Bi-Criteria Approximation Algorithms for Bounded-Degree Subset TSP

Zachary Friggstad ✉

Department of Computing Science, University of Alberta, Edmonton, Canada

Ramin Mousavi ✉

Department of Computing Science, University of Alberta, Edmonton, Canada

Abstract

We initiate the study of the Bounded-Degree Subset Traveling Salesman problem (BDSTSP) in which we are given a graph $G = (V, E)$ with edge cost $c_e \geq 0$ on each edge e, degree bounds $b_v \geq 0$ on each vertex $v \in V$ and a subset of terminals $X \subseteq V$. The goal is to find a minimum-cost closed walk that spans all terminals and visits each vertex $v \in V$ at most $\frac{b_v}{2}$ times. In this paper, we study bi-criteria approximations that find tours whose cost is within a constant-factor of the optimum tour length while violating the bounds b_v at each vertex by additive quantities.

If $X = V$, an adaptation of the Christofides-Serdyukov algorithm yields a $(3/2, +4)$-approximation, that is the tour passes through each vertex at most $b_v/2 + 2$ times (the degree of v in the multiset of edges on the tour being at most $b_v + 4$). This is enabled through known results in bounded-degree Steiner trees and integrality of the bounded-degree Y-join polytope. The general case $X \neq V$ is more challenging since we cannot pass to the metric completion on X. However, it is at least simple to get a $(5/3, +4)$-bicriteria approximation by using ideas similar to Hoogeveen's TSP-Path algorithm.

Our main result is an improved approximation with marginally worse violations of the vertex bounds: a $(13/8, +6)$-approximation. We obtain this primarily through adapting the bounded-degree Steiner tree approximation to ensure certain "dangerous" nodes always have even degree in the resulting tree which allows us to bound the cost of the resulting degree-bounded Y-join. We also recover a $(3/2, +8)$-approximation for this general case.

2012 ACM Subject Classification Mathematics of computing → Approximation algorithms

Keywords and phrases Linear programming, approximation algorithms, combinatorial optimization

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2022.8

Funding *Zachary Friggstad*: Supported by an NSERC Discovery Grant and Accelerator Supplement.

1 Introduction

Consider the problem of having a very disruptive vehicle travel about a road network to serve some locations X. As with classic TSP, one could be interested in minimizing the total distance of this tour. However, we may want to restrict the number of times the vehicle passes through a location due to its disruptive nature. Or perhaps the driver does not wish to pass through a location too many times, e.g. it is difficult to traverse. In this paper, we consider approximations for this problem that induce only mild violations on these restrictions.

Without these traversal restrictions, the problem is equivalent to classic Traveling Salesman Problem (TSP), e.g. by considering the metric completion of the underlying graph and then restricting it to X. In an instance of TSP, we are given a graph $G = (V, E)$ with edge cost $c_e \geq 0$ on each edge $e \in E$. The goal is to find a shortest tour that visits all the vertices at least once. The Chistofides-Serdyukov algorithm [2, 17] from 1976 gives a simple $\frac{3}{2}$-approximation for TSP; this was only recently improved to $(\frac{3}{2} - \delta)$-approximation for some constant $\delta > 0$ by Karlin, Klein, and Oveis Gharan [6]. On the other hand, it is NP-hard to approximate TSP within a constant factor smaller than $\frac{123}{122}$ [7].

33rd International Symposium on Algorithms and Computation (ISAAC 2022).
Editors: Sang Won Bae and Heejin Park; Article No. 8; pp. 8:1–8:17
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In reality, TSP problems are concerned with visiting a subset of nodes of some larger graph. Thus, we consider the BOUNDED-DEGREE SUBSET TRAVELING SALESMAN problem (BDSTSP) in which we are given an undirected graph $G = (V, E)$ with edge costs $c_e \geq 0, e \in E$, a subset of terminals $X \subseteq V$ we are to visit ($|X| \geq 2$), and even integer bounds $b_v \geq 0$ for all nodes $v \in V$. The goal is to find a minimum-cost closed walk $\mathcal{Q}$ spanning all terminals such that $d_{\mathcal{Q}}(v) \leq b_v$ (i.e., the number of edges in the multiset $\mathcal{Q}$ incident to v is at most b_v). Note b_v should be thought of as a degree bound, thus the tour should pass through v at most $b_v/2$ times. We call a special case of BDSTSP where $X = V$, BOUNDED-DEGREE TRAVELING SALESMAN problem (BDTSP).

To the best of our knowledge (and to our surprise), BDTSP or BDSTSP have not been studied before. However, finding special subgraphs whose vertices satisfy some degree bounds has been an active research area in computer science and operations research, e.g., BOUNDED-DEGREE SPANNING TREES [5, 18], BOUNDED-DEGREE STEINER NETWORKS [9, 11, 12, 13], BOUNDED-DEGREE ELEMENT-CONNECTIVITY and BOUNDED-DEGREE VERTEX-CONNECTIVITY [4, 8].

Throughout this paper by a *tour* we mean a closed walk that spans all the required vertices, i.e., a tour in BDTSP is a closed walk that contains all the vertices and a tour in BDSTSP is a closed walk that contains all the terminals and possibly other vertices. The cost of a tour is the sum of the cost of the edges in the tour (counting with multiplicity). Note the edge set of a tour is potentially a multiset, we may be required to span an edge multiple times. All of our algorithms are based on linear-programming relaxations: if the LPs are not feasible then we report there is no feasible solution. Otherwise, we will find an $(\alpha, +d)$-approximate solution: the cost of the tour will be at most α times the value of the LP relaxation and will visit each node at most $(b_v + d)/2$ times (i.e. the degree of the tour at v will be at most $b_v + d$). We note that if there is a feasible solution, then the LP relaxations we use will be feasible and will have value at most the optimum solution value.

1.1 Our Results and Techniques

As a warm up, in Section 3 we present a simple $(3/2, +4)$-approximation algorithm for BDTSP (i.e. if $X = V$).

▶ **Theorem 1** (BDTSP). *There is a $(3/2, +4)$-approximation algorithm for BDTSP.*

Since a feasible solution is an Eulerian graph and b_v's are even, if there is an approximation algorithm whose degree violation is better than additive factor of 2, then this algorithm can decide the Hamiltonian cycle problem. Hence, assuming P $\neq$ NP, the additive factor of 2 violation on degree bounds is necessary. Furthermore, the same integrality gap example for Held-Karp relaxation where the degree bound on every vertex is 2, see Figure 1, shows the integrality gap of the natural LP formulation (**BDTSP-LP**) is at least $(4/3, +2)$, meaning any tour has cost at least $4/3$ times the LP optimum and any tour violates the degree bound of at least one vertex by at least $+2$.

The proof is a straightforward adaptation of Wolsey's analysis [19] of the Christofides-Serdyukov algorithm for TSP so we sketch it here to discuss our techniques. Let x^* be an optimal solution for the natural LP formulation of BDTSP. Step (1): using the natural cut-based LP formulation (augmented with degree bounds) for spanning trees and the fact that $\frac{x^*}{2}$ is feasible for this LP, using the rounding technique in [11] one can obtain a spanning tree T of cost at most $\sum_{e \in E} c_e \cdot x^*_e$ whose degree on vertex v is at most $\frac{b_v}{2} + 3$. Step (2): fix the degree parities of T using Y-join polytope augmented with degree bounds at most $\frac{b_v}{2} + 1$ (depending on the parity of v's degree in the tour) and show $\frac{x^*}{2}$ is feasible for this LP.

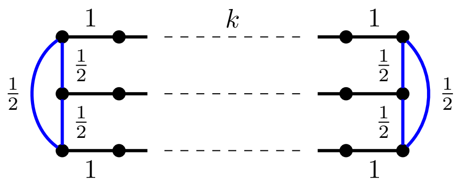

Figure 1 This is the graph $G = (V, E)$ that shows the $(\frac{4}{3}, +2)$ integrality gap of the natural LP for BDTSP (**BDTSP-LP**). All vertices are terminals, the cost of blue edges is zero and the cost of black edges is 1. The LP value on blue edges are $\frac{1}{2}$, and 1 on all the other edges. Also $b_v = 2$ for all $v \in V$. Note that the cost of the LP is $3 \cdot k$ and satisfies all the degree bounds. However, in any integer solution we must cross one of the path of length k at least twice. Therefore, any integer solution will violate the degree constraint by at least an additive factor of 2 and its cost is at least $4 \cdot k$ which give the desired integrality gap.

