

On Approximability of Steiner Tree in ℓ_p -metrics

Received Oct 11, 2023
Revised Sep 12, 2024
Accepted Dec 3, 2024
Published Jan 20, 2025

Key words and phrases
Hardness of Approximation,
Steiner Tree, Graph Coloring, Set
Cover

Henry Fleischmann^a ✉ 

Surya Teja Gavva^b ✉ 

Karthik C. S.^b ✉ 

^a Carnegie Mellon University, USA

^b Rutgers University, USA

ABSTRACT. In the Continuous Steiner Tree problem (CST), we are given as input a set of points (called *terminals*) in a metric space and asked for the minimum-cost tree connecting them. Additional points (called *Steiner points*) from the metric space can be introduced as nodes in the solution. In the Discrete Steiner Tree problem (DST), we are given in addition to the terminals, a set of facilities, and any solution tree connecting the terminals can only contain the Steiner points from this set of facilities.

Trevisan [SICOMP'00] showed that CST and DST are APX-hard when the input lies in the ℓ_1 -metric (and Hamming metric). Chlebík and Chlebíková [TCS'08] showed that DST is NP-hard to approximate to factor of $96/95 \approx 1.01$ in the graph metric (and consequently ℓ_∞ -metric). Prior to this work, it was unclear if CST and DST are APX-hard in essentially every other popular metric.

In this work, we prove that DST is APX-hard in every ℓ_p -metric. We also prove that CST is APX-hard in the ℓ_∞ -metric. Finally, we relate CST and DST, by observing a gap preserving reduction from CST to DST in ℓ_p -metrics.

It is known that the APX-hardness of DST in ℓ_0 , ℓ_1 , and ℓ_∞ -metrics can be obtained from the APX-hardness of covering problems (with additional structure). Our main conceptual insight is that for certain ranges of p (such as $p = 2$), the soundness guarantees of covering problems might be insufficient to show that DST in the ℓ_p -metric is APX-hard, but the soundness guarantees of a packing problem (with requisite additional structure) is enough. Equipped with this insight,

Henry Fleischmann's work was supported by NSF grant CNS-2150186 and an NSF Graduate Research Fellowship under Grant No. 2140739. Part of this research was conducted while he was affiliated with the University of Michigan. The work of Karthik C. S. was supported by Rutgers University's Research Council Individual Fulcrum Award (#AWD00010234), by the National Science Foundation under Grant CCF-2313372, and a grant from the Simons Foundation, Grant Number 825876, Awardee Thu D. Nguyen. A preliminary version of this article appeared at SODA 2024 [28].

we are then able to embed set systems into every ℓ_p -metric, and, depending on the value of p , our soundness analysis of the corresponding DST instance in the ℓ_p -metric will use either the packing property or the covering property (or both) of the set system in the soundness case.

Due to the discrete structure of the Hamming space, the APX-hardness of CST in the ℓ_0 and ℓ_1 metrics follow from similar methods. However, these techniques do not extend to CST in other ℓ_p -metrics. To show APX-hardness of CST in the ℓ_∞ -metric, we instead rely on the robust hardness guarantees of graph coloring problems. Concretely, we present a reduction from a graph G on n vertices to a point-set P in the ℓ_∞ -metric space, where the cost of the optimal Steiner tree of P is exactly $\frac{n+\chi(G)}{2}$, where $\chi(G)$ is the chromatic number of G .

1. Introduction

Given a set of points (called *terminals*) in a metric space, a *Steiner tree* is defined to be a minimum length tree connecting them, with the tree possibly including additional points (called *Steiner points*) as nodes. Computing the Steiner tree is one of the classic and most fundamental problems in Computer Science, Combinatorial Optimization, and Operations Research, with both great theoretical and practical relevance and interest [45]. This problem emerges in a number of contexts, such as network design problems, design of integrated circuits, location problems [14, 17, 35, 43, 52, 47] and more recently even in machine learning, systems biology, and bioinformatics [5, 36, 55, 61]. For example, in VLSI circuits, wire routing is carried out by wires running only in vertical and horizontal directions, due to high computational complexity of the task. Therefore, the wire length is the sum of the lengths of vertical and horizontal segments, and the distance between two pins of a net is actually the ℓ_1 -metric distance between the corresponding geometric points in the design plane [57].

Formally, the Steiner tree problem can be formulated in two ways: discrete and continuous. In the Discrete Steiner Tree problem (DST), we are given as input two sets of points in a metric space, called *terminals* and *facilities* respectively. The goal is to find the minimum-cost tree connecting the terminals, possibly introducing new points (*Steiner points*) from the set of facilities as nodes in the solution. This problem is well-defined even in general metric spaces since the relevant metric space can be fully described as part of the input. In the Continuous Steiner Tree problem (CST), we are only given a set of terminals as input, and we are allowed to use any point in the metric space as a Steiner point in the Steiner tree connecting all the terminals. For CST, we assume knowledge of the ambient metric space structure (e.g., $(\{0, 1\}^n, \ell_1)$ or (\mathbb{R}^n, ℓ_2)). This is not part of the input.

General Metrics. In general metrics, where the metric space is specified as part of the input (for example, by specifying all pairwise distances), it does not make sense to discuss the CST problem,

so we restrict our attention to DST. The Steiner tree problem is one of Karp's 21 NP-complete problems, i.e., in his seminal work [41], Karp showed that DST is NP-complete. Therefore, the attention of the algorithmic community turned towards obtaining good approximation algorithms for DST in general metrics. It is well-known that a minimum-cost tree only containing the terminals as nodes is a 2-approximation for DST (and CST) [33, 62].

A sequence of improved approximation algorithms for DST in general metrics appeared in the literature [64, 42, 49, 53], culminating with the famous $\ln(4) + \varepsilon < 1.39$ factor approximation algorithm by Byrka, Grandoni, Rothvoß, and Sanitá [9] (where $\varepsilon > 0$ is an arbitrarily small constant). On the hardness of approximation front, Chlebík and Chlebíková [16] showed that DST in general metrics is hard to approximate to a factor of $\frac{96}{95} > 1.01$. Thus, there is still quite a substantial gap in our understanding of the approximability of DST in general metrics.

OPEN QUESTION 1. *What is the tight inapproximability ratio of DST in general metrics?*

Geometric Metrics. Special cases of CST in Euclidean metric were studied by Fermat, Torricelli and other mathematicians as early as the 17th century and were also discussed in a letter from Gauss to Schumacher in 1836 (see [8] for more details on the history of this problem). On the other hand, phylogeny reconstruction is a long-standing (dating from Darwin's evolutionary theory) and intensively studied problem in computational biology, and computing the phylogenetic tree can be modeled as solving CST in the Hamming metric [29]. We have already discussed the importance of CST in the ℓ_1 -metric in VLSI design. Therefore, there has been a lot of interest in understanding the complexity of CST in ℓ_p -metrics.

Garey, Graham, and Johnson showed that Euclidean CST is NP-hard, even in the plane [31]. Garey and Johnson also showed NP-hardness of CST in the ℓ_1 -metric [32] (also in the plane). Foulds and Graham extended this further, showing NP-hardness CST in the Hamming metric [29]. These intractability results lead to the search for efficient approximation algorithms.

In a remarkable collection of works, a PTAS was established for CST in ℓ_1 and ℓ_2 -metrics in constant dimensions (indeed, Arora's PTAS extends to all ℓ_p -metrics) [4, 46, 51]. However, it was unclear if CST in ℓ_1 and ℓ_2 -metrics in high dimensions admitted a PTAS. Trevisan [60] answered this question in the negative for the ℓ_1 -metric, by showing that CST in ℓ_1 -metric is APX-hard [12, 34, 23, 63]. He left it as an open problem to show a similar hardness for the Euclidean CST in high dimensions.

OPEN QUESTION 2. *Does Euclidean CST on n terminals in $\Omega(\log n)$ dimensions admit a PTAS or is it APX-hard?*

If a geometric optimization problem is APX-hard, it is often easiest to establish that hardness in the ℓ_∞ -metric. Intuitively, this arises due to the existence of the Fréchet embedding, an isometric embedding of discrete metric spaces into the ℓ_∞ -metric. However, it remained unknown whether high-dimensional CST in the ℓ_p -metric is APX-hard for *any* $p > 1$, including

$p = \infty$. Morally, we should not be able to understand hardness of approximation of Euclidean CST until at least understanding the hardness of approximation of CST in the ℓ_∞ -metric.

OPEN QUESTION 3. *Does CST in the ℓ_∞ -metric on n terminals in $\Omega(\log n)$ dimensions admit a PTAS or is it APX-hard?*

We now shift our focus to DST in ℓ_p -metrics. Most approximation algorithms for CST in the ℓ_p -metric work by reducing them to an instance of DST in the ℓ_p -metric (including, for example, Arora's PTAS [4]). Therefore, from a hardness point-of-view, studying DST in ℓ_p -metrics can also be seen as a stepping stone to understanding CST in ℓ_p -metrics. Bartal and Gottlieb [6] showed that a PTAS exists for DST in spaces of bounded doubling dimension (including, for example, DST in ℓ_p -metrics in constant dimensions) and Euclidean space up to $O(\sqrt{\log \log n})$ dimensions, where n is the number of input points. Nonetheless, the inapproximability in higher dimensions remains open. For example, we can ask the following discrete variant of Open Question 2.

OPEN QUESTION 4. *Does Euclidean DST on n terminals in $\Omega(\log n)$ dimensions admit a PTAS or is it APX-hard?*

String metrics are another important family of metrics for DST, particularly given applications of the Steiner tree problem to the study of phylogenetic trees in computational biology and to computational linguistics. Several works have studied the computational hardness of Steiner tree in the Hamming metric, establishing APX-hardness in high-dimensions and NP-hardness for low dimensions with large alphabets [63, 3, 56, 29, 27]. Other popular string metrics less well understood in the context of the Steiner tree problem include the edit distance metric [44] and the Ulam metric.

OPEN QUESTION 5. *Does DST in the edit distance metric admit a PTAS or is it APX-hard? What about the Ulam metric?*

We remark that exploring the discrete avenue to better understand the continuous version has been pursued for another celebrated geometric problem: *clustering*. The discrete and continuous versions of minimizing popular clustering objectives such as k -means and k -median has been studied by the algorithmic community extensively [10, 2, 13, 20, 21].

1.1 Our results

Our first contribution is the resolution of Open Question 3.

THEOREM 1.1. *There is an efficiently computable function mapping a graph G of order n to an instance of CST in the ℓ_∞ -metric such that the optimal cost of the Steiner tree is $(n + \chi(G))/2$, where $\chi(G)$ is the chromatic number of G . Consequently, CST in the ℓ_∞ -metric is APX-hard.*

Our result also establishes a neat connection between the chromatic number of a graph and the cost of the optimal Steiner tree of the mapping of the graph. We remark that our current hardness results are derived from the hardness of graph coloring on graphs with linear chromatic number. Although the current hardness is only enough to imply APX-hardness of CST in the ℓ_∞ -metric, stronger inapproximability results for graph coloring immediately translate to more robust results for CST. For example, if the best known approximation algorithm for clique cover on cubic graphs were shown to be optimal, this would yield hardness of approximation within a factor of 17/16 (they show a 5/4-approximation algorithm in [11]).

Our second contribution is the resolution of Open Question 4. In fact, we prove the APX-hardness of DST for all ℓ_p -metrics.

THEOREM 1.2 (Implied by Theorem 5.3). *Let $p \in \mathbb{R}_{\geq 1} \cup \{\infty\}$. There is some constant $\varepsilon > 0$ such that DST on n terminals is NP-hard to approximate to within a $(1 + \varepsilon)$ factor in the ℓ_p -metric, even in $O(\log n)$ dimensions.*

In fact, it is sufficient to prove that DST is APX-hard in the Euclidean metric, and then simply use the near isometric embedding of ℓ_2 -metric to other ℓ_p -metrics [50] to obtain that DST is APX-hard in every ℓ_p -metric.

That said, our result is stronger than suggested in Theorem 1.2, in the following sense. Our reduction is from a hybrid variant of the set cover and set packing problems (both restricted to the case when the input is a collection of sets each of size exactly 3). Although current known hardness of bounded set cover and set packing yield very small constant factor inapproximability bounds for DST in the ℓ_p -metric, these problems are expected to be harder than currently known. We explicitly compute inapproximability factors throughout so that improvements for those problems will also immediately translate to improved hardness of approximation factors for DST in ℓ_p -metrics.

Elaborating, our reduction is from the (ε, δ) -SP₃ problem, where given as input a collection of sets \mathcal{S} over universe $[n]$, we would like to distinguish between the completeness case where there exists $n/3$ pairwise disjoint sets in \mathcal{S} that cover the whole universe $[n]$, and the soundness case where every set cover of $[n]$ from \mathcal{S} must be of size at least $(1 + \delta)(n/3)$ and every subcollection of pairwise disjoint sets in \mathcal{S} can cover at most a $(1 - \varepsilon)$ fraction of the universe. Note that if (ε, δ) -SP₃ problem is NP-hard then we may assume $\delta \in [\varepsilon/2, 2\varepsilon]$ (see Remark 2.1 for an explanation).

We derive the following hardness of approximation factors, in terms of the parameters of the reduction from SP₃ to DST, from our proof of Theorem 5.3.

THEOREM 1.3. Let $p \in \mathbb{R}_{\geq 1} \cup \{\infty\}$. Assuming (ε, δ) -SP₃ is NP-hard, we have that DST in ℓ_p -metric is NP-hard to approximate to $(1 + \gamma)$ factor, where

$$\gamma := \begin{cases} \delta/4 & \text{if } p = \infty \\ \frac{\varepsilon}{2} \left(1 - \frac{1}{3^{1/p}}\right) + 2\delta \left(\frac{1}{2 \cdot 3^{1/p}} - \frac{3}{8}\right) & \text{if } p > 1/\log_3(4/3) \\ \varepsilon/8 & \text{if } p = 1/\log_3(4/3) \approx 3.8 \\ \varepsilon/26 & \text{if } p \in [1, 1/\log_3(4/3)) \end{cases}$$

We have only tried to optimize our reduction for $p = 2$ and $p > 1/\log_3(4/3) \approx 3.8$, and, even then, we only optimize fully in the case of $\delta = \varepsilon/2$. At each of these thresholds, the scaling of our DST hard instances changes. We optimize in the range $p > 1/\log_3(4/3)$ because the optimal scaling of our DST hard instance is clear. Indeed, $p = 1/\log_3(4/3)$ is exactly the threshold where the interactions between the parameters become far more complex. It is possible to obtain a better relationship between ε , δ , and γ for $p \in (1, 1/\log_3(4/3)) \setminus \{2\}$ by carefully analyzing the constraints mentioned in Theorem 5.3.

In Figure 1, we capture the interplay between ε , δ , and γ . The purple region indicates that, between the set packing and set cover (actually vertex cover) problems, the reduction (i.e., embedding) of the set system only goes through by starting from the covering problem and not from the packing problem. The green region on the other hand is the exact opposite, and the reduction of the set system to the DST only goes through by starting from the packing problem and not from the covering problem. In the yellow region, the behavior is unclear, in part because we did not carefully analyze the dependency between ε and δ . Finally, in the red region, we benefit from both the covering and packing hardness. This is elaborated further in Section 5.2.

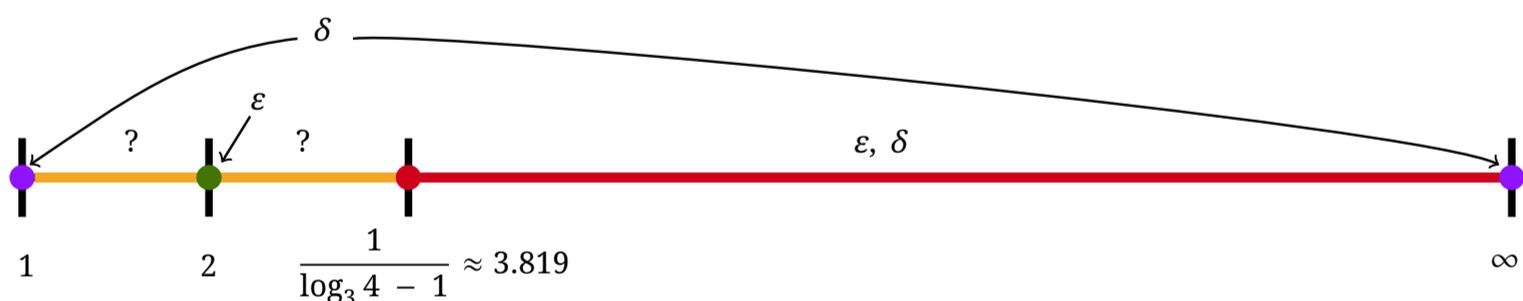


Figure 1. Color-coded range of ℓ_p hardness of approximation dependency on ε and δ .

Moreover, we note that we can derive explicit constants for the hardness of approximation factors, and we do for DST in the Hamming (ℓ_0), ℓ_1 , ℓ_2 -metrics in Appendix A.

Our third contribution is observing the relative computational hardness of DST with respect to CST. Although solving DST instances is often used as a subroutine in algorithms for

CST, prior to this work it was unclear whether solving DST problems was computationally easier or harder in general. We show that DST is essentially at least as hard as CST.

THEOREM 1.4 (Implied by Theorem 6.1). *Let $p \in [1, \infty]$ and $\alpha > 0$. If CST in the ℓ_p -metric is NP-hard to approximate within a factor of $(1 + \alpha)$, then, for all $\varepsilon > 0$, DST in the ℓ_p -metric is NP-hard to approximate within a factor of $(1 + \alpha - \varepsilon)$.*

We prove this by showing that we can efficiently compute a collection of candidate Steiner points for a DST instance such that the optimal tree for the DST instance is near-optimal for the corresponding CST instance. Hence, efficient approximation algorithms for DST yield similarly effective approximation algorithms for CST. As an immediate consequence, this yields a 1.39-approximation polynomial time algorithm for high dimensional CST in ℓ_p -metrics (see Corollary 6.4). The proof heavily relies on structural results about near-optimal Steiner trees from [24, 7, 6].

String Metrics. By applying the isometric embedding of the Hamming metric to the Ulam metric (for example see Lemma 4.5 in [1]) we can prove APX-hardness of DST in the Ulam metric.

THEOREM 1.5. *There is some constant $\varepsilon > 0$ such that DST in Ulam metric on n terminals is NP-hard to approximation to within a $(1 + \varepsilon)$ factor, even on permutations of length $O(\log n)$.*

Similarly, by applying the near-isometric embedding of the Hamming metric to the edit distance metric (for example, see Section 4.1 in [54]), we can prove APX-hardness for DST in the edit distance metric as well.

THEOREM 1.6. *There is some constant $\varepsilon > 0$ such that DST in edit distance metric on n terminals is NP-hard to approximation to within a $(1 + \varepsilon)$ factor, even on Boolean strings of length $O(\log n \log \log n)$.*

In combination, these results resolve Open Question 5. See Appendix B for details.

Dimensionality Reduction. Our proof of Theorem 1.2 holds when the dimension of the point set has linear dependency in the size of the point set. However, we can use the dimensionality reduction technique introduced in [20] and generalized in [21] to obtain the same inapproximability result as in Theorem 1.2, even when the dimension of the point set has logarithmic dependency on the size of the point set. We discuss these dimensionality reduction details (for constant p) in Appendix C.¹ Assuming the *Exponential Time Hypothesis* [37, 38], the dimension cannot be reduced much below logarithmic in size of point set as the runtime of the approximation scheme of [4, 46, 51] would then be subexponential.

¹ Alternatively, it is possible to apply the Johnson-Lindenstrauss lemma [40] in the ℓ_2 -metric and then use a near-isometric embedding into the ℓ_p -metric for $p \in [1, \infty]$ to achieve similar guarantees. We thank the anonymous referee for this suggestion.

1.2 Our Techniques

In this subsection, we contextualize our techniques and provide a high level description of our technical contributions. We emphasize that the conceptual difficulty in the proofs of Theorems 1.1, 1.2, and 1.3 is devising the appropriate starting problem and a clean embedding. The main technical challenge is to prove the soundness property of the reduction. This involves proving several structural properties of optimal Steiner trees in the corresponding metrics.

1.2.1 Proof Overview of Theorem 1.1: Hardness of Approximation of CST in ℓ_∞ -metric

First, we discuss the context of Open Problems 2 and 3 and describe our techniques for proving Theorem 1.1 (the resolution of Open Problem 3).

Prior Work on CST and Technical Difficulties. Prior to this work, the only APX-hardness results for high dimensional CST were in the Hamming (ℓ_0) and Rectilinear (ℓ_1) metrics [63, 60]. The discrete structure of Hamming CST makes it far more amenable to existing tools, such as reductions from hard graph problems, than other variants of CST. Indeed, Trevisan’s proof of APX-hardness of Rectilinear CST is a direct reduction from Hamming CST. While he appeals to the integrality property of the Min Cut LP relaxation, Hanan showed decades earlier that instances of Rectilinear CST admit optimal Steiner trees using only Steiner points on the grid formed by the intersection of axis parallel lines passing through the terminals [34]. This *Hanan grid* arises from the “coordinate independence” in the ℓ_1 -metric, i.e., the fact that a change in one coordinate changes the ℓ_1 distance by the same amount. This property is unique to the ℓ_1 -metric among ℓ_p -metrics.

In other ℓ_p -metrics, we know much less. From [31], we know that Euclidean CST is NP-hard. The proof is shown in the plane, making heavy use of the fact that optimal Euclidean Steiner trees have non-intersecting edges and those edges meet at at least 120 degree angles [33]. However, the non-intersecting edge property is not especially useful for point configurations outside of the plane. Without this property, we have very few strong tools for restricting the structure of optimal Steiner trees. Indeed, we do not know a proof of NP-hardness of Euclidean CST which does not involve directly appealing to hardness in the plane. This poses a major barrier to showing APX-hardness of Euclidean CST. In low dimensions, we know a PTAS exists [4], so we cannot appeal to NP-hardness techniques. Hence, showing APX-hardness in high-dimensions will ultimately require a brand new NP-hardness argument.

Euclidean CST has also been heavily studied from the perspective of discrete geometry. The central quantity of interest in that setting is the *Steiner ratio* of a point configuration: the ratio of the cost of an optimal Steiner tree and a minimum spanning tree. For example, the Steiner ratio of the three vertices of a unit equilateral triangle is $\sqrt{3}/2$ with the optimal Steiner tree achieved by adding the center of the triangle as a Steiner point and connecting it to the

vertices. Perhaps the most basic question about the Steiner ratio is the following: what is the minimum Steiner ratio over point configurations in the Euclidean plane? Gilbert and Pollak conjectured that it is exactly $\sqrt{3}/2 \approx 0.866$, coming from the equilateral triangle [33]. This conjecture, the Gilbert-Pollak Steiner ratio conjecture, remains open after more than half a century [39]. The best known lower bound on this quantity is ≈ 0.824 , leaving a major gap, even in the plane [18].

Predictably, in higher dimensions we know even less about the Steiner ratio and the construction of Steiner trees. Although extrapolating from the plane one might guess that the vertices of a regular simplex might admit the minimum Steiner ratio in higher dimensions (the Generalized Gilbert-Pollak Steiner ratio conjecture), Du and Smith disproved this conjecture without providing a satisfying new candidate point configuration [25]. Moreover, we do not know how to construct optimal Steiner trees of high-dimensional point configurations, not even of the vertices of regular simplices (although [19] gives a candidate optimal Steiner tree for some dimensions). As in the plane, we also have a considerable gap in the bounds on the Steiner ratio: the minimum lies somewhere between 0.669 and 0.62 [19, 25, 26].

Given our weak bounds, we cannot appeal to the Steiner ratio in analyzing delicate inapproximability arguments. The fact that we cannot even construct high-dimensional optimal Steiner trees compounds this difficulty further. The state of affairs is similar in other ℓ_p -metrics. New results about the inapproximability of high-dimensional CST require novel structural insights.

Reduction from Graph Coloring to CST in the ℓ_∞ -metric. In this work, we show APX-hardness for CST in the ℓ_∞ -metric. Our proof is unlike the results in [63, 60]. Unable to reduce the problem from a discrete metric space (e.g., Hamming space), we rely instead on the observation that the distance between any two points in the ℓ_∞ -metric is determined by a single coordinate. We use this observation to introduce structural guarantees on Steiner trees we construct in the ℓ_∞ -metric.

Given a graph, we arbitrarily orient its edges to form the digraph $G(V, A)$. We construct an instance of CST in the ℓ_∞ -metric in $\mathbb{R}^{|A|}$ as follows. Each coordinate of the terminals corresponds to an edge in A . The terminals themselves correspond to vertices, along with the usual addition of a root terminal $\mathbf{0}$, the all-zeroes vector. Given $v \in V$ and $a \in A$, the terminal t corresponding to v is 1 in the coordinate corresponding to a if a is an outgoing edge from v , -1 if a is an incoming edge to v , and 0 otherwise. This ensures that for each coordinate, there are exactly two terminals that are nonzero in that coordinate. The terminals corresponding to adjacent vertices are at distance 2, and the terminals corresponding to non-neighboring vertices are at distance 1.

Then, a minimum proper vertex coloring $\pi : G \rightarrow [\chi(G)]$ induces a low-cost Steiner tree of these terminals. For each color class of vertices, there is a Steiner point s at distance $1/2$

from the set of corresponding terminals of that color class. Among those terminals, for each coordinate i , there is at most one terminal nonzero in that coordinate (using that each color class is an independent set). Setting s_i to be 0 or half of that nonzero coordinate yields the desired Steiner point (thus, we have $\|s\|_\infty = 1/2$). The tree is then constructed by connecting the terminals corresponding to each color class to the Steiner point described above and then connecting each of those Steiner points to $\mathbf{0}$. Each edge in the tree is of cost exactly $1/2$ and there are $n + \chi(G) + 1$ total vertices, yielding a total cost of $(n + \chi(G))/2$.

It turns out that this simple tree is actually the optimal Steiner tree for this configuration! We prove this in Theorem 1.1. Upon showing this fact, the hardness of graph coloring (even in graphs with linear chromatic number) yields APX-hardness of CST in the ℓ_∞ -metric.

To do this, we show that there exists an optimal Steiner tree exhibiting many complementary properties and then prove that this tree is precisely the tree described above. At a high level, the objective is to prove two main claims: nonzero terminals are leaf nodes in the Steiner tree, and no Steiner point is adjacent to another Steiner point. Upon showing these two claims, it is easy to check that the Steiner points in the optimal tree are of the desired form.

Underlying the arguments required for each of these claims is an important observation about the structure of optimal Steiner trees of this CST instance. Let T be an optimal Steiner tree of this CST instance, and let t, t' be terminals such that $t_i = 1$ and $t'_i = -1$. Observe that t and t' are the only terminals in T with nonzero i^{th} coordinates. We can show that there exists an optimal Steiner tree such that the i^{th} coordinate of Steiner points on the path between t and t' decreases maximally away from t until the coordinate reaches 0 (in the sense of determining the distance between adjacent points). Likewise, the i^{th} coordinate increases maximally toward 0 along edges away from t' . Throughout, the coordinates reduce in magnitude maximally away from the path linking t and t' . If, for example, the path between t and some Steiner point s passes through $\mathbf{0}$, then s cannot have any coordinates of the same sign as t . Adopting this view of the points in T being projected onto the i^{th} coordinate is valuable intuition for the proofs.

For example, we use this intuition in proving that we may assume nonzero terminals are leaf nodes. It is easy to show that in some optimal Steiner tree, the nonzero terminals are non-adjacent, but, a priori, it is unclear whether they may be adjacent to many Steiner points. Suppose that t is a nonzero terminal adjacent to multiple Steiner points. Consider removing the edges from t to Steiner points such that their resultant connected components do not contain $\mathbf{0}$. The idea is that, in these new connected components, coordinates of Steiner points with the same sign as coordinates of t are no longer useful, so we can set them all to zero. Upon doing so, connecting the Steiner points originally adjacent to t to $\mathbf{0}$ reconnects the tree without increasing the cost. This amounts to ensuring that $\mathbf{0}$ is included in the path between these Steiner points and t .

Showing that the Steiner points in the tree are non-adjacent is the most technical part of the proof. To do so, we first consider *Steiner leaves*, the Steiner points in the Steiner tree that

become leaf nodes upon removing all nonzero terminals in the tree. One important property of these Steiner points is that they are necessarily adjacent to more terminals than Steiner points (it is easy to show that Steiner points must be degree at least 3 by the triangle inequality). It is straightforward to show that terminals adjacent to the same Steiner point cannot overlap in nonzero coordinates. Then, to minimize the total cost of the edges incident to the Steiner leaf, its coordinates must have magnitude at most $1/2$, with magnitude $1/2$ in each of the nonzero coordinates of the adjacent terminals (with the same sign as that nonzero coordinate). Using this fact we can show that the terminals adjacent to adjacent Steiner leaves cannot overlap on nonzero coordinates either (by removing the edge between the Steiner leaves and adding a new edge to $\mathbf{0}$). Indeed, an even stronger fact is true: if two Steiner leaves are adjacent to a common Steiner point, then the terminals adjacent to the Steiner leaves cannot overlap on nonzero coordinates. This follows from the usual trick of dropping an edge and connecting one of the Steiner leaves directly to $\mathbf{0}$.

By leveraging these facts and the coordinate-wise view of the Steiner tree, we can show that each non-leaf Steiner point in the tree is adjacent to at least one other non-leaf Steiner point. To do this, consider s and s' , two Steiner leaves adjacent to a common Steiner point s'' . We can remove all edges from s' to its adjacent terminals and connect them instead to s . We can then modify s to ensure that all neighboring terminals are at distance $1/2$ away and then remove s' entirely since it is a Steiner leaf. Repeating this process will result in one of two things. Either s'' will become degree 2 and can be removed via the triangle inequality or s'' is adjacent to another non-leaf Steiner point. After repeating for all Steiner points adjacent to Steiner leaves, all remaining non-leaf Steiner points will be adjacent to at least one other non-leaf Steiner point.

Finally, consider removing all nonzero terminals and Steiner leaves. The resulting tree has at most vertex of degree 1, $\mathbf{0}$, and, hence, is only a tree if all Steiner points were Steiner leaves. That is, all Steiner points in the tree must have only been adjacent to terminals.

1.2.2 Proof Overview of Theorems 1.2 and 1.3: Hardness of Approximation of DST in ℓ_p -metrics

We present below the techniques needed to prove Theorems 1.2 and 1.3, first establishing the appropriate technical context.

Known Reductions from Set Cover and Vertex Cover. The following is a gap-preserving reduction from Set Cover to DST in general metrics due to Karp [41]. Given a set system $([n], \mathcal{S})$, where \mathcal{S} is a collection of m subsets of $[n]$, we construct a DST instance as follows. Our terminals will be $[n] \cup \{r\}$ (a special root vertex is introduced). Our facilities will be the subsets in \mathcal{S} . We connect each subset to r and also have edges between subsets/facilities and terminals/universe elements based on their membership. See Figure 2. The metric is simply the shortest path metric

on the graph. In the completeness case² there is a set cover of size k , and that corresponds to a Steiner tree of cost $n + k$. In the soundness case, if every set cover of $[n]$ must be of size $(1 + \delta)k$, (for some $\delta > 0$), then one can show that the Steiner tree cost is at least $n + k(1 + \delta)$. By then considering set systems where each set contains exactly 3 elements and in which the set cover is a partition in the completeness case, we obtain $k = n/3$, and thus DST is hard to approximate to within a $(1 + \delta/4)$ factor.³

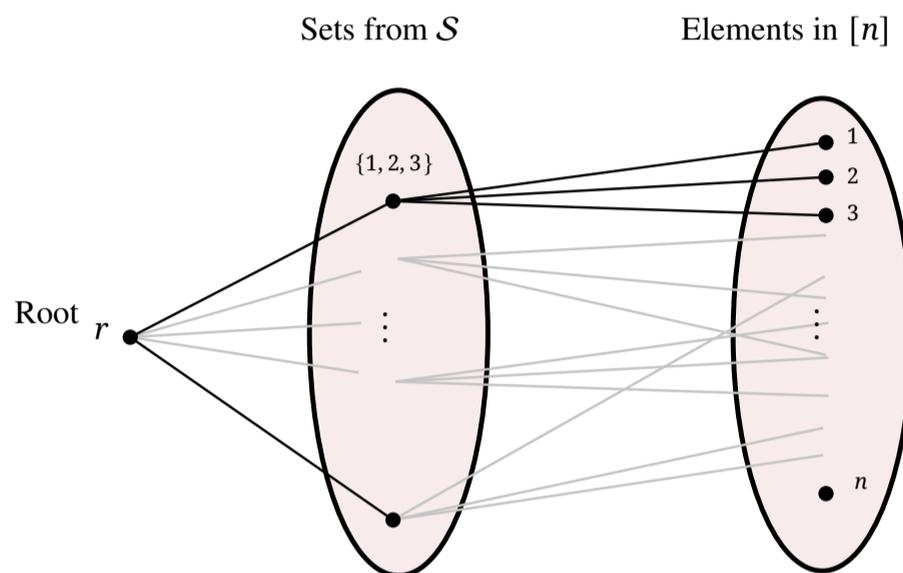


Figure 2. Instance of DST in general metrics from Karp's reduction from Set Cover

To obtain APX-hardness for DST in the Hamming (or ℓ_1) metric, we can simply start from a special set system, namely the vertex cover problem on cubic graphs, and then embed the graphs into the Hamming space as follows. Given a vertex cover instance, $G(V, E)$, our facilities are $\mathbf{e}_1, \dots, \mathbf{e}_{|V|}$, and for every edge (u, v) we have the terminal $\mathbf{e}_u + \mathbf{e}_v$ (where \mathbf{e}_i is a standard basis vector). We have a special additional terminal (much like the reduction from Set Cover in the previous paragraph) which is the all-zeroes vector. The completeness and soundness case go through like in the set cover instance, although some care needs to be taken in the completeness case. In particular, we cannot assume that the minimum vertex cover forms an independent set in the completeness case as the problem of deciding whether that is the case is in P.

We can then try to similarly reduce from vertex cover to DST in the Euclidean metric, but the reduction fails in a suggestive manner. Consider using the same choice of facilities and terminals as in the Hamming case. Given a vertex cover of size k , there is a Steiner tree of cost $m + k$ where we choose the facilities corresponding to the vertices in the vertex cover. Each

2 Here we use terminology from the hardness of approximation literature to refer to the two cases of our decision problems. The completeness case is the case of our transformed problem corresponding to the “yes” case of our original instance. The soundness case corresponds to the “no” case.

3 Although better hardness of approximation results are known for set cover for sets of size $B \geq 4$, reductions based on these set systems translate to hardness of approximation factors of roughly $\frac{B+1+\delta}{B+1}$, where we have hardness of approximation within a factor of $(1 + \delta)$. The value of δ here grows much more slowly than B [59], so it is generally better to consider sets of size 3, the minimum size where we can achieve hardness.

non-root terminal is a leaf node connected either to a facility at distance 1 or directly to the all-zeroes vector. Each facility is then connected to the all-zeroes vector.

However, unlike in the Hamming metric, such a Steiner tree might be far from optimal. Crucially, if a facility is adjacent to two non-root terminals, it is also more efficient to connect each of those adjacent terminals directly to the all-zeroes vector, incurring a saving of $3 - 2\sqrt{2}$ in the Steiner tree cost.

It is relatively straightforward to show that there are minimal cost Steiner trees in which all non-root terminals are leaf nodes and either adjacent to a facility of maximum degree or the all-zeroes vector. This is notably the same as the construction outlined above upon removing facilities adjacent to two non-root terminals. However, the minimum cost of a Steiner tree on these facilities and terminals no longer only depends on the minimum size of a vertex cover in G . Instead, the cost of the Steiner tree is minimized by finding a maximum independent set, then taking the corresponding Steiner points, and connecting all terminals corresponding to edges not incident to the independent set directly to the all-zeroes vector. Our reduction from minimum vertex cover naturally transformed into a reduction from maximum independent set!

Importantly, when we consider this reduction for even larger p , larger than around $p \approx 2.409$, the Steiner points lose all of their utility and even facilities adjacent to three non-root terminals do not help reduce the cost. By using a more robust reduction from set packing, the analog of maximum independent set on graphs, we manage to avoid this issue.

Reduction from Set Packing to Euclidean DST. Our main insight is that while we cannot reduce from the vertex cover problem to Euclidean DST, we can however reduce from the independent set problem on cubic graphs to Euclidean DST.

As such, we work with set packing problem, a generalization of the independent set problem. In particular, we work with SP_3 , the set packing problem where each set has cardinality 3, and it becomes much easier to limit the adjacencies of Steiner points when they depend on fewer coordinates. We also assume that there is a set partition in the completeness case. Given a set system $([n], \mathcal{S})$ where \mathcal{S} is a collection of subsets of $[n]$ of cardinality 3, we construct an instance of Euclidean DST as follows. The terminals are $\{\mathbf{e}_i : i \in [n]\} \cup \{\mathbf{0}\}$ and correspond exactly to the elements of the universe $[n]$, besides the addition of a special root terminal $\mathbf{0}$. The facilities are $\{\frac{1}{6}(\mathbf{e}_i + \mathbf{e}_j + \mathbf{e}_k) : \{i, j, k\} \in \mathcal{S}\}$. Each facility corresponds exactly to a set in \mathcal{S} ; it is a scaling of the characteristic vector of the set.

In the completeness case, there is a set packing. Choosing the facilities corresponding to the sets in the packing yields a Steiner tree of cost $(\sqrt{3}/2 + 1/\sqrt{108})n$. Each universe element is connected to the facility such that the corresponding universe element is contained in the corresponding set in the packing. All facilities are connected to the root terminal $\mathbf{0}$. In the soundness case, suppose that every set packing of $([n], \mathcal{S})$ can cover at most $(1 - \epsilon)n$ elements.

One can show that there exists a minimum cost tree such that: (1) the choice of facilities correspond to a maximal set packing $\mathcal{S}' \subseteq \mathcal{S}$, (2) each facility is adjacent exactly to $\mathbf{0}$ and the three terminals corresponding to universe elements contained in the facility's set, and (3) each non-root terminal is a leaf node and is either adjacent to a Steiner point or the root terminal $\mathbf{0}$. Given this structure, the cost of the tree is at least $((1 - \varepsilon)(\sqrt{3}/2 + 1/\sqrt{108}) + \varepsilon)n$, yielding hardness of approximation within a factor of $1 + \varepsilon/26$. The scale factor of $1/6$ in the construction of the facilities is chosen carefully to ensure (1), (2), and (3) and optimize this hardness of approximation factor (See Section 5 for details.).

This reduction is natural when the terminals are universe elements and the Steiner points are sets. By having our facilities depend on multiple coordinates, we keep the terminals far apart and also keep Steiner points relatively far apart (at least relative to our root) so that we can limit their adjacencies. Thus, upon removing the zeroes vector (special terminal), the Steiner tree splits into many connected components which can be interpreted as a choice of sets for our set packing.

Reduction from Set Packing to DST in ℓ_p -metric Spaces. It turns out that the reduction from SP_3 to Euclidean DST generalizes to a reduction from SP_3 to DST in ℓ_p -metric spaces. There are two main differences.

First, we modify our hard instance of SP_3 slightly to become $(\varepsilon, \delta)\text{-SP}_3$. The completeness case is identical to that outlined in the Euclidean DST case above. In the soundness case, in addition to each set packing of $([n], \mathcal{S})$ covering at most $(1 - \varepsilon)n$ elements, we also have that, for any set cover $\mathcal{S}' \subseteq \mathcal{S}$ such that $[n] \subseteq \cup_{S \in \mathcal{S}'} S$, $|\mathcal{S}'| \geq (1 + \delta)(n/3)$. We add this restriction on the size of minimum set covers since we observe that, for certain p , the hardness of approximation of DST in ℓ_p -metric spaces resultant from this reduction depends on both ε and δ (or even only δ). See Figure 1 and the surrounding discussion.

Additionally, in the reduction to DST in the ℓ_p -metric, the scale factor of $1/6$ in the Euclidean setting becomes some constant depending on p . In part, this is to sufficiently restrict the structure of optimal Steiner trees in the soundness case. Different choices of scale factor also yield improved hardness of approximation ratios. For example, it is optimal to use a scale factor close to $1/2$ for very large p .

1.2.3 Proof Overview of Theorem 1.4: Reduction from CST to DST

Finally, we discuss the techniques needed to prove Theorem 1.4. CST and DST have each independently garnered considerable interest, but, prior to this work, the relative computational hardness of the problems was unclear.⁴ On the one hand, searching over a possibly uncountable collection of candidate Steiner points in CST appears daunting. On the other hand, perhaps

⁴ In fact, to the best of our knowledge, the distinction between the discrete and continuous variants was never formalized and systematically studied in the literature.

it is possible to select collections of candidate Steiner points for DST that more effectively emulate other computationally hard problems. However, as the other side of the same coin, the continuous version allows the algorithm designer the freedom to pick Steiner points anywhere in space, whereas the discrete variant restricts the use of the geometry of the space by only allowing Steiner points from the set of facilities.⁵

In this paper we show that, up to an arbitrarily small factor, DST is at least as hard to approximate as CST (see Theorem 1.4 above). The proof uses an important structural result of optimal Steiner trees proved in a recent work of Bartal and Gottlieb [6]. (Similar structural results were proved in the context of k -restricted Steiner trees by Du, Zhang, and Feng, and Du and Borchers [24, 7] in the past and these results would suffice for our purpose as well). Namely, there exist near-optimal Steiner trees composed of optimal Steiner trees of constant size, linked by edges between terminals. By embedding all small subsets of terminals in low dimensional spaces and computing optimal Steiner trees using existing PTAS's, we can then construct a set of candidate Steiner points for DST (by unioning all of the Steiner points in the optimal Steiner trees of these small subsets). Then, the optimal Steiner tree on the DST instance approximates the cost of the Steiner tree on the CST instance arbitrarily well. As a consequence, this yields the best known approximation algorithm for high-dimensional CST (Corollary 6.4).

1.3 Organization of the Paper

In Section 2 we define the problems of interest to this paper. In Section 3 we prove APX-hardness of CST in the ℓ_∞ -metric via a reduction from graph coloring. In Section 4, we provide our generic framework of moving from a set system to an abstract and yet parameterized metric space to obtain hardness of DST. In Section 5, we detail our ℓ_p -metric embedding of the abstract metric space described in the previous section. In Section 6, we relate CST to DST, describing the reduction between the two problems. In Appendix A, we prove explicit factors of inapproximability for DST in ℓ_0 , ℓ_1 , and ℓ_2 -metrics. In Appendix B, we explain how to derive inapproximability of DST in string metrics using known results from the literature. In Appendix C, we show that our hardness results for DST hold even in $O(\log n)$ -dimensions using a specialized near-isometric embedding.

2. Preliminaries

In this section, we detail some notations and definitions related to ℓ_p -metric spaces, the Steiner tree problem, the graph coloring problem, the set packing problem, and the vertex cover problem.

⁵ These remarks also hold (in spirit) for other geometric optimization problems. For a useful comparison, consider clustering. Rather surprisingly, in [22], the authors showed that, in the ℓ_∞ -metric space, the continuous version of clustering is possibly harder than the discrete version.

ℓ_p -metrics. Let $p \in \mathbb{R}_{\geq 1} \cup \{0, \infty\}$. For any two points $a, b \in \mathbb{R}^d$ we denote the distance between them in the ℓ_p -metric by

$$\|a - b\|_p := \begin{cases} |\{i \in [d] \mid a_i \neq b_i\}| & \text{if } p = 0, \\ \left(\sum_{i \in [d]} |a_i - b_i|^p \right)^{1/p} & \text{if } p \in \mathbb{R}_{\geq 1}, \\ \max_{i \in [d]} |a_i - b_i| & \text{if } p = \infty. \end{cases}$$

Steiner tree problem. Let (\mathcal{X}, Δ) be a metric space on set \mathcal{X} and distance function Δ . Given a set P of points in (\mathcal{X}, Δ) , a *Steiner tree* of those points is a minimal spanning tree of $P \cup S$ for some $S \subseteq \mathcal{X}$. The initial points in P are called *terminals* and the points in S are called *Steiner points*.

Given a tree T with vertices as points in \mathcal{X} and edge weights induced by the distance function Δ , the cost of the tree $T = (P \cup S, E)$, denoted by $\text{cost}_\Delta(T)$ is

$$\text{cost}_\Delta(T) := \sum_{e=(u,v) \in E} \Delta(u, v).$$

In the continuous Steiner tree problem (CST) in the metric space (\mathcal{X}, Δ) , we are given as input $P \subseteq \mathcal{X}$, and the goal is to find a minimum cost Steiner tree of P . In the discrete Steiner tree problem (DST) in the metric space (\mathcal{X}, Δ) , we are given a pair (P, X) as input, with $P \subseteq \mathcal{X}$ the set of n terminals and $X \subseteq \mathcal{X}$ with $|X| = O(\text{poly}(n))$ as the set of facilities. The goal is then to find a minimum cost of a Steiner tree $T = (P \cup S, E)$ of P with $S \subseteq X$.

DST can also be considered as a graph problem. In this setting it is called the metric Steiner Tree problem on graphs. Given a complete graph $G = (P \cup X, E)$ with weight function $w : E \rightarrow \mathbb{R}^+$ satisfying the triangle inequality, the metric Steiner tree problem is to find the minimum cost of a Steiner tree of P using Steiner points from the set X .

Graph coloring problem. In the k -coloring problem, we are given a graph $G(V, E)$ and a positive integer k , and the goal is to determine whether there is a *proper coloring* $\pi : V \rightarrow [k]$ such that, for each $i \in [k]$, $\pi^{-1}(i)$ forms an independent set in G . That is, the vertices in G can be partitioned into at most k independent sets. The minimum k such that G admits a proper k -coloring is the *chromatic number* of G and is denoted $\chi(G)$.

In the (a, b) -coloring, the question is to instead decide which of the following two cases hold:

- *Completeness:* G admits a proper a -coloring ($\chi(G) \leq a$).
- *Soundness:* Every proper coloring of G uses at least b colors ($\chi(G) \geq b$).

Set packing problem. We say that $(\mathcal{U}, \mathcal{S})$ is a set system if \mathcal{S} is a collection of subsets of \mathcal{U} . An instance of the Set Packing problem (SP) is given by the triple $([n], \mathcal{S}, k)$, where $([n], \mathcal{S})$ is a set system with $|\mathcal{S}| = m$ and $k \in \mathbb{N}$, and the goal is to determine if there exists some $\mathcal{S}' \subseteq \mathcal{S}$ such that for all distinct $S, S' \in \mathcal{S}'$, we have $S \cap S' = \emptyset$ and $|\cup_{S \in \mathcal{S}'} S| \geq k$? When all of the sets in \mathcal{S} are fixed to be of size B , we abbreviate the problem as SP_B .

We define (ε, δ) - SP_3 to be the related decision problem in which we have as input $([n], \mathcal{S})$ as above and decide which of the following two cases hold.

- *Completeness:* $[n]$ may be partitioned into $n/3$ sets all of which are in \mathcal{S} .
- *Soundness:* Any collection of disjoint sets in \mathcal{S} covers at most $(1 - \varepsilon)n$ elements in $[n]$, and for any collection $\mathcal{S}' \subseteq \mathcal{S}$ (sets in \mathcal{S}' need not be pairwise disjoint) such that $[n] \subseteq \cup_{S \in \mathcal{S}'} S$ satisfies $|\mathcal{S}'| \geq (1 + \delta)(n/3)$.

REMARK 2.1. In the soundness case, the parameters ε and δ characterize the maximum coverage of a set packing and the minimum size of a set cover, respectively. These parameters are related. Indeed, given ε , we should have $\delta \in [\varepsilon/2, 2\varepsilon]$. If any collection of disjoint sets in \mathcal{S} covers at most $(1 - \varepsilon)n$ elements in $[n]$, then the smallest possible size of a set cover is $(1 + \varepsilon/2)(n/3)$ (if every additional set covers 2 elements). Assuming a set packing covering $(1 - \varepsilon)n$ elements is possible, the largest possible minimum set cover is $(1 + 2\varepsilon)(n/3)$ (covering the remaining elements one at a time).

Our reduction will use the following NP-hardness of (ε, δ) - SP_3 .

THEOREM 2.2 (Theorem 4.4, [48]). *There is some $\varepsilon > 0$ such that $(\varepsilon, \varepsilon/2)$ - SP_3 is NP-hard.*

In [48], they actually prove Theorem 2.2 for the Max 3-Dimensional Matching problem, but Max 3-Dimensional Matching is a special case of SP_3 . This correspondence is apparent by viewing each edge of a 3-uniform hypergraph as a subset of $[n]$.

Vertex cover problem. In the minimum Vertex Cover problem (VC), we are given a graph $G(V, E)$ and a positive integer k , and the goal is to determine if there exists $C \subseteq V$ with $|C| \leq k$ such that, for all $e \in E$, at least one endpoint of e is in C . Such a C , containing an endpoint of each edge, is called a *vertex cover* of G .

In (a, b) -VC, we are given a graph $G(V, E)$, and the goal is to decide whether there exists a vertex cover of G of size at most $a \cdot |V|$ or if every vertex cover of G has size at least $b \cdot |V|$.

3. APX-hardness of CST in ℓ_∞ -metric

In this section we prove Theorem 1.1. First, we recall some inapproximability results on graph coloring.

In [11], they study the problem of partitioning a cubic graph into the minimum number of independent cliques. The minimum size of such a partition is the minimum size of a partition of the complement graph into independent sets, the chromatic number of the complement graph. In particular, they show the following.

THEOREM 3.1 ([11], Theorem 7). *There exist constants $\varepsilon_1 < \varepsilon_2$ such that $(\varepsilon_1 n, \varepsilon_2 n)$ -coloring is NP-hard.*

Below we provide an efficiently computable function mapping a graph G of order n to an instance of CST in the ℓ_∞ -metric such that the optimal cost of the Steiner tree is $(n + \chi(G))/2$. Then the APX-hardness of CST in the ℓ_∞ -metric follows as a corollary by applying Theorem 3.1.⁶

Given an undirected simple graph, we arbitrarily direct/orient the edges in the graph to form a digraph $G(V, A)$. Define the function $\Gamma : V \rightarrow \mathbb{R}^{|A|}$ such that, for all $i \in V$,

$$\Gamma(i) = \left(\sum_{j: (i,j) \in A} \mathbf{e}_{(i,j)} \right) - \left(\sum_{j: (j,i) \in A} \mathbf{e}_{(j,i)} \right),$$

where $\mathbf{e}_{(i,j)}$ is the standard basis vector in $\mathbb{R}^{|A|}$ with 1 in the coordinate corresponding to the directed edge (i, j) and 0 elsewhere. Let $P = \{\Gamma(i) \mid i \in V\}$. Our corresponding instance of CST in the ℓ_∞ -metric is on the set of terminals $\tilde{P} = P \cup \{\mathbf{0}\}$. We will refer to $\mathbf{0}$ as the *root* terminal.

In the spirit of the language of gap-preserving reductions, we refer to the case of showing that the optimal cost of a Steiner tree of \tilde{P} is at most $(n + \chi(G))/2$ as the completeness case. Similarly, the case of showing that the optimal cost of a Steiner tree of \tilde{P} is at least $(n + \chi(G))/2$ is the soundness case.

3.1 Completeness

Let $\pi : V \rightarrow [a]$ be a proper a -coloring of the vertices in G . Let A^+ and A^- be two distinct partitions of A where for every $k \in [a]$, the k^{th} part of A^+ is $A^+(k) := \{(i, j) \in A \mid \pi(i) = k\}$ and the k^{th} part of A^- is $A^-(k) = \{(j, i) \in A \mid \pi(i) = k\}$. Define $\sigma : [a] \rightarrow \mathbb{R}^{|A|}$ as follows. The function σ will map each color class to a Steiner point. We have

$$\forall k \in [a], \sigma(k) = \left(\sum_{f \in A^+(k)} \frac{1}{2} \cdot \mathbf{e}_f \right) - \left(\sum_{f \in A^-(k)} \frac{1}{2} \cdot \mathbf{e}_f \right).$$

Now, our set of Steiner points will be exactly the set $\{\sigma(k) \mid k \in [a]\}$. Note that, since π is a proper coloring, for all $k \in [a]$, $A^+(k) \cap A^-(k) = \emptyset$. We create our Steiner tree as follows. Add an edge from each Steiner point $\sigma(k)$ to the root terminal $\mathbf{0}$. Additionally, add an edge from each non-root terminal $t \in P$ to the Steiner point $\sigma(\pi(\Gamma^{-1}(t)))$. Note that $\sigma(\pi(\Gamma^{-1}(t)))$ is the

⁶ Theorem 3.1 is stated without explicit values of ε_1 and ε_2 since the current derivable parameters come from a line of works proving APX-hardness without trying to optimize the hardness of approximation factor. If the best known approximation algorithm for partitioning cubic graphs into cliques were shown to be optimal, that would immediately imply CST in the ℓ_∞ -metric is NP-hard to approximate within a factor of 17/16.

Steiner point corresponding to the color class corresponding to the vertex mapped to terminal t . See Figure 3.

Note that the resultant Steiner tree is connected and every edge has length $1/2$. This length is clear for the edges from Steiner points to the root terminal. Now fix some $t \in P$. In particular, if $t_f = 1$, then $f \in A^+(\pi(\Gamma^{-1}(t)))$. That is, in this case, $\sigma(\pi(\Gamma^{-1}(t)))_f = 1/2$ since π is a proper coloring. Similarly, if $t_f = -1$, then $f \in A^-(\pi(\Gamma^{-1}(t)))$. Hence, in this case, $\sigma(\pi(\Gamma^{-1}(t)))_f = -1/2$. Note that $\|\sigma(\pi(\Gamma^{-1}(t)))\|_\infty = 1/2$ by definition, so, namely, if $t_f = 0$, $|\sigma(\pi(\Gamma^{-1}(t)))_f| \leq 1/2$. Altogether, this implies that $\|\sigma(\pi(\Gamma^{-1}(t))) - t\|_\infty = 1/2$.

Hence, this constructed tree has total cost $(n + a)/2$. In particular, the optimal cost of a Steiner tree in this instance of CST in the ℓ_∞ -metric is at most $(n + \chi(G))/2$.

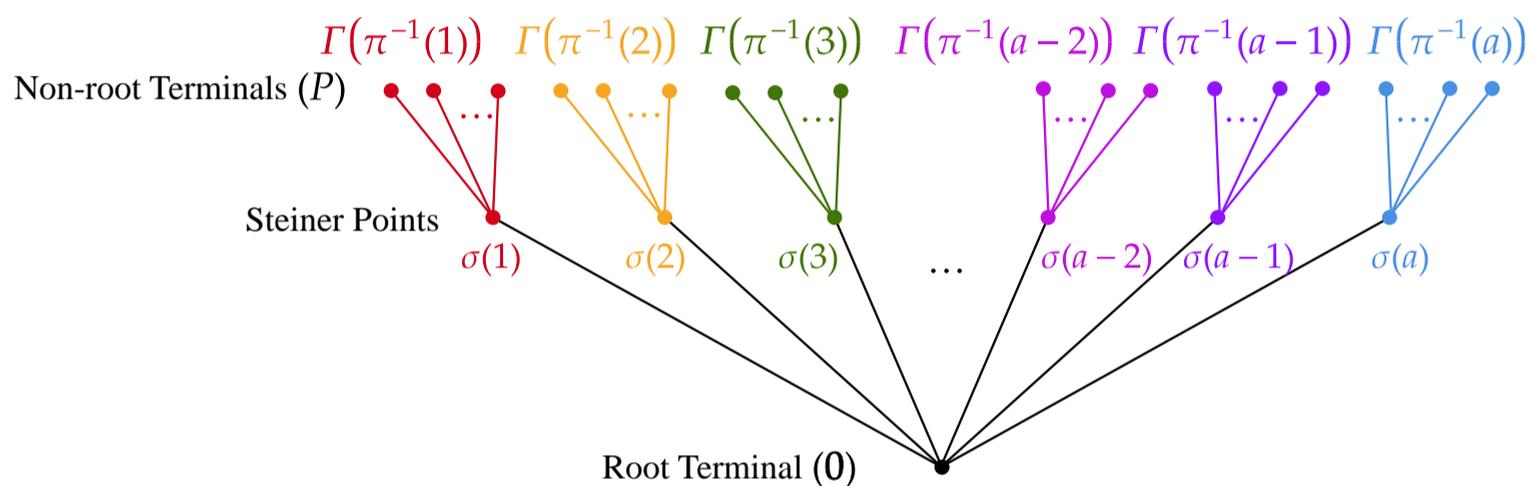


Figure 3. Steiner tree of P constructed from an a -coloring of G

3.2 Soundness

It remains to show that the optimal cost of a Steiner tree of \tilde{P} in the ℓ_∞ -metric is at least $(n + \chi(G))/2$. To do so, we prove a series of lemmas which in tandem show that the Steiner tree constructed in the completeness case is in fact optimal. We emphasize that the function of these lemmas is to prove increasingly rigid structural guarantees about the optimal Steiner trees in the soundness case, building off of weaker guarantees proved in previous lemmas.

LEMMA 3.2. *Let $T(\tilde{P} \cup X, E)$ be an optimal Steiner tree of \tilde{P} . Then, for all $(u, v) \in E$, we have $\|u - v\|_\infty \leq 1$.*

PROOF. Suppose T has some edge of length greater than 1. Remove the edge from T . Now, the resultant graph T' is split into two connected components. This yields a partition of the terminals in \tilde{P} into two parts. If one part contains no terminals, dropping that part yields a lower cost tree.

Otherwise, consider the minimum spanning tree of the terminals in \tilde{P} . All edges have length 1. Namely, since it is a spanning tree, it has some edge crossing the cut induced by the partition of terminals in T' . Adding that edge to T' forms a new Steiner tree of cost less than T , contradicting minimality of T and proving the claim. ■

For ease of notation, in the remainder of this section we will make use of the following function $\text{sgn} : \mathbb{R} \rightarrow \{-1, 0, 1\}$, where

$$\text{sgn}(x) = \begin{cases} -1, & \text{if } x < 0; \\ 0, & \text{if } x = 0; \\ 1, & \text{if } x > 0. \end{cases}$$

With Lemma 3.3, we begin by showing that there are optimal Steiner trees which are highly structured.

LEMMA 3.3. *There exists an optimal Steiner tree $T(\tilde{P} \cup X, E)$ with the following two properties.*

- (P1) *Rapid decay: For all edges $(u, v) \in E$ with $u \in P \cup X$ and $v \in X$, if $u_i \neq 0, v_i \neq 0$, and $\text{sgn}(u_i) = \text{sgn}(v_i)$, then $|u_i - v_i| = \|u - v\|_\infty$.*
- (P2) *Short edges: Each edge (u, v) for $u, v \in P \cup X$ satisfies $\|u - v\|_\infty < 1$.*

PROOF. Let $T(\tilde{P} \cup X, E)$ be an optimal Steiner tree of \tilde{P} . Consider the i^{th} coordinate (corresponding to some arc in $G(V, A)$). First, note that there are exactly two terminals in P with nonzero i^{th} coordinate. Let t be the terminal such that $t_i = 1$. Root T at t_i in the sense of directing each edge in T away from t_i . Suppose that (u, v) is a directed edge in T such that $u_i, v_i > 0$, and $\|u - v\|_\infty \neq u_i - v_i$. Note that this is only possible for $u \in P \cup X$ and $v \in X$.

We construct a new Steiner tree T' with $\text{cost}_\infty(T') \leq \text{cost}_\infty(T)$. The tree T' will be the same as T except for the following modifications. Set $v'_i := \max(0, u_i - \|u - v\|_\infty)$. Since $v_i > v'_i \geq 0$, and, for all terminals other than t we have $t_i \leq 0$, this cannot increase the length of edges to terminals (compared to T). This is because, by Lemma 3.2, since $v_i > 0$, v cannot be adjacent to terminals other than t .

Now, if for all directed edges (v, w) in T with $w \in X$ we have $v'_i \geq w_i$, then we have $\text{cost}_\infty(T') \leq \text{cost}_\infty(T)$ (since $v_i \geq v'_i$). Otherwise, suppose there exists some w such that $v'_i < w_i$. Then, set $w'_i := v'_i$. Then, if for all (w, x) we have $w'_i = v'_i \geq x_i$, we have $\text{cost}_\infty(T') \leq \text{cost}_\infty(T)$. Otherwise, repeat this process, setting $x'_i := v'_i$, checking the i^{th} coordinate of the out-neighbors of x' in T' and setting their i^{th} coordinates to v'_i if they are greater. Since T' is acyclic and finite, this process will eventually terminate.

Repeat this process for all other directed edges (u, v) with $u_i, v_i > 0$, and $\|u - v\|_\infty \neq u_i - v_i$, treating the newly created optimal Steiner tree T' as T and the resultant tree as T' . Then, further repeat this process for all other coordinates i . Afterward, repeat this process for all coordinates i , except this time for all edges (u, v) with $u_i, v_i < 0$ and $\|u - v\|_\infty \neq v_i - u_i$. For

negative coordinates and a directed edge (u, v) such that $u_i, v_i < 0$ and $\|u - v\|_\infty \neq v_i - u_i$, we will set $v_i^* := \min(0, u_i + \|u - v\|_\infty)$. Since each round of modifying the tree only affects a single coordinate (and only Steiner points which are nonzero and of a given sign in those coordinates), this yields (P1).

Before we modify the tree further to ensure (P2), we observe that (P1) implies a useful property of the Steiner points in the Steiner tree.

CLAIM 3.4. *Let $T(\tilde{P} \cup X, E)$ be an optimal Steiner tree satisfying (P1). Then, for all Steiner points $s \in X$, $\|s\|_\infty < 1$.*

Proof. Let $t \in P$ be the unique terminal with $t_i = 1$. Then, suppose that $s \in X$ is a Steiner point adjacent to t . From Lemma 3.2, $s_i \geq 0$. Since the distance between any two points in T is greater than 0 (assuming no degenerate Steiner points), from (P1), $s_i < 1$. Then, consider rooting T at t . (P1) immediately implies that the i^{th} coordinate of any Steiner point not adjacent to t is at most the maximum i^{th} coordinate of the Steiner points adjacent to t . This observation holds for any coordinate i and holds similarly for negatively signed coordinates, yielding the claim. \blacklozenge

Now, suppose that T' is the optimal Steiner tree resultant from the modifications for ensuring (P1). Let (u, v) be an edge with $u, v \in P \cup X$. By Lemma 3.2, $\|u - v\|_\infty \leq 1$. If $\|u - v\|_\infty = 1$, consider removing the edge (u, v) , thereby disconnecting the tree. One of u or v will be in the same resultant connected component as $\mathbf{0}$. Connect the other to $\mathbf{0}$, reconnecting the tree with an edge of cost at most 1. Repeating this process yields (P2) since all edges added in each iteration of this process involve $\mathbf{0}$ and each iteration removes an edge in violation of (P2). Since this process only removes edges and adds edges adjacent to $\mathbf{0}$, the resultant tree still satisfies (P1). \blacksquare

Optimal Steiner trees with the properties of Lemma 3.3 have some other useful properties.

COROLLARY 3.5. *Let $T(\tilde{P} \cup X, E)$ be an optimal Steiner tree satisfying the properties of Lemma 3.3. Then, for all $t, t' \in P$, $(t, t') \notin E$.*

PROOF. Note that for all $t, t' \in P$ distinct, $\|t - t'\|_\infty \geq 1$. (P2) then implies the desired result. \blacksquare

COROLLARY 3.6. *Let $T(\tilde{P} \cup X, E)$ be an optimal Steiner tree satisfying the properties of Lemma 3.3. Then, for all coordinates i and $t, t' \in P$ such that $t_i = 1$ and $t'_i = -1$, there does not exist Steiner point $s \in X$ such that $(t_i, s), (t'_i, s) \in E$. That is, t and t' do not share a common neighboring Steiner point.*

PROOF. Suppose otherwise. Then, consider s_i . (P2) implies that $s_i > 0$ and $s_i < 0$, a contradiction. \blacksquare

We now proceed by a series of lemmas extending the structural constraints established in Lemma 3.3.

LEMMA 3.7. *There exists an optimal Steiner tree $T(\tilde{P} \cup X, E)$ satisfying (P1), (P2), and the following additional property:*

(P3) *Terminal connectivity: For all $t \in P$, t is a leaf node.*

PROOF. Let $T(\tilde{P} \cup X, E)$ be an optimal Steiner tree satisfying (P1) and (P2) (such a tree exists from Lemma 3.3). We modify T without increasing the cost to yield the desired Steiner tree.

First, by Corollary 3.5, there are no edges between non-root terminals in T . Now, suppose that $t \in P$ is adjacent to $\mathbf{0}$ and some Steiner point s . Note that t may be adjacent to other Steiner points. Drop the edge from t to $\mathbf{0}$ and replace it with an edge from $\mathbf{0}$ to s . From Claim 3.4, this decreases the cost of T , contradicting minimality.

Finally, suppose that t is adjacent to at least two Steiner points and is not adjacent to $\mathbf{0}$. Consider the subtrees of T that would contain these Steiner points if we removed t from the tree. In particular, one of the subtrees must contain $\mathbf{0}$ and the others cannot contain $\mathbf{0}$. Let s be a Steiner point adjacent to t whose resultant subtree does not contain $\mathbf{0}$. Call the subtree T_s . Now, drop the edge from s to t . For each coordinate i such that $t_i \neq 0$, for all Steiner points s' in T_s , if $\text{sgn}(s'_i) = \text{sgn}(t_i)$, set $s'_i = 0$. This cannot increase the cost of T_s —the length of edges between Steiner points cannot increase, the length of edges between Steiner points and terminals with i^{th} coordinate 0 cannot increase, and, by (P2), any Steiner points adjacent to the terminal with opposite sign i^{th} coordinate are unaffected. Now add an edge from s to $\mathbf{0}$, reconnecting the tree. In particular, we now have $s_j \neq 0$ only if $t_j = 0$, so $\|s\|_\infty$ is at most the former distance between s and t . Hence, these modifications did not increase the cost of the tree and removed a neighbor of t without adding any neighbors to a nonzero terminal. Repeating this process—for other Steiner points s adjacent to t such that their subtrees resultant from removing t do not contain $\mathbf{0}$ —yields (P3). Crucially, these modifications are consistent with (P1) and (P2). After making these modifications, the tree remains optimal and we can re-apply the modifications detailed in the proof of Lemma 3.3 to ensure (P1) and (P2) since those modifications do not add edges to any terminals in P . ■

Optimal Steiner trees with properties (P1), (P2), and (P3) reflect some the structure of the Steiner trees constructed in the completeness case. In particular, Corollary 3.6 combined with (P3) shows that these optimal Steiner trees induce a proper coloring of V . The color classes are defined by the (unique) Steiner points adjacent to non-root terminals, and each non-root terminal adjacent to $\mathbf{0}$ has its own color class. Given some optimal Steiner tree with properties (P1), (P2), and (P3), the induced coloring of V is given by the function π_T . Denote the set of non-root terminals adjacent to a Steiner point s by $P(s)$ (note that $\Gamma^{-1}(P(s))$ is a color class defined by π_T).

Equipped with these structural observations, additional completeness structure emerges in the following lemma. In the statement of the lemma, σ is defined exactly as in Section 3.1,

induced by the coloring π_T . In other words, for all Steiner points s only adjacent to non-root terminals and $\mathbf{0}$, if s is adjacent to some terminal t with $t_i \neq 0$, then $s_i = t_i/2$.

LEMMA 3.8. *There exists an optimal Steiner tree $T(\tilde{P} \cup X, E)$ satisfying (P1), (P2), and (P3), and the following additional property.*

(P4) *Completeness structure: For all Steiner points s only adjacent to non-root terminals and $\mathbf{0}$, we have $s = \sigma(\pi_T(\Gamma^{-1}(P(s))))$.*

PROOF. Let $T(\tilde{P} \cup X, E)$ be a Steiner tree satisfying (P1), (P2), and (P3). Consider some such Steiner point s in T adjacent to only non-root terminals and $\mathbf{0}$. Note that, by (P3), since nonzero terminals are leaf nodes, s must be adjacent to $\mathbf{0}$ so that T is connected.

First, we may assume that if $s_i \neq 0$, then there exists $t \in P(s)$ such that $t_i \neq 0$. Otherwise setting $s_i = 0$ does not increase the cost of the tree.

Next, if there exists $t \in P(s)$ with $t_i = 1$, then, by Corollary 3.6, t' such that $t'_i = -1$ is not an element of $P(s)$. Moreover, (P2) then implies that $s_i > 0$. The analogous facts are true if there exists $t \in P(s)$ with $t_i = -1$.

Now suppose that $\|s\|_\infty = 1/2 + \varepsilon$ for some $\varepsilon > 0$. Then, for each i such that $s_i > 0$, set $s_i = \min(s_i, 1/2)$, and, for each j such that $s_j < 0$, set $s_j = \max(s_j, -1/2)$. Note that, for at most one $t \in P(s)$, $\|s - t\|_\infty$ will increase as a result of this operation (and it will increase by at most ε). However, $\|s\|_\infty$ will decrease by ε , implying that the overall cost of the tree does not increase as a result of this operation (due to the edge to $\mathbf{0}$).

Finally, suppose that there is some i such that $0 < |s_i| < 1/2$. Let

$$\min_{\{i: 0 < |s_i| < 1/2\}} |s_i| = 1/2 - \varepsilon$$

for some $\varepsilon > 0$ and let

$$i_0 = \arg \min_{\{i: 0 < |s_i| < 1/2\}} |s_i|.$$

Let \tilde{t} be a terminal in $P(s)$ such that $|s_{i_0}| = 1/2 - \varepsilon$ and $\tilde{t}_{i_0} \neq 0$. Then, for each i such that $s_i > 0$, set $s_i = 1/2$ and for each j such that $s_j < 0$, set $s_j = -1/2$. Now, this modification decreases the distance from s to \tilde{t} by ε while increasing the distance from $\mathbf{0}$ to s by at most ε . Moreover, for all $t \in P(s)$, this does not increase $\|t - s\|_\infty$ since, from the previous modification, $\|t - s\|_\infty \geq 1/2$. Hence, this does not increase the cost of T and yields (P4). Note that (P1), (P2), and (P3) are all preserved under these operations. ■

Lemma 3.8 shows that when Steiner points are not adjacent to other Steiner points, they appear exactly like Steiner points in the completeness case. Our next objective is then to show that we may assume there are no edges between Steiner points. To do so, we will first consider the most “extremal” Steiner points which are not accounted for by Property (P4). Let $T(\tilde{P} \cup X, E)$ be an optimal Steiner tree satisfying (P1), (P2), (P3), and (P4). Then, removing all terminals in P

will result in a smaller (connected) tree since all such terminals are leaf nodes by (P3). Call the non-terminal leaf nodes in the resultant tree *Steiner leaves*. We show that Steiner leaves are highly structured in Lemma 3.9.

LEMMA 3.9. *There exists an optimal Steiner tree $T(\tilde{P} \cup X, E)$ satisfying (P1), (P2), (P3), (P4), and the following additional property.*

(P5) *Steiner leaf structure: For all Steiner leaves $s \in X$, $\|s\|_\infty = 1/2$. Moreover, for each coordinate i such that there exists $t \in P(s)$ with $t_i \neq 0$, $|s_i| = 1/2$.*

PROOF. Let $T(\tilde{P} \cup X, E)$ be a Steiner tree satisfying (P1), (P2), (P3), and (P4). Suppose that s is a Steiner leaf in T not satisfying (P5). Using the triangle inequality, we may assume that s has degree at least 3 (otherwise, the Steiner point may be replaced by an edge directly between its two neighbors without increasing the cost of the tree). Such an application of the triangle inequality is consistent with (P1), (P3), and (P4). The argument in Lemma 3.3 to achieve (P2) can be reapplied if this results in an edge of length 1. Then, the fact that s is a Steiner leaf implies that $|P(s)| \geq 2$. If s is only adjacent to terminals, we may apply (P4).

Otherwise, s must be adjacent to one Steiner point s' and at least two nonzero terminals. Suppose that $\|s\|_\infty = 1/2 + \varepsilon$ for some $\varepsilon > 0$. There are two cases to consider.

First, suppose that for all i such that $|s_i| = 1/2 + \varepsilon$, there exists $t \in P(s)$ such that $\text{sgn}(t_i) = \text{sgn}(s_i)$. Let $\gamma = \max(\{|s_i| : \forall t \in P(s), t_i = 0\} \cup \{1/2\})$. That is, γ is a maximum magnitude of a coordinate of s not corresponding to a coordinate of a terminal in $P(s)$ (or $1/2$ if no such coordinates are greater than $1/2$). Now, for all coordinates i such that $t_i = 1/2 + \varepsilon$, set $t_i = \gamma$. Likewise, for all coordinates j such that $t_j = -1/2 - \varepsilon$, set $t_j = -\gamma$. This increases $\|s - t\|_\infty$ for at most one $t \in N(s)$ by at most $1/2 + \varepsilon - \gamma$, but for all other $t' \in N(s)$, $\|s - t'\|_\infty$ decreases by $1/2 + \varepsilon - \gamma$. Hence, this can only increase the cost of tree if the tree if $\|s - s'\|_\infty$ increases. But, we may propagate this reduction in magnitude of coordinates through s' as in Lemma 3.3 in order to reestablish (P1) and ensure that this change does not increase $\|s - s'\|_\infty$, thereby not increasing the cost of the tree. Note that this relies on our assumption that, for all i such that $|s_i| = 1/2 + \varepsilon$, there exists $t \in P(s)$ such that $\text{sgn}(t_i) = \text{sgn}(s_i)$. Afterward, we may apply the modifications outlined in Lemma 3.3 to re-establish (P1) and (P2) since they will not increase the magnitude of coordinates of Steiner leaves.

Now, if we still have $\|s\|_\infty = 1/2 + \varepsilon' > 1/2$, there exists some coordinate i such that $|s_i| = \|s\|_\infty$ and for all $t \in P(s)$, $t_i = 0$. Now, for coordinates j such that $s_j > 1/2$, set $s_j = 1/2$ and, likewise, if $s_j < -1/2$, set $s_j = -1/2$. Then, for all $t \in P(s)$, $\|s - t\|_\infty$ decreases by ε' and $\|s - s'\|_\infty$ increases by at most ε' . Since $|P(s)| \geq 2$, this actually decreases the cost of the tree, contradicting minimality. Hence, this case cannot occur.

It remains to show that, for each coordinate i such that there exists $t \in P(s)$ with $t_i \neq 0$, $|s_i| = 1/2$. Call such coordinates *corresponding* coordinates of s (with the terminal $t \in P(s)$ corresponding to such a coordinate referred to as a *corresponding terminal*). Let i be a

corresponding coordinate of s such that $|s_i|$ is minimal; say $|s_i| = 1/2 - \varepsilon$. By (P2), $\text{sgn}(s_i) = \text{sgn}(t_i)$ for corresponding terminal t . By the above, $\|s\|_\infty \leq 1/2$. Now, increasing the magnitude of all corresponding coordinates of s (while retaining their sign) will decrease the distance to t by ε and increase the distance from s to s' by at most ε . Hence, repeating this operation does not increase the cost of the Steiner tree while ensuring Property (P5). If (P1) or (P2) are ever violated, they be restored by the processes described in Lemma 3.3 without violating (P5). Properties (P3) and (P4) are unaffected by the modifications outlined above. ■

The language introduced in Lemma 3.9 will be useful in the following lemmas. For s a Steiner point, call $C(s) = \{i : t \in P(s), t_i \neq 0\}$ the set of corresponding coordinates of s . We observe an implication of Lemma 3.9.

LEMMA 3.10. *There exists a optimal Steiner $T(\tilde{P} \cup X, E)$ satisfying (P1), (P2), (P3), (P4), (P5), and the following additional property:*

(P6) *Steiner leaf independence: Let s be a Steiner leaf. Then, if s'' is a Steiner point adjacent to s , then $C(s) \cap C(s'') = \emptyset$. Moreover, for all other Steiner leaves s' adjacent to s'' , $C(s) \cap C(s') = \emptyset$.*

PROOF. Let $T(\tilde{P} \cup X, E)$ be a Steiner tree satisfying (P1), (P2), (P3), (P4), and (P5). Let s be a Steiner leaf and s'' be a Steiner point neighboring s . First, suppose that $i \in C(s) \cap C(s'')$. Then, $s_i = 1/2$ (using (P5)). Note that $s''_i < 0$ by (P2). Then, dropping the edge (s, s'') and adding the edge $(s, \mathbf{0})$ maintains the connectivity of the tree while decreasing its cost, contradicting optimality of T .

Now, let s' be a second Steiner leaf neighboring s'' and let $i \in C(s) \cap C(s')$. Without loss of generality assume $s_i = 1/2$ and $s'_i = -1/2$. Then, clearly either $\|s'' - s\|_\infty \geq 1/2$ or $\|s'' - s'\|_\infty \geq 1/2$. Assume that $\|s'' - s\|_\infty \geq 1/2$. Removing edge (s, s'') and adding edge $(s, \mathbf{0})$ then removes this pair of Steiner points violating (P6) while keeping the tree connected and not increasing its cost. The fact that the cost does not increase follows from (P5). This modification may affect (P1), (P4), (P5), but following the algorithms described in Lemmas 3.3, and 3.8 and 3.9 can re-establish these properties without increasing the number of pairs of Steiner leaves violating (P6). Repeating this process then yields the result. ■

We are now ready to prove the final and most important structural property of optimal Steiner trees: there exist optimal Steiner trees with no edges between Steiner points.

LEMMA 3.11. *There exists a optimal Steiner $T(\tilde{P} \cup X, E)$ satisfying (P1), (P2), (P3), (P4), (P5), (P6), and the following additional property:*

(P7) *Steiner adjacency: There are no edges between Steiner points in T .*

PROOF. Let $T(\tilde{P} \cup X, E)$ be an optimal Steiner tree satisfying (P1), (P2), (P3), (P4), (P5), and (P6). To prove this result, we begin with an important claim.

CLAIM 3.12. *We may assume that, for all non-leaf $s \in X$, s is adjacent (in T) to $\mathbf{0}$ and at least one other non-leaf Steiner point or at least two non-leaf Steiner points.*

Proof. Suppose some non-leaf Steiner point s in T does not have this property. Since s is not a Steiner leaf, it must be adjacent to at least two points that are not nonzero terminals. Hence, it is adjacent to at least one Steiner leaf s' . By (P6), we have $C(s) \cap C(s') = \emptyset$ and, for any two Steiner leaves s' and s'' adjacent to s , s' and s'' , $C(s') \cap C(s'') = \emptyset$.

Note that by (P1), all coordinates in $C(s')$ behave the same in the following sense. Upon rooting the tree at s' and directing all edges away from s' , these coordinates decrease maximally in magnitude until reaching 0. So, until they reach 0, they all have the same magnitude. Now, fix a Steiner leaf s' adjacent to s . Remove all edges from s to its neighboring nonzero terminals, add edges from those nonzero terminals to s' , and treat each coordinate previously in $C(s)$ like the coordinates previously in $C(s')$ (setting them to have the same magnitudes and appropriate signs), propagating through the tree until the coordinates reach 0. This does not increase the cost of the tree since s' is a Steiner leaf, the tree satisfies property (P1), and we had $C(s) \cap C(s') = \emptyset$.

Do the same process outlined above for all other Steiner leaves adjacent to s (replace their edges to nonzero terminals with edges from those terminals to s' and have their corresponding coordinates emulate the coordinates of other nonzero terminals adjacent to s'). Also drop those Steiner leaf Steiner points.

After this process, s has degree 2, so, by the triangle inequality, it can be removed and an edge can be added directly between its two neighbors. We can repeat this process until s' is either adjacent to $\mathbf{0}$ or a non-leaf Steiner point satisfying the desired property. At some points, we may need to restore (P6), but we can do so by following the algorithm described in Lemma 3.10.

Repeating for all non-leaf Steiner points without this property then ensures the property holds in the resultant tree. After each step, we can restore properties (P1) and (P2) by applying the algorithm described in Lemma 3.3. The other properties are unaffected by these modifications. ◆

Given the claim, consider removing non-root terminals and Steiner leaf nodes from T . By the claim, all remaining Steiner points are degree at least two. But then, by (P3), the only possible leaf node in the resultant graph is $\mathbf{0}$. The only tree with only one leaf node is a singleton vertex, implying all Steiner points in T are Steiner leaves. Hence, we have (P7). ■

Lemma 3.8 in fact show that there exist optimal Steiner trees with exactly the structