text{ is an } st\text{-setpair in } G\}$$

Theorem 1.1 (Lando & Nutov [35]) $\rho_{\mathsf{dir}}(k,n) \le \rho_{\mathsf{und}}(k+n, 2n)$*, and this is so also if both directed and undirected problems are restricted to rooted demands, or to rooted uniform demands, or to* $\{0,1\}$*-costs.*

In the rest of this section we prove Theorem 1.1. Let $G = (V, E)$ be a directed graph. The **bipartite graph of** G has node set $V' \cup V''$ where V', V'' are copies of V, and edge set $\{u'v'' : uv \in E\}$, where for $v \in V$ we denote by v' and v'' the copies of v in V' and V'', respectively. The **padded graph of** G is obtained by adding to the bipartite graph of G a **padding edge set** of cliques on each of V' and V'', and the matching $M = \{v'v'' : v \in V\}$ (see Fig. 1.1). Given an instance $\hat{G}, c, r$ of directed k-Survivable Network,

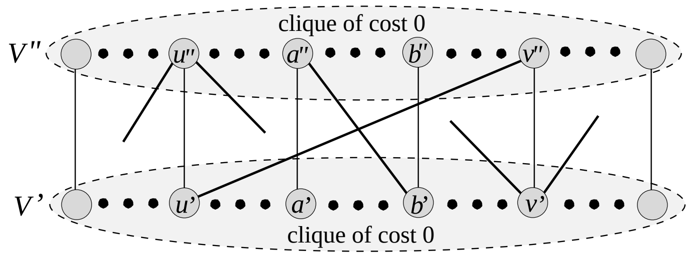

Figure 1.1: The padded bipartite graph; edges in M are shown by thin lines.

obtain an instance of undirected $(k+n)$-Survivable Network $\hat{H}, c', r'$, where $\hat{H}$ is the padded graph of $\hat{G}$. The demands and the costs are derived via the natural bijection between ordered node pairs (u,v) of $\hat{G}$ and (u',v'') of $\hat{H}$, while keeping costs but shifting the demands by n; namely, the demand of the pair (s',t'') is $r_{st}+n$ and the costs are: $c'(u'v'')=c(uv)$ whenever $uv \in \hat{E}$ and the padding edges have cost 0. Note that if G is a spanning subgraph of $\hat{G}$ then the padded graph H of G is a spanning subgraph of $\hat{H}$, and G and H have the same cost. Thus to finish the proof of Theorem 1.1 it is sufficient to prove the following.

Lemma 1.3 ([35]) *If H is the padded graph of a directed graph $G=(V,E)$ on n nodes, then $\kappa_H(s',t'') = \kappa_G(s,t)+n$ for all $s,t \in V$.*

Proof: Let (A,B) be an st-tight setpair in G. Then (A',A'') is an $s't''$-setpair in H, where A' is the copy of A in V' and B'' is the copy of B in V''. Note that $d_G(A,B)=d_H(A',B'')$. This implies $\kappa_H(s',t'') \leq d_H(A',B'')+2n-|A' \cup B''| = \kappa_G(s,t)+n$. We prove the inverse inequality. Let $(A' \cup A'', B' \cup B'')$ be an $s't''$-tight setpair in H with $|A'|+|A''|+|B'|+|B''|$ minimal, where $A',B' \subseteq V'$ and $A'',B'' \subseteq V''$. We claim that $A''=\emptyset$ and $B'=\emptyset$; otherwise, if say there is $a'' \in A''$ then excluding a'' from A'' decreases A'' by exactly 1 and decreases $d_H(A' \cup A'', B' \cup B'')$ by at least $|B''| \geq 1$. Furthermore, no node has a copy both in A' and B'', by a similar argument. Hence A' is a copy in V' of some $A \subseteq V$, B'' is a copy in V'' of some $B \subseteq V$, and A and B are disjoint; thus (A,B) is an st-setpair. Consequently, $\kappa_H(s',t'') = d_H(A',B'')+2n-|A' \cup B''| = d_G(A,B)+n+n-|A \cup B| \geq n+\kappa_G(s,t)$. $\square$

1.3 General demands and costs

Since for directed graphs Survivable Network has strong inapproximability results already for small values of k, we will focus on undirected graphs and on variants that have good ratios for small values of k.

Theorem 1.2 (Chuzhoy & Khanna [9]) *Undirected* Survivable Network *admits ratio* $O(k^3 \ln |T|)$.

The idea is to decompose the problem into $O(k^3 \ln |T|)$ element connectivity Survivable Network problems.

Definition 1.1 *A set family $\mathcal{F}$ on a groundset T is k-**resilient** if for any pair (K, L) of disjoint subsets of T with $|L| = 2$ and $|K| \le k$, there exists $S \in \mathcal{F}$ such that $L \subseteq S$ and $K \subseteq T \setminus S$. Let $\rho_n(k)$ denote the minimum size of a k-resilient set-family on a groundset of n elements.*

Lemma 1.4 *Suppose that for a* Survivable Network *instance we are given on the set T of terminals a $(k-1)$-resilient family $\mathcal{F}$ of size ρ. Then such instance admits ratio 2ρ.*

Proof: For each $T_i \in \mathcal{F}$, compute a 2-approximate solution E_i using the algorithm of [12] to the instance of element connectivity Survivable Network with $Q_i = V \setminus T_i$ and the demands restricted to the pairs in T_i. Let $E = \bigcup_{i=1}^{\rho} E_i$ be the union of the solutions computed. Any feasible solution to the Survivable Network instance is also feasible for each element connectivity instance, hence, $c(E)$ is at most 2ρ times the optimal solution value of the Survivable Network instance. To prove that $G = (V, E)$ is a feasible solution, we show that for any $st \in D$ and any st-setpair (A, B), $d_E(A, B) + n - |A \cup B| \ge r_{st}$. If $n - |A \cup B| \ge k$ we are done, so assume that $n - |A \cup B| \le k - 1$. Let $T_i \in \mathcal{F}$ satisfy $\{s, t\} \subseteq T_i$ and $V \setminus (A \cup B) \subseteq T \setminus T_i$, so $V \setminus (A \cup B) \subseteq V \setminus T_i = Q_i$. Then $d_{E_i}(A, B) + n - |A \cup B| \ge r_{st}$, since E_i is a feasible solution to the corresponding element-connectivity Survivable Network instance. However, $E_i \subseteq E$, and thus $d_E(A, B) \ge d_{E_i}(A, B)$. Consequently, $d_E(A, B) + n - |A \cup B| \ge r_{st}$, as required. $\square$

We now show that there exists a k-resilient family of size $O(k^3 \ln |T|)$ and how to find such a family.

Theorem 1.3 (Chuzhoy & Khanna [9]) $\rho_n(k) = O(k^3 \ln n)$, *and there exists a randomized polynomial time algorithm that computes a k-resilient family within this bound, with high probability.*

Proof: If $n \le 3k + 1$ then the subsets of T of size 2 form a k-resilient family of size $O(k^2)$, so assume that

$n \geq 3k+2$. Also, by augmenting the groundset T by a set of k "dummy" elements and revising the notation to $n \leftarrow n+k$, we may consider sets K of size exactly k.

We will seek a "small" resilient family in which all sets have the same size p. We cast this problem as a Hitting-Set problem in a hypergraph $(\mathcal{V}, \mathcal{E})$ where $\mathcal{V}$ is the family of the subsets of T of size p and $\mathcal{E}$ is the family of pairs (K, L) of disjoint subsets of T with $|L| = 2$ and $|K| = k$, namely

$$\mathcal{V} = \{S : S \subseteq T, |S| = p\} \qquad \mathcal{E} = \{(K, L) : K, L \subseteq V, K \cap L = \emptyset, |L| = 2, |K| = k\} \ .$$

We say that $S \in \mathcal{V}$ **hits** $(K, L) \in \mathcal{E}$ if $L \subseteq S$ and $K \subseteq T \setminus S$; $\mathcal{F} \subseteq \mathcal{V}$ is a **hitting set** of $\mathcal{E}$ if for every $(K, L) \in \mathcal{E}$ there is $S \in \mathcal{V}$ that hits (K, L). By Definition 1.1, $\mathcal{F} \subseteq \mathcal{V}$ is k-resilient if and only if $\mathcal{F}$ is a hitting set of $\mathcal{E}$. Thus $\rho_n(k)$ is bounded by the minimum size of a hitting set of $\mathcal{E}$. A **fractional hitting set** of $\mathcal{E}$ is a function $h : \mathcal{V} \longrightarrow [0, 1]$ such that $\sum\{h(S) : S \text{ hits } (K, L)\} \geq 1$ for every $(K, L) \in \mathcal{E}$; the value of h is $\sum_{S \in \mathcal{V}} h(S)$. It is known that if $\mathcal{E}$ has a fractional hitting set of value τ, then $\mathcal{E}$ has a hitting set of size $O(\tau \ln |\mathcal{E}|)$. Our next goal is to show that $\mathcal{E}$ has a low value fractional hitting set, and to bound $\ln |\mathcal{E}|$.

Claim 1.1 *$|\mathcal{V}| = \binom{n}{p}$, $|\mathcal{E}| = \binom{n}{k}\binom{n-k}{2}$, and the size of each hyperedge $(K, L) \in \mathcal{E}$ is $s = \binom{n-k-2}{p-2}$.*

Proof: $|\mathcal{V}|$ equals the number $\binom{n}{p}$ of choices of p elements from n elements. $|\mathcal{E}|$ equals the number $\binom{n}{k}$ of choices of a set K of size k multiplied by the number $\binom{n-k}{2}$ of choices of a pair L from the set $V \setminus K$ of size $n-k$. For every $(K, L) \in \mathcal{E}$ the size of the set $\{S \subseteq V : |S| = p, L \subseteq S, K \subseteq V \setminus S\}$ equals the number $\binom{n-k-2}{p-2}$ of choices of the set $S \setminus L$ of size $p-2$ from the set $T \setminus (K \cup L)$ of size $n-k-2$. $\square$

We have $\left(\frac{n}{k}\right)^k \leq \binom{n}{k} \leq \left(\frac{n}{k}\right)^k \cdot e^k$. In particular, $\ln |\mathcal{E}| = O(k \ln n)$.

