The Open University of Israel
Department of Mathematics and Computer Science

Minimum Power Connectivity Problems

Thesis submitted as partial fulfillment of the requirements
towards an M.Sc. degree in Computer Science
The Open University of Israel
Computer Science Division

By
Yuval Lando

Prepared under the supervision of Prof. Zeev Nutov

April 2007

Abstract

Given a (directed or undirected) graph with costs on the edges, the power of a node is the maximum cost of an edge leaving it, and the power of the graph is the sum of the powers of its nodes. Motivated by applications for wireless networks, we consider some fundamental network design problems under the power minimization criteria. Let $\mathcal{G} = (V, \mathcal{E})$ be a graph with edge-costs $\{c_e : e \in \mathcal{E}\}$ and let k be an integer. We consider finding a min-power subgraph G of $\mathcal{G}$ that satisfies some prescribed connectivity requirements. The **Min-Power *k* Edge-Disjoint Paths** (MPk-EDP) problem requires that G contains k pairwise edge-disjoint st-paths for given $s, t \in V$; the **Min-Power *k*-Edge-Outconnected Subgraph** (MPk-EOS) problem requires that G contains k pairwise edge-disjoint sv-paths for all $v \in V - s$, for given $s \in V$; and the **Min-Power *k*-Edge-Connected Subgraph** (MPk-ECS) problem requires that G is spanning and k-connected. When the paths are required to be internally disjoint, we get the problems **Min-Power *k* Disjoint Paths** (MPk-DP), **Min-Power *k*-Outconnected Subgraph** (MPk-OS), and **Min-Power *k*-Connected Subgraph** (MPk-CS), respectively. We survey the currently best known approximation algorithms for these problems, mainly for directed graphs.

We then present our original results as follows. We give an evidence that the undirected MPk-EDP and MPk-ECS and directed MPk-EOS and MPk-ECS are unlikely to admit a polylogarithmic approximation ratio even for unit costs. On the other hand, for both directed and undirected graphs we give a polynomial time algorithm for finding a min-power augmenting edge set that increases the st-edge-connectivity by 1; this implies a k-approximation algorithm for undirected MPk-EDP. We also give a $\min\{k + 4, O(\log n)\}$-approximation algorithm for node-connectivity version of undirected MPk-EOS; this improves the previously best known ratio of $2k - 1/3$.

Acknowledgements

I wish to thank my thesis advisor, Dr. Zeev Nutov, for his friendly, knowledgeable and invaluable guidance, throughout the gestation and development of the results presented in this thesis. Dr. Nutov's patience, encouragement, inspiration and time investment in guiding me through this process are all much appreciated.

I wish to thank my parents for their support and encouragement through my studying toward the M.Sc.

Contents

1 Introduction and preliminaries **4**

1.1 Problems considered, motivation, and previous work 4

1.2 Results in this Thesis 6

1.3 Notation 8

2 Approximation algorithms for directed min-power connectivity problems **8**

2.1 Covering intersecting families (Proof of Theorem 2.2) 11

2.2 Directed min-power connectivity (Proof of Theorem 2.1) 16

2.2.1 Part (i): Algorithm for directed MPk-OS/MPk-EOS 16

2.2.2 Part (ii): Algorithm for directed MPk-ECS 18

2.2.3 Part (iii): Algorithm for directed MPk-CS 19

3 Undirected MPk-EDP (Proof of Theorem 1.3) **19**

3.1 Part (i): Approximation hardness of undirected MPk-EDP 20

3.2 Part (ii): Polynomial algorithm for MPk-EDPA 21

4 Hardness of directed min-power edge-connectivity problems (Proof of Theorem 1.4) **22**

4.1 Arbitrary costs 23

4.2 Symmetric costs 24

5 Approximation algorithm for MPk-OS (Proof of Theorem 1.5) **26**

6 Conclusions **28**

1 Introduction and preliminaries

1.1 Problems considered, motivation, and previous work

Wireless networks are an important subject of study due to their extensive applications. A large research effort focused on performing network tasks while minimizing the power consumption of the radio transmitters of the network. In wired networks, one wants to find a subgraph of the minimum cost instead of the minimum power. This is the main difference between the optimization problems for wired versus wireless networks. In wireless networks, a range (power) assignment to radio transmitters determines the resulting communication network. We consider finding a power assignment to the nodes of a network such that the resulting communication network satisfies prescribed connectivity properties, and such that the total power is minimized. For motivation and applications to wireless networks (which is the same as of their min-cost variant for wired networks), see, e.g., [15, 1, 3, 16, 20].

Let $G = (V, E)$ be a (possibly undirected) graph with edge costs $\{c_e : e \in E\}$. For $v \in V$, the *power* $p(v) = p_c(v)$ *of* v in G (w.r.t. c) is the maximum cost of an edge leaving v in G (or zero, if no such edge exists). The power $p(G) = \sum_{v \in V} p(v)$ of G is the sum of powers of its nodes. Note that $p(G)$ differs from the ordinary cost $c(G) = \sum_{e \in E} c(e)$ of G even for unit costs; for unit costs, if G is undirected then $c(G) = |E|$ and $p(G) = |V|$. For example, if E is a perfect matching on V then $p(G) = 2c(G)$. If G is a clique then $p(G)$ is roughly $c(G)/\sqrt{|E|/2}$. For directed graphs, the ratio of cost over the power can be equal to the maximum outdegree of a node in G, e.g., for stars with unit costs. The following statement shows that these are the extremal cases for general costs.

Proposition 1.1 ([16]) $c(G)/\sqrt{|E|/2} \leq p(G) \leq 2c(G)$ *for any undirected graph* $G = (V, E)$*, and if* G *is a forest then* $c(G) \leq p(G) \leq 2c(G)$*. For any directed graph* G *holds:* $c(G)/d_{\max}(G) \leq p(G) \leq c(G)$*, where* $d_{\max}(G)$ *is the maximum outdegree of a node in* G*.*

The simplest connectivity requirements is when there should be an st-path for a specified node pair s, t. Another relatively simple case is when there should be a path from s to any other node. Min-cost variants are the (directed/undirected) Shortest st-Path problem and the Min-Cost Spanning Tree problem. In the min-power case, the directed/undirected Min-Power st-Path is solvable in polynomial time by a simple reduction to the min-cost case. The Min-Power Spanning Tree problem is APX-hard for undirected graphs and admits a 5/3-approximation algorithm [1], while the directed case is at least as hard as the Set-Cover problem, and thus has an $\Omega(\log n)$-approximation threshold (namely, it cannot be approximated within $C \ln n$ for some universal constant $C < 1$, unless P=NP). For the directed min-power spanning tree problem a $2H(n)$-approximation algorithm is given in [3] ($H(n) = \sum_{i=1}^{n} 1/i$ is the nth harmonic number). However, the "reverse" directed min-power spanning tree problem, when we require a path from every node to s, is equivalent to the min-cost case, and thus is solvable in polynomial time.

An important network property is fault-tolerance. A graph G is *k-outconnected from s* if it has k internally disjoint sv-paths for any $v \in V$; G is *k-inconnected to s* if its reverse graph is k-outconnected from s (for undirected graphs these two concepts are the same); G is *k-connected* if it has k internally disjoint uv-paths for all $u, v \in V$. When the paths are required only to be edge-disjoint, the graph is *k-edge outconnected from s*, *k-edge inconnected to s*, and *k-edge-connected*, respectively (for undirected graphs these concepts are the same).

We consider the following generalizations of the problems from [1, 3], that were already studied, e.g., [16, 20, 25]. These problems are defined for both directed and undirected graphs.

Min-Power k Edge-Disjoint Paths (MPk-EDP)
Instance: A graph $\mathcal{G} = (V, \mathcal{E})$, edge-costs $\{c_e : e \in \mathcal{E}\}$, $s, t \in V$, and an integer k.
Objective: Find a min-power subgraph G of $\mathcal{G}$ that contains k pairwise edge-disjoint st-paths.

Min-Power k-Edge-Inconnected Subgraph (MPk-EIS):
Instance: A graph $\mathcal{G} = (V, \mathcal{E})$, edge-costs $\{c_e : e \in \mathcal{E}\}$, $s \in V$, and an integer k.
Objective: Find a min-power k-inconnected to s spanning subgraph G of $\mathcal{G}$.

Min-Power k-Edge-Outconnected Subgraph (MPk-EOS):
Instance: A graph $\mathcal{G} = (V, \mathcal{E})$, edge-costs $\{c_e : e \in \mathcal{E}\}$, $s \in V$, and an integer k.
Objective: Find a min-power k-outconnected from s spanning subgraph G of $\mathcal{G}$.

Min-Power k-Edge-Connected Subgraph (MPk-ECS):
Instance: A graph $\mathcal{G} = (V, \mathcal{E})$, edge-costs $\{c_e : e \in \mathcal{E}\}$, and an integer k.
Objective: Find a min-power k-connected spanning subgraph G of $\mathcal{G}$.

When the paths are required to be internally disjoint we get the problems
Min-Power k Disjoint Paths (MPk-DP) (instead of MPk-EDP);
Min-Power k-Inconnected Subgraph (MPk-IS) (instead of MPk-EIS);
Min-Power k-Outconnected Subgraph (MPk-OS) (instead of MPk-EOS);
Min-Power k-Connected Subgraph (MPk-CS) (instead of MPk-ECS).

Min-cost versions of these problems were studied extensively for both directed and undirected graphs, see, e.g., [8, 13, 14, 11, 12], and surveys in [10, 18, 23]. Min-cost version of directed/undirected MPk-EDP/MPk-DP is polynomially solvable (this is the incapacitated Min-Cost k-Flow problem). For directed graphs the min-cost versions of MPk-EOS/MPk-OS are polynomially solvable, see [8] and [13, 12], respectively; more efficient algorithms are given in [14, 11]. This implies a 2-approximation algorithm for undirected graphs. Directed MPk-DP is solvable in polynomial time by a straightforward reduction to the min-cost case, and this implies a 2-approximation algorithm for undirected MPk-DP. On the other hand, for directed MPk-EDP we have the following approximation threshold:

Theorem 1.2 ([16]) *Directed* MPk-EDP *cannot be approximated within* $O(2^{\log^{1-\varepsilon} n})$ *for any fixed* $\varepsilon > 0$*, unless* NP $\subseteq$ DTIME$(n^{\mathrm{polylog}(n)})$.

1.2 Results in this Thesis

As was mentioned, for $k = 1$ directed/undirected MPk-EDP are easily reduced to the min-cost case. We are not aware of any nontrivial exact/approximation algorithms for undirected MPk-EDP for arbitrary k. We give a strong evidence that a polylogarithmic approximation algorithm for undirected MPk-EDP/MPk-ECS may not exist even for highly restricted instances. For that, we show a reduction from the following extensively studied problem to the undirected MPk-EDP/MPk-ECS. For a graph $\mathcal{J} = (V, \mathcal{I})$ and $X \subseteq V$ let $\mathcal{I}(X)$ denote the edges in $\mathcal{I}$ with both ends in X.

Densest ℓ-Subgraph (Dℓ-S)
Instance: A graph $\mathcal{J} = (V, \mathcal{I})$ and an integer ℓ.
Objective: Find $X \subseteq V$ with $|X| \leq \ell$ and $|\mathcal{I}(X)|$ maximum.

The best known approximation ratio for Dℓ-S is roughly $|V|^{-1/3}$ [9] even for the case of bipartite graphs (which up to a factor of 2 is as hard to approximate as the general case), and in spite of numerous attempts to improve it, this ratio holds for almost 10 years. We also consider the following "augmentation" version of undirected MPk-EDP (the directed case is easy, c.f., [25]), which already generalizes the case $k = 1$ considered in [1].

