Approximating minimum power edge-multi-covers

Nachshon Cohen and Zeev Nutov

The Open University of Israel,
nachshonc@gmail.com, nutov@openu.ac.il

Abstract. Given a graph with edge costs, the *power* of a node is the maximum cost of an edge incident to it, and the power of a graph is the sum of the powers of its nodes. Motivated by applications in wireless networks, we consider the following fundamental problem in wireless network design. Given a graph $G = (V, E)$ with edge costs and degree bounds $\{r(v) : v \in V\}$, the Minimum-Power Edge-Multi-Cover (MPEMC) problem is to find a minimum-power subgraph J of G such that the degree of every node v in J is at least $r(v)$. Let $k = \max_{v \in V} r(v)$. For $k = \Omega(\log n)$, the previous best approximation ratio for MPEMC was $O(\log n)$, even for uniform costs [3]. Our main result improves this ratio to $O(\log k)$ for general costs, and to $O(1)$ for uniform costs. This also implies ratios $O(\log k)$ for the Minimum-Power k-Outconnected Subgraph and $O\left(\log k \log \frac{n}{n-k}\right)$ for the Minimum-Power k-Connected Subgraph problems; the latter is the currently best known ratio for the min-cost version of the problem. In addition, for small values of k, we improve the previously best ratio $k + 1$ to $k + 1/2$.

1 Introduction

1.1 Motivation and problems considered

Wireless networks are studied extensively due to their wide applications. The power consumption of a station determines its transmission range, and thus also the stations it can send messages to; the power typically increases at least quadratically in the transmission range. Assigning power levels to the stations (nodes) determines the resulting communication network. Conversely, given a communication network, the power required at v only depends on the farthest node reached directly by v. This is in contrast with wired networks, in which every pair of stations that communicate directly incurs a cost. An important network property is fault-tolerance, which is often measured by minimum degree or node-connectivity of the network. Node-connectivity is much more central here than edge-connectivity, as it models stations failures. Such power minimization problems were vastly studied; see for example [1, 2, 5, 8, 9] and the references therein for a small sample of papers in this area. The first problem we consider is finding a low power network with specified lower degree bounds. The second problem is the Min-Power k-Connected Subgraph problem. We give approximation algorithms for these problems, improving the previously best known ratios.

Definition 1. *Let (V, J) be a graph with edge-costs $\{c(e) : e \in J\}$. For a node $v \in V$ let $\delta_J(v)$ denote the set of edges incident to v in J. The* power $p_J(v)$ *of v is the maximum cost of an edge in J incident to v, or 0 if v is an isolated node of J; i.e., $p_J(v) = \max_{e \in \delta_J(v)} c(e)$ if $\delta_J(v) \neq \emptyset$, and $p_J(v) = 0$ otherwise. For $V' \subseteq V$ the power of V' w.r.t. J is the sum $p_J(V') = \sum_{v \in V'} p_J(v)$ of the powers of the nodes in V'.*

Unless stated otherwise, all graphs are assumed to be undirected and simple. Let $n = |V|$. Given a graph $G = (V, E)$ with edge-costs $\{c(e) : e \in E\}$, we seek to find a low power subgraph (V, J) of G that satisfies some prescribed property. One of the most fundamental problems in Combinatorial Optimization is finding a minimum-cost subgraph that obeys specified degree constraints (sometimes called also "matching problems") c.f. [10]. Another fundamental property is fault-tolerance (connectivity). In fact, these problems are related, and we use our algorithm for the former as a tool for approximating the latter.

Definition 2. *Given degree bounds $r = \{r(v) : v \in V\}$, we say that an edge-set J on V is an* r-edge cover *if $d_J(v) \geq r(v)$ for every $v \in V$, where $d_J(v) = |\delta_J(v)|$ is the degree of v in the graph (V, J).*

Minimum-Power Edge-Multi-Cover (MPEMC):
Instance: A graph $G = (V, E)$ with edge-costs $\{c(e) : e \in E\}$, degree bounds $r = \{r(v) : v \in V\}$.
Objective: Find a minimum power r-edge cover $J \subseteq E$.

Given an instance of MPEMC, let $k = \max_{v \in V} r(v)$ denote the maximum requirement.

We now define our connectivity problems. A graph is *k-outconnected from* s if it contains k internally-disjoint sv-paths for all $v \in V \setminus \{s\}$. A graph is *k-connected* if it is k-outconnected from every node, namely, if it contains k internally-disjoint uv-paths for all $u, v \in V$.

Minimum-Power k-Outonnected Subgraph (MPkOS):
Instance: A graph $G = (V, E)$ with edge-costs $\{c(e) : e \in E\}$, a root $s \in V$, and an integer k.
Objective: Find a minimum-power k-outconnected from s spanning subgraph J of G.

Minimum-Power k-Connected Subgraph (MPkCS):
Instance: A graph $G = (V, E)$ with edge-costs $\{c(e) : e \in E\}$ and an integer k.
Objective: Find a minimum-power k-connected spanning subgraph J of G.

1.2 Our Results

For large values of $k = \Omega(\log n)$, the previous best approximation ratio for MPEMC was $O(\log n)$, even for uniform costs [3]. Our main result improves this ratio to $O(\log k)$ for general costs, and to $O(1)$ for uniform costs.

Theorem 1. MPEMC *admits an $O(\log k)$-approximation algorithm. For uniform costs,* MPEMC *admits a randomized approximation algorithm with expected approximation ratio $\rho < 2.16851$, where ρ is the real root of the qubic equation $e(\rho-1)^3 = 2\rho$.*

For small values of k, the problem admits also the ratios $k+1$ for arbitrary k [2], while for $k = 1$ the best known ratio is $k + 1/2 = 3/2$ [4]. Our second result extends the latter ratio to arbitrary k.

Theorem 2. MPEMC *admits a $(k + 1/2)$-approximation algorithm.*

For small values of k, say $k \leq 6$, the ratio $(k + 1/2)$ is better than $O(\log k)$ because of the constant hidden in the $O(\cdot)$ term. And overall, our paper gives the currently best known ratios for all values $k \geq 2$.

In [5] it is proved that an α-approximation for MPEMC implies an $(\alpha + 4)$-approximation for MPkOS. The previous best ratio for MPkOS was $O(\log n) + 4 = O(\log n)$ [5] for large values of $k = \Omega(\log n)$, and $k + 1$ for small values of k [9]. From Theorem 1 we obtain the following.

Theorem 3. MPkOS *admits an $O(\log k)$-approximation algorithm.*

In [2] it is proved that an α-approximation for MPEMC and a β-approximation for Min-Cost k-Connected Subgraph implies a $(\alpha+2\beta)$-approximation for MPkCS. Thus the previous best ratio for MPkCS was $2\beta + O(\log n)$ [3], where β is the best ratio for MCkCS (for small values of k better ratios for MPkCS are given in [9]). The currently best known value of β is $O\left(\log k \log \frac{n}{n-k}\right)$ [7], which is $O(\log k)$, unless $k = n - o(n)$. From Theorem 1 we obtain the following.

Theorem 4. MPkCS *admits an $O(\beta+\log k)$-approximation algorithm, where β is the best ratio for* MCkCS*. In particular,* MPkCS *admits an $O\left(\log k \log \frac{n}{n-k}\right)$-approximation algorithm.*

1.3 Overview of the techniques

Let the *trivial solution* for MPEMC be obtained by picking for every node $v \in V$ the cheapest $r(v)$ edges incident to v. It is known and easy to see that this produces an edge set of power at most $(k+1) \cdot \mathsf{opt}$, see [2].

Our $O(\log k)$-approximation algorithm uses the following idea. Extending and generalizing an idea from [3], we show how to find an edge set $I \subseteq E$ of power $O(\mathsf{opt})$ such that for the residual instance, the trivial solution value is reduced by a constant fraction. We repeatedly find and add such an edge set I to the constructed solution, while updating the degree bounds accordingly to $r(v) \leftarrow \max\{r(v) - d_I(v), 0\}$. After $O(\log k)$ steps, the trivial solution value is reduced to opt, and the total power of the edges we picked is $O(\log k) \cdot \mathsf{opt}$. At this point we add to the constructed solution the trivial solution of the residual problem, which at this point has value opt, obtaining an $O(\log k)$-approximate solution.

Our algorithm for uniform costs has two phases. In the first phase we compute an optimal solution x to a certain LP-relaxation for the problem and round it to 1 with probability $\min\{\rho \cdot x, 1\}$. In the second phase we add to the obtained partial solution the trivial solution to the residual problem.

Our $(k+1/2)$-approximation algorithm uses a two-stage reduction. The first reduction reduces MPEMC to a constrained version of MPEMC with $k = 1$, where we also have lower bounds ℓ_v on the power of each node $v \in V$; these lower bounds are determined by the trivial solution to the problem. We will show that a ρ-approximation algorithm to this constrained version implies a $(k-1+\rho)$-approximation algorithm for MPEMC. The second reduction reduces the constrained version to the Minimum-Cost Edge Cover problem with a loss of $3/2$ in the approximation ratio. As Minimum-Cost Edge Cover admits a polynomial time algorithm, we get a ratio $\rho = 3/2$ for the constrained problem, which in turn gives the ratio $k-1+\rho = k+1/2$ for MPEMC.

2 Proof of Theorem 1

2.1 Reduction to bipartite graphs

Let Bipartite MPEMC be the restriction of MPEMC to instances for which the input graph $G = (V, E)$ is a bipartite graph with sides A, B, and with $r(a) = 0$ for every $a \in A$ (so, only the nodes in B may have positive degree bound).

As in [3], we can reduce MPEMC to Bipartite MPEMC, by taking two copies $A = \{a_v : v \in V\}$ and $B = \{b_v : v \in V\}$ of V, for every edge $e = uv \in E$ adding the two edges $a_u b_v$ and $a_v b_u$ of cost $c(e)$ each, and for every $v \in V$ setting $r(b_v) = r(v)$ and $r(a_v) = 0$. It is proved in [3] that this reduction invokes a factor of 2 in the approximation ratio, namely, that a ρ-approximation for bipartite MPEMC implies a 2ρ-approximation for general MPEMC.

In the case of uniform costs, we can save a factor of 2 using a different reduction.

Proposition 1. *Ratio ρ for* Bipartite MPEMC *with unit costs implies ratio ρ for* MPEMC *with uniform costs.*

Proof. Clearly, the case of uniform costs is equivalent to the case of unit costs. Now we show that for unit costs, MPEMC can be reduced to Bipartite MPEMC. Let $G = (V, E), r$ be an instance of MPEMC with unit costs. If there is an edge $e = uv \in E$ with $r(u), r(v) \geq 1$ or with $r(u) = r(v) = 0$, then we can obtain an equivalent instance by removing e from G, and in the case $r(u), r(v) \geq 1$ also decreasing each of $r(u), r(v)$ by 1. Hence we may assume that every $e \in E$ has one end in $A = \{a \in V : r(a) = 0\}$ and the other end in $B = \{b \in V : r(b) \geq 1\}$. The statement follows. ☐

2.2 An $O(\log k)$-approximation algorithm for general costs

Let opt denote the optimal solution value of a problem instance at hand. For $v \in V$, let w_v be the cost of the $r(v)$-th least cost edge incident to v in E if

$r(v) \geq 1$, and $w_v = 0$ otherwise. Given a partial solution J to Bipartite MPEMC let $r_J(v) = \max\{r(v) - d_J(v), 0\}$ be the *residual bound* of v w.r.t. J. Let

$$R_J = \sum_{b \in B} w_b r_J(b) \ .$$

The main step in our algorithm is given in the following lemma, which will be proved later.