Assigning value $1/s$ to each $S \in \mathcal{V}$ gives a fractional hitting set of value $|\mathcal{V}|/s$. Denote $m = n-k$. Then:

$$\frac{|\mathcal{V}|}{s} = \frac{\binom{n}{p}}{\binom{m-2}{p-2}} = \frac{n!}{p!(n-p)!} \cdot \frac{(p-2)!(m-p)!}{(m-2)!} = \frac{m(m-1)}{p(p-1)} \cdot \frac{n!}{(n-p)!} \cdot \frac{(m-p)!}{m!} \leq \frac{m^2}{(p-1)^2} \prod_{i=1}^{p} \frac{n-i+1}{m-i+1}$$

Note that for $1 \leq i \leq p$ we have $\frac{n-i+1}{m-i+1} = 1 + \frac{n-m}{m-i+1} \leq 1 + \frac{k}{n-k-p}$. Let us choose p such that $\frac{k}{n-k-p} = \frac{1}{p}$, so $p = \frac{n-k}{k+1}$; assume that p is an integer, as adjustment to floors and ceilings only affects by a small amount the constant hidden in the $O(\cdot)$ term. Since $(1+1/p)^p \leq e$ we obtain

$$\prod_{i=1}^{p} \frac{n-i+1}{m-i+1} \leq \left(1 + \frac{1}{p}\right)^p \leq e \ .$$

Since we assume that $n \geq 3k+2$, we have $\frac{n-k}{k+1} \geq 2$ and thus $\frac{m}{p-1} \leq 2(k+1)$. Consequently, we get that $\frac{|\mathcal{V}|}{s} \cdot \ln |\mathcal{E}| = O(k^3 \ln n)$. This implies that a standard greedy algorithm for Hitting Set, produces a k-resilient family of size $O\left(k^3 \ln n\right)$, There is some difficulty to implement this algorithm in time polynomial in n; thus we use a randomized algorithm for Hitting-Set, by rounding each entry to 1 with probability determined by our fractional hitting set. It is known that repeating this rounding $2\lceil \ln |\mathcal{E}| \rceil$ times gives a hitting set w.h.p., and clearly its expected size is $2\lceil \ln |\mathcal{E}| \rceil$ times the value of the fractional hitting set. In our case, the value of every $S \in \mathcal{V}$ is $1/s = 1/\binom{n-k-2}{p-2}$. Thus we just need to sample a set S of size $p = \frac{n-k}{k+1}$ with probability $1/s$, independently, $2\lceil \ln |\mathcal{E}| \rceil = O(k \ln n)$ times. □

1.4 Biset functions and Survivable Network Augmentation problems

1.4.1 Biset formulation of Survivable Network problems

In Section 1.2 we formulated Mengers's Theorem in terms of setpairs. It would be more convenient to consider instead of a setpair (A,B) the pair of sets $(A, V \setminus B)$ called a "biset", defined as follows.

Definition 1.2 *An ordered pair $\mathbb{A} = (A, A^+)$ of subsets of V with $A \subseteq A^+$ is called a* **biset***; A is the* **inner part** *and A^+ is the* **outer part** *of $\mathbb{A}$, and $\partial\mathbb{A} = A^+ \setminus A$ is the* **boundary** *of $\mathbb{A}$. The* **co-set** *of $\mathbb{A}$ is $A^* = V \setminus A^+$; the* **co-biset** *of $\mathbb{A}$ is $\mathbb{A}^* = (A^*, V \setminus A)$. Let $\mathcal{V}$ denote the family of bisets over V.*

A **biset function** assigns to every $\mathbb{A} \in \mathcal{V}$ a real number; in our context, it will always be an integer.

Definition 1.3 *An* **edge covers a biset** *$\mathbb{A}$ if it goes from A to A^*. For a biset $\mathbb{A}$ and an edge-set/graph J let $\delta_J(\mathbb{A})$ denote the set of edges in J covering $\mathbb{A}$ and let $d_J(\mathbb{A}) = |\delta_J(\mathbb{A})|$. The* **residual function** *of a biset function f w.r.t. a partial f-cover J is defined by $f^J(\mathbb{A}) = f(\mathbb{A}) - d_J(\mathbb{A})$ for all $\mathbb{A} \in \mathcal{V}$. We say that an edge set/graph J* **covers** *a biset function f, or that J is an* **f-cover***, if $d_J(\mathbb{A}) \geq f(\mathbb{A})$ for all $\mathbb{A} \in \mathcal{V}$.*

We say that $\mathbb{A}$ is an st**-biset** if $s \in A$ and $t \in A^*$. By the node-capacitated version of Menger's Theorem we have that if a (directed or undirected) graph $G = (V,E)$ has node capacities $\{q(v) : v \in V\}$ then:

$$\lambda^q_G(s,t) = \min\{q(\partial\mathbb{A}) + d_G(\mathbb{A}) : \mathbb{A} \text{ is an } st\text{-biset}\} \ .$$

Thus $\lambda^q_G(s,t) \geq r_{st}$ if and only if $d_G(\mathbb{A}) \geq r_{st} - q(\partial\mathbb{A})$ for every st-biset $\mathbb{A}$ Consequently, we get that G satisfies r if and only if G covers the biset function defined by $f^q_r(\mathbb{A}) = \max\limits_{st\in\delta_D(\mathbb{A})} r_{st} - q(\partial\mathbb{A})$ for all $\mathbb{A} \in \mathcal{V}$, where we treat demand pairs as edges and denote by $\delta_D(\mathbb{A})$ the set of demand pairs that cover $\mathbb{A}$. In the Q-connectivity case we have $q(v) = 1$ if $v \in Q$ and $q(v) = \infty$ otherwise, so G satisfies r iff G covers the biset function $f_r = f^Q_r$ defined by

$$f_r(\mathbb{A}) = \begin{cases} \max\limits_{st\in\delta_D(\mathbb{A})} r_{st} - |\partial\mathbb{A}| & \text{if } \partial\mathbb{A} \subseteq Q \\ 0 & \text{otherwise} \end{cases}$$

We thus often will consider the following generic problem:

<u>Biset-Function Edge-Cover</u>

Input: A graph $\hat{G} = (V, \hat{E})$ with edge costs $\{c_e : e \in \hat{E}\}$ and a biset function f on V.

Output: A minimum cost edge-set $E \subseteq \hat{E}$ that covers f.

Here f may not be given explicitly, and an efficient implementation of algorithms requires that certain queries related to f can be answered in time polynomial in n. We will not consider implementation details. In most applications relevant polynomial time oracles are available via min-cut computations. In particular, we have a polynomial time separation oracle for the LP-relaxation due to Frank and Jordan [16]:

$$\textbf{(Biset-LP)} \qquad \begin{array}{rlll} \tau(f) = & \min & c\cdot x & \\ & \text{s.t.} & x(\delta_E(\mathbb{A})) \geq f(\mathbb{A}) & \forall \mathbb{A} \in \mathcal{V} \\ & & 0 \leq x_e \leq 1 & \forall e \in E \end{array}$$

Note that any $0,1$-valued biset function f corresponds to the **biset family** $\mathcal{F} = \{\mathbb{A} \in \mathcal{V} : f(\mathbb{A}) = 1\}$. We thus use for biset families the same terminology and notation as for biset functions, e.g., $\tau(\mathcal{F})$ is the optimal value of the Biset-LP for covering $\mathcal{F}$, and $\mathcal{F}^J = \{\mathbb{A} \in \mathcal{F} : \delta_J(\mathbb{A}) = \emptyset\}$ denotes the residual family of $\mathcal{F}$. In the Biset-Family Edge-Cover problem we seek a minimum cost edge-set $E \subseteq \hat{E}$ that covers $\mathcal{F}$.

1.4.2 The backwards augmentation method

We now describe the **backwards augmentation method** due to [20] for reducing a Biset-Function Edge-Cover problem to $\max(f) = \max\limits_{\mathbb{A}\in\mathcal{V}} f(\mathbb{A})$ instances of the Biset-Family Edge-Cover problem, but with a loss of only $O(\ln \max(f))$ in the approximation ratio. Let us say that Biset-Function Edge-Cover **admits LP-ratio** ρ if there exists a polynomial time algorithm that computes an f-cover of cost at most $\rho \cdot \tau(f)$.

Lemma 1.5 *Suppose that for a* Biset-Function Edge-Cover *instance* $(f, \hat{G}, c)$ *the following holds: for any* $E \subseteq \hat{E}$*, the* Biset-Family Edge-Cover *instance* $(\mathcal{F}, \hat{G} \setminus E, c)$ *with* $\mathcal{F} = \{\mathbb{A} \in \mathcal{V} : f^E(\mathbb{A}) = \max(f^E)\}$ *admits LP-ratio* ρ*. Then such* Biset-Function Edge-Cover *instance admits LP-ratio* $\rho \cdot H(\max(f))$*.*

Proof: Let x be an optimal solution for the Biset-LP for f. Consider the following algorithm.

Algorithm 1: BACKWARDS AUGMENTATION$(f, \hat{G}, c)$

$E \leftarrow \emptyset$
while $\max(f^E) > 0$ **do**
  find a cover $J \subseteq \hat{E}$ of the family $\mathcal{F} = \{\mathbb{A} \in \mathcal{V} : f^E(\mathbb{A}) = \max(f^E)\}$ of cost $c(J) \leq \rho\tau(\mathcal{F})$
  $E \leftarrow E \cup J$, $\hat{E} \leftarrow \hat{E} \setminus J$
return E

After each iteration $\max(f^E)$ decreases by at least 1, hence the algorithm terminates and computes a feasible solution. Consider some iteration with $\max(f^E) = \ell$. Note that $\frac{x}{\ell}$ is a feasible solution to the Biset-LP for the family $\mathcal{F}$ covered at this iteration. This implies $\tau(\mathcal{F}) \leq \frac{\tau(f)}{\ell}$, and hence $c(J) \leq \rho\tau(\mathcal{F}) \leq \rho\frac{\tau(f)}{\ell}$. Consequently, at the end of the algorithm $c(E) \leq \rho\tau(f) \sum_{\ell=k}^{1} \frac{1}{\ell} = \rho\tau(f)H(k)$. □

Applying the backward augmentation method means that at each iteration we cover all bisets with maximum residual demand. This is equivalent to increasing the connectivity by 1 between demand pairs for which $r_{st} - \lambda_G^Q(s,t)$ is maximum, where $G = (V, E)$ is the partial solution computed so far. Then we get the following problem, which is a restriction of Survivable Network to instances when $\hat{G}$ contains a subgraph $G = (V, E)$ of cost 0 and the demands are $r_{st} = \lambda_G^Q(s,t) + 1$ for all $st \in D$.

