Power Optimization for Connectivity Problems

Mohammad T. Hajiaghayi[1], Guy Kortsarz[2], Vahab S. Mirrokni[1], and Zeev Nutov[3]

[1] Computer Science and Artificial Intelligence Laboratory, MIT
{mirrokni, hajiagha}@csail.mit.edu
[2] Department of Computer Science, Rutgers University-Camden
guyk@crab.rutgers.edu
[3] Computer Science Division, The Open University of Israel
nutov@openu.ac.il

Abstract. Given a 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 the nodes of this graph. Motivated by applications in wireless multi-hop networks, we consider four fundamental problems under the power minimization criteria: the *Min-Power b-Edge-Cover* problem (MPb-EC) where the goal is to find a min-power subgraph so that the degree of every node v is at least some given integer $b(v)$, the *Min-Power k-node Connected Spanning Subgraph problem* (MPk-CSS), *Min-Power k-edge Connected Spanning Subgraph problem* (MPk-ECSS), and finally the *Min-Power k-Edge-Disjoint Paths* problem in directed graphs (MPk-EDP). We give an $O(\log^4 n)$-approximation algorithm for MPb-EC. This gives an $O(\log^4 n)$-approximation algorithm for MPk-CSS for most values of k, improving the best previously known $O(k)$-approximation guarantee. In contrast, we obtain an $O(\sqrt{n})$ approximation algorithm for MPk-ECSS, and for its variant in directed graphs (i.e., MPk-EDP), we establish the following inapproximability threshold: MPk-EDP cannot be approximated within $O(2^{\log^{1-\varepsilon} n})$ for any fixed $\varepsilon > 0$, unless NP-hard problems can be solved in quasi-polynomial time.

1 Introduction

Wireless multihop networks are an important subject of study due to their extensive applications (see e.g., [8, 24]). A large research effort focused on performing network tasks while minimizing the power consumption of the radio transmitters of the network. In ad-hoc networks, a range assignment to radio transmitters means to assign a set of powers to mobile devices. We consider finding a range assignment for the nodes of a network such that the resulting communication network satisfies some prescribed properties, and such that the total power is minimized. Specifically, we consider "min-power" variants of three extensively studied "min-cost" problems: the b-Edge Cover problem and the k-Connected Spanning Subgraph Problem in undirected networks, and the k-Edge-Disjoint Paths problem in directed networks.

M. Jünger and V. Kaibel (Eds.): IPCO 2005, LNCS 3509, pp. 349–361, 2005.

In wired networks, generally we want to find a subgraph with the minimum cost instead of the minimum power. This is the main difference between the optimization problems for wired versus wireless networks. The power model for undirected graphs corresponds to the static symmetric multi-hop ad-hoc wireless networks with omnidirectional transmitters. This model is justified and used in several other papers [3, 4, 14].

An important network task is assuring high fault-tolerance ([1–4, 11, 18]). The simplest version 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 NP-hard even in the Euclidean plane [9]. There are several localized and distributed heuristics to find the range assignment to keep the network connected [18, 24, 25]. Constant approximation guarantees for the min-power spanning tree problem are given in [4, 14]. For general k, the best previously known approximation ratio for MPk-CSS was $2k = O(k)$ [5, 11, 19].

Min-cost k-connected and k-edge connected spanning subgraph problems were extensively studied [7, 10, 12, 13, 16, 17]. While the min-cost k-edge connected spanning subgraph problem admits a 2-approximation algorithm [12, 13], no constant approximation guarantee is known for the min-cost k-connected spanning subgraph problem. The best known approximation ratios for the latter are $O(\ln k \cdot \min\{\sqrt{k}, \frac{n+k}{n-k} \ln k\})$ [17] and $O(\ln k)$ for $n \geq 2k^2$ [7].

The notation and preliminaries used in the paper are as follows. 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. We will often omit the subscripts G, E if they are clear from the context. For brevity, $\delta_E(X) = \delta_E(X, V - X)$, and $\deg_E(X) = |\delta_E(X)|$ is the *degree of* X. For a function g on a groundset U and $S \subseteq U$ let $g(S) = \sum_{u \in S} g(u)$. 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, $p(v) = \max_{e \in \delta_E(v)} c(e)$. 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 $c(G)/\sqrt{m/2}$[1]. The following statement whose proof is presented in Section 3 shows that these are the extremal cases also for general edge costs.

Lemma 1. *For any graph* $G = (V, E)$ *holds:* $c(G)/\sqrt{|E|/2} \leq p(G) \leq 2c(G)$. *For a forest* T, $c(T) \leq p(T) \leq 2c(T)$.

