

Approximate Min-Power Strong Connectivity *

G. Calinescu †

Abstract

Given a directed simple graph $G = (V, E)$ and a cost function $c : E \rightarrow R_+$, the *power* of a vertex u in a directed spanning subgraph H is given by $p_H(u) = \max_{uv \in E(H)} c(uv)$, and corresponds to the energy consumption required for wireless node u to transmit to all nodes v with $uv \in E(H)$. The *power* of H is given by $p(H) = \sum_{u \in V} p_H(u)$.

Power Assignment seeks to minimize $p(H)$ while H satisfies some connectivity constraint. In this paper, we assume E is bidirected (for every directed edge $e \in E$, the opposite edge exists and has the same cost), while H is required to be strongly connected. This is the original power assignment problem introduced by Chen and Huang in 1989, who proved that a bidirected minimum spanning tree has approximation ratio at most 2 (this is tight). In Approx 2010, we introduced a greedy approximation algorithm and claimed a ratio of 1.992. Here we improve the algorithm's analysis to 1.85, combining techniques from Robins-Zelikovsky (2000) for Steiner Tree, and Caragiannis, Flammini, and Moscardelli (2007) for the broadcast version of Power Assignment, together with a simple idea inspired by Byrka, Grandoni, Rothvoß, and Sanità (2010).

The proof also shows that a natural linear programming relaxation, introduced by Calinescu and Qiao in Infocom 2012, has integrality gap at most 1.85.

1 Introduction

There has been a surge of research in Power Assignment problems since 2000 (among the earlier papers are [17, 21, 12]) This class of problems take as input a directed simple graph $G = (V, E)$ and a cost function $c : E \rightarrow R_+$. The *power* of a vertex u in a directed spanning simple subgraph H of G is given by $p_H(u) = \max_{uv \in E(H)} c(uv)$, and corresponds to the energy consumption required for wireless node u to transmit to all nodes v with $uv \in E(H)$. The *power* (or *total power*) of H is given by $p(H) = \sum_{u \in V} p_H(u)$.

The study of the min-power power assignment was started by Chen and Huang [8], which consider, as we do, the case when E is bidirected, (that is, $uv \in E$ if and only if $vu \in E$, and if weighted, the two edge have the same cost; this case was sometimes called “symmetric” or “undirected” in the literature) while H is required to be strongly connected. We call this problem *Min-Power Strong Connectivity*. We use with the same name both the (bi)directed and the undirected version of G . [8] prove that the bidirected version of a minimum (cost) spanning tree (MST) of the input graph G has

*These results were announced in <http://arxiv.org/abs/1205.3397> and in a preliminary version in [4]

†Department of Computer Science, Illinois Institute of Technology, Chicago, IL 60616, USA. calinescu@iit.edu. Research supported in part by NSF grants CCF-0515088 and NeTS-0916743, and the ERC StG project PAAI no. 259515.

power at most twice the optimum, and therefore the MST algorithm has approximation ratio at most 2. This is known to be tight (see Section 2).

We improve the approximation ratio to 1.85 by combining techniques from Robins-Zelikovsky [18, 19] for Steiner Tree, Caragiannis, Flammini, and Moscardelli [6] for the broadcast version of symmetric Power Assignment (assuming a bidirected $G = (V, E, c)$ and a “root” $u \in V$ is given, H must contain a directed path from u to every vertex of G), together with a simple idea inspired by the journal submission of Byrka, Grandoni, Rothvoß, and Sanità (this idea was not explicit enough in the conference version [3]).

Very restricted versions of Min-Power Strong Connectivity have been proven NP-Hard [14, 10, 7]. Other than [4], we are not aware of a approximation ratio better than 2, except for [7], (where $c : E \rightarrow \{A, B\}$, for $0 \leq A < B$; see also [4]), [2] (where c is assumed to be a metric), and the exact (dynamic programming) algorithms of [14] for the specific case where each vertex of G maps to a point on a line, and $c(uv)$ is an increasing function of the Euclidean distance between the images of u and v . A related version, also NP-Hard, asks for H to be bidirected (also called “undirected” or “symmetric” in previous papers). This problem is called Min-Power Symmetric Connectivity, and the best known ratio of $5/3 + \epsilon$ [1] is obtained with techniques first applied to Steiner Tree; when $c : E \rightarrow \{A, B\}$ one gets $3/2$ with the same method [16]. In fact, many but not all power assignment algorithms use techniques from Steiner Tree variants, or (sometimes non-trivial) reductions to Steiner Tree variants. For example, Caragiannis et al [6] uses the relative greedy heuristic of Zelikovsky [22]. New interesting techniques were also developed for power assignment problems, as in [15], an improvement over [13].

The existing lower bound of the optimum, which we use, is the cost of the minimum spanning tree of G . Indeed (argument from [8]), the optimum solution OPT contains an in-arborescence rooted at v , for some $v \in V$, and then, for all $u \in V \setminus \{v\}$, $p_{OPT}(u)$ is at least the cost of the directed edge connecting u to its parent in this in-arborescence, whose total cost is at least the cost of the minimum spanning tree of G .

We also use a relative greedy method as in [22, 18]. The techniques of Robins-Zelikovsky [18] were rarely used, even though this paper has hundreds of citations, the vast majority only citing the approximation ratio (improved by now in [3]). We use the natural structures of [6] to improve over the minimum spanning tree: these are *stars*, directed trees of height 1. Our second lower bound (improved over [4]) comes from “covering” the edges of the a spanning tree by the stars of the optimum solution. With precise definitions later, we just mention that an edge of a tree is covered by a star if it is on a path of the tree between two vertices of the star. A low-cost fractional covering can be easily obtained from either optimum or the linear programming relaxation. In our earlier work [4], we used an integer cover which was extremely hard to obtain. Using fractional covers (inspired by the submitted journal version of [3]) is the only significant difference of this version versus [4]. Also, interestingly, the submodularity of the covering function is only used implicitly.

2 Preliminaries

In directed graphs, we use *arc* to denote a directed edge. In a directed graph K , an *incoming arborescence* rooted at $x \in V(K)$ is a spanning subgraph T of K such that the underlying undirected graph of T is a tree and every vertex of T other than x has exactly one outgoing arc in T .

