Approximating Source Location and Star Survivable Network Problems

Guy Kortsarz[a], Zeev Nutov[b]

[a]*Rutgers University, Camden*
[b]*The Open University of Israel*

Abstract

In Source Location (SL) problems the goal is to select a minimum cost source set $S \subseteq V$ such that the connectivity (or flow) $\psi(S, v)$ from S to any node v is at least the demand d_v of v. In many SL problems $\psi(S, v) = d_v$ if $v \in S$, so the demand of nodes selected to S is completely satisfied. In a variant suggested recently by Fukunaga [7], every node v selected to S gets a "bonus" $p_v \leq d_v$, and $\psi(S, v) = p_v + \kappa(S \setminus \{v\}, v)$ if $v \in S$ and $\psi(S, v) = \kappa(S, v)$ otherwise, where $\kappa(S, v)$ is the maximum number of internally disjoint (S, v)-paths. While the approximability of many SL problems was seemingly settled to $\Theta(\ln d(V))$ in [20], for his variant on undirected graphs Fukunaga achieved ratio $O(k \ln k)$, where $k = \max_{v \in V} d_v$ is the maximum demand. We improve this by achieving ratio $\min\{p^* \ln k, k\} \cdot O(\ln k)$ for a more general version with node capacities, where $p^* = \max_{v \in V} p_v$ is the maximum bonus. In particular, for the most natural case $p^* = 1$ we improve the ratio from $O(k \ln k)$ to $O(\ln^2 k)$. To derive these results, we consider a particular case of the Survivable Network (SN) problem when all edges of positive cost form a star. We obtain ratio $O(\min\{\ln n, \ln^2 k\})$ for this variant, improving over the best ratio known for the general case $O(k^3 \ln n)$ of Chuzhoy and Khanna [3].

In addition, we show that directed SL with unit costs is $\Omega(\log n)$-hard to approximate even for $0, 1$ demands, while SL with uniform demands can be solved in polynomial time. Finally, we obtain a logarithmic ratio for a generalization of SL where we also have edge-costs and flow-cost bounds $\{b_v : v \in V\}$, and require that the minimum cost of a flow of value d_v from S to every node v is at most b_v.

Email addresses: `guyk@camden.rutgers.edu` (Guy Kortsarz), `nutov@openu.ac.il` (Zeev Nutov)

Keywords: Source location; Survivable network; Submodular cover

1. Introduction

In Source Location (SL) problems, the goal is to select a minimum cost source set $S \subseteq V$ such that the connectivity from S to any node v is at least the demand d_v of v. Formally, the generic version of this problem is as follows.

Source Location (SL)
Instance: A graph $G = (V, E)$ with node-costs $c = \{c_v : v \in V\}$, *connectivity demands* $d = \{d_v : v \in V\}$, and a *source connectivity function* $\psi : 2^V \times V \to \mathbb{Z}_+$, where $\mathbb{Z}_+$ denotes the set of non-negative integers.
Objective: Find a minimum cost source node set $S \subseteq V$ such that $\psi(S, v) \geq d_v$ for every $v \in V$.

Several source connectivity functions ψ appear in the literature. To avoid considering many cases, we suggest two generic types, that include previous particular cases.

Definition 1.1. *An integer set-function f on a groundset U is submodular if $f(A)+f(B) \geq f(A\cap B)+f(A\cup B)$ for all $A, B \subseteq U$, and f is non-decreasing if $f(A) \leq f(B)$ for all $A \subseteq B \subseteq U$.*

Definition 1.2. *Let $G = (V, E)$ be a graph with node-capacities $\{q_u : u \in V\}$. For $S \subseteq V$ and $v \in V$ the (S, v)-q-*connectivity *$\lambda^q_G(S, v)$ is the maximum number of edge-disjoint paths from $S \setminus \{v\}$ to v in G such that every node $u \in V$ is an internal node in at most q_u paths. Given* connectivity bonuses *$\{p_u \geq q_u : u \in V\}$, the (S, v)-(p, q)-*connectivity *$\lambda^{p;q}_G(S, v)$ is defined by: $\lambda^{p;q}_G(S, v) = p_v + \lambda^q_G(S, v)$ if $v \in S$, and $\lambda^{p;q}_G(S, v) = \lambda^q_G(S, v)$ otherwise.*

We will say that a source connectivity function $\psi(S, v)$ is submodular if for every $v \in V$, the function $f_v(S) = \psi(S, v)$ is submodular and non-decreasing; $\psi(S, v)$ is survivable if it is of the type $\psi(S, v) = \lambda^{p;q}_G(S, v)$. The concept of q-connectivity is essentially "mixed connectivity" (the case $q_u \in \{0, k\}$) introduced by Frank, Ibaraki, and Nagamochi [5], while (p, q)-connectivity combines it with the connectivity function introduced recently by Fukunaga [7] (the case $q \equiv 1$). The case of arbitrary node capacities includes additional connectivity versions compared to [7], e.g., the edge-connectivity case.

It is not hard to see that every survivable source connectivity function $\psi(S,v)$ is submodular (see Section 4), but the inverse is not true in general. This gives only two types of SL problems.

Submodular SL:	The connectivity function $\psi(S,v)$ is submodular.
Survivable SL:	The connectivity function $\psi(S,v)$ is survivable.

We list four source connectivity functions that appear in the literature. All of them are submodular, and three of them are also survivable. Given an SL instance let $k = \max_{v \in V} d_v$ denote the maximum demand, and in the case of Survivable SL let $p^* = \max_{u \in V} p_u$ denote the maximum connectivity bonus and $q^* = \min_{u \in V} q_u$ denote the minimum node capacity. In what follows assume that $1 \leq q_u \leq p_u \leq k$ for all $u \in V$, and thus $1 \leq p^* \leq k$ and $1 \leq q^* \leq k$ holds.

1. λ-SL: $\lambda_G(S,v)$ is the maximum number of pairwise edge-disjoint (S,v)-paths if $v \notin S$ and $\lambda_G(S,v) = \infty$ otherwise.
   This is Survivable SL with $p_u = q_u = k$ for every $u \in V$.
2. κ-SL: $\kappa(S,v)$ is the maximum number of (S,v)-paths no two of which have a common node in $V \setminus (S \cup v)$ if $v \notin S$, and $\kappa(S,v) = \infty$ otherwise.
3. $\hat{\kappa}$-SL: $\hat{\kappa}(S,v)$ is the maximum number of (S,v)-paths no two of which have a common node in $V \setminus \{v\}$ if $v \notin S$, and $\hat{\kappa}(S,v) = \infty$ otherwise.
   This is Survivable SL with $p_u = k$ and $q_u = 1$ for every $u \in V$.
4. κ'-SL: $\kappa'(S,v) = \hat{\kappa}(S,v)$ if $v \notin S$ and $\kappa'(S,v) = p_v + \hat{\kappa}(S \setminus \{v\}, v)$ if $v \in S$.
   This is Survivable SL with $q_u = 1$ for every $u \in V$.