Lemma 1. *There exists a polynomial time algorithm that given an edge set $J \subseteq E$, an integer τ, and a parameter $\gamma > 1$, either correctly establishes that $\tau < \mathsf{opt}$, or returns an edge set $I \subseteq E \setminus J$ such that $p_I(V) \leq (1+\gamma)\tau$ and $R_{J \cup I} \leq \theta R_J$, where $\theta = 1 - \left(1 - \frac{1}{\gamma}\right)\left(1 - \frac{1}{e}\right)$.*

Lemma 2. *Let $J \subseteq E$ and let $F \subseteq E \setminus J$ be an edge set obtained by picking $r_J(b)$ least cost edges in $\delta_{E \setminus J}(b)$ for every $b \in B$. Then $J \cup F$ is an r-edge-cover and: $p_F(B) \leq \mathsf{opt}$, $p_F(A) \leq R_J \leq k \cdot \mathsf{opt}$.*

Proof. Since F is an r_J-edge-cover, $J \cup F$ is an r-edge-cover. By the definition of F, for any r-edge-cover I, $p_F(b) \leq w_b \leq p_I(b)$ for all $b \in B$. In particular, if I is an optimal r-edge-cover, then

$$p_F(B) \leq \sum_{b \in B} w_b \leq \sum_{b \in B} p_I(b) = p_I(B) \leq \mathsf{opt} \ .$$

Also,

$$R_J = \sum_{b \in B} w_b r_J(b) \leq k \cdot \sum_{b \in B} w_b \leq k \cdot \mathsf{opt} \ .$$

Finally, $p_F(A) \leq R_J$ since

$$p_F(A) = \sum_{a \in A} p_F(a) \leq \sum_{a \in A} \sum_{e \in \delta_F(a)} c(e) = \sum_{e \in F} c(e) \leq \sum_{b \in B} w_b r_J(b) = R_J \ .$$

This concludes the proof of the lemma. ◻

Theorem 1 is deduced from Lemmas 1 and 2 as follows. We set γ to be constant strictly greater than 1, say $\gamma = 2$. Then $\theta = 1 - \frac{1}{2}\left(1 - \frac{1}{e}\right)$. Using binary search, we find the least integer τ such that the following procedure computes an edge set J satisfying $R_J \leq \tau$.

Initialization: $J \leftarrow \emptyset$.
Loop: Repeat $\lceil \log_{1/\theta} k \rceil$ times:
Apply the algorithm from Lemma 2:
- If it establishes that $\tau < \mathsf{opt}$ then return "ERROR" and STOP.
- Else do $J \leftarrow J \cup I$.

After computing J as above, we compute an edge set $F \subseteq E \setminus J$ as in Lemma 2. The edge-set $J \cup F$ is a feasible solution, by Lemma 2. We claim that for any $\tau \geq \mathsf{opt}$ the above procedure returns an edge set J satisfying $R_J \leq \tau$; thus binary search indeed applies. To see this, note that $R_\emptyset \leq k \cdot \mathsf{opt}$ and thus

$$R_J \leq R_\emptyset \cdot \theta^{\lceil \log_{1/\theta} k \rceil} \leq k \cdot \mathsf{opt} \cdot 1/k = \mathsf{opt} \leq \tau \ .$$

Consequently, the least integer τ for which the above procedure does not return "ERROR" satisfies $\tau \leq \mathsf{opt}$. Thus $p_J(V) \leq \lceil \log_{1/\theta} k \rceil \cdot (1+\gamma) \cdot \tau = O(\log k) \cdot \mathsf{opt}$. Also, by Lemma 2, $p_F(V) \leq \mathsf{opt} + R_J \leq 2\mathsf{opt}$. Consequently,

$$p_{J \cup F}(V) \leq p_J(V) + p_F(V) = O(\log k) \cdot \mathsf{opt} + 2\mathsf{opt} = O(\log k) \cdot \mathsf{opt} \ .$$

In the rest of this section we prove Lemma 1. It is sufficient to prove the statement in the lemma for the residual instance $((V, E \setminus J), r_J)$ with edge-costs restricted to $E \setminus J$; namely, we may assume that $J = \emptyset$. Let $R = R_\emptyset = \sum_{b \in B} w_b r(b)$.

Definition 3. *An edge $e \in E$ incident to a node $b \in B$ is* τ-cheap *if $c(e) \leq \frac{\tau\gamma}{R} \cdot w_b r(b)$.*

Lemma 3. *Let F be an r-edge-cover, let $\tau \geq p_F(B)$, and let*

$$I = \bigcup_{b \in B} \{e \in \delta_E(b) : c(e) \leq \frac{\tau\gamma}{R} \cdot w_b r(b)\}$$

be the set of τ-cheap edges in E. Then $R_{I \cap F} \leq R/\gamma$ and $p_I(B) \leq \gamma\tau$.

Proof. Let $D = \{b \in B : \delta_{F \setminus I}(b) \neq \emptyset\}$. Since for every $b \in D$ there is an edge $e \in F \setminus I$ incident to b with $c(e) > \frac{\tau\gamma}{R} \cdot w_b r(b)$, we have $p_{F \setminus I}(b) \geq \frac{\tau\gamma}{R} \cdot w_b r(b)$ for every $b \in D$. Thus

$$\tau \geq p_F(B) \geq p_{F \setminus I}(B) = \sum_{b \in D} p_{F \setminus I}(b) \geq \tau \cdot \frac{\gamma}{R} \sum_{b \in D} w_b r(b) \ .$$

This implies $\sum_{b \in D} w_b r(b) \leq R/\gamma$. Note that for every $b \in B \setminus D$, $\delta_F(b) \subseteq \delta_I(b)$ and hence $r_{I \cap F}(b) = r_F(b) = 0$. Thus we obtain:

$$R_{I \cap F} = \sum_{b \in B} w_b r_{I \cap F}(b) = \sum_{b \in D} w_b r_{I \cap F}(b) \leq \sum_{b \in D} w_b r(b) \leq R/\gamma \ .$$

To see that $p_I(B) \leq \gamma\tau$ note that

$$p_I(B) = \sum_{b \in B} p_I(b) \leq \frac{\tau\gamma}{R} \sum_{b \in B} w_b r(b) = \frac{\tau\gamma}{R} \cdot R = \tau\gamma \ .$$

This concludes the proof of the lemma. $\square$

In [3] it is proved that the following problem, which is a particular case of submodular function minimization subject to matroid and knapsack constraint (see [6]) admits a $\left(1 - \frac{1}{e}\right)$-approximation algorithm.