Survivable Network Augmentation

Input: A graph G, and edge set $\hat{E}$ with costs $\{c_e : e \in \hat{E}\}$ and a set $D \subseteq V \times V$ of demand pairs.

Output: A minimum cost edge-set $J \subseteq \hat{E}$ such that $\lambda_{G \cup J}^Q(s,t) \geq \lambda_G^Q(s,t) + 1$ for all $st \in D$.

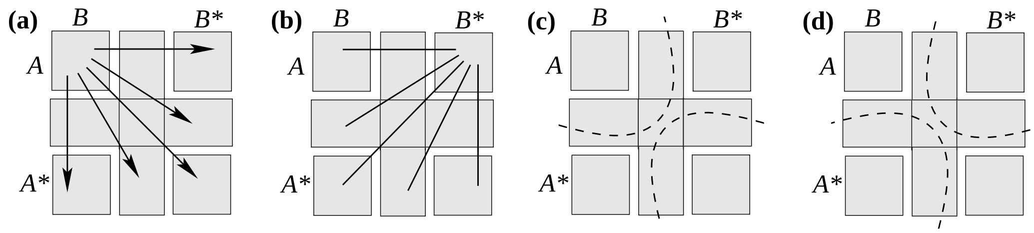

Figure 1.2: Illustration to Fact 1.2; the sets $A, \partial\mathbb{A}, A^*$ are the "rows" and $B, \partial\mathbb{B}, B^*$ are the "columns" of a 3×3 "matrix". (a) Types of directed edges that cover $\mathbb{A} \cap \mathbb{B}$. (b) Types of undirected edges that cover $\mathbb{A} \setminus \mathbb{B}$. (c) Nodes appearing in $\partial(\mathbb{A} \cap \mathbb{B})$ and in $\partial(\mathbb{A} \cup \mathbb{B})$. (d) Nodes appearing in $\partial(\mathbb{A} \setminus \mathbb{B})$ and in $\partial(\mathbb{B} \setminus \mathbb{A})$.

Definition 1.4 *A ts-biset $\mathbb{A}$ is ts-**tight** in a graph G if $|\partial\mathbb{A}| + d_G(\mathbb{A}) = \kappa_G(t,s)$; in the case of Q-connectivity we also require $\partial\mathbb{A} \subseteq Q$. Given an instance of* Survivable Network Augmentation*, we say that $\mathbb{A}$ is* **tight** *if it is ts-tight for some $ts \in D$.*

By Menger's Theorem we have:

Fact 1.1 *An edge set $J \subseteq \hat{E}$ is a feasible solution to the* Survivable Network Augmentation *problem if and only if J covers the family of tight bisets.*

1.4.3 Properties of tight bisets

By Fact 1.1, the Survivable Network Augmentation problem is equivalent to the Biset-Family Edge-Cover problem with $\mathcal{F}$ being the family of tight bisets; in this section we will establish some properties of this family.

Definition 1.5 *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 $A^+ \subseteq B^+$.*

The following properties of bisets are easy to verify (see Fig. 1.2).

Fact 1.2 *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 $\mathbb{A}, \mathbb{B}$ (see Fig. 1.2); 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 $|\partial\mathbb{A}| + |\partial\mathbb{B}| = |\partial(\mathbb{A} \cap \mathbb{B})| + |\partial(\mathbb{A} \cup \mathbb{B})| = |\partial(\mathbb{A} \setminus B)| + |\partial(\mathbb{B} \setminus \mathbb{A})|$.*

For a biset function f and bisets $\mathbb{A}, \mathbb{B}$ the **supermodular inequality** is:

$$f(\mathbb{A} \cap \mathbb{B}) + f(\mathbb{A} \cup \mathbb{B}) \geq f(\mathbb{A}) + f(\mathbb{B})$$

A biset function f is: **supermodular** if the supermodular inequality holds for all $\mathbb{A}, \mathbb{B} \in \mathcal{V}$, **submodular** if $-f$ is supermodular, and **modular** if f is both submodular and supermodular. f is **symmetric** if $f(\mathbb{A}) = f(\mathbb{A}^*)$ for all $\mathbb{A} \in \mathcal{V}$. It is easy to see that for symmetric f, if the supermodular inequality holds for $\mathbb{A}, \mathbb{B}^*$ then $f(\mathbb{A} \setminus \mathbb{B}) + f(\mathbb{B} \setminus \mathbb{A}) \geq f(\mathbb{A}) + f(\mathbb{B})$. From Fact 1.2 one can deduce the following.

- For any (directed or undirected) graph G the function $d_G(\cdot)$ is submodular, and if G is an undirected graph then $d_G(\cdot)$ is symmetric.
- The function $|\partial(\cdot)|$ is modular.

The following lemma from [41] gives a general "uncrossing property" of tight bisets, namely, a useful characterization of those pairs $\mathbb{A}, \mathbb{B} \in \mathcal{F}$ for which $\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cup \mathbb{B} \in \mathcal{F}$ or $\mathbb{A} \setminus \mathbb{B}, \mathbb{B} \setminus \mathbb{A} \in \mathcal{F}$ holds.

Lemma 1.6 *Let $\mathbb{A}, \mathbb{B}$ be bisets in a graph G such that $\mathbb{A}$ is aa'-tight, $\mathbb{B}$ is bb'-tight, and $\kappa_G(a, a') \geq \kappa_G(b, b')$.*

(i) *If $\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cup \mathbb{B}$ are both aa'-bisets then they are both aa'-tight. Otherwise, the following holds.*

  (a) *If $\mathbb{A} \cap \mathbb{B}$ is an aa'-biset and $\mathbb{A} \cup \mathbb{B}$ is a bb'-biset then $\mathbb{A} \cap \mathbb{B}$ is aa'-tight and $\mathbb{A} \cup \mathbb{B}$ is bb'-tight.*

  (b) *If $\mathbb{A} \cap \mathbb{B}$ is a bb'-biset and $\mathbb{A} \cup \mathbb{B}$ is an aa'-biset then $\mathbb{A} \cap \mathbb{B}$ is bb'-tight and $\mathbb{A} \cup \mathbb{B}$ is aa'-tight.*

(ii) *Suppose that G is an undirected graph. If $\mathbb{A} \setminus \mathbb{B}$ is an aa' biset and $\mathbb{B} \setminus \mathbb{A}$ is an $a'a$-biset then $\mathbb{A} \setminus \mathbb{B}$ is aa'-tight and $\mathbb{B} \setminus \mathbb{A}$ is $a'a$-tight. Otherwise, the following holds.*

  (a) *If $\mathbb{A} \setminus \mathbb{B}$ is an aa'-biset and $\mathbb{B} \setminus \mathbb{A}$ is a bb'-biset then $\mathbb{A} \setminus B$ is aa'-tight and $\mathbb{B} \setminus \mathbb{A}$ is bb' tight.*

  (b) *If $\mathbb{A} \setminus \mathbb{B}$ is a $b'b$-biset and $\mathbb{B} \setminus \mathbb{A}$ is an $a'a$-biset then $\mathbb{A} \setminus B$ is $b'b$-tight and $\mathbb{B} \setminus \mathbb{A}$ is $a'a$ tight.*

Proof: Let us use the notation $\kappa(s, t) = \kappa_G(s, t)$ and $\psi(\mathbb{A}) = |\partial \mathbb{A}| + d_G(\mathbb{A})$. Note that ψ is submodular, namely, $\psi(\mathbb{A}) + \psi(\mathbb{B}) \geq \psi(\mathbb{A} \cap \mathbb{B}) + \psi(\mathbb{A} \cup \mathbb{B})$ for all $\mathbb{A}, \mathbb{B} \in \mathcal{V}$. Also note that for any st-biset $\mathbb{C}$ we have $\psi(\mathbb{C}) \geq \kappa(s, t)$ and $\kappa(s, t) = \psi(\mathbb{C})$ iff $\mathbb{C}$ is st-tight. If $\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cup \mathbb{B}$ are both aa'-bisets then:

$$\kappa(a, a') + \kappa(a, a') \geq \kappa(a, a') + \kappa(b, b') = \psi(\mathbb{A}) + \psi(\mathbb{B}) \geq \psi(\mathbb{A} \cap \mathbb{B}) + \psi(\mathbb{A} \cup \mathbb{B}) \geq \kappa(a, a') + \kappa(a, a') \ .$$

Hence equality holds everywhere, so $\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cup \mathbb{B}$ are both aa'-tight (and $\kappa(a,a') = \kappa(b,b')$).

In case (ia) we have $\kappa(a,a') + \kappa(b,b') = \psi(\mathbb{A}) + \psi(\mathbb{B}) \geq \psi(\mathbb{A} \cap \mathbb{B}) + \psi(\mathbb{A} \cup \mathbb{B}) \geq \kappa(a,a') + \kappa(b,b')$. Hence equality holds everywhere, so $\mathbb{A} \cap \mathbb{B}$ is aa'-tight and $\mathbb{A} \cup \mathbb{B}$ is bb'-tight.

In case (ib) we have $\kappa(a,a') + \kappa(b,b') = \psi(\mathbb{A}) + \psi(\mathbb{B}) \geq \psi(\mathbb{A} \cap \mathbb{B}) + \psi(\mathbb{A} \cup \mathbb{B}) \geq \kappa(b,b') + \kappa(a,a')$. Hence equality holds everywhere, so $\mathbb{A} \cap \mathbb{B}$ is aa'-tight and $\mathbb{A} \cup \mathbb{B}$ is bb'-tight.

