Quasi-Polynomial Time Approximation Algorithm for Low-Degree Minimum-Cost Steiner Trees

J. Könemann and R. Ravi

GSIA, Carnegie Mellon University,
Pittsburgh PA 15213, USA.
{jochen,ravi}@cmu.edu

Abstract. In a recent paper [5], we addressed the problem of finding a minimum-cost spanning tree T for a given undirected graph $G = (V, E)$ with maximum node-degree at most a given parameter $B > 1$. We developed an algorithm based on Lagrangean relaxation that uses a repeated application of Kruskal's MST algorithm interleaved with a combinatorial update of approximate Lagrangean node-multipliers maintained by the algorithm.

In this paper, we show how to extend this algorithm to the case of Steiner trees where we use a primal-dual approximation algorithm due to Agrawal, Klein, and Ravi [1] in place of Kruskal's minimum-cost spanning tree algorithm. The algorithm computes a Steiner tree of maximum degree $O(B + \log n)$ and total cost that is within a constant factor of that of a minimum-cost Steiner tree whose maximum degree is bounded by B. However, the running time is quasi-polynomial.

1 Introduction

We consider the minimum-degree Steiner tree problem (B-ST) where we are given an undirected graph $G = (V, E)$, a non-negative cost c_e for each edge $e \in E$ and a set of terminal nodes $R \subseteq V$. Additionally, the problem input also specifies positive integers $\{B_v\}_{v \in V}$. The goal is to find a minimum-cost Steiner tree T covering R such each node v in T has degree at most B_v, i.e. $\deg_T(v) \leq B_v$ for all $v \in V$. We present an algorithm for the problem and give a proof of the following theorem.

Theorem 1. *There is a primal-dual approximation algorithm that, given a graph $G = (V, E)$, a set of terminal nodes $R \subseteq V$, a nonnegative cost function $c : E \to \mathcal{R}^+$, integers $B_v > 1$ for all $v \in V$, and an arbitrary $b > 1$ computes a Steiner tree T that spans the nodes of R such that*

1. $\deg_T(v) \leq 12b \cdot B_v + \lceil 4 \log_b n \rceil + 1$ *for all* $v \in V$, *and*
2. $c(T) \leq 3\texttt{OPT}$

where OPT *is the minimum cost of any Steiner tree whose degree at node v is bounded by B_v for all v. Our method runs in $O(n \log(|R|) \cdot |R|^{\lceil 4 \log n \rceil})$ iterations each of which can be implemented in polynomial time.*

The algorithm combines ideas from the primal-dual algorithm for Steiner trees due to Agrawal et al. [1] with local search elements from [2].

2 A linear programming formulation

The following natural integer programming formulation models the problem. Let $\mathcal{U} = \{U \subseteq V : U \cap R \neq \emptyset, R \setminus U \neq \emptyset\}$.

$$\begin{aligned} \min \quad & \sum_{e \in E} c_e x_e & \text{(IP)} \\ \text{s.t} \quad & x(\delta(U)) \geq 1 \quad \forall U \in \mathcal{U} & (1) \\ & x(\delta(v)) \leq B_v \quad \forall v \in V & (2) \\ & x \text{ integer} \end{aligned}$$

The dual of the linear programming relaxation (LP) of (IP) is given by

$$\begin{aligned} \max \quad & \sum_{U \in \mathcal{U}} y_U - \sum_{v \in V} \lambda_v \cdot B_v & \text{(D)} \\ \text{s.t} \quad & \sum_{U : e \in \delta(U)} y_U \leq c_e + \lambda_u + \lambda_v \quad \forall e = uv \in E & (3) \\ & y, \lambda \geq 0 \end{aligned}$$

We also let (IP-ST) denote (IP) without constraints of type (2). This is the usual integer programming formulation for the Steiner tree problem. Let the LP relaxation be denoted by (LP-ST) and let its dual be (D-ST).

3 An algorithm for the Steiner tree problem

Our algorithm is based upon previous work on the generalized Steiner tree problem by Agrawal et al. [1]. A fairly complete description of their algorithm to compute an approximate Steiner tree can be found in the extended version of this paper. We refer to the algorithm by `AKR`.

Before proceeding with the description of an algorithm for minimum-cost degree-bounded Steiner trees, we present an alternate view of Algorithm `AKR` which simplifies the following developments in this paper.

3.1 An alternate view of `AKR`

Executing `AKR` on an undirected graph $G = (V, E)$ with terminal set $R \subseteq V$ and costs $\{c_e\}_{e \in E}$ is – in a certain sense – equivalent to executing `AKR` on the complete graph H with vertex set R where the cost of edge $e = uv \in R \times R$ is equal to the minimum cost of any u, v-path in G.