▶ Remark 2. Notice we did not use the +1 algorithm for degree-bounded spanning trees by [18]. This is because dividing x by 2 is only guaranteed to satisfy the weaker cut-based LP relaxation for spanning trees.

The main focus of this paper is on BDSTSP (i.e. $X \neq V$). Our results present different approximation/violation tradeoffs.

▶ **Theorem 3.** *There is a* $(5/3, +4)$*-approximation algorithm for BDSTSP.*

▶ **Theorem 4.** *There is a* $(13/8, +6)$*-approximation algorithm for BDSTSP.*

▶ **Theorem 5.** *There is a* $(3/2, +8)$*-approximation algorithm for BDSTSP.*

In each of these, we first adapt step (1) from BDSTP discussed above, compute a Steiner tree (instead of spanning tree) T using $\frac{x^*}{2}$ as a fractional solution to the Bounded-Degree Steiner Tree polytope. However, step (2) is not applicable since $\frac{x^*}{2}$ might not be feasible for the Y-join polytope.

To prove Theorem 3, it is easy to show that $\frac{1}{3} \cdot (\chi_T + x^*)$ is feasible for the degree-bounded Y-join polytope where χ_T is the characteristic vector of T which yields $(5/3, +4)$-approximation factor[1]. In order to improve the cost factor, we first augment the natural LP for BDSTSP with non-trivial constraints asserting the degree of Steiner cuts should be at least the degree of any Steiner node in the cut. Then, we modify the iterative rounding algorithm of [11] using splitting off techniques by Mader to obtain a more "structured" Steiner tree. Namely, some Steiner nodes are designated *dangerous* because they have low fractional degree in our LP solution: our modification ensures dangerous nodes will have even degree in the resulting tree. Finally, we show how this Steiner tree helps us to obtained a better bounded-degree Y-join to fix the degree parity of odd-degree vertices.

2 Preliminaries

In this section, we state definition and recall previous work that will be used throughout the paper. All graphs may be multigraphs and all subsets of edges may be multisubsets, we adopt this convention now so we do not have to use the prefix *multi* on every set or graph. In particular, when we discuss the degree of a vertex with respect to a set of edges or with

[1] Interestingly, this is basically the same fractional join from [1] that could be formed to analyze Hoogeveen's TSP-Path algorithm.

respect to a graph, we mean its degree if we count all edges with the same multiplicity that they appear in the set/graph. However, all subsets of vertices will be actual sets: each vertex will be in the set at most once.

Given a graph (or subgraph) $H = (V, E)$, denote by $d_H(v)$ the degree of vertex v in H. For a subset $S \subseteq V$, we denote by $\delta_H(S)$ the set of edges in $E(H)$ with exactly one endpoint in S while $E_H[S]$ is the set of all edges in $E(H)$ having both endpoints in S. We define $cost_c(H) := \sum_{e \in E(H)} c_e$. We may drop the subscripts in the above notation if the underlying graph/cost is clear from the context. Denote by $odd(H)$ the set of all vertices with odd degree in H, i.e., $odd(H) = \{v \in V(H) : \ d_H(v) \ is \ odd\}$. For a subset of edges (or a subgraph) F, we denote by χ_F the characteristic vector of the edges in F (i.e., $\chi_F(e) = 1$ if $e \in F$ and zero otherwise). For a solution x of an LP by $cost(x)$ we mean the value of the objective function given solution x. We sometime use notation $|A| = odd$ which means $|A| \equiv 1 \pmod 2$, similarly we define the notation $|A| = even$.

We use the Bounded-Degree Steiner Tree problem (BDSTP) and the Bounded-Degree Y-Join problem (BD-Y-join) in our results. In BDSTP, the input is a tuple $\big(G = (V, E), X, c, b\big)$, where c is the edge cost, i.e., $c_e \geq 0$ on each edge $e \in E$, $X \subseteq V$ is a set of terminals, and b is a degree bound for a subset of vertices $W \subseteq V$, i.e., $b_v \in \mathbb{Z}_{\geq 0}$ for all $v \in W$. Non-terminal vertices are called *Steiner* nodes. The goal is to compute a minimum cost connected subgraph T that spans all the terminals and respects the degree bounds, i.e., $d_T(v) \leq b_v$ for all $v \in W$.

There is a natural cut-based LP relaxation for this problem. Let $W \subseteq V$:

$$\text{minimize:} \quad \sum_{e \in E} c_e \cdot x_e \qquad \textbf{(SNDP-LP)}$$

$$\text{subject to:} \quad x(\delta(v)) \leq b_v \quad \forall v \in W \qquad (1)$$

$$x(\delta(S)) \geq 1 \quad \forall S \neq X, \ S \cap X \neq \emptyset \qquad (2)$$

$$x \geq 0 \qquad (3)$$

Lau and Singh [11] presented an iterative rounding algorithm for this LP.

▶ **Theorem 6** (Theorem 1.1 in [11]). *There exists a polynomial time algorithm for Bounded-Degree Steiner Tree which returns a Steiner tree of cost at most* $2 \cdot opt$ *with additive degree violation of at most* 3*, where opt is the cost of an optimal Steiner tree. In our notation, this is a* $(2, +3)$*-approximation algorithm.*

The above theorem is based on an iterative rounding method that rounds an extreme point of (**SNDP-LP**) to an integral solution. The following is an immediate consequence of the above theorem.

▶ **Corollary 7.** *Let* $\bar{x}$ *be a feasible solution to* (**SNDP-LP**)*. Then, in polynomial time, one could get a Steiner tree* T *of cost at most* $2 \cdot cost(\bar{x})$ *and* $d_T(v) \leq b_v + 3$ *for all* $v \in W$.

In the BD-Y-join problem, we are given a graph $G = (V, E)$, a non-negative cost c_e on each edge $e \in E$, a subset $Y \subseteq V$ with even size, and a degree bound $b_v \in \mathbb{Z}_{\geq 0}$ on each vertex such that b_v is odd for all $v \in Y$ and even otherwise. The goal is to find a minimum cost subset of edges $J \subseteq E$ such that $|d_J(v)| = odd$ if and only if $v \in Y$. Furthermore, we want $|d_J(v)| \leq b_v$ for all $v \in V$. There is a natural LP relaxation for this problem as well, which is just augmenting the integral Y-join polyhedron with the degree constraints.

$$\text{minimize:} \quad \sum_{e \in E} c_e \cdot x_e \qquad \textbf{(BD-}Y\textbf{-join LP)}$$

$$\text{subject to:} \quad x(\delta(v)) \leq b_v \quad \forall v \in V \tag{4}$$

$$x(\delta(S)) \geq 1 \quad \forall S \subsetneq V, \ |S \cap Y| = odd \tag{5}$$

$$x \geq 0 \tag{6}$$

It is known that (**BD-Y-join LP**) is integral.

▶ **Theorem 8** (Theorem 36.8 in [16]). (**BD-Y-join LP**) *is integral, if and only if, b_v is odd if $v \in Y$ and even otherwise for all $v \in V$.*

The above result is a corollary of Theorem 36.8 in [16] but it is not trivial to see at the first sight. So for the sake of completeness, we present a self-contained proof of Theorem 8 in Appendix C. Our proof is based on the iterative rounding technique which is different and simpler than the proof stated in [16] for Theorem 36.8, as we are trying to prove a special case of Theorem 36.8.

Also note that (**BD-Y-join LP**) admits a polynomial-time separation oracle since the odd-cut constraints can be separated just like with the classic Y-join polyhedron (eg. by using Gomory-Hu trees). So we can find a minimum-cost degree-bounded Y-join in polynomial time or determine one does not exist.

3 Bounded-Degree TSP (Warm Up!)

We quickly present the simple result for BDTSP in order to warm the reader up to how we use the results cited in the last section. Fix an instance of BDTSP: $G = (V, E)$, edge cost $c_e \geq 0$ for all $e \in E$, even degree bound $b_v \geq 0$ for all $v \in V$. The following is a natrual LP formulation for this problem. For each edge e, there is a variable x_e indicating whether e is in the solution or not. Note that in an optimal solution for the problem, we might need to pick an edge twice but not more than twice since, otherwise, one can reduce its occurrence by two and retain connectivity.

$$\text{minimize:} \quad \sum_{e \in E} c_e \cdot x_e \qquad \textbf{(BDTSP-LP)}$$

$$\text{subject to:} \quad x(\delta(S)) \geq 2 \quad \forall \, \emptyset \neq S \subsetneq V \tag{7}$$

$$x(\delta(v)) \leq b_v \quad \forall \, v \in V \tag{8}$$

$$0 \leq x_e \leq 2 \quad \forall \, e \in E \tag{9}$$

One can separate the constraints using a minimum-cut algorithm, so we can find an optimal solution (or determine (**BDTSP-LP**) is infeasible) in polynomial time. If the LP is infeasible, we report there is no feasible solution and terminate. Otherwise, we proceed as follows.