Throughout the paper, let $\mathcal{G} = (V, \mathcal{E})$ denote the input graph with nonnegative costs on the edges; n denotes the number of nodes in $\mathcal{G}$, and m the number of edges in $\mathcal{G}$. Let opt denote the optimal solution value of an instance at hand. Given $\mathcal{G}$, our goal is to find a minimum power spanning subgraph G of $\mathcal{G}$ that satisfies some prescribed property. In undirected graphs, we consider the following two variants. Given an integral function b on V, we say that G (or E) is

[1] In this paper, we ignore that some numbers might not be integers, since the adaption to floors and ceilings is immediate.

a *b-edge cover* if $\deg_G(v) \geq b(v)$ for every $v \in V$, where $\deg_G(v) = \deg_E(v)$ is the degree of v in G. In the *Minimum Power b-Edge Cover Problem* (MPb-EC), G is required to be a b-edge cover; the *Minimum Power k-Edge Cover Problem* (MPk-EC) is a particular case when $b(v) = k$ for all $v \in V$. It is easy to see that the greedy algorithm that for every $v \in V$ picks the lightest $b(v)$ edges incident to v is a $(k+1)$-approximation algorithm for MPb-EC, where $k = \max_{v \in V} b(v)$. The following simple example shows that the $(k+1)$-approximation ratio is tight for this greedy algorithm. Take $k+1$ stars with k leaves each, and add edges of a clique on their centers. All edges have unit costs. Set $b(v) = k$ if v is a star center and $b(v) = 0$ otherwise. The greedy algorithm may pick the edges of the stars, thus getting a solution of value $(k+1)^2$. The optimal solution is obtained by picking the clique, and has power $k+1$. This example easily extends to MPk-EC. We prove:

Theorem 1. MPb-EC *is APX-hard. It admits an $O(\log^4 n)$-approximation algorithm.*

A graph G is *k-(node) connected* if there are k internally disjoint paths between every pair of its nodes. In the *Minimum Power k-Connected Spanning Subgraph Problem* (MPk-CSS) G is required to be k-connected. The motivation of the "min-power" variant for wireless networks is similar to the one of the "min-cost" variants for wired networks, e.g., for MPk-CSS we require that the network remains connected even in failure of up to $k-1$ terminals. The problem admits an $O(k)$-approximation algorithm [5, 11]. We prove:

Theorem 2. MPk-CSS *is APX-hard. Unless $k = n - o(n)$,* MPk-CSS *admits an $O(\log^4 n)$-approximation algorithm. For $k = n - o(n)$,* MPk-CSS *admits an $O(\sqrt{k})$-approximation algorithm.*

Theorem 2 is proved by combining Theorem 1 with part (i) of the following theorem, and using the currently best known approximation guarantees [7, 16] for the Min-Cost k-Connected Spanning Subgraph problem.

Theorem 3. *(i) If there exists an α-approximation algorithm for the Min-Cost k-Connected Spanning Subgraph problem and a β-approximation algorithm for* MPk-EC *then there exists a $(2\alpha + \beta)$-approximation algorithm for* MPk-CSS*.*

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

Note that part (ii) of Theorem 3 implies that MPk-CSS is almost as hard to approximate as the Min-Cost k-Connected Spanning Subgraph problem.

We also consider the *Min-Power k-Edge Connected Spanning Subgraph* (MPk-ECSS) problem where G is required to be k-edge connected. This problem admits a $O(k)$-approximation algorithm [11]. We prove:

Theorem 4. MPk-ECSS *is APX-hard and admits an $O(\sqrt{n})$-approximation algorithm.*

Power optimization problems were considered in asymmetric networks as well [14]. This setting is mainly motivated for the purpose of broadcasting or multicasting in multihop wireless networks. In this case, the power of a node v is the maximum cost of an edge outgoing from v. We give some evidence that minimum-power connectivity problems in directed graphs are hard by showing a strong inapproximability result for a simple variant: the problem of finding the minimum-power subgraph that contains k edge-disjoint directed (s,t)-paths. We call it the *Min-Power k-Edge-Disjoint Paths* (MPk-EDP) problem, since it is the "min-power variant" of the Min-Cost k-Edge-Disjoint Paths problem. We prove the following strong inapproximability result for MPk-EDP, in contrast to the polynomial solvability of the "min-cost" case.

Theorem 5. MPk-EDP *cannot be approximated within $O(2^{\log^{1-\varepsilon} n})$ for any fixed $\varepsilon > 0$, unless NP-hard problems can be solved in quasi-polynomial time.*

We also note that, in contrast, the problem of finding Minimum Power k-Vertex-Disjoint Paths (from s to t) in directed graphs can be solved in polynomial time as follows. First, we can assume that we know the power p of s (there are at most n possible values) and thus we know all optimum edges incident to s. Now, we give zero cost to all these edge (whose original costs were at most p), delete all the other edges incident to s, and compute the minimum cost k internally vertex-disjoint paths using the polynomial-time min-cost k-flow algorithm of Orlin [21], and a flow decomposition. As the outdegree of every internal node is one, and the outdegree of t is zero, this is an optimal solution to our minimum power vertex-disjoint case.