Min-Power k Edge-Disjoint Paths Augmentation (MPk-EDPA)
Instance: A graph $\mathcal{G} = (V, \mathcal{E})$ with edge-costs $\{c_e : e \in \mathcal{E}\}$, $s, t \in V$, an integer k, and a subgraph $G_0 = (V, E_0)$ of $\mathcal{G}$ that contains $k-1$ pairwise edge-disjoint st-paths.
Objective: Find $F \subseteq \mathcal{E} - E_0$ so that $G_0 + F$ contains k pairwise edge-disjoint st-paths and with $p(G_0 + F) - p(G_0)$ minimum.

Theorem 1.3
(i) *If there exists a ρ-approximation algorithm for undirected* MPk-EDP/MPk-ECS*, then there exists a $1/(2\rho^2)$-approximation algorithm for* Dℓ-S *on bipartite graphs.*
(ii) *Undirected* MPk-EDPA *is in P; thus* MPk-EDP *admits a k-approximation algorithm.*

In [25] is given an $O(k \ln n)$-approximation algorithm for directed MPk-OS, MPk-EOS, and MPk-ECS, and a k-approximation algorithm for directed MPk-EIS; these ratios are tight up to constant factor if k is "small", but may seem weak if k is large. We prove that for these four problems a polylogarithmic approximation ratio is unlikely to exist even when the costs are symmetric.

Theorem 1.4 *Directed* MPk-EDP/MPk-EOS/MPk-EIS/MPk-ECS *cannot be approximated within* $O(2^{\log^{1-\varepsilon} n})$ *for any* $\varepsilon > 0$ *even for symmetric costs, unless* $\mathrm{NP} \subseteq \mathrm{DTIME}(n^{\mathrm{polylog}(n)})$.

We also improve the best known ratio of $2k - 1/3$ [20] for undirected MPk-OS as follows:

Theorem 1.5 *Undirected* MPk-OS *admits a* $\min\{k+4, O(\log n)\}$*-approximation algorithm.*

The following table summarizes the currently best known approximation ratios and thresholds for the problems considered (note that *directed* MPk-EOS and directed MPk-EIS are *not* equivalent).

Problem	**Edge-Connectivity**		**Node-Connectivity**	
	Undirected	*Directed*	*Undirected*	*Directed*
MPk-DP	k $\Omega(1/\sqrt{\sigma})$	k [25] $\Omega(2^{\log^{1-\varepsilon} n})$ [16]	2 [16] $--$	in P [16]
MPk-IS	$2k-1/3$ [20] $\Omega(1/\sqrt{\sigma})$	k [25] $\Omega(2^{\log^{1-\varepsilon} n})$	$\min\{k+4, O(\ln n)\}$ APX for $k=1$ [16]	k [25] $--$
MPk-OS	$2k-1/3$ [20] $\Omega(1/\sqrt{\sigma})$	$O(k \ln n)$ [25] $\Omega(2^{\log^{1-\varepsilon} n})$	$\min\{k+4, O(\ln n)\}$ APX for $k=1$ [16]	$O(k \ln n)$ [25] $\Omega(\ln n)$ for $k=1$ [3]
MPk-CS	$2k-1/3$ [20] $\Omega(1/\sqrt{\sigma})$	$O(k \ln n)$ [25] $\Omega(2^{\log^{1-\varepsilon} n})$	$O(\alpha + \ln n)$ [20] $\Omega(\alpha)$ unless $\alpha = O(\ln n)$ [20]	$O(k(\ln n + k))$ [25] $\Omega(\ln n)$ for $k=1$ [3]

Table 1: Currently best known approximation ratios and thresholds for min-power connectivity problems. Results without references are the ones proved in this paper. σ is the best ratio for Dℓ-S; currently σ is roughly $O(n^{1/3})$ [9]. α is the best ratio known for the Min-Cost k-Connected subgraph problem; currently, $\alpha = \lceil (k+1)/2 \rceil$ for $2 \leq k \leq 7$ (see [2] for $k = 2, 3$, [7] for $k = 4, 5$, and [21] for $k = 6, 7$); $\alpha = k$ for $k = O(\ln n)$ [21], $\alpha = 6H(k)$ for $n \geq k(2k-1)$ [5], and $\alpha = O(\ln k \cdot \min\{\sqrt{k}, \frac{n}{n-k} \ln k\})$ for $n < k(2k-1)$ [22].

We note that for min-cost connectivity problems, a ρ-approximation algorithm for directed graphs usually implies a 2ρ-approximation for undirected graphs, c.f., [23]. For min-power problems we do not see such a reduction. For min-cost problems a standard reduction to reduce the undirected variant to the directed one is: replace every undirected edge uv by two anti-parallel directed edges uv, vu of the same cost as e, find a solution G to the directed variant and take the underlying graph of G. This reduction does not work for min-power problems, since the power of the underlying graph of G can be much larger than that of G, e.g., if G is a star. The approximation algorithm for the directed case might select only one of the two anti-parallel edges, and this does not correspond to a solution for the undirected case.

Theorems 1.3, 1.4, and 1.5 are proved in Sections 3, 4, and 5, respectively.

1.3 Notation

Here is some notation used in the paper. Let $G = (V, E)$ be a directed graph. For disjoint $X, Y \subseteq V$ let $\delta_G(X, Y) = \delta_E(X, Y)$ be the set of edges from X to Y in E. For brevity, $\delta_E(X) = \delta_E(X, V - X)$, $d_E(X) = |\delta_E(X)|$, and $\delta^+_E(X) = \delta_E(V - X, X)$. Given edge costs $\{c(e) : e \in E\}$, the power of v in G is $p(v) = \max_{e \in \delta_E(v)} c(e)$, and the power of G is $p(G) = p_E(V) = \sum_{v \in V} p(v)$. Throughout the paper, $\mathcal{G} = (V, \mathcal{E})$ denotes the input graph with nonnegative costs on the edges. Let $n = |V|$ and $m = |\mathcal{E}|$. Given $\mathcal{G}$, our goal is to find a minimum power spanning subgraph $G = (V, E)$ of $\mathcal{G}$ that satisfies some prescribed property. We assume that a feasible solution exists; let opt denote the optimal solution value of an instance at hand.

2 Approximation algorithms for directed min-power connectivity problems

In this section we we consider only directed graphs, so, unless stated otherwise, "graph" means "directed graph". We survey the approximation algorithms from [25] for min-power directed connectivity problems $\mathsf{MP}k$-EOS,$\mathsf{MP}k$-OS,$\mathsf{MP}k$-ECS, and $\mathsf{MP}k$-CS; [25] gives approximation algorithms for the following augmentation versions of these problems. Suppose that $\mathcal{G}$ has a subgraph $G_0 = (V, E_0)$ of power zero which is k_0-outconnected from r, and the goal is to augment G_0 by a min-power edge-set $F \subseteq \mathcal{E} - E_0$ so that the resulting graph $G = G_0 + F$ is k-outconnected from r. Formally:

Min-Power (k_0, k)-Outconnectivity Augmentation ($\mathsf{MP}(k_0, k)$-OA):
Instance: A graph $G_0 = (V, E_0)$ which is k_0-outconnected from r, an edge set $\mathcal{I}$ on V with costs $\{c_e : e \in I\}$, and an integer $k > k_0$.
Objective: Find min-power $I \subseteq \mathcal{I}$ so that $G = G_0 + I$ is k-outconnected from r.

In a similar way define the augmentation versions of MPk-EOS, MPk-CS and MPk-ECS, respectively:

Min-Power (k_0, k)-Edge-Outconnectivity Augmentation (MP(k_0, k)-EOA);
Min-Power (k_0, k)-Connectivity Augmentation (MP(k_0, k)-CA);
Min-Power (k_0, k)-Edge-Connectivity Augmentation (MP(k_0, k)-ECA).

In [3] are given approximation algorithms for $k_0 = 0$ and $k = 1$: a $2H(n)$-approximation for the Min-Power Directed Tree problem and a $(2H(n) + 1)$-approximation for the Min-Power Strongly Connected Subgraph problem ($H(n)$ denotes the nth Harmonic number). As each one of these problems generalizes the Set-Cover problem (c.f., [3]), the results in [3] are essentially tight up to a constant factor. For arbitrary k_0, k, [25] proves:

Theorem 2.1 *There exist approximation algorithms with approximation ratios:*
(i) $2(k - k_0)H(n) = O(k \ln n)$ *for directed* MP(k_0, k)-OA *and* MP(k_0, k)-EOA*;*
(ii) $(k - k_0)(2H(n) + 1) = O(k \ln n)$ *for directed* MP(k_0, k)-ECA*;*
(iii) $(k - k_0)(2H(n) + (k + k_0 + 1)/2) = O(k(\ln n + k))$ *for directed* MP(k_0, k)-CA.

The approximation ratios in Theorem 2.1 are $O(\ln n)$ for any fixed k, which is tight (up to a constant factor) if k is "small" (usually, $k \leq 3$ in practical networks), but may seem weak if k is large. However, it might be that a much better approximation algorithm does not exists: Theorem 1.4 states that (for large k) MPk-EOS cannot be approximated within $O(2^{\log^{1-\varepsilon} n})$ for any fixed $\varepsilon > 0$, unless NP $\subseteq$ DTIME$(n^{\text{polylog}(n)})$. The same hardness result is valid for the "reverse" problem MPk-EIS when there should be k edge-disjoint vs-paths for every $v \in V$; however, unlike MPk-EOS, this problem admits a k-approximation algorithm [24], and, in particular, is in P for $k = 1$. In contrast, for *undirected* MPk-OS [24] gives an $O(\ln n)$-approximation algorithm for any k.

In fact, Theorem 2.1 is just a summary of (some) applications of a general approximation algorithm for finding a min-power (edge-)cover of an intersecting family. A family $\mathcal{F}$ of subsets of a groundset V is an *intersecting family* if $X \cap Y, X \cup Y \in \mathcal{F}$ for any intersecting $X, Y \in \mathcal{F}$. An edge set I covers $\mathcal{F}$ if for every $X \in \mathcal{F}$ there is an edge in I entering X, that is, there is $uv \in I$ with $u \in V - X$ and $v \in X$. In [25] is given a $2H(n)$-approximation algorithm for the problem of finding a min-power cover of an intersecting family $\mathcal{F}$, but its polynomial implementation (in case $\mathcal{F}$ is not given explicitly) requires that certain queries related to $\mathcal{F}$ can be answered in polynomial time. Given an edge set I on V, the *residual family* $\mathcal{F}_I$ *of* $\mathcal{F}$ (w.r.t. I) consists of all members of $\mathcal{F}$ that are uncovered by edges of I. It is well known that if $\mathcal{F}$ is intersecting so is $\mathcal{F}_I$ for any I. A set $C \in \mathcal{F}$ is an $\mathcal{F}$-*core*, or simply a *core* if $\mathcal{F}$ is understood, if C does not contain two disjoint members of $\mathcal{F}$. Clearly, the maximal $\mathcal{F}$-cores are pairwise disjoint if $\mathcal{F}$ is intersecting. Given a maximal core C let $\mathcal{F}(C) = \{X \in \mathcal{F} : X \subseteq C\}$. For any edge set I on V, make the following two assumptions:

Assumption 1:
The maximal $\mathcal{F}_I$-cores can be computed in polynomial time.

Assumption 2:
For any maximal $\mathcal{F}_I$-core C, a min-cost $\mathcal{F}_I(C)$-cover can be computed in polynomial time.