Bipartite Power-Budgeted Maximum Edge-Multi-Coverage (BPBMEM):
Instance: A bipartite graph $G = (A \cup B, E)$ with edge-costs $\{c(e) : e \in E\}$ and node-weights $\{w_v : v \in B\}$, degree bounds $\{r(v) : v \in B\}$, and a budget τ.
Objective: Find $I \subseteq E$ with $p_I(A) \leq \tau$ that maximizes

$$\mathsf{val}(I) = \sum_{v \in B} w_v \cdot \min\{d_I(v), r(v)\} \ .$$

The following algorithm computes an edge set as in Lemma 1.

1. Among the τ-cheap edges, compute a $\left(1 - \frac{1}{e}\right)$-approximate solution I to BPBMEM.
2. If $R_I \leq \theta R$ then return I, where $\theta = 1 - \left(1 - \frac{1}{\gamma}\right)\left(1 - \frac{1}{e}\right)$;
   Else declare "$\tau < \mathsf{opt}$".

Clearly, $p_I(A) \leq \tau$. By Lemma 3, $p_I(B) \leq \gamma\tau$. Thus $p_I(V) \leq p_I(A)+p_I(B) \leq (1+\gamma)\tau$.

Now we show that if $\tau \geq \mathsf{opt}$ then $R_I \leq \theta R$. Let F be the set of cheap edges in some optimal solution. Then $p_F(A) \leq \mathsf{opt} \leq \tau$. By Lemma 3 $R_F \leq R/\gamma$, namely, F reduces R by at least $R\left(1 - \frac{1}{\gamma}\right)$. Hence our $\left(1 - \frac{1}{e}\right)$-approximate solution I to BPBMEM reduces R by at least $R\left(1 - \frac{1}{e}\right)\left(1 - \frac{1}{\gamma}\right)$. Consequently, we have $R_I \leq R - R\left(1 - \frac{1}{e}\right)\left(1 - \frac{1}{\gamma}\right) = \theta R$, as claimed.

The proof of Theorem 1 for the case of general costs is complete.

2.3 A constant ratio approximation algorithm for unit costs

Bipartite MPEMC with unit costs is closely related to the Set-Multicover problem, that can be casted as follows.

Set-Multicover
Instance: A bipartite graph $G = (A \cup B, E)$ and demands $\{r(b) : b \in B\}$.
Objective: Find a subgraph H of G with $\deg_H(b) \geq r(b)$ for every $b \in B$; minimize $|H \cap A|$.

In fact, it is easy to see that Bipartite MPEMC with unit costs is equivalent to the following modification of Set-Multicover, where instead of minimizing $|H \cap A|$ we seek to minimize $|H \cap A| + |B|$; namely, the problem we consider is as follows.

Set-Multicover+
Instance: A bipartite graph $G = (A \cup B, E)$ and demands $\{r(b) : b \in B\}$.
Objective: Find a subgraph H of G with $\deg_H(b) \geq r(b)$ for every $b \in B$; minimize $|H \cap A| + |B|$.

Clearly, ratio ρ for Set-Multicover implies ratio ρ for Set-Multicover+. As Set-Multicover admits ratio $H(|B|)$, so is Set-Multicover+. On the other hand, Set-Multicover+ is APX-hard even for instances with $\max_{a \in A} \deg_G(a) = 3$, by a reduction from 3-Set-Cover. If $A = O(|B|)$ then the problem is clearly approximable within a constant; but we may have $|A| >> |B|$, if $k = \max_{b \in B} r(b)$ is large. We prove the following theorem that implies the second part of Theorem 1, and is also of independent interest.

Theorem 5. Set-Multicover+ *admits a randomized approximation algorithm with expected approximation ratio* ρ*, where* $\rho < 2.16851$ *is the real root of the qubic equation* $e(\rho-1)^3 = 2\rho$.

Let $\Gamma(a)$ denote the set of neighbors of a in G. Consider the following LP-relaxation for both Set-Multicover and Set-Multicover+

$$\min\left\{\sum_{a\in A} x_a : \sum_{a\in\Gamma(b)} x_a \geq r(b)\ \forall b\in B, 0\leq x_a\leq 1\ \forall a\in A\right\}. \qquad (1)$$

The value of a solution x to LP (1) is $x(A) = \sum_{a\in A} x_a$ in the Set-Multicover case, and $x(A) + |B|$ in the Set-Multicover+ case. Given a partial cover $S \subseteq A$, the residual demand of $b \in B$ is $r_S(b) = \max\{r(b) - |\Gamma(b)\cap S|, 0\}$. Let $\rho > 1$ be a parameter eventually set to be as in Theorem 5. Let $\gamma = \gamma(\rho) = \frac{(\rho-1)^2}{2\rho}$. Note that $\gamma = 1$ if, and only if, $\rho = 2+\sqrt{3}$, and that the value of ρ in Theorem 5 is less than $2+\sqrt{3}$. Let x be a feasible solution to LP (1), and let $S \subseteq A$ be obtained by choosing every $a \in A$ with probability $\min\{\rho\cdot x_a, 1\}$.

Lemma 4. *If* $x_a < 1/\rho$ *for all* $a \in A$ *then* $\Pr[r_S(b) \geq 1] \leq e^{-\gamma\cdot r(b)}$ *for every* $b \in B$.