The known approximability status of SL problems with source connectivity functions $\lambda, \kappa, \hat{\kappa}, \kappa'$, is summarized in Table 1; see also a survey in [16]. The approximability of $\lambda, \kappa, \hat{\kappa}$-SL problems was settled to $O(\ln d(V))$ in [20] (where $d(V) = \sum_{v \in V} d_v$), while Fukunaga [7] showed that undirected κ'-SL admits ratio $O(k \ln k)$. We prove the following.

Theorem 1.3. Submodular SL *admits ratio* $O(\ln d(V))$. *Undirected* Survivable SL *admits ratio* $\min\{p^* \ln k, k\} \cdot O(\ln(k/q^*))$*; furthermore, if* $q^* = k$ *(this is the edge-connectivity case) then the ratio is exactly* k.

Theorem 1.3 has several consequences. While ratio $O(\ln(d(V))$ was known for source connectivity functions λ,κ, $\hat{\kappa}$ [20], our proof of a more general result

$\boldsymbol{c},\boldsymbol{d}$	$\boldsymbol{\lambda}$ $(p,q \equiv k)$		$\boldsymbol{\kappa}$	
	Undirected	*Directed*	*Undirected*	*Directed*
GC,GD	$\Theta(\ln d(V))$ [2, 20]	$\Theta(\ln d(V))$ [2, 20]	$\Theta(\ln d(V))$ [2, 20]	$\Theta(\ln d(V))$ [2, 20]
GC,UD	in P [1]	$O(\ln d(V))$ [2]	$O(\ln d(V))$ [2]	$O(\ln d(V))$ [2]
UC,GD	in P [1]	$O(\ln d(V))$ [2]	$O(\ln d(V))$ [2]	$O(\ln d(V))$ [2]
UC,UD	in P [22]	in P [10]	$O(\ln d(V))$ [2]	$O(\ln d(V))$ [2]
	$\hat{\boldsymbol{\kappa}}$ $(p \equiv k,\, q \equiv 1)$		$\boldsymbol{\kappa'}$ $(q \equiv 1)$	
GC,GD	$\Theta(\ln d(V))$ [20] $O(k \ln k)$ [7]	$\Theta(\ln d(V))$ [20]	$O(\ln d(V))$ [7] $O(k \ln k)$ [7]	$O(\ln d(V))$ [7]
GC,UD	in P [17]	in P [17]		
UC,GD	$O(\ln d(V))$ [20] $O(k)$ [9]	$O(\ln d(V))$ [20]		
UC,UD	in P [17]	in P [17]		

Table 1: Previous approximation ratios and lower bounds for SL problems. GC and UC stand for general and uniform costs, GD and UD stand for general and uniform demands, respectively.

is simpler and shorter than the proof of each particular case. For undirected graphs, the second part of Theorem 1.3 implies that Survivable SL problems admit ratio $O(k\ln(k/q^*))$ if $p^* \geq k/\ln k$ (e.g., $p^* = k$ in λ-SL and $\hat{\kappa}$-SL), and ratio $O(p^* \ln k \ln(k/q^*))$ if $p^* < k/\ln k$ (e.g., κ'-SL with $p^* = 1$). In the case of λ-SL we have $q^* = k$ which implies ratio exactly k. We note that ratio k for λ-SL can be achieved by decomposing the problem into k problems with demands in $\{0, \ell\}$, $\ell = 1 \ldots, k$; each of these problems can be solved in polynomial time. However, this algorithm is just a particular case of our algorithm.

Summarizing, we get the following results for connectivity functions λ, κ'.

Corollary 1.4. λ-SL *admits ratio k and* κ'-SL *admits ratio $O(p^* \ln^2 k)$.*

To prove Theorem 1.3, we consider the following known problem.

> Survivable Network (SN)
> *Instance:* A graph $G = (V, E)$ with edge-costs $\{c_e : e \in E\}$ and node capacities $\{q_u : u \in V\}$, and connectivity requirements $r = \{r_{sv} : sv \in D\}$ on a set D of demand edges on V.
> *Objective:* Find a min-cost subgraph G' of G such that $\lambda^q_{G'}(s,v) \geq r_{sv}$ for every $sv \in D$.

Let $k = \max_{sv \in D} r_{sv}$ denote the maximum requirement. For $q \equiv k$ we get the edge-connectivity version which admits ratio 2 due to Jain [11], while for $q \equiv 1$ we get the node-connectivity version. SN admits a folklore ratio $O(|D|)$, and for directed graphs no better ratio is known. Undirected SN admits ratios $O\left(k^3 \log n\right)$ [3] for edge-costs, and $O\left(k^4 \log^2 n\right)$ for node-costs [18, 23], and has an $\Omega(\max\{\{k^{1/4}, |D|^{1/6}\})$ approximation lower bound [14]. We consider the following particular case of SN, studied previously in [13, 7].

Star-SN: the set F of edges in E of positive cost is a star with center a.

The Star-SN problem was defined in [13], where it was shown to admit ratio $O(\ln n)$ for unit edge-costs. The study of this problem in [13] is motivated by the observation that directed SN instances when (V, F) is a complete graph with unit edge costs (so called Connectivity Augmentation problem) can be reduced to Star-SN with a loss of a factor of 2 in the approximation ratio. Fukunaga [7] observed that κ'-SL is a special case of Star-SN. Hence the Star-SN problem is important, as it generalizes several well known problems, and it is also a particular interesting case of the SN problem. Our results for Star-SN, summarized in the following theorem, substantially improve over the best known ratios for SN. These results are of independent interest, as they show that Star-SN admits much better ratios than general SN.

Theorem 1.5. Star-SN *admits approximation ratios $O(\ln n)$ for directed graphs, and $O(\min\{\ln n, \ln k \ln(k/q^*)\})$ for undirected graphs.*

We further study SL problems and prove the following.

Theorem 1.6. *Directed* Survivable SL *for $k = 1$ and unit costs is $\Omega(\log n)$-hard to approximate. Directed/undirected κ'-*SL *with uniform demands and with $p \equiv 1$ can be solved in polynomial time.*

Finally, we consider the following generalization of Survivable SL. Given an instance of Survivable SL and edge-costs $c = \{c_e : e \in E\}$, let $\mu_G^{p,q}(S, v)$ denote the minimum cost of an edge set $F \subseteq E$ such that $\lambda_{(V,F)}^{p,q}(S, v) \geq d_v$, where $\mu_G^{p,q}(S, v) = \infty$ if no such edge set F exists (namely, if $\lambda_G^{p,q}(S, v) < d_v$).

Survivable SL with Flow-Cost Bounds *Instance:* As in Survivable SL, but in addition we are also given edge-costs $\{c_e : e \in E\}$ and flow-cost bounds $\{b_v \leq c(E) : v \in V\}$. *Objective:* As in Survivable SL, with an additional constraint $\mu_G^{p,q}(S, v) \leq b_v$ for every $v \in V$.

Theorem 1.7. Survivable SL with Flow-Cost Bounds *admits approximation ratio* $H(d(V)) + H(nc(E) - b(V))$.

