

A Unified Algorithm for Degree Bounded Survivable Network Design

Lap Chi Lau and Hong Zhou

Department of Computer Science and Engineering
The Chinese University of Hong Kong
{chi,hzhou}@cse.cuhk.edu.hk

Abstract. We present an approximation algorithm for the minimum bounded degree Steiner network problem that returns a Steiner network of cost at most two times the optimal and the degree on each vertex v is at most $\min\{b_v + 3r_{\max}, 2b_v + 2\}$, where r_{\max} is the maximum connectivity requirement and b_v is the given degree bound on v . This unifies, simplifies, and improves the previous results for this problem.

1 Introduction

In the *minimum bounded degree Steiner network* problem, we are given an undirected graph $G = (V, E)$, a cost c_e on each edge $e \in E$, a degree bound b_v on each vertex $v \in V$, and a connectivity requirement r_{uv} for each pair of vertices $u, v \in V$. A subgraph H of G is called a *Steiner network* if there are at least r_{uv} edge-disjoint paths in H for all $u, v \in V$. The task of the minimum bounded degree Steiner network problem is to find a Steiner network H with minimum total cost such that $d_H(v) \leq b_v$ for each $v \in V$. This is a general problem of interest to algorithm design, computer networks, graph theory, and operations research.

It is NP-hard to determine whether there is a Steiner network satisfying all the degree bounds, even if we do not consider the cost of the Steiner network, as the Hamiltonian cycle problem is a special case. Thus, researchers focus on designing bicriteria approximation algorithms for the problem that minimize both the total cost and the degree violation. We say an algorithm is an $(\alpha, f(b_v))$ -approximation algorithm for the minimum bounded degree Steiner network problem if it returns a Steiner network H of cost at most $\alpha \cdot \text{opt}$ and $d_H(v) \leq f(b_v)$ for each $v \in V$, where opt is the cost of an optimal Steiner network that satisfies all the degree bounds.

The first bicriteria approximation algorithm for this problem is a $(2, 2b_v + 3)$ -approximation algorithm by Lau, Naor, Salavatipour, and Singh [12], and it was improved to $(2, 2b_v + 2)$ by Louis and Vishnoi [15]. There are also bicriteria approximation algorithms with additive violation on the degrees in terms of the maximum connectivity requirement. Let $r_{\max} = \max_{u,v} \{r_{u,v}\}$. Lau and Singh [14] gave a $(2, b_v + 6r_{\max} + 3)$ -approximation algorithm for the problem, and a $(2, b_v + 3)$ -approximation algorithm in the special case when $r_{\max} = 1$. The special case when $r_{\max} = 1$ is known as the *minimum bounded degree Steiner forest*

problem. In this paper, we present a $(2, \min\{b_v + 3r_{\max}, 2b_v + 2\})$ -approximation algorithm for the problem.

Theorem 1. *There is a polynomial time algorithm for the minimum bounded degree Steiner network problem that returns a Steiner network H of cost at most 2opt and $\deg_H(v) \leq \min\{b_v + 3r_{\max}, 2b_v + 2\}$ for all v .*

Theorem 1 improves the $(2, b_v + 6r_{\max} + 3)$ result in [14] when $r_{\max} \geq 2$ and recovers the $(2, b_v + 3)$ result in [14] for the minimum bounded degree Steiner forest problem. Besides, it achieves the $(2, 2b_v + 2)$ result in [15] simultaneously, while previously there was no such guarantee¹. Furthermore, both our algorithm and its analysis are simpler² as we will discuss in Section 2. We believe that our result unifies what can be achieved using existing techniques. We show an example where our algorithm fails to give a $(2, b_v + 2)$ -approximation algorithm for the minimum bounded degree Steiner forest problem in Section 3.2.

1.1 Related Work

Jain [9] introduced the iterative rounding method to give a 2-approximation algorithm for the minimum Steiner network problem, improving on a line of research that applied primal-dual methods to these problems. Later, the iterative rounding method has been applied to obtain the best known approximation algorithms for network design problems for element-connectivity [4,3], vertex-connectivity [3,2], and directed edge-connectivity [7].

The iterative relaxation method was introduced in [12] to adapt Jain's method to degree bounded network design problems, which are well-studied especially in the special case of spanning trees [6,8]. Later, this method has also been applied to achieve the best known approximation algorithms for the degree bounded network design problems, including spanning trees [17,1,11], Steiner networks [12,14,15], directed edge-connectivity [12,1], element-connectivity and vertex-connectivity [10,16,5]. See [13] for a survey on this approach.