Proof. Let $C(b) = \Gamma(b)\cap S$ be a random variable that counts the number of times b is "covered" by S. Clearly, $r_S(b) \geq 1$ if, and only if, $C(b) < r(b)$. The expectation of $C(b)$ is $\mu_b = \mathbb{E}[C(b)] = \sum_{a\in\Gamma(b)} \rho\cdot x_a \geq \rho\cdot r(b)$. Since the nodes in $\Gamma(b)$ are chosen independently, $C(b)$ is a sum of independent Bernoulli random variables. The statement now follows by applying the Chernoff bound:

$$\begin{aligned}\Pr[C(b) < r(b)] &= \Pr\left[C(b) < \left(1-\frac{\rho-1}{\rho}\right)\cdot\rho\cdot r(b)\right] \leq \\ &\leq \Pr\left[C(b) < \left(1-\frac{\rho-1}{\rho}\right)\cdot\mu_b\right] \leq \\ &\leq e^{-\frac{1}{2}\left(\frac{\rho-1}{\rho}\right)^2\mu_b} \leq e^{-\gamma\cdot r(b)}\ .\end{aligned}$$

□

Corollary 1. $\mathbb{E}[r_S(B)] \leq f(\rho)|B|$*, where* $f(\rho) = \frac{1}{e\gamma}$ *if* $1 < \rho \leq 2+\sqrt{3}$*. and* $f(\rho) = e^{-\gamma}$ *if* $\rho \geq 2+\sqrt{3}$.

Proof. Let $S' = \{a : x_a \geq 1/\rho\}$, let $r'(b) = r_{S'}(b) = \max\{r(b) - |\Gamma(b)\cap S'|, 0\}$, and let x' be defined by $x'_a = 0$ if $a \in S'$ and $x'_a = x_a$ otherwise. Note that x' is a feasible solution to LP (1) with the residual requirements r', and that $x'_a < 1/\rho$ for all $a \in A$. Thus by Lemma 4 we have

$$\mathbb{E}\left[r'(B)\right] = \sum_{b \in B} \mathbb{E}\left[r'(b)\right] \leq \sum_{b \in B} \Pr\left[r'(b) \geq 1\right] \cdot r'(b) \leq \sum_{b \in B} e^{-\gamma \cdot r'(b)} \cdot r'(b) \ .$$

Let $z = r'(b) \geq 1$ and $f(z) = e^{-\gamma z} \cdot z$. Then $f'(z) = e^{-\gamma z}(1 - \gamma z)$. Hence in the range $z \geq 1$, the function $f(z)$ has maximum value:

- $\frac{1}{e\gamma}$ if $\gamma \leq 1$ (namely, if $1 < \rho \leq 2 + \sqrt{3}$), attained at $z = 1/\gamma$.
- $e^{-\gamma}$ if $\gamma \geq 1$ (namely, if $\rho \geq 2 + \sqrt{3}$), attained at $z = 1$.

The statement follows. □

Now we finish the proof of Theorem 5. The algorithm is as follows. We compute an optimal solution x to LP (1), and then an edge set S as in Corollary 1. For every $b \in B$ let A_b be a set of $r_S(b)$ neighbors in $\Gamma(b) \setminus S$, and let $S' = \bigcup_{b \in B} A_b$. The solution returned is $S \cup S'$. Note that $\mathbb{E}[|S'|] \leq \mathbb{E}[r_S(B)]$. Thus by Corollary 1, the expected size of our solution is bounded by

$$\mathbb{E}(|S|) + \mathbb{E}(r_S(B)) + |B| \leq \rho x(A) + f(\rho)|B| + |B| \leq \max\{\rho, f(\rho) + 1\}(x(A) + |B|) \ .$$

Consequently, as $x(A) + |B|$ is a lower bound on the optimal solution value, the approximation ratio is bounded by $\max\{\rho, f(\rho) + 1\}$. Solving the equation $\rho = f(\rho) + 1$ for $f(\rho) = \frac{1}{e\gamma} = \frac{2\rho}{e(\rho-1)^2}$ gives the result.

The proof of Theorem 5, and thus also of Theorem 1 for the case of uniform cost, is complete.

3 Proof of Theorem 2

We say that an edge set $F \subseteq E$ *covers* a node set $U \subseteq V$, or that F is a U-*cover*, if $\delta_F(v) \neq \emptyset$ for every $v \in U$. Consider the following auxiliary problem:

Restricted Minimum-Power Edge-Cover
Instance: A graph $G = (V, E)$ with edge-costs $\{c(e) : e \in E\}$, $U \subseteq V$, and degree bounds $\{\ell_v : v \in U\}$.
Objective: Find a power assignment $\{\pi(v) : v \in V\}$ that minimizes $\sum_{v \in V} \pi(v)$, such that $\pi(v) \geq \ell_v$ for all $v \in U$, and such that the edge set $F = \{e = uv \in E : \pi(u), \pi(v) \geq c(e)\}$ covers U.

Later, we will prove the following lemma.

Lemma 5. Restricted Minimum-Power Edge-Cover *admits a 3/2-approximation algorithm.*

Theorem 2 is deduced from Lemma 5 and the following statement.

Lemma 6. *If* Restricted Minimum-Power Edge-Cover *admits a ρ-approximation algorithm, then* Minimum-Power Edge-Multi-Cover *admits a $(k-1+\rho)$-approximation algorithm.*

Proof. Consider the following algorithm.

1. Let $\pi(v)$ be the power assignment computed by the ρ-approximation algorithm for $\mathsf{Restricted\ Minimum\text{-}Power\ Edge\text{-}Cover}$ with $U = \{v \in V : r(v) \geq 1\}$ and bounds $\ell_v = w_v$ for all $v \in U$. Let $F = \{e = uv \in E : \pi(u), \pi(v) \geq c(e)\}$.
2. For every $v \in V$ let I_v be the edge-set obtained by picking the least cost $r_F(v)$ edges in $\delta_{E \setminus F}(v)$ and let $I = \cup_{v \in V} I_v$.