Theorem 2.2 *The problem of finding a min-power cover of an intersecting family on n elements admits a $2H(n)$-approximation algorithm under Assumptions 1 and 2.*

Theorem 2.2 extends to so called "crossing families". Two sets $X, Y \subset V$ *cross* if $X \cap Y, X - Y, Y - X, V - (X \cup Y)$ are all nonempty. A set family $\mathcal{F}$ is a *crossing family* if $X \cap Y, X \cup Y \in \mathcal{F}$ for any crossing $X, Y \in \mathcal{F}$. Let us say that an edge set I is a *reverse cover* of $\mathcal{F}$ if for every $X \in \mathcal{F}$ there is an edge in I entering X. It is known that (c.f., [24]):

Fact 2.3 *Let $\mathcal{F}$ be an intersecting family. If I is an inclusion miniaml reverse cover of $\mathcal{F}$ then $d_I(v) \leq 1$ for every $v \in V$, and thus the power of I equals it cost. In particular, I is a min-power reverse cover of $\mathcal{F}$ if, and only if, I is a min-cost reverse cover of $\mathcal{F}$.*

Any crossing family $\mathcal{F}$ can be naturally represented by two intersecting families as follows: fix $r \in V$ and define $\mathcal{F}_r^{in} = \{X \in \mathcal{F} : r \notin X\}$ and $\mathcal{F}_r^{out} = \{V - X : X \in \mathcal{F} - \mathcal{F}_r^+\}$. Then I covers $\mathcal{F}$ if, and only if, I is a cover of $\mathcal{F}_r^{in}$ and I is reverse cover of $\mathcal{F}_r^{out}$. Combining with Fact 2.3, we get:

Corollary 2.4 *The problem of finding a min-power cover of a crossing family $\mathcal{F}$ on V admits a $(2H(n)+1)$-approximation algorithm, if for some $r \in V$ Assumptions 1 and 2 are valid for $\mathcal{F}_r^{in}$ and if the min-cost reverse cover of $\mathcal{F}_r^{out}$ can be computed in polynomial time.*

A set function f defined on subsets of V is *intersecting supermodular* if $f(X) + f(Y) \leq f(X \cap Y) + f(X \cup Y)$ for any intersecting $X, Y \subset V$. An edge set I covers f if in the graph (V, I) the indegree of every $X \subset V$ is at least $f(X)$. A $\{0,1\}$-valued set function is intersecting supermodular if, and only if, its support is an intersecting family. A natural question is whether Theorem 2.2 extends to intersecting supermodular set functions. As $\mathsf{MP}k\text{-}\mathsf{EOS}$ is a particular case of the problem of finding a min-power cover of an intersecting supermodular set function, such an extension is unlikely due to the hardness result in Theorem 1.4.

The proof of Theorem 2.2 combined "set-cover approximation techniques" of [19] used in [3] for $k = 1$ that are based on "density" considerations (c.f., [17]) with the techniques used for min-cost connectivity problems. However, unlike [19, 3], one cannot use specific graph properties. To prove that one can find an edge set of appropriate density, [25] uses the method of "uncrossing" sets (c.f., [27]). It defines an analogue of spiders called "star-covers": unlike [19, 3] a star cover is not necesarilly a tree. Showing that any inclusion minimal $\mathcal{F}$-cover can be decomposed into such star-covers is harder than showing a decomposition of a tree into spiders; recently, such a decomposition was extended for set families related to the undirected Node Weighted Steiner Network problem – a generalization of the Node Weighted Steiner Forest problem, see [26].

Proofs of Theorems 2.2 and 2.1 are given in Sections 2.1 and 2.2, respectively.

2.1 Covering intersecting families (Proof of Theorem 2.2)

The following result about the performance of an *Approximate Greedy Algorithm* for a certain type of covering problems is widely used to prove polylogarithmic approximation ratios. This type of problems can be defined as follows:

Covering Problem
Instance: A groundset $\mathcal{I}$ and integral functions ν, p on $2^{\mathcal{I}}$ given by an evaluation oracle.
Objective: Find $I \subseteq \mathcal{I}$ with $\nu(I) = \nu(\mathcal{I})$ and with $p(I)$ minimized.

We call ν the *deficiency function* and p the *payment function*; ν is assumed to be decreasing and measures how far I is from being a feasible solution; p is assumed to be increasing and subadditive, namely $p(I_1 \cup I_2) \leq p(I_1) + p(I_2)$ for all $I_1, I_2 \subseteq \mathcal{I}$. In our case p is just the power function, and clearly, it is decreasing and subadditive. Let $\rho > 1$ and let opt be the optimal solution value for the Covering Problem. The *ρ-Approximate Greedy Algorithm* starts with $I = \emptyset$ and iteratively adds subsets of $\mathcal{I} - I$ to I one after the other using the following rule. As long as $\nu(I) > \nu(\mathcal{I})$ it adds to I a set $F \subseteq \mathcal{I} - I$ so that

$$\sigma_I(F) = \frac{\nu(I) - \nu(I + F)}{p(F)} \geq \frac{\nu(I) - \nu(\mathcal{I})}{\rho \cdot \mathsf{opt}}. \tag{1}$$

The following known statement is proved using the same methods as in [17] where the Set-Cover problem was considered. We give a proof sketch of a slightly weker version.

Theorem 2.5 *For any covering problem so that the payment function p is increasing and subadditive and the deficiency function ν is monotone decreasing, the ρ-Approximate Greedy Algorithm computes a solution I with $p(I) \leq \rho H(\nu(\emptyset) - \nu(\mathcal{I})) \cdot \mathsf{opt}$, where $H(n)$ denotes the nth Harmonic number.*

Proof sketch: We may assume that $\nu(\mathcal{I}) = 0$. We prove a slightly weaker result, namely: $p(I) \leq \rho(1 + \ln \nu(\emptyset)) \cdot \mathsf{opt}$. Let I_j be the partial solution at the end of iteration j, where $I_0 = \emptyset$, and let F_j be the set added at iteration j, thus $I_j = I_{j-1} + F_j$, $j = 1, \ldots, \ell$. Let $\nu_j = \nu(I_j)$ and $p_j = p(F_j)$. Since ν is decresing, then by (1) we have $(\nu_{j-1} - \nu_j)/w_j \geq \nu_{j-1}/(\rho \cdot \mathsf{opt})$. Thus

$$\nu_j \leq \nu_{j-1} \left(1 - \frac{p_j}{\rho \cdot \mathsf{opt}}\right) .$$

We have $\nu_\ell = 0$ while $\nu_{\ell-1} \geq 1$. Unraveling the last inequality we obtain:

$$\frac{\nu_{\ell-1}}{\nu_0} \leq \prod_{j=1}^{\ell-1} \left(1 - \frac{p_j}{\rho \cdot \mathsf{opt}}\right) .$$

Taking natural logarithms from both sides and using the inequality $\ln(1+x) \leq x$ we obtain:

$$\rho \cdot \mathsf{opt} \cdot \ln\left(\frac{\nu_0}{\nu_{\ell-1}}\right) \geq \sum_{j=1}^{\ell-1} p_j .$$

Finally, using the subadditivity of p and observing that $p_\ell \leq \rho \cdot \mathsf{opt}$ and $\nu_{\ell-1} \geq 1$ we get:

$$p(I) = p\left(\bigcup_{j=1}^{\ell} F_j\right) \leq p_\ell + \sum_{j=1}^{\ell-1} p_j \leq \rho \cdot \mathsf{opt} + \rho \cdot \mathsf{opt} \cdot \ln \nu_0 = \rho(1 + \ln \nu_0) \cdot \mathsf{opt} .$$

□

In the rest of this section we prove the following Lemma:

Lemma 2.6 *Let $\nu(I)$ be the number of minimal cores in $\mathcal{F}_I$. Then an edge set F satisfying (1) with $\rho = 2$ can be found in polynomial time under Assumptions 1 and 2.*

For simplicity of exposition, let us revise our notation and use $\mathcal{F}$ instead of $\mathcal{F}_I$, and let $\nu = \nu(\emptyset)$. We assume that $\mathcal{I}$ is a feasible solution, thus $\nu(\mathcal{I}) = 0$. Then we need to show that under Assumptions 1 and 2 one can find in polynomial time an edge set F so that:

$$\sigma(F) = \frac{\nu - \nu(F)}{p(F)} \geq \frac{\nu}{2 \cdot \mathsf{opt}}. \tag{2}$$

Before presenting a formal proof of Lemma 2.6, we give a sketch. Let $\mathcal{C}$ be the set of maximal $\mathcal{F}$-cores. For $C \in \mathcal{C}$ let $E(C) = \{uv \in E : u, v \in C\}$ be the edges in E with both endpoints in C, and let $\mathcal{F}(C) = \{X \in \mathcal{F} : X \subseteq C\}$. Let E be a minimal $\mathcal{F}$-cover. We prove that (Corollary 2.8 and Lemma 2.9):
(i) $d_E(v) \leq 1$ for any $v \in C$ and $d_E^+(C) = 1$.
(ii) $E(C)$ plus the unique edge e_C in E that enters C cover $\mathcal{F}(C)$.
An edge set F is a *star-cover with root s* (an analogue of [19, 3] spiders) if for every $e \in \delta_F(s)$ there exists $C \in \mathcal{C}$ with $\delta_F(C) = \{e\}$ such that $e + F(C)$ is a minimal $\mathcal{F}(C)$-cover

(Definition 2.1). We prove that adding a star cover F decreases the number of cores by at least $\Delta(F)$, where $\Delta(F) = d_F(s) - 1$ if $d_F(s) \geq 2$ and $\Delta(F) = 1$ if $d_F(s) = 1$ (Lemma 2.11). By (ii), the set E' of edges in E which head lies in some core is decomposed into star-covers $F_1, \ldots, F_t$, and adding all these star-covers decreases ν by at least $\sum_{i=1}^{t} \Delta(F_i) \geq \nu/2$. As $p(E') \leq \mathsf{opt}$, we use an averaging argument as in [19, 3] to conclude that there exists a star-cover F for which (2) holds (Lemma 2.12). By (i), the power of a star-cover equals the power of s plus the *cost* of its edges that are not incident to s (Corollary 2.10). This, together with Assumptions 1 and 2 enables us to find in polynomial time a star-cover F that maximizes $\Delta(F)/p(F)$ (Lemma 2.13).

A formal proof of Lemma 2.6 follows. We need to establish some properties of minimal $\mathcal{F}$-covers of an intersecting family $\mathcal{F}$. Let E be a minimal $\mathcal{F}$-cover. By the minimality of E, for every $e \in E$ there exists $W_e \in \mathcal{F}$ such that $\delta_E^{in}(W_e) = \{e\}$; we call such W_e a *witness set for* e; note that e might have several distinct witness sets.

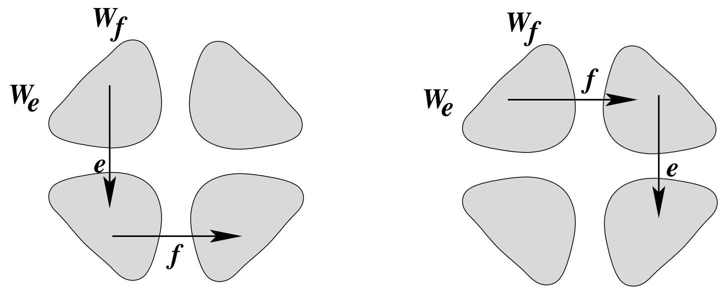

Figure 1: Two possible cases of intersecting witness sets W_e, W_f.