2. Relations between SL and SN problems

To explain the relation between SL and SN problems it would be convenient to consider the augmentation version of the SN problem, with arbitrary connectivity functions and allowing node-costs. Given a function $w = \{w_u : u \in U\}$ on a groundset U and $U' \subseteq U$, let $w(U') = \sum_{u \in U'} w_u$. If w is a cost function on U and I is an edge-set on U, then the cost (or the node-costs) $w(I)$ of I is the cost of the set of the endnodes of I. Formally, we define the problem we need as follows.

Network Augmentation (NA) *Input:* A graph $G = (V, E)$, an edge-set F on V, a cost function c on F or on V, connectivity requirements $r = \{r_{sv} : sv \in D\}$ on a set D of demand edges on V, and a family $\{f_{sv} : 2^F \to \mathbb{Z}_+ : sv \in D\}$ of connectivity functions. *Output:* A min-cost edge-set $I \subseteq F$ such that $f_{sv}(I) \geq r_{sv}$ for every $sv \in D$.

Note that here the connectivity functions $f_{sv}(I)$ differ from the source connectivity functions in SL problems. As in the case of SL problems, we consider two types of NA problems:

Submodular NA: connectivity functions $f_{sv}(I)$ are submodular and non-decreasing. Survivable NA: connectivity functions are $f_{sv}(I) = \lambda^q_{G+I}(s, v)$.

SN is a particular case of Survivable NA when $E = \emptyset$, but for edge-costs the problems are equivalent. Here a survivable connectivity function may not be submodular; indeed, we will obtain a logarithmic ratio for Submodular NA, while Survivable NA has a polynomial approximation threshold. To see this, consider the following simple example: $V = \{s, u, v\}$, $E = \emptyset$, $F = \{su, uv\}$, and $f(I) = f_{sv}(I) = \lambda_{(V,I)}(s, v)$ is just the edge connectivity function. Let $A = \{su\}$ and $B = \{sv\}$. Then $f(A) = f(B) = 0$ and $f(A \cup B) = 1$, and the submodular inequality in Definition 1.1 does not hold. However, we will show that if F is a star, then in the case of directed graphs every survivable connectivity function is submodular.

Let Rooted NA be a particular case of NA when D is a star with center s. As we shall see, SL is equivalent to the node-costs version of the following particular case of both Star-NA and Rooted NA.

Centered-NA: D, F are both stars with a common center s.

Fukunaga [7] made an important observation that κ'-SL is equivalent (via an approximation ratio preserving reduction) to Survivable Centered-NA with *edge-costs* and $q \equiv 1$. Here we further observe the following. For an edge-set/graph J let $\delta_J(X)$ denote the set of edges in J from X to $V \setminus X$.

Lemma 2.1. *For both directed and undirected graphs,* Survivable SL *is equivalent to* Survivable Centered-NA *with node-costs such that* $\delta_G(s) = \emptyset$ *and* $c(s) = 0$.

Proof. Given a Survivable SL instance construct a Survivable Centered-NA instance as follows: add to G a new node s of cost 0, and for every $v \in V$ set $r_{sv} = d_v$ and put p_v edges from s to v into F. Conversely, given a Survivable Centered-NA instance construct a Survivable SL instance as follows. Remove s from G, and for every $v \in V$ set p_v to be the number of edges in F from s to v and $d_v = r_{sv}$. In both directions, it is easy to see that S is a solution to the Survivable SL instance, if, and only if, the edge set I of all edges in F from s to S is a solution to the Survivable Centered-NA instance, and clearly I and S have the same node-cost. □

It is not hard to see that for Survivable Star-NA, approximation ratio ρ for directed graphs implies ratio ρ for undirected graphs. This is achieved by a standard reduction of bidirecting the edges of the undirected instance, removing the directed edges entering the center a, and solving the problem on the obtained directed instance. The same reduction works for Submodular Star-NA problems. We omit the somewhat standard proof details.

The best known ratios for Survivable NA are $O(k^3 \log n)$ for edge-costs [3], and $O\left(k^4 \log^2 n\right)$ for node-costs [18, 23]. The best known ratio for Survivable Rooted NA are $O(k \log k)$ for edge-costs [18] and $O\left(k^2 \log n\right)$ for node-costs [18, 23], and no better ratios were known even for Survivable Centered-NA, see [7] where ratio $O(k \log k)$ for undirected Survivable Centered-NA was deduced in two ways: from the ratio $O(k \log k)$ for Survivable Rooted NA [18], and via iterative rounding. Our results for Star-NA, that imply Theorems 1.3 and 1.5, are summarized in the following three statements.

Let $H(j)$ denote the jth Harmonic number. The following lemma says that Submodular Star-NA problems admit approximation ratio that is logarithmic in terms of certain parameters α and β. These parameters are the maximum total increase (namely, the sum of the increases) in connectivity of all pairs in D as a result of taking a single edge (the parameter α) or a single node (the parameter β) to the solution.

Lemma 2.2. *For directed graphs,* Submodular NA *with edge costs admits ratio $H(\alpha)$, and* Submodular Star-NA *with node costs admits ratio $H(\beta)$, where*

$$\begin{aligned}\alpha &= \max_{e \in F} \sum_{sv \in D} [\min\{f_{sv}(\{e\}), r_{sv}\} - f_{sv}(\emptyset)] \\ \beta &= \max_{z \in V} \sum_{sv \in D} [\min\{f_{sv}(\delta_F(z)), r_{sv}\} - f_{sv}(\emptyset)]\ .\end{aligned}$$

The next lemma says that Survivable Star-NA is a particular case of Submodular Star-NA, and thus the previous lemma can be applied. Moreover, the lemma bounds the parameters α and β as above in terms of the Survivable Star-NA instance ingredients r, D, and F.

Lemma 2.3. *For directed graphs, any* Survivable Star-NA *problem is a* Submodular NA *problem, for which $\alpha \leq |D|$ and $\beta \leq \min\{r(D), p^*|D|\}$ holds, where here p^* denotes the maximum number of parallel edges in F.*

The above two lemmas imply ratio no better than $O(\ln|D|) = O(\ln n)$ for Survivable Star-NA. The next theorem, which is our main technical contribution, says that for undirected graphs we can achieve ratio roughly $O(p^* \ln^2 k)$, which may be much better that $O(\ln n)$ if the maximum requirement $k = \max_{sv \in D} r_{sv}$ and the maximum number p^* of parallel edges in F are small.

Theorem 2.4. *Undirected* Survivable Star-NA *admits ratio $O(\ln k \ln(k/q^*))$ for edge-costs and $\min\{p^* \ln k, k\} \cdot O(\ln(k/q^*))$ for node-costs; furthermore, in the case of node costs and $q^* = k$ the ratio is exactly k.*

The above three statements imply Theorem 1.5; they also imply Theorem 1.3, when combined with Lemma 2.1.

Our ratios for Star-NA and SL are summarized Table 2.