Table 1 summarizes our main results.

Table 1. Our approximation ratios and hardness results (α is the best approximation ratio for the Min-Cost k-Connected Spanning Subgraph problem)

Problem	Approximation Ratio	Hardness
MPb-EC	$\min(O(\log^4 n), k+1)$	APX-hard
MPk-CSS	$\min(O(\log^4 n) + 2\alpha, k(1+o(1))$	APX-hard, $\Omega(\alpha)$
MPk-ECSS	$O(\sqrt{n})$	APX-hard
MPk-EDP	–	$\Omega\left(2^{\log^{1-\varepsilon} n}\right)$

Theorem 1 is proved in Section 2, Lemma 1 and Theorems 2, 3, and 4 are proved in Section 3, and finally Theorem 5 is proved in Section 4. In the rest of this section we show that already very restricted instances of MPb-EC, MPk-CSS, and MPk-ECSS are APX hard, thus proving the hardness results of Theorem 1, 2, and 4.

Theorem 6. MPk-EC*,* MPk-CSS*, and* MPk-ECSS *are APX-hard even for $k = 1$.*

Proof. To prove the theorem, we will use the following well-known formulation of the Set-Cover Problem (SCP); in this formulation, J is the incidence graph

of sets and elements, where A is the family of sets and B is the universe (we denote the edge set by I).
Input: A bipartite graph $J = (A \cup B, I)$ without isolated nodes.
Output: A minimum size subset $T \subseteq A$ such that every node in B has a neighbor in T.

The reduction is as follows. Given an instance $J = (A \cup B, I)$ for the SCP, we construct a graph $G = (V \cup \{r\}, E)$ with edge cost function c by setting $c(e) = 1$ for every $e \in E$, adding a new node r and edges of cost zero from r to every $a \in A$; for MPk-EC we set $b(v) = 1$ for every $v \in V$, and for MPk-CSS and MPk-ECSS we set $k = 1$. It is easy to see that SCP has a solution of size τ if and only if the obtained instance of MPk-EC (MPk-CSS/MPk-ECSS) has a solution of size $|B| + \tau$.

A 4-bounded instance SC-4 of SC is one in which all sets have size at most 4, that is $\deg_J(a) \leq 4$ for every $a \in A$. Any solution to SC-4 has size $\geq |B|/4$. Thus any solution of power $|B| + \tau$ in our MPk-EC (MPk-CSS/MPk-ECSS) instance of power at most $(1 + \varepsilon)(|B| + \tau)$ gives us a solution to SC-4 of size at most $\tau + \tau\varepsilon + |B|\varepsilon \leq (1+5\varepsilon)\tau$. Consequently, a $(1+\varepsilon)$-approximation to MPk-EC gives a $(1 + 5\varepsilon)$-approximation to SC-4. Since SC-4 is APX-hard [22], APX-hardness of MPk-EC (MPk-CSS/MPk-ECSS) follows.

Finally to obtain APX-hardness of MPk-EC (MPk-CSS/MPk-ECSS) for $k > 1$, we add vertices $t_1, \ldots, t_{k-1}$ to graph G constructed above and we add edges of zero cost from them to all previous vertices. The proof again follows from the fact that each vertex corresponding to an element should be adjacent to at least one edge of cost one. □

2 Proof of Theorem 1

In this section, we present the proof of Theorem 1. Throughout this section we assume that $c(e) \in \{1, \ldots, n^4\}$ for every $e \in \mathcal{E}$. In particular, opt $\leq n^6$. Indeed, let c be the least integer so that $\{e \in E : u_e \leq c\}$ is a b-edge cover. Edges of cost $\geq cn^2$ do not belong to any optimal solution, and thus deleted from the graph. Edges of cost $\leq c/n^2$ in fact get zero costs, as adding all of them to the solution affects only the constant in the approximation ratio (we also update $b(v)$'s, $v \in V$, accordingly). This gives an instance with $c_{\max}/c_{\min} \leq n^4$, where $c_{\max}$ and $c_{\min}$ denote the maximum and the minimum nonzero cost of an edge in $\mathcal{E}$, respectively. Further, for every $e \in \mathcal{E}$ set $c(e) \leftarrow \lceil c(e)/c_{\min} \rceil$. It is easy to see that the loss incurred in the approximation ratio is only a constant, which is negligible in our context.