Given an arc xy , its *undirected version* is the undirected edge with endpoints x and y . Arcs xy and yx are *antiparallel*, and the antiparallel arcs resulting from the undirected edge uv are uv and vu ;




Figure 1: Tight example for the performance ratio of the MST algorithm. In both cases, the solution is bidirected and the undirected version of the arcs of the solution is given by solid edges. For each vertex, its power in the solution is written next to it. (a) The MST-based power assignment needs total power $2n$. (b) Optimum power assignment has total power $n(1 + \epsilon)^2 + (n - 1)\epsilon^2 + 1 \rightarrow n + 1$.

if undirected edge uv has cost, then each of the two antiparallel arcs resulting from undirected edge uv have this cost. We sometimes identify a spanning tree T with its set of edges.

An alternative definition of our problem (how it was originally posed) is: we are given a simple undirected graph $G = (V, E)$ and a cost function $c : E \rightarrow R_+$. A power assignment is a function $p : V \rightarrow R_+$, and it induces a simple directed graph $H(p)$ on vertex set V given by xy being an arc of $H(p)$ if and only if $\{x, y\} \in E$ and $p(x) \geq c(\{x, y\})$. The problem is to minimize $\sum_{u \in V} p(u)$ subject to $H(p)$ being strongly connected. To see the equivalence of the definition, given directed spanning subgraph H , define for each $u \in V$ the power assignment $p(u) = p_H(u)$.

The following known example (see Figure 1) shows that the ratio of 2 for the MST algorithm is tight. Consider $2n$ points located on a single line such that the distance between consecutive points alternates between 1 and $\epsilon < 1$, and let the cost function c be the square of the Euclidean distance. Then the minimum spanning tree MST connects consecutive neighbors and has power $p(MST) = 2n$. On the other hand, the bidirected tree T' with costly arcs connecting each other node (see Figure 1(b)) has a power equal to $p(T') = n(1 + \epsilon)^2 + (n - 1)\epsilon^2 + 1$. When $n \rightarrow \infty$ and $\epsilon \rightarrow 0$, we obtain that $p(MST)/p(T') \rightarrow 2$. On the other hand (this argument of [8] gives their ratio of 2), for any input graph, the power of the bidirected minimum spanning tree T is at most

$$p(T) = \sum_{v \in V} \max_{u|vu \in E(T)} c(vu) \leq \sum_{v \in V} \sum_{u|vu \in E(T)} c(vu) = 2c(T) \leq 2opt$$

where $opt = p(OPT)$ for an optimum solution OPT (the last inequality is the first lower bound from the introduction).

The example above may give intuition on how Power Assignment (and even more specifically, Min-Power Symmetric Connectivity - the variant when H must be bidirected mentioned in the introduction) relates to the k -restricted Steiner trees, with stars (trees of height 1) taking the place of restricted components. Another example from [1], (see Figure 2, and the following paragraph), shows how Min-Power Strong Connectivity differs from Min-Power Symmetric Connectivity.

The power of a Min-Power Strong Connectivity optimum solution can be almost half the power of a Min-Power Symmetric Connectivity optimum solution for the same instance: we present a series




Figure 2: Total power for Min-Power Strong Connectivity can be half the total power for Min-Power Symmetric Connectivity. In both cases, the solution is represented by solid segments (arrows meaning not bidirected), and next to each vertex its power is given. (a) Minimum power assignment ensuring strong connectivity has total power $n + n^2\epsilon^2 = n + n^2\frac{1}{n^2} = n + 1$. (b) Minimum power assignment ensuring symmetric connectivity has total power $(2n - 2) + (n(n - 1) + 2)\epsilon^2 = 2n - 1 - \frac{1}{n} + \frac{2}{n^2}$, the solution is bidirected and the undirected version of the arcs of the solution is given by solid segments.

of examples illustrated in Figure 2. The $n(n + 1)$ vertices are embedded in the plane in n groups of $n + 1$ points each. Each group has two “terminals” (represented as thick circles in Figure 2), and the $2n$ terminals are the corners of a regular $2n$ -gon with sides of length 1. Each group has another $n - 1$ equally spaced points (dashes in Figure 2) on the line segment between the two terminals. The cost function c is the square of the Euclidean distance. It is easy to see that a minimum power assignment ensuring strong connectivity assigns a power of 1 to one thick terminal in each group and a power of $\epsilon^2 = (1/n)^2$ to all other points in the group - the arcs going clockwise. The total power then equals $n + 1$. For symmetric connectivity it is necessary to assign a power of 1 to all but two of the thick points (a less than optimal solution uses in each of all but two groups one single vertex to connect to the two neighboring groups, but then the power of this vertex is almost $1.5^2 > 2$ for large n) and of ϵ^2 to the remaining points, which results in total power $(2n - 2) + (n(n - 1) + 2)\epsilon^2 = 2n - 1 - \frac{1}{n} + \frac{2}{n^2}$. Also, keep in mind that the minimum spanning tree solution is a symmetric solution.

3 The Approximation Algorithm

This section is dedicated to proving the main result of this paper:

Theorem 1 *There exists a polynomial time algorithm for Min-Power Strong Connectivity with approximation ratio 1.85.*

The outline of the proof of Theorem 1, is as follows:

1. In the first subsection, we present the greedy algorithm, preceded by necessary notation.
2. Then we establish (Lemma 3) that the algorithm’s output is strongly connected.
3. A second subsection gives the analysis. It starts with the new lower bound, the fractional cover of tree edges by stars.

4. Finally, Lemma 5 shows how the Greedy method of [18] (applied with new parameters) combines the two lower bounds (the one above, and the cost of the minimum spanning tree) to obtain the claimed approximation ratio.

3.1 The algorithm

Our algorithm uses a greedy approach similar to [22, 18, 6]. Let T be the undirected minimum spanning tree of G . Let $opt = p(OPT)$ for an optimum solution OPT . The fact that T has minimum cost will be not further used, except to note $opt \geq c(T)$. Let \tilde{T} be the bidirected version of T .