2 Technical Overview

Since this work is tightly connected to previous work, we give a high level overview to describe the previous work and highlight where the improvement comes from.

2.1 Iterative Rounding and Relaxation

All the previous results on this problem are based on the iterative rounding method introduced by Jain [9] for the minimum Steiner network problem. This

¹ For instance, it was not known how to combine the results in [15,14] to obtain a $(2, \min\{b_v + 6r_{\max} + 3, 2b_v + 2\})$ -approximation algorithm for the problem.

² In particular, the analysis of the $(2, b_v + 3)$ result is significantly simpler than that in [14].

method is based on analyzing the extreme point solutions to a linear programming relaxation for the problem. Let us first formulate the linear programming relaxation for the minimum bounded degree Steiner network problem. For a subset $S \subseteq V$, we let $\delta(S)$ be the set of edges with one endpoint in S and one endpoint in $V - S$ in the graph and let $d(S) := |\delta(S)|$. In the linear program, there is one variable x_e for each edge, where the intended value is one if this edge is used in the solution and zero if this edge is not used. For a subset of edges $E' \subseteq E$, we write $x(E') = \sum_{e \in E'} x_e$. For a subset of vertices $S \subseteq V$, we define $f(S) := \max_{u \in S, v \notin S} \{r_{uv}\}$ to be the maximum requirement crossing S . To satisfy the connectivity requirement, we should have $x(\delta(S)) \geq f(S)$ for each $S \subseteq V$. The following is a linear programming relaxation for the minimum bounded degree Steiner network problem. It has exponentially many constraints, but there is a polynomial time separation oracle to determine whether a solution is feasible or not, and thus it can be solved in polynomial time by the ellipsoid method.

$$\begin{aligned}
 \text{(LP)} \quad & \text{minimize} \quad \sum_{e \in E} c_e x_e \\
 & \text{subject to} \quad x(\delta(S)) \geq f(S) \quad \forall S \subset V \\
 & \quad \quad \quad x(\delta(v)) \leq b_v \quad \forall v \in V \\
 & \quad \quad \quad x_e \geq 0 \quad \forall e \in E
 \end{aligned}$$

When there are no degree constraints, Jain [9] proved that there exists an edge e with $x_e \geq \frac{1}{2}$ in any extreme point solution to the above linear program. We call such an edge a *heavy* edge. He used this to obtain an *iterative rounding* algorithm for the minimum Steiner network problem, by repeatedly picking a heavy edge and recomputing an optimal extreme point solution to the residual problem. When there are degree constraints, Lau et.al. [12] showed that either there is a heavy edge or there is a degree constrained vertex with at most four nonzero edges incident to it. They then introduced an extra relaxation step to remove the degree constrained vertex in the latter case, leading to a $(2, 2b_v + 3)$ -approximation algorithm for the minimum bounded degree Steiner network problem.

Roughly speaking, all the later improvements are based on proving the existence of a heavy edge with additional properties. To improve the degree violation, Louis and Vishnoi [15] proved that in any extreme point solution either there is an edge of integral value (zero or one), or a vertex v with at most $2b_v + 2$ edges incident to it, or a heavy edge *with no endpoint having a degree bound at most one*. They showed that using this iteratively would imply a $(2, 2b_v + 2)$ -approximation algorithm for the problem. Note that in the above algorithms, after we pick a heavy edge, we need to decrease the degree bound by half in order to achieve the guarantee on the degree violation, and thus they have to consider a slightly more general problem where the degree bounds are half-integral and some subtle issue arose as we will discuss later.

To obtain additive violation on the degree bounds, Lau and Singh [14] proved that in any extreme point solution either there is an edge of integral value, or a vertex v with at most four edges incident to it, or a heavy edge *between two vertices with*

degree bounds at most $6r_{\max}$, where the last condition guarantees that the degree violation is bounded when we picked edges with value at least half. For the minimum bounded degree Steiner forest problem, they proved that in any extreme point solution either there is an integral edge, or a vertex v with at most $b_v + 3$ edges incident to it, or a heavy edge with no degree constraint on its endpoints. They showed that these would lead to a $(2, b_v + 6r_{\max} + 3)$ -approximation algorithm for the Steiner network problem and a $(2, b_v + 3)$ -approximation algorithm for the Steiner forest problem. The algorithm for Steiner forest is simpler, as it just removes the degree constraint on a vertex when it has at most $b_v + 3$ edges, and does not need to update the degree constraint to a half-integral value, as it only picks edges with value at least half when both endpoints have no degree constraints.

Our algorithm is very similar to that for the Steiner forest problem in [14] (see Algorithm 1). We prove that either there is an edge of integral value, or there is a vertex with at most $\min\{b_v + 3r_{\max}, 2b_v + 2\}$ edges incident to it, or there is a heavy edge with no degree constraints on its endpoints. The resulting algorithm is quite simple, in the first case we delete an edge when $x_e = 0$ or pick an edge when $x_e = 1$, in the second case we remove the degree constraint on that vertex, and in the final case we pick such a heavy edge. Note that we only update the degree constraints when we pick an edge with $x_e = 1$, and thus we can maintain the invariant that the degree bounds are integral, and this will simplify the analysis for the $2b_v + 2$ bound.

2.2 Analysis

To analyze the extreme point solutions, an uncrossing technique is used to show that the extreme point solutions are defined by a set of constraints with a special structure. A function $f : 2^V \rightarrow \mathbb{R}$ is skew supermodular if for any $X, Y \subseteq V$ either $f(X) + f(Y) \leq f(X \cup Y) + f(X \cap Y)$ or $f(X) + f(Y) \leq f(X - Y) + f(Y - X)$. It is known that the function f defined by the connectivity requirements is a skew supermodular function. For a set $S \subseteq V$, the corresponding constraint $x(\delta(S)) \geq f(S)$ defines a vector in $\mathbb{R}^{|E|}$: the vector has a one corresponding to each edge $e \in \delta(S)$ and a zero otherwise. We call this vector the characteristic vector of $\delta(S)$ and denote it by $\chi_{\delta(S)}$. A family of sets \mathcal{L} is laminar if $X, Y \in \mathcal{L}$ implies that either $X \cap Y = \emptyset$, or $X \subseteq Y$, or $Y \subseteq X$. Using the assumption that f is skew supermodular, it follows from standard uncrossing technique that any extreme point solution of (LP) is characterized by a laminar family of tight constraints.

Lemma 1 ([12]). *Suppose that the requirement function f of (LP) is skew supermodular. Let x be an extreme point solution of (LP) such that $0 < x_e < 1$ for all edges $e \in E$. Then there exist a laminar family \mathcal{L} of sets and set $T \subseteq V$ such that x is the unique solution to*

$$\{x(\delta(S)) = f(S) \mid S \in \mathcal{L}\} \cup \{x(\delta(v)) = b_v \mid v \in T\}$$

that satisfies the following properties:

1. The characteristic vectors $\chi_{\delta(S)}$ for $S \in \mathcal{L}$ and $\chi_{\delta(v)}$ for $v \in T$ are linearly independent.
2. $|E| = |\mathcal{L}| + |T|$.

The structure of a laminar set family \mathcal{L} can be seen as a forest in which nodes correspond to sets in \mathcal{L} and there is an edge from set R to set S if R is the smallest set containing S . We call R the *parent* of S and S is a *child* of R . A set without any parent is a *root* and a set without any child is a *leaf*. The *subtree rooted at a set S* consists of S and all its descendants.

Lau et.al. [12], following Jain [9], used this forest structure in a counting argument to prove that either there exists an edge with integral value, or a vertex with degree at most four, or an edge with value at least $\frac{1}{2}$. First, each edge is assigned two tokens, one to each endpoint, for a total of $2|E|$ tokens. Assuming none of the conditions hold, i.e. $0 < x_e < \frac{1}{2}$ for each edge and every vertex is of degree at least five. Then, it can be shown that the tokens can be redistributed such that each member of \mathcal{L} and each vertex in T get at least two tokens, and there are still some extra tokens left. This implies that there are more than $2|\mathcal{L}| + 2|T|$ tokens, contradicting property 2 of Lemma 1. The redistribution is done inductively using the following lemma.

Lemma 2 ([12]). *For any subtree of \mathcal{L} rooted at node S , we can reassign tokens collected from child nodes of S and endpoints owned by S such that each vertex in $T \cap S$ gets at least two tokens, each node in the subtree gets at least two tokens, and the root S gets at least three tokens. Moreover, S gets exactly three tokens only if $\text{coreq}(S) = \frac{1}{2}$, where $\text{coreq}(S) := \sum_{e \in \delta(S)} (\frac{1}{2} - x_e)$.*

The proof in [12] is almost the same as Jain’s proof, but with the presence of degree constraints. Note that each vertex with degree constraint gets at least five tokens in the initial token assignment. Intuitively, we can think of each degree constraint as a singleton set (a leaf set in the laminar family), and thus having five tokens is more than enough for Jain’s proof to go through. And one may think that it is already enough if every degree constraint gets at least four tokens to satisfy the induction hypothesis in Lemma 2, and this would imply a $(2, 2b_v + 2)$ -approximation algorithm. Unfortunately, the subtle point here is that the degree bounds are half-integral, but for Jain’s proof to work they need to be integral. To overcome this problem, Louis and Vishnoi [15] needed to modify the algorithm and the analysis to obtain a $(2, 2b_v + 2)$ -approximation algorithm.