Lemma 2.7 *Let $\mathcal{F}$ be an intersecting family and let E be a minimal $\mathcal{F}$-cover. Let W_e, W_f be intersecting witness sets of two distinct edges $e, f \in E$. Then $W_e \cap W_f$ is a witness for one of e, f and $W_e \cup W_f$ is a witness for the other. (See Fig. 1.)*

Proof: Note that there is an edge in E entering $W_e \cap W_f$ and there is an edge in E entering $W_e \cup W_f$; this is since $W_e, W_f \in \mathcal{F}$ implies that $W_e \cap W_f, W_e \cup W_f$ belong to $\mathcal{F}$ and thus each of them is covered by some edge in E. However, if for arbitrary sets X, Y an edge covers one of $X \cap Y, X \cup Y$ then it also covers one of X, Y, and if some edge covers both $X \cap Y$ and $X \cup Y$ then it must cover both X and Y. Thus no edge in $E - \{e, f\}$ can cover $W_e \cap W_f$ or $W_e \cup W_f$, so one of e, f covers $W_e \cap W_f$, and thus the other must cover $W_e \cup W_f$. $\square$

Corollary 2.8 *Let X be a minimal core of an intersecting family $\mathcal{F}$ and let E be a minimal $\mathcal{F}$-cover. Then $d_E^{in}(X) = 1$.*

Proof: Clearly $d_E^{in}(X) \geq 1$, since E is an $\mathcal{F}$-cover and $X \in \mathcal{F}$. Assume that there are distinct $e, f \in \delta_E^{in}(X)$, and let W_e, W_f be their witness sets. Then $X \subseteq W_e \cap W_f$ (in particular, W_e, W_f intersect), and thus $e, f \in \delta_E^{in}(W_e \cap W_f)$. This contradicts Lemma 2.7. $\square$

Lemma 2.9 *Let C be a maximal core of an intersecting family $\mathcal{F}$ and let E be a minimal $\mathcal{F}$-cover. Let $E(C)$ be the set of edges in E with both endpoints in C, let X be the minimal core of $\mathcal{F}_{E(C)}$ contained in C (possibly $X = C$), and let e_C be the unique edge in E that enters X. Then $E(C) + e_C$ covers $\mathcal{F}(C) = \{X \in \mathcal{F} : X \subseteq C\}$, and $d_{E(C)}(v) \leq 1$ for every $v \in C$; thus $p(E(C)) = c(E(C))$, namely, the power of $E(C)$ equals its cost.*

Proof: Let X_1 be the minimal $\mathcal{F}$-core contained in C. By Corollary 2.8 there is a unique edge in E entering X_1, say e_1. If e_1 covers C, then $E(C) = \emptyset$, and it is easy to see that the statement holds. Otherwise, let X_2 be the minimal $\mathcal{F}_{e_1}$-core contained in C and let e_2 be the unique edge in E entering X_2, and so on, until C is covered by some edge e_q. In such a way we obtain sequences $e_1, \ldots, e_{q-1}$ of edges in $E(C)$ together with an additional edge e_q that enters C, and $X_1 \subset X_2 \cdots \subset X_q \subseteq C$ of sets in $\mathcal{F}(C)$ so that X_{i+1} is the core of $\mathcal{F}_{E_i}$ where $E_i = \{e_1, \ldots, e_i\}$ and e_i is the unique edge in E entering X_i. The statement follows, since we must have $E_{q-1} = E(C)$, $X = X_q$, and $e_q = e_C$. In particular, $E(C) + e_C = E_q$ covers $\mathcal{F}(C)$ and no two edges in $E(C)$ share a tail. $\square$

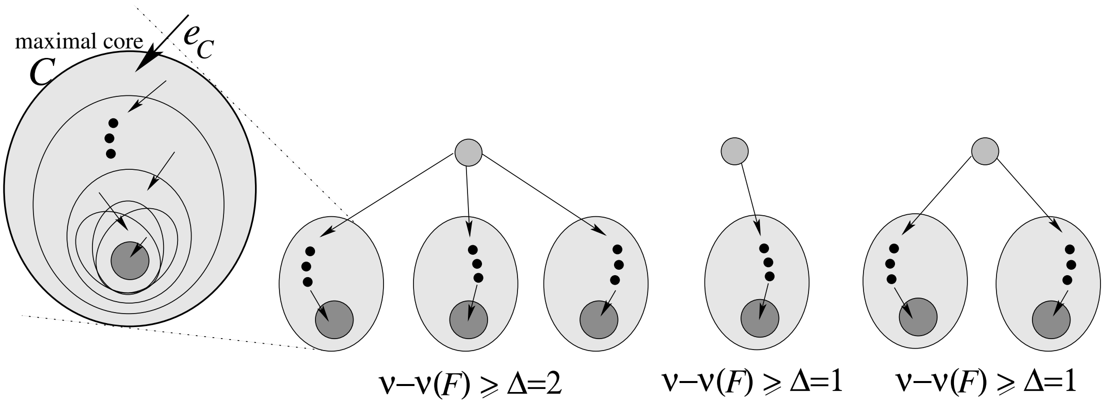

Figure 2: Star-covers.

Definition 2.1 *An edge set F is a* star-cover *(with root s) if for every $e \in \delta_F(s)$ there exists $C \in \mathcal{C}$ with $\delta_F(C) = \{e\}$ such that $e + F(C)$ is a minimal $\mathcal{F}(C)$-cover (see Fig. 2).*

As the family $\mathcal{F}(C)$ is intersecting for any maximal $\mathcal{F}$-core C, by Lemma 2.9 we get:

Corollary 2.10 *Let F be a star-cover with root s. Then $\delta_F(v) \leq 1$ for any $v \neq s$, and thus $p(F) = p_F(s) + c(F - \delta_F(s))$.*

Lemma 2.11 *For a star-cover F with root s let $\Delta(F) = d(s)-1$ if $d_F(s) \geq 2$ and $\Delta(F) = 1$ if $d_F(s) = 1$. Then $\nu - \nu(F) \geq \Delta(F)$.*

Proof: Let $\mathcal{F}'$ be the residual family of the sets that are uncovered by F. The minimal $\mathcal{F}$-cores not covered by F are also minimal $\mathcal{F}'$-cores, while any other minimal $\mathcal{F}'$-core X' must contain at least one $\mathcal{F}$-core covered by F. We claim that $s \in X'$ must hold for any such X', and thus: if $d(s) = 1$ no such X' exists, and if $d(s) \geq 2$ there is at most one such X', since the minimal $\mathcal{F}'$-cores are disjoint. To see that $s \in X'$, let X be a minimal $\mathcal{F}$-core contained in X' and let C be the maximal $\mathcal{F}$-core containing X. Let $Y = X \cap C$. Then $Y \in \mathcal{F}(C)$, thus there is an edge $uv \in F$ entering Y. Since uv does not cover X', we must have $u \in X' - C$. But then uv covers C, implying $u = s$. $\square$

Lemma 2.12 *There exists a star-cover F for which (2) holds.*

Proof: Let E be an inclusion minimal optimal $\mathcal{F}$-cover. For every maximal core C of $\mathcal{F}$ let E_C and e_C be as in Lemma 2.9. Let E' be the union taken over all maximal cores $C \in \mathcal{C}$ of the edge sets $F_C = E_C + e_C$. Then E' is decomposed into node disjoint star-covers $F_1, \ldots, F_t$. Now the statement follows by a simple averaging argument. Let $p_i = p(F_i)$ and let $\Delta_i = \Delta(F_i)$. We have $\sum_{i=1}^t p_i = p(E') \leq p(E) = \mathsf{opt}$ and $\sum_{i=1}^t \Delta_i \geq \nu/2$. Thus:

$$\frac{\sum_{i=1}^t \Delta_i}{\sum_{i=1}^t p_i} \geq \frac{\nu/2}{p(E')} \ .$$

From number theory we know that there must be index i so that $\Delta_i/p_i \geq \nu/(2p(E'))$. Let $F = F_i$. Then $\nu - \nu(F) \geq \Delta_i$, by Lemma 2.11. Consequently

$$\sigma(F) = \frac{\nu - \nu(F)}{p(F)} \geq \frac{\nu}{2 \cdot p(E')} \geq \frac{\nu}{2 \cdot p(E)} = \frac{\nu}{2 \cdot \mathsf{opt}} \ .$$

$\square$

Lemma 2.13 *A star-cover F that maximizes $\Delta(F)/p(F)$ can be found in polynomial time under Assumptions 1 and 2.*

Proof: We first compute the maximal cores; this can be done in polynomial time by Assumption 1. Second, for every node v that belongs to a maximal core C we define the weight $w(v)$ of v to be the minimum cost of an $\mathcal{F}_e(C)$-cover, where $e = uv$ is an arbitrary edge that has head v and enters C. This can be done in polynomial time by Assumption 2. Let us say that a star F is proper if every its edge enters some maximal $\mathcal{F}$-core. Given a proper star F with root s, let $w(F) = p(s) + w(L_F)$ where L_F is the set of leaves of F. We now see that our goal is to compute a proper star F that maximizes $\max\{|L_F| - 1, 1\}/w(F)$. We may assume that we know the root s and its power $p = p_F(s)$ in F; there are $O(n^2)$ distinct choices.

Delete all the edges, except that for every core $C \in \mathcal{C}$ among the edges sv with $v \in C$ and $p(sv) \leq p(s)$, if any, choose one with $w(v)$ minimal. This defines an auxiliary star T. Let $v_1, \ldots, v_q$ be the leaves of T sorted by increasing weight, so $w(v_1) \leq w(v_2) \leq \ldots \leq w(v_q)$. Let $W_j = \sum_{i=1}^{j} w(v_i)$, and let $\sigma_1 = 1/(p + W_1)$ and $\sigma_j = (j-1)/(p + W_j)$, $j = 1, \ldots, q$. We find the index j for which σ_j is maximum, which will determine the required star-cover. □

2.2 Directed min-power connectivity (Proof of Theorem 2.1)

2.2.1 Part (i): Algorithm for directed MPk-OS/MPk-EOS

We give a $2H(n)$-approximation algorithm for MP$(\ell, \ell+1)$-OA (rep., MP$(\ell, \ell+1)$-EOA). We apply this algorithm sequentially for $\ell = k_0, \ldots, k-1$ to produce edge sets $F_{k_0}, \ldots, F_{k-1}$ so that $G_0 + (F_{k_0} + \cdots + F_\ell)$ is $(\ell+1)$-outconnected (resp., $(\ell+1)$-edge-outconnected) from r, and $p(F_\ell) \leq 2H(n) \cdot \mathsf{opt}$, $\ell = k_0, \cdots, k-1$. Consequently, $F = F_{k_0} + \cdots + F_{k-1}$ is k-outconnected from r, and

$$p(F) \leq \sum_{\ell=k_0}^{k-1} p(F_\ell) \leq \sum_{\ell=k_0}^{k-1} 2H(n) \cdot \mathsf{opt} = 2(k-k_0)H(n) \cdot \mathsf{opt} \ .$$

A graph $G = (V, E)$ is *ℓ-edge-outconnected from r to T* if it has ℓ pairwise edge-disjoint rt-paths for every $t \in T$. Using Theorem 2.2, we give a $2H(n)$-approximation algorithm for the following problem, that includes both MP$(\ell, \ell+1)$-OA and MP$(\ell, \ell+1)$-EOA.

Instance: A graph $G_0 = (V, E_0)$ which is ℓ-edge-outconnected from r to T and an edge set $\mathcal{I}$ on V with costs $\{c_e : e \in \mathcal{I}\}$ so that every edge in $\mathcal{I}$ has its head in T.

Objective: Find a min-power edge-set $I \subseteq \mathcal{I}$ so that $G = G_0 + I$ is $(\ell+1)$-edge-outconnected from r to T.