For $u \in V$ and $r \in \{c(uv) \mid uv \in E\}$, let $S(u, r)$ be the directed star with center u containing all the arcs uv with $c(uv) \leq r$; note that r is the power of S , also denoted by $p(S)$. For a directed star S , let $E(S)$ be its set of arcs and $V(S)$ be its set of vertices.

For a given $S(u, r)$, let $Q(u, r)$ be the set of edges e of T such that there exist $x, y \in V(S(u, r))$ with e on the path from x to y in T . Let $\hat{Q}(u, r)$ be the set of arcs e of \tilde{T} such that there exist $x \in V(S(u, r))$ with e on the directed path from u to x in \tilde{T} ; it is easy to verify that the undirected version of $\hat{Q}(u, r)$ is $Q(u, r)$.

For a collection \mathcal{A} of directed stars $S(u_i, r_i)$, define $Q(\mathcal{A}) = \bigcup_{S(u_i, r_i) \in \mathcal{A}} Q(u_i, r_i)$ and $f(\mathcal{A}) = \sum_{e \in Q(\mathcal{A})} c(e)$. We sometimes write $Q(S)$ instead of $Q(\{S\})$. As a remark (we do not use this explicitly in the proofs) The function $f(\mathcal{A})$ is known to be monotone and submodular (see an example in [20], pages 768-769). For $S = S(u, r)$, $f_{\mathcal{A}}(S) := f(\mathcal{A} \cup \{S\}) - f(\mathcal{A}) = \sum_{e \in Q(u, r) \setminus Q(\mathcal{A})} c(e) = \sum_{e \in I_{\mathcal{A}}(S)} c(e)$, where $I_{\mathcal{A}}(S)$ is defined to be those arcs of $\hat{Q}(u, r)$ for which the undirected version is not in $Q(\mathcal{A})$.

The algorithm starts with $M = \tilde{T}$ as the set of arcs, and adds directed stars to the collection \mathcal{A} (initially empty) replacing some arcs from M to reduce the sum of costs of the arcs in M plus the sum of the powers of the stars in \mathcal{A} . For intuition, we mention that this sum is our upper bound on the power of the algorithm's output. To simplify later proofs, the algorithm makes changes (adding directed stars and removing arcs from M) even if our sum stays the same. Assume below that $0/0 = 1$. To be precise:

Algorithm Greedy:

```

 $\mathcal{A} \leftarrow \emptyset, M \leftarrow \tilde{T}$ 
While ( $f(\mathcal{A}) < c(T)$ ) do
   $(u, r) \leftarrow \operatorname{argmax}_{(u', r')} f_{\mathcal{A}}(S(u', r'))/r'$ 
   $M \leftarrow M \setminus I_{\mathcal{A}}(S(u, r))$ 
   $\mathcal{A} \leftarrow \mathcal{A} \cup \{S(u, r)\}$ 
Output  $\bigcup_{S \in \mathcal{A}} E(S) \cup M$ 

```

Each of the figures 3 and 4 shows two iterations of the algorithm. For intuition, we mention that this algorithm “covers” undirected edges of the minimum spanning tree by “stars” and when implemented, it is a variant of Chvatal’s [9] greedy algorithm for Set Cover.

Fact 2 *Note that, unless $f(\mathcal{A}) = c(T)$, a star $S(u, r)$ always exists for which $f_{\mathcal{A}}(S(u, r)) > 0$ and $f_{\mathcal{A}}(S(u, r))/r \geq 1$. Indeed, as long as a pair of antiparallel arcs e' and e'' are in M , we can pick as next star $S(u, r)$ the one given by u being the tail of e' and $r = c(e')$.*

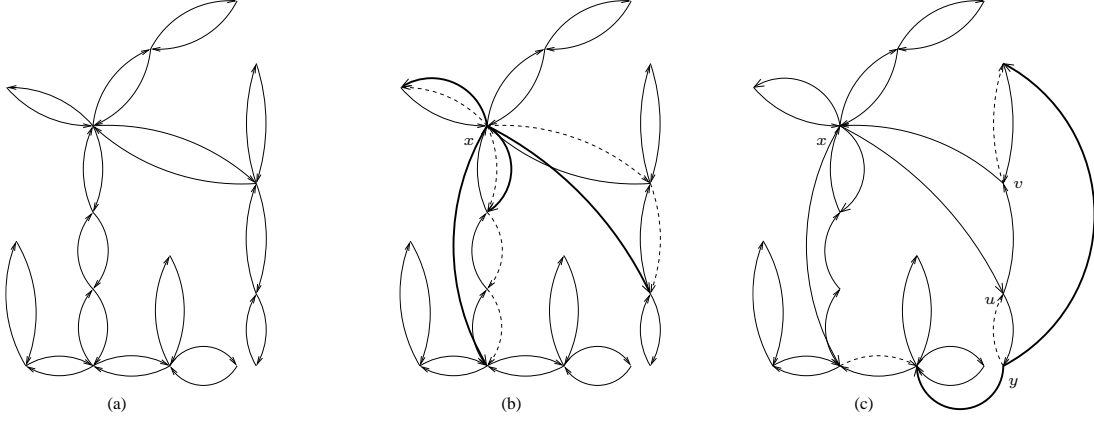


Figure 3: Costs are not relevant here. (a) Initial M . (b) Solid arcs give $\bigcup_{S \in \mathcal{A}} E(S) \cup M$ after adding to \mathcal{A} the star S_1 (with thick solid arcs) centered at x . The arcs removed from M are dashed. (c) Solid arcs give $\bigcup_{S \in \mathcal{A}} E(S) \cup M$ after adding to \mathcal{A} the star S_2 (with thick solid arcs) centered at y . The arcs removed from M are dashed. Note that the algorithm does not remove arc uv . Once an arc from a pair of antiparallel arcs of \tilde{T} is removed, the algorithm keeps the other arc, since there are cases (as in Figure 4) when not all arcs of $\hat{Q}(u, r)$ can be removed while keeping strong connectivity, and benefiting from removing arcs whose antiparallel arc has already been removed from \tilde{T} (when possible) destroys the submodularity implicitly needed in the proof of the approximation ratio.