Let $b(V) = \sum_{v \in V} b(v)$. For an edge set F and $v \in V$, let

$$b_F(v) = \max\{b(v) - \deg_F(v), 0\}$$

be the *residual deficiency* of v w.r.t. F (so $b(v) = b_\emptyset(v)$). Also, $b_F(V) = \sum_{v \in V} b_F(v)$. Our algorithm runs with a parameter τ that should be set to $\tau = $ opt to achieve the claimed approximation ratio. Specifically, we will prove:

Lemma 2. *There exists a polynomial time algorithm that given an instance of* MPb-EC *and an integer τ, either returns an edge set $E' \subseteq \mathcal{E}$ such that*

$$p_{E'}(V) = \tau \cdot O(\lg^4 n) \tag{1}$$

$$b_{E'}(V) \le \lg^3 n \tag{2}$$

or establishes that $\tau <$ opt.

Note that if $\tau <$ opt, the algorithm may return an edge set E' that satisfies (1) and (2). Let us now show that Lemma 2 implies Theorem 1. Since opt is not known, we apply binary search to find the minimum integer τ so that an edge set E' satisfying (1) and (2) is returned; then $p_{E'}(V) = \text{opt} \cdot O(\lg^4 n)$ (note that binary search for appropriate τ requires $O(\lg n^6) = O(\lg n)$ iterations). Then we apply the greedy algorithm on $\mathcal{G} - E'$ to compute a $b_{E'}$-edge cover E'' of power $\le$ opt $\cdot$ $(\lg^3 n + 1)$. Then $E = E' \cup E''$ is a feasible solution, and $p_{E' \cup E''}(V) \le p_{E'}(V) + p_{E''}(V) = \text{opt} \cdot O(\lg^4 n)$. Thus Lemma 2 implies Theorem 1.

The proof of Lemma 2 follows. Let $D(F) = \{v \in V : b_F(v) > 0\}$ be the set of *deficient* nodes w.r.t. F, and $D = D(\emptyset) = \{v \in V : b(v) > 0\}$. Let $\mu = \min\{b(v) : v \in D\}$.

Lemma 3. *There exists a polynomial time algorithm that given an instance of* MPb-EC *with* $\max\{b(v) : v \in D\} \le r\mu$ *and integers W, T, and τ, returns an edge set F such that*

$$p(F) \le 2\left(W|D| + b(V)\lg W/T\right), \tag{3}$$

and if $\tau \ge$ opt then

$$b_F(V) \le \tau\left(T \lg W + \mu r/(2W)\right). \tag{4}$$

Proof. Let $E_0 = \{e \in E : 1 \le c(e) \le 2\}$ and $E_i = \{e \in E : 2^i + 1 \le c(e) \le 2^{i+1}\}$ for $i = 1, \ldots, \lg W$. Consider the following algorithm that starts with $F = \emptyset$:
For $i = 0$ to $\lg W$ do:
While there is $v \in V$ with $|\delta_{E_i}(v, D(F))| \ge 2^i T$ do $F \leftarrow F + \delta_{E_i}(v, D(F))$
End For
It is easy to see that the algorithm is polynomial. Let F be the edge set computed by the algorithm. Let $p' = \sum_{v \in D} p_F(v)$ and $p'' = \sum_{v \in V-D} p_F(v) = p_F(V) - p'$. The following two claims show that (3) holds.

Claim: $p' \le 2W|D|$.

Proof: $p_F(v) \le 2^{i+1} \le 2W$ for every $v \in D$. Thus $p' \le 2W|D|$. □

Claim: $p'' \le 2b(V)\lg W/T$.

Proof: If at iteration i we added to F edges incident to v, then the deficiency of v drops by at least $2^i T$. Thus the total number of nodes in $V - D$ incident to edges added at iteration i is at most $b(V)/(2^i T)$. Since every added edge has cost at most 2^{i+1} the total increase in the power at iteration i is at most $2^{i+1}b(V)/(2^i T) = 2b(V)/T$. The claim follows. □

Assume that $\tau \geq opt$. Let O be a feasible solution with $p(O) \leq \tau$. Let $A = \{e \in O - F : c(e) \leq 2W\}$, $B = (O - F) - A$. The following two claims show that $O - F$ decreases the deficiency of $D(F)$ by at most $\tau\,(T \lg W + r\mu/2W)$. This implies (4), since $b_F(V) = b_F(D(F))$.

Claim: A decreases the deficiency of $D(F)$ by at most $\tau T \lg W$.

Proof: Fix some $i \leq \lg W$. Let $A_i = E_i \cap A$. Since $p_{A_i}(V) \leq \tau$, the edges in A_i are incident to at most $\tau/2^i$ nodes. Note that $|\delta_{A_i}(v, D(F))| \leq T2^i$ for every $v \in V$. Thus each A_i reduces the deficiency of $D(F)$ by at most τT. The claim follows. □

Claim: B decreases the deficiency of $D(F)$ by at most $\mu r\tau/(2W)$.