We briefly mention the techniques we use to prove these statements. Lemma 2.2 is essentially an easy application of the greedy algorithm of Wolsey [24] for the Submodular Cover problem. Parts of Lemma 2.3 were

	submodular		survivable	
	directed	*undirected*	*directed*	*undirected*
Star-NA (edge-costs)	$H(\alpha)$	$H(\alpha)$	$H(\|D\|)$	$H(\|D\|)$ $O(\ln k \ln(k/q^*))$
Star-NA (node-costs)	$H(\beta)$	$H(\beta)$	$H(\min\{r(D), p^*\|D\|\})$	$H(\min\{r(D), p^*\|D\|\})$ $\min\{p^* \ln k, k\} \cdot O(\ln(k/q^*))$
SL	$H(d(V))$	$H(d(V))$	$H(\min\{d(V), p^*\|V\|\})$	$H(\min\{d(V), p^*\|V\|\})$ $\min\{p^* \ln k, k\} \cdot O(\ln(k/q^*))$

Table 2: Approximation ratios for Star-NA and SL problems proved in this paper.

implicitly proved in [13], but our proof is both more general and substantially simpler. Our main technical contribution is Theorem 2.4. To prove this theorem, we consider the augmentation version of Survivable Star-NA with edge-costs where the goal is to increase the connectivity by one between the pairs in D. Using LP-scaling we show that ratio ρ for the augmentation version implies ratio $O(\rho \ln k)$ for the edge-costs version of the general problems, and ratio $\min\{p^* \ln k, k\} \cdot O(\rho)$ for the node-costs version. Then we design an $O(\ln(k/q^*))$-approximation algorithm for the augmentation version. This is achieved by formulating the augmentation problem as a Biset-Family Edge-Cover problem, reducing the later problem to the problem of finding a minimum cost vertex cover in a hypergraph, and using a theorem from [19] to show that the maximum degree in the obtained hypergraph is $O\left((k/q^*)^2\right)$.

3. Directed Submodular NA problems (Lemma 2.2)

All graphs in this and the next sections are assumed to be directed. To prove Lemma 2.2 we use a result due to Wolsey [24] about a performance of a greedy algorithm for submodular covering problems. In a generic covering problem we are given by a value oracle two set functions on a groundset U: a cost-function $c : 2^U \to \mathbb{R}$ and a progress function $g : 2^U \to \mathbb{Z}$. The goal is to find $S \subseteq U$ of minimum cost such that $g(S) = g(U)$. The Submodular Cover problem is a special case when the function g is submodular and non-decreasing, and $c(S) = \sum_{v \in S} c(v)$ for some $c : U \to \mathbb{R}^+$. Wolsey [24] proved that then, the greedy algorithm, that starts with $S = \emptyset$ and as long as $g(S) < g(U)$ repeatedely adds to A an element $u \in U \setminus S$ with maximum $\frac{g(S \cup \{u\}) - g(S)}{c_u}$, has approximation ratio $H\left(\max_{u \in U} g(\{u\}) - g(\emptyset)\right)$.

We start with the case of edge-costs. Then the function g is defined in the same way as in [13, 20]: $U = F$ and for $I \subseteq F$

$$g(I) = \sum_{sv \in D} \min\{f_{sv}(I), r_{sv}\}.$$

It is not hard to verify that g is non-decreasing, and that I is a feasible solution to an NA instance if and only if $g(I) = g(F) = r(D)$. Also, for any $e \in F$

$$g(\{e\}) - g(\emptyset) = \sum_{sv \in D} \left[\min\{f_{sv}(\{e\}), r_{sv}\} - f_{sv}(\emptyset)\right] .$$

We show that g is submodular. It is known (c.f. [21]) that if h is submodular, then $\min\{h, r\}$ is submodular for any constant r. Thus the function $h_{sv}(I) = \min\{f_{sv}(I), r_{sv}\}$ is submodular. As a sum of submodular functions is also submodular, we obtain that g is submodular.

Now let us consider node-costs. For $S \subseteq V$ let F_S denote the set of edges in F from a to S, and let $f'_{sv}(S) = f_{sv}(F_S)$. We have $U = V$ and for $S \subseteq V$ let

$$g'(S) = \sum_{sv \in D} \min\{f'_{sv}(S), r_{sv}\} .$$

As in the edge-costs case, it is not hard to verify that g' is non-decreasing and that S is a feasible solution to an NA instance if and only if $g'(S) = g'(V) = r(D)$. Also, for any $z \in V$

$$g'(\{z\}) - g'(\emptyset) = \sum_{sv \in D} \left[\min\{f_{sv}(\delta_F(z)), r_{sv}\} - f_{sv}(\emptyset)\right] .$$

We show that g' is submodular. We claim that the submodularity of $f(I)$ implies that $f'(S)$ is submodular. This is not true in general, but holds if F is a star, and hence for Star-NA instances. More precisely, it is not hard to verify the following statement, that finishes the proof of Lemma 2.2.

Lemma 3.1. *Let (V, F) be a graph and let f be a submodular set function on F. If F is a star with center a, then the set function $f'(S) = f(F_S)$ defined on $V \setminus \{a\}$ is also submodular.*

Proof. Let $A, B \subseteq V \setminus \{a\}$. It is easy to see that since F is a star then

$$F_A \cap F_B = F_{A \cap B} \qquad F_A \cup F_B = F_{A \cup B} .$$

Thus by the definition of f' and the submodularity of f we have

$$\begin{aligned} f'(A) + f'(B) &= f(F_A) + f(F_B) \geq f(F_A \cap F_B) + f(F_A \cup F_B) \\ &= f(F_{A\cap B}) + f(F_{A\cup B}) = f'(A \cap B) + f'(A \cup B) \ . \end{aligned}$$

□

4. Survivable Star-NA is a Submodular NA problem (Lemma 2.3)

We start by showing that in the case of edge-costs, directed Survivable Star-NA is a particular case of Submodular NA. Let $s, v \in V$ and let $f : 2^F \to \mathbb{Z}$ be defined by $f(I) = \lambda^q_{G+I}(s, v)$, $I \subseteq F$. It is easy to see that f is non-decreasing and we prove that if F is a star then f is submodular. For that, we use the following known characterization of submodularity, c.f. [21]:
A set-function f on F is submodular if, and only if

$$f(I_0 \cup \{e\}) + f(I_0 \cup \{e'\}) \geq f(I_0) + f(I_0 \cup \{e, e'\}) \quad \forall I_0 \subset F, e, e' \in F \setminus I_0$$

Let us fix $I_0 \subseteq F$. Revising our notation to $G \leftarrow G + I_0$, $F \leftarrow F \setminus I_0$, and denoting $h(I) = f(I_0 \cup I) - f(I_0)$, we get that f is submodular if, and only if

$$h(\{e\}) + h(\{e'\}) \geq h(\{e, e'\}) \quad \forall e, e' \in F \ .$$

In our setting, F is a star and $h(I) = \lambda^q_{G+I}(s, v) - \lambda^q_G(s, v)$ is the increase in the (s, v)-q-connectivity as a result of adding I to G. Thus $0 \leq h(I) \leq |I|$ for any $I \subseteq F$, so $0 \leq h(\{e, e'\}) \leq 2$. If $h(\{e, e'\}) = 0$, then we are done; if $h(\{e, e'\}) = 1$, then we need to show that $h(\{e\}) = 1$ or $h(\{e'\}) = 1$; and if $h(\{e, e'\}) = 2$, then we need to show that $h(\{e\}) = 1$ and $h(\{e'\}) = 1$. We prove the following general statement, that implies the above; it says that if an augmenting edge set I is a star that increases the st-connectivity by h, then there are h edges in I that cover all minimum st-cuts, and thus each of these edges increases the st-connectivity by 1.