Thus, as written, the algorithm can have iterations that do not change the output, i.e. the star $S(u, r)$ above could have just the edge e' and be added to \mathcal{A} while e' is removed from M .

Lemma 3 *The output of Greedy is a spanning strongly connected subgraph of G .*

Proof. We prove the following invariant: $X := \bigcup_{S \in \mathcal{A}} E(S) \cup M$ gives a spanning strongly connected subgraph whenever the **while** condition is checked by the algorithm. Moreover, suppose we remove from T all edges for which both antiparallel arcs appear in M , splitting T in components with vertex sets T_i , for some range of i . We prove that for every i and every $x, y \in T_i$, there exists a directed path P from x to y using only vertices of T_i and arcs from X .

Proving that the invariant holds is done as always by induction on the number of iterations. The invariant is true before the first iteration, when each T_i has just one vertex, so consider the moment when a star $S = S(u, r)$ is added to \mathcal{A} . Figure 5 may provide intuition. We add the arcs uz , for $z \in V(S) \setminus \{u\}$, while removing from M (and from X) the arcs xy for which $yx \in M$ and there is some z such that xy is on the directed simple path from u to z in \tilde{T} . The same effect is obtained if we do this change for each $z \in V(S) \setminus \{u\}$ one after the other, instead of all such z at the same time.

Let P be the simple path in T from u to z , and let $x_i y_i$, for $1 \leq i \leq k$, be, in order, the arcs of M on P such that also $y_i x_i \in M$. Thus the change to X consists of adding the arc uz and removing all the arcs $x_i y_i$; note that if $k = 0$ no arc is removed and our induction step is complete. Let $M' = M \setminus \{x_1 y_1, \dots, x_k y_k\}$ and $X' = X \setminus \{x_1 y_1, \dots, x_k y_k\} \cup \{uz\}$. We need to show that X' and M' satisfy the conditions from the induction hypothesis.

Let us split T into components by removing all the undirected edges xy with both antiparallel arcs xy and yx in M (in particular all the undirected edges $x_i y_i$, for $1 \leq i \leq k$), resulting in the components T_i .

By induction, X contains the following directed paths: P_1 from x_1 to u , P_2 from x_2 to y_1 , \dots ,

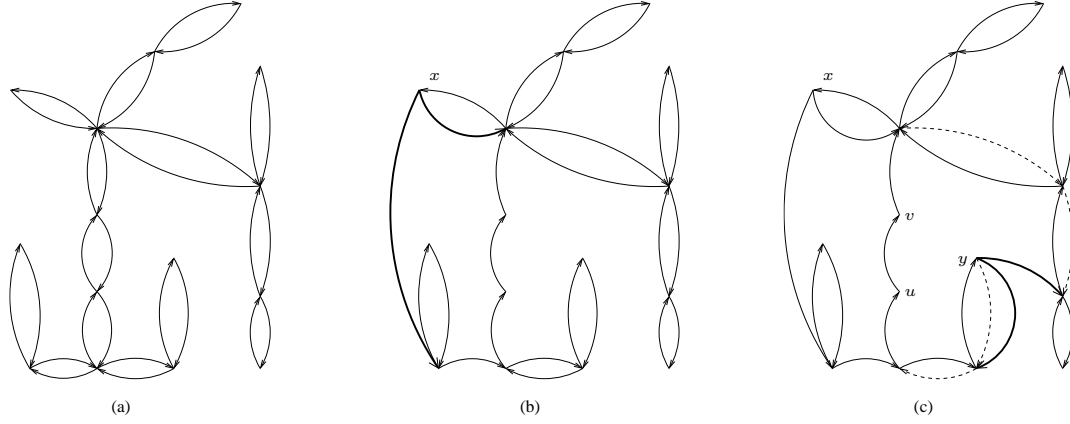


Figure 4: Costs are not relevant here. (a) Initial M . (b) $\bigcup_{S \in \mathcal{A}} E(S) \cup M$ after adding to \mathcal{A} the star S_1 (with thick solid arcs) centered at x , and removing arcs from M . (c) Solid arcs give $\bigcup_{S \in \mathcal{A}} E(S) \cup M$ after adding to \mathcal{A} the star S_2 (with thick solid arcs) centered at y . The arcs removed from M are dashed. Note that the algorithm cannot remove arc uv even though $uv \in \hat{Q}(S_2)$, since u and v become disconnected.

P_k from x_k to y_{k-1} , and P_{k+1} from z to y_k , and each of these paths stays in the same component T_i of the split (done in the previous paragraph) of T by M . Thus none of these paths uses $x_i y_i$ or $y_i x_i$. Putting together these paths, the arcs $y_i x_i$, for $1 \leq i \leq k$, and the arc uz , we have a directed cycle C containing none of the arcs $x_i y_i$, for $1 \leq i \leq k$. Any arc removed can be replaced, when discussing connectivity, with a path around the cycle C , and so (V, X') is strongly connected, as required.

We now split T into components by removing all the undirected edges xy with both xy and yx in M' , obtaining components T'_i . Note that none of P_i , $1 \leq i \leq k+1$, from above, has an arc with endpoints in two distinct components T'_i (as T'_i is the union of several T_j). As all the edges on the path from u to z in T do not have anymore both antiparallel arcs in M' , all the vertices on this path including u, z and all x_i, y_i are in the same component T'_j of the split of T by M' . Thus all the arcs of C have their endpoints in the same component of the split of T by M' .