The new idea in the Steiner network algorithm in [14] is to only pick heavy edges when both endpoints are of low degree. In the analysis, with the presence of heavy edges that are not allowed to be picked (when some endpoint is of high degree), there could be some sets S with $d(S) = 2$ and $x(\delta(S)) = 1$, and thus the counting argument as above would not work in the base case for those sets. For the same induction hypothesis in Lemma 2 to work, a new rule is added to the initial token assignment: if (w, v) is a heavy edge with $b_w \geq 6r_{\max}$ and v is not degree constrained, then v gets two tokens from the edge (w, v) while w gets none. The counting argument would work in the base case with this new rule, but w may not receive any token for the induction step to work. The assumption

$b_w \geq 6r_{\max}$ is used to ensure that w can *get back* the tokens in the induction step. To illustrate this, consider a worst case scenario in Figure 1(a). In the figure, there is a degree constrained vertex w where all its incident edges are heavy, and we need to collect two tokens for w and four tokens for S . Each child contributes only one token but “consumes” the degree bound of w by r_{\max} . This is where the assumption that $b_w \geq 6r_{\max}$ is used to guarantee that S has at least five children, so that each can contribute at least one token for w and S (and use some additional argument to collect one more token).




Fig. 1. Worst cases for counting arguments

In this paper, we slightly change the algorithm to remove any vertex with degree at most $\min\{b_v + 3r_{\max}, 2b_v + 2\}$. We used the same initial token assignment rule as in [14] and the same induction hypothesis in Lemma 2 for the counting argument. First, using a simple argument (see Lemma 6), we show that any vertex with degree at least $2b_v + 3$ can receive at least four tokens, and this allows the induction to work and recovers the $(2, 2b_v + 2)$ result by Louis and Vishnoi. As mentioned before, this is possible because our algorithm maintains the invariant that all the degree bounds are integral.

Our improvement to $b_v + 3r_{\max}$ comes from the concept of the *integrality gap* of the degree of a set S , defined as $d(S) - x(\delta(S))$. To illustrate this, consider the same scenario in Figure 1. In the new algorithm, the vertex w with degree bound b_w but having more than $b_w + 3r_{\max}$ edges incident to it has an integrality gap of $3r_{\max}$ on its degree. The main observation is that the heavy edges from a child having only three tokens (one extra token) can only contribute $\frac{1}{2}$ to the integrality gap of w (see Lemma 7), while a child having at least four tokens could contribute r_{\max} to the integrality gap of w . This observation basically allows us to rule out children with only three tokens in the worst case scenarios of the counting argument, as it contributes one token but only consumes $\frac{1}{2}$ of the integrality gap. This allows us to assume that each child can contribute two tokens instead of one token, and this reduces the degree violation from $6r_{\max}$ to $3r_{\max}$ (see Figure 1(b)), with some additional arguments. An additional advantage of our algorithm is that we also avoid the additive term $+3$ in the

previous algorithm [14]. The proof for $r_{\max} \geq 2$ is quite short (see Section 3.1). The proof for $r_{\max} = 1$ still has some case analysis (see the full version), but is considerably shorter than that in [14], as a complicated induction hypothesis was used in [14] that caused many more case analyses.

3 Algorithm and Analysis

In the following, let W be the set of vertices with degree constraints.

Algorithm 1. Minimum Bounded Degree Steiner Network

- 1 Initialization: $H = (V, \emptyset), W \leftarrow V, f'(S) \leftarrow f(S)$ for all $S \subseteq V$.
- 2 **while** H is not a Steiner Network **do**
 - (a) Compute an optimal extreme point solution x to (LP) and remove all edges e with $x_e = 0$.
 - (b) For each vertex $v \in W$ with degree at most $\min\{b_v + 3r_{\max}, 2b_v + 2\}$, remove v from W .
 - (c) For each edge $e = (u, v)$ with $x_e = 1$, add e to H and remove e from G , and decrease b_u, b_v by one.
 - (d) For each edge $e = (u, v)$ with $x_e \geq \frac{1}{2}$ and $u, v \notin W$, add e to H and remove e from G .
 - (e) For each $S \subset V$, set $f'(S) \leftarrow f(S) - d_H(S)$.
- 3 **Return** H .

Given that f is initially a skew supermodular function, it is known that f' in any later iteration is still a skew supermodular function [9]. So, the residual LP in any iteration is still in the original form and it has a polynomial time separation oracle [9]. Assuming the algorithm always makes progress in each iteration, then we can prove Theorem 1 by a standard inductive argument as in [12]. It remains to prove the following lemma to complete the proof of Theorem 1.

Lemma 3. *Let x be an extreme point solution to (LP) and W be the set of vertices with degree constraints. Then at least one of the following is true.*

1. *There exists a vertex $v \in W$ with $d(v) \leq \min\{b_v + 3r_{\max}, 2b_v + 2\}$.*
2. *There exists an edge e with $x_e = 0$ or $x_e = 1$.*
3. *There exists an edge $e = (u, v)$ with $x_e \geq \frac{1}{2}$ and $u, v \notin W$.*

We prove Lemma 3 by contradiction. Assuming none of the three conditions holds, then we have

1. $d(v) \geq \min\{b_v + 3r_{\max} + 1, 2b_v + 3\}$ for $v \in W$,
2. $0 < x_e < 1$ for $e \in E$,
3. if $x_{uv} \geq \frac{1}{2}$ then either u or v is in W .

We will use a token counting argument to derive a contradiction with Lemma 1. Let \mathcal{L} be the laminar family and T be the set of vertices with tight degree constraints as defined in Lemma 1. In the token counting argument, we first assign

two tokens to each edge, for a total of $2|E|$ tokens. Then, using the assumptions above, we show that these tokens can be redistributed such that each member in \mathcal{L} and each vertex in T gets at least two tokens and there are some tokens left, but this contradicts with $|E| = |\mathcal{L}| + |T|$ from Lemma 1.

Initial Token Assignment Rule. Each edge receives two tokens. If $e = (u, v)$ is a heavy edge with $u \in W$ and $v \notin W$, then v gets two tokens from e and u gets no token. For every other edge e , each endpoint of e gets one token.

We will redistribute the tokens inductively using the forest structure of the laminar family \mathcal{L} . We need some definitions to state the induction hypothesis. We say a vertex v is *owned by* a set $S \in \mathcal{L}$ if S is the smallest set in \mathcal{L} that contains v . Given an extreme point solution x , we say an edge e is a *heavy edge* if $x_e \geq 1/2$, otherwise we say e is a *light edge*. Let $\delta^h(S) = \{e \in \delta(S), x_e \geq 1/2\}$ ($\delta^l(S) = \{e \in \delta(S), x_e < 1/2\}$) be the set of heavy edges (light edges) in $\delta(S)$. The *corequirement* of a set S is defined as

$$\text{coreq}(S) = \sum_{e \in \delta^h(S)} (1 - x_e) + \sum_{e \in \delta^l(S)} (1/2 - x_e) = |\delta^h(S)| + \frac{|\delta^l(S)|}{2} - x(\delta(S)).$$

We will prove the following lemma which shows that the tokens can be redistributed to obtain a contradiction, proving Lemma 3.

Lemma 4. *For any subtree of \mathcal{L} rooted at node S , we can reassign tokens collected from child nodes of S and endpoints owned by S such that each vertex in $T \cap S$ gets at least two tokens, each node in the subtree gets at least two tokens, and the root S gets at least three tokens. Moreover, S gets exactly three tokens only if $\text{coreq}(S) = \frac{1}{2}$.*

We focus on the case when S owns some vertices in T , and show that in such case S can get at least four tokens. In the induction step, we assume that the induction hypothesis holds for each child of S . We say a child of S is a *rich* child if it gets at least four tokens, and say a child of S is a *poor* child if it gets exactly three tokens. Note that a child only needs two tokens and thus has some *excess* tokens, i.e., each rich child of S has at least two excess tokens and each poor child of S has one excess token. The following lemma is the technical core of this paper.

Lemma 5. *Let $S \in \mathcal{L}$. Suppose that the induction hypothesis in Lemma 4 holds for each child of S and S owns $k \geq 1$ vertices in T . Then the number of excess tokens from the child nodes of S , plus the number of tokens collected from endpoints owned by S is at least $2k + 4$.*

Lemma 5 handles the cases when S owns some vertex in T , to guarantee that S gets at least four tokens and each vertex in T owned by S gets two tokens for the induction hypothesis. The remaining cases can be handled exactly as in the proof in [14]. Please refer to the full version for the proof of Lemma 4 assuming Lemma 5, which follows the proof in [14].

We will prove Lemma 5 in the remainder of this paper. We present the proof of Lemma 5 when $r_{\max} \geq 2$ in Section 3.1, which improves the result in [14] about Steiner networks. Due to space limit, please refer to the full version for the proof of Lemma 5 when $r_{\max} = 1$, which recovers the result in [14] about Steiner forest with a considerably simpler proof.

3.1 Proof of Lemma 5

Before we assume $r_{\max} \geq 2$, we prove two useful lemmas. The first lemma takes care of those vertices $w \in W$ with $d(w) \geq 2b_w + 3$.

Lemma 6. *If $d(w) \geq 2b_w + 3$ for $w \in W$, then w receives at least four tokens in the initial token assignment.*

Proof. Assume there are h heavy edges incident and l light edges incident to w . If $l \geq 4$, then w receives at least four tokens in the initial token assignment. Suppose to the contrary that $l \leq 3$. Then $h \geq 2b_w$ as $d(w) = h + l \geq 2b_w + 3$. If $h > 2b_w$, then $x(\delta(w)) \geq \frac{h}{2} > b_w$, contradicting that x is a feasible solution to (LP). Otherwise, if $h = 2b_w$, since each light edge has positive value, we have $x(\delta(w)) > \frac{h}{2} = b_w$, again contradicting that x is a feasible solution to (LP). \square

Lemma 6 says that any vertex $w \in W$ with $d(w) \geq 2b_w + 3$ gets at least four tokens in the initial assignment. Together with the fact that b_w is an integer, then any $w \in T$ with $d(w) \geq 2b_w + 3$ is a singleton set $\{w\}$ with $x(\delta(w))$ integral and has at least four tokens, and thus it behaves the same as a rich child in the proof of Lemma 5. Henceforth, we can assume that $b_w + 3r_{\max} + 1 < 2b_w + 3$ for each $w \in W$, and thus

$$b_w > 3r_{\max} - 2 \text{ and } d(w) \geq b_w + 3r_{\max} + 1 \text{ for } w \in W. \tag{1}$$

The second lemma takes care of the poor children using the concept of integrality gap of the degree of a set. For an edge e with $0 < x_e < 1$, let $y_e = 1 - x_e \in (0, 1)$ be the *integrality gap of e* . For a subset of edges $E' \subseteq E$, let $y(E') := \sum_{e \in E'} y_e$.

Lemma 7. *Suppose the induction hypothesis in Lemma 4 holds for each child of $S \in \mathcal{L}$. Let $R \in \mathcal{L}$ be a poor child of S . Then*

$$y(\delta^h(R)) \leq \frac{1}{2}.$$

Proof. Note that

$$\frac{1}{2} = \text{coreq}(R) = \sum_{e \in \delta^h(R)} (1 - x_e) + \sum_{e \in \delta^l(R)} \left(\frac{1}{2} - x_e\right) = \sum_{e \in \delta^h(R)} y_e + \sum_{e \in \delta^l(R)} \left(\frac{1}{2} - x_e\right).$$

Since $1/2 - x_e > 0$ for each light edge e , we have $\sum_{e \in \delta^h(R)} y_e \leq 1/2$. \square

Let w_1, \dots, w_k be the vertices in T owned by S . The main idea in the proof of Lemma 5 is to consider $Y := \sum_{i=1}^k y(\delta(w_i))$. Since $w_i \in T$, it follows from (1) that $y(\delta(w_i)) = d(w_i) - x(\delta(w_i)) = d(w_i) - b_{w_i} \geq 3r_{\max} + 1$ and thus

$$Y \geq (3r_{\max} + 1)k.$$

By Lemma 7, the heavy edges from a poor child can only contribute very little to this sum, and this will allow us to rule out the existence of a poor child in S .

Proof of Lemma 5 when $r_{\max} \geq 2$: First, we count the number of tokens that S can collect. Consider the edges in $F := \cup_{i=1}^k \delta(w_i)$. Let $F_2 \subseteq F$ be the subset of edges of F with both endpoints in S , and $F_1 := F - F_2$ be the subset of edges of F with exactly one endpoint in S . Note that each edge in F_2 can contribute two tokens to S , regardless of whether it is heavy or light. Let ℓ be the number of light edges in F_1 . Then each such edge can contribute one token to S . Let γ be the number of rich children of S and ρ be the number of poor children of S . By the induction hypothesis, the children can contribute at least $2\gamma + \rho$ tokens to S . Therefore, S can collect at least $2\gamma + \rho + \ell + 2|F_2|$ tokens from its children and the endpoints that it owns. If $2\gamma + \rho + \ell + 2|F_2| \geq 2k + 4$, then we are done. So we assume to the contrary that

$$2\gamma + \rho + \ell + 2|F_2| \leq 2k + 3. \tag{2}$$

Next, we consider the contribution to Y . Each endpoint of an edge e can contribute strictly less than one to Y , as $y_e = 1 - x_e < 1$ for each edge. So, the edges in F_2 and the light edges in F_1 can contribute strictly less than $2|F_2| + \ell$ to Y . It remains to count the contribution from the heavy edges in F_1 . The heavy edges from a rich child R can contribute at most r_{\max} to Y , because each heavy edge can contribute at most $\frac{1}{2}$ to Y and $|\delta^h(R)| \leq 2r_{\max}$ as $x(\delta(R)) \leq r_{\max}$. The heavy edges from a poor child can contribute at most $\frac{1}{2}$ to Y by Lemma 7. Finally, the heavy edges in $F_1 \cap \delta^h(S)$ can contribute at most r_{\max} to Y , because $|\delta^h(S)| \leq 2r_{\max}$ as $x(\delta(S)) \leq r_{\max}$. These count all the contributions to Y . Therefore, we must have

$$(3r_{\max} + 1)k \leq Y \leq (\gamma + 1) \cdot r_{\max} + \frac{1}{2}\rho + \ell + 2|F_2|. \tag{3}$$

To satisfy (2) as an equality, we must have $\rho + \ell + 2|F_2| \geq 1$ since γ is an integer. If $\rho + 2|F_2| \geq 1$, then the second inequality in (3) is a strict inequality. Otherwise, if $\rho \geq 1$ or (2) is not an equality, by plugging in (2), we also have the following strict inequality

$$(3r_{\max} + 1)k < (\gamma + 1) \cdot r_{\max} + 2k + 3 - 2\gamma \iff (3k - 1)r_{\max} < k + 3 + \gamma(r_{\max} - 2). \tag{4}$$

As $\gamma \leq k + 1$ by (2) and assuming $r_{\max} \geq 2$, we have

$$(3k - 1)r_{\max} < k + 3 + (k + 1)(r_{\max} - 2) \iff r_{\max}(2k - 2) < 1 - k,$$

which is impossible for $k \geq 1$. This contradiction shows that (2) cannot hold, and thus S can collect $2k + 4$ tokens as claimed. \square

The proof for $r_{\max} = 1$ will be shown in the full version.

3.2 A Hard Example for the Algorithm

A natural question is that, for the minimum bounded degree Steiner forest problem, whether we can improve our algorithm further by only relaxing vertices with degree at most $b_v + 2$. This would imply a $(2, b_v + 2)$ -approximation algorithm for the problem, matching the known integrality gap for this linear program [12].

In the example shown in Figure 2, some vertices have a degree bound equal to two, but there are five edges incident to these vertices. This is an extreme point solution to (LP) as the characteristic vectors are linearly independent. Our algorithm will get stuck in this example, and it is not clear to us how to modify the algorithm to deal with it. We believe that some new ideas are needed to obtain a $(2, b_v + 2)$ -approximation algorithm for this problem.