The algorithm is very similar to Wolsey's analysis of Christofides-Serdyukov algorithm. First we compute at spanning tree T using an optimal solution to (**BDTSP-LP**) and then we fix the degree parities using the $odd(T)$-join polytope. However, we need to respect (approximately) the degree bounds. For a node v, let $b'(v,T)$ be the smallest integer at least $\frac{b_v}{2}$ whose parity is the same as $|d_T(v)|$: note $b'(v,T) \in \{\frac{b_v}{2}, \frac{b_v}{2} + 1\}$.

Algorithm 1 $(3/2, +4)$-approximation algorithm for BDTSP.

Input: Graph $G = (V, E)$ with edge costs $c_e \geq 0$ for every $e \in E$ and even degree bounds b_v for every $v \in V$.
Output: A tour that spans V.

Let T be a Steiner tree (in this case spanning tree) obtained from applying Theorem 6 with input $G = (V, E)$, edge cost c_e for $e \in E$, $X := V$ and degree bounds $\frac{b_v}{2}$ for every $v \in V$.
Let $odd(T)$ be the set of vertices with odd degrees with respect to T. Compute an $odd(T)$-join J in G using Theorem 8 with degree bounds $b'(v, T)$ for all $v \in V$.
Output a closed spanning walk in $T \cup J$.

We show Algorithm 1 works correctly and this proves Theorem 1.

Proof of Theorem 1. Let x^* be an optimal solution to (**BDTSP-LP**). Note that $\bar{x} := \frac{x^*}{2}$ is feasible for (**SNDP-LP**) where $X := V$ and degree bounds $\frac{b_v}{2}$ for every $v \in V$. Hence, by Theorem 6 we have $cost(T) \leq 2 \cdot cost(\bar{x}) = cost(x^*) \leq opt$. Furthermore, $d_T(v) \leq \frac{b_v}{2} + 3$.

Consider vertex v in the graph, note that $\bar{x}(\delta(v)) \leq \frac{b_v}{2}$. By using degree bounds $b'(T, v)$ for $v \in V$, we ensure (**BD-Y-join LP**) is integral and that $\bar{x}$ is a feasible solution. Thus, a minimum-cost degree-bounded $odd(T)$-join J has cost at most $cost(\bar{x}) = cost(\frac{x^*}{2})$ and $d_J(v) \leq \frac{b_v}{2} + 1$ for every $v \in V$.

Putting the bounds on T and J together we have an Eulerian subgraph $T \cup J$ with cost $\frac{3}{2} \cdot cost(x^*)$ and $d_{T \cup J}(v) \leq b_v + 4$. ◀

4 Bounded-Degree Subset TSP

In this section, we prove our main results, i.e., Theorems 3, 4 and 5. Fix an instance of BDSTSP: $G = (V, E)$, a edge cost $c_e \geq 0$ for each $e \in E$, a set of terminals $X \subseteq V$, and an even integer degree bound $b_v \geq 0$ on each vertex $v \in V$. We refer to vertices in $V \setminus X$ as *Steiner* nodes.

The algorithm for BDSTSP is to find a "good" Steiner tree T that spans the terminals and then fix the degree parity using $odd(T)$-join. However, finding a "good" $odd(T)$-join is not as trivial as it was for BDTSP since $\frac{x^*}{2}$ may no longer be feasible for $odd(T)$-join polytope since some cuts S involving only Steiner nodes may have very low $x^*(\delta(S))$. Nevertheless, for all the approximation factors in this section, we show combining x^* and T itself with appropriate ratios is sufficient to construct a "good" (fractional) solution for the $odd(T)$-join polytope with degree constraints. We start with a simple application of this idea by proving Theorem 3.

We begin with the natural LP for BDSTSP. As before, we assume the LP has an optimal solution, otherwise the BTSTSP instance has no feasible solution.

$$\text{minimize:} \quad \sum_{e \in E} c_e \cdot x_e \qquad (\textbf{BDSTSP-Natural-LP})$$

$$\text{subject to:} \quad x(\delta(S)) \geq 2 \quad \forall S \neq X,\ S \cap X \neq \emptyset \qquad (10)$$

$$x(\delta(v)) \leq b_v \quad \forall v \in V \qquad (11)$$

$$0 \leq x_e \leq 2 \quad \forall e \in E \qquad (12)$$

Proof of Theorem 3. Let x^* be an optimal solution to (**BDSTSP-Natural-LP**). Since $\frac{x^*}{2}$ is feasible for (**SNDP-LP**) where degree bounds are $\frac{b_v}{2}$ for every $v \in V$, we can obtain a Steiner tree T of cost at most $2 \cdot cost(\frac{x^*}{2}) \leq opt$ and $d_T(v) \leq \frac{b_v}{2} + 3$, see Corollary 7. Furthermore, we iteratively prune leaf nodes that are Steiner nodes so all leaves of T are terminals. With abuse of notation, we denote the resulting tree by T.

Next, we show that $\bar{y} := \frac{\chi_T}{3} + \frac{x^*}{3}$ is feasible for (**BD-Y-join LP**) when $odd(T)$ is the set of odd degree vertices and the RHS of degree constraint is either $\frac{b_v}{2} + 1$ or $\frac{b_v}{2} + 2$ (whichever has the same parity as $d_T(v)$). Note that by definition of $\bar{y}$, and the fact that $d_T(v) \leq \frac{b_v}{2} + 3$, $\bar{y}$ respects the degree constraints in (**BD-Y-join LP**). Now consider a cut S that contains a terminal. Then T crosses the cut at least once and $x^*(\delta(S)) \geq 2$ so $\bar{y}(\delta(S)) \geq 1$.

Now consider an odd cut S, i.e. $|S \cap odd(T)| = odd$, that contains only Steiner nodes. Since $\sum_{v \in S} d(v) = 2 \cdot |E_T[S]| + |\delta_T(S)|$ and $\sum_{v \in V} d(v)$ is odd, we must have $|\delta_T(S)| = odd$. We claim that $|\delta_T(S)| > 1$, otherwise $|\delta_T(S)| = 1$ means S contains a leaf node which is impossible since (the pruned version) of T has only terminals as leaf nodes. Therefore, $|\delta_T(S)| \geq 3$ and by definition of $\bar{y}$ we have $\bar{y}(\delta(S)) \geq 1$, as desired.

So we have proved $\bar{y}$ is feasible for (**BD-Y-join LP**). By Theorem 8, there is an $odd(T)$-join J of cost at most $cost(\bar{y}) = \frac{1}{3} \cdot cost(T) + \frac{1}{3} \cdot cost(x^*) \leq \frac{2}{3} \cdot opt$ and $d_J(v) \leq \frac{b_v}{2} + 2$ for all $v \in V$. Finally we output a closed walk in subgraph $\mathcal{Q} := T \cup J$. Note $cost(\mathcal{Q}) \leq (1 + \frac{2}{3}) \cdot opt = \frac{5}{3} \cdot opt$ and $d_{\mathcal{Q}}(v) \leq b_v + 5$ for all $v \in V$. Since $\mathcal{Q}$ is an Eulerian graph and b_v's are even, it must be that $d_{\mathcal{Q}}(v) \leq b_v + 4$ for all $v \in V$. This finishes the proof of Theorem 3. ◀

To improve on the approximation factor of Theorem 3, i.e. the results in Theorem 4 & 5, we consider a slight strengthening of (**BDSTSP-Natural-LP**) for BDSTSP. We first make an observation about the structure of an optimal solution and then we add a constraint based on this observation.

We use the following definition throughout this section. Let v be a Steiner node, we say S is a v, X-cut if $v \in S \subseteq V \setminus X$.