We prove that for every i and every $x, y \in T'_i$, there exists a directed path P' from x to y using only vertices of T'_i and arcs from X' . First, let us describe a path P from x to y using only arcs of X : find the path from x to y in T , and let $z_i w_i$, for $1 \leq i \leq q$, be, in order, the arcs of M on P such that also $z_i w_i \in M$. When $q = 0$, $x, y \in T_j$ for some j (same component of the split of T by M) and by induction, a path P from x to y exists in X using only vertices inside T_j . We pick $P' = P$, and indeed P' only uses arcs of X' since the arcs of $X \setminus X'$ (same set as $M \setminus M'$) cross from one component to another of the split of T by M . Assume now $q > 0$. Notice that all unordered pairs $z_i w_i$ belong in the set of unordered pairs $x_j y_j$ on the simple path from u to v mentioned earlier, or else we cannot have that x and y belong to the same T'_i of the split of T by M' . Also, by induction, X contains the following directed paths: P_1 from x to z_1 , P_2 from w_1 to z_2 , \dots , P_q from w_{q-1} to z_q , and P_{q+1} from w_q to y , and each of these paths stays in the same component of the split of T by M . Thus none of these paths uses $z_i w_i$ or $w_i z_i$. Next, obtain P' by replacing in P , if necessary, arcs of $X \setminus X'$ (same set as $M \setminus M'$) by arcs of C , staying, as shown in the previous paragraph, in the same component of T split by M' . Note that P' indeed uses only vertices of T'_i . This completes the induction step. ■




Figure 5: Rounded rectangles show the components T_i , dashed before S is added to \mathcal{A} (the split of T by M) and solid afterward (the split of T by M'). Arcs of M crossing from one component to another are given by thin arcs, S by the four thick arcs uz_1, uz_2, uz_3, uz_4 , and the dashed arcs are those removed from M when S is added.

3.2 Approximation ratio analysis

For a collection \mathcal{A} of directed stars $S(u_i, r_i)$, define $p(\mathcal{A}) = \sum_{S(u_i, r_i) \in \mathcal{A}} r_i$, the total power used by the stars in \mathcal{A} .

Lemma 4 *Let \mathcal{B} be an arbitrary collection of stars, and T be an arbitrary spanning tree. There exist non-negative coefficients x_S (over the collection of all possible stars $S(u, r)$) such that $\sum_S x_S f_{\mathcal{B}}(S) \geq c(T) - f(\mathcal{B})$ and $\sum_S x_S p(S) \leq (1/2) \text{opt}$.*

Proof. We assume that $c(T) - f(\mathcal{B}) > 0$, or else $x_S = 0$ for all S will do. Assign $x_S = (1/2)$ for every star of OPT , and $x_S = 0$ otherwise. Therefore $\sum_S x_S p(S) = (1/2) \text{opt}$.

Recall that $Q(\mathcal{B}) = \bigcup_{S(u_i, r_i) \in \mathcal{B}} Q(u_i, r_i)$ and let $e \in T \setminus Q(\mathcal{B})$. If we remove e from T , we create two subtrees T_u and T_v , where u and v are the endpoints of e . OPT , being strongly connected, has at least one star S_u with the center in $V(T_u)$ and one of its other vertices in $V(T_v)$. Then $e \in Q(\mathcal{B} \cup \{S_u\})$. Similarly, OPT has at least one star S_v with the center in $V(T_v)$ and one of its other vertices in $V(T_u)$. Then $e \in Q(\mathcal{B} \cup \{S_v\})$. Note that $S_v \neq S_u$ (the two centers are in two disjoint vertex sets).

We have:

$$\begin{aligned}
\sum_S x_S f_{\mathcal{B}}(S) &= \sum_S x_S \sum_{e \in Q(\mathcal{B} \cup \{S\}) \setminus Q(\mathcal{B})} c(e) \\
&= \sum_{e \in T} c(e) \sum_{S \mid e \in Q(\mathcal{B} \cup \{S\}) \setminus Q(\mathcal{B})} x_S \\
&= \sum_{e \in T \setminus Q(\mathcal{B})} c(e) \sum_{S \mid e \in Q(S)} x_S \\
&\geq \sum_{e \in T \setminus Q(\mathcal{B})} c(e) = c(T) - f(\mathcal{B}),
\end{aligned}$$

where the inequation is given by the two distinct stars S_u and S_v described above for every edge $e \in T \setminus Q(\mathcal{B})$. This completes the proof. ■

The example from Figure 1 shows that a constant better (smaller) than $1/2$ is not possible in the statement of Lemma 4: with \mathcal{B} being empty, we have that $c(T)$ and opt are circa n , and for each star S of optimum, $p(S)$ is circa 1 and $f(S)$ circa 2.

Now we need the following lemma, whose proof is obtained from Robins-Zelikovsky as presented in [11] by changing what quantities represent and some parameters, together with a “fractional cover” idea from the submitted journal version of [3].

Lemma 5 *Assuming that for the minimum spanning tree T , constant $0 < \alpha < 1$, and for any collection of stars \mathcal{B} , there exist non-negative coefficients (x_S) such that $\sum_S x_S f_{\mathcal{B}}(S) \geq c(T) - f(\mathcal{B})$ and $\sum_S x_S p(S) \leq \alpha opt$, where opt is the power of the optimum solution, the Greedy algorithms’ output has power at most βopt where $\beta = 1 + \alpha + \alpha \ln(1/\alpha)$.*

Proof. First, if $c(T) \leq \alpha opt$, then before any improvement we have a solution of cost at most $2\alpha opt$ and $2\alpha < \beta$. Thus in the following we assume $opt \geq c(T) > \alpha opt > 0$ (the first inequality is due to T being a minimum spanning tree).

Note that at the end of the algorithm, M contains exactly one of the two antiparallel arcs for each edge of T . Then, for the final collection of stars \mathcal{A} , the output H satisfies

$$p(H) \leq c(T) + p(\mathcal{A}) \tag{1}$$

as it follows by summation over $u \in V$ from

$$p_H(u) = \max_{uv \in H} c(uv) \leq \sum_{uv \in M} c(uv) + \sum_{S \in \mathcal{A}} p(S),$$

which holds for every vertex u (recall that $E(H) = \bigcup_{S \in \mathcal{A}} E(S) \cup M$).

Let S_1, S_2, \dots, S_q be the stars picked by our algorithm and let \mathcal{A}_i , for $1 \leq i \leq q$, be the collection of the first i stars; also let for convenience \mathcal{A}_0 be the empty collection. For $1 \leq i \leq q$, let $p_i = p(S_i)$, and let $f_i = f_{\mathcal{A}_{i-1}}(S_i)$. Note then that for all i , since $f_{\mathcal{A}_{i-1}}(S_i) = f(\mathcal{A}_i) - f(\mathcal{A}_{i-1})$, we have $f(\mathcal{A}_i) = \sum_{j=1}^i f_j$.

If $f_i = 0$, then $p_i = 0$ and Equation 2 below holds. Otherwise, the greedy choice of the algorithm and the assumptions of the theorem for $\mathcal{B} = \mathcal{A}_{i-1}$ give:

$$p_i \leq f_i \frac{\alpha opt}{c(T) - \sum_{j=1}^{i-1} f_j}. \tag{2}$$

Define the function $g : [0..c(T)] \rightarrow [0..1]$ by $g(x) = \alpha opt / (c(T) - x)$ for $x \leq c(T) - \alpha opt$, and $g(x) = 1$ for $x > c(T) - \alpha opt$. Then from Equation 2 and Fact 2 (that $p_i \leq f_i$), we obtain:

$$p_i \leq \int_{\sum_{j=1}^{i-1} f_j}^{\sum_{j=1}^i f_j} g(x) dx.$$

Therefore (see Figure 6):




Figure 6: The function $g(x)$ is given by the solid curve. $\sum_{i=1}^q p_i$ is the shaded area, in rectangles of width f_i and height $\frac{\alpha opt}{c(T) - \sum_{j=1}^{i-1} f_j}$. In this particular picture, $\alpha opt = (2/3)c(T)$ and therefore the integral is circa $0.94c(T)$.

$$\begin{aligned}
\sum_{i=1}^q p_i &\leq \int_0^{c(T)} g(x) dx = \int_0^{c(T) - \alpha opt} \frac{\alpha opt}{c(T) - x} dx + \int_{c(T) - \alpha opt}^{c(T)} 1 dx \\
&= (-\alpha opt) \ln(c(T) - x) \Big|_0^{c(T) - \alpha opt} + (c(T) - (c(T) - \alpha opt)) \\
&= (-\alpha opt) (\ln(c(T) - (c(T) - \alpha opt)) - \ln c(T)) + \alpha opt = \alpha opt \left(1 + \ln \frac{c(T)}{\alpha opt} \right).
\end{aligned}$$

Using this and $c(T) \leq opt$ and Equation 1 (recall that $p(\mathcal{A}) = \sum_{i=1}^q p_i$), we obtain that the power of the output is at most

$$c(T) + \alpha opt \left(1 + \ln \frac{c(T)}{\alpha opt} \right) \leq opt (1 + \alpha + \alpha \ln(1/\alpha))$$

finishing the proof. ■

Based on Lemmas 3 and 4, Theorem 1 follows immediately from the fact that $\alpha = 1/2$ makes $\beta \leq 1.85$. Note also that $\alpha < 1$ implies $\beta < 2$, which follows from $\alpha(1 + \ln(1/\alpha)) < 1$, which is equivalent to $\ln(1/\alpha) < 1/\alpha - 1$, a fact that holds for all $\alpha < 1$. Our earlier work [4] had $\alpha = 7/8$.

4 Linear Programming Relaxation

While not improving the approximation ratio of Greedy, this may allow for further LP-based algorithm. The following natural Integer Linear Program is called *IP2* as in [5]. We adapted the notation, and have variables y_S for every star $S = S(u, r)$. The idea is that y_S being 1 represents that S is a star of the optimum solution. We say that a star $S = S(u, r) \in \delta^-(X)$, for $X \subset V, \emptyset \neq X \neq V$ iff $u \notin X$ and $V(S) \cap X \neq \emptyset$.

$$\text{minimize } \sum_S y_S p(S) \text{ subject to}$$

$$\sum_{S \in \delta^-(X)} y_S \geq 1 \quad \forall X \subset V, \emptyset \neq X \neq V \quad (3)$$

$$y_S \geq 0 \quad \forall S \quad (4)$$

$$y_S \in \mathbb{Z} \quad \forall S \quad (5)$$

LP2 is the linear relaxation of IP2, that is, the linear program given by exactly the same constraints except the last one. LP2 has exponentially many “cut” constraints, but using standard methods they can be replaced them by “flow” constraints, Indeed, [5] shows the details and obtains $O(mn^2)$ non-zero entries in the linear programming matrix. Thus LP2 can be solved in polynomial-time.

Let opt^* be the optimum of the linear program LP2 for a given instance. Then clearly $opt^* \leq opt$ and [5] proves that $opt \leq 2opt^*$. We do better here. First, Lemma 4 has a straightforward adaptation:

Lemma 6 *Let \mathcal{B} be an arbitrary collection of stars, and T be an arbitrary spanning tree. There exist non-negative coefficients x_S (over the collection of all possible stars $S(u, r)$) such that $\sum_S x_S f_{\mathcal{B}}(S) \geq c(T) - f(\mathcal{B})$ and $\sum_S x_S p(S) \leq (1/2)opt^*$.*

Proof. We assume that $c(T) - f(\mathcal{B}) > 0$, or else $x_S = 0$ for all S will do. Let y_S (for all $S = S(u, r)$) be an optimum solution of LP2. Assign $x_S = (1/2)y_S$ for all S ; therefore $\sum_S x_S p(S) = (1/2)opt^*$.

Recall that $Q(\mathcal{B}) = \bigcup_{S(u_i, r_i) \in \mathcal{B}} Q(u_i, r_i)$ and let $e \in T \setminus Q(\mathcal{B})$. If we remove e from T , we create two subtrees T_u and T_v , where u and v are the endpoints of e . Constraints 3 give:

$$\sum_{S \in \delta^-(V(T_u))} x_S \geq 1/2$$

and

$$\sum_{S \in \delta^-(V(T_v))} x_S \geq 1/2.$$

Note that $\delta^-(V(T_u))$ and $\delta^-(V(T_v))$ are disjoint sets since a star in the first set has its center in $V(T_v)$, while a star in the second set has its center in $V(T_u)$. Moreover, any star $S \in \delta^-(V(T_u)) \cup \delta^-(V(T_v))$ has $e \in Q(S)$. Thus

$$\sum_{S \mid e \in Q(S)} x_S \geq 1. \quad (6)$$

From now on we copied from the previous proof:

$$\begin{aligned} \sum_S x_S f_{\mathcal{B}}(S) &= \sum_S x_S \sum_{e \in Q(\mathcal{B} \cup \{S\}) \setminus Q(\mathcal{B})} c(e) \\ &= \sum_{e \in T} c(e) \sum_{S \mid e \in Q(\mathcal{B} \cup \{S\}) \setminus Q(\mathcal{B})} x_S \\ &= \sum_{e \in T \setminus Q(\mathcal{B})} c(e) \sum_{S \mid e \in Q(S)} x_S \\ &\geq \sum_{e \in T \setminus Q(\mathcal{B})} c(e) = c(T) - f(\mathcal{B}), \end{aligned}$$

where the inequation is from Inequality (6). This completes the proof. ■

The example from Figure 1 again shows that a constant smaller than $1/2$ is not possible in Lemma 6: with \mathcal{B} being empty, we that for each star S with $f(S) > \epsilon^2$, $p(S) \geq (1/2)f(S)$. Thus, if $\sum_S x_S f(S) \geq c(T)$, then $\sum_S x_S p(S) \geq (1/2)(c(T) - n\epsilon^2) \geq (1/2)(1 - \gamma)opt \geq (1/2 - \gamma)opt^*$, where we used that the example can be made to have $opt(1 - \gamma) \leq (c(T) - n\epsilon^2)$, for any $\gamma > 0$, by first making n large and then ϵ very small. As γ can be made arbitrarily small, we can see that indeed a constant smaller than $1/2$ is not possible.

It is also proven in Section II of [5] (see also below) that $c(T) \leq opt^*$; then Lemma 5 goes through with opt^* instead of opt . Therefore we conclude that the output of Greedy is at most $1.85opt^*$.

We include for completeness:

Lemma 7 *Let T be a minimum spanning tree in G . Then $c(T) \leq opt^*$.*

Proof. The dual of LP2, denoted by **DP2**, has a variable α_X for every $X \subset V$ with $\emptyset \neq X \neq V$.

$$\text{maximize } \sum_X \alpha_X \text{ subject to}$$

$$\sum_{X:S \in \delta^-(X)} \alpha_X \leq p(S) \quad \forall S \tag{7}$$

$$\alpha_X \geq 0 \quad \forall X \tag{8}$$

Pick an arbitrary $r \in V$ as a root, and consider $F \subseteq E$ (E is treated as a directed set of edges from now until the end of the proof) a minimum-cost r -rooted incoming arborescence in (V, E) ; note that $c(F) = c(T)$.

From Edmonds' minimum-cost arborescence algorithm, we obtain that $c(F)$ equals the optimum of the linear program **DP3**, given below. DP3 has a variable β_X for every $X \subset V$ with $r \in X$ and $X \neq V$.

$$\text{maximize } \sum_X \beta_X$$

$$\text{subject to } \sum_{X:vu \in \delta^-(X)} \beta_X \leq c(vu) \quad \forall vu \in E \tag{9}$$

$$\beta_X \geq 0 \quad \forall X \tag{10}$$

Moreover, it is known [20] (Subsection 52.4a in page 899, and precisely the statement 52.14) that there is a DP3 optimum solution such that the family of set U with $\beta_{V \setminus U} > 0$ is laminar (in the linear programs of [20], the sets U not containing r give the variables of the DP3 equivalent). By definition, in a laminar family of sets, any two sets are disjoint, or contained one in the other. For us, this laminar property implies that, for any $v \in V \setminus \{r\}$, the family of those X with $v \notin X$ and $\beta_X > 0$ is a *chain family*, meaning that for all such X, X' , either $X \subset X'$ or $X' \subset X$. If a DP3 feasible solution is such that for all $v \in V \setminus r$, the family of sets X with $v \notin X$ and $\beta_X > 0$ is a chain family, we call β a *chain solution*.

Claim 8 *For any instance, a feasible chain DP3 solution gives rise to a feasible DP2 solution by setting $\alpha_X = \beta_X$ if $r \in X$, and $\alpha_X = 0$ otherwise.*

Proof. We only need to check that Constraints (7) are satisfied. Pick an arbitrary star $S = S(v, r)$ and consider the sets X with $\alpha_X > 0$ and $S \in \delta^-(X)$. Using the chain property for v , among these X , there is one, say Y , that is included in all. Note that $r \in Y$, and also there is an $z \in V(S)$ with $z \in Y$ ($z = r$ is possible). Then $z \in X$ for all the sets X with $\alpha_X > 0$ and $S \in \delta^-(X)$. We have

$$\begin{aligned}
\sum_{X: S \in \delta^-(X)} \alpha_X &= \sum_{X: r \in X \wedge S \in \delta^-(X)} \beta_X \\
&= \sum_{X: \{r, z\} \subseteq X \wedge S \in \delta^-(X)} \beta_X \\
&\leq \sum_{X: z \in X \wedge v \notin X} \beta_X = \sum_{X: vz \in \delta^-(X)} \beta_X \\
&\leq c(vz) \leq r,
\end{aligned}$$

with the last inequality following from $z \in V(S)$ and the previous inequality from Constraints (9). ■

Thus $\text{opt}(DP2) \geq \text{opt}(DP3)$. We conclude the proof of Lemma 7:

$$\text{opt}^* = \text{opt}(LP2) = \text{opt}(DP2) \geq \text{opt}(DP3) = c(F) = c(T),$$

with the two middle equalities given by linear programming duality and Edmonds' primal-dual algorithm for constructing arborescences. ■

We also include a series of examples showing that the integrality gap of LP2 is at least $8/7$. See Figure 7. We have $n + 1$ vertices s_0, s_1, \dots, s_n , and another $4n$ vertices: $t_1, \dots, t_n, u_1, \dots, u_n, v_1, \dots, v_n$, and z_1, \dots, z_n . We have $3n$ edges cost 0: $t_i u_i, u_i v_i$, and $v_i t_i$. Another $6n$ edges have cost 1: for all $1 \leq i \leq n$, edges $t_i s_{i-1}, u_i s_{i-1}, u_i z_i, v_i z_i, t_i s_i$, and $v_i s_i$.




Figure 7:

An example for the lower bound on the integrality gap. With $n = 2$, we use solid segments to represent the edges of cost 1, and dotted segments to represent the edges of cost 0.

An integral solution has cost $1 + 4n$, assigning power 1 to each s_i and each z_i , as well as, for each $1 \leq i \leq n$, two of the three vertices t_i, u_i, v_i (indeed, if only one of t_i, u_i, v_i has power 1, then either s_{i-1} or s_i or z_i is not reachable from t_i).

A fractional solution assigns coefficients $y_S = 1$ for all S with power 0, or for S having power 1 and center s_i or z_i , for all i . Also, if $S = S(t_i, 1)$, or $S = S(u_i, 1)$, or $S = S(v_i, 1)$, $y_S = 1/2$. One can easily check this fractional solution is feasible, and it has objective $2n + 1 + 3n(1/2) = (7/2)n + 1$. Letting $n \rightarrow \infty$, we obtain an integrality gap of at least $8/7$.

5 Conclusions

This work greatly simplifies and at the same time improves our earlier work [4]. Instead of Greedy we could have used the iterative randomized rounding of Byrka et. al [3] for Steiner Tree, with the same approximation ratio. However, we do not see further improvements coming from using their full range of techniques, since we do not see the equivalent of the concept of “loss” used explicitly by [18] and implicitly by [3].

As part of the simplification of proofs, the submodularity of the “covering” function $f(\mathcal{A})$ is not used explicitly in the proofs (it is implicitly proved and used in Lemma 4). However, it should be noted that the relative greedy method of Zelikovsky [22] relies on the more general concept of submodularity (plus the lower bounds) and may have wider applicability.

6 Acknowledgments

We would like to thank Alexander Zelikovsky for a streamlined explanation of his algorithms, and to Gianpiero Monaco and Anna Zych for a group reading of [3] at the University of Warsaw.

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] Christoph Ambühl, Andrea E. F. Clementi, Paolo Penna, Gianluca Rossi, and Riccardo Silvestri. On the approximability of the range assignment problem on radio networks in presence of selfish agents. *Theor. Comput. Sci.*, 343(1-2):27–41, 2005.
- [3] Jaroslaw Byrka, Fabrizio Grandoni, Thomas Rothvoß, and Laura Sanità. An improved LP-based approximation for Steiner tree. In *STOC '10: Proceedings of the 42nd ACM symposium on Theory of computing*, pages 583–592, New York, NY, USA, 2010. ACM.
- [4] G. Calinescu. Min-power strong connectivity. In M. Serna, K. Jansen, and J. Rolin, editors, *Proceedings of the International Workshop on Approximation Algorithms for Combinatorial Optimization*, number 6302 in Lecture Notes in Computer Science, pages 67–80. Springer, 2010.
- [5] G. Calinescu and K. Qiao. Asymmetric topology control: exact solutions and fast approximations. In *Infocom*, 2012.
- [6] Ioannis Caragiannis, Michele Flammini, and Luca Moscardelli. An Exponential Improvement on the MST Heuristic for Minimum Energy Broadcasting in Ad Hoc Wireless Networks. In *ICALP*, pages 447–458, 2007.
- [7] Paz Carmi and Matthew J. Katz. Power assignment in radio networks with two power levels. *Algorithmica*, 47(2):183–201, 2007.

- [8] W.T. Chen and N.F. Huang. The strongly connecting problem on multihop packet radio networks. *IEEE Transactions on Communications*, 37(3):293–295, 1989.
- [9] V. Chvátal. A greedy heuristic for the set-covering problem. *Mathematics of Operations Research*, 4(3):233–235, 1979.
- [10] A.E.F. Clementi, P. Penna, and R. Silvestri. On the power assignment problem in radio networks. *Electronic Colloquium on Computational Complexity (ECCC)*, (054), 2000.
- [11] C. Gröpl, S. Hougardy, T. Nierhoff, and H.J. Prömel. Approximation algorithms for the Steiner tree problem in graphs. In D.-Z. Du and X. Cheng, editors, *Steiner Trees in Industries*, pages 235–279. Kluwer Academic Publishers, Dordrecht, 2001.
- [12] Mohammad Taghi Hajiaghayi, Nicole Immorlica, and Vahab S. Mirrokni. Power optimization in fault-tolerant topology control algorithms for wireless multi-hop networks. In *MOBICOM*, pages 300–312, 2003.
- [13] Mohammad Taghi Hajiaghayi, Guy Kortsarz, Vahab S. Mirrokni, and Zeev Nutov. Power optimization for connectivity problems. *Math. Program.*, 110(1):195–208, 2007.
- [14] L. M. Kirousis, E. Kranakis, D. Krizanc, and A. Pelc. Power consumption in packet radio networks. *Theoretical Computer Science*, 243:289–305, 2000.
- [15] Guy Kortsarz, Vahab S. Mirrokni, Zeev Nutov, and Elena Tsanko. Approximating minimum-power degree and connectivity problems. In *LATIN*, pages 423–435, 2008.
- [16] Zeev Nutov and Ariel Yaroshevitch. Wireless network design via 3-decompositions. *Inf. Process. Lett.*, 109(19):1136–1140, 2009.
- [17] Ram Ramanathan and Regina Rosales-Hain. Topology control of multihop wireless networks using transmit power adjustment. In *INFOCOM (2)*, pages 404–413, 2000.
- [18] G. Robins and A. Zelikovsky. Improved Steiner tree approximation in graphs. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 770–779, 2000.
- [19] G. Robins and A. Zelikovsky. Tighter bounds for graph Steiner tree approximation. *SIAM Journal of Discrete Mathematics*, 19(1):122–134, 2005.
- [20] A. Schrijver. *Combinatorial Optimization*. Springer, 2003.
- [21] Roger Wattenhofer, Li Li, Paramvir Bahl, and Yi-Min Wang. Distributed topology control for wireless multihop ad-hoc networks. In *IEEE INFOCOM'01*, 2001.
- [22] A. Zelikovsky. Better approximation bounds for the network and Euclidean Steiner tree problems. Technical Report CS-96-06, Department of Computer Science, University of Virginia, 1996.