Let $\mathcal{P}_{S,T}$ denote the set of s,t-Steiner paths for all $s \in R \cap S$ and $t \in R \cap T$. For an s,t-Steiner path P, we define

$$c^\lambda(P) = c(P) + \lambda_s + \lambda_t + 2 \cdot \sum_{v \in \mathtt{int}(P)} \lambda_v$$

and finally, let

$$\mathtt{dist}_{c^\lambda}(S,T) = \min_{P \in \mathcal{P}_{S,T}} c^\lambda(P)$$

be the minimum c^λ-cost for any S,T-Steiner path. We now work with the following dual:

$$\max \quad \sum_{U \subset R} y_U - \sum_{v \in V} \lambda_v \cdot B_v \tag{D2}$$

$$\text{s.t} \quad \sum_{U \subset R, s \in U, t \notin U} y_U \leq c^\lambda(P) \quad \forall s,t \in R, P \in \mathcal{P}_{s,t} \in E \tag{4}$$

$$y, \lambda \geq 0$$

Let H be a complete graph on vertex set R. We let the length of edge $(s,t) \in E[H]$ be $\mathtt{dist}_{c^\lambda}(s,t)$. Running AKR on input graph H with length function $\mathtt{dist}_{c^\lambda}$ yields a tree in H that corresponds to a Steiner tree spanning the nodes of R in G in a natural way. We also obtain a feasible dual solution for (D2). The following lemma shows that (D) and (D2) are equivalent (the proof is deferred to the extended version of this paper).

Lemma 1. *(D) and (D2) are equivalent.*

4 An algorithm for the B-ST problem

In this section, we propose a modification of AKR in order to compute a feasible degree-bounded Steiner tree of low total cost. We start by giving a rough overview over our algorithm.

4.1 Algorithm: Overview

We define the normalized degree $\mathtt{ndeg}_T(v)$ of a node v in a Steiner tree T as

$$\mathtt{ndeg}_T(v) = \max\{0, \deg_T(v) - \beta_v \cdot B_v\} \tag{5}$$

where $\{\beta_v\}_{v \in V}$ are parameters to be defined later.

Algorithm B-ST goes through a sequence of Steiner trees $E^0, \ldots, E^t$ and associated pairs of primal (infeasible) and dual feasible solutions x^i, (y^i, λ^i) for $0 \leq i \leq t$. The goal is to reduce the maximum normalized degree of at least one node in V in the transition from one Steiner tree to the next.

In the i^{th} iteration our algorithm passes through two main steps:

Compute Steiner tree. Compute an approximate Steiner tree spanning the nodes of R for our graph $G = (V, E)$ using a modified version of AKR. Roughly speaking, this algorithm implicitly assumes a cost function $\widetilde{c}$ that satisfies

$$c(P) \leq \widetilde{c}(P) \leq c^{\lambda^i}(P) \tag{6}$$

for all Steiner paths P.

When the algorithm finishes, we obtain a primal solution x^i together with a corresponding dual solution y^i. In the following we use $\mathcal{P}^i$ to denote the set of paths used by AKR to connect the terminals in iteration i.

Notice that using the cost function $\widetilde{c}$ that satisfies (6) ensures that (y^i, λ^i) is a feasible solution for (D2). The primal solution x^i may induce high normalized degree at some of the vertices of V and hence may not be feasible for (IP).

Update node multipliers. The main goal here is to update the node multipliers λ^i such that another run of AKR yields a tree in which the normalized degree of at least one node decreases. Specifically, we continue running our algorithm as long as the maximum normalized node-degree induced by x^i is at least $2 \log_b n$ where $b > 1$ is a positive constant to be specified later.

Let Δ^i be the maximum normalized degree of any node in the tree induced by x^i. The algorithm then picks a threshold $d^i \geq \Delta^i - \lceil 4 \log_b n \rceil + 2$. Subsequently we raise the λ values of all nodes that have normalized degree at least $d^i - 2$ in the tree induced by x^i by some $\epsilon^i > 0$. We also implicitly increase the $\widetilde{c}$ cost of two sets of Steiner paths:

1. those paths $P \in \mathcal{P}^i$ that contain nodes of degree at least d^i and
2. those paths $P \notin \mathcal{P}^i$ that contain nodes of degree at least $d^i - 2$.

We denote to the set of all such paths by $\mathcal{L}^i$.

Rerunning AKR replaces at least one Steiner path whose $\widetilde{c}$-cost increased with a Steiner path whose length stayed the same. In other words, a path that touches a node of normalized degree at least d^i is replaced by some other path that has only nodes of normalized degree less than $d^i - 2$.

Throughout the algorithm we will maintain that the cost of the current tree induced by x^i is within a constant factor of the dual objective function value induced by (y^i, λ^i). By weak duality, this ensures that the cost of our tree is within a constant times the cost of any Steiner tree that satisfies the individual degree bounds. However, we are only able to argue that the number of iterations is quasi-polynomial.