▶ **Lemma 9.** *There exists an optimal solution $\mathcal{Q}^*$ such that for any Steiner node $\bar{v}$ and any $\bar{v}, X$-cut S, we have $|\delta_{\mathcal{Q}^*}(S)| \geq d_{\mathcal{Q}^*}(\bar{v})$.*

Proof. Among all optimal solutions, let $\mathcal{Q}^*$ be one with the minimum number of edges. For this proof, every degree or cut is with respect to $\mathcal{Q}^*$ unless stated otherwise. We show $|\delta_{\mathcal{Q}^*}(S)| \geq d_{\mathcal{Q}^*}(\bar{v})$ for every $\bar{v}$ and any $\bar{v}, X$-cut S. Suppose otherwise, that for $\bar{v}$ there is some $\bar{v}, X$-cut S with $|\delta_{\mathcal{Q}^*}(S)| < d_{\mathcal{Q}^*}(\bar{v})$. We take S to be a minimum-cardinality $\bar{v}, X$-cut.

Let $k := |\delta(S)|$. Note $\sum_{v \in S} d(v) = 2 \cdot |E_{\mathcal{Q}^*}[S]| + |\delta(S)|$ and since all vertices have even degree, k must be even. We say $u \in S$ is a *boundary node* with respect to S if $\delta(u) \cap \delta(S) \neq \emptyset$.

Contract $V \setminus S$ to a single vertex and call it t. Since S is a minimum cardinality $\bar{v}, t$-cut, there are k edge-disjoint simple paths $P_1, ..., P_k$ from $\bar{v}$ to t, in particular, for every boundary node $u \in S$, $|\delta(u) \cap \delta(S)|$ of the paths $P_1, ..., P_k$ have u as their second-last node (just before t).

Construct a graph G' obtained from $\mathcal{Q}^*$ as follows: remove all the edges in $E_{\mathcal{Q}^*}[S] \setminus \cup_{i=1}^k P_i$ which is non-empty as we assumed $d_{\mathcal{Q}^*}(v) > k$ and then remove all the isolated vertices. Note that $d_{G'}(\bar{v}) = k$ which is even. Also for every boundary vertex $u \in S$, $|\delta(u) \cap \delta(S)|$ of P_i's contains u, so the degree of boundary vertices is even. The vertices inside S except $\bar{v}$ and the boundary vertices are internal vertices of some paths so their degrees are even. Finally, degree of the vertices outside of S did not change, so their degree is even as well. Hence, G' is Eulerian.

We claim G' connects all the terminals, which would contradict the minimality of $\mathcal{Q}^*$. Hence, $|\delta_{\mathcal{Q}^*}(S)| \geq d_{\mathcal{Q}^*}(\bar{v})$, as desired. It is easy to prove the claim. Consider two terminals x and x'. If a $x - x'$ path in $\mathcal{Q}^*$ does not use any vertices in S then this path exists in $E(G')$. So suppose every $x - x'$ paths in $\mathcal{Q}^*$ crosses S, let u and u' (possibly $u = u'$) be the two boundary vertices (with respect to S) on a $x - x'$ path in $\mathcal{Q}^*$. Note that there is a path between u and $\bar{v}$, and a path between u' and $\bar{v}$ in $E(G')$. Therefore, x and x' are connected in $E(G')$. ◀

We use Lemma 9 to get a slightly stronger LP relaxation for BDSTSP.

$$\text{minimize:} \quad \sum_{e \in E} c_e \cdot x_e \qquad \text{(BDSTSP-LP)}$$

$$\text{subject to:} \quad x(\delta(S)) \geq 2 \qquad \forall S \neq X,\ S \cap X \neq \emptyset \tag{13}$$

$$x(\delta(v)) \leq b_v \qquad \forall v \in V \tag{14}$$

$$x(\delta(v)) \leq x(\delta(S)) \qquad \forall S \subseteq V \setminus X, \forall v \in S \tag{15}$$

$$0 \leq x_e \leq 2 \qquad \forall e \in E \tag{16}$$

Constraint (15) is valid because there is an optimal solution that has this property according to Lemma 9. Furthermore, this constraint can be separated by computing a minimum weight v, X-cut (with respect to weight x on the edges) for every Steiner node v. Let x^* be an optimal solution to this LP.

Now we are ready to discuss how to get a more well-structured Steiner "tree"[2] T using a slightly modified version of the algorithm in [11] for computing a degree-bounded Steiner tree, which in turn, results in a cheaper $odd(T)$-join.

The tweak in the algorithm of [11] for BDSTP is to completely "split off" a predetermined subset of Steiner nodes when the algorithm decides to drop the degree constraint corresponding to these Steiner nodes. This will ensure that we get a feasible solution for BDSTP (not necessarily a tree) such that all the Steiner nodes in the predetermined subset have even degree (counting with multiplicities). First we explain the splitting-off procedure and then present the tweak in the algorithm for BDSTP.

Splitting-off procedure. We begin with some definition. Fix a multigraph $G = (V, E)$. We say edge e is a *cut-edge* if there is a set S such that $\delta(S)$ contains only e. We denote the minimum cardinality of a u, v-cut by $\lambda_G(u, v)$. We say a pair of edges (su, sv) is an *splittable pair* if in $G' = \big(V \cup \{s\}, (E \setminus \{su, sv\}) \cup \{uv\}\big)$, we have $\lambda_G(u, v) = \lambda_{G'}(u, v)$ for all $u, v \in V$. In other words, removing su, sv and adding an edge between u and v preserves the minimum u, v-cut value for all $u, v \in V$. This process is called *splitting off* pair (su, sv). Recall a classical splitting-off result by Mader.

▶ **Theorem 10** (Mader, [14, 3]). *Let $G = (V \cup \{s\}, E)$ be a connected graph, perhaps with parallel edges. Assume there is no cut-edge incident to s and $d(s)$ is even. Then, there is a splittable pair (su, sv) for some u, v (possibly $u = v$) adjacent to s.*

When we remove all the splittable pairs so that there is no incident edge to s, we say we *completely split off* s. Since after applying Mader's theorem, the conditions of the theorem still holds for s, one can repeatedly apply Mader's theorem to completely split off s while preserving the local connectivities between any pair of vertices in V.

[2] We put tree in the quotation as we will see later that T might contain cycles to allow the degree parity of some Steiner nodes to be even, so T might not be a proper tree.

The following is the result that allows us to tweak the iterative rounding algorithm of [11] for BDSTP to ensure Steiner nodes with small fractional degree have even degree in the resulting tree. The proof uses Mader theorem as the subroutine. However, just applying Mader theorem repeatedly when the connectivities are based on weighted edges does not run in polynomial time. Here we use the idea used by Post and Swamy [15] to make the complete splitting off procedure polytime. In [15] they work with the directed version of Mader theorem. Although everything will be translated to our setting in a straightforward fashion, we present the proof in Appendix A for completeness.

▶ **Lemma 11.** *Let $\mathcal{I} = \big(G = (V, E), X, c, b\big)$ be an instance of the BDSTP and let $\bar{x}$ be a feasible solution for* (**SNDP-LP**) *of this instance. Let $s \in V$ be a Steiner node. Then, in polynomial time, one can obtain an instance of BDSTP $\mathcal{I}' = \big(G' = (V \setminus \{s\}, E'), X, c', b\big)$ and a feasible solution x' for* (**SNDP-LP**) *of this instance such that*

1. *$cost_{c'}(x') \leq cost_c(x)$.*
2. *An integral solution T' for $\mathcal{I}'$ can be transformed to an integral solution T for $\mathcal{I}$ whose cost is at most $cost_{c'}(T')$, $d_T(s)$ is even, and $d_T(v) = d_{T'}(v)$ for all $v \in V \setminus s$.*

Next, we modify one step of $(2, +3)$-approximation of [11] for BDSTP. For the sake of space, we only sketch the adaptation and defer the full description to Appendix B.

The iterative rounding algorithm for BDSTP. Consider an instance $\big(G = (V, E), X, c, b\big)$ of BDSTP and a subset of Steiner nodes $A \subseteq V \setminus X$ where $b_v = 1$ for all $v \in A$. The goal is to find a minimum cost solution T for this instance such that $d_T(v)$ is even for all $v \in A$ while violating the degree bounds by a small additive constant for these vertices, and $d_T(v) \leq b_v$ for all $v \in V \setminus A$ (see Lemma below). In the context of our BDSTSP, A will be the set of all Steiner nodes with "small" fractional degree (with respect to an optimal solution of (**BDSTSP-LP**)). Then, this more structured Steiner tree T helps us to find a cheaper $odd(T)$-join.