MP$(\ell, \ell+1)$-EOA is a special case of this problem when $T = V$. For MP$(\ell, \ell+1)$-OA apply the following approximation ratio preserving reduction. Given an instance $G_0 = (V, E_0), \ell, r, \mathcal{I}$ of MP$(\ell, \ell+1)$-OA obtain an instance $G'_0 = (V', E'_0), T', \ell, r, \mathcal{I}', c'$ of the above problem as follows. Replace every node $v \in V$ by the two nodes v_t, v_h connected by the edge $v_t v_h$ of cost zero, and replace every edge $uv \in E_0 \cup \mathcal{I}$ by the edge $u_h v_t$ having the same cost as uv (which is zero if $uv \in E_0$). Let $r' = r_h$, $T' = \{v_t : v \in V\}$, and

$$E'_0 = \{u_h v_t : uv \in E_0\} + \{v_t v_h : v \in V\}, \quad \mathcal{I}' = \{u_h v_t : uv \in \mathcal{I}\} \ .$$

This establishes a bijective correspondence between edges in $\mathcal{I}$ and the edges in $\mathcal{I}'$. It is not hard to verify (see [11] for details) that $G'_0 = (V', E'_0)$ is ℓ-edge-connected from r' to T'. Furthermore, if $I' \subseteq \mathcal{I}$ corresponds to $I \subseteq \mathcal{I}$ then:

(i) I is a feasible solution if, and only if, I' is a feasible solution.
(ii) $d_I(v) = d_{I'}(v_h)$ and $d_{I'}(v_t) = 0$ for every $v \in V$; thus $p(I) = p(I')$.

We now show that above problem can be reduced to the min-power intersecting family cover problem, so that Assumptions 1 and 2 are valid. We say that $X \subseteq V - s$ is *tight* in G_0 if $X \cap T \neq \emptyset$ and $d^{in}(X) = \ell$. From Menger's Theorem we have:

Fact 2.14 *Let $G_0 = (V, E_0)$ be ℓ-edge-outconnected from r to T. Then $G = G_0 + I$ is $(\ell+1)$-edge-outconnected from r to T if, and only if, I covers all the tight sets in G_0.*

We now see that the augmentation problem is equivalent to the problem of finding a cover of the family of tight sets. However, since only edges with head in T can be added, this is equivalent to covering the family:

$$\mathcal{F} = \{X \cap T : X \text{ is tight in } G_0\} . \tag{3}$$

It is well known (c.f. [11]) that:

Fact 2.15 *The family $\mathcal{F}$ defined in (3) is intersecting.*

It remains to show that Assumptions 1 and 2 are valid for $\mathcal{F}$ defined by (3). For Assumption 1 we need to show that given a graph, the maximal $\mathcal{F}$-cores can be found in polynomial time (if some edges were added at previous steps, we consider the graph after these edges were added). We first show how to find the minimal $\mathcal{F}$-cores. Then, for Assumption 1, we will show that finding maximal $\mathcal{F}$-cores can be done using n max-flow computations; for Assumption 2 we will show that finding a min-cost $\mathcal{F}(C)$-cover for a given maximal core C can be done using one min-cost $(\ell+1)$-flow computation.

The minimal cores can be found using $|T|$ max-flow computations as follows. For every $t \in T$, compute a maximum rt-flow. If its value is ℓ, then in the corresponding residual network the set of nodes $\{v \in T : t \text{ is reachable from } v\}$ is the minimal core containing t; otherwise, no minimal core containing t exists. After the minimal $\mathcal{F}$-cores are found, to find the maximal cores, for every minimal core X do the following. Add an edge from r to every minimal core distinct from X. Then choose $t \in X$ and compute a maximum rt-flow; in the corresponding residual network the set of nodes $T - \{v \in T : v \text{ is reachable from } r\}$ is the maximal core containing X. Now we show how to find a min-cost $\mathcal{F}(C)$-cover for a maximal core C that contains a minimal core X. The construction is similar to the previous one: construct a network $H = G_0 + \mathcal{I}$, assigning zero costs to edges in E_0. Then add an edge from r to every minimal core distinct from X, and compute a min-cost $(\ell+1)$-flow f from r to some $t \in X$. The edge set $\{e \in \mathcal{I} : f(e) = 1\}$ is the desired $\mathcal{F}(C)$-cover.

2.2.2 Part (ii): Algorithm for directed $\mathsf{MP}k$-ECS

We give a $(2H(n){+}1)$-approximation algorithm for $\mathsf{MP}(\ell,\ell{+}1)$-$\mathsf{ECA}$. We apply this algorithm sequentially for $\ell = k_0,\ldots,k{-}1$ to produce edge sets $F_{k_0},\ldots,F_{k-1}$ so that $G_0{+}(F_{k_0}{+}\cdots{+}F_\ell)$ is $(\ell+1)$-edge-connected, and $p(F_\ell) \leq (2H(n)+1)\cdot \mathsf{opt}$, $\ell = k_0,\ldots,k-1$. Consequently, for $F = F_{k_0}+\cdots+F_{k-1}$, we get that G_0+F is k-edge-connected, and

$$p(F) \leq \sum_{\ell=k_0}^{k-1} p(F_\ell) \leq \sum_{\ell=k_0}^{k-1} (2H(n)+1)\cdot \mathsf{opt} = (k-k_0)(2H(n)+1)\cdot \mathsf{opt} \ .$$

Let us say that a graph is *ℓ-edge-inconnected to r* (resp., *ℓ-inconnected to r*) if its reverse graph is ℓ-edge-outconnected from r (resp., ℓ-outconnected from r). The problem of finding a min-cost augmenting edge set that increases the inconnectivity (or edge-inconnectivity) of a given directed graph from ℓ to $\ell+1$ can be solved in polynomial time, c.f., [11]. Using Fact 2.3 and methods as in the previos section, one can easily deduce (see [24] for details):

Proposition 2.16 *Finding a min-power augmenting edge set to increase the edge-inconnectivity (or the inconnectivity) of a given directed graph by* 1 *can be done in polynomial time.*

Now a $(2H(n)+1)$-approximation algorithm for $\mathsf{MP}(\ell,\ell+1)$-ECA can be deduced from Corollary 2.4, and explicitly is as follows. Let r be an arbitrary node of G.

1. Using the algorithm as in part (i) of Theorem 2.1 compute an edge set F' so that G_0+F' is $(\ell+1)$-edge-outconnected from r.

2. Compute a min-power edge set F'' so that G_0+F'' is $(\ell+1)$-edge-inconnected to r.

3. Output $F = F'+F''$.

Note that $G = G_0+F$ is both $(\ell+1)$-edge-outconnected from r and $(\ell+1)$-edge-inconnected to r. This implies that G is $(\ell+1)$-edge connected, so F is a feasible solution. To bound its power, let OPT be an optimal solution for $\mathsf{MP}(\ell,\ell+1)$-ECA. Since G_0+OPT is $(\ell+1)$-edge-outconnected from r we have $p(F') \leq 2H(n)p(OPT) \leq \mathsf{opt}$. Since G_0+OPT is $(\ell+1)$-edge-inconnected to r we have $p(F'') \leq p(OPT) \leq \mathsf{opt}$. Consequently,

$$p(F) = p(F'+F'') \leq p(F')+p(F'') \leq 2H(n)\cdot\mathsf{opt}+\mathsf{opt} = (2H(n)+1)\cdot\mathsf{opt}.$$

The proof of part (ii) of Theorem 2.1 is complete.

2.2.3 Part (iii): Algorithm for directed MPk-CS

We give a $(2H(n)+\ell+1)$-approximation algorithm for MP$(\ell,\ell+1)$-CA. We apply this algorithm sequentially for $\ell = k_0, \ldots, k-1$ to produce edge sets $F_{k_0}, \ldots, F_{k-1}$ so that $G_0 + (F_{k_0} + \cdots + F_\ell)$ is $(\ell+1)$-connected, and $p(F_\ell) \leq (2H(n)+\ell+1) \cdot \mathsf{opt}$, $\ell = k_0, \ldots, k-1$. Consequently, for $F = F_{k_0} + \cdots + F_{k-1}$, we get that $G_0 + F$ is k-connected, and

$$p(F) \leq \sum_{\ell=k_0}^{k-1} p(F_\ell) \leq \sum_{\ell=k_0}^{k-1} (2H(n)+\ell+1) \cdot \mathsf{opt} = (k-k_0)(2H(n)+(k+k_0+1)/2) \cdot \mathsf{opt} \ .$$

The algorithm for MP$(\ell,\ell+1)$-CA is the "augmentation power variant" of the $(k+1)$-approximation algorithm of [21] for the Min-Cost k-Connected Subgraph problem, and is as follows. Let $S \subseteq V$ be a subset of $\ell+1$ nodes (so $|S| = \ell+1$).

1. Construct a graph $\mathcal{G}_r$ by adding to $\mathcal{G}$ a new node r, and edges $\{rs : s \in S\}$ of cost 0. Using the algorithm as in Theorem 2.1 (i) compute an augmenting edge set F_r so that $G_0 + s + F_r$ is $(\ell+1)$-outconnected from r, and delete from F_r the edges leaving r.

2. For every $s \in S$ compute an optimal min-power augmenting edge set F_s so that $G_0 + F_s$ is $(\ell+1)$-inconnected to s.

3. Output $F_r + \bigcup_{s \in S} F_s$.

The fact that the algorithm computes a feasible solution was proved in [21] (this fact is independent from the cost/power of the edge sets computed). For every $s \in S$ we have $p(F_s) \leq \mathsf{opt}$, by a similar argument as in the proof of part (ii). We also have $p(F_r) \leq 2H(n) \cdot \mathsf{opt}$. Consequently,

$$p\left(F_r + \bigcup_{s \in S} F_s\right) \leq (2H(n)+|S|) \cdot \mathsf{opt} = (2H(n)+\ell+1) \cdot \mathsf{opt} \ .$$

The proof of part (ii) of Theorem 2.1 is complete.

3 Undirected MPk-EDP (Proof of Theorem 1.3)

In this section we give a proof of Theorem 1.3: a hardness result showing that MPk-EDP is unlikely to admit a polylogarithmic approximation, and a polynomial time algorithm for the partical case of augmentation version MPk-EDPA of MPk-EDP.

3.1 Part (i): Approximation hardness of undirected MPk-EDP

To prove part (i) of Theorem 1.3 we need the following known statement:

Lemma 3.1 *There exists a polynomial time algorithm that given a graph $G = (V, E)$ and an integer $1 \leq \ell \leq n = |V|$ finds a subgraph $G' = (V', E')$ of G with $|V'| = \ell$ and $|E'| \geq |E| \cdot \frac{\ell(\ell-1)}{n(n-1)}$.*

Proof: While G has more than ℓ nodes, repeatedly delete the minimum degree node from G. At the beginning of iteration $i + 1$, G has $n_i = n - i$ nodes and m_i edges, where $n_0 = n$ and $m_0 = m$. The average degree is $2m_i/n_i$, thus after iteration $i + 1$ the number m_{i+1} of edges in G is at least:

$$m_{i+1} \geq m_i - \frac{2m_i}{n_i} = m_i \cdot \frac{n-i-2}{n-i} \ .$$

The statement follows since the above recursive formula implies that after $i = n-\ell$ iterations:

$$\frac{m_i}{m} \geq \frac{(n-2)\cdots(n-i+1)(n-i)(n-i-1)}{n(n-1)(n-2)\cdots(n-i+1)} = \frac{(n-i)(n-i-1)}{n(n-1)} = \frac{\ell(\ell-1)}{n(n-1)} \ .$$

□

Given an instance $\mathcal{J} = (A + B, \mathcal{I})$ and ℓ of bipartite Dℓ-S, define an instance of (undirected) unit-cost MPk-EDP/MPk-ECS by adding new nodes $\{s, t\}$, a set of edges $E_0 = \{aa' : a \in A + s\} \cup \{bb' : b \in B + t\}$ of capacity $|A| + |B|$ and cost 0 each, and setting $c(e) = 1$ for all $e \in \mathcal{I}$. It is easy to see that any $E \subseteq \mathcal{I}$ determines $|E|$ edge-disjoint st-paths, and that $(A + B + \{s, t\}, E_0 + |E|)$ is k-connected if, and only if, $|E| \geq k$. Thus for any integer $k \in \{1, \ldots, |\mathcal{I}|\}$, if we have a ρ-approximation algorithm for undirected MPk-EDP/MPk-ECS, then we have a ρ-approximation algorithm for

$$\min\{|X| : X \subseteq A + B, |\mathcal{I}(X)| \geq k\} \ .$$

We show that this implies a $1/(2\rho^2)$-approximation algorithm for the original instance of bipartite Dℓ-S, which is $\max\{|\mathcal{I}(X)| : X \subseteq A + B, |X| \leq \ell\}$).

For every $k = 1, \ldots, |\mathcal{I}|$, use the ρ-approximation algorithm for MPk-EDP/MPk-ECS to compute a subset $X_k \subseteq A + B$ so that $|\mathcal{I}(X_k)| \geq k$, or to determine that no such X_k exists. Set $X = X_k$ where k is the largest integer so that $|X_k| \leq \min\{\lfloor \rho \cdot \ell \rfloor, |A| + |B|\}$ and $|\mathcal{I}(X_k)| \geq k$. Let X^* be an optimal solution for Dℓ-S. Note that $|\mathcal{I}(X)| \geq |\mathcal{I}(X^*)|$ and that $\frac{\ell(\ell-1)}{|X|(|X|-1)} \geq 1/(2\rho^2)$. By Lemma 3.1 we can find in polynomial time $X' \subseteq X$ so that $|X'| = \ell$ and $|\mathcal{I}(X')| \geq |\mathcal{I}(X)| \cdot \frac{\ell(\ell-1)}{|X|(|X|-1)} \geq |\mathcal{I}(X^*)| \cdot 1/(2\rho^2)$. Thus X' is a $1/(2\rho^2)$-approximation for the original bipartite Dℓ-S instance.

3.2 Part (ii): Polynomial algorithm for MPk-EDPA

We now prove part (ii) of Theorem 1.3. It would be convenient to describe the algorithm using "mixed" graphs that contain both directed and undirected edges. Given such mixed graph with weights on the nodes, a minimum weight path between two given nodes can be found in polynomial time using Dikjstra's algorithm and elementary constructions (namely, replacing every undirected edge by two opposite directed edges, and converting node weights to edge weights). The algorithm for undirected MPk-EDPA is as follows (see Fig. 3).

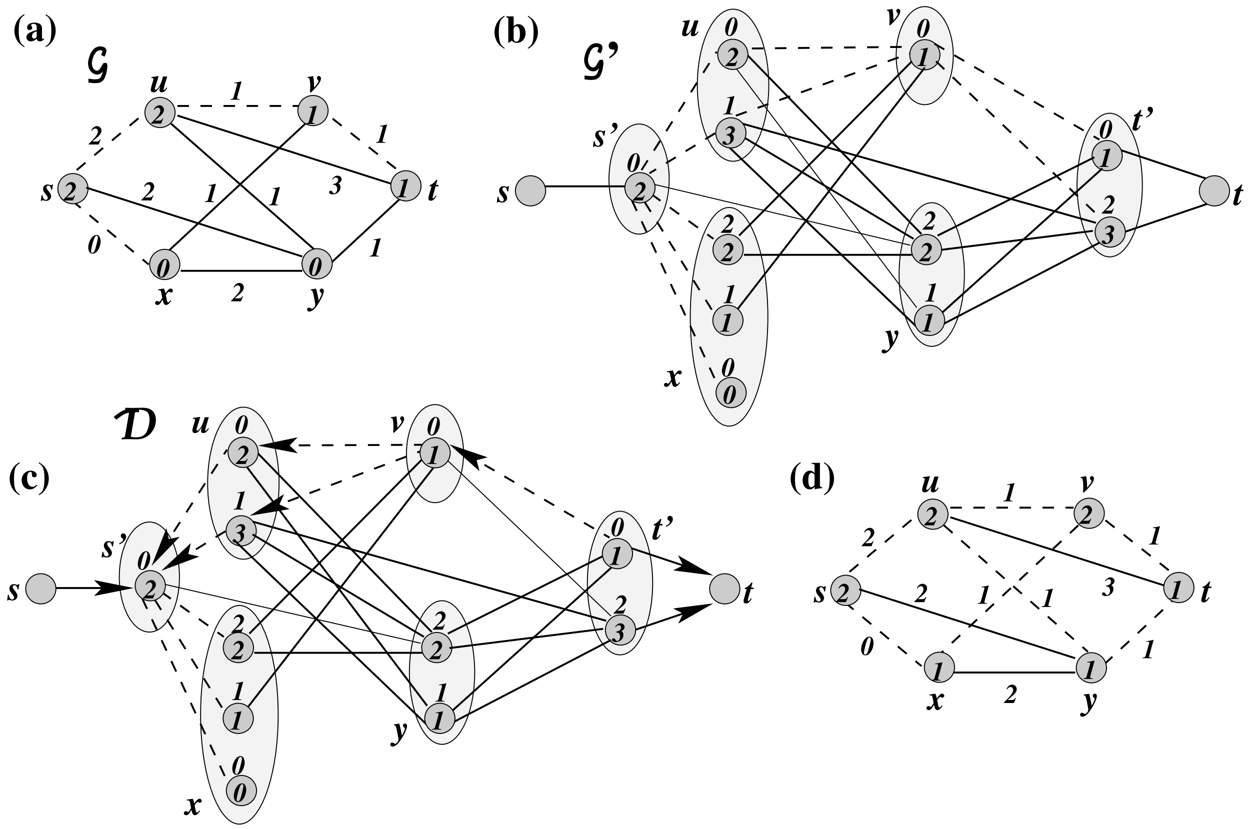

Figure 3: (a) The input graph $\mathcal{G}$ and its subgraph G_0; $k = 2$, edges of G_0 are shown by dashed lines, the number in each node v is $p_0(v)$. (b) The graph $\mathcal{G}'$; the number above each node is its weight $w(v)$. (c) The graph $\mathcal{D}'$; the optimal augmenting path is $P = (s, s', x, b, a, y, t', t)$ has weight 2, the power of each one of x, y is increased by 1. (d) An optimal solution $E_0 + F$ is shown by dashed lines, $F = \{xv, yu\}$, the two edge-disjoint paths are $(s, x, v, t), (s, u, y, t)$.

1. Construct a graph $\mathcal{G}'$ from $\mathcal{G}$ as follows (see Fig. 3a,b). Let $p_0(v)$ be the power of v in G_0. For every $v \in V$ do the following. Let $p_0(v) \le c_1 < c_2 < \cdots$ be the costs of the edges in $\delta_{\mathcal{E}}(v)$ of cost at least $p_0(v)$ sorted in increasing order. For every c_j add a

node v_j of the weight $w(v_j) = c_j - p_0(v)$. Then for every $u_{j'}, v_{j''}$ add an edge $u_{j'}v_{j''}$ if $w(u_{j'}), w(v_{j''}) \geq c(uv)$. Finally, add two nodes s, t and an edge from s to every s_j and from every t_j to t.

2. Construct a mixed graph $\mathcal{D}$ from $\mathcal{G}'$ as follows (see Fig. 3c). Let I be an inclusion minimal edge set in G_0 that contains $k-1$ pairwise edge-disjoint st-paths. Direct those paths from t to s, and direct accordingly every edge of $\mathcal{G}'$ that corresponds to an edge in I.

3. In $\mathcal{D}$, compute a minimum weight st-path P (see Fig. 3c,d). Return the set of edges of $\mathcal{G}$ that correspond to P that are not in E_0.

We now explain why the algorithm is correct. It is known that the following "augmenting path" algorithm solves the **Min-Cost k Edge-Disjoint Paths Augmentation** problem (the min-cost version of MPk-EDPA, where the edges in G_0 have cost 0) in undirected graphs (c.f., [6]).

1. Let I be an inclusion minimal edge set in G_0 that contains $k-1$ pairwise edge-disjoint st-paths. Construct a mixed graph $\mathcal{D}$ from $\mathcal{G}$ by directing these paths from t to s.

2. Find a min-cost path P in $\mathcal{D}$. Return $P - E_0$.

Our algorithm for MPk-EDPA does the same but on the graph $\mathcal{G}'$. The key point is that in $\mathcal{G}'$ the weight of a node is the increase of its power caused by taking an edge incident to this node. For example, in Fig. 3a, if we choose the edge xy, then the corresponding edge in $\mathcal{G}'$ and in $\mathcal{D}$ in Fig. 3b,c connects two nodes of the weight 2. It can be shown that for any feasible solution F corresponds a unique path P in $\mathcal{D}$ so that $p(G_0+F) - p(G_0) = w(P)$, and vice versa. As we choose the minimum weight path in $\mathcal{D}$, the returned solution is optimal.

The proof of theorem 1.3 is complete.

4 Hardness of directed min-power edge-connectivity problems (Proof of Theorem 1.4)

In this section we give the proof of Theorem 1.4 which establishes that many directed min-power edge-connectivity problems are unlikely to admit a polylogarithmic approximation ratio.

4.1 Arbitrary costs

For simplicity of exposition, we first prove the statement for arbitrary costs, not necessarily symmetric. For that, we use the hardness result for MPk-EDP given in Theorem 1.2, to show a similar hardness for the other three problems MPk-EOS, MPk-EIS, and MPk-ECS. Loosely speaking, we show that each of directed MPk-EOS/MPk-EIS is at least as hard as MPk-EDP, and that MPk-ECS is at least as hard as MPk-EOS.

We start by describing how to reduce directed MPk-EDP to directed MPk-EOS (see Fig 4(a)). Given an instance $\mathcal{G} = (V, \mathcal{E}), c, (s,t), k$ of MPk-EDP construct an instance of $\mathcal{G}' = (V', \mathcal{E}'), c', s, k$ of directed MPk-EOS as follows. Add to $\mathcal{G}$ a set $U = \{u_1, \ldots, u_k\}$ of k new nodes, and then add an edge set E_0 of cost zero: from t to every node in U, and from every node in U to every node $v \in V - \{s,t\}$. That is

$$V' = V + U = V + \{u_1, \ldots, u_k\},$$

$$\mathcal{E}' = \mathcal{E} + E_0 = \mathcal{E} + \{tu : u \in U\} + \{uv : u \in U, v \in V' - \{s,t\}\},$$

$$c'(e) = c(e) \text{ if } e \in \mathcal{E} \text{ and } c'(e) = 0 \text{ otherwise .}$$

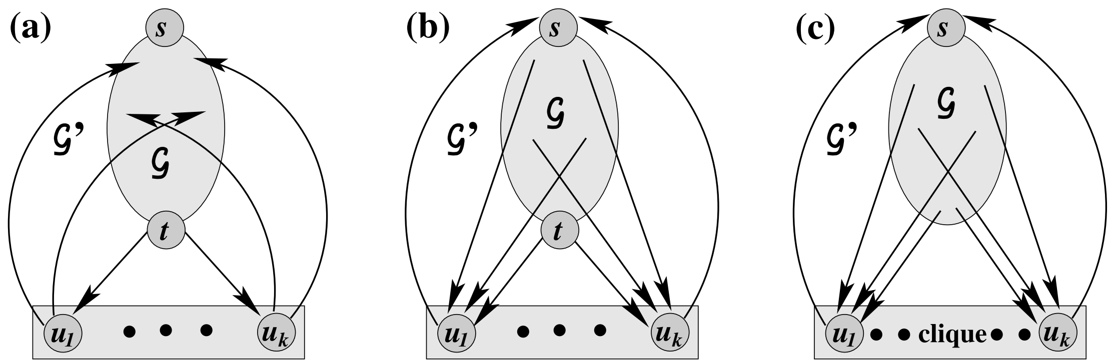

Figure 4: Reductions for asymmetric costs: (a) directed MPk-EDP to directed MPk-EOS; (b) directed MPk-EDP to directed MPk-EIS; (c) directed MPk-EOS to directed MPk-ECS.

Claim 4.1 *$G = (V, E)$ is a solution to the* MPk-EDP *instance if, and only if, $G' = (V, E' = E + E_0)$ is a solution to the constructed* MPk-EOS *instance.*

Proof: Let E be a solution to the MPk-EDP instance and let $\Pi = \{P_1, \ldots, P_k\}$ be a set of k pairwise edge disjoint st-paths in E. Then in $G' = (V', E + E_0)$ for every $v \in V' - s$ there is a set $\Pi' = \{P'_1, \ldots, P'_k\}$ of k pairwise edge-disjoint sv-paths as follows. If $v = t$ then set $\Pi' = \Pi$. If $v \neq s$ then set $P'_j = P_j + tu_j + u_jv$ for every $j = 1, \ldots, k$.

Now let $E' = E + E_0$ be a solution to constructed MPk-EOS instance. In particular, (V', E') contains a set Π of k-edge disjoint st-paths, none of which has t as an internal node. Consequently, no path in Π passes through U, as t is the tail of every edge entering U. Thus Π is a set k-edge disjoint st-paths in G, namely, $G = (V, E)$ is a solution to the original MPk-EDP instance. $\square$

Since in the construction $|V'| = |V| + k \leq |V| + |V|^2 \leq 2|V|^2$, Theorem 1.3 together with Claim 4.1 implies the first part of Theorem 1.4 for MPk-EOS.

Asymmetric MPk-EIS: The reduction of asymmetric MPk-EIS to MPk-EDP is similar to the one of MPk-EOS described above, except that here set $E_0 = \{us : u \in U\} + \{vu : v \in V - \{s,t\}, u \in U\}$ (see Fig 4(b)); namely, connect every $u \in U$ to s, and every $v \in V - \{s,t\}$ to every $u \in U$. Then in the obtained MPk-EIS instance, require k internally edge-disjoint vt-paths for every $v \in V$, namely, we seek a graph that is k-edge-inconnected to t. The other parts of the proof for MPk-EIS are identical to those for MPk-EOS described above.

Asymmetric MPk-ECS: Reduce the directed MPk-EOS to the directed MPk-ECS as follows (see Fig 4(c)). Let $\mathcal{G} = (V, \mathcal{E}), c, s, k$ be an instance of MPk-EOS. Construct an instance of $\mathcal{G}' = (V', \mathcal{E}'), c', s, k$ of MPk-EOS as follows. Add to $\mathcal{G}$ a set $U = \{u_1, \ldots, u_k\}$ of k new nodes, and then add an edge set $E_0 = \{uu' : u, u' \in U\} + \{vu : v \in V - s, u \in U\} + \{us : u \in U\}$ of cost 0; namely, E_0 is obtained by taking a complete graph on U and adding all edges from $V - s$ to U and all edges from U to s. It is not hard to verify that if $E \subseteq \mathcal{E}$, then $G = (V, E)$ is k-edge-outconnected from s if, and only if, $G' = (V', E_0 + E)$ is k-edge-connected.

4.2 Symmetric costs

We now show that these directed problems MPk-EDP, MPk-EOS, MPk-EIS, MPk-ECS is hard to approximate even for symmetric costs. We start with directed symmetric MPk-EDP. We use a refinement of Theorem 1.2 of [16]. In [16] it is shown that the hardness result in Theorem 1.2 for directed MPk-EDP holds for simple graphs with costs in $\{0, n^3\}$, where $n = |V|$. If we change the cost of every edge of cost 0 to 1, it will add no more then n^2/n^3 to the total cost of any solution that uses at least one edge of cost n^3. Thus we have the following refinement of Theorem 1.2:

Corollary 4.2 ([16]) *Directed* MPk-EDP *on simple graphs with costs in* $\{1, n^3\}$ *cannot be approximated within* $O(2^{\log^{1-\varepsilon} n})$ *for any fixed* $\varepsilon > 0$, *unless* NP $\subseteq$ DTIME$(n^{\mathrm{polylog}(n)})$.

We show that a ρ-approximation algorithm for directed symmetric $\mathsf{MP}k\text{-}\mathsf{EDP}$ implies a ρ-approximation algorithm for directed $\mathsf{MP}k\text{-}\mathsf{EDP}$ with costs in $\{1, n^3\}$, for any $\rho < n^{1/7}$. Let $\mathcal{G} = (V, \mathcal{E}), c, (s,t), k$ be an instance of $\mathsf{MP}k\text{-}\mathsf{EDP}$ with costs in $\{1, n^3\}$. Let opt be an optimal solution value for this instance. Note that $\mathsf{opt} \leq n^4$. Let $N = n^5$. Define an instance $\mathcal{G}' = (V', \mathcal{E}'), c', (s,t), k' = kN$ for directed symmetric $\mathsf{MP}k\text{-}\mathsf{EDP}$ as follows. To obtain $\mathcal{G}', c'$ from $\mathcal{G}, c$ do the following. First, obtain $\mathcal{G}^+ = (V', \mathcal{E}^+), c^+$ by replacing every edge $e = uv \in E$ by N internally-disjoint uv-paths of the length 2 each, where the cost of the first edge in each paths is $c(e)$ and the cost of the second edge is 0 (see Fig. 5). Second, to obtain a symmetric instance $\mathcal{G}', c'$, for every edge $ab \in \mathcal{E}^+$ add the opposite edge ba of the same cost as ab.

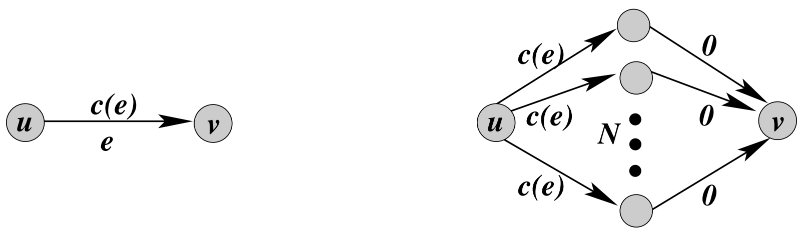

Figure 5: Reductions for symmetric costs: transforming an edge into N paths.

For a path P^+ in $\mathcal{E}^+$, let $\psi(P^+)$ denote the unique path in $\mathcal{E}$ corresponding to P^+. For any path P in $\mathcal{E}$, the paths in the set $\psi^{-1}(P)$ of the paths in $\mathcal{E}^+$ that corresponds to P are edge-disjoint. Hence, any set Π of paths in $\mathcal{E}$ is mapped by ψ^{-1} to a set $\Pi^+ = \psi^{-1}(\Pi)$ of exactly $N|\Pi|$ edge-disjoint paths in $\mathcal{E}^+$ of the same power, namely $|\Pi^+| = N|\Pi|$ and $p_c(\Pi) = p_{c^+}(\Pi^+)$. Conversely, any set Π^+ of paths in $\mathcal{E}^+$ is mapped by ψ to a set $\Pi = \psi(\Pi^+)$ of at least $\lceil |\Pi^+|/N \rceil$ edge-disjoint paths in $\mathcal{E}$ of the same power, namely, $|\Pi| = \lceil |\Pi|/N \rceil$ and $p_c(\Pi) = p_{c^+}(\Pi^+)$. In particular:

Corollary 4.3 $\mathsf{opt}' \leq \mathsf{opt} \leq n^4$, *where* opt' *is an optimal solution value for the obtained instance.*

Note that $|V'| = n' \leq n^7$, hence to prove Theorem 1.4 for directed symmetric $\mathsf{MP}k\text{-}\mathsf{EDP}$ it is sufficient to prove that a $\rho(n')$-approximation algorithm for $\mathcal{G}', c', (s,t), k'$ with $\rho(n') < n'^{1/7}$ implies a $\rho(n)$-approximation algorithm for the original instance. Suppose that we have a $\rho(n')$-approximation algorithm that computes an edge set $E' \subseteq \mathcal{E}'$ that contains a set Π' of kN edge-disjoint paths in $\mathcal{G}'$ of power $p_{c'}(E') \leq \rho \cdot \mathsf{opt}'$, where $\rho = \rho(n') < n'^{1/7} \leq n$. Then $|E' - \mathcal{E}^+| \leq \rho \cdot \mathsf{opt}' \leq \rho \cdot n^4$, since every edge in $|E' - \mathcal{E}^+|$ adds at least one to $p_{c'}(E')$. Consequently, there is a set $\Pi^+ \subseteq \Pi'$ of at least $kN - \rho \cdot n^4$ paths in Π that are contained in $E^+ = E' \cap \mathcal{E}^+$. Hence, since $\rho = \rho(n') > n'^{1/7} \geq n$, the number of paths in $\Pi = \psi(\Pi^+)$ is at least

$$|\Pi| \geq \left\lceil \frac{kN - \rho \cdot n^4}{N} \right\rceil \geq \left\lceil k - \frac{\rho \cdot n^4}{N} \right\rceil = \left\lceil k - \frac{\rho}{n} \right\rceil \geq k \ .$$

Consequently, the set E of edges of Π is a feasible solution for $\mathcal{G}, c, (s,t), k$ of power at most

$$p_c(E) \leq p_{c'}(E') \leq \rho \mathsf{opt}' \leq \rho \mathsf{opt} \ .$$

Since in the construction $|V'| \leq |V|^7$, Corollary 4.2 implies Theorem 1.4 for directed symmetric MPk-EDP.

The proof for the other problems MPk-EOS, MPk-EIS, and MPk-ECS, is similar, with the help of reductions described for the asymmetric case. E.g., for MPk-EOS, we start with an instance of MPk-EDP with costs in $\{0, n^3\}$, and reduce it to an instance of (asymmetric) MPk-EOS with costs in $\{0, n^3\}$. Then we change the cost of every edge of cost 0 to 1, "blow-up" every edge e into sufficiently large number of paths of length 2, set the costs of the first edge to $c(e)$ and of the second to 0 in each path, add "symmetric completion", and continue in the same way as for the symmetric MPk-EDP.

The proof of theorem 1.3 is complete.

5 Approximation algorithm for MPk-OS (Proof of Theorem 1.5)

In this section we consider undirected graphs only. An edge set F on V is an *ℓ-cover* (of V) if $\deg_F(v) \geq \ell$ for every $v \in V$, where $\deg_F(v)$ is the degree of v w.r.t. F. We prove the following general statement, which is of independent interest, as it shows that undirected MPk-OS and Min-Power $(k-1)$-Edge Cover are almost equivalent w.r.t. approximation.

Theorem 5.1

(i) *If there exist a ρ-approximation algorithm for the Min-Power $(k-1)$-Edge-Cover problem, then there exists a $(\rho+4)$-approximation algorithm for undirected* MPk-OS.

(ii) *If there exists a ρ-approximation algorithm for undirected* MPk-OS/MPk-CS *then there exists a ρ-approximation algorithm for the Min-Power $(k-1)$-Edge Cover problem.*

Theorem 1.5 follows from Part (i) of Theorem 5.1 and the fact that Min-Power $(k-1)$-Edge-Cover admits a $\min\{k, O(\log n)\}$-approximation algorithm [16, 20].