In the following, we will give a detailed description of the algorithm. In particular, we elaborate on the choice of ϵ^i and d^i in the node-multiplier update and on the modification to AKR that we have alluded to in the previous intuitive description.

4.2 Algorithm: A detailed top-level description

We first present the pseudo-code of our B-ST-algorithm. In the description of the algorithm we use the abbreviation $\mathtt{ndeg}^i(v)$ in place of $\mathtt{ndeg}_{E^i}(v)$ for the normalized degree of vertex v in the Steiner tree E^i. We also let Δ^i denote the maximum normalized degree of any vertex in E^i, i.e.

$$\Delta^i = \max_{v \in R} \texttt{ndeg}^i(v).$$

Furthermore, we adopt the notation of [2] and let

$$S^i_d = \{v \in V \ : \texttt{ndeg}^i(v) \geq d\}$$

be the set of all nodes whose normalized degrees are at least d in the i^{th} solution.

The following lemma is proved easily by contradiction

Lemma 2. *There is a* $d^i \in \{\Delta^i - \lceil 4 \log_b n \rceil + 2, \ldots, \Delta^i\}$ *such that*

$$\sum_{v \in S^i_{d^i-2}} B_v \leq b \cdot \sum_{v \in S^i_{d^i}} B_v$$

for a given constant $b > 1$.

The *low expansion* of $S^i_{d^i-2}$ turns out to be crucial in the analysis of the performance guarantee of our algorithm.

Finally, we let `mod-AKR` denote a call to the modified version of the Steiner tree algorithm `AKR`. Algorithm 1 has the pseudo code for our method.

Algorithm 1 An algorithm to compute an approximate minimum-cost degree-bounded Steiner tree.

1: Given: primal feasible solution $x^0, \mathcal{P}^0$ to (LP-ST) and dual feasible solution y^0 to (D-ST)
2: $\lambda^0_v \leftarrow 0, \forall v \in V$
3: $i \leftarrow 0$
4: **while** $\Delta^i > 4\lceil \log_b n \rceil$ **do**
5: Choose $d^i \geq \Delta^i - \lceil 4 \log_b n \rceil + 2$ s.t. $\sum_{v \in S^i_{d^i-2}} B_v \leq b \cdot \sum_{v \in S^i_{d^i}} B_v$
6: Choose $\epsilon^i > 0$ and identify swap pair $\langle P^i, \overline{P}^i \rangle$.
7: $\lambda^{i+1}_v \leftarrow \lambda^i_v + \epsilon^i$ if $v \in S^i_{d^i-2}$ and $\lambda^{i+1}_v \leftarrow \lambda^i_v$ otherwise
8: $y^{i+1} \leftarrow$ `mod-AKR`$(\mathcal{P}^i, \epsilon^i, y^i, \langle P^i, \overline{P}^i \rangle)$
9: $\mathcal{P}^{i+1} \leftarrow \mathcal{P}^i \setminus \{P^i\} \cup \{\overline{P}^i\}$
10: $i \leftarrow i + 1$
11: **end while**

Step 6 of Algorithm 1 hides the details of choosing an appropriate ϵ^i. We lengthen all Steiner paths in $\mathcal{L}^i$. Our choice of ϵ^i will ensure that there exists at least one point in time during the execution of a slightly modified version of `AKR` in step 8 at which we now have the choice to connect two moats using paths P^i and $\overline{P}^i$, respectively. We show that there is a way to pick ϵ^i such that

$$P^i \cap S^i_{d^i} \neq \emptyset \text{ and } \overline{P}^i \cap S^i_{d^i-2} = \emptyset.$$

We now break ties such that $\overline{P}^i$ is chosen instead of P^i and hence, we end up with a new Steiner tree E^{i+1}.

In mod-AKR, we prohibit including alternate paths that contain nodes from $S^i_{d^i-2}$ and argue that the dual load that such a non-tree path P' sees does not go up by more than ϵ^i. Hence, we preserve dual feasibility.

We first present the details of Algorithm mod-AKR and discuss how to find ϵ^i afterwards.

4.3 Algorithm: mod-AKR

Throughout this section and the description of mod-AKR we work with the modified dual (D2) as discussed in Section 3.1.

For a r_1, r_2-Steiner path P we let $R_P \subseteq 2^R$ denote all sets $S \subset R$ that contain exactly one of $r_1, r_2 \in R$. For a dual solution y, λ we then define the cut-metric $l_y(P) = \sum_{S \in R_P} y_s$. From here it is clear that (y, λ) is a feasible dual solution iff $l_y(P) \leq c^{\lambda}(P)$ for all Steiner paths P. We use $l^i(P)$ as an abbreviation for $l_{y^i}(P)$.

At all times during the execution of Algorithm 1 we want to maintain dual feasibility, i.e. we maintain

$$l^i(P) \leq c^{\lambda^i}(P) \tag{7}$$

for all Steiner paths P and for all i. Moreover, we want to maintain that for all i, the cost of any path $P \in \mathcal{P}^i$ is bounded by the dual load that P sees. In other words, we want to enforce that

$$c(P) \leq l^i(P) \tag{8}$$

for all $P \in \mathcal{P}^i$ and for all i. It is easy to see that both (7) and (8) hold for $i = 0$ from the properties of AKR.

First, let $\mathcal{P}^i = \mathcal{P}^i_1 \cup \mathcal{P}^i_2$ be a partition of the set of Steiner paths used to connect the terminal nodes in the i^{th} iteration. Here, a path $P \in \mathcal{P}^i$ is added to $\mathcal{P}^i_1$ iff $P \cap S^i_{d^i} \neq \emptyset$ and we let $\mathcal{P}^i_2 = \mathcal{P}^i \setminus \mathcal{P}^i_1$.

mod-AKR first constructs an auxiliary graph G^i with vertex set R. We add an edge (s,t) to G^i for each s,t-path $P \in \mathcal{P}^i \setminus \{P^i\}$. The edge (s,t) is then assigned a length of $l^{i+1}_{st} = l^i(P) + \epsilon^i$ if $P \in \mathcal{P}^i_1$ and $l^{i+1}_{st} = l^i(P)$ otherwise.

Assume that $\overline{P}^i$ is an s', t'-path. We then also add an edge connecting s' and t' to G^i and let its length be the maximum of $l^i(\overline{P}^i)$ and $c(\overline{P}^i)$. Observe, that since $\mathcal{P}^i \setminus \{P^i\} \cup \{\overline{P}^i\}$ is tree, G^i is a tree as well.

Subsequently, mod-AKR runs AKR on the graph G^i and returns the computed dual solution. We will show that this solution together with λ^{i+1} is feasible for (D2). A formal definition of mod-AKR is given in Algorithm 2.

We defer the proof of invariants (7) and (8) to the end of the next section.

Algorithm 2 `mod-AKR`$(\mathcal{P}^i, \epsilon^i, y^i, \langle P^i, \overline{P}^i \rangle)$: A modified version of `AKR`.

1: Assume $\overline{P}^i$ is an s', t'-Steiner path
2: $G^i = (R, E^i)$ where
$E^i = \{(s,t) : \exists s,t-\text{path } P \in \mathcal{P}^i \setminus \{P^i\}\} \cup \{(s',t')\}$
3: For all s, t Steiner paths $P \in \mathcal{P}^i \setminus \{P^i\}$:

$$l^{i+1}_{st} = \begin{cases} l^i(P) + \epsilon^i & : \quad P \in \mathcal{P}^i_1 \\ l^i(P) & : \quad \text{otherwise} \end{cases}$$

4: $l^{i+1}_{s't'} = \max\{c(\overline{P}^i), l^i(\overline{P}^i)\}$
5: $y^{i+1} \leftarrow$ `AKR`(G^i, l^{i+1})
6: `return` y^{i+1}

4.4 Algorithm: Choosing ϵ^i

In this section, we show how to choose ϵ^i. Remember that, intuitively, we want to increase the cost of currently used Steiner paths that touch nodes of normalized degree at least d^i. The idea is to increase the cost of such paths by the smallest possible amount such that other non-tree paths whose length we did not increase can be used at their place. We make this idea more precise in the following.

We first define $\mathcal{K}^i$ to be the set of connected components of

$$G\left[\bigcup_{P \in \mathcal{P}^i_2} P\right].$$

Let H^i be an auxiliary graph that has one node for each set in $\mathcal{K}^i$. Moreover, H^i contains edge (K', K'') iff there is a K', K''-Steiner path in the set $\mathcal{P}^i_1$. It can be seen that each path $P \in \mathcal{P}^i_1$ corresponds to unique edge in H^i. It then follows from the fact that $G[E^i]$ is a tree that H^i must also be a tree.

For $K', K'' \in \mathcal{K}^i$ such that (K', K'') is not an edge of H^i, let C be the unique cycle in $H^i + (K', K'')$. We then use $\mathcal{P}^i(C)$ to denote the set of Steiner paths from $\mathcal{P}^i$ corresponding to edges on C.

For any two connected components $K', K'' \in \mathcal{K}^i$ we let

$$d^i(K', K'') = \min_{\substack{P \in \mathcal{P}_{K',K''} \\ P \cap S^i_{d^i-2} = \emptyset}} c(P). \tag{9}$$

be the cost of the minimum-cost K', K''-Steiner path that *avoids* nodes from $S^i_{d^i-2}$. For a pair of components $K', K'' \in \mathcal{K}^i$ we denote the path that achieves the above minimum by $P_{K',K''}$.

Definition 1. *We say that a path $\overline{P} \notin \mathcal{P}^i$ that contains no nodes from $S^i_{d^i-2}$ is ϵ-swappable against $P \in \mathcal{P}^i_1$ in iteration i if*

1. *$P \in \mathcal{P}^i(C)$ where C is the unique cycle created in H^i by adding the edge corresponding to $\overline{P}$, and*

2. $c(\overline{P}) \leq l^i(P) + \epsilon$

We are now looking for the smallest ϵ^i such that there exists a *witness pair* of paths $\langle P^i, \overline{P}^i \rangle$ where $\overline{P}^i$ is ϵ^i-swappable against P^i.

Formally consider all pairs $K', K'' \in \mathcal{K}^i$ such that (K', K'') is not an edge of H^i. Inserting the edge corresponding to $P_{K',K''}$ into H^i creates a unique cycle C. For each such path $P \in \mathcal{P}^i(C)$, let $\epsilon^i_{K',K''}(P)$ be the smallest non-negative value of ϵ such that

$$d^i(K', K'') \leq l^i(P) + \epsilon. \tag{10}$$

We then let $\epsilon^i_{K',K''} = \min_{P \in \mathcal{P}^i(C)} \epsilon^i_{K',K''}(P)$ and define

$$\epsilon^i = \min_{K',K'' \in \mathcal{K}^i} \epsilon^i_{K',K''}.$$

We let $\langle P^i, \overline{P}^i \rangle$ be the pair of Steiner paths that defines ϵ^i, i.e. $\overline{P}^i$ is a K', K''-Steiner path such that

1. inserting edge (K', K'') into H^i creates a cycle C and $P^i \in \mathcal{P}^i(C)$, and
2. $c(\overline{P}^i) \leq l^i(P^i) + \epsilon^i$.

We are now in the position to show that (7) and (8) are maintained for our choice of $(P^i, \overline{P}^i)$ and ϵ^i. The following Lemma whose proof is deferred to the full version of this paper shows that `mod-AKR` produces a feasible dual solution (y^{i+1}, λ^{i+1}) for (D2) provided that (y^i, λ^i) was dual feasible.

Lemma 3. *Algorithm 2 produces a feasible dual solution (y^{i+1}, λ^{i+1}) for (D2) given that (y^i, λ^i) is dual feasible for (D2).*

This shows (7). It is clear from the choice of ϵ^i that we include a Steiner path $\overline{P}^i$ into $\mathcal{P}^{i+1}$ only if $l^{i+1}(\overline{P}^i) \geq c(\overline{P}^i)$. (8) now follows since the dual load on any path is non-decreasing as we progress.

4.5 Analysis: Performance guarantee

In this section we show that the cost of the tree computed by Algorithm 1 is within a constant factor of any Steiner tree satisfying all degree bounds. We ensure this by way of weak duality. In particular, our goal is to prove the inequality

$$\sum_{P \in \mathcal{P}^i} c(P) \leq 3 \sum_{S \subset R} y^i_S - 3 \sum_{v \in V} B_v \cdot \lambda^i_v \tag{11}$$

for all iterations i of our algorithm.

First, we observe the following simple consequence of the `AKR` algorithm.

Lemma 4. *Assume that Algorithm 1 terminates after t iterations. For iteration $0 \leq i \leq t$, let $l^i_{\max} = \max_{P \in \mathcal{P}^i} l^i(P)$. We then must have*

$$\sum_{P \in \mathcal{P}^i} l^i(P) = 2 \sum_{S \subset R} y^i_S - l^i_{\max}.$$

Proof. Let $r = |R|$ and let $\mathcal{P}^i = \{P_1^i, \ldots, P_{r-1}^i\}$ be the paths computed by `mod-AKR` in iteration $i-1$. Also let y^i be the corresponding dual solution returned by `mod-AKR`. W.l.o.g. we may assume that

$$l^i(P_1^i) \leq \ldots \leq l^i(P_{r-1}^i).$$

From the `AKR`algorithm it is not hard to see that

$$\begin{aligned}
\sum_{S \subset R} y_S^i &= \frac{1}{2} \cdot \sum_{j=1}^{r-1} (l^i(P_j^i) - l^i(P_{j-1}^i)) \cdot (r - j + 1) \qquad (12) \\
&= \frac{1}{2} \cdot \sum_{j=1}^{r-1} l^i(P_j^i)\,((r-j+1) - (r-j)) + \frac{1}{2} l^i(P_{r-1}^i) \\
&= \frac{1}{2} \cdot \sum_{j=1}^{r-1} l^i(P_j^i) + \frac{1}{2} l^i(P_{r-1}^i)
\end{aligned}$$

where we define $l^i(P_0^i) = 0$. The last equality (12) can be restated as

$$\sum_{P \in \mathcal{P}^i} l^i(P) = 2 \sum_{S \subset R} y_S^i - l_{\max}^i$$

and that yields the correctness of the lemma.

We now proceed with proving (11) for all $1 \leq i \leq t$. Notice that Lemma 4 together with (8) implies (11) for $i = 0$. We concentrate on the case $i \geq 1$.

The proof is based on the following invariant that we maintain inductively for all $0 \leq i \leq t$:

$$3 \cdot \sum_{v \in V} B_v \lambda_v^i \leq \sum_{S \subset R} y_S^i. \qquad \text{(Inv)}$$

Since, $\lambda_v^0 = 0$ for all $v \in V$ by definition, (Inv) holds for $i = 0$.

Growing λ_v^i by ϵ^i at nodes $v \in S_{d^i-2}^i$ decreases the right hand side of (11) by $3 \cdot \epsilon \cdot \sum_{v \in S_{d^i-2}^i} B_v$. Still the cost of the Steiner tree E^{i+1} is potentially higher than the cost of the old tree E^i. We must show that the first term on the right hand-side of (11), i.e. $3 \cdot \sum_{S \subset R} y_S^i$ grows sufficiently to compensate for the decrease in the second term and the increased Steiner tree cost. In order to show this we need the following technical lemma that lower-bounds the number of paths that contain nodes of degree at least d^i in terms of the number of nodes of normalized degree at least $d^i - 2$.

Lemma 5. *In each iteration $1 \leq i \leq t$ we must have*

$$|\mathcal{P}_1^i| \geq \alpha \cdot \sum_{v \in S_{d^i-2}^i} B_v$$

for an arbitrary parameter $\alpha > 0$ by setting $\beta_v \geq 2\alpha b + 1/B_v$ for all $v \in V$ in the definition of $\texttt{ndeg}_T(v)$ *in (5).*

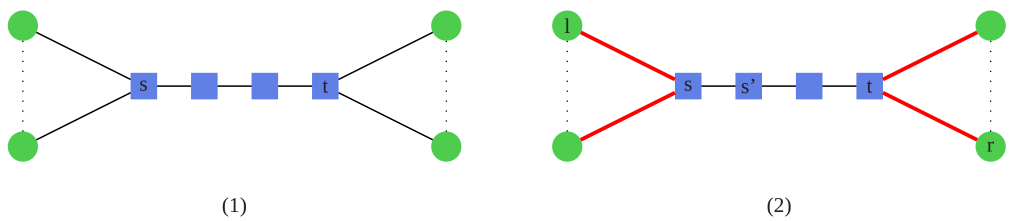

Fig. 1. Figure (1) shows a Steiner tree where circles represent terminals and squares represent Steiner nodes. We assume that there are exactly two nodes of high normalized degree: s and t. Figure (2) shows the set M of marked edges in red. Notice that the edge between Steiner nodes s and s' is not marked since there must be a Steiner path connecting a terminal node l on the left side and a terminal node r on the right side. This Steiner path has the form $\langle P_{ls}, ss', P_{s'r}\rangle$ and P_{ls} contains node s which has high normalized degree.

Proof. We first define a set of marked edges

$$M \subseteq \bigcup_{v \in S^i_{d^i}} \delta(v)$$

and then show that each Steiner path that contains nodes from $S^i_{d^i}$ has at most two marked edges. This shows that the cardinality of the set of marked edges is at most twice the number of paths in $\mathcal{P}^i_1$, i.e.

$$|M| \leq 2 \cdot |\mathcal{P}^i_1|. \tag{13}$$

In the second part of the proof we argue that M is sufficiently large.

First, we include all edges that are incident to terminal nodes from $S^i_{d^i}$ into M. Secondly, we also mark edges $uv \in E^i$ that are incident to non-terminal nodes in $S^i_{d^i}$ and that in addition satisfy that there is no Steiner path

$$P = \langle P_1, uv, P_2\rangle \in \mathcal{P}^i$$

such that both P_1 and P_2 contain nodes from $S^i_{d^i}$.

It is immediately clear from this definition that each Steiner path $P \in \mathcal{P}^i$ has at most two edges from M.

We now claim that M contains at least

$$b\alpha \cdot \sum_{v \in S^i_{d^i}} B_v \tag{14}$$

edges. To see this, we let T be the tree on node set $S^i_{d^i}$ that is induced by E^i: For $s, t \in S^i_{d^i}$ we insert the edge st into T iff the unique s,t-path in E^i has no other nodes from $S^i_{d^i}$. We let $P_e \subseteq E^i$ be the path that corresponds to an edge $e \in E[T]$.

Define $E^i_{d^i} \subseteq E^i$ to be the set of tree edges that are incident to nodes of normalized degree at least d^i, i.e.

$$E^i_{d^i} = \bigcup_{v \in S^i_{d^i}} \delta(v).$$

Now let $U \subseteq E^i$ be the set of unmarked tree edges that are incident to nodes of normalized degree at least d^i, i.e. $U = E^i_{d^i} \setminus M$.

First observe that, by definition of M, for each unmarked edge $e \in U$ there must be an edge $e^t \in E[T]$ such that e is an edge on the path P_{e^t}. Moreover, for all $e_t \in E[T]$ there are at most two unmarked edges on the path P_{e^t}. Since T has $|S^i_{d^i}| - 1$ edges we obtain

$$|U| \leq 2 \cdot (|S^i_{d^i}| - 1). \tag{15}$$

Each node in $S^i_{d^i}$ has at least $\beta_v B_v + d^i$ edges incident to it. On the other hand, since E^i is a tree, at most $(|S^i_{d^i}| - 1)$ of the edges in $E^i_{d^i}$ are incident to exactly two nodes from $S^i_{d^i}$. Hence, we obtain

$$|E^i_{d^i}| \geq \left(\sum_{v \in S^i_{d^i}} \beta_v B_v + d^i \right) - (|S^i_{d^i}| - 1) = \left(2\alpha b \cdot \sum_{v \in S^i_{d^i}} B_v \right) + d^i \cdot |S^i_{d^i}| + 1 \tag{16}$$

where the last equality uses the definition of β_v.

Now observe that $|M| = |E^i_{d^i}| - |U|$ and hence

$$|M| \geq \left(2\alpha b \cdot \sum_{v \in S^i_{d^i}} B_v \right) + |S^i_{d^i}|(d^i - 2) - 1. \tag{17}$$

using (15) and (16). Notice that $d^i \geq \Delta^i - \lceil 4 \log_b n \rceil + 2$ and $\Delta^i > \lceil 4 \log_b n \rceil$ and hence $d^i \geq 3$. This together with (17) and the fact that $S^i_{d^i}$ is non-empty implies

$$|M| \geq 2\alpha b \cdot \sum_{v \in S^i_{d^i}} B_v. \tag{18}$$

Combining (13) and (18) yields $|\mathcal{P}^i_1| \geq \alpha b \cdot \sum_{v \in S^i_{d^i}} B_v$. Using the fact that $\sum_{v \in S^i_{d^i - 2}} B_v \leq b \cdot \sum_{v \in S^i_{d^i}} B_v$ finishes the proof of the lemma.

The following claim now presents the essential insight that ultimately yields the validity of (11).

Lemma 6. *Let α be as in Lemma 5. We then must have*

$$\sum_{S \subset R} y^{i+1}_S \geq \sum_{S \subset R} y^i_S + \frac{\alpha}{2} \epsilon^i \cdot \sum_{v \in S^i_{d^i - 2}} B_v$$

for all $0 \leq i \leq t$.

Proof. We can use (12) to quantify the change in dual in iteration i.

$$\sum_{S \subset R} (y_S^{i+1} - y_S^i) = \frac{1}{2} \cdot \sum_{j=1}^{r-1} (l^{i+1}(P_j^i) - l^i(P_j^i)) + \frac{1}{2}(l^{i+1}(P_{r-1}^i) - l^{i+1}(P_{r-1}^i))$$
$$\geq \frac{\epsilon^i}{2} \cdot |\mathcal{P}_1^i|$$

where the inequality follows from the fact that we increase the length of all paths in $\mathcal{P}_1^i$ by ϵ^i and the length of all other paths are non-decreasing as we progress. An application of Lemma 5 finishes the proof.

As `mod-AKR` finishes with cut metric l^{i+1}, we obtain

$$l^{i+1}(\mathcal{P}^{i+1}) = \sum_{P \in \mathcal{P}^{i+1}} l^{i+1}(P) \leq 2 \sum_{S \subset R} y_S^{i+1} \tag{19}$$

from Lemma 4. Observe that the real cost of the Steiner tree E^{i+1} is much smaller than $l^{i+1}(\mathcal{P}^{i+1})$. In fact, notice that we have

$$\begin{aligned} c(\mathcal{P}^{i+1}) &\leq l^{i+1}(\overline{P}^i) + c(\mathcal{P}^i \setminus \{P^i\}) \\ &\leq l^{i+1}(\overline{P}^i) + l^i(\mathcal{P}^i \setminus \{P^i\}) \end{aligned} \tag{20}$$

where the last inequality follows from (8), i.e. the l-cost of a Steiner path in $\mathcal{P}^i$ always dominates its c-cost. Also, observe that

$$\begin{aligned} l^{i+1}(\mathcal{P}^i \setminus \{P^i\}) &= l^i(\mathcal{P}^i \setminus \{P^i\}) + \epsilon^i \cdot |\mathcal{P}_i^1| \\ &\geq l^i(\mathcal{P}^i \setminus \{P^i\}) + \alpha\epsilon^i \cdot \sum_{v \in S^i_{d^i-2}} B_v \end{aligned} \tag{21}$$

using Lemma 5. Combining (19), (20) and (21) yields

$$\begin{aligned} c(\mathcal{P}^{i+1}) &\leq l^{i+1}(\mathcal{P}^{i+1}) - \alpha\epsilon^i \cdot \sum_{v \in S^i_{d^i-2}} B_v \\ &\leq 2 \cdot \sum_{S \subset R} y_S^{i+1} - \alpha\epsilon^i \cdot \sum_{v \in S^i_{d^i-2}} B_v. \end{aligned}$$

We can now add (Inv) to the last inequality and get

$$c(\mathcal{P}^{i+1}) \leq 3 \sum_{S \subset R} y_S^{i+1} - 3 \cdot \sum_{v \in V} B_v \lambda_v^i - \alpha\epsilon^i \cdot \sum_{v \in S^i_{d^i-2}} B_v.$$

Finally notice that $\lambda_v^{i+1} = \lambda_v^i + \epsilon^i$ if $v \in S^i_{d^i-2}$ and $\lambda_v^{i+1} = \lambda_v^i$ otherwise. Now choose $\alpha \geq 3$ and it follows that

$$c(\mathcal{P}^{i+1}) \leq 3 \sum_{S \subset R} y_S^{i+1} - 3 \cdot \sum_{v \in V} B_v \lambda_v^{i+1}.$$

We have to show that (Inv) is maintained as well. Observe that the left hand side of (Inv) increases by $3\epsilon^i \cdot \sum_{v \in S^i_{d^i-2}} B_v$. We obtain from Lemma 6 that

$$\sum_{S \subset R} y_S^{i+1} - y_S^i \geq \frac{\alpha}{2} \cdot \epsilon^i \cdot \sum_{v \in S^i_{d^i-2}} B_v .$$

Choosing $\alpha \geq 6$ shows that the right hand side of (Inv) increases sufficiently and (Inv) holds in iteration $i+1$ as well.

4.6 Analysis: Running time

For a Steiner tree $\mathcal{P}$ in path representation, we define its potential value as

$$\Phi(\mathcal{P}) = \sum_{P \in \mathcal{P}} |R|^{\max_{v \in P} \mathtt{ndeg}_{\mathcal{P}}(v)}$$

where $\mathtt{ndeg}_{\mathcal{P}}(v)$ is the normalized degree of node v in the Steiner tree defined by $\mathcal{P}$. The proof of the following lemma is a direct adaptation of the arguments in [8] via the above potential function and is omitted.

Lemma 7. *Algorithm 1 terminates after $O(n \log(|R|) \cdot |R|^{\lceil 4 \log n \rceil})$ iterations.*

References

1. A. Agrawal, P. Klein and R. Ravi, "When trees collide: an approximation algorithm for the generalized Steiner problem on networks," *SIAM Journal on Computing*, vol.24, (1995), pp. 440-456.
2. M. Fürer, B. Raghavachari: Approximating the minimum degree Steiner tree to within one of optimal. Journal of Algorithms, **17** (1994) 409–423
3. M. X. Goemans and D. P. Williamson, "A general approximation technique for constrained forest problems," *SIAM J. Computing*, Vol. 24, 1995, pp. 296–317.
4. J. Könemann, R. Ravi: A matter of degree: improved approximation algorithms for degree-bounded minimum spanning trees. Proceedings of the 32nd Symp. on the Theory of Comput. (2000) 537–546. A complete version appears in the SIAM J. on Comput, 31(6), 1783–1793 (2002)
5. J. Könemann, R. Ravi: Primal-dual algorithms come of age: Approximating MST's with nonuniform degree bounds. Proceedings, ACM Symposium on Theory of Computing (2003)
6. R. Ravi, M. V. Marathe, S. S. Ravi, D. J. Rosenkrantz, H. B. Hunt III: Many Birds with One Stone: Multi-Objective Approximation Algorithms (Extended Abstract). Proceedings of the 25th Symp. on the Theory of Comput. (1993) 438–447
7. R. Ravi, M. V. Marathe, S. S. Ravi, D. J. Rosenkrantz, H. B. Hunt III: Approximation Algorithms for Degree-Constrained Minimum-Cost Network-Design Problems. Algorithmica **31(1)** (2001) 58–78
8. R. Ravi, B. Raghavachari, and P. N. Klein, "Approximation through local optimality: Designing networks with small degree," *Proc., 12th Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FST & TCS)*, December (1992), LNCS 652, pp. 279-290.