In the iterative rounding algorithm for BDSTP whenever we drop the degree constraints corresponding to a Steiner node v in A, we completely split off v as well. After the tree is constructed, we restore any edges produced by the splitting off procedure to the original set of edges and further prune some edges if necessary (i.e., if their multiplicity is > 2 after this restoration step)[3]. Again, see Appendix B for the full description of this algorithm.

Putting together Corollary 7 and Lemma 11 we have the following result for BDSTP. It is important to note that while we use the letter T to denote the solution, it is not necessarily a tree because we cannot necessarily drop edges from cycles while preserving the parity of nodes in A.

▶ **Lemma 12.** *Given an instance $\big(G = (V, E), X, c, b\big)$ of BDSTP, a feasible solution $\bar{x}$ of* (**SNDP-LP**) *and a set $A \subseteq V \setminus X$ where $b_v = 1$ for all $v \in A$, there is a polynomial time algorithm that computes a feasible solution T of BDSTP such that*

1. *$cost(T) \leq 2 \cdot cost(\bar{x})$.*
2. *$d_T(v) \leq b_v + 3$ for all $v \in V \setminus A$.*
3. *$d_T(v) \leq b_v + 7$ for all $v \in A$.*
4. *$d_T(v)$ is even for all $v \in A$.*
5. *$|\delta_T(S)|$ is even for any $S \subseteq A$.*

[3] Note we might have edges with multiplicity 2 in this solution but nevertheless we call the solution a Steiner tree.

6. *Let $T^{(m)}$ be a minimal (inclusion-wise) subset of edges of T such that $T^{(m)}$ is still a feasible solution for the BDSTP instance. Then, $T^{(m)}$ satisfies property 1 and $d_T(v) \leq b_v + 3$ for all $v \in V$.*

Proof. We run the $(2, +3)$-approximation of [11] with the above mentioned tweak, see Algorithm 3 and apply Lemma 11 (property 2) repeatedly to obtain a feasible solution T for BDSTP on G.

Properties 1 and 2 are trivial because of the approximation factor of Algorithm 3 and Lemma 11.

Property 3: we completely split off a Steiner node $s \in A$ in Algorithm 3 when the algorithm decides to drop the degree constraint corresponding to s (see step b in Algorithm 3). Hence, we must have $d(s) \leq b_s + 3 \leq 4$ in the current graph in that iteration. Once we apply Lemma 11 (property 2) to obtain a solution for the current instance of BDSTP (i.e., by putting s back) we might use these (up to) four edges incident to s multiple times; however, we do not need to use any edge more than twice (otherwise we can drop two copies of the edge and preserve both connectivity and parities) so the degree of s in the solution will be at most $8 = b_v + 7$ since $b_v = 1$ for $v \in A$. This proves property 3.

Property 4: the degree of a node $v \in A$ that was split off in our iterative rounding algorithm is even simply because each edge used in T that was produced in the splitting off procedure is then subdivided to re-integrate v so the degree of v is even. Further, any parallel edges that were pruned maintain the parity of the degree of v.

Property 5: for any set $S \subseteq V$ we have

$$\sum_{v \in S} d_T(v) = 2 \cdot |E_T[S]| + |\delta_T(S)|. \tag{17}$$

If $S \subseteq A$, by property 4 the LHS of (17) is even and therefore $|\delta_T(S)|$ must be even.

Property 6: by property 2, for all $v \in V \setminus A$ we have $d_T(v) \leq b_v + 3$. In $T^{(m)}$ we keep only one copy of each parallel edge so $d_T(v) \leq b_v + 3$ for all $v \in A$ (see the argument for property 3) and the resulting solution is still feasible which implies the last property for $T^{(m)}$. ◀

Finally, we can present our $(13/8, +6)$-approximation for BDSTSP, see Algorithm 2.

Algorithm 2 $(13/8, +6)$-approximation algorithm for BDSTSP.

Input: Graph $\big(G = (V, E), X, c, b\big)$.
Output: A closed walk $\mathcal{Q}$ in G that spans X.

(a) Compute an optimal solution x^* of (**BDSTSP-LP**). Let $A := \{v \in V \setminus X : x^*(\delta(v)) < 2\}$, and let $b'_v := \frac{b_v}{2}$ if $v \notin A$ and $b'_v := 1$ if $v \in A$.
(b) Apply Lemma 12 with instance $\big(G = (V, E), X, c, b'\big)$, feasible solution $\bar{x} := \frac{x^*}{2}$, and set A to obtain a Steiner tree T with properties 1-5 in the lemma.
(c) Compute a Steiner tree $T^{(m)}$ using T according to property 6 of Lemma 12.
(d) Compute a minimum cost $odd(T^{(m)})$-join J such that $d_J(v) \leq b'_v + 3$ for all $v \in V$ (cf. Lemma 13).
(e) Output a closed walk $\mathcal{Q}$ in $T^{(m)} \cup J$ that spans all the terminals.

Analysis. For the rest of this section, let x^* be an optimal solution for (**BDSTSP-LP**) computed in step (a) and $\bar{x} := \frac{x^*}{2}$. Let A be the subset of Steiner nodes and b'_v's the degree bounds constructed based on x^* in step (a). Also let T and $T^{(m)}$ be the Steiner trees computed in steps (b) and (c) of the algorithm, respectively. Note that $\bar{x}$ is feasible for (**SNDP-LP**) when the degree bounds are according to b'. Combining this fact and Lemma 12 we have:

$$cost(T^{(m)}) \leq cost(T) \leq 2 \cdot cost(\bar{x}) = cost(x^*). \tag{18}$$

▶ **Lemma 13.** *There is an $odd(T^{(m)})$-join J with cost at most $\frac{5}{8} \cdot opt$ and $d_J(v) \leq \frac{b_v}{2} + 3$ for all $v \in V$.*

Proof. We claim $y := \frac{\chi_T}{4} + \frac{3 \cdot x^*}{8}$ is feasible for (**BD-Y-join LP**) when the set of odd degree vertices is $odd(T^{(m)})$ and the RHS of degree constraint is at most $b'_v + 3$ for all $v \in V$ (in fact this fractional solution is feasible when degree bounds are $\frac{b_v}{2} + 2$ but similar to the proof of Theorems 1 and 3, we might need to consider $\frac{b_v}{2} + 3$ for some vertices as (**BD-Y-join LP**) is integral if and only if the parity of the degree bounds match the parity of the degree of vertices in an integral solution). We prove the claim after we show how the lemma follows from this claim. By Theorem 8, there is an integral $odd(T^{(m)})$-join J whose cost is at most

$$cost(y) = cost(\frac{1}{4} \cdot \chi_T) + cost(\frac{3}{8} \cdot x^*) \leq \frac{1}{4} \cdot cost(x^*) + \frac{3}{8} \cdot cost(x^*) \leq \frac{5}{8} \cdot opt,$$

where the first inequality follows from (18). Furthermore, $d_J(v) \leq b'_v + 3 \leq \frac{b_v}{2} + 3$.

So it remains to prove the claim. As (19) shows, y satisfies the degree constraints of (**BD-Y-join LP**) when the degree bound is at most $b'_v + 3$ for all $v \in V$.

$$y(\delta(v)) = \frac{1}{4} \cdot d_T(v) + \frac{3}{8} \cdot x^*(\delta(v)) \leq \frac{1}{4} \cdot (b'_v + 7) + \frac{3}{8} \cdot b_v \leq \frac{b_v}{2} + 3, \tag{19}$$

where the first inequality follows from properties 2 & 3 of Lemma 12.

Next we show cut constraints (5) in (**BD-Y-join LP**) hold under solution y. Consider a subset $S \subseteq V$ such that $|S \cap odd(T^{(m)})| = odd$. There are three cases to consider:

- Case 1: If $S \cap X \neq \emptyset$. Then, $x^*(\delta(S)) \geq 2$ and since T is connected we have $|\delta_T(S)| \geq 1$ which implies $y(\delta(S)) \geq 1$.
- Case 2: If $S \cap X = \emptyset$ and $S \cap (V \setminus A) \neq \emptyset$. Then, there is a Steiner node $s \in S$ such that $s \notin A$. By definition of set A we have $x^*(\delta(s)) \geq 2$. By constraint (15) in (**BDSTSP-LP**) this implies $x^*(\delta(S)) \geq 2$ as well. Again since T is a connected so $|\delta_T(S)| \geq 1$ and this implies $y(\delta(S)) \geq 1$.
- Case 3: If $S \subseteq A$. Since (17) holds for any subset of vertices and $|S \cap odd(T^{(m)})| = odd$, the LHS of (17) is odd and so is $|\delta_{T^{(m)}}(S)|$. Furthermore, if $|\delta_{T^{(m)}}(S)| = 1$ then we can remove S from $T^{(m)}$ and still have a feasible solution, contradicting the inclusion-wise minimality of $T^{(m)}$. Hence $|\delta_{T^{(m)}}(S)| \geq 3$. So we have

  $$|\delta_T(S)| \geq |\delta_{T^{(m)}}(S)| \geq 3.$$

  On the other hand, since $S \subseteq A$, by property 5 Lemma 12 we have $|\delta_T(S)| = even$; together with above inequality we get $|\delta_T(S)| \geq 4$. Therefore, $y(\delta(S)) \geq 1$ in this case as well. ◀

Now the proof of Theorem 4 follows easily.

Proof of Theorem 4. Since $T^{(m)} \cup J$ is an Eulerian subgraph, there is a closed walk Q in it that spans X. By (18) we have $cost(T^{(m)}) \leq cost(x^*)$. By Lemma 13 $cost(J) \leq \frac{5}{8} \cdot opt$ and this implies $cost(T^{(m)} \cup J) \leq (1 + \frac{5}{8}) \cdot opt = \frac{13}{8} \cdot opt$, as desired.

By property 6 of Lemma 12, $d_{T^{(m)}}(v) \leq b'_v + 3 \leq \frac{b_v}{2} + 3$ and by Lemma 13 we have $d_J(v) \leq \frac{b_v}{2} + 3$. So $d_{T^{(m)} \cup J}(v) \leq b_v + 6$ for all $v \in V$, as desired. ◀

To prove Theorem 5, we modify Algorithm 2 slightly as follows: remove step (c), in step (d) compute a minimum cost $odd(T)$-join J (instead of $odd(T^{(m)})$-join) and in step (e) output a closed walk in $T \cup J$ that spans all the terminals. Similar to Lemma 13, we have the following bound on the bounded-degree $odd(T)$-join LP.

▶ **Lemma 14.** *There is an $odd(T)$-join J with cost at most $\frac{opt}{2}$ and $d_J(v) \leq \frac{b_v}{2} + 1$ for all $v \in V$.*

Proof. We just need to show $\frac{x^*}{2}$ is feasible for degree-bounded $odd(T)$-join polytope when the degree bound is $\frac{b_v}{2}$ for all $v \in V$. Then, whenever necessary we replace $\frac{b_v}{2}$ by $\frac{b_v}{2} + 1$ for $v \in V$ so that Theorem 8 to be applicable. It is trivial that the degree constraints hold. So we show constraints (5) in (**BD-Y-join LP**) holds. Let $S \subsetneq V$ such that $|S \cap odd(T)| = odd$.

- Case 1: If $S \cap X \neq \emptyset$. Then, $x^*(\delta(S)) \geq 2$ which implies $\frac{x^*}{2}(\delta(S)) \geq 1$.
- Case 2: If $S \cap X = \emptyset$ and $S \cap (V \setminus A) \neq \emptyset$. Then, there is a Steiner node $s \in S$ such that $s \notin A$. By definition of set A we have $x^*(\delta(s)) \geq 2$. By constraint (15) in (**BDSTSP-LP**) this implies $x^*(\delta(S)) \geq 2$ as well which implies $\frac{x^*}{2}(\delta(S)) \geq 1$.
- Case 3: If $S \subseteq A$. Thus, we have $S \subseteq A$ and note that $d_T(v)$ is even for all $v \in A$ by property 4 of Lemma 12. This is a contradiction because we assumed $|S \cap odd(T)|$ is odd. So the constraints for $S \subseteq A$ are not present in the LP. ◀

Proof of Theorem 5. By (18) we have $cost(T) \leq cost(x^*) \leq opt$ and by Lemma 14 we have $cost(J) \leq \frac{cost(x^*)}{2} \leq \frac{opt}{2}$ which implies $cost(T \cup J) \leq \frac{3}{2} \cdot opt$.

By properties 2 and 3 of Lemma 12, we have $d_T(v) \leq \frac{b_v}{2} + 7$ and by Lemma 14 we have $d_J(v) \leq \frac{b_v}{2} + 1$ for all $v \in V$. Thus, $d_{T \cup J}(v) \leq b_v + 8$, as desired. ◀

5 Conclusion

We gave a $(5/3, +4)$, a $(13/8, +6)$ and a $(3/2, +8)$ approximations for BDSTSP. It would be interesting to see if there is any $O(1)$-approximation that violates the degree bounds by at most $+2$. On the other hand, a demonstration that any LP-based $O(1)$-approximation requires a $+4$ violation would be interesting as well. Though, we suspect an $O(1)$-approximation with $+2$ violation is possible. Further improvements to our $\frac{13}{8}$-approximation for low violations (e.g. $+4$ or $+6$) would also be interesting. For example, perhaps a best-of-many Christofides approach could help to show that, on average, the Y-join can be constructed more cheaply than in our analysis. As a starting point, we note it is possible to extend the bounded-degree Steiner tree iterated rounding result to show that for feasible BDST solution x^*, $2 \cdot x^*$ dominates a convex combination of trees each of which violates degree bounds by $+3$.

Finally, we did not consider edge bounds b_e because if we even allow a $+1$ violation it becomes quite trivial: no BDSTSP solution will use any edge more than twice because we could remove two copies of any edge used more than twice. So what about satisfying edge bounds exactly? Obviously cut edges in the original graph that have Steiner nodes on both sides of the cut must be used twice. Even if there are no cut edges it is still hard to find a spanning Eulerian tour that uses each edge at most once: This would still model the Hamiltonian cycle problem in cubic graphs. Still, is there some interesting extent to which we could bound the use of some edges and still expect some approximation algorithm to respect most of these bounds?

References

1 Hyung-Chan An, Robert Kleinberg, and David B Shmoys. Improving christofides' algorithm for the st path tsp. *Journal of the ACM (JACM)*, 62(5):1–28, 2015.
2 Nicos Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical report, Carnegie-Mellon Univ Pittsburgh Pa Management Sciences Research Group, 1976.
3 András Frank. On a theorem of mader. *Discret. Math.*, 101:49–57, 1992.
4 Takuro Fukunaga, Zeev Nutov, and R Ravi. Iterative rounding approximation algorithms for degree-bounded node-connectivity network design. *SIAM Journal on Computing*, 44(5):1202–1229, 2015.
5 Michel X Goemans. Minimum bounded degree spanning trees. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 273–282. IEEE, 2006.
6 Anna R. Karlin, Nathan Klein, and Shayan Oveis Gharan. A (slightly) improved approximation algorithm for metric tsp. *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, 2021.
7 Marek Karpinski, Michael Lampis, and Richard Schmied. New inapproximability bounds for tsp. *Journal of Computer and System Sciences*, 81(8):1665–1677, 2015.
8 Rohit Khandekar, Guy Kortsarz, and Zeev Nutov. On some network design problems with degree constraints. *Journal of Computer and System Sciences*, 79(5):725–736, 2013.
9 Lap Chi Lau, Joseph Naor, Mohammad R Salavatipour, and Mohit Singh. Survivable network design with degree or order constraints. *SIAM Journal on Computing*, 39(3):1062–1087, 2009.
10 Lap Chi Lau, Ramamoorthi Ravi, and Mohit Singh. *Iterative methods in combinatorial optimization*, volume 46. Cambridge University Press, 2011.
11 Lap Chi Lau and Mohit Singh. Additive approximation for bounded degree survivable network design. *SIAM Journal on Computing*, 42(6):2217–2242, 2013.
12 Lap Chi Lau and Hong Zhou. A unified algorithm for degree bounded survivable network design. *Mathematical Programming*, 154(1):515–532, 2015.
13 Anand Louis and Nisheeth K Vishnoi. Improved algorithm for degree bounded survivable network design problem. In *Scandinavian Workshop on Algorithm Theory*, pages 408–419. Springer, 2010.
14 Wolfgang Mader. A reduction method for edge-connectivity in graphs. *Annals of discrete mathematics*, 3:145–164, 1978.
15 Ian Post and Chaitanya Swamy. Linear programming-based approximation algorithms for multi-vehicle minimum latency problems. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 512–531. SIAM, 2014.
16 Alexander Schrijver. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer Science & Business Media, 2003.
17 AI Serdyukov. O nekotorykh ekstremal'nykh obkhodakh v grafakh. *Upravlyayemyye sistemy*, 17:76–79, 1978.
18 Mohit Singh and Lap Chi Lau. Approximating minimum bounded degree spanning trees to within one of optimal. *Journal of the ACM (JACM)*, 62(1):1–19, 2015.
19 Laurence A Wolsey. Heuristic analysis, linear programming and branch and bound. In *Combinatorial Optimization II*, pages 121–134. Springer, 1980.

