The Open University of Israel
Department of Mathematics and Computer Science

Approximating minimum power network design 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
Elena Tsanko

Prepared under the supervision of **Dr. Zeev Nutov**

April 2006

Acknowledgments

I wish to thank my thesis advisor, Zeev Nutov, for his guidance, enthusiasm, valuable advices, and support, throughout development of the results presented in this Thesis. I benefited a lot from his professionalism and perspective in the field of approximation algorithms, and learned how to write mathematics in a laconic and clear way.

I am also grateful to Guy Kortsarz for his help in obtaining one of the main results of this work: the algorithm for underected Min-Power Edge-Multi-Cover Problem. Guy's role in overcoming the difficulties in this part of the research was crucial.

I am also thankful to my family for their support, patience and love.

Contents

Abstract

Given a (possibly directed) 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 in wireless networks, we consider several network design problems under the power minimization criteria. Given a graph $\mathcal{G} = (V, \mathcal{E})$ with costs on the edges and requirements $r(v)$ for each $v \in V$, the *Min-Power Edge-Multi-Cover* problem (MPEMC) is to find a min-power subgraph so that the degree (indegree, in the case of directed graphs) of every node v is at least $r(v)$; the *Power Budgeted Maximum Edge-Multi-Coverage* (PBMEMC) is to find a subgraph of total power budget at most P to satisfy the maximum amount of requirements. Our main result is an $O(\log n)$-approximation algorithms for MPEMC on both directed and undirected graphs. For directed graphs our ratio is tight, since the problem generalizes the min-cost Set-Multicover problem. For undirected graphs our result improves the previously best known $O(\log^4 n)$-approximation, and implies an $O(\log n + \alpha)$-approximation algorithm for the undirected *Min-Power k-Connected Subgraph* (MPk-CS) problem, where α is the best known approximation for the min-cost variant of the problem. (Currently, $\alpha = O(\ln k)$ for $n \geq 2k^2$ and $\alpha = O(\ln^2 k \cdot \min\{\frac{n}{n-k}, \frac{\sqrt{k}}{\ln n}\})$ otherwise.) We also give a $(1 - 1/e)$-approximation algorithm for the directed PBMEMC, and show that the undirected PBMEMC is at least as hard to approximate as the k-Densest Subgraph problem. Finally, we give a $4r_{\max}$-approximation algorithm for the undirected min-power Steiner Network problem: find a min-power subgraph that contains $r(u, v)$ pairwise edge-disjoint paths for every pair of nodes u, v.

1 Introduction

Wireless networks are an important subject of study due to their extensive applications. A large research effort focused on designing fault-tolerant networks 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 assignment to radio transmitters means to assign a set of powers to the nodes of the network. We consider finding a range assignment for the nodes of a network such that the resulting communication network satisfies some prescribed connectivity or degree properties, and such that the total power is minimized. The motivation for wireless networks of these problems is the same as of their min-cost variant for wired networks.

An important network property is fault-tolerance, which is usually measured by node-connectivity or the minimum degree/indegree of the network. These variants of fault-tolerant power-minimization problems were already extensively studied [1, 2, 5, 3, 4, 13, 14, 19]. The simplest undirected connectivity problem is when we require the network to be connected. In this case, the min-cost variant is just the min-cost spanning tree problem, while the min-power variant is APX-hard [14]. A 5/3-approximation algorithm for the min-power spanning tree problem is given in [1]. For recent results on undirected and directed graphs see [23] and [24], respectively.

This work is organized as follows: Section 2 defines the problems under consideration, highlights the relationships between them, reviews previously published results and shows that Min-Power k-Connected Subgraph Problem admits an $O(\alpha + \log n)$-approximation, where α is the approximation ratio for the undirected Min-Cost k-Connected Subgraph Problem. Section 3 presents a $(1 - 1/e)$-approximation algorithm for directed Power Budgeted Maximum Edge-Multi-Coverage Problem. Section 4 presents an $H(\Delta)$-approximation algorithm for directed Min-Power Edge-Multi-Cover Problem, where Δ is the maximum outdegree of a node in $\mathcal{G}$, and $H(k)$ denotes the kth Harmonic number. Section 5 presents an $O(\log n)$-approximation algorithm for undirected Min-Power Edge-Multi-Cover Problem. Finally, in Section 6 we discuss the Min-Power Steiner Network Problem and present a $4r_{\max}$-approximation algorithm fot it.

2 Problems considered in this research

We start by defining some necessary notation. Let $G = (V, E)$ be a (possibly directed) graph with cost $\{c(e) : e \in E\}$ on the edges. 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 in G leaving v. The power of the graph is the sum of powers of its nodes. Given an integral requirement function r on V, we say that G (or E) is an *r-edge cover* if for every $v \in V$: $d_G(v) \geq r(v)$ if G is undirected, and $d_G^+(v) \geq r(v)$ if G is directed, where $d_G(v) = d_E(v)$ is the *degree* of v in G if G is undirected, and $d_G^+(v) = d_E^+(v)$ is the *indegree* of v in G if G is directed.

2.1 Min-Power Edge-Multi-Cover and Power Budgeted Maximum Edge-Multi-Coverage

We start by considering the following two related problems:

Min-Power Edge-Multi-Cover (MPEMC):

Instance: A (possibly directed) graph $\mathcal{G} = (V, \mathcal{E})$ with cots on the edges and a requirement $r(v)$ for every $v \in V$.

Objective: Find a min-power subgraph $G = (V, E)$ of $\mathcal{G}$ so that G is an r-edge cover.

Power Budgeted Maximum Edge-Multi-Coverage (PBMEMC):

Instance: A graph $\mathcal{G} = (V, \mathcal{E})$ with costs $\{c(e) : e \in \mathcal{E}\}$ on the edges, weights $w(v)$ and requirements $r(v)$ for each $v \in V$, and budget P.

Objective: Find $E \subseteq \mathcal{E}$ with $p(E) \leq P$ and maximum $\mathsf{val}(E) = \sum_{v \in V} \min\{d_E^+(v), r(v)\} \cdot w(v)$.

For practical applications of MPEMC or PBMEMC consider the following scenario. Given a set A of "transmitters" and a set B of "clients", set the power of transmitters so that every client can receive a message. To ensure reliability of communication (fault-tolerance), we require that each $b \in B$ will be able to receive the message from at least $r(b)$ transmitters.

The directed MPEMC was mentioned in [14], but it seems that it was not studied before, although it is a fundamental problem that generalizes the classic Min-Cost Set-Multicover problem; the later is a particular case when for every node $v \in V$ the costs of the edges leaving v are the same. Thus by the hardness result of [25] the directed MPEMC has an $\Omega(\log n)$-approximation threshold, that is, it cannot be approximated within $C \ln n$ for some universal constant C. In the same way directed PBMEMC generalizes the Cost-Budgeted Maximum Coverage problem, that admits a $(1 - 1/e)$-approximation algorithm [18], which is tight unless P=NP. For the generalization considered here we prove:

Theorem 2.1 *Directed* PBMEMC *admits a* $(1-1/e)$*-approximation algorithm.*

Theorem 2.2 *The directed* MPEMC *admits an* $H(\Delta)$*-approximation algorithm, where* Δ *is the maximum outdegree of a node in* $\mathcal{G}$*, and* $H(k)$ *denotes the* k*th Harmonic number.*

We note that Theorem 2.1 easily implies an $H(R)$-approximation algorithm for directed MPEMC, where $R = \sum_{v \in V} r(v)$. However, the result proved in Theorem 2.2 is sharper: the approximation ratio is better and the algorithm is faster. Furthermore, we show a way to formulate the directed MPEMC as an integer program, and using the dual fitting method show that the solution computed is within an $H(\Delta)$ factor from a solution to the corresponding LP-relaxation, see Section 4.

Our result for undirected MPEMC, Theorem 2.3 to follow, was obtained in a joint research with G. Kortsarz. It is easy to see that the greedy algorithm which for every $v \in V$ picks the lightest $r(v)$ edges entering v is an $r_{\max}$-approximation algorithm for undirected MPEMC, where $r_{\max} = \max_{v \in V} r(v)$. In [14] it is proved that the undirected MPEMC is APX-hard, and that it admits an $O(\log^4 n)$-approximation algorithm. We prove:

Theorem 2.3 *The undirected* MPEMC *admits an* $O(\log n)$*-approximation algorithm.*

Our algorithm for the undirected MPEMC uses as a subroutine our algorithm for the directed PBMEMC, but this is not straightforward, and for power problems we do not see an easy way to deduce the undirected case (Theorem 2.3) from the directed one (Theorem 2.2). For example, 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, e.g., for MPEMC, 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.

The following statement shows that for undirected PBMEMC a good approximation algorithm (e.g., with a constant or a polylogarithmic approximation ratio) might not exist even for unit costs and unit weights. The *Densest* k*-Subgraph* problem is given a graph $\mathcal{G} = (V, \mathcal{E})$ to find a subgraph of $\mathcal{G}$ with k nodes and maximum number of edges. The best known approximation ratio for the Densest k-subgraph problem is roughly $n^{-1/3}$ [10], and in spite of numerous attempts to improve it, this ratio holds for almost 10 years.

Proposition 2.4 *If there exists a* ρ*-approximation algorithm for undirected* PBMEMC *with unit costs and unit weights, then there exist a* ρ*-approximation algorithm for the Densest* k*-Subgraph problem.*

Proof: Given an instance $\mathcal{G} = (V, \mathcal{E})$ of the Densest k-Subgraph problem, define an instance $\mathcal{G}, r, P$ with unit costs and unit weights for PBMEMC as follows: $r(v) = k - 1$ for all $v \in V$ and the power budget is $P = k$. Then the problem is to find a node subset $U \subseteq V$ with $|U| = k$ so that the number of edges in the subgraph induced by U in $\mathcal{G}$ is maximum. The later is the Densest k-Subgraph problem. □

2.2 Min-Power k-Connected Subgraph

A (simple) graph is k-connected if it contains k internally disjoint uv-paths between every pair u, v of its nodes. We also consider is the min-power variant of the undirected Min-Cost k-Connected Subgraph (MCk-CS) problem, namely (for results on directed graphs see [24]):

Min-Power k-Connected Subgraph (MPk-CS):
Instance: A graph $\mathcal{G} = (V, \mathcal{E})$ with costs on the edges, and an integer k.
Objective: Find a min-power k-connected spanning subgraph G of $\mathcal{G}$.

Min-cost connectivity problems were extensively studied, see surveys in [17] and [21]. The best known approximation ratios for MCk-CS is $O(\ln^2 k \cdot \min\{\frac{n}{n-k}, \frac{\sqrt{k}}{\ln k}\})$ for both directed and undirected graphs [20], and $O(\ln k)$ for undirected graphs with $n \geq 2k^2$ [7]. It turns out that approximating MPk-CS is closely related to approximating MCk-CS and MPEMC, as shows the following theorem from [14] (for its proof see the Appendix).

Theorem 2.5 ([14]) *(i) If there exists an α-approximation algorithm for the undirected Min-Cost k-Connected Subgraph problem and a β-approximation algorithm for undirected* MPEMC *then there exists a $(2\alpha + \beta)$-approximation algorithm for* MPk-CS.

(ii) If there exists a ρ-approximation for undirected MPk-CS *then there exists a $(2\rho + 1)$-approximation for the Min-Cost k-Connected Subgraph problem.*

In [14] is given an $O(\log^4 n)$-approximation algorithm for undirected MPEMC. Thus, an $O(\alpha + \log^4 n)$-approximation algorithm for undirected MPk-CS is derived from Theorem 2.5. We use our result for undirected MPEMC (Theorem 2.3) to conclude:

Corollary 2.6 *An α-approximation algorithm for the undirected Min-Cost k-Connected Subgraph problem implies an $O(\alpha + \log n)$-approximation algorithm for the undirected* MPk-CS.

Combined with part (ii) of Theorem 2.5, we get that for undirected graphs, MPk-CS and MCk-CS are equivalent with respect to approximation (up to constants), unless MCk-CS admits better than $O(\log n)$ approximation ratio. For most values of k the best known approximation ratio for MCk-CS is $\Theta(\log^2 n)$ [20], for $k \leq \sqrt{n/2}$ in undirected graphs an $O(\log n)$ ratio is known [7], while for $k > \log n$ a better than $O(\log n)$ ratio is not known.

2.3 Min-Power Steiner Network

The last problem we study is:

Min-Power Steiner Network (MPSN):

Instance: A graph $\mathcal{G} = (V, \mathcal{E})$ with costs on the edges and requirement $r(u, v)$ for every node pair $u, v \in V$.

Objective: Find a min-power subgraph G of $\mathcal{G}$ so that G contains $r(u, v)$ pairwise edge-disjoint uv-paths for every $u, v \in V$.

Williamson et. al. [26] gave a $2r_{\max}$-approximation algorithm for the min-cost Steiner Network problem, and then this was improved to $2H(r_{\max})$ in [12]. The currently best known approximation ratio for the min-cost Steiner Network problem is 2 due to Jain [16]. We show that the algorithm of [26, 12] for the min-cost case, has approximation ratio $4r_{\max}$ for the min-power variant MPSN.

Theorem 2.7 *Undirected* MPSN *admits a* $4r_{\max}$*-approximation algorithm.*

If $r_{\max}$ is a constant, say $r_{\max} \in \{0, 1, 2\}$, then the approximation ratio in Theorem 2.7 is a constant. The approximation ratio may seem weak if $r_{\max}$ is large. However, it might be that a polylogarithmic approximation algorithm does not exist: in [14] it was shown that the directed min-power variant cannot be approximated within $O(2^{\log^{1-\varepsilon} n})$ for any fixed $\varepsilon > 0$, unless NP $\subseteq$ DTIME$(n^{\text{polylog}(n)})$. This hardness result is valid even when there is one pair u, v with $r(u, v) > 0$, a case that can be solved optimally in the min-cost variant, c.f., [8].

2.4 Notation and Preliminaries

In the rest of this section we give some notation and preliminaries used in the work. Let $G = (V, E)$ be a 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, and let $d(X, Y) = |\delta_G(X, Y)|$ be the number of edges in G going from X to Y. We sometimes omit the subscripts G, E if they are clear from the context. For brevity, $\delta_E(X) = \delta_E(X, V - X)$, and $d_E(X) = |\delta_E(X)|$ is the *degree of* X. Let $\Gamma_G(X) = \{u \in V - X : v \in X, vu \in E\}$ be the set of *neighbors* of X. Given edge costs $\{c_e : e \in E\}$, the *power* $p_G(v) = p_E(v)$ *of a node* v in G is the maximum cost of an edge incident to v in E, that is, $p(v) = \max_{e \in \delta_E(v)} c(e)$. Note that in the undirected case $p(v)$ is the maximum cost of an edge touching v while in the directed case it is the maximum cost of an edge leaving v. The *power of* G is $p(G) = p_E(V) = \sum_{v \in V} p(v)$. Note that $p(G)$ differs from the ordinary cost $c(G) = \sum_{e \in E} c(e)$ of G even for unit costs. In this case, if G has no isolated nodes 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 $c(G)/\sqrt{m/2}$. The following statement shows that these are the extremal cases also for general edge costs.

Proposition 2.8 ([14]) *For any graph $G = (V, E)$ holds: $c(G)/\sqrt{|E|/2} \leq p(G) \leq 2c(G)$. For a forest T, $c(T) \leq p(T) \leq 2c(T)$.*

Proposition 2.8 and the Theorem 2.5 were proved in [14]. For completeness of exposition, we restate the proofs in the Appendix.

Throughout the work, let $\mathcal{G} = (V, \mathcal{E})$ denote 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; otherwise our algorithms can be easily modified to return an error message. Let opt denote the optimal solution value of an instance at hand.

3 A $(1-1/e)$-approximation for directed Power Budgeted Maximum Edge-MultiCoverage

In this section we prove Theorem 2.1.

Given an instance of directed PBMEMC, apply the following transformation. For each node v with $r(v) > 0$ add to $\mathcal{G}$ a copy v' of v and redirect all the edges entering v to enter v', keeping their costs. Furthermore, for every $v \in V$ do the following. Let $\{e_1, \ldots, e_k\}$ be the edges in $\delta_{\mathcal{E}}(v)$ sorted by increasing costs. For every e_i add a node a_i of cost $c(e_i)$ and for every edge vu' of cost $\leq c(e_i)$ add an edge $a_i u'$. Finally, consider the corresponding underlying graph. Thus our problem can be reformulated as follows:

Instance: A bipartite graph $\mathcal{G} = (A + B, \mathcal{E})$, costs $\{c(a) : a \in A\}$, budget P, requirements $\{r(b) : b \in B\}$, weights $\{w(b) : b \in B\}$, and a partition $\mathcal{A}$ of A satisfying

(*) for every $A^i \in \mathcal{A}$ there exists an ordering $a_1, a_2, \ldots$ of A^i so that $\Gamma(a_{j-1}) \subseteq \Gamma(a_j)$.

Objective: Find $S \subseteq A$ with $c(S) \leq P$ and maximum $\mathsf{val}(S) = \sum_{b \in B} \min\{d(S, b), r(b)\} \cdot w(b)$ so that

(**) $|S \cap A^i| \leq 1$ for every $A^i \in \mathcal{A}$.

Remark: The same problem with unit requirements but without Property (*) is well known, and a 2-approximation algorithm for it was given by Chekuri and Kumar [6]. We do not know whether without property (*) one can achieve a constant approximation ratio for arbitrary requirements. In the special case when Property $(*)$ holds, we give a $(1-1/e)$-approximation algorithm for arbitrary requirements, and, in particular, improve the 2-approximation of [6] for unit requirements.

Our algorithm and the proof of the approximation ratio is similar to the ones in [18] where the ordinary (Cost) Budgeted Maximum Coverage was considered. There is a difference in that our algorithm is a local search (replacement) algorithm, while [18] only adds elements. But taking the analysis of [18] with replacing the costs of element added by their costs minus costs of elements deleted, the analysis carries through. For completeness of exposition, we give a full proof.

Note that in each part the costs defined by the ordering in (*) are strictly increasing. Clearly, we may assume that $c(a) \leq P$ for every $a \in A$. For $S \subseteq A$ and $b \in B$ let $r_S(b) = \max\{r(b) - d(S, b), 0\}$ be the *residual requirement* of b w.r.t. S (so $r(b) = r_\emptyset(b)$). $S \subseteq A$ is a feasible solution if $c(S) \leq P$ and (**) holds for S.

Let S satisfy (**), and set $s^i = A^i \cap S$ (possibly $s^i = \emptyset$). Let $B_S = \{b \in B : r_S(b) > 0\}$ be the set of *deficient* nodes w.r.t. S. For $a \in A^i$ with $c(a) > c(s^i)$ the *density of a w.r.t. S* is:

$$\sigma_{c,w}(S, a) = \frac{\mathsf{val}(S - s^i + a) - \mathsf{val}(S)}{c(a) - c(s^i)} = \frac{w((\Gamma(a) - \Gamma(s^i)) \cap B_S|)}{c(a) - c(s^i)} \ .$$

The algorithm uses the following procedure, which receives a feasible solution $S_0 \subseteq A$ and returns a feasible solution $S \subseteq A$ that contains S_0.

Procedure GREEDY(S_0)

Initialization: $S \leftarrow S_0$, $r \leftarrow r_{S_0}$, and remove from A the parts corresponding to S_0.

While $A \neq \emptyset$ do:

1. Find $a \in A$ of maximum density, and let A^i be the part with $a \in A^i$.
2. If $c(S - s^i + a) \leq P$ then $S \leftarrow S - s^i + a$, where $s^i = A^i \cap S$ (possibly $s^i = \emptyset$).
3. $A \leftarrow A - a$.

End While

The algorithm for directed PBMEMC is as follows. Let $k > e$ be some fixed integer.

Algorithm for PBMEMC

1. For every feasible $S_0 \subseteq A$ with $|S_0| \leq k$ do GREEDY(S_0).
2. Among the sets S returned, output one with maximum $\mathsf{val}(S)$.

Clearly, the algorithm can be implemented in polynomial time for any fixed integer k (we set $k = 3$). We now prove that the approximation ratio is $(1 - 1/e)$.

Remark: It may seem futile to start with some "best" triplet of elements going over all possible triplets. The goal of these three elements is to overcome a "knapsack type" difficulty the algorithm encounters. The fact that the elements have costs and the budget bound P creates a problem with the last element GREEDY tries to add. Thus, if we are able to add this element (say at exactly cost P) there would be no need for "guessing" the "correct" first three elements. However, since the last element may create a budget overflow, it can not be taken. The selection of the "correct" three first elements compensate for the last element not being added. We remark that with the choice of $k = 0$ the ratio is *unbounded*, with $k = 1$ the ratio is $\frac{1}{2}(1 - 1/e)$ and with $k = 2$ the ratio is $1/2$.

Let OPT be an optimal solution. Clearly, if $|\mathsf{OPT}| \leq k$ the algorithm returns an optimal solution. Henceforth assume $|\mathsf{OPT}| > k$. Let $s_1, s_2, \ldots$ be an ordering of nodes of OPT such that the total weight of uncovered elements, covered by the nodes is not increased.

Consider the computation at Step 1 of the algorithm when $S_0 = \{s_1, s_2, \ldots, s_k\}$ was considered. Let $\mathsf{OPT}' = \mathsf{OPT} - S_0$ and $P' = P - c(S_0)$. Let ℓ be the number of nodes added by GREEDY to S_0 until first node from OPT' is considered but not added to S because its addition would violate the budget P; let $a \in A^i$ be this node. Let S_j be the set of the first j nodes added to S_0 by GREEDY, where we set $S_{\ell+1} = S_\ell - s^i + a$. Note that $c(S_{\ell+1}) = c(S_\ell - s^i + a) > P'$, since a was not added. Let $\Delta_i \mathsf{val}(S) = \mathsf{val}(S_i) - \mathsf{val}(S_{i-1})$ and $\Delta_i c(S) = c(S_i) - c(S_{i-1})$, $i = 1, \ldots, \ell + 1$. The following two statements are similar to the ones used in [18].

Lemma 3.1 *For each* $j = 1, \ldots, \ell + 1$,

$$\frac{\Delta_j \mathsf{val}(S)}{\Delta_j c(S)} \geq \frac{\mathsf{val}(\mathsf{OPT}') - \mathsf{val}(S_{j-1})}{P'}.$$

Proof: At least $\mathsf{val}(\mathsf{OPT}') - \mathsf{val}(S_{j-1})$ worth of elements not covered by nodes of S_{j-1} are covered by nodes of OPT'. For each node in $\mathsf{OPT}' - S_{j-1}$ the ratio of weight to cost is at most $\frac{\Delta_j \mathsf{val}(S)}{\Delta_j c(S)}$, since the node that the algorithm picks in each iteration maximizes this ratio. Since the total cost of the nodes in $\mathsf{OPT}' - S_{j-1}$ is bounded by the residual budget P', we get

$$\mathsf{val}(\mathsf{OPT}') - \mathsf{val}(S_{j-1}) \leq P' \cdot \frac{\Delta_j \mathsf{val}(S)}{\Delta_j c(S)}.$$

□

Lemma 3.2 *For every* $j = 1, \ldots, \ell + 1$

$$\mathsf{val}(S_j) \geq \left[1 - \prod_{i=1}^{j} \left(1 - \frac{\Delta_i c(S)}{P'}\right)\right] \cdot \mathsf{val}(\mathsf{OPT}').$$

Proof: The proof is by induction on j. For $j = 1$ the statement holds by Lemma 3.1. Suppose the statement holds for $1, \ldots, j - 1$ and we prove it for j. Note that $\mathsf{val}(S_j) = \mathsf{val}(S_{j-1}) + \Delta_j \mathsf{val}(S)$. Thus by Lemma 3.1 we get:

$$\mathsf{val}(S_j) = \mathsf{val}(S_{j-1}) + \Delta_j \mathsf{val}(S) \;\; \geq \;\; \mathsf{val}(S_{j-1}) + \frac{\Delta_j c(S)}{P'} \cdot (\mathsf{val}(\mathsf{OPT}') - \mathsf{val}(S_{j-1}))$$

$$\geq \left(1 - \frac{\Delta_j c(S)}{P'}\right) \cdot \mathsf{val}(S_{j-1}) + \frac{\Delta_j c(S)}{P'} \cdot \mathsf{val}(\mathsf{OPT}').$$

Consequently, by the induction hypothesis

$$\begin{aligned}
\mathsf{val}(S_j) &\geq \left(1 - \frac{\Delta_j c(S)}{P'}\right) \cdot \mathsf{val}(S_{j-1}) + \frac{\Delta_j c(S)}{P'} \cdot \mathsf{val}(\mathsf{OPT}') \\
&\geq \left(1 - \frac{\Delta_j c(S)}{P'}\right) \cdot \left[1 - \prod_{i=1}^{j-1} \left(1 - \frac{\Delta_i c(S)}{P'}\right)\right] \cdot \mathsf{val}(\mathsf{OPT}') + \frac{\Delta_j c(S)}{P'} \cdot \mathsf{val}(\mathsf{OPT}') \\
&= \left[1 - \prod_{i=1}^{j} \left(1 - \frac{\Delta_i c(S)}{P'}\right)\right] \cdot \mathsf{val}(\mathsf{OPT}').
\end{aligned}$$

□

Applying Lemma 3.2 with $j = \ell + 1$ we get:

$$\mathsf{val}(S_{\ell+1}) \geq \left[1 - \prod_{i=1}^{\ell+1} \left(1 - \frac{\Delta_i c(S)}{P'}\right)\right] \cdot \mathsf{val}(\mathsf{OPT}') \geq \left[1 - \prod_{i=1}^{\ell+1} \left(1 - \frac{\Delta_i c(S)}{c(S_{\ell+1})}\right)\right] \cdot \mathsf{val}(\mathsf{OPT}')$$

Since for $a_1, ...a_n \in R^+$ such that $\sum_{i=1}^n a_i = A$, the function $1 - \prod_{i=1}^n (1 - \frac{a_i}{A})$ achieves its minimum when $a_1 = a_2 = ... = a_n = \frac{A}{n}$ we get:

$$\mathsf{val}(S_{\ell+1}) \geq \left[1 - \left(1 - \frac{1}{\ell+1}\right)^{\ell+1}\right] \cdot \mathsf{val}(\mathsf{OPT}') \geq \left(1 - \frac{1}{e}\right) \cdot \mathsf{val}(\mathsf{OPT}').$$

Let S be the set returned by GREEDY(S_0), and let $S' = S - S_0$. Then

$$\mathsf{val}(S') + \Delta_{\ell+1}\mathsf{val}(S) \geq \mathsf{val}(S_{\ell+1}) \geq \left(1 - \frac{1}{e}\right) \cdot \mathsf{val}(\mathsf{OPT}').$$

In addition $\Delta_{\ell+1}\mathsf{val}(S) \leq \frac{1}{k}\mathsf{val}(S_0)$ by the way the nodes in OPT were ordered. Thus:

$$\begin{aligned}
\mathsf{val}(S) = \mathsf{val}(S_0) + \mathsf{val}(S') &\geq \mathsf{val}(S_0) + \left(1 - \frac{1}{e}\right) \cdot \mathsf{val}(\mathsf{OPT}') - \Delta_{\ell+1}\mathsf{val}(S) \\
&\geq \mathsf{val}(S_0) + \left(1 - \frac{1}{e}\right) \cdot \mathsf{val}(\mathsf{OPT}') - \frac{1}{k}\mathsf{val}(S_0) \\
&\geq \left(1 - \frac{1}{k}\right) \cdot \mathsf{val}(S_0) + \left(1 - \frac{1}{e}\right) \cdot \mathsf{val}(\mathsf{OPT}') \\
&\geq \left(1 - \frac{1}{e}\right) \cdot (\mathsf{val}(S_0) + \mathsf{val}(\mathsf{OPT}')) \\
&= \left(1 - \frac{1}{e}\right) \cdot (\mathsf{val}(S_0) + \mathsf{val}(\mathsf{OPT} - S_0)) \\
&= \left(1 - \frac{1}{e}\right) \cdot \mathsf{val}(\mathsf{OPT})
\end{aligned}$$

The last inequality follows from the fact that $k > e$.

The proof of Theorem 2.1 is complete.

4 An $H(\Delta)$-approximation for directed Min-Power Edge-Multicover

In this section we prove Theorem 2.2.

Its not hard to see that the natural LP for MPEMC has $\Omega(\sqrt{\Delta})$ integrality gap or worse (see [13] for similar integrality gap for $\mathsf{MP}k\text{-}\mathsf{CS}$). The significance of our proof is showing a simple reduction to an equivalent problem that admits a logarithmic integrality gap. The analysis is similar to the standard dual fitting algorithm for the min-cost Set-Multicover problem (see [15]). However, as far as we can see, Theorem 2.2 cannot be deduced from the logarithmic approximation for the min-cost Set-Multicover problem, and thus our result seems to be a strict generalization of the later. Furthermore, the analysis carries extra terms that are required to guarantee Property (**).

By a similar transformation as for the directed PBMEMC in the previous section, the directed MPEMC can be formulated as follows:

Instance: A bipartite graph $\mathcal{G} = (A+B, \mathcal{E})$, costs $\{c(a) : a \in A\}$, requirements $\{r(b) : b \in B\}$, and a partition $\mathcal{A}$ of A satisfying (*).
Objective: Find min-cost $S \subseteq A$ so that $d_{\mathcal{E}}(S, b) \geq r(b)$ for every $b \in B$ and (**) holds.

Remark: Without property (*), even the feasibility version of the problem is NP-complete, e.g., see [9].

Similarly to the previous section, for S satisfying (**) let $B_S = \{b \in B : r_S(b) > 0\}$ be the set of *deficient* nodes w.r.t. S, let $s^i = A^i \cap S$ (possibly $s^i = \emptyset$). For $a \in A^i$ with $c(a) > c(s^i)$ the *density of a w.r.t. S* is:

$$\sigma_c(S, a) = \frac{|(\Gamma(a) - \Gamma(s^i)) \cap B_S|}{c(a) - c(s^i)}$$

and *cost-effectiveness of a w.r.t. S*

$$c_S(a) = \frac{1}{\sigma_c(S, a)} = \frac{c(a) - c(s^i)}{|(\Gamma(a) - \Gamma(s^i)) \cap B_S|}.$$

The following algorithm starts with $S = \emptyset$ and generalizes the greedy algorithm for the Min-Cost Set Multicover problem. In the algorithm, we keep variables π^i_b for each $b \in B$, initially set to zero, that indicate the amount paid by part A^i for covering b.

While $B_S \neq \emptyset$ do:

1. Find $a \in A$ with $\sigma_c(S,a)$ maximal, and let A^i be the part with $a \in A^i$;
2. $\pi^i_b \leftarrow c_S(a)$ for every $b \in (\Gamma(a) - \Gamma(s^i)) \cap B_S$;
3. $S \leftarrow S - s^i + a$.

End While

For the analysis, we use the corresponding LP-relaxation and its dual:

$$\begin{array}{lll} \min & \sum_{a\in A} c_a x_a & \\ \text{s.t.} & \sum_{a\in\Gamma(b)} x_a \geq r_b & \forall b \in B \\ & \sum_{a\in A_i} x_a \leq 1 & \forall A_i \in A \\ & x_a \geq 0 & \forall a \in A \end{array} \tag{1}$$

$$\begin{array}{lll} \max & \sum_{b\in B} r_b y_b - \sum_{A^i\in\mathcal{A}} z_{A^i} & \\ \text{s.t.} & \sum_{b\in\Gamma(a)} y_b - z_{A_i} \leq c_a & \forall a \in A, a \in A_i \\ & y_b, z_{A_i} \geq 0 & \forall b \in B, A^i \in \mathcal{A} \end{array} \tag{2}$$

Claim 4.1 *The sequence of cost effectiveness of nodes picked by the algorithm is nondecreasing, that is: if a'' was added to $S = S''$ after a' was added to $S = S'$ then $c_{S''}(a'') \geq c_{S'}(a')$.*

Henceforth, let $B^i = \{b \in B : \pi^i_b \neq 0\}$ be the set of nodes covered by A^i. At the end of the algorithm set:

$$\alpha_b = \max_i \pi^i_b \quad \forall b \in B \qquad \beta_{A^i} = \sum_{b\in B^i} (\alpha_b - \pi^i_b) \quad \forall A^i \in \mathcal{A}\ .$$

Note that α_b is the maximum amount (which is the last amount) paid for covering b. It is easy to see that $\alpha_b, \beta_{A^i} \geq 0$.

Lemma 4.2

$$\sum_{b\in B} r_b\alpha_b - \sum_{A^i\in A} \beta_{A^i} = c(S).$$

Proof: The total cost of the solution is

$$\sum_{b\in B} r_b\alpha_b - \sum_{A^i\in A} \beta_{A^i} = \sum_{b\in B} r_b\alpha_b - \sum_{A^i\in A}\sum_{b\in B^i} (\alpha_b - \pi^i_b) = \sum_{A^i\in A}\sum_{b\in B^i} \pi^i_b = c(S).$$

□

Let $\Delta \leq |B|$ be the maximal degree of a node in A. The dual variables are set as follows:

$$y_b = \frac{\alpha_b}{H(\Delta)} \quad \forall b \in B \qquad z_{A^i} = \frac{\beta_{A^i}}{H(\Delta)} \quad \forall A^i \in \mathcal{A}.$$

Theorem 2.2 now follows from the following lemma:

Lemma 4.3 *The pair (y, z) is a feasible solution for the dual problem (2).*

Proof: We need to prove that for every $a \in A^i \in \mathcal{A}$:

$$\sum_{b\in\Gamma(a)} \alpha_b - \beta_{A^i} = \sum_{b\in\Gamma(a)} \alpha_b - \sum_{b\in B^i} (\alpha_b - \pi_b^i) \leq c(a)H(\Delta) \ .$$

Let $k = |\Gamma(a)|$ and let $b_1, b_2, ..., b_k$ be an ordering of $\Gamma(a)$ in which these nodes left B_S.

Suppose that $A^i \cap S = \emptyset$. Then when the algorithm is about to r-cover b_j, $|\Gamma(a) \cap B_S| \geq k - j + 1$, so $\alpha_{b_j} \leq \frac{c(a)}{k-j+1}$. Since $\beta_{A^i} = 0$, in this case:

$$\sum_{b\in\Gamma(a)} \alpha_b - \beta_{A^i} = \sum_{b\in\Gamma(a)} \alpha_b \leq c(a)(\frac{1}{k} + \frac{1}{k-1} + \cdots + 1) \leq c(a)H(\Delta) \ .$$

Henceforth assume that $A^i \cap S \neq \emptyset$. Note that:

$$\sum_{b\in\Gamma(a)} \alpha_b - \sum_{b\in B^i} (\alpha_b - \pi_b^i) = \sum_{b\in\Gamma(a)-B^i} \alpha_b + \sum_{b\in\Gamma(a)\cap B^i} \pi_b^i - \sum_{b\in B^i-\Gamma(a)} (\alpha_b - \pi_b^i) \ .$$

By the definition of α_b, we have:

$$\sum_{b\in B^i-\Gamma(a)} (\alpha_b - \pi_b^i) \geq 0 \ .$$

Thus, it would be sufficient to prove that:

$$\sum_{b\in\Gamma(a)-B^i} \alpha_b + \sum_{b\in\Gamma(a)\cap B^i} \pi_b^i \leq c(a)H(\Delta) \ . \tag{3}$$

Number the nodes in $\Gamma(a) - B^i$ in the order they left B_S, say $b_1, b_2, ..., b_t$. Consider some $b_j \in \Gamma(a) - B^i$. Then b_j had requirement zero when s^i was picked, since $b_j \notin B^i$. Let a_j be the last node of A^i that was chosen when b_j had positive requirement, if such a_j exists, or $a_j = \emptyset$ otherwise. Note that $b_j \in \Gamma(a) - \Gamma(a_j)$ since a_j *does not* cover b_j. If a_j does not exist, then $c(a_j) = 0$. Let $C_j = \Gamma(a_j) \cap B^i$. Since after choice of a_j, $\Gamma(a) - \Gamma(a_j)$ contains at least $k - |C_j| - j + 1$ deficient nodes we have:

$$\alpha_{b_j} \leq \frac{c(a) - c(a_j)}{k - |C_j| - j + 1} \leq \frac{c(a)}{k - |C_j| - j + 1} \ .$$

For each b_j the denominator values are different and decrease with increase of j, so

$$\sum_{b\in\Gamma(a)-B^i} \alpha_b \leq \sum_{j=1}^{t-1} \frac{c(a)}{k - |C_j| - j + 1} + \frac{c(a) - c(a_t)}{k - |C_t| - t + 1} \ . \tag{4}$$

If $k - |C_t| - t + 1 > 1$ then

$$\sum_{b \in \Gamma(a) - B^i} \alpha_b \le \sum_{j=1}^{t} \frac{c(a)}{k - |C_j| - j + 1} \le H(\Delta) \cdot c(a) - c(a),$$

and

$$\sum_{b \in \Gamma(a) \cap B^i} \pi_b^i \le c(a) \ .$$

Thus

$$\sum_{b \in \Gamma(a) - B^i} \alpha_b + \sum_{b \in \Gamma(a) \cap B^i} \pi_b^i \le (H(\Delta) \cdot c(a) - c(a)) + c(a) = c(a) H(\Delta) \ .$$

Assume therefore that $k - |C_t| - t + 1 = 1$. Then all nodes in $\Gamma(a) \cap B^i$ are covered by a_t, thus

$$\sum_{b \in \Gamma(a) \cap B^i} \pi_b^i = c(a_t)$$

and together with (4) this implies

$$\sum_{b \in \Gamma(a) - B^i} \alpha_b + \sum_{b \in \Gamma(a) \cap B^i} \pi_b^i \le \sum_{j=1}^{t-1} \frac{c(a)}{k - |C_j| - j + 1} + (c(a) - c(a_t)) + c(a_t) \le c(a) H(\Delta) \ .$$

□

From Lemmas 4.2 and 4.3 we conclude:

$$c(S) = \sum_{b \in B} r_b \alpha_b - \sum_{A^i \in A} \beta_{A^i} = H(\Delta) \left(\sum_{b \in B} r_b y_b - \sum_{A^i \in A} z_{A^i} \right) \le H(\Delta) \cdot \mathsf{opt}.$$

The proof of Theorem 2.2 is complete.

5 An $O(\log n)$-approximation for undirected Min-Power Edge-Multicover

In this section we prove Theorem 2.3.

We show an $O(\log n)$-approximation algorithm for (undirected) *bipartite* MPEMC where $\mathcal{G} = (A + B, \mathcal{E})$ is bipartite and $r(a) = 0$ for every $a \in A$.

Lemma 5.1 *If there exists a ρ-approximation algorithm for bipartite* MPEMC *then there exists a 2ρ-approximation algorithm for general* MPEMC.

Proof: Given an instance $\mathcal{G} = (V, \mathcal{E}), c, r$ of MPEMC, construct an instance $\mathcal{G}' = (V' = A + B, \mathcal{E}'), c', r'$ of bipartite MPEMC as follows. Let $A = \{a_v : v \in V\}$ and $B = \{b_v : v \in V\}$ (so each of A, B is a copy of V) and for every $uv \in \mathcal{E}$ add two edges: $a_u a_v$ and $a_v a_u$ each with cost $c(uv)$. Also, set $r'(b_v) = r(v)$ for every $b_v \in B$ and $r'(a_v) = 0$ for every $a_v \in A$. Given $F' \subseteq \mathcal{E}'$ let $F = \{uv \in \mathcal{E} : a_u b_v \in F' \text{ or } a_v b_u \in F'\}$ be the edge set in $\mathcal{E}$ that corresponds to F'. Now compute an r'-edge cover E' in $\mathcal{G}'$ using the ρ-approximation algorithm and output the edge set $E \subseteq \mathcal{E}$ that corresponds to E', namely $E = \{uv \in \mathcal{E} : a_u b_v \in E' \text{ or } a_v b_u \in E'\}$.

It is easy to see that if F' is an r'-edge cover then F is an r-edge cover. Furthermore, if for every edge in F correspond two edges in F' ($|F'| = 2|F|$), then F is an r-edge cover if, and only if, F' is an r'-edge cover. The later implies that $\mathsf{opt}' \leq 2\mathsf{opt}$, where opt and opt' is the optimal solution value to $\mathcal{G}, c, r$ and $\mathcal{G}', c', r'$, respectively. Consequently, E is an r-edge cover, and $p_E(V) \leq p_{E'}(V') \leq \rho\mathsf{opt}' \leq 2\rho\mathsf{opt}$. $\square$

We henceforth prove that bipartite MPEMC admits an $O(\log n)$-approximation algorithm. The *residual requirement* of $v \in V$ w.r.t. an edge set E is

$$r_E(v) = \max\{r(v) - d_E(v), 0\} .$$

Lemma 5.2 *For bipartite* MPEMC *there exists a polynomial time algorithm that given an integer τ and $\alpha > 1$ either establishes that $\tau < \mathsf{opt}$ or returns an edge set $I \subseteq \mathcal{E}$ such that*

$$p_I(V) \leq (\alpha + 1)\tau \tag{5}$$

$$r_I(B) \leq (1 - \beta) r(B) , \tag{6}$$

where $\beta = (1 - 1/e)(1 - 1/\alpha)$.

Note that if $\tau < \mathsf{opt}$ the algorithm may return a edge set I that satisfies (5) and (6); if the algorithm declares "$\tau < \mathsf{opt}$" then this is correct. An $O(\log n)$-approximation algorithm for the bipartite MPEMC easily follows from Lemma 5.2:

While $r(B) > 0$ do

Find the least integer τ so that the algorithm in Lemma 5.2 returns an edge set I so that (5) and (6) holds.

$E \leftarrow E + I$, $\mathcal{E} \leftarrow \mathcal{E} - I$, $r \leftarrow r_I$.

End While

We note that the least integer τ as in the main loop can be found in polynomial time using binary search. For any constant $\alpha > 1$, say $\alpha = 2$, the number of iterations is $O(\log r(B))$, and at every iteration an edge set of power at most $(1+\alpha)\mathsf{opt}$ is added. Thus the algorithm can be implemented to run in polynomial time, and has approximation ratio $O(\log r(B)) = O(\log(n^2)) = O(\log n)$.

In the rest of this section we prove Lemma 5.2. Let τ be an integer and let $R = r(B) = \sum_{b \in B} r(b)$. An edge $ab \in \mathcal{E}$ with $b \in B$ is *dangerous* if $c(ab) \geq \alpha\tau \cdot r(b)/R$. Let $\mathcal{I}$ be the set of non-dangerous edges in $\mathcal{E}$.

Lemma 5.3 *Let F be a set of dangerous edges with $p_F(B) \leq \tau$. Then $r_F(B) \geq R(1 - 1/\alpha)$. Thus if $\tau \geq \mathsf{opt}$ then $r_{\mathcal{I}}(B) \leq R/\alpha$.*

Proof: Let $D = \{b \in B : d_F(b) > 0\}$. We show that $r(D) \leq R/\alpha$, implying $r_F(V) \geq R - r(D) \geq R(1 - 1/\alpha)$. Since all the edges in F are dangerous, $p_F(b) \geq \alpha\tau \cdot r(b)/R$ for every $b \in D$. Thus

$$\tau \geq \sum_{b \in D} p_F(b) \geq \sum_{b \in D} (\alpha\tau \cdot r(b)/R) = \frac{\alpha\tau}{R} \sum_{b \in D} r(b) = \frac{\alpha\tau}{R} r(D) .$$

For the second statement, note that if $\tau \geq \mathsf{opt}$ then there exists $E \subseteq \mathcal{E}$ with $p_E(V) \leq \tau$ so that $r_E(B) = 0$. Thus for the set I of non-dangerous edges in E we have $r_I(B) \leq R/\alpha$. Since $I \subseteq \mathcal{I}$, the statement follows. $\square$

Lemma 5.4 *$p_{\mathcal{I}}(B) \leq \alpha\tau$.*

Proof: Note that $p_{\mathcal{I}}(b) \leq \alpha\tau \cdot r(b)/R$ for every $b \in B$. Thus:

$$p_{\mathcal{I}}(B) = \sum_{b \in B} p_{\mathcal{I}}(b) \leq \sum_{b \in B} (\alpha\tau \cdot r(b)/R) = \frac{\alpha\tau}{R} \sum_{b \in B} r(b) = \alpha\tau .$$

$\square$

The algorithm is as follows:

1. With budget τ, compute an edge set $I \subseteq \mathcal{I}$ using the $(1 - 1/e)$-approximation algorithm for directed PBMEMC (see Theorem 2.1).
2. *If* $r_I(B) \leq (1 - \beta)R$ (recall that $\beta = (1 - 1/e)(1 - 1/\alpha)$) then output I;
   Else declare "$\tau < \mathsf{opt}$".

We show that if $\tau \geq \mathsf{opt}$ then the algorithm outputs an edge set I that satisfies (5) and (6). By Lemma 5.3, if the algorithm returns an edge set I then (5) holds for I, and if the algorithm declares "$\tau < \mathsf{opt}$" then this is correct. All the edges in I are not dangerous, thus $p_I(B) \leq \alpha\tau$ by Lemma 5.4. As we used budget τ, $p_I(A) \leq \tau$. Thus $p_I(V) = p_I(A) + p_I(B) \leq (1+\alpha)\tau$.

The proof of Theorem 2.3 is complete.

6 A $4r_{\max}$-approximation for undirected Min-Power Steiner Network

In this section we prove Theorem 2.7.

We need some definitions and a description of certain results from [26, 12].

Min-cost/power Steiner Network problem can be formulated as a *set-function edge-cover problem*. Let $p : 2^V \rightarrow Z_+$ be a *set-function* defined on a groundset V. An edge set E on V is a *p-cover*, if $d_E(X) \geq p(X)$ for every $X \subseteq V$. For Steiner Network problems, an appropriate choice of p is as follows. By Menger's Theorem, E is a feasible solution to min-cost/power Steiner network problem if, and only if, $d_E(X) \geq R(X)$ for all $\emptyset \subset X \subset V$, where $R(X) = \max\{r(u,v) : u \in X, v \in V - X\}$ (and $R(\emptyset) = R(V) = 0$). That is

$$d_E(X) \geq p(X) \equiv \max\{0, R(X)\} \quad \forall\, \emptyset \subseteq X \subseteq V. \tag{7}$$

The function p defined above is *skew-supermodular*, that is $p(\emptyset) = 0$ and for every $X, Y \subseteq V$ with $p(X) > 0$, $p(Y) > 0$ at least one of the following holds:

$$p(X) + p(Y) \leq p(X \cap Y) + p(X \cup Y) \tag{8}$$

$$p(X) + p(Y) \leq p(X - Y) + p(Y - X) \tag{9}$$

Note that p is also *symmetric*, that is, $p(X) = p(V - X)$ for all $X \subseteq V$.

Several connectivity problems can be formulated as (min-cost/power) edge cover problems of a skew-supermodular function, see [21]. A seminal paper of Jain [16] gives a 2-approximation algorithm for finding a min-cost edge-cover of an arbitrary skew-supermodular set function p, provided certain queries related to p can be answered in polynomial time (note that p is usually not given explicitly). For p defined in (7) these queries can be realized via max-flows, which implies a 2-approximation algorithm for the min-cost Steiner network problem. Earlier, Williamson et. al [26] gave an algorithm with approximation ratio $2p_{\max}$, which was improved later to $2H(p_{\max})$ by Goemans et. al [11].

Given a set function q, let $\hat{q}(X) = 1$ if $q(X) = q_{\max}$ and $hq(X) = 0$ otherwise, where $q_{\max} = \max_{X \subseteq V} q(X)$. It is easy to see that any inclusion minimal edge-cover of a $\{0,1\}$-valued set function is a forest. For an edge set E, let p_E be defined as follows: $p_E(X) = \max\{p(X) - d_E(X), 0\}$. It is well known that if p is skew supermodular, so is p_E (for any edge set E), see [16]. Consider the following algorithm that applies on an arbitrary set-function p, and begins with $E = \emptyset$.

While there is $X \subseteq V$ with $p_E(X) > 0$ do:

1. Find a $\hat{p}_E$-cover $F \subseteq \mathcal{E} - E$;
2. $E \leftarrow E + F$.

End While

The approximation ratio of the algorithm depends on step 1. A set function is called *uncrossable* if it is $\{0,1\}$-valued skew supermodular. It is easy to see that if q is skew supermodular, so is $\hat{q}$, that is $\hat{q}$ is uncrossable. Williamson et. al [26] gave an algorithm that finds an edge cover of an uncrossable function $\hat{q}$ of cost at most twice the optimum of the following LP-relaxation:

$$\min\{\sum_{e\in\mathcal{E}-E} c(e)x_e : \sum_{e\in\delta(X)} x_e \geq \hat{q}(X)\ \forall X \subseteq V, x_e \geq 0\} \ . \tag{10}$$

Williamson et. al [26] proved:

Theorem 6.1 ([26]) *For p defined by (7) the above algorithm can be implemented in polynomial time, so that at any iteration for $q = p_E$ the forest F found has cost at most twice the optimal value of (10).*

Note that the number of iterations of the algorithm is at most $p_{\max}$. Thus Theorem 6.1 implies that for the min-cost Steiner network problem the algorithm has approximation ratio $2p_{\max} \leq 2r_{\max}$. Later, Goemans et. al [12] used linear programming scaling techniques to show that the approximation ratio is in fact $2H(r_{\max})$. This scaling method does not work for the min-power variant.

We can show that for the min-power variant, the algorithm of [26] has approximation ratio $4r_{\max}$. This follows from Theorem 6.1 and the second part of Proposition 2.8. Indeed, the algorithm of [26] constructs the solution from at most $r_{\max}$ forests, where each forest has cost at most $2\mathsf{opt}_\mathsf{c}$, where opt_c is the optimal solution value to the min-cost variant. By Proposition 2.8, each forest has power at most $2 \cdot 2\mathsf{opt}_\mathsf{p} = 4\mathsf{opt}_\mathsf{p}$, where opt_p is the otimal solution value to the min-power variant.

The proof of Theorem 2.7 is complete.

7 Conclusions

One of the main results in this work is an $O(\log n)$-approximation algorithm for the undirected Min-Power Edge-Multicover (MPEMC) problem, improving the previosely best known $O(\log^4 n)$-approximation. This implies an $O(\alpha + \log n)$-approximation for the Min-Power k-Connected Subgraph (MPk-CS) problem, where α is the best known approximation ratio for the Min-Cost k-Connected Subgraph (MCk-CS) problem. Consequently, for undirected graphs, MPk-CS and MCk-CS are equivalent with respect to approximation (up to constants), unless MCk-CS admits a better than $O(\log n)$ approximation ratio. We note that for $k = n - o(n)$ the best known approximation ratio for MCk-CS is $O(\sqrt{n} \log n)$, while the best known lower bound on approximation is APX-hardness. We believe that the established equivalence between MPk-CS and MCk-CS will enable to reduce the gap.

We also gave an $H(\Delta)$-approximation algorithm for the directed Min-Power Edge-Multi-Cover MPEMC problem; the approximation ratio is proved using the dual fitting method applied on an appropriately defined LP-relaxation. This result seems a strict extension of the classic results for the min-cost case. We believe that the "node duplicating technique" used can be useful to other min-power problems. For example, we used the same method to extend the $(1 - 1/e)$-approximation algorithm of Khuller, Moss, and Naor [18] for the Cost-Budgeted Maximum Coverage problem to its min-power "multi" variant PBMEMC.

Finally, we gave a $4r_{\max}$-approximation algorithm for (undirected) Min-Power Steiner Network MPSN problem. We note that even for $r_{\max} = 1$, the best known approximation ratio for MPSN is 4, which follows immediately from the 2-approximation for the min-cost case and Propositin 2.8. Our result is an extension to arbitrary requirements.

Some open problems that arise from this research are as follows.

- Improving the 4-approximation for MPSN with $r_{\max} = 1$.
- Approximation status of undirected MPEMC: the problem is APX-hard [14], and in this work we gave an $O(\log n)$-approximation algorithm. Can this gap be closed?
- Approximation status of undirected PBMEMC. The best known approximation ratio for the Densest k-subgraph problem is roughly $n^{-1/3}$ [10], and in spite of numerous attempts to improve it, this ratio holds for almost 10 years. We showed that for $r_{\max} = k$ undirected PBMEMC is at least as hard as the Densest k-Subgraph Problem. We leave an open question whether for small $r_{\max}$ a constant approximation ratio is possible for the undirected PBMEMC.

8 Appendix

Here we restate the proofs of Proposition 2.8 and Theorem 2.5 from [14].

8.1 Proof of Proposition 2.8

The inequality $p(G) \le 2c(G)$ follows from

$$p(G) = \sum_{v\in V} p(v) \le \sum_{v\in V}\sum_{e\in\delta(v)} c(e) = 2\sum_{e\in E} c(e) = 2c(G).$$

If T is a tree, root it at an arbitrary node r. Then $c(T) \le p(T)$ since for each $v \ne r$, $p(v)$ is at least the cost of the parent edge of v.

We now show that $c(G) \le \sqrt{|E|/2}p(G)$ It is sufficient to prove that

$$\sum_{xy\in E} \min\{p(x), p(y)\} \le \sqrt{|E|/2}\sum_{v\in V} p(v) \tag{11}$$

for any graph $G = (V, E)$ with nonnegative weights $p(v)$ on the nodes. Suppose to the contrary that the statement is false, and let $G = (V, E)$ with p be a counterexample to (11) so that $\max_{v\in V} p(v) - \min_{v\in V} p(v)$ is minimal. Let $\mu = \min_{v\in V} p(v)$, let $U = \{v \in V : p(v) = \mu\}$, and let E_U be the set of edges in E with at least one endpoint in U. If $|E_U| \le \sqrt{|E|/2}|U|$ then the statement is also false for $G' = (V', E') = (V - U, E - E_U)$ and p' being the restriction of p to V' since

$$\begin{aligned}
\sum_{xy\in E'} \min\{p'(x), p'(y)\} &\ge \sum_{xy\in E} \min\{p(x), p(y)\} - \sqrt{|E|/2}|U|\mu > \\
&> \sqrt{|E|/2}\sum_{v\in V} p(v) - \sqrt{|E|/2}|U|\mu = \sqrt{|E|/2}\sum_{v\in V'} p'(v) > \\
&> \sqrt{|E'|/2}\sum_{v\in V'} p'(v).
\end{aligned}$$

In particular, this implies a contradiction if $U = V$. Else, let $\mu' = \min\{p(v) : v \in V - U\}$ be the second minimum value of p. Then by setting $p(v) \leftarrow p(v) + \mu' - \mu$ for every $v \in U$ we obtain again a counterexample to (11). This contradicts our choice of G, p.

8.2 Proof of Theorem 2.5

To prove Theorem 2.5 we use the following fundamental statement due to Mader.

Theorem 8.1 ([22]) *In a k-connected graph G, any cycle in which every edge is critical contains a node whose degree in G is k.*

Here an edge e of a k-connected graph G is *critical* (w.r.t. k-connectivity) if $G - e$ is not k-connected.

The following corollary (e.g., see [22]) is used to get a relation between $(k-1)$-edge covers and k-connected spanning subgraphs.

Corollary 8.2 *If $\deg_J(v) \geq k - 1$ for every node v of a graph J, and if F is an inclusion minimal edge set such that $J \cup F$ is k-connected, 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 G, contradicting Mader's Theorem. $\square$

Proof of Theorem 2.5: By the assumption, we can find a subgraph J with $\deg_J(v) \geq k-1$ of power at most $p(J) \leq \beta\mathsf{opt}$. We reset the costs of edges in J to zero, and apply an α-approximation algorithm for the Min-Cost k-Connected Spanning Subgraph problem to compute an (inclusion) minimal edge set F so that $J + F$ is k-connected. By Corollary 8.2, F is a forest. Thus $p(F) \leq 2c(F) \leq 2\alpha\mathsf{opt}$, by Lemma 2.8. Combining, we get the desired statement.

The proof of the other direction is similar. We find a min-cost $(k - 1)$-edge cover J in polynomial time, and reset the costs of its edges to zero. Then we use the ρ-approximation algorithm for $\mathsf{MP}k\text{-}\mathsf{CS}$ with the new cost function. The edges with nonzero cost in this new graph form a forest F, by Corollary 8.2. Then clearly $c(J)$ is at most the minimum cost of a k-connected spanning subgraph, and $c(F)$ is at most 2ρ times the minimum cost of a k-connected spanning subgraph, by Lemma 2.8. This gives a $(2\rho + 1)$-approximation algorithm for the Min-Cost k-Connected Spanning Subgraph problem. $\square$

References

[1] E. Althaus, G. Calinescu, I. Mandoiu, S. Prasad, N. Tchervenski, and A. Zelikovsky. Power efficient range assignment in ad-hoc wireless networks. In *Proceedings of IEEE Wireless Communications and Networking Conference (WCNC)*, pages 1889–1894, 2003.

[2] D. M. Blough, M. Leoncini, G. Resta, and P. Santi. On the symmetric range assignment problem in wireless ad hoc networks. In *Proceedings of the IFIP World Computer*

Congress- TC1 Stream/2nd IFIP International Conference on Theoretical Computer Science, 2002.

[3] J. L. Bredin, E. D. Demaine, M. T. Hajiaghayi, and D. Rus. Deploying sensor networks with guaranteed capacity and fault tolerance. In *Proceedings of the 6th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC 2005)*, pages 309–319, 2005.

[4] G. Calinescu1, S. Kapoor, A. Olshevsky, and A. Zelikovsky. Network lifetime and power assignment in ad hoc wireless networks. In *European Symposium on Algorithms (ESA)*, pages 114–126, 2003.

[5] G. Calinescu1, I. I. Mandoiu, and A. Zelikovsky. Symmetric connectivity with minimum power consumption in radio networks. In *TCS '02: Proceedings of the IFIP 17th World Computer Congress - TC1 Stream / 2nd IFIP International Con- ference on Theoretical Computer Science*, pages 119–130, 2002.

[6] C. Chekuri and A. Kumar. Maximum coverage problem with group budget constraints and applications. In *APPROX-RANDOM*, pages 72–83, 2004.

[7] 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.

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

[9] M. Elkin and G. Kortsarz. An approximation algorithm for directed broadcast. *Algorithmica*, 2006. to appear. Preliminary version in the thirtieth International Colloquium on Automata, Languages and Programming (ICALP) 2003, pages 212-223.

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

[11] H. N. Gabow, M. X. Goemans, E. Tardos, and D. P. Williamson. Approximating the smallest k-edge connected spanning subgraph by LP-rounding. In *Symposium on Discrete Algorithms (SODA)*, pages 562–571, 2005.

[12] M. X. Goemans, A. V. Goldberg, S. Plotkin, D. B.Shmoys, E. Tardos, and D. P. Williamson. Improved approximation algorithms for network design problems. In *Symposium on Discrete Algorithms (SODA)*, pages 223–232, 1994.

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

[14] 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.

[15] D. Hochbaum. *Approximating covering and packing problems: set cover, vertex cover, independent set and related problems,* Chapter 6 in *Approximation Algorithms for NP-hard problems,* D. S. Hochbaum Ed., pages 94-143. PWS, 1995.

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

[17] 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.

[18] S. Khuller, A. Moss, and J. Naor. The budgeted maximum coverage problem. *Information Processing Letters*, 70(1):39–45, 1999.

[19] L. M. Kirousis, E. Kranakis, D. Krizanc, and A. Pelc. Power consumption in packet radio networks. *Theoret. Comput. Sci.*, 243:289–305, 2000.

[20] G. Kortsarz and Z. Nutov. Approximation algorithm for k-node connected subgraphs via critical graphs. In *Symposium on the Theory of Computing (STOC)*, pages 138–145, 2004.

[21] G. Kortsarz and Z. Nutov. Approximating minimum cost connectivity problems, in *Approximation Algorithms and Metaheuristics,* T. F. Gonzalez ed.,. 2005. To appear.

[22] W. Mader. Ecken vom grad n in minimalen n-fach zusammenhängenden graphen. *Archive der Mathematik*, 23:219–224, 1972.

[23] Z. Nutov. Approximating min-power connectivity problems. *Manuscript*, 2006.

[24] Z. Nutov. Approximating minimum power covers of intersecting families and directed connectivity problems. *Manuscript*, 2006.

[25] R. Raz and S. Safra. A sub-constant error-probability low-degree test and a sub-constant error-probability PCP characterization of NP. In *Symposium on the Theory of Computing (STOC)*, pages 475–484, 1997.

[26] 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.