We prove (ii). If G is undirected then ψ is symmetric, and thus $\psi(\mathbb{A}) + \psi(\mathbb{B}) \geq \psi(\mathbb{A} \setminus \mathbb{B}) + \psi(\mathbb{B} \setminus \mathbb{A})$. Also note that $\kappa(s,t) = \kappa(t,s)$ for all $s,t \in V$. If $\mathbb{A} \setminus \mathbb{B}$ is an aa' biset and $\mathbb{B} \setminus \mathbb{A}$ is an $a'a$-biset then

$$\kappa(a,a') + \kappa(a',a) \geq \kappa(a,a') + \kappa(b,b') = \psi(\mathbb{A}) + \psi(\mathbb{B}) \geq \psi(\mathbb{A} \setminus \mathbb{B}) + \psi(\mathbb{B} \setminus \mathbb{A}) \geq \kappa(a,a') + \kappa(a',a) \ .$$

Hence equality holds everywhere, so $\mathbb{A} \setminus \mathbb{B}$ is aa'-tight and $\mathbb{B} \setminus \mathbb{A}$ is $a'a$-tight (and $\kappa(a,a') = \kappa(b,b')$).

In case (iia) we have $\kappa(a',a) + \kappa(b',b) = \psi(\mathbb{A}) + \psi(\mathbb{B}) \geq \psi(\mathbb{A} \setminus \mathbb{B}) + \psi(\mathbb{B} \setminus \mathbb{A}) \geq \kappa(a',a) + \kappa(b',b)$. Hence equality holds everywhere, so $\mathbb{A} \setminus \mathbb{B}$ is $b'b$-tight and $\mathbb{B} \setminus \mathbb{A}$ is $a'a$-tight.

In case (iib) we have $\kappa(a,a') + \kappa(b,b') = \psi(\mathbb{A}) + \psi(\mathbb{B}) \geq \psi(\mathbb{A} \setminus \mathbb{B}) + \psi(\mathbb{B} \setminus \mathbb{A}) \geq \kappa(a',a) + \kappa(b',b)$. Hence equality holds everywhere, so $\mathbb{A} \setminus \mathbb{B}$ is aa'-tight and $\mathbb{B} \setminus \mathbb{A}$ is bb'-tight. □

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. From Lemma 1.6 we have the following:

Corollary 1.1 *Let $\mathbb{A}, \mathbb{B}$ be distinct cores of the family of tight bisets, where $\mathbb{A}$ is aa'-tight and $\mathbb{B}$ is bb'-tight, and $aa', bb' \in D$. If $A \cap B \neq \emptyset$ then $\{a,a'\} \cap \partial\mathbb{B} \neq \emptyset$ or $\{b,b'\} \cap \partial\mathbb{A} \neq \emptyset$, and if $a' = b'$ then $a \in \partial\mathbb{B}$ or $b \in \partial\mathbb{A}$.*

1.5 Rooted and subset uniform demands

Let us say that a graph G is **k-(T,s)-connected** if $\kappa_G(t,s) \geq k$ for all $t \in T$. In the k-(T,s)-Connectivity Augmentation problem the goal is to augment a k-(T,s)-connected graph G by a minimum cost edge set J such that $G \cup J$ is $(k+1)$-(T,s)-connected. We survey the following result.

Theorem 1.4 (Nutov [39, 40]) *Undirected k-(T,s)-Connectivity Augmentation admits LP-ratio $O(k)$.*

Theorem 1.4 can be applied to compute a solution to Rooted Survivable Network in k iterations, where at iteration $\ell = 1, \ldots, k$ we increase the connectivity between nodes in T and s from $\ell - 1$ to ℓ. For general rooted demands we get ratio $\sum_{\ell=1}^{k} O(\ell) = O(k^2)$. For demands $r_{st} = k$ for all $t \in T$, this is equivalent to the backward augmentation method, so we get ratio $O(k \ln k)$ in this case. Summarizing, we have:

Corollary 1.2 *Undirected* Rooted Survivable Network *with demands $r_{st} = k$ for all $t \in T$ admits LP-ratio $O(k \ln k)$. For rooted general demands the problem admits LP-ratio $O(k^2)$.*

We now prove Theorem 1.4. The idea of the proof is similar to the one of Theorem 1.2 – to decompose the problem into $O(k)$ element-connectivity Survivable Network problems. We will prove a more general result stated in biset families terms. The family of tight bisets in our problem is

$$\mathcal{F} = \{\mathbb{A} \in \mathcal{V} : \mathbb{A} \text{ is a } ts\text{-biset for some } t \in T, |\partial\mathbb{A}| + d_G(\mathbb{A}) = k\}$$

Definition 1.6 *A biset family $\mathcal{F}$ is* **uncrossable** *if $\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cap \mathbb{B} \in \mathcal{F}$ or $\mathbb{A} \setminus \mathbb{B}, \mathbb{B} \setminus \mathbb{A} \in \mathcal{F}$ for any $\mathbb{A}, \mathbb{B} \in \mathcal{V}$. Let us say that bisets $\mathbb{A}, \mathbb{B}$:* T**-intersect** *if $A \cap B \cap T \neq \emptyset$, and* T**-co-cross** *if $A \cap B^* \cap T \neq \emptyset$ and $B \cap A^* \cap T \neq \emptyset$. A biset family $\mathcal{F}$ is* T**-uncrossable** *if $A \cap T \neq \emptyset$ for all $\mathbb{A} \in \mathcal{F}$ and 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.*

Lemma 1.7 *If G is a k-(T, s)-connected undirected graph then the family $\mathcal{F}$ of tight bisets is T-uncrossable.*

Proof: Let $\mathbb{A}, \mathbb{B} \in \mathcal{F}$. If $\mathbb{A}, \mathbb{B}$ T-intersect then $\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cup \mathbb{B} \in \mathcal{F}$ by Lemma 1.6(i). If $\mathbb{A}, \mathbb{B}$ T-co-cross then $\mathbb{A} \setminus \mathbb{B}, \mathbb{B} \setminus \mathbb{A} \in \mathcal{F}$ by Lemma 1.6(iia). $\square$

Recall that by Fact 1.1 any Survivable Network Augmentation problem is equivalent to the corresponding Biset-Family Edge-Cover problem with $\mathcal{F}$ being the family of tight bisets. By Lemma 1.7, in the k-(T, s)-Connectivity Augmentation problem the family of tight bisets of the input graph G is T-uncrossable. Also note that $|\partial\mathbb{A}| \leq k$ for every tigh biset $\mathbb{A}$ (see Definition 1.4). Thus the following theorem implies Theorem 1.4.

Theorem 1.5 (Nutov [39, 40]) *Undirected* Biset-Family Edge-Cover *with T-uncrossable biset family $\mathcal{F}$ admits LP-ratio $\frac{4}{3}\left(\ell + \frac{4\gamma^2}{\gamma+1}\right) + 2 = O(\gamma + 1)$, where $\gamma = \max_{\mathbb{A},\mathbb{B} \in \mathcal{F}} |\partial\mathbb{A} \cap B \cap T|$ and ℓ is the least integer such that $2^\ell \geq \gamma + 1$.*

In the rest of this section we prove Theorem 1.5. For an $\mathcal{F}$-core $\mathbb{C} \in \mathcal{C}(\mathcal{F})$, 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 $\mathbb{C}$.

Lemma 1.8 *Let $\mathcal{F}$ be a T-uncrossable biset family and let $p = \min_{\mathbb{A} \in \mathcal{F}} |A \cap T|$. Then there exists a polynomial time algorithm that computes a partition Π of $\mathcal{C}(\mathcal{F})$ with at most $2\lfloor \gamma/p \rfloor + 1$ parts such that for each $\mathcal{C} \in \Pi$ the family $\bigcup_{\mathbb{C} \in \mathcal{C}} \mathcal{F}(\mathbb{C})$ is uncrossable. Furthermore, if $p \geq \gamma + 1$ then $\mathcal{F}$ is uncrossable.*

Proof: It is easy to see that if $p \geq \gamma + 1$ then any $\mathbb{A}, \mathbb{B} \in \mathcal{F}$ must T-intersect or T-co-cross; thus $\mathcal{F}$ is uncrossable in this case. We prove the first statement. For $\mathbb{C}_i \in \mathcal{C}(\mathcal{F})$ let $\mathbb{B}_i = \bigcup_{\mathbb{A} \in \mathcal{F}(\mathbb{C}_i)} \mathbb{A}$ be the union of the bisets in the halo family of $\mathbb{C}_i$, namely, $\mathbb{B}_i$ is the inclusionwise maximal biset in $\mathcal{F}(\mathbb{C}_i)$. Note that since $\mathcal{F}$ is T-uncrossable, then for any $\mathbb{A}_i \in \mathcal{F}(\mathbb{C}_i)$ and $\mathbb{A}_j \in \mathcal{F}(\mathbb{C}_j)$ we have:

(i) $\mathbb{A}_i, \mathbb{A}_j$ T-intersect if and only if $i = j$.

(ii) If $C_i \cap B_j^* \cap T$ and $C_j \cap B_i^* \cap T$ are both nonempty then $\mathbb{A}_i, \mathbb{A}_j$ T-co-cross.

Construct an auxiliary directed graph $\mathcal{J}$ as follows. The node set of $\mathcal{J}$ is $\mathcal{C}(\mathcal{F})$. Add an arc $\mathbb{C}_i\mathbb{C}_j$ if $C_i \cap T \subseteq \partial \mathbb{B}_j$. The indegree of every node in $\mathcal{J}$ is at most $\lfloor \gamma/p \rfloor$, by (i). This implies that every subgraph of the underlying graph of $\mathcal{J}$ has a node of degree $2\lfloor \gamma/p \rfloor$. A graph is d-degenerate if every subgraph of it has a node of degree $\leq d$. It is known that any d-degenerate graph can be colored with $d+1$ colors, in polynomial time. Hence $\mathcal{J}$ is $(2\lfloor \gamma/p \rfloor + 1)$-colorable, and such a coloring can be computed in polynomial time. Consequently, we can compute in polynomial time a partition Π of $\mathcal{C}(\mathcal{F})$ into at most $2\lfloor \gamma/p \rfloor + 1$ independent sets. For each independent set $\mathcal{C} \in \Pi$, the family $\bigcup_{\mathbb{C} \in \mathcal{C}} \mathcal{F}(\mathbb{C})$ is uncrossable, by (ii). Hence Π is a partition as required, and the proof of the lemma is complete. $\square$

The best known ratio for covering an uncrossable biset family is 2. However, ratio 4/3 is known for the case when the uncrossable family is a union of its halo families [19].

Consider the following algorithm.