A Proof of Lemma 11

Here we present the proof of Lemma 11.

Recall $\bar{x}$ is an optimal solution to (**SNDP-LP**) and s is a Steiner node such that $\bar{x}(\delta(s)) \leq 1$. First we show how to obtain x' and c' with an easy application of Mader's theorem but the running time of this procedure might be exponential and then at the end we show how to turn this procedure to run in polynomial time.

Since $\bar{x}$ is rational and the number bits needed to represent it is polynomial in the size of the input, there is a positive integer Δ such that $\Delta \cdot \bar{x}_e$ is an even integer for all $e \in E$. We replace each edge e, with $\Delta \cdot x_e$ copies of parallel edges. Note that degree of s is even and there is no cut-edge in the graph. By applying Mader's theorem repeatedly, we can split off s completely. Denote the resulting graph by $G' = (V \setminus \{s\}, E')$. Finally, we define a new solution x' for the resulting graph as follows: $x'_e = \frac{\#\ copies\ of\ e}{\Delta}$ for all $e \in E'$. Note that by construction, x' respects the degree bounds and we preserve the connectivity between each pairs of nodes (except s), hence x' is feasible for (**SNDP-LP**) corresponding to G'. Next, we define the cost function c'. It is defined naturally, i.e., for the existed edges in G, c and c' agree with each other and for a new introduced edge uv we set $c'_{uv} := c_{su} + c_{sv}$. Let $\mathcal{I}' := (G', X, c', b)$ be the resulting instance.

We show that $cost_{c'}(x') \leq cost_c(\bar{x})$. This follows from the fact that if we introduce a new edge uv, then its cost is $c_{su} + c_{sv}$ and we decrease the LP weight on su and sv by the same amount we increase the LP value on uv. If uv exists in G we claim that $c_{su} + c_{sv} > c_{uv}$ because otherwise one can increase the LP value on su and sv and decrease the LP value on uv by the same amount and get a cheaper feasible solution, contradicting the fact that $\bar{x}$ is optimal.

Finally, for any integral solution T' for $\mathcal{I}'$ we construct an integral solution T for $\mathcal{I}$ by replacing every edge uv in T' that is not in G with the two edges su and sv (keep the edges with multiplicities). Note we might have more than 2 copies of an edge e incident to s in T. In this case, we reduce the occurrences of e as much as possible so that the resulting solution is feasible for $\mathcal{I}$ AND the parity of the degree of the endpoints of e does not change. By definition of c' we have $cost_c(T) \leq cost_{c'}(T')$.

Here we show by a straightforward adaptation of techniques in [15], we can construct x' and c' that satisfy the property of the lemma efficiently.

Let $\bar{x}$ be the LP weight on the edges. Splitting off a pair (su, sv) to the extent of $\alpha > 0$ means reducing $\bar{x}_{sv}$ and $\bar{x}_{sv}$ by α, and increasing $\bar{x}_{uv}$ by α such that all the connectivities between any pair of nodes (except pairs involving s) are preserved (the connectivity between two nodes u and v is defined as the minimum weight $u - v$ cut when edges have weight according to $\bar{x}$). We say we split of (su, sv) to the *maximum extent* if value α above is the maximum value possible.

Note that by the first procedure we explained earlier, there exists always a pair (su, sv) and value $\alpha > 0$ such that we can split of (su, sv) to the extent of α. We can find such pairs by brute force ($O(n^2)$ pairs) and find α by binary search. Note that it is possible $u = v$ which in that case it means reducing x_{su} by α.

The algorithm is simple, we find a pair of edges (su, sv) and a value $\alpha > 0$ and split off (su, sv) to the extent of α and repeat this procedure.

In the next claim, we show if we split off a pair of edges to the maximum extent, that pair never becomes splittable again which in turn implies a polynomial running time of above procedure.

▷ Claim 15. Consider an splittable pair (su, sv) according to Mader theorem (Theorem 10). If we split off (su, sv) to the maximum extent, then (su, sv) will not become splittable again.

Proof of Claim 15. This follows from Claim 3.1 in [3] which states the following:

▷ Claim 16 (Claim 3.1 in [3]). A pair (su, sv) is splittable if and only if there is no set Y such that $u, v \in Y$, $s \notin Y$, and there are two nodes $w \in Y$, $s \neq z \notin Y$ such that $|\delta_G(Y)| \leq \lambda_G(w, z) + 1$.

Multiply $\bar{x}$ by a suitable integer Δ such that $\Delta \cdot \bar{x}_e$ is even for all $e \in E$. Replace each edge e by $\Delta \cdot \bar{x}_e$ many parallel edges in G.

Suppose we split as much copy of (su, sv) as possible. Then, there must be a set Y that satisfies the properties of Claim 16. Now assume we split of a pair of edges (e, f) incident to s and let G' be the resulting graph. We show that Y still satisfies the properties of Claim 16; hence, non of the copy of (su, sv) are splittable after splitting off (e, f).

Note $s \notin Y$ so $|\delta_{G'}(Y)| \leq |\delta_G(Y)|$. Furthermore, since (e, f) is a splittable pair we have $\lambda_{G'}(w, z) = \lambda_G(w, z)$. Therefore, $|\delta_{G'}(Y)| \leq |\delta_G(Y)| \leq \lambda_G(w, z) = \lambda_{G'}(w, z) + 1$. Hence, Y satisfies the properties of Claim 16 in G' as well so (e, f) is not splittable in G'. $\lhd$

B Iterative Rounding Algorithm for BDSTP

Algorithm 3 Iterative rounding algorithm of [11] for BDTSP with small change in step (b).

Input: Graph $\big(G = (V, E), X, c, b\big)$, a subset of Steiner nodes A where $b_v = 1$ for all $v \in A$.
Output: A connected subgraph T of G that spans X such that $d_T(v) \leq b_v + 7$ is even for all $v \in A$ and $d_T(v) \leq b_v + 3$ for all $v \in V \setminus A$.

Initialize $T' \leftarrow \emptyset$ and $W \leftarrow V$ (W is the set of vertices with degree constraints present in the LP formulation).

while T' is not a feasible solution for BDSTP **do**

(a) Compute an optimal extreme point solution x of (**SNDP-LP**) and remove every edge e with $x_e = 0$.

(**b, with modification**). Removing a degree constraint: For every $v \in W$ with degree at most $b_v + 3$ in G, remove v from W. Furthermore, if $v \in A$, completely split off v and compute an optimal solution for the resulting instance (cf. Lemma 11). Redefine G to be this new graph and x to be the new optimal solution after splitting off procedure.

(c) Picking 1-edge: For each edge $e = uv$ with $x_e = 1$, add e to T, remove e from G, and decrease b_u, b_v by 1.

(d) Picking a heavy edge with no degree constraints: For each edge $e = uv$ with $x_e \geq \frac{1}{2}$ and $u, v \notin W$, add e to T' and remove e from G.

Updating the connectivity requirements: For every set $S \neq X$, $S \cap X \neq \emptyset$ update the RHS of constraint (2) in (**SNDP-LP**) to be $\max\{1 - |\delta'_T(S)|, 0\}$.

Let T be the resulting Steiner tree obtained from T' by applying Lemma 11 repeatedly.

Note that the algorithm works correctly for any subset A of Steiner nodes (i.e., A does not need to contain only Steiner nodes v with $b_v = 1$). The performance guarantee of this algorithm follows from [11] and the fact that after splitting off a Steiner node, by Lemma 11 we still have a feasible solution for the (**SNDP-LP**) of the resulting instance with cost at most the cost of the original LP. So T' is a Steiner tree for the instance $\big(G' = (V \setminus A, E'), X, c', b\big)$. The fact that the Steiner tree T for the original instance obtained from T' has the desired properties of Lemma 12 follows from Lemma 11.