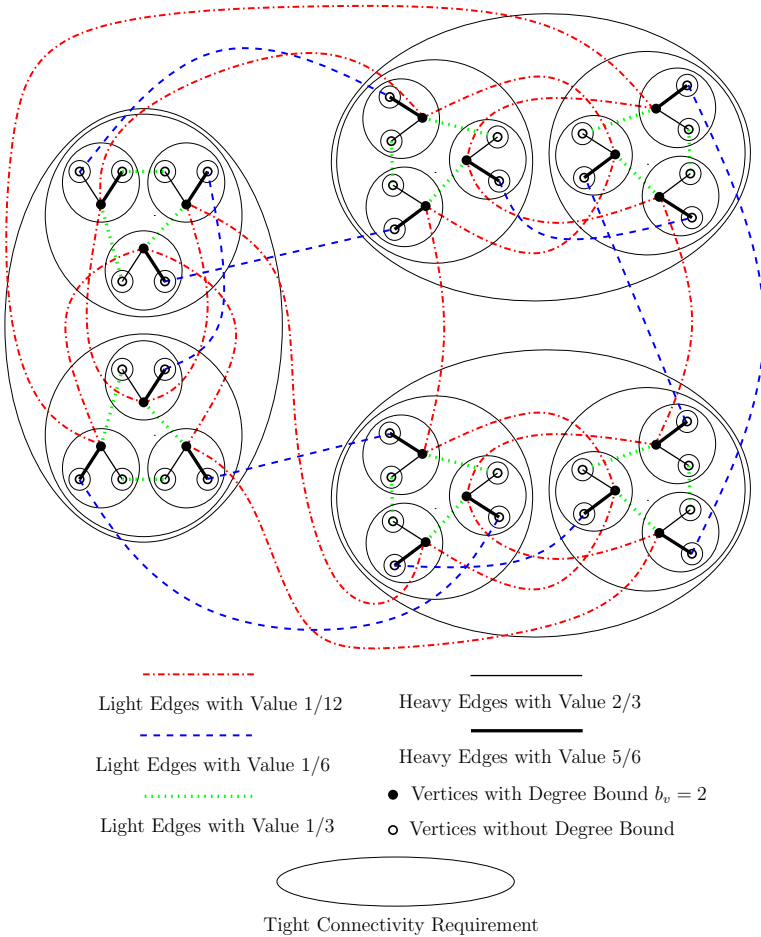


Fig. 2. A hard example

Acknowledgement. We thank the anonymous reviewers for comments that improve the presentation of the paper. This research is partially supported by HK RGC grant 2150701.

References

1. Bansal, N., Khandekar, R., Nagarajan, V.: Additive guarantees for degree bounded directed network design. *SIAM Journal on Computing* 29, 1413–1431 (2009)
2. Cheriyan, J., Vegh, L.: Approximating minimum-cost k -node connected subgraphs via independence-free graphs. In: *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)* (2013)
3. Cheriyan, J., Vempala, S., Vetta, A.: Network design via iterative rounding of setpair relaxations. *Combinatorica* 26(3), 255–275 (2006)
4. Fleischer, L., Jain, K., Williamson, D.P.: Iterative rounding 2-approximation algorithms for minimum-cost vertex connectivity problems. *J. Comput. Syst. Sci.* 72(5), 838–867 (2006)
5. Fukunaga, T., Ravi, R.: Iterative rounding approximation algorithms for degree-bounded node-connectivity network design. In: *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 263–272 (2012)
6. Fürer, M., Raghavachari, B.: Approximating the minimum-degree Steiner tree to within one of optimal. *J. of Algorithms* 17(3), 409–423 (1994)
7. Gabow, H.N.: On the L_∞ -norm of extreme points for crossing supermodular directed network LPs. In: Jünger, M., Kaibel, V. (eds.) *IPCO 2005*. LNCS, vol. 3509, pp. 392–406. Springer, Heidelberg (2005)
8. Goemans, M.X.: Minimum Bounded-Degree Spanning Trees. In: *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pp. 273–282 (2006)
9. Jain, K.: A factor 2-approximation algorithm for the generalized steiner network problem. *Combinatorica* 21(1), 39–60 (2001)
10. Khandekar, R., Kortsarz, G., Nutov, Z.: On some network design problems with degree constraints. *Journal of Computer and System Sciences* 79(5), 725–736 (2013)
11. Király, T., Lau, L.C., Singh, M.: Degree Bounded Matroids and Submodular Flows. In: Lodi, A., Panconesi, A., Rinaldi, G. (eds.) *IPCO 2008*. LNCS, vol. 5035, pp. 259–272. Springer, Heidelberg (2008)
12. Lau, L.C., Naor, S., Salavatipour, M., Singh, M.: Survivable network design with degree or order constraints. *SIAM Journal on Computing* 39(3), 1062–1087 (2009)
13. Lau, L.C., Ravi, R., Singh, M.: *Iterative Methods in Combinatorial Optimization*. Cambridge University Press (2011)
14. Lau, L.C., Singh, M.: Additive approximation for bounded degree survivable network design. *SIAM Journal on Computing* 42(6), 2217–2242 (2014)
15. Louis, A., Vishnoi, N.K.: Improved algorithm for degree bounded survivable network design problem. In: *Proceedings of the 12th Scandinavian Symposium and Workshops on Algorithm Theory*, pp. 408–419 (2010)
16. Nutov, Z.: Degree-constrained node-connectivity. In: *Proceedings of the 10th Latin American Symposium on Theoretical Informatics (LATIN)*, pp. 582–593 (2012)
17. Singh, M., Lau, L.C.: Approximating minimum bounded degree spanning trees to within one of optimal. In: *Proceedings of the 39th ACM Symposium on Theory of Computing (STOC)*, pp. 661–670 (2007)