Algorithm 2: T-Uncrossable Edge-Cover($\hat{G}, c, \mathcal{F}$)

$J \leftarrow \emptyset$

while $p := \min_{\mathbb{A} \in \mathcal{C}(\mathcal{F}^J)} |A \cap T| \leq \gamma/2$ **do**

  find a partition Π of $\mathcal{C}(\mathcal{F}^J)$ as in Lemma 1.8 with at most $2\lfloor \gamma/p \rfloor + 1$ parts

  for every $\mathcal{C} \in \Pi$ find a 4/3-approximate edge-cover $J_{\mathcal{C}}$ of the family $\bigcup_{\mathbb{C} \in \mathcal{C}} \mathcal{F}^J(\mathbb{C})$

  for every $\mathcal{C} \in \Pi$ do: $J \leftarrow J \cup J_{\mathcal{C}}$

find a 2-approximate edge-cover of J' of $\mathcal{F}^J$ and add J' to J

return J

Let p_i denote the value of p at the beginning of iteration i in the while loop. Initially, $p_1 \geq 1$. Note that for any T-uncrossable family $\mathcal{F}$, if an $\mathcal{F}$-core $\mathbb{C}$ and $\mathbb{A} \in \mathcal{F}$ T-intersect then $\mathbb{C} \subseteq \mathbb{A}$; this implies that if I is an edge set that covers all halo families of $\mathcal{F}$ then every $\mathcal{F}^I$-core $\mathbb{A}$ contains at least two $\mathcal{F}$-cores. From this it follows that $p_i \geq 2p_{i-1}$ for all i. Thus the number of iterations in the while loop is at most $\ell - 1$, where ℓ is the least integer such that $2^\ell \geq \gamma + 1$. Consequently, the number of simple uncrossable biset families covered in the while loop is bounded by

$$\sum_{i=0}^{\ell-1} (2\lfloor \gamma/2^i \rfloor + 1) \leq \ell + 2\gamma \sum_{i=0}^{\ell-1} (1/2)^i = \ell + 4\gamma(1 - 1/2^\ell) \leq \ell + \frac{4\gamma^2}{\gamma + 1} .$$

Thus the approximation ratio is bounded by $\frac{4}{3}\left(\ell + \frac{4\gamma^2}{\gamma+1}\right) + 2$, as claimed in Theorem 1.5.

Let us now consider the Subset k-Connected Subgraph problem. Let us say that G is **k-T-connected** if $\kappa_G(s,t) \geq k$ for all $s,t \in T$. In the k-T-Connectivity Augmentation problem the goal is to augment a k-T-connected graph G by a minimum cost edge set J such that $G \cup J$ is $(k+1)$-T-connected. The ratios for this problem can derived from those of the k-(T,s)-Connectivity Augmentation problem via the following relation (the proof is omitted).

Theorem 1.6 (Nutov [42]) *For $|T| > k$, if k-(T,s)-Connectivity Augmentation admits ratio (LP-ratio) ρ then k-T-Connectivity Augmentation admits the following ratios (LP-ratios):*

(i) $b(\rho + k) + O(\mu^2 \ln \mu)$, *where $b = 1$ for undirected graphs and $b = 2$ for directed graphs, and* $\mu = \frac{T}{|T|-k}$.

(ii) $\rho \cdot O(\mu \ln k)$ *for both directed and undirected graphs; furthermore, this is so also for $\{0,1\}$-costs.*

For $|T| > k$, the best known values of ρ on undirected graphs are $O(k)$ for arbitrary costs and $O(\ln k)$ for $\{0,1\}$-costs. For directed graphs $\rho = |T|$ for arbitrary costs and $\rho = O(\log |T|)$ for $\{0,1\}$-costs. These are LP-ratios, so the backward augmentation method can be applied. Thus Theorem 1.6 implies the following.

Corollary 1.3 *For $|T| > k$, k-T-*Connectivity Augmentation *admits the following LP-ratios.*

- *For undirected graphs: $O(k + \mu^2 \ln \mu)$ for arbitrary costs, and $O(\mu \ln^2 k)$ for $\{0,1\}$-costs, $\mu = \frac{T}{|T|-k}$.*
- *For directed graphs: $2(|T| + k) + O(\mu^2 \ln \mu)$ for arbitrary costs, and $O(\mu \ln k \ln |T|)$ for $\{0,1\}$-costs.*

For Subset k-Connected Subgraph*, the ratios are larger by a factor of $O(\log k)$.*

1.6 $\{0,1\}$-Costs Survivable Network **problems**

Here we consider the $\{0,1\}$-Costs Survivable Network problem, that can be formulated as follows.

$\{0,1\}$-Costs Survivable Network
Input: A directed/undirected graph $G = (V, E)$, $Q \subseteq V$, and connectivity demands $\{r_{st} : st \in D\}$.
Output: A minimum size set J of new edges (any edge is allowed) such that $G \cup J$ satisfies r.

Theorem 1.7 (Kortsarz & Nutov [31], Nutov [41]) *Suppose that for an instance of* $\{0,1\}$-Costs Survivable Network *we are given a node subset $Z \subseteq V$ such that $Z \setminus \partial\mathbb{A} \neq \emptyset$ for every biset $\mathbb{A} \in \mathcal{V}$ with $f_r(\mathbb{A}) > 0$. Then the problem admits ratio $O(|Z| \ln |D|)$. For undirected graphs the problem also admits ratio $O(|Z| \ln^2 k)$.*

Note that the following choices of Z satisfy the assumption of Theorem 1.7: (i) $Z = \{z\}$ for some $z \in V \setminus Q$; (ii) $Z = \{s\}$ in the case of rooted demands; (iii) any $Z \subseteq V$ with $|Z| = k$. Thus we have:

Corollary 1.4 $\{0,1\}$-Costs Survivable Network *admits ratios $O(\ln |D|)$ if $Q \neq V$ and $O(k \ln |D|)$ if $Q = V$. For undirected graphs the problem also admits ratios $O(\ln^2 k)$ if $Q \neq V$ or if the demands are rooted, and ratio $O(k \ln^2 k)$ otherwise.*

For $Q \neq V$ the ratio $O(\ln |D|)$ is tight since the problem is Set-Cover hard [37]; for directed graphs this is so even for rooted $0,1$ requirements [13]. For $Q = V$, the ratio $O(k \ln |D|)$ is tight if k is small, but may seem weak if k is large. However, a much better ratio might not exist, see Theorem 1.10.

The proof of Theorem 1.7 follows. An undirected edge set (or a graph) S is a **star** if all its edges are incident to the same node, called the **center** of the star. In the case of directed graphs, S should be either an **outstar** – all edges in S leave the center, or an **instar** – all edges in S enter the center. Let Star Survivable Network be the restriction of Survivable Network to instances when the set of positive cost edges in $\hat{E}$ is a star S with center $z \in V \setminus Q$. Note that here the edges incident to z may have arbitrary costs.

Lemma 1.9 *If* Star Survivable Network *admits ratio* ρ *then any* $\{0,1\}$-Costs Survivable Network *instance* $\mathcal{I}$ *with* Z *as in Theorem 1.7 admits ratio* $2|Z|\rho$.

Proof: For every $z \in Z$ obtain from $\mathcal{I}$ an instance $\mathcal{I}_z$, in which z is excluded from Q, and except the edges in E, only edges incident to z are allowed, by cost 1 each. In the undirected case, $\mathcal{I}_z$ is a Star Survivable Network instance; for each $z \in Z$ we compute a ρ-approximate solution I_z for $\mathcal{I}_z$ and return $I = \bigcup\limits_{z \in Z} I_z$. In the directed case we further split the instance $\mathcal{I}_z$ into two Star Survivable Network instances $\mathcal{I}_z^{out}$ and $\mathcal{I}_z^{in}$.

- $\mathcal{I}_z^{out}$ is obtained from $\mathcal{I}_z$ by adding to E zero cost edges from every $v \in V \setminus \{z\}$ to z, and allowing only additional edges leaving z, by cost 1 each.

- $\mathcal{I}_z^{in}$ is obtained from $\mathcal{I}_z$ by adding to E zero cost edges from z to every $v \in V \setminus \{z\}$, and allowing only additional edges entering z, by cost 1 each.

Each of $\mathcal{I}_z^{out}, \mathcal{I}_z^{in}$ is a Star Survivable Network instance; for each $z \in Z$ we compute a ρ-approximate solution I_z^{out} and I_z^{in} for $\mathcal{I}_z^{out}$ and $\mathcal{I}_z^{in}$, respectively, set $I_z = I_z^{out} \cup I_z^{in}$, and return $I = \bigcup\limits_{z \in Z} I_z$.

We show that I is a feasible solution to $\mathcal{I}$. We prove this for the directed case, and the proof for the undirected case is similar. Suppose to the contrary that there is $st \in D$ such that $\lambda^Q_{G \cup I}(s,t) < r_{st}$. Then there is an st-biset $\mathbb{A}$ such that $|\partial \mathbb{A}| + d_{G \cup I}(\mathbb{A}) < r_{st}$. By the assumption on Z, there exists $z \in Z$ such that $z \in A$ or $z \in A^*$. Since $I_z^{out}, I_z^{in} \subseteq I$, $|\partial \mathbb{A}| + d_{G \cup I^{out}}(\mathbb{A}) < r_{st}$ and $|\partial \mathbb{A}| + d_{G \cup I^{in}}(\mathbb{A}) < r_{st}$. If $z \in A$ then the former inequality contradicts the feasibility of I_z^{out} for $\mathcal{I}_z^{out}$, since adding to G links entering z does affect the inequality. The contradiction for the case $z \in A^*$ is obtained in a similar way.

Now let J be an optimal solution to $\mathcal{I}$. We will show that $I_z \leq 2\rho|J|$ for every $z \in Z$, and thus $I \leq 2|Z|\rho$. Obtain from J a solution J_z for the instance $\mathcal{I}_z$ by subdividing every edge in J by a new node, and then identifying all new nodes into z. It is easy to see that J_z is a feasible solution for $\mathcal{I}_z$, and clearly $|J_z| = 2|J|$. Thus in the undirected case our ρ-approximate solution I_z satisfies $|I_z| \leq 2\rho|J_z| = 2\rho|J|$. In the directed case we further split J_z into sets J_z^{out} and J_z^{in} of edges leaving and entering z, respectively. Then $|I_z^{out}| \leq \rho|J_z^{out}|$ and $|I_z^{in}| \leq \rho|J_z^{in}|$, implying $|I_z| = |I_z^{out}| + |I_z^{in}| \leq \rho(|J_z^{out}| + |J_z^{in}|) = \rho|J_z| = 2\rho|J|$. This concludes the proof of the lemma. $\square$

We prove the following theorem that together with Lemma 1.9 implies Theorem 1.7.

Theorem 1.8 ([31, 41, 32]) *Directed* Star Survivable Network *admits ratio* $H(|D|)$, *where* $H(j) = \sum_{i=1}^{j} 1/i$ *is the* j*th harmonic number. For undirected graphs the problem admits ratio* $H(k) \cdot H(8k^2 - 14k + 7)$, $k \geq 2$.

1.6.1 Ratio $H(|D|)$ for directed Star Survivable Network

Here we prove the directed part of Theorem 1.8. Since it concerns directed graphs only, we may consider only the edge connectivity case by applying Lemma 1.1, since the reduction in the proof of Lemma 1.1 preserves our problem type and does not change the parameter $|D|$. We will also assume that F is an outstar.

We use a result due to Wolsey [46] about the performance of a greedy algorithm for the Submodular Cover problem. Here we are given two set functions on subsets of a groundset S: a cost-function c and and an integer valued "progress function" p. Each of the functions may be given by a value oracle, meaning that there exists an oracle that given a subset of S return the corresponding function value. The goal is to find $J \subseteq S$ of minimum cost such that $p(J) = p(S)$. In the Submodular Cover problem the function p is submodular and non-decreasing, and the cost of $J \subseteq S$ is $c(J) = \sum_{e \in J} c_e$ for some non-negative costs $\{c_e : e \in S\}$. Wolsey [46] proved that the greedy algorithm, that starts with $J = \emptyset$ and as long as $p(J) < p(S)$ repeatedly adds to J an element $e \in S \setminus J$ with maximum $\frac{p(J \cup \{e\}) - p(J)}{c_e}$, has ratio $H\left(\max_{e \in S} p(\{e\}) - p(\emptyset)\right)$.

In our case S is an outstar, and for $J \subseteq S$ let

$$p(J) = \sum_{st \in D} \min\{p_{st}(J), r_{st}\} \quad \text{where} \quad p_{st}(J) = \lambda_{G \cup J}(s,t)$$

It is not hard to verify that p is non-decreasing, and that J is a feasible solution to a Survivable Network instance if and only if $p(J) = r(D)$. It is also easy to see that $\lambda_{G \cup \{e\}}(s,t) - \lambda_G(s,t) \leq 1$ for any graph G and any edge e, as adding a single edge can increase the st-connectivity by at most 1. Thus for any edge e we have $p_{st}(\{e\}) - p_{st}(\emptyset) \leq 1$, which implies $p(\{e\}) - p(\emptyset) \leq |D|$.

Note that if each function $p_{st}(J) = \lambda_{G \cup J}(s,t)$ is submodular then so is p. Indeed, it is known (c.f. [44]) that if f is submodular and non-decreasing, then $\min\{f, \ell\}$ is submodular for any constant ℓ; thus $\min\{p_{st}(J), r_{st}\}$ is submodular. As a sum of submodular functions is also submodular, we obtain that p is submodular.

Fix $s,t \in V$ and for $J \subseteq S$ let $f(J) = \lambda_{G \cup J}(s,t)$. We show that if S is a star then f is submodular (this may not be so if S is not a star). We use the following known characterization of submodularity, c.f. [44]: *A set-function f on a groundset S is submodular iff*

$$f(J_0 \cup \{e\}) + f(J_0 \cup \{e'\}) \geq f(J_0) + f(J_0 \cup \{e,e'\}) \quad \forall J_0 \subseteq S, e, e' \in F \setminus J_0$$

Let $J_0 \subseteq S$. Revising our notation to $G \leftarrow G \cup J_0$, $S \leftarrow S \setminus J_0$, and $f(J) \leftarrow f(J_0 \cup J) - f(J_0)$, we get that the above condition is equivalent to

$$f(\{e\}) + f(\{e'\}) \geq f(\{e,e'\}) \quad \forall e, e' \in S \ .$$

As before, S is an outstar and $f(J) = \lambda_{G \cup J}(s,t) - \lambda_G(s,t)$ is the increase in the (s,t)-edge-connectivity as a result of adding J to G. We prove the following general statement, that implies the above; it says that if an augmenting edge set J is an outstar that increases the st-edge-connectivity by h, then there are h edges in J that cover all st-tight sets, and thus each of these edges increases the st-edge-connectivity by 1.

Lemma 1.10 *Let $G = (V,E)$ be a directed graph and J an outstar with center z on V. Let $s,t \in V$, and let $h = \lambda_{G \cup J}(s,t) - \lambda_G(s,t)$. Then there is $I \subseteq J$ of size h such that $\lambda_{G \cup \{e\}}(s,t) = \lambda_G(s,t) + 1$ for every $e \in I$.*

Proof: Let $\mathcal{F}$ be the family of st-tight sets. Then $\mathcal{F}$ is non-empty, and by Lemma 1.6(i) $A \cap B, A \cup B \in \mathcal{F}$ for any $A, B \in \mathcal{F}$. As the intersection of all sets in $\mathcal{F}$ contains s and thus nonempty, $\mathcal{F}$ has a unique inclusion-minimal set $A_{\min}$ and a unique inclusion-maximal set $A_{\max}$, and $A_{\min} \subseteq A_{\max}$.

Let $I = \{zv \in J : a \in A_{\min}, v \in V \setminus A_{\max}\}$ be the set of edges in J that go from $A_{\min}$ to $V \setminus A_{\max}$. Each edge in I covers $\mathcal{F}$, hence by Menger's Theorem $\lambda_{G \cup \{e\}}(s,t) = \lambda_G(s,t) + 1$ for every $e \in I$.

It remains to prove that $|I| \geq h$. We claim that since J is an outstar, then $\lambda_{G \cup J}(s,t) \leq \lambda_G(s,t) + |I|$, hence $|I| \geq \lambda_{G \cup J}(s,t) - \lambda_G(s,t) = h$. Note that from Menger's Theorem we have

$$\lambda_{G \cup J}(s,t) \leq \lambda_G(s,t) + |\delta_J(A_{\min})| \qquad \lambda_{G \cup J}(s,t) \leq \lambda_G(s,t) + |\delta_J(A_{\max})|$$

The first inequality implies that if $\delta_J(A_{\min}) = \emptyset$, then $\lambda_{G \cup J}(s,t) = \lambda_G(s,t)$, and then we are done. Else, we must have $z \in A_{\min}$. In this case $I = \delta_J(A_{\max})$, since J is a star. Then the second inequality implies $\lambda_{G \cup J}(s,t) \leq \lambda_G(s,t) + |I|$, as claimed. $\square$

1.6.2 Ratio $O(\ln^2 k)$ for undirected Star Survivable Network

Here we prove the undirected part of Theorem 1.8. For simplicity of exposition, we consider the case $Q = V \setminus \{z\}$ where z is the center of the star of the available edges.

As an intermediate problem, we consider the augmentation version of the problem, namely:

Star Survivable Network Augmentation
Input: A graph $G = (V,E)$, a star S with center z and costs $\{c_e : e \in S\}$, and a set D of demand pairs.
Output: A minimum cost edge set $J \subseteq S$ such that $\lambda^Q_{G \cup J}(s,t) \geq \lambda^Q_G(s,t) + 1$ for all $st \in D$, $Q = V \setminus \{z\}$.

Here the family of tight bisets is

$$\{\mathbb{A} \in \mathcal{V} : |\partial \mathbb{A}| + d_G(\mathbb{A}) = \lambda^Q_G(s,t) \text{ for some } st \in \delta_D(\mathbb{A}), z \notin \partial \mathbb{A}\}$$

Note that this family is symmetric, since if $\mathbb{A}$ is tight then so is $\mathbb{A}^*$. Let us say that U is a **transversal** (hitting set) of a biset family $\mathcal{C}$ or that U is a **$\mathcal{C}$-transversal** if $U \cap C \neq \emptyset$ for all $\mathbb{C} \in \mathcal{C}$. A **fractional $\mathcal{C}$-transversal** is a function $h : V \longrightarrow [0,1]$ such that $h(C) \geq 1$ for all $\mathbb{C} \in \mathcal{C}$. Given node-weights $\{w_v : v \in V\}$ let $\tau_w(\mathcal{C})$ denote the optimal value of a fractional $\mathcal{C}$-transversal, namely

$$\tau_w(\mathcal{C}) = \min\{\sum_{v \in V} w_v x_v : x(C) \geq 1 \ \ \forall \mathbb{C} \in \mathcal{C}, \ x_v \geq 0 \ \forall v \in V\} \ .$$

Let $\Delta(\mathcal{C})$ denote the maximum degree in the (multi)hypergraph $\{C : \mathbb{C} \in \mathcal{C}\}$ formed by the inner parts of the members of $\mathcal{C}$. The problem of finding a minimum weight $\mathcal{C}$-transversal is the Hitting Set problem, and it is a particular case of the Submodular Cover problem. In this case the greedy algorithm computes a solution of weight at most $H(\Delta(\mathcal{C}))\tau_w(\mathcal{C})$. It is easy to see that if J covers a biset family $\mathcal{F}$ then the set of endnodes of J is a $\mathcal{C}(\mathcal{F})$-transversal. The following lemma shows that in our case the inverse is also true.

Lemma 1.11 *Let U be a transversal of a symmetric biset family $\mathcal{F}$ and let J be a star on U with center z such that $z \notin \partial\mathbb{A}$ for every $\mathbb{A} \in \mathcal{F}$. Then J covers $\mathcal{F}$.*

Proof: Let $\mathbb{A} \in \mathcal{F}$. Since $z \notin \partial\mathbb{A}$, $z \in A$ or $z \in V \setminus A^+$. If $z \in V \setminus A^+$, then since U is an $\mathcal{F}$-transversal and since $\mathcal{F}$ is symmetric, there is $u \in U \cap A$. If $z \in A$, then there $u \in U \cap A^*$, by a similar argument. In both cases, the edge zu belongs to the star J and covers $\mathbb{A}$. $\square$

Lemma 1.11 implies that Star Survivable Network Augmentation is equivalent to finding a minimum weight $\mathcal{C}$-transversal, where $\mathcal{C}$ is the family of cores of tight bisets and the weights are $w_v = c_{zv}$ if $v \neq z$ and $w_z = 0$. Thus the problem admits LP-ratio $H(\Delta(\mathcal{C}))$. Hence if Δ_k is a bound on $\Delta(\mathcal{C})$ for all instances with maximum demand k, then we can apply the backward augmentation method as described in Lemma 1.5 and achieve an LP-ratio $H(k) \cdot H(\Delta_k)$. Consequently, the following theorem finishes the proof of Theorem 1.8.