Lemma 4.1. *Let $G = (V, E)$ be a directed graph with node capacities $\{q_v : v \in V\}$, let I be a set of edges on V disjoint to E such that I is a star with center a, let $s, t \in V$, and let $h = \lambda^q_{G+I}(s, t) - \lambda^q_G(s, t)$. Then there is $J \subseteq I$ of size $|J| \geq h$ such that $\lambda^q_{G+\{e\}}(s, t) = \lambda^q_G(s, t) + 1$ for every $e \in J$.*

Proof. Since we consider directed graphs, it is sufficient to prove the lemma for the case of edge-connectivity. For that, apply the following standard reduction that eliminates node capacities: replace every $v \in V \setminus \{s,t\}$ by two nodes v^{in}, v^{out} connected by q_v parallel edges from v^{in} to v^{out} and replace every $uv \in E \cup I$ by an edge from u^{out} to v^{in}. Hence we will prove the lemma for the edge connectivity function λ.

Let us say that $S \subseteq V$ is *tight* if $s \in S$, $t \notin S$, and $|\delta_G(S)| = \lambda_G(s,t)$, namely, if $\delta_G(S)$ is a minimum st-cut. Let $\mathcal{F}$ be the family of tight sets. By Menger's Theorem $\mathcal{F}$ is non-empty. It is also known that $\mathcal{F}$ is a ring family, namely, the intersection of all the sets in $\mathcal{F}$ is nonempty, and if $X, Y \in \mathcal{F}$ then $X \cap Y, X \cup Y \in \mathcal{F}$. Thus $\mathcal{F}$ has a unique inclusion-minimal set $S_{\min}$ and a unique inclusion-maximal set $S_{\max}$, and $S_{\min} \subseteq S_{\max}$ holds.

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

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

$$\lambda_{G+I}(s,t) \leq \lambda_G(s,t) + |\delta_I(S_{\min})| \qquad \lambda_{G+I}(s,t) \leq \lambda_G(s,t) + |\delta_I(S_{\max})|$$

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

Note that Lemma 4.1 does not hold if I is an arbitrary edge set. To see this, consider the following example (this is the example given at the beginning of Section 2): $V = \{s,u,t\}$, $E = \emptyset$, and $I = \{su, ut\}$. Then $h = \lambda_{G+I}(s,t) - \lambda_G(s,t) = 1 - 0 = 1$, but $\lambda_{G+\{e\}}(s,t) = 0$ for every $e \in I$.

We now bound the parameters α and β. The bound $\beta \leq r(D)$ is obvious, while the other bounds on α and β follow from the simple observation that for any $s, v \in V$, the set-function on F defined by $f(I) = \lambda^q_{G+I}(s,v)$ has the following properties: $f(\{e\}) \leq 1$ for any $e \in F$ and $f(\delta_F(z)) \leq |\delta_F(z)| \leq p^*$ for any $z \in V$.

The proof of Lemma 2.3 is now complete.

5. Undirected Survivable Star-NA (Theorem 2.4)

All graphs in this and the next section are assumed to be undirected. We start by considering the edge-costs case, and then will show that it implies the node-costs case by reductions. We need several definitions.

Definition 5.1. *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* $\partial\mathbb{A} = A^+ \setminus A$ *is the* boundary *of* $\mathbb{A}$*. An edge* e *covers a biset* $\mathbb{A}$ *if it has one endnode in* A *and the other in* $V \setminus A^+$*. For a biset* $\mathbb{A}$ *and an edge-set/graph* J *let* $\delta_J(\mathbb{A})$ *denote the set of edges in* J *covering* $\mathbb{A}$*.*

Given an instance of Survivable NA and a biset $\mathbb{A}$ on V, let the requirement of $\mathbb{A}$ be $r(\mathbb{A}) = \max\{r_{uv} : uv \in \delta_D(\mathbb{A})\}$ if $\delta_D(\mathbb{A}) \neq \emptyset$ and $r(\mathbb{A}) = 0$ otherwise. By the q-connectivity version of Menger's Theorem (c.f. [12]), $I \subseteq F$ is a feasible solution to an Survivable NA instance if, and only if, $|\delta_I(\mathbb{A})| \geq h(\mathbb{A})$ for every bisets $\mathbb{A}$ on V, where h is a biset-function defined by

$$h(\mathbb{A}) = \max\{r(\mathbb{A}) - (q(\partial\mathbb{A}) + |\delta_G(\mathbb{A})|), 0\} \tag{1}$$

Let $\mathcal{P}_h$ denote the polytope of "fractional edge-covers" of h, namely,

$$\mathcal{P}_h = \left\{x \in \mathbb{R}^F : x\left(\delta_F(\mathbb{A})\right) \geq h(\mathbb{A}) \ \forall \text{ biset } \mathbb{A} \text{ on } V,\ 0 \leq x_e \leq 1 \ \forall e \in F\right\} .$$

Let $\tau(h)$ denote the optimal value of a standard LP-relaxation for edge covering h by a minimum cost edge set, namely, $\tau(h) = \min\left\{\sum_{e \in F} c_e x_e : x \in \mathcal{P}_h\right\}$.

As an intermediate problem, we consider Survivable NA instances when we seek to increase the connectivity by 1 for every $uv \in D$, namely, when $r_{uv} = \lambda^q_G(u,v) + 1$ for all $uv \in D$.

D-Survivable NA (the edge-costs version) *Input:* A graph $G = (V, E)$ with node-capacities $\{q_v : v \in V\}$, an edge set F on V, a cost function c on F, and a set D of demand edges on V. *Output:* Find a min-cost edge-set $I \subseteq E$ such that $\lambda^q_{G+I}(u,v) \geq \lambda^q_G(u,v) + 1$ for all $uv \in D$.

Given a D-Survivable NA instance, let us say that a biset $\mathbb{A}$ is *tight* if $h(\mathbb{A}) = 1$, where h is defined by (1). The D-Survivable NA problem is equivalent to the problem of finding a minimum cost edge-cover of the biset family $\mathcal{F} = \{\mathbb{A} : h(\mathbb{A}) = 1\}$ of tight bisets. Thus the following generic problem includes the D-Survivable NA problem.

Biset-Family Edge-Cover *Input:* A graph (V, F) with edge-costs and a biset family $\mathcal{F}$ on V. *Output:* Find a min-cost $\mathcal{F}$-cover $I \subseteq F$.