In the rest of this section we prove Theorem 5.1. Given a graph G which is k-outconnected from s, let us say that an edge e of G is *critical* if $G - e$ is not k-outconnected from s. We need the following fundamental statement:

Theorem 5.2 ([4]) *In a k-outconnected from s undirected graph G, any cycle in which every edge is critical contains a node $v \neq s$ whose degree in G is exactly k.*

The following corollary (e.g., see [4]) is used to get a relation between $(k-1)$-edge covers and k-outconnected subgraphs.

Corollary 5.3 *If $\deg_J(v) \geq k-1$ for every node v of an undirected graph J, and if F is an inclusion minimal edge set such that $J \cup F$ is k-outconnected from s, then F is a forest.*

Proof: If not, then F contains a cycle C of critical edges, but every node of this cycle is incident to 2 edges of C and to at least $k-1$ edges of J, contradicting Theorem 5.2. $\square$

Proof of Theorem 5.1: We start by proving Part (i). By the assumption, we can find a subgraph J with $\deg_J(v) \geq k-1$ of power at most $p(J) \leq \rho\mathsf{opt}$. We reset the costs of edges in J to zero, and apply a 2-approximation algorithm for the Min-Cost k-Outconnected Subgraph problem (c.f., [13]) to compute an (inclusion) minimal edge set F so that $J+F$ is k-outconnected from s. By Corollary 5.3, F is a forest. Thus $p(F) \leq 2c(F) \leq 4\mathsf{opt}$, by Proposition 1.1. Combining, we get Part (i).

We now prove Part (ii). The reduction for MPk-CS is a s follows. Let $\mathcal{G} = (V, \mathcal{E}), c$ be an instance of Min-Power $(k-1)$-Edge Cover with $|V| \geq k$. Construct an instance $\mathcal{G}', c'$ for MPk-CS as follows. Add a copy V' of V and the edges $\{vv' : v \in V\}$ of cost 0 ($v' \in V'$ is the copy of $v \in V$), and then add a clique of cost 0 on V'. Let $\mathcal{E}'$ be the edges of $\mathcal{G}' - \mathcal{E}$. We claim that $E \subseteq \mathcal{E}$ is a $(k-1)$-Edge Cover in $\mathcal{G}$ if, and only if, $G' = (V+V', E+\mathcal{E}')$ is k-connected.

Suppose that G' is k-connected. Then $\deg_{E+\mathcal{E}'}(v) \geq k$ and $\deg_{E'}(v) = 1$ for all $v \in V$. Hence $\deg_E(v) \geq k-1$ for all $v \in V$, and thus E is a $(k-1)$-edge-cover in G.

Suppose that $E \subseteq \mathcal{E}$ is a $(k-1)$-edge cover in $\mathcal{G}$. We will show that G' has k internally disjoint vu-paths for any $u, v \in V+V'$. It is clear that $G'-E$, and thus also G, has k internally disjoint vu-paths for any $u, v \in V'$. Let $v \in V$. Consider two cases: $u \in V'$ and $u \in V$. Assume that $u \in V'$. Every neighbor v_i of v in (V, E) defines the vu path (v, v_i, v_i', u) (possibly $v_i' = u$), which gives $\deg_E(v) \geq k-1$ internally disjoint vu-paths. An additional path is (v, v', u). Now assume that $u \in V$. Every common neighbor a of u and v defines the vu-path (v, a, u), and suppose that there are q such common neighbors. Each of v and u has at least $k-1-t$ more neighbors in G, say $\{v_1, \ldots, v_{k-1-q}\}$ and $u_1, \ldots, u_{k-1-t}$, respectively. This gives $k-1-q$ internally disjoint vu-paths (v, v_i, v_i', u_i', u), $i = 1, \ldots, k-1-q$. An additional path is (v, v', u', u). It is easy to see that these k vu-paths are internally disjoint. The proof for MPk-CS is complete.

The reduction for $\mathsf{MP}k$-OS is the same, except that in the construction of $\mathcal{G}'$ we also add a node s and edges $\{sv' : v' \in V'\}$ of cost 0. $\square$

The proof of theorem 1.5 is complete.

6 Conclusions

In this paper we considered four fundamental min-power edge connectivity problems: $\mathsf{MP}k$-EDP, $\mathsf{MP}k$-EOS/$\mathsf{MP}k$-EIS, and $\mathsf{MP}k$-ECS. For both directed and undirected graphs, we have shown that none of these problems is unlikely to admit a polylogarithmic approximation ratio, and for directed graphs this is so even if the costs are symmetric. In the undirected case, we showed that a polylogarithmic approximation ratio for one of these problems implies a polylogarithmic approximation for the ℓ-Densest Subgraph problem. For directed graphs, we showed that a polylogarithmic approximation is unlikely unless NP-hard problems can be solved in quasi-polynomial time. In contrast, we showed that for undirected graphs the augmentation version $\mathsf{MP}k$-EDPA of $\mathsf{MP}k$-EDP, where the goal is only to increase the st-connectivity by 1, can be reduced to the shortest path problem. The same result holds for directed graphs, see [25].

We now list some open problems, that follow from Table 1. Most of them concern node-connectivity problems. One of the main open problem seems to determine whether the *undirected* $\mathsf{MP}k$-DP is in P or is NP-hard (as was mentioned, the directed $\mathsf{MP}k$-DP is in P, c.f., [16]). In fact, we do not even know whether the augmentation version $\mathsf{MP}k$-DPA of undirected $\mathsf{MP}k$-DP is in P, but we conjecture this is so. We note that a polynomial algorithm for undirected $\mathsf{MP}k$-DPA can be used to improve the currently best known ratios for the undirected Min-Power k-Connected Subgraph problem: from 9 (follows from [21]) to $8\frac{1}{3}$ for $k = 4$, and from 11 (follows from [7]) to $10\frac{1}{3}$ for $k = 5$. This is achieved as follows. In [7] it is shown that any graph G that is k-outconnected from a node r of degree k is $(\lceil k/2 \rceil + 1)$-connected; furthermore, for $k = 4, 5$, G contains two nodes s, t so that increasing the connectivity between them by one results in a k-connected graph. Hence for $k = 4, 5$, we can get approximation ratio $\alpha + \beta$, where α is the best known ratio for undirected $\mathsf{MP}k$-OS, and β is the best known ratio for undirected $\mathsf{MP}k$-DPA. As $\mathsf{MP}k$-OS can be approximated within $2(k-1/3)$ [20], then if $\mathsf{MP}k$-DP augmentation is in P, we can get approximation ratio $2(k-1/3)+1$ for $\mathsf{MP}k$-CS with $k = 4, 5$, which is $8\frac{1}{3}$ for $k = 4$ and $10\frac{1}{3}$ for $k = 5$.

Another interesting question is the approximability of the directed $\mathsf{MP}k$-IS. Currently, we are not aware of any hardness result, while the best known approximation ratio is k. Except directed MPk-DP, there is still a large gap between upper and lower bounds of approximation for all the other min-power node connectivity problems, for both directed and undirected graphs.

References

[1] E. Althaus, G. Calinescu, I. Mandoiu, S. Prasad, N. Tchervenski, and A. Zelikovsky. Power efficient range assignment for symmetric connectivity in static ad-hoc wireless networks. *Wireless Networks*, 12(3):287–299, 2006.

[2] V. Auletta, Y. Dinitz, Z. Nutov, and D. Parente. A 2-approximation algorithm for finding an optimum 3-vertex-connected spanning subgraph. *Journal of Algorithms*, 32(1):21–30, 1999.

[3] G. Calinescu, S. Kapoor, A. Olshevsky, and A. Zelikovsky. Network lifetime and power assignment in ad hoc wireless networks. In *Proc. European Symposium on Algorithms (ESA), LNCS* 2832, pages 114–126, 2003.

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

[5] J. Cheriyan, S. Vempala, and A. Vetta. An approximation algorithm for the minimum-cost k-vertex connected subgraph. *SIAM Journal on Computing*, 32(4):1050–1055, 2003.

[6] W. J. Cook, W. H. Cunningham, W. R. Pulleyblank, and A. Scrijver. *Combinatorial Optimization*. Wiley, 1998.

[7] Y. Dinitz and Z. Nutov. A 3-approximation algorithm for finding an optimum 4,5-vertex-connected spanning subgraph. *Journal of Algorithms*, 32(1):31–40, 1999.

[8] J. Edmonds. Matroid intersection. *Annals of discrete Math.*, pages 185–204, 1979.

[9] U. Feige, G. Kortsarz, and D. Peleg. The dense k-subgraph problem. *Algorithmica*, (29):410–421, 2001.

[10] A. Frank. *Connectivity and network flows,* in *Handbook of Combinatorics,* eds. R. Graham, M. Grötschel, and L. Lovász, pages 111–177. Elsvier Science, 1995.

[11] A. Frank. Increasing the rooted-connectivity of a digraph by one. *Mathematical Programming*, 84(3):565–576, 1999.

[12] A. Frank. Rooted k-connections in digraphs, EGRES TR No 2006-07. 2001.

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

[14] H. N. Gabow. A representation for crossing set families with application to submodular flow problems. In *Symposium on Discrete Algorithms (SODA)*, pages 202–211, 1993.

[15] M. T. Hajiaghayi, N. Immorlica, and V. S. Mirokni. Power optimization in fault-tolerant topology control algorithms for wireless multi-hop networks. In *Proc. of the 9th conference on mobile computing and networking (MOBICOM)*, pages 300–312, 2003.

[16] M. T. Hajiaghayi, G. Kortsarz, V. S. Mirokni, and Z. Nutov. Power optimization for connectivity problems. In *Integer Programming and Combinatorial Optimization (IPCO)*, pages 349–361, 2005.

[17] D. S. Johnson. Approximation algorithms for combinatorial problems. *Joural of Computing and System Sciences*, 9(3):256–278, 1974.

[18] S. Khuller. *Approximation algorithms for for finding highly connected subgraphs,* Chapter 6 in *Approximation Algorithms for NP-hard problems,* D. S. Hochbaum Ed., pages 236-265. PWS, 1995.

[19] C. Klein and R. Ravi. A nearly best-possible approximation algorithm for node-weighted steiner trees. *Journal of Algorithms*, 19(1):104–115, 1995.

[20] G. Kortsarz, V. S. Mirokkni, Z. Nutov, and E. Tsanko. Approximation algorithms for minimum power wireless network design problems. *Manuscript*, 2006.

[21] G. Kortsarz and Z. Nutov. Approximating node-connectivity problems via set covers. *Algorithmica*, 37:75–92, 2003.

[22] G. Kortsarz and Z. Nutov. Approximating k-node connected subgraphs via critical graphs. *SIAM J. on Computing*, 35(1):247–257, 2005.

[23] G. Kortsarz and Z. Nutov. Approximating minimum cost connectivity problems, in *Approximation Algorithms and Metaheuristics,* T. F. Gonzalez ed.,. 2007. To appear.

[24] Z. Nutov. Approximating minimum power connectivity problems. *Manuscript*, 2006.

[25] Z. Nutov. Approximating minimum power covers of intersecting families and directed connectivity problems. In *Workshop on Approximation algorithms (APPROX), LNCS* 4110, pages 236–247, 2006.

[26] Z. Nutov. Approximating node weighted steiner networks. *Manuscript*, 2007.

[27] D. P. Williamson, M. X. Goemans, M. Mihail, and V. V. Vazirani. A primal-dual approximation algorithm for generalized steiner network problems. *Combinatorica*, 15:435–454, 1995.