Theorem 1.9 ([41]) *Let $\mathcal{C}$ be the family of cores of the family of tight bisets and let $\gamma = \max\limits_{\mathbb{A} \in \mathcal{C}} |\partial\mathbb{A} \cap T|$. Then $\Delta(\mathcal{C}) \leq 8\gamma^2 + 2\gamma + 1$; thus $\Delta(\mathcal{C}) \leq 8(k-1)^2 + 2(k-1) + 1 = 8k^2 - 14k + 7$.*

Proof: Fix some $v \in V$ and let $\mathcal{C}' = \{\mathbb{C} \in \mathcal{C} : v \in C\}$. For every $\mathbb{C}_i \in \mathcal{C}'$ choose one pair $s_i t_i \in D$ such that $\mathbb{C}_i$ is $s_i t_i$-tight. Let (T', D') be the directed graph induced by the chosen pairs. Also construct an auxiliary directed graph $(\mathcal{C}', \mathcal{J}')$ with edge labels where for every $u \in \{s_i, t_i\} \cap \partial\mathbb{C}_j$ there is an edge $\mathbb{C}_i\mathbb{C}_j \in \mathcal{J}'$ with label u. Note that for any $i \neq j$ we have $v \in C_i \cap C_j$, hence by Corollary 1.1 $\{s_i, t_i\} \cap \partial\mathbb{C}_j \neq \emptyset$ or $\{s_j, t_j\} \cap \partial\mathbb{C}_i \neq \emptyset$. Thus the (simple) underlying graph of $(\mathcal{C}', \mathcal{J})$ is a clique.

Claim 1.2 *For any $u \in T'$, the number of edges in D' incident to u is at most $2\gamma + 1$. Consequently, for any $st \in D'$ there are at most $4\gamma + 1$ edges in D' that are incident to s or to t.*

Proof: Let D'_u be the set of edges in D' incident to u and $\mathcal{C}'_u$ the corresponding set of cores. Let $d = |D'_u|$. For every $s_i t_i \in D'_u$ let $v_i = \{s_i, t_i\} \setminus \{u\}$, namely $v_i = t_i$ if $u = s_i$ and $v_i = s_i$ if $u = t_i$. Consider a pair $s_i t_i, s_j t_j \in D'_u$. Since $v \in C_i \cap C_j$, then by Lemma 1.6 and the minimality of $\mathbb{C}_i, \mathbb{C}_j$ we must have $v_i \in \partial\mathbb{C}_j$

or $v_j \in \partial\mathbb{C}_i$. Now consider the subgraph $(\mathcal{C}'_u, \mathcal{J}'_u)$ of $(\mathcal{C}', \mathcal{J}')$ induced by $\mathcal{C}'_u$. The underlying graph of this subgraph is a clique, hence $|\mathcal{J}'_u| \geq \frac{d(d-1)}{2}$. Consequently, $(\mathcal{C}'_u, \mathcal{J}'_u)$ has a node $\mathbb{C}$ of indegree at least $\frac{d-1}{2}$. Note that each edge in $\mathcal{J}'_u$ that enters $\mathbb{C}$ contributes the node of its label to $\partial\mathbb{C} \cap T'$, and no two edges entering $\mathbb{C}$ have the same label. Hence $\frac{d-1}{2} \leq |\partial\mathbb{C} \cap T'| \leq \gamma$, implying $d \leq 2\gamma + 1$. □

We now finish the proof of Theorem 1.9. Let $\Delta = |\mathcal{C}'|$. The graph $(\mathcal{C}', \mathcal{J}')$ has a node $\mathbb{C}$ of indegree $\geq (\Delta - 1)/2$, since its underlying graph is a clique. Now consider the labels of the arcs entering $\mathbb{C}$. By Claim 1.2, there are at least $(\Delta-1)/(8\gamma+2)$ edges that enter $\mathbb{C}$, such that no two edges have the same label. Each one of these edges contributes a node to $\partial\mathbb{C} \cap T$. Consequently, we must have $(\Delta - 1)/(8\gamma + 2) \leq \gamma$, which implies $\Delta \leq 8\gamma^2 + 2\gamma + 1$. □

1.6.3 Hardness of approximation of the $\{0,1\}$-costs case

We will show that Survivable Network with $\{0,1\}$-costs is almost as hard to approximate as the following variant of the Label-Cover problem:

Min-Rep
Instance: A bipartite graph $H = (A \cup B, I)$ and partitions $\mathcal{A}$ of A and $\mathcal{B}$ of B into parts of equal size.
Objective: Find a minimum size node set $A' \cup B'$, where $A' \subseteq A$ and $B' \subseteq B$, such that for any $A_i \in \mathcal{A}, B_j \in \mathcal{B}$ with $\delta_I(A_i, B_j) \neq \emptyset$ there are $a \in A' \cap A_i, b \in B' \cap B_j$ such that $ab \in I$.

By the Parallel Repetition Theorem [43] (see [28] for details) Min-Rep cannot be approximated within $O(2^{\log^{1-\varepsilon} n})$ unless NP has quasi-polynomial time algorithms. The Min-Rep problem was defined by Kortsarz [28], and was used later by Dodis and Khanna [10] to show a $2^{\log^{1-\varepsilon} n}$-approximation hardness of the Directed Steiner Forest problem. This was extended to high connectivity undirected Survivable Network problems by Kortsarz, Krauthgamer, and Lee [29], and further generalized and simplified in [35]. Here we will describe a slightly more complicated variant, that proves the same hardness already for the $\{0,1\}$-costs case.

Theorem 1.10 (Kortsarz, Krauthgamer, & Lee [29], Nutov [38])

Directed Survivable Network *with* $\{0,1\}$*-costs cannot be approximated within* $2^{\ln^{1-\epsilon} n}$ *for any fixed* $\epsilon > 0$*, unless NP has quasi-polynomial time algorithms.*

Proof: Given an instance of Min-Rep construct an instance $G=(V,E),r$ of Survivable Network with $\{0,1\}$-costs as follows, where edges in E have cost 0. Let $\mathcal{E}=\{ij : A_i \in \mathcal{A}, B_j \in \mathcal{B}, \delta_H(A_i,B_j) \neq \emptyset\}$. Direct the edges of H from A to B. Then the graph $G=(V,E)$ is obtained from H as follows.

1. Add to H: a set $\{a_1,\ldots,a_{|\mathcal{A}|},b_1,\ldots,b_{|\mathcal{B}|}\}$ of $|\mathcal{A}|+|\mathcal{B}|$ nodes, and for every $ij \in \mathcal{E}$ a pair of nodes a_{ij},b_{ij} (so a total number of nodes added to H is $|\mathcal{A}|+|\mathcal{B}|+2|\mathcal{E}|$). Thus

$$V = A \cup B \cup \{a_1,\ldots,a_{|\mathcal{A}|},b_1,\ldots,b_{|\mathcal{B}|}\} \cup \{a_{ij} : ij \in \mathcal{E}\} \cup \{b_{ij} : ij \in \mathcal{E}\} \ .$$

2. For every $ij \in \mathcal{E}$: connect a_{ij} to every node that is not in $\bar{A}_{ij}=A_i \cup B_j \cup \{b_j,b_{ij}\}$, and connect every node that is not in $\bar{B}_{ij}=A_i \cup B_j \cup \{a_i,a_{ij}\}$ to b_{ij}. Thus

$$E = I \cup \{a_{ij}v : ij \in \mathcal{E}, v \in V \setminus \bar{A}_{ij}\} \cup \{ub_{ij} : ij \in \mathcal{E}, u \in V \setminus \bar{B}_{ij}\}.$$

For $ij \in \mathcal{E}$ let $C_{ij}=\{v \in V : a_{ij}v, vb_{ij} \in E\}$. By the construction

$$C_{ij} = V \setminus (\bar{A}_{ij} \cup \bar{B}_{ij}) = V \setminus (A_i \cup B_j \cup \{a_i,b_j,a_{ij},b_{ij}\})$$

Since the sets A_i have the same size and the sets B_j have the same size, the sets C_{ij} are also all of the same size, say $k-1$. Every node in C_{ij} is an internal node of an $a_{ij}b_{ij}$-path of length 2. Let $G_{ij}=G \setminus C_{ij}$ be the subgraph of G induced by $A_i \cup B_j \cup \{a_i,b_j,a_{ij},b_{ij}\}$. By the construction, G_{ij} has no $a_{ij}b_{ij}$-path. Thus $\kappa_G(a_{ij},b_{ij})=k-1$ for every $ij \in \mathcal{E}$. We will ask to increase the connectivity by 1 between pairs $\{a_{ij}b_{ij} : ij \in \mathcal{E}\}$, so the demands are:

$$r(a_{ij},b_{ij}) = k \ \text{ for every } \ ij \in \mathcal{E} \ .$$

For $ij \in \mathcal{E}$ let F_{ij} be the set of edges in F with both endnodes in G_{ij}. Note that for every $ij \in \mathcal{E}$:

(i) G_{ij} has no $a_{ij}b_{ij}$-path.

(ii) $\kappa_{G \cup F}(a_{ij},b_{ij}) \geq k$ iff $G_{ij} \cup F_{ij}$ has an $a_{ij}b_{ij}$-path.

Let us say that a new edge is **proper** if it goes from a_i to A_i for some i, or if it goes from B_j to b_j for some j. Let F be a feasible solution to the constructed Survivable Network instance and let $e \in F$. Assume that e belongs to F_{ij} for some $ij \in \mathcal{E}$, as otherwise e can be removed. We claim if e is non-proper, then e can be replaced by at most two proper edges e',e'' while keeping F feasible, as follows:

- If e goes from $A_i \cup \{a_i, a_{ij}\}$ to $B_j \cup \{b_j, b_{ij}\}$ then $\{e', e''\} = \{a_i a, bb_j\}$ for some $ab \in \delta_I(A_i, B_j)$.
- If $e = a_{ij}a$ for some $a \in A_i$ or $e = bb_{ij}$ for some $b \in B_j$ then $e' = e'' = a_i a$ or $e' = e'' = bb_j$, respectively.
- If $e = a'a''$ for some $a', a'' \in A_i$ or $e = b'b''$ for some $b', b'' \in B_j$, then $\{e', e''\} = \{a_i a', a_i a''\}$ or $\{e', e''\} = \{b' b_j, b'' b_j\}$, respectively.