Clearly, $F \cup I$ is a feasible solution to Minimum-Power Edge-Multi-Cover. Let opt denote the optimal solution value for Minimum-Power Edge-Multi-Cover. In what follows note that $\pi(V) \leq \rho \cdot \mathsf{opt}$ and that $\sum_{v \in V} w_v \leq \mathsf{opt}$.

We claim that

$$p_{I \cup F}(V) \leq \pi(V) + (k-1) \cdot \mathsf{opt} .$$

As $\pi(V) \leq \rho \cdot \mathsf{opt}$, this implies $p_{I \cup F}(V) \leq (\rho + k - 1) \cdot \mathsf{opt}$.

For $v \in V$ let Γ_v be the set of neighbors of v in the graph (V, I_v). The contribution of each edge set I_v to the total power is at most $p_{I_v}(\Gamma_v) + p_{I_v}(v)$. Note that $\pi(v) \geq p_{I_v}(v)$ and $\pi(v) \geq p_F(v)$ for every $v \in V$, hence $p_{F \cup I_v}(v) \leq \pi(v)$. This implies

$$p_{F \cup I}(V) \leq \sum_{v \in V} (\pi(v) + p_{I_v}(\Gamma_v)) = \pi(V) + \sum_{v \in V} p_{I_v}(\Gamma_v) .$$

Now observe that $|\Gamma_v| = |I_v| = r_F(v) \leq k - 1$ and that $p_{I_v}(u) \leq w_v$ for every $u \in \Gamma_v$. Thus

$$p_{I_v}(\Gamma_v) \leq (k-1) \cdot w_v \qquad \forall v \in V .$$

Finally, using the fact that $\sum_{v \in V} w_v \leq \mathsf{opt}$, we obtain

$$p_{F \cup I}(V) \leq \pi(V) + \sum_{v \in V} p_{I_v}(\Gamma_v) \leq \pi(V) + (k-1) \sum_{v \in V} w_v \leq \pi(V) + (k-1) \cdot \mathsf{opt} .$$

This finishes the proof of the lemma. □

In the rest of this section we prove Lemma 5.

We reduce Restricted Minimum-Power Edge-Cover to the following problem that admits an exact polynomial time algorithm, c.f. [10].

Minimum-Cost Edge-Cover:
Instance: A multi-graph (possibly with loops) $G = (U, E)$ with edge-costs $\{c(e) : e \in E\}$.
Objective: Find a minimum cost edge-set $F \subseteq E$ that covers U.

Our reduction is not approximation ratio preserving, but incurs a loss of $3/2$ in the approximation ratio. That is, given an instance (G, c, U, ℓ) of Restricted Minimum-Power Edge-Cover, we construct in polynomial time an instance (G', c') of Minimum-Cost Edge-Cover such that:

(i) For any U-cover I' in G' corresponds a feasible solution π to (G, c, U, ℓ) with $\pi(V) \leq c'(I')$.

(ii) $\mathsf{opt}' \le 3\mathsf{opt}/2$, where opt is an optimal solution value to Restricted Minimum-Power Edge-Cover and opt' is the minimum cost of a U-cover in G'.

Hence if I' is an optimal (min-cost) solution to (G', c'), then $\pi(V) \le c'(I') \le 3\mathsf{opt}/2$.

Clearly, we may set $\ell_v = 0$ for all $v \in V \setminus U$. For $I \subseteq E$ let

$$D(I) = \sum_{v \in V} \max\{p_I(v) - \ell_v, 0\} \ .$$

Here is the construction of the instance (G', c'), where $G' = (U, E')$ and E' consists of the following three types of edges, where for every edge $e' \in E'$ corresponds a set $I(e') \subseteq E$ of one edge or of two edges.

1. For every $v \in U$, E' has a loop-edge $e' = vv$ with $c'(vv) = \ell_v + D(\{vu\})$ where vu is is an arbitrary chosen minimum cost edge in $\delta_E(v)$.
   Here $I(e') = \{vu\}$.
2. For every $uv \in E$ such that $u, v \in U$, E' has an edge $e' = uv$ with $c'(uv) = \ell_u + \ell_v + D((\{uv\})$.
   Here $I(e') = \{uv\}$.
3. For every pair of edges $ux, xv \in E$ such that $c(ux) \ge c(xv)$, E' has an edge $e' = uv$ with $c'(uv) = \ell_v + \ell_u + D(\{ux, xv\})$.
   Here $I(e') = \{ux, xv\}$.

Lemma 7. *Let $I' \subseteq E'$ be a U-cover in G', let $I = \cup_{e \in I'} I(e)$, and let π be a power assignment defined on V by $\pi(v) = \max\{p_I(v), \ell_v\}$. Then $\pi(V) \le c'(I')$, I is a U-cover in G, and π is a feasible solution to (G, c, U, ℓ).*

Proof. We have that I is a U-cover in G, by the definition of I and since $I(e')$ covers both endnodes of every $e' \in E'$. By the definition of π, we have that $I \subseteq \{e = uv \in E : \pi(u), \pi(v) \ge c(e)\}$. Hence π is a feasible solution to (G, c, U, ℓ), as claimed.

We prove that $\pi(V) \le c'(I')$. For $e' = uv \in E'$ let $\ell(e') = \ell_v$ if e' is a type 1 edge, and $\ell(e') = \ell_u + \ell_v$ otherwise. Note that $\pi(v) = \max\{p_I(v), \ell(v)\} = \ell_v + \max\{p_I(v) - \ell(v), 0\}$, hence

$$\pi(V) = \sum_{v \in U} \ell_v + \sum_{v \in V} \max\{p_I(v) - \ell(v), 0\} = \sum_{v \in U} \ell_v + D(I) \ .$$

By the definition of $\ell(e')$ and since I' is a U-cover $\sum_{v \in U} \ell_v \le \sum_{e' \in I'} \ell(e')$. Also, $D(I) = D\left(\cup_{e' \in I'} I(e')\right)$, by the definition of I. Thus we have

$$\sum_{v \in U} \ell_v + D(I) \le \sum_{e' \in I'} \ell(e') + D\left(\cup_{e' \in I'} I(e')\right) \ .$$

It is easy to see that

$$D\left(\cup_{e' \in I'} I(e')\right) \le \sum_{e' \in I'} D(I(e')) \ .$$

Finally, note that $\ell(e') + D(I(e')) = c'(e')$ for every $e' \in I'$ (if e' is a type 1 edge, this follows from our assumption that $\ell_v \geq \min\{c(e) : e \in \delta_E(v)\}$). Combining we get

$$\begin{aligned}
\pi(V) &= \sum_{v \in U} \ell_v + D(I) \leq \\
&\leq \sum_{e' \in I'} \ell(e') + D\left(\cup_{e' \in I'} I(e')\right) \leq \\
&\leq \sum_{e' \in I'} \ell(e') + \sum_{e' \in I'} D(I(e')) = \\
&= \sum_{e' \in I'} \left(\ell(e') + D(I(e'))\right) = \\
&= \sum_{e' \in I'} c'(e') = c'(I') \; .
\end{aligned}$$

□

Lemma 8. *Let $\{\pi(v) : v \in V\}$ be a feasible solution to an instance (G, c, U, ℓ) of* Restricted Minimum-Power Edge-Cover*. Then there exists a U-cover I' in G' such that $c'(I') \leq 3\pi(V)/2$.*

Proof. Let $I \subseteq \{e = uv \in E : \pi(u), \pi(v) \geq c(e)\}$ be an inclusion minimal U-cover. We may assume that $\pi(v) = \max\{p_I(v), \ell_v\}$ for every $v \in V$. Since any inclusion minimal U-cover is a collection of node disjoint stars, it is sufficient to prove the statement for the case when I is a star. Then I has at most one node not in U, and if there is such a node, then it is the center of the star, if $|I| \geq 2$; in the case I consists of a single edge e, then we define the center of I to be the endnode of e in $V \setminus U$ if such exists, or an arbitrary endnode of e otherwise.

We define a U-cover I' in G', and show that

$$c'(I') \leq \frac{3}{2} \sum_{v \in V} \max\{p_I(v), \ell_v\} = \frac{3}{2}\pi(V) \; . \tag{2}$$

Let v_0 be the center of I and let $\{v_i : 1 \leq i \leq d\}$ be the leaves of I ordered by descending order of costs $c(v_0 v_i) \geq c(v_0 v_{i+1})$. The U-cover $I' \subseteq E'$ is defined as follows. We cover each pair v_{2i-1}, v_{2i}, $i = 1, \ldots, \lfloor d/2 \rfloor$, by a type 3 edge. This covers all the nodes except v_0, and maybe v_d if d is odd. We add an additional edge f of type 1 or 2, if there are nodes in U (v_0 and/or v_d) that remain uncovered by the picked type 3 edges. Formally, we have the following 4 cases, see Figure 1.

1. d is even and $v_0 \notin U$, see Figure 1(a). Then U is covered by type 3 edges.
2. d is odd, and $v_0 \notin U$, see Figure 1(b). Then we add a type 1 edge f to cover v_d.
3. d is odd and $v_0 \in U$, see Figure 1(c). Then we add a type 2 edge f to cover v_0, v_d.

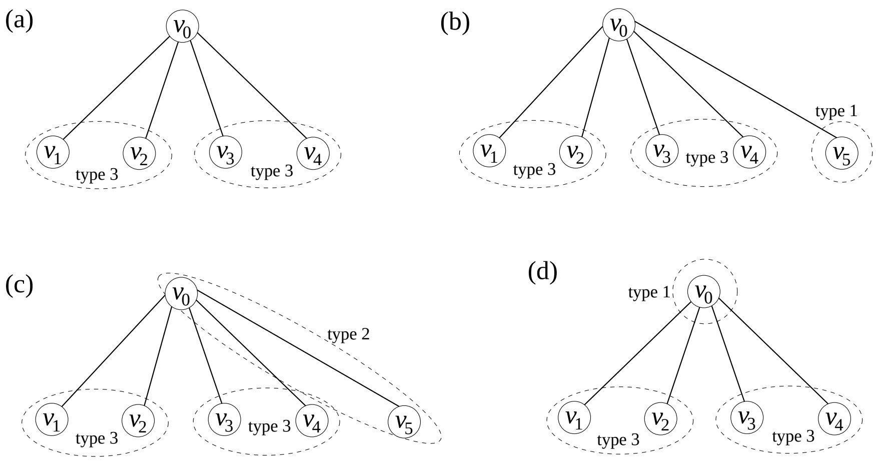

Fig. 1. Illustration to the definition of the U-cover I'.

4. d is even and $v_0 \in U$, see Figure 1(d). Then we add a type 1 edge f to cover v_0.

Consider a type 3 edge $v_{2i-1}v_{2i} \in I'$. Let $q_i = \max\{c(v_{2i-1}v_0) - \ell_{v_0}, 0\}$. Note that $c'(v_{2i-1}v_{2i}) \leq \pi(v_{2i-1}) + \pi(v_{2i}) + q_i$. The key point is that

$$q_i \leq \frac{1}{2}(\pi(v_{2i-3}) + \pi(v_{2i-2})) \quad i = 2, \ldots, \lfloor d/2 \rfloor \ .$$

This is since $q_i \leq c(v_0v_{2i-1}) \leq \frac{1}{2}\left(c(v_0v_{2i-3}) + c(v_0v_{2i-2})\right)$ while $c(v_0v_j) \leq \pi(v_j)$. Therefore,

$$\sum_{i=1}^{d/2} c'(v_{2i-1}v_{2i}) \leq \sum_{i=1}^{d/2}[\pi(v_{2i-1}) + \pi(v_{2i}) + q_i] \leq \sum_{i=1}^{2\lfloor d/2 \rfloor} \pi(v_i) + q_1 + \frac{1}{2}\sum_{i=1}^{d-2}\pi(v_i)$$

Now, we prove that (2) hold in each one of our four cases.

1. $v_0 \notin U$ and d is even. Note that $q_1 \leq c(v_0v_1) \leq \pi(v_0)$. Then:

$$c'(I') = \sum_{i=1}^{d/2} c'(e_i) \leq \frac{3}{2}\sum_{i=1}^{d}\pi(v_i) + q_1 \leq \frac{3}{2}\sum_{i=1}^{d}\pi(v_i) + \pi(v_0) \leq \frac{3}{2}\sum_{i=0}^{d}\pi(v_i)$$

2. $v_0 \notin U$ and d is odd. In this case $f = v_dv_d$ is a loop type 1 edge, so $c'(f) \leq \pi(v_d) + \max(c(v_0v_d) - \ell_{v_0}, 0)$. This implies

$$\begin{aligned} q_1 + c'(f) &\leq c(v_0v_1) + c(v_0v_d) + \pi(v_d) \leq \pi(v_0) + \frac{1}{2}[\pi(v_0) + \pi(v_d)] + \pi(v_d) \\ &= \frac{3}{2}\left(\pi(v_0) + \pi(v_d)\right) \ . \end{aligned}$$

Thus

$$c'(I') = \sum_{i=1}^{d/2} c'(e_i) + c'(f) \le \frac{3}{2}\sum_{i=1}^{d-1} \pi(v_i) + c'(f) + q_1 \le \frac{3}{2}\sum_{i=0}^{d} \pi(v_i)$$

3. $v_0 \in U$ and d is odd. In this case $f = v_0 v_d$, so $c'(f) \le \max(\ell_{v_0}, c(v_0 v_d)) + \pi(v_d)$. This implies $q_1 + c'(f) \le c(v_0v_1) + c(v_0v_d) + \pi(v_d) \le \frac{3}{2}(\pi(v_0) + \pi(v_d))$. Thus

$$c'(I') = \sum_{i=1}^{d/2} c'(e_i) + c'(f) \le \frac{3}{2}\sum_{i=1}^{d-1} \pi(v_i) + c'(f) + q_1 \le \frac{3}{2}\sum_{i=0}^{d} \pi(v_i) \ .$$

4. $v_0 \in U$ and d is even. In this case $f = v_0 v_0$ is a loop type 1 edge, so $c'(f) \le \ell_{v_0} + c(v_0 v_d) \le \ell_{v_0} + \frac{1}{2}(\pi(v_{d-1}) + \pi(v_d))$. This implies $q_1 + c'(f) \le c(v_0v_1) + \frac{1}{2}(\pi(v_{d-1}) + \pi(v_d))$. Thus

$$\begin{aligned} c'(I') &= \sum_{i=1}^{d/2} c'(e_i) + c'(f) \le \sum_{i=1}^{d} \pi(v_i) + \frac{1}{2}\sum_{i=1}^{d-2} \pi(v_i) + q_1 + c'(f) \\ &\le \frac{3}{2}\sum_{i=1}^{d} \pi(v_i) + \pi(v_0) \le \sum_{i=0}^{d} \pi(v_i) \ . \end{aligned}$$

This concludes the proof of the lemma. □

As was mentioned, Lemmas 7 and 8 imply Lemma 5. Lemmas 5 and 6 imply Theorem 2, hence the proof of Theorem 2 is now complete.

4 Conclusions and open problems

The main results of this paper are two new approximation algorithm for MPEMC: one with ratio $O(\log k)$ for general costs, and the other with constant ratio for uniform costs. This improves the ratio $O(\log(nk)) = O(\log n)$ of [3]. We also gave a $(k + 1/2)$-approximation algorithm, which is better than our $O(\log k)$-approximation algorithm for small values of k (roughly $k \le 6$).

The main open problem is whether for general costs, the ratio $O(\log k)$ shown in this paper is tight, or the problem admits a constant ratio approximation algorithm.

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. M. Hajiaghayi, G. Kortsarz, V. Mirrokni, and Z. Nutov. Power optimization for connectivity problems. *Math. Programming*, 110(1):195–208, 2007.

3. G. Kortsarz, V. Mirrokni, Z. Nutov, and E. Tsanko. Approximating minimum-power degree and connectivity problems. *Algorithmica*, 58, 2010.
4. G. Kortsarz and Z. Nutov. Approximating minimum-power edge-covers and $2, 3$-connectivity. *Discrete Applied Mathematics*, 157:1840–1847, 2009.
5. Y. Lando and Z. Nutov. On minimum power connectivity problems. *J. of Discrete Algorithms*, 8(2):164–173, 2010.
6. J. Lee, V. Mirrokni, V. Nagarajan, and M. Sviridenko. Maximizing nonmonotone submodular functions under matroid or knapsack constraints. *SIAM J. Discrete Mathematics*, 23(4):20532078, 2010.
7. Z. Nutov. An almost $O(\log k)$-approximation for k-connected subgraphs. In *SODA*, pages 912–921, 2009.
8. Z. Nutov. Approximating minimum power covers of intersecting families and directed edge-connectivity problems. *Theoretical Computer Science*, 411(26-28):2502–2512, 2010.
9. Z. Nutov. Approximating minimum power k-connectivity. *Ad Hoc & Sensor Wireless Networks*, 9(1-2):129–137, 2010.
10. A. Schrijver. *Combinatorial Optimization, Polyhedra and Efficiency.* Springer-Verlag Berlin, Heidelberg New York, 2004.