Proof: The number of nodes in $D(F)$ adjacent to the edges in B is at most $\tau/(2W)$. The deficiency of each $v \in D(F)$ is at most $r\mu$. The claim follows. □

The proof of Lemma 3 is complete. □

Corollary 1. *There exists a polynomial time algorithm that given an instance of* MPb-EC *with* $\max\{b_v : v \in D\} \leq r\mu$ *and an integer* τ*, returns an edge set* F *such that:* $p(F) = \tau \cdot O(r + \lg^2 n)$ *and if* $\tau \geq opt$ *then* $b_F(V) \leq b(V)/2$.

Proof. For $W = 2\tau\mu r/b(V)$ and $T = b(V)/(4\tau \lg W)$, the algorithm from Lemma 3 computes an edge set F such that (note that $W = 2\tau \cdot (\mu r)/b(V) \leq 2\tau \leq 2n^6$):

$$p(F) \leq 2\left(2\tau r\frac{\mu|D|}{b(V)} + 4\tau \lg^2 W\right) \leq 4\tau\left(r + 2\lg^2 W\right) = \tau \cdot O(r + \lg^2 n).$$

If $\tau \geq opt$ then:

$$b_F(V) \leq \tau\left(\frac{b(V)/4}{\tau} + \frac{\mu r b(V)}{4\tau\mu r}\right) = b(V)\left(\frac{1}{4} + \frac{1}{4}\right) = \frac{b(V)}{2}.$$

□

Proof of Lemma 2: Consider the following algorithm that starts with $E' = \emptyset$:

Algorithm b-Edge-Cover(τ)
While $b(V) \geq \lg^3 n$ do:

- Let $V_0 = \{v \in V : 1 \leq b(v) \leq 2\}$ and $V_j = \{v \in V : 2^j + 1 \leq b(v) \leq 2^{j+1}\}$, $j = 1, \ldots, \lg n$.
- Let q be an index so that $b(V_q) \geq b(V)/\lg n$.
- Compute F as in Corollary 1 with $b'(v) = b(v)$ if $v \in V_q$ and $b'(v) = 0$ otherwise.
- *If* $b_F(V_q) \leq b(V_q)/2$ then: $E' \leftarrow E' \cup F$, $G \leftarrow G - F$, $b \leftarrow b_F$; *Else* declare "$\tau <$ opt" and STOP.

End While

If the algorithm declares "$\tau < \text{opt}$" then this is correct, by Corollary 1. Let us assume therefore that this is not so.

Claim: The algorithm calls to the algorithm from Corollary 1 $O(\lg^2 n)$ times.

Proof: Let B_t be the total residual deficiency before iteration $t+1$ of the while loop, where $B_0 = b(V) \leq n^2$. We have $B_{t+1} \leq B_t(1 - 1/(2\lg n))$, so $B_t \leq B_0(1 - 1/(2\lg n))^t$. Thus after at most

$$\frac{\lg(n^2/\lg^3 n)}{-\lg(1 - 1/(2\lg n))} = O(\lg^2 n)$$

iterations the condition in the while loop is met, and the iterations stop. □

From the last claim and Corollary 1, we obtain that $p_{E'}(V) = \tau \cdot O(\lg^4 n)$ (note that when **Algorithm *b*-Edge-Cover(τ)** calls Corollary 1, we can set $r = 2$ since $b(v)$'s of all $v \in V_q$ are within a factor 2 of each other). □

The proof of Theorem 1 is now complete.

3 Proof of Lemma 1 and Theorems 2–4

We first present the proof of Lemma 1 which is a basis to our results.

Proof of Lemma 1: Except the inequality $c(G)/\sqrt{|E|/2} \leq p(G)$ the statement was proved in [11, 19]. We restate the proof for completeness of exposition. The inequality $p(G) \leq 2c(G)$ follows from

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

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

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

$$\sum_{xy \in E} \min\{p(x), p(y)\} \leq \sqrt{|E|/2} \sum_{v \in V} p(v) \tag{5}$$

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 (5) 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 $|E_U| \leq \sqrt{|E|/2}|U|$ 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

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

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 (5). This contradicts our choice of G, p. □

We now prove Theorems 2 and 3. We need the following fundamental statement due to Mader.

Theorem 7 ([20]). *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 [20]) is used to get a relation between $(k-1)$-edge covers and k-connected spanning subgraphs.

Corollary 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. □

Proof of Theorem 3: By the assumption, we can find a subgraph J with $\deg_J(v) \geq k-1$ of power at most $p(J) \leq \beta\mathrm{opt}$. 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 2, F is a forest. Thus $p(F) \leq 2c(F) \leq 2\alpha\mathrm{opt}$, by Lemma 1. 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 $\mathsf{MP}k\text{-}\mathsf{CSS}$ with the new cost function. The edges with nonzero cost in this new graph form a forest F, by Corollary 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 1. This gives a $(2\rho + 1)$-approximation algorithm for the Min-Cost k-Connected Spanning Subgraph problem. □

We can combine the various existing approximation algorithms for the Min-Cost k-Connected Spanning Subgraph problem [7, 16, 17] to get better approximation for $\mathsf{MP}k\text{-}\mathsf{CSS}$. The currently best approximation ratios for the former are $O(\ln k \cdot \min\{\sqrt{k}, \frac{n+k}{n-k} \ln k\})$ [17] and $O(\ln k)$ for $n \geq 6k^2$ [7].

In particular, we set $\beta = k$ in Theorem 3 to get a $k(1+o(1))$-approximation for any non-constant k. Using $\alpha = O(\ln k \cdot \min\{\sqrt{k}, \frac{n+k}{n-k} \ln k\})$ from [17] gives the bound in Theorem 2.

Remark: In [16] a $(2 + k/n)$-approximation was given for k-CS with metric costs. This does not imply that for metric costs we can set $\alpha = 2 + k/n$ in Theorem 3. Note that our algorithm first resets the costs of the edges in a k-edge cover to zero, and thus when applying an algorithm for min-cost k-CS the triangle inequality property *does not* hold for the obtained k-CS instance.

To prove Theorem 4, we combine Lemma 1 with the following theorem due to Cheriyan and Thurimella [6], which is the edge-connectivity counterpart of Corollary 2.

Theorem 8 ([6]). *If $deg_J(v) \geq k$ for every node v of a graph J, and if F is an inclusion minimal edge set such that $G \cup F$ is k-edge connected, then $|F| \leq n-1$.*

Proof of Theorem 4: We use the $O(\log^4)$-approximation for MPb-EC. Then we change the cost of these edges to zero and find the minimum cost k-edge connected subgraph using the known 2-approximation algorithms for the minimum cost k-edge connected subgraph problem [13]. From Lemma 1 and Theorem 8, the power of this augmentation is at most $2\sqrt{n/2}$ of the minimum power k-edge connected subgraph. This gives an $O(\sqrt{n})$-approximation algorithm. □

4 Proof of Theorem 5

To prove Theorem 5, we will show that approximating MPk-EDP is at least as hard as approximating the following problem, which is an alternative formulation, of the LabelCover-Max Problem defined in [15].

The MaxRep Problem:
Instance: A bipartite graph $H = (A \cup B, I)$, and equitable partitions $\mathcal{A}$ of A and $\mathcal{B}$ of B into q sets of same size each.
Objective: Choose $A' \subseteq A$ and $B' \subseteq B$ with $|A' \cap A_i| = |B' \cap B_j| = 1$ for each $i, j = 1, \ldots, q$ such that the subgraph induced by $A' \cup B'$ has maximum number of edges.

The bipartite graph and the partition of A and B induce a super-graph Γ in the following way: The vertices in Γ are the sets A_i and B_j. Two sets A_i and B_j are connected by a (super) edge in Γ if and only if there exist $a_i \in A_i$ and $b_j \in B_j$ which are adjacent in G. For our purposes, it is convenient (and possible) to assume that the graph Γ is regular. Say that every vertex in Γ has degree d, and hence, the number of super-edges is $h = qd$. Raz [23] proved:

Theorem 9. [23] *Let I be an instance of any NP-complete problem. For any $0 < \epsilon < 1$, there exists a (quasi-polynomial) reduction that maps I to an instance G of* MaxRep *with n vertices so that: 1) If I corresponds to a yes instance then there exists a feasible solution covering all super-edges, and 2) If I corresponds to a no instance, then every* MaxRep *feasible solution covers at most $h/2^{\log^{1-\epsilon} n}$ super-edges.*

In the above reduction the size n of the MaxRep instance G is quasi-polynomial in the size of the NP-complete instance. The following is implied from Theorem 9.

Theorem 10. *Unless $NP \subseteq DTIME(n^{\mathrm{polylog} n})$, the* MaxRep *Problem admits no $2^{\log^{1-\epsilon} n}$-approximation algorithm, for any constant $\epsilon > 0$.*

The Reduction. We reduce MaxRep to MPk-EDP. Let H be the bipartite instance of MaxRep. Form an instance G for MPk-EDP as follows. First we put H into G and give all the edges of H directions from the A_i vertices to the B_j vertices. The edges of H are assigned cost n^3. Add a source s and a sink t. For each set A_i (B_i), $1 \leq i \leq q$, we also add a local source s_i (a local sink t_i). We add d edge-disjoint paths of length 2 from s_i, $1 \leq i \leq q$, into every $a_i^j \in A_i$, $1 \leq j \leq N$, and d edge-disjoint paths of length 2 from s to every s_i. These edges are given cost 0. Finally, we add d edge-disjoint paths of length 2 from every $b_i^j \in B_i$, $1 \leq i \leq q, 1 \leq j \leq N$ into t_i, and d edge-disjoint paths of length 2 from t_i into t. The first edge in every one of paths from a vertex in B_i to t_i gets cost n^3 while the rest of edges get cost 0.

A direct inspection shows that there exists $h = dq$ edge-disjoint paths from s to t and indeed we pick $k = h$ for the MPk-EDP instance.

Let H be a MaxRep instance resulting from a yes instance of the NP-complete instance and let G be the resulting MPk-EDP instance.

Lemma 4. *The graph G admits a subgraph G' of power-cost $2qn^3$ so that in G' there exist $k = h$ edge-disjoint paths from s to t.*

Proof. We select the following edges as a solution F to MPk-EDP. Let $a_i \in A_i$, $b_j \in B_j$ be a MaxRep solution covering all the superedges as guaranteed in Theorem 9. Add all the a_i to b_j edges into the solution. Note that the edge (a_i, b_j) exists as the chosen representatives cover all the super-edges. Include all the edges which are on a path from s to a_i, $1 \leq i \leq q$, and all edges which are on a path from b_j, $1 \leq j \leq q$, to t. Clearly the solution F admits h edge-disjoint $s-t$ paths. The solution pays n^3 per every a_i because of the A_i to B_j edges and n^3 per every b_j because of the d paths to t_j. □

Lemma 5. *If G corresponds to a no instance of* MaxRep *then the cost of any* MPk-EDP *solution is at least $0.4qn^3 2^{\frac{\log^{1-\epsilon} n}{4}}$.*

Proof. The idea of the proof is to start with a solution for MPk-EDP and use it to build a MaxRep solution that covers a number of superedges which is related to the cost of this solution. Let F be the solution to MPk-EDP. Call a vertex v *active* (with respect to F) if at least one edge in F touches v. Let A_i' (respectively, B_j') be the collection of active vertices in A_i (respectively, B_j).

We may clearly assume that the outdegree of A_i' and B_j' vertices is nonzero (vertices that do not obey this can be discarded). The power-cost is thus at least $(\sum_i |A_i'| + \sum_j |B'|_j)n^3$.

Let $(\sum_i |A_i'| + \sum_j |B'|_j) = 2q\rho$. The average size of A_i' (respectively, $|B_j'|$) is at most 2ρ. Call an A_i *sparse* if $|A_i'| > 8\rho$. Similarly, B_j is *sparse* if $|B_j'| > 8\rho$.

Remove from the super-graph Γ all the sparse sets A_i and B_j. Clearly, the number of sparse A_i sets is no larger than $q/4$ and the same holds for B_j. Now we update the number of $s-t$ paths discarding paths of sparse sets. The loss of paths incurred by the removal of a sparse A_i or sparse B_j is at most d. Hence, the removal of sparse A_i and B_j sets incurs a loss of at most $2q/4d = h/2$ paths. Hence, at least $h/2$ $s-t$ edge-disjoint paths still exist after this update.

We now dilute the path collections so that at most one path remains between every pair of sets A_i, B_j. Since the remaining sets A_i and B_j are not sparse, the number of active vertices in each set is bounded by 8ρ. Hence, the total number of paths between every pair of sets A_i and B_j is at most $(8\rho)^2$. Therefore, the dilution results in a total number of paths of at least $\frac{h}{128\rho^2}$. Let F' be the subset of $\bigcup A'_i \cup \bigcup B'_j$ restricted to the non-sparse A_i, B_j.

We now create a feasible MaxRep solution by drawing a single vertex in every non-sparse A'_i and B'_i with all elements being equally likely to be chosen. Let F'' be the resulting set of unique representatives; Clearly F'' is a feasible MaxRep solution. Observe that a super-edge covered by F' has probability exactly $1/64\rho^2$ to be covered by F''. The expected number of superedges covered by F'' is $\frac{h}{8192\rho^4}$. This implies the existence of a MaxRep solution that covers this many superedges. By Theorem 9, $8192\rho^4 \geq 2^{\log^{1-\epsilon} n}$. Finally, we note that the probabilistic construction of F'' can be easily de-randomized using the method of conditional expectation and thus the claim follows. □

By Lemma 4 and 5, it is hard to approximate MPk-EDP within $2^{\frac{\log^{1-\epsilon} n}{4}}/10$. Since ϵ can be chosen to be any arbitrary constant, the hardness result follows.

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), 2003, pp. 1889–1894.
2. M. Bahramgiri, M. Hajiaghayi, and V. Mirrokni, *Fault-tolerant and 3-dimensional distributed topology control algorithms wireless multi-hop networks*, in Proceedings of the 11th IEEE International Conference on Computer Communications and Networks (ICCCN), IEEE Press, 2002, pp. 392–398.
3. D. M. Blough, M. Leoncini, G. Resta, and P. Santi, *On the symmetric range assignment problem in wireless ad hoc networks*, in TCS '02: Proceedings of the IFIP 17th World Computer Congress - TC1 Stream / 2nd IFIP International Conference on Theoretical Computer Science, Kluwer, B.V., 2002, pp. 71–82.
4. G. Calinescu, 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 Conference on Theoretical Computer Science, Kluwer, B.V., 2002, pp. 119–130.
5. G. Calinescu and P.-J. Wan, *Range assignment for high connectivity in wireless ad hoc network*, in Adhoc-Now, 2003, pp. 235–246.
6. J. Cheriyan and R. Thurimella, *Fast algorithms for k-shredders and k-node connectivity augmentation*, J. Algorithms, 33 (1999), pp. 15–50.

7. J. Cheriyan, S. Vempala, and A. Vetta, *An approximation algorithm for the minimum-cost k-vertex connected subgraph*, SIAM J. Comput., 32 (2003), pp. 1050–1055 (electronic).
8. A. Clementi, G. Huiban, P. Penna, G. Rossi, and Y. Verhoeven, *Some recent theoretical advances and open questions on energy consumption in ad-hoc wireless networks*, 2002, pp. 23–38.
9. A. E. F. Clementi, P. Penna, and R. Silvestri, *Hardness results for the power range assignmet problem in packet radio networks*, in RANDOM-APPROX '99: Proceedings of the Third International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, Springer-Verlag, 1999, pp. 197–208.
10. A. Frank and É. Tardos, *An application of submodular flows*, Linear Algebra Appl., 114/115 (1989), pp. 329–348.
11. M. Hajiaghayi, N. Immorlica, and V. S. Mirrokni, *Power optimization in fault-tolerant topology control algorithms for wireless multi-hop networks*, in MobiCom '03: Proceedings of the 9th annual international conference on Mobile computing and networking, ACM Press, 2003, pp. 300–312.
12. S. Khuller and B. Raghavachari, *Improved approximation algorithms for uniform connectivity problems*, J. Algorithms, 21 (1996), pp. 434–450.
13. S. Khuller and U. Vishkin, *Biconnectivity approximations and graph carvings*, J. Assoc. Comput. Mach., 41 (1994), pp. 214–235.
14. L. M. Kirousis, E. Kranakis, D. Krizanc, and A. Pelc, *Power consumption in packet radio networks*, Theoret. Comput. Sci., 243 (2000), pp. 289–305.
15. G. Kortsarz, *On the hardness of approximating spanners*, Algorithmica, 30 (2001), pp. 432–450.
16. G. Kortsarz and Z. Nutov, *Approximating node connectivity problems via set covers*, Algorithmica, 37 (2003), pp. 75–92.
17. G. Kortsarz and Z. Nutov, *Approximation algorithm for k-node connected subgraphs via critical graphs*, in STOC '04: Proceedings of the thirty-sixth annual ACM symposium on Theory of computing, ACM Press, 2004, pp. 138–145.
18. L. Li, J. Y. Halpern, P. Bahl, Y.-M. Wang, and R. Wattenhofer, *Analysis of a cone-based distributed topology control algorithm for wireless multi-hop networks*, in PODC '01: Proceedings of the twentieth annual ACM symposium on Principles of distributed computing, ACM Press, 2001, pp. 264–273.
19. E. L. Lloyd, R. Liu, M. V. Marathe, R. Ramanathan, and S. S. Ravi, *Algorithmic aspects of topology control problems for ad hoc networks*, Mob. Netw. Appl., 10 (2005), pp. 19–34.
20. W. Mader, *Ecken vom grad n in minimalen n-fach zusammenhangenden Graphen*, Arch. Math. (Basel), 23 (1972), pp. 219–224.
21. J. B. Orlin, *A faster strongly polynomial minimum cost flow algorithm*, Oper. Res., 41 (1993), pp. 338–350.
22. C. H. Papadimitriou and M. Yannakakis, *Optimization, approximation, and complexity classes*, J. Comput. System Sci., 43 (1991), pp. 425–440.
23. R. Raz, *A parallel repetition theorem*, SIAM J. Comput., 27 (1998), pp. 763–803 (electronic).
24. V. Rodoplu and T. H. Meng, *Minimum energy mobile wireless networks*, IEEE J. Selected Areas in Communications, 17 (1999), pp. 1333–1344.
25. R. Wattenhofer, L. Li, V. Bahl, and Y. Wang, *Distributed topology control for power efficient operation in multihop wireless ad hoc networks*, in Proceedings of twentieth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM), IEEE Press, 2001, pp. 1388–1397.