In each one of the cases, it is not hard to verify that $(F \setminus \{e\}) \cup \{e', e''\}$ is a feasible solution as well. This implies that there exists a proper feasible solution F' with $|F'| \leq 2|F|$.

Note that for every $v \in A \cup B$ corresponds a unique proper edge, namely, $a_i v$ if $v \in A_i$ and vb_j if $v \in B_j$. Thus there is a bijective correspondence between proper edge sets F' and node subsets $A' \cup B'$ of $A \cup B$, where $A' \subseteq A$ and $B' \subseteq B$. Let F' be a proper edge set and let $ij \in \mathcal{E}$. Then there are edges $a_i a, bb_j \in F'_{ij}$ such that $ab \in I$ iff $G_{ij} \cup F'_{ij}$ has an $a_{ij}b_{ij}$-path (the path $(a_{ij}, a_i, a, b, b_j, b_{ij})$ of the lenght 5). Thus F' is a feasible solution for the Survivable Network instance if and only if the set $A' \cup B'$ of end-nodes in $A \cup B$ of the edges in F' is a feasible solution to the Min-Rep instance. Since in the construction $|V| = O(n^2)$, where $n = |A| + |B|$, the theorem follows. □

1.7 Open problems

In this section we list some open problems in the field.

Undirected Survivable Network **with general demands and costs.** Best known ratio: $O(k^3 \ln |T|)$. Does the problem admit ratio that depends on k only, as in the case of rooted demands?

Rooted Survivable Network **with rooted subset uniform demands.** Best known ratio: $O(k \ln k)$ for undirected graphs, and $|D|$ for directed graphs. Does the directed variant admits ratio $n^{1-\epsilon}$ for some $\epsilon > 0$? A related (harder) problem is whether the undirected variant admits ratio $k^{1-\epsilon}$.

Undirected Rooted Survivable Network **with** $\{0,1\}$**-costs.** Best known ratio: $O(\min\{\ln^2 k, \ln n\})$. The problem also has an $\Omega(\ln k)$ approximation threshold when $k = \Theta(n)$. Does the problem admits ratio

$O(\ln k)$? This is open even for Star Survivable Network with rooted demands.

Undirected Survivable Network **with metric costs.** Best known ratio: $O(\ln k)$ for general demands, and 24 for subset uniform and rooted subset uniform demands. Is the general demands case indeed $\Omega(\ln k)$ hard to approximate, or can we get a constant ratio? Can we substantially improve the constant ratio 24 for subset uniform and rooted subset uniform demands?

References

[1] A. Agrawal, P. Klein, and R. Ravi. When trees collide: An approximation algorithm for the generalized steiner problem on networks. *SIAM J. Comput.*, 24(3):440–456, 1995.

[2] P. Berman, A. Bhattacharyya, K. Makarychev, S. Raskhodnikova, and G. Yaroslavtsev. Approximation algorithms for spanner problems and directed steiner forest. *Inf. Comput.*, 222:93–107, 2013.

[3] J. Byrka, F. Grandoni, T. Rothvoß, and L. Sanitá. Steiner tree approximation via iterative randomized rounding. *J. ACM*, 60(1):6, 2013.

[4] M. Charikar, C. Chekuri, T. Cheung, Z. Dai, A. Goel, S. Guha, and M. Li. Approximation algorithms for directed Steiner problems. *J. of Algorithms*, 33:73–91, 1999.

[5] C. Chekuri, G. Even, A. Gupta, and D. Segev. Set connectivity problems in undirected graphs and the directed steiner network problem. *ACM Transactions on Algorithms*, 7(2):18, 2011.

[6] J. Cheriyan, T. Jordán, and Z. Nutov. On rooted node-connectivity problems. *Algorithmica*, 30(3):353–375, 2001.

[7] J. Cheriyan, B. Laekhanukit, G. Naves, and A. Vetta. Approximating rooted steiner networks. In *SODA*, pages 1499–1511, 2012.

[8] J. Cheriyan and A. Vetta. Approximation algorithms for network design with metric costs. *SIAM J. Discrete Math*, 21(3):612–636, 2007.

[9] J. Chuzhoy and S. Khanna. An $O(k^3 \log n)$-approximation algorithms for vertex-connectivity survivable network design. *Theory of Computing*, 8(1):401–413, 2012.

[10] Y. Dodis and S. Khanna. Design networks with bounded pairwise distance. In *STOC*, pages 750–759, 1999.

[11] M. Feldman, G. Kortsarz, and Z. Nutov. Improved approximation algorithms for directed steiner forest. *J. Comput. Syst. Sci.*, 78(1):279–292, 2012.

[12] L. Fleischer, K. Jain, and D. Williamson. Iterative rounding 2-approximation algorithms for minimum-cost vertex connectivity problems. *J. Comput. Syst. Sci*, 72(5):838–867, 2006.

[13] A. Frank. Augmenting graphs to meet edge-connectivity requirements. *SIAM J. on Discrete Math.*, 5(1):25–53, 1992.

[14] A. Frank. Rooted k-connections in digraphs. *Discrete Applied Mathematics*, 157(6):1242–1254, 2009.

[15] A. Frank. *Connections in Combinatorial Optimization*. Oxford University Press, 2011.

[16] A. Frank and T. Jordán. Minimal edge-coverings of pairs of sets. *J. on Comb. Theory B*, 65:73–110, 1995.

[17] A. Frank and T. Kiraly. A survey on covering supermodular functions. In W. Cook, L. Lovsz, and J. Vygen, editors, *Research Trends in Combinatorial Optimization*, pages 87–126. Springer, Berlin, 2009.

[18] A. Frank and E. Tardos. An application of submodular flows. *Linear Algebra and its Applications*, 114/115:329–348, 1989.

[19] T. Fukunaga. Approximating the generalized terminal backup problem via half-integral multiflow relaxation. In *STACS*, pages 316–328, 2015.

[20] 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.

[21] M. Goemans, N. Olver, T. Rothvoß, and R. Zenklusen. Matroids and integrality gaps for hypergraphic steiner tree relaxations. In *STOC*, pages 1161–1176, 2012.

[22] M. Goemans and D. Williamson. A general approximation technique for constrained forest problems. *SIAM J. Comput.*, 24(2):296–317, 1995.

[23] E. Halperin and R. Krauthgamer. Polylogarithmic inapproximability. In *STOC*, pages 585–594, 2003.

[24] C. Helvig, G. Robins, and A. Zelikovsky. Improved approximation scheme for the group steiner problem. *Networks*, 37(1):820, 2001.

[25] K. Jain. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica*, 21(1):39–60, 2001.

[26] S. Khuller. Approximation algorithms for finding highly connected subgraphs. In D. Hochbaum, editor, *Approximation Algorithms for NP-hard problems*, chapter 6, pages 236–265. PWS, 1995.

[27] Z. Király, B. Cosh, and B. Jackson. Local edge-connectivity augmentation in hypergraphs is NP-complete. *Discrete Applied Mathematics*, 158(6):723–727, 2010.

[28] G. Kortsarz. On the hardness of approximating spanners. *Algorithmica*, 30(3):432–450, 2001.

[29] G. Kortsarz, R. Krauthgamer, and J. R. Lee. Hardness of approximation for vertex-connectivity network design problems. *SIAM J. on Computing*, 33(3):704–720, 2004.

[30] G. Kortsarz and Z. Nutov. Approximating minimum cost connectivity problems. In T. Gonzalez, editor, *Approximation Algorithms and Metaheuristics*, chapter 58. Chapman & Hall, 2007.

[31] G. Kortsarz and Z. Nutov. Tight approximation for connectivity augmentation problems. *J. Comput. Syst. Sci.*, 74(5):662–670, 2008.

[32] G. Kortsarz and Z. Nutov. Approximating source location and star survivable network problems. Manuscript, to appear in *Theoretical Computer Science*, 2014.

[33] G. Kortsarz and D. Peleg. Approximating the weight of shallow Steiner trees. *Discrete Applied Math.*, 93(2-3):265–285, 1999.

[34] B. Laekhanukit. Parameters of two-prover-one-round game and the hardness of connectivity problems. In *SODA*, pages 1626–1643, 2014.

[35] Y. Lando and Z. Nutov. Inapproximability of survivable networks. *Theoretical Computer Science*, 410(21-23):2122–2125, 2009.

[36] Z. Nutov. The k-connected subgraph problem. In T. Gonzalez, editor, *Approximation Algorithms and Metaheuristics*, chapter ?? Chapman & Hall, 20??

[37] Z. Nutov. Approximating rooted connectivity augmentation problems. *Algorithmica*, 44(3):213–231, 2006.

[38] Z. Nutov. Approximating connectivity augmentation problems. *ACM Transactions on Algorithms*, 6(1), 2009.

[39] Z. Nutov. A note on rooted survivable networks. *Inf. Process. Lett.*, 109(19):1114–1119, 2009.

[40] Z. Nutov. Approximating minimum cost connectivity problems via uncrossable bifamilies. *ACM Transactions on Algorithms*, 9(1):1, 2012.

[41] Z. Nutov. Approximating node-connectivity augmentation problems. *Algorithmica*, 63(1-2):398–410, 2012.

[42] Z. Nutov. Approximating subset k-connectivity problems. *J. Discrete Algorithms*, 17:51–59, 2012.

[43] R. Raz. A parallel repitition theorem. *SIAM J. on Comput.*, 27(3):763–803, 1998.

[44] A. Schrijver. *Combinatorial Optimization, Polyhedra and Efficiency.* Springer-Verlag,

[45] A. Vetta. Approximating the minimum strongly connected subgraph via a matching lower bound. In *SODA*, pages 417–426, 2001.

[46] L. A. Wolsey. An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica*, 2:385–393, 1982.

[47] A. Zelikovsky. A series of approximation algorithms for the acyclic directed steiner tree problem. *Algorithmica*, 18(1):99–110, 1997.

Index