For a biset-family $\mathcal{F}$ let $\tau(\mathcal{F})$ denote the optimal value of a standard LP-relaxation for edge covering $\mathcal{F}$ by a minimum cost edge set, namely, $\tau(\mathcal{F}) = \tau(h)$ where h is defined by $h(\mathbb{A}) = 1$ if $\mathbb{A} \in \mathcal{F}$ and $h(\mathbb{A}) = 0$ otherwise.

The following statement considers the approximation factor invoked by applying the so called "backward augmentation" method due to [8]. Some parts of this staement are known, but we will provide a proof for completeness of exposition.

Proposition 5.2. *Suppose that D-*Survivable Star-NA *with edge-costs admits a polynomial time algorithm that computes a solution of cost at most $\rho(k)\tau(\mathcal{F})$, where $\mathcal{F}$ is the family of tight bisets. Then* Survivable Star-NA *admits a polynomial time algorithm that computes a solution I such that:*

- *For edge-costs, $c(I) \leq \tau(h) \cdot \sum_{\ell=1}^{k} \frac{\rho(\ell)}{k-\ell+1}$, where h is defined by (1).*
- *For node-costs, $c(I) \leq \mathsf{opt} \cdot \sum_{\ell=1}^{k} \rho(\ell) \cdot \min\left\{\frac{p^*}{k-\ell+1}, 1\right\}$.*

Proof. We start with the edge-costs case. Consider the following sequential algorithm. Start with $I = \emptyset$. At iteration $\ell = 1, \ldots, k$, add to I and remove from F an edge-set $I_\ell \subseteq F$ that increases by 1 the q-connectivity of $G + I$ on the set of demands

$$D_\ell = \{sv : \lambda^q_{G+I}(s,v) = r(s,v) - k + \ell - 1, sv \in D\} \ ,$$

by covering the corresponding biset-family $\mathcal{F}_\ell$ using the ρ-approximation algorithm. After iteration ℓ, we have $\lambda^q_{G+I}(s,v) \geq r(s,v) - k + \ell$ for all $sv \in D$. Consequently, after k iterations $\lambda^q_{G+I}(s,v) \geq r(s,v)$ holds for all $sv \in D$, thus the computed solution is feasible. The approximation ratio follows from the following two observations.

(i) $c(I_\ell) \leq \rho(\ell) \cdot \tau(\mathcal{F}_\ell)$. This is so since $\lambda(s,v) \leq \ell - 1$ for every $sv \in D_\ell$, hence the maximum requirement at iteration ℓ is at most ℓ.

(ii) $\tau(\mathcal{F}_\ell) \leq \frac{\tau(h)}{k-\ell+1}$. To see this, note that if $\mathbb{A} \in \mathcal{F}_\ell$ and $x \in \mathcal{P}_h$ then $x(\delta(\mathbb{A})) \geq k - \ell + 1$, by Menger's Theorem. Thus $x/(k - \ell + 1)$ is a feasible solution for the LP-relaxation for edge-covering $\mathcal{F}_\ell$, of value $c \cdot x/(k - \ell + 1)$.

Consequently, $c(I) = \sum_{\ell=1}^{k} c(I_\ell) \leq \sum_{\ell=1}^{k} \rho(\ell) \cdot \frac{\tau(h)}{k-\ell+1} = \tau(h) \cdot \sum_{\ell=1}^{k} \frac{\rho(\ell)}{k-\ell+1}$.

Now let us consider the case of node-costs. Then we convert node-costs into edge-costs by assigning to every edge $e = av$ the cost $c'(e) = c(v)$. Let opt' denote the optimal solution value of the edge-costs instance obtained. Clearly, $\mathsf{opt} \leq \mathsf{opt}' \leq p^* \cdot \mathsf{opt}$. Note that any inclusion minimal solution to a D-Survivable NA instance has no parallel edges. This implies that $c(I_\ell) \leq \rho(\ell) \cdot \mathsf{opt}$ and that $c(I_\ell) = c'(I_\ell)$. The latter implies $c(I_\ell) = c'(I_\ell) \leq \rho(\ell) \cdot \frac{\mathsf{opt}'}{k-\ell+1} \leq \rho(\ell) \cdot \mathsf{opt} \cdot \frac{p^*}{k-\ell+1}$, and the statement for the node-costs case follows. □

In the next section we prove the following theorem, that together with Proposition 5.2 finishes the proof of Theorem 2.4.

Theorem 5.3. *Undirected D-Survivable Star-NA with edge-costs admits a polynomial time algorithm that computes a solution I of cost $\tau(\mathcal{F}) \cdot O(\ln(k/q^*))$. Furthermore, if D is a star then $c(I) \leq \tau(\mathcal{F}) \cdot H\left(2\left\lfloor \frac{k-1}{q^*} \right\rfloor + 1\right)$.*

6. Proof of Theorem 5.3

Recall that D-Survivable NA reduces to Biset-Family Edge-Cover with $\mathcal{F}$ being the family of tight bisets; in the case of rooted requirements, when D is a star with center s, it is sufficient to cover the biset-family

$$\mathcal{F}^s = \{\mathbb{A} \in \mathcal{F} : s \in V \setminus A^+\} .$$

Biset-families arising from Survivable NA instances have some special properties, that are summarized in the following definitions.

Definition 6.1. *The intersection and the union of two bisets $\mathbb{A}, \mathbb{B}$ is 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 write $\mathbb{A} \subseteq \mathbb{B}$ and say that $\mathbb{B}$* contains *$\mathbb{A}$ if $A \subseteq B$ and $A^+ \subseteq B^+$. Let $\mathcal{C}_{\mathcal{F}}$ denote the inclusion-minimal bisets in $\mathcal{F}$.*

Definition 6.2. *Two bisets $\mathbb{A}, \mathbb{B}$ covered by an edge-set D are D-*independent *if for any $xx', yy' \in D$ such that xx' covers $\mathbb{A}$ and yy' covers $\mathbb{B}$, $\{x, x'\} \cap \partial\mathbb{B} \neq \emptyset$ or $\{y, y'\} \cap \partial\mathbb{A} \neq \emptyset$; otherwise, $\mathbb{A}, \mathbb{B}$ are D-*dependent. *We say that a biset family $\mathcal{F}$ is D-*uncrossable *if D covers $\mathcal{F}$ and if for any D-dependent $\mathbb{A}, \mathbb{B} \in \mathcal{F}$ the following holds:*

$$\mathbb{A} \cap \mathbb{B}, \mathbb{A} \cup \mathbb{B} \in \mathcal{F} \text{ or } \mathbb{A} \setminus \mathbb{B}, \mathbb{B} \setminus \mathbb{A} \in \mathcal{F} . \tag{2}$$

Similarly, given a set $T \subseteq V$ of terminals, we say that $\mathbb{A}, \mathbb{B}$ are T-independent *if $A \cap T \subseteq \partial\mathbb{B}$ or if $B \cap T \subseteq \partial\mathbb{A}$, and $\mathbb{A}, \mathbb{B}$ are T*-dependent *otherwise. We say that $\mathcal{F}$ is T-uncrossable if T covers the set-family of the inner parts of $\mathcal{F}$, and if (2) holds for any T-dependent $\mathbb{A}, \mathbb{B} \in \mathcal{F}$.*

A biset-family $\mathcal{F}$ is symmetric if $\mathbb{A} \in \mathcal{F}$ implies $(V \setminus A^+, V \setminus A) \in \mathcal{F}$. We will use the the following statement, that was implicitly proved in [19].

Lemma 6.3 ([19])**.** *The family $\mathcal{F}$ of tight bisets is symmetric and D-uncrossable; if D is a star with leaf-set T then $\{\mathbb{A} \in \mathcal{F} : s \notin A^+\}$ is T-uncrossable.*

For a biset-family $\mathcal{F}$ let $\gamma_\mathcal{F} = \max\{|\partial\mathbb{A}| : \mathbb{A} \in \mathcal{F}\}$ denote the maximum size of the boundary of a biset in $\mathcal{F}$. Note that if $\mathcal{F}$ is the family of tight bisets then $\gamma_\mathcal{F} \leq (k-1)/q^*$. Given an instance of Biset-Family Edge-Cover, we will assume that the family $\mathcal{C}$ of the inclusion members of $\mathcal{F}$ can be computed in polynomial time. We note that for $\mathcal{F}$ being the family of tight bisets, this step can be implemented in polynomial time, c.f. [19]. Under this assumption, we prove the following generalization of Theorem 5.3.

Theorem 6.4. *For edge/node-costs,* Biset-Family Edge-Cover *with F being a star admits a polynomial time algorithm that computes a cover I of $\mathcal{F}$ such that:*

(i) *$c(I) \leq H\left((4\gamma_\mathcal{C}+1)^2\right) \cdot \tau(\mathcal{F})$ if $\mathcal{F}$ is symmetric and D-uncrossable.*

(ii) *$c(I) \leq H(2\gamma_\mathcal{C}+1) \cdot \tau(\mathcal{F})$ if $\mathcal{F}$ is T-uncrossable and $a \in V \setminus X^+$ for all $\mathbb{A} \in \mathcal{F}$.*

In the rest of this section we prove Theorem 6.4.

Definition 6.5. *A set $U \subseteq V$ of nodes is a $\mathcal{C}$-transversal of a hypergraph (set-family) $\mathcal{C}$ on V if U intersects every set in $\mathcal{C}$; if $\mathcal{C}$ is a biset-family then U should intersect the inner part of every member of $\mathcal{C}$. Given node costs $\{c_v : v \in V\}$, let $t^*(\mathcal{C})$ denote the minimum value of a fractional $\mathcal{C}$-transversal, namely:*

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

In [19], the following is proved.

Theorem 6.6 ([19])**.** *let $\mathcal{C}$ be the family of the inclusion members of a biset family $\mathcal{F}$. Then the maximum degree in the hypergraph $\{C : \mathbb{C} \in \mathcal{C}\}$ is at most:*

(i) $(4\gamma_{\mathcal{C}} + 1)^2$ *if $\mathcal{F}$ is D-uncrossable.*

(ii) $2\gamma_{\mathcal{C}} + 1$ *if $\mathcal{F}$ is T-uncrossable.*

Given a hypergraph $(V, \mathcal{C})$ with node-costs, the greedy algorithm computes in polynomial time a $\mathcal{C}$-transversal $U \subseteq V$ of cost $c(U) \leq H(\Delta(\mathcal{C}))t^*(\mathcal{C})$, where $\Delta(\mathcal{C})$ is the maximum degree of the hypergraph (c.f. [15]).

The following statement is obvious.

Lemma 6.7. *If an edge-set I covers a biset-family $\mathcal{F}$ then the set of endnodes of I is a transversal of $\mathcal{F}$.*

Lemma 6.8. *Let $\mathcal{F}$ be a biset family on V and I a star with center a on a transversal $U \subseteq V$ of $\mathcal{F}$. Then I covers $\mathcal{F}$ in each one of the following cases.*

(i) *$\mathcal{F}$ is symmetric and $a \notin \Gamma(\mathbb{A})$ for all $\mathbb{A} \in \mathcal{F}$.*

(ii) *$a \in V \setminus A^+$ for all $\mathbb{A} \in \mathcal{F}$.*

Proof. Let $\mathbb{A} \in \mathcal{F}$. Then $a \in A$ or $a \in V \setminus A^+$. If $a \in V \setminus A^+$, then since U is a transversal of $\mathcal{C}$, there is $u \in U \cap A$. If $a \in A$, then if $\mathcal{F}$ is symmetric, then there $u \in U \cap (V \setminus X^+)$. In both cases, there is an edge $au \in I$, and this edge covers $\mathbb{A}$. $\square$

The algorithm as in Theorem 6.4, for both edge-costs and node-costs is as follows, where in the case of node-costs we may assume that the cost of a is zero.

1. For every $v \in V \setminus \{a\}$, let e_v be the minimum-cost edge incident to v, and in the case of edge-costs define node-costs $c_v = \min_{e \in \delta_F(v)} c_e$ if $\delta_F(v) \neq \emptyset$, and $c_v = \infty$ otherwise.
2. Let $\mathcal{C}$ be the family of the inclusion members of $\mathcal{F}$. With node-costs $\{c_v : v \in V\}$, compute a transversal U of $\mathcal{C}$ of cost $c(U) \leq H(\Delta(\mathcal{C}))t^*(\mathcal{C})$.
3. Return $I = \{e_v : v \in U\}$.

The solution computed is feasible by Lemma 6.8. The approximation ratio follows from Theorem 6.6 and Lemma 6.7.

This concludes the proof of Theorem 5.3, and thus also the proof of Theorem 2.4 is now complete.

7. Proof of Theorem 1.6

Note that in the reduction in Lemma 2.1 we have the following.

- Uniform demands $d_v = k$ for all $v \in V$ in Survivable SL correspond to requirements $r_{sv} = k$ for all $v \in V \setminus \{s\}$ in Survivable Centered-NA.
- κ'-SL with $p \equiv 1$ corresponds to Survivable Centered-NA with edge costs.
- Unit node-costs in Survivable SL correspond to unit node-costs in Survivable Centered-NA.

Directed Rooted Survivable NA with edge-costs and requirements $r_{sv} = k$ for all $v \in V \setminus \{s\}$ can be solved in polynomial time [6]; this implies that also *undirected* Survivable Centered-NA with edge-costs and requirements $r_{sv} = k$ for all $v \in V \setminus \{s\}$ can be solved in polynomial time. Thus the same holds for κ'-SL with $p \equiv 1$ and uniform demands.

Frank [4] showed that *directed* Survivable Centered-NA with $\delta_G(s) = \emptyset$ and $k = 1$ is NP-hard. Using a slight modification of his reduction we can show that the problem is in fact Set-Cover hard to approximate, and thus is $\Omega(\log n)$-hard to approximate. Given an instance of Set-Cover, where a family A of sets needs to cover a set B of elements, construct the corresponding directed bipartite graph $G' = (A \cup B, E')$, by putting an edge from every set to each element it contains. The graph $G = (V, E)$ is obtained from G' by adding M copies of B, connecting A to each copy in the same way as to B, and adding a new node s. Let $F = \{sv : v \in V\}$, $c(e) = 1$ for every $e \in F$, and $r_{sv} = 0$ if $v \in A$ and $r_{sv} = 1$ otherwise. It is easy to see that if $I \subseteq F$ is a feasible solution to the obtained Survivable Centered-NA instance, then either I corresponds to a feasible solution to the Set-Cover instance, or $|I| \geq M$. The $\Omega(\log n)$-hardness follows for M large enough, say $|M| = (|A| + |B|)^2$, and $|A| = |B|$. Since for $k = 1$ all connectivity functions of Survivable NA are equivalent, we get $\Omega(\log n)$ hardness for directed Survivable NA with $k = 1$ and unit costs.

8. Survivable SL with Flow-Cost Bounds (Theorem 1.7)

Survivable SL with Flow-Cost Bounds is a special case of the following generalization of the Submodular Cover problem, where we have two progress functions:

$$f(S) = \sum_{v \in V} \min\{\lambda_G^{p;q}(S, v), d_v\} \quad \text{and} \quad g(S) = \sum_{v \in V} \min\{-\mu_G^{p;q}(S, v), -b_v\}. \quad (3)$$

It is easy to see that S is a feasible solution to Submodular SL with Flow-Cost Bounds if and only if both

$$f(S) = f(V) = \sum_{v \in V} d_v \text{ and } g(S) = g(V) = -\sum_{v \in V} b_v \ .$$

For f, g defined by (3) we have $\max_{u \in U} f(\{u\}) - f(\emptyset) \leq d(V)$, but note that $\max_{u \in U} g(\{u\}) - g(\emptyset) = \infty$ may hold. The function f is submodular since for any $v \in V$ the function $f_v(S) = \lambda_G^{p,q}(S, v)$ is submodular, as can be deduced from Lemmas 3.1 and 2.3. The function g is submodular since for any $v \in V$ the function $g_v(S) = \lambda_G^{p,q}(S, v)$ is submodular; this is proved in [2] for the case of edge-connectivity, and the proof for (p, q)-connectivity is similar. Also, both functions are non-decreasing and admit a polynomial time value oracle.

Double Submodular Cover
Instance: A groundset V with costs $\{c_v : v \in V\}$ and submodular non-decreasing functions $f : 2^V \to \mathbb{Z}$ and $g : 2^V \to \mathbb{Z} \cup \{-\infty\}$ given by a value oracle.
Objective: Find $S \subseteq V$ of minimum cost such that $f(S) = f(V)$ and $g(S) = g(V)$.

There are several natural approaches to solve the Double Submodular Cover problem using the greedy algorithm of Wolsey [24]. One is to apply the greedy algorithm with the function $f+g$. Another possibility is to solve two instances of Submodular Cover, one with function f and the other with function g, returning the union of the solutions S_f and S_g computed. However, in both cases the ratio may be unbounded if $g(\emptyset) = -\infty$, which may happen for g defined by (3).

The idea is to compute S_f and then to compute S_g for the residual problem. Note that for f, g defined by (3) we have the following property: if $f(S_f) = f(U)$ then $g(S) \geq -n \cdot c(E)$ for any $S \supseteq S_f$. Therefore, the following approach works. We take the set S_f into our solution, and consider the residual Submodular Cover problem with groundset $V \setminus S_f$ and the set function $h(S) = g(S_f \cup S)$, $S \subseteq V \setminus S_f$. The function h is submodular if g is. Note that for g defined by (3), $\max_{u \in U} h(\{u\}) - h(\emptyset) \leq n \cdot c(E) - b(V)$, and we get approximation ratio

$$H\left(\max_{v \in V} f(\{v\}) - f(\emptyset)\right) + H\left(\max_{v \in V} h(\{v\}) - h(\emptyset)\right) \leq H(d(V)) + H(nc(E) - b(V)) \ .$$

Clearly, the approach described can be generalized to the case when we have many non-decreasing submodular functions, under the assumption that there exists an ordering $f_1, f_2, \ldots$ of the functions such that for any i, if $f_j(S) = f(U)$ for every $j \leq i$, then $f_{j+1}(S') \neq -\infty$ for any $S' \supseteq S$.

Acknowledgment: The second author thank Takuro Fukunaga and an anonymous referee for many useful comments.

References

[1] K. Arata, S. Iwata, K. Makino, and S. Fujishige. Locating sources to meet flow demands in undirected networks. *J. Algorithms*, 42:54–68, 2002.

[2] J. Bar-Ilan, G. Kortsarz, and D. Peleg. Generalized submodular cover problems and applications. *Theoretical Computer Science*, 250(1-2):179–200, 2001.

[3] J. Chuzhoy and S. Khanna. An $O(k^3 \log n)$-approximation algorithms for vertex-connectivity survivable network design. In *FOCS*, pages 437–441, 2009.

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

[5] A. Frank, T. Ibaraki, and H. Nagamochi. On sparse subgraphs preserving connectivity properties. *J. Graph Theory*, 17(3):275–281, 1993.

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

[7] T. Fukunaga. Approximating minimum cost source location problems with local vertex-connectivity demands. In *TAMC*, pages 428–439, 2011.

[8] 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.

[9] T. Ishii. Greedy approximation for source location problem with vertex-connectivity requirements in undirected graphs. *J. Discrete Algorithms*, 7:570–578, 2009.

[10] H. Ito, K. Makino, K. Arata, S. Honami, Y. Itatsu, and S. Fujishige. Source location problem with flow requirements in directed networks. *Optimization Methods and Software*, 18(4):427–435, 2003.

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

[12] G. Kortsarz and Z. Nutov. *Approximating minimum cost connectivity problems,* Ch. 58 in *Approximation Algorithms and Metahueristics,* T. F. Gonzales ed. 2007.

[13] G. Kortsarz and Z. Nutov. Tight approximation algorithm for connectivity augmentation problems. *J. Computer and System Sciences*, 74(5):662–670, 2008.

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

[15] L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13:383–390, 1975.

[16] H. Nagamochi and T. Ibaraki. *Algorithmic Aspects of Graph Connectivity, Ch. 9.* Cambridge University Press, 2008.

[17] H. Nagamochi, T. Ishii, and H. Ito. Minimum cost source location problem with vertex-connectivity requirements in digraphs. *IPL*, 80:287–294, 2001.

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

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

[20] M. Sakashita, K. Makino, and S. Fujishige. Minimum cost source location problems with flow requirements. *Algorithmica*, 50:555–583, 2008.

[21] A. Schrijver. *Combinatorial Optimization Polyhedra and Efficiency.* Springer, 2004.

[22] H. Tamura, M. Sengoku, S. Shinoda, and T. Abe. Location problems on undirected flow networks. *IEICE Trans.*, E73:1989–1993, 1990.

[23] A. Vakilian. Node-weighted prize-collecting network design problems. M.Sc. Thesis, 2013.

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