C Iterative Rounding Algorithm for BD Y-join

In this section, we prove Theorem 8 using an iterative rounding algorithm. The algorithm and its analysis are a simple adaptation of the iterative rounding algorithm for maximum matching in [10].

Algorithm 4 An iterative algorithm for bounded-degree Y-join problem.

Input: Undirected graph $G = (V, E)$ with edge costs $c_e \geq 0, e \in E$, degree bounds $b_v, v \in V$, and $Y \subseteq V$ where $|Y| = even$.
Output: A Y-join.

$J \leftarrow \emptyset$
while $E \neq \emptyset$ **do**
  Find an optimal extreme point solution x to **BD-Y-join LP** defined on $G = (V, E)$.
  Find an edge $e = uv$ with either $x_e = 0$ or $x_e = 1$ (cf. Lemma 18). In the former case remove e, and in the latter case add e to J and do the following:
  Update $Y \leftarrow Y \Delta \{u, v\}$. Update $E \leftarrow E \setminus \{e\}$. Update $b(v) \leftarrow b(v) - 1$, and $b(u) \leftarrow b(u) - 1$.
output J.

We only need to show that there is always an edge with LP value 0 or 1. This shows the LP is integral. To do this, we state some properties of an extreme point solution to **BD-Y-join LP** which follows from rank lemma (Lemma 2.1.4 in [10]) and standard uncrossing techniques.

▶ **Lemma 17** (Properties of an extreme point). *Consider an extreme point x and suppose $0 < x_e < 1$ for all $e \in E$. Then, there exists a laminar family $\mathcal{L}$ of Y-odd sets and $W \subseteq V$ such that*

(i) *For any $S \in \mathcal{L}$ we have $x(\delta(S)) = 1$ and for any $v \in W$ we have $x(\delta(v)) = b(v)$.*
(ii) *The vectors in $\{\chi(\delta(S)) : S \in \mathcal{L}\} \cup \{\chi(\delta(v)) : v \in W\}$ are linearly independent.*
(iii) $|\mathcal{L}| + |W| = |E|$.
(iv) *$E[S]$ is connected for each $S \in \mathcal{L}$.*

Proof. Properties (i), (ii), and (iii) follow from a standard application of uncrossing technique and Rank Lemma in polyhedral theory.

We show that one can further modify the laminar family to obtain (iv). Consider $S \in \mathcal{L}$ and suppose $E[S]$ is not connected. Since $|S \cap Y| = odd$, there must be a connected component C of $E[S]$ such that $|C \cap Y| = odd$. Since $x(\delta(C)) \geq 1$ and the fact that S is a tight set we have $\chi(\delta(C)) = \chi(\delta(S))$ and C is a tight set. By property (ii) $C \notin \mathcal{L}$. Now consider the laminar family $(\mathcal{L} \setminus \{S\}) \cup \{C\}$. Repeating this procedure until there is no set in the laminar family that violates (iv) finishes the proof. ◀

▶ **Lemma 18.** *Given any extreme point x of* **BD-Y-join LP** *there must exist an edge e with $x_e = 0$ or $x_e = 1$.*

Proof. Suppose not. So we have $0 < x_e < 1$ for each $e \in E$. Let $\mathcal{L}$ be a laminar family and let W be a subset of V that satisfy properties (i)-(iv) in Lemma 17. We show a contradiction with property (iii) using a token-based argument. Let $\mathcal{L}' := \mathcal{L} \cup W$ be the extended laminar family. We assign one token to each edge in the support of x. Then we distribute the tokens inductively among the sets in the laminar family such that each member of $\mathcal{L}'$ receives at least one token and we show there are some extra tokens left which shows the contradiction with the fact that $|E| = |\mathcal{L}'|$.

We use the natural forest structure that the laminar family $\mathcal{L}'$ imposes, recall that each component of this forest is a rooted tree. We use the following claim to redistribute the tokens of $E[S]$ among the laminar sets inside S. We use G_S to denote the graph after contracting the children of S in $G[S]$. We label the contracted vertices with the same label as the contracted set, i.e., if R_i is a child of S then we call the contracted vertex of R_i by R_i as well.

▷ Claim 19. For any rooted subtree of $\mathcal{L}'$ with root S, the tokens of edges in $E[S]$ can be distributed such that for every $S' \in \mathcal{L}'$ where $S' \subsetneq S$ receives at least one token.

Proof. We prove this by induction. If $S \in W$ the claim holds vacuously. So let $S \in \mathcal{L}$ and let $R_1, ..., R_k$ be the children of S. We call a child of S *good* if a token is already (before considering S) assigned to it and call it *bad* otherwise. Note that none of the tokens in $E(G_S)$ is assigned to any sets yet because of the way the inductive procedure works.

By property (iv) of Lemma 17, G_S is a connected graph and if it is not a tree then $E(G_S) \geq k$ and we can assign the tokens of these (at least) k edges in $E(G_S)$ to $R_1, ..., R_k$ and we are done. So suppose G_S is a tree. If one of $R_1, ..., R_k$ is a good child then again we have enough tokens in $E(G_S)$ to assign to bad children of S. So from now on we assume all the children of S are bad.

Since $|S \cap Y| = odd$ and the fact that if $x(\delta(R_i)) = odd$ then $|R_i \cap Y| = odd$, we conclude the number of children of S that have an odd fractional degree is odd. Therefore, $\sum_{i=1}^{k} x(\delta(R_i)) = odd.$

Since G_S is a tree so it is a bipartite graph. Let V_1 and V_2 be the two parts of G_S and since $\sum_{i=1}^{k} x(\delta(R_i)) = odd$, wlog, we can assume $\sum_{R_i \in V_1} x(\delta(R_i)) \leq \frac{\sum_{i=1}^{k} x(\delta(R_i))-1}{2}$. From this inequality and the fact that $x(\delta(S)) = 1$ we get

$$\frac{\sum_{i=1}^{k} x(\delta(R_i)) - 1}{2} = x(E(G_S)) \leq \sum_{R_i \in V_1} x(\delta(R_i)) \leq \frac{\sum_{i=1}^{k} x(\delta(R_i)) - 1}{2}. \tag{20}$$

So all the inequalities in (20) must hold as equality. Therefore, $x(E(G_S)) = \frac{\sum_{i=1}^{k} x(\delta(R_i))-1}{2}$ which implies $\chi(\delta(S)) = \sum_{R_i \in V_2} \chi(\delta(R_i)) - \sum_{R_i \in V_1} \chi(\delta(R_i))$ and this contradicts the linear independence of the characteristic vectors in $\mathcal{L}'$. This finishes the proof of the claim. ◁

We continue the proof of Lemma 18. By Claim 19 every non-root vertices of the forest (obtained from $\mathcal{L}'$) is assigned a token. Let $S_1, ..., S_k$ be the root nodes of the forest. Denote by G_R the graph obtained from contracting S_i's. Note that none of the tokens of edges in $E(G_R)$ is assigned to any member of $\mathcal{L}'$ by our token assignment given in Claim 19. Since $x_e < 1$, $|\delta(S_i)| \geq 2$. We show that at least one root node has degree at least 3. But first let us show if this holds then the lemma follows. Since one non-root node has degree at least 3, we have $|E(G_R)| \geq k + 1$ which implies we can assign one token to each root node and still have at least one token left unassigned. This implies $|E(G)| > |\mathcal{L}'|$, a contradiction with property (iii) of Lemma 17. Therefore, there must be an edge e with $x_e = 0$ or $x_e = 1$.

Now it remains to prove at least one root node has degree at least 3. Suppose not. Then, every vertex in G_R has degree 2 and so G_R is a collection of cycles. Consider a cycle component C of G_R. Note $|C|$ cannot be odd otherwise there should be at lest fractionally one edge going outside of the nodes in C. So let $S_1, ..., S_l$ be the vertices of C and note l is even. Then we have $\sum_{i=1}^{l}(-1)^i \chi(\delta(S_i)) = 0$ which contradicts the linear independence of characteristic vectors in $\mathcal{L}'$.

We finish the proof by showing G_R itself cannot be a cycle. Note $x(\delta(S_i)) = 1$ for $1 \leq i \leq k$ which implies $|S_i \cap Y| = odd$ and so the number of vertices in G_R is even. And if G_R forms a cycle, then the same argument as the previous paragraph shows a linear dependence of characteristic vectors in $\mathcal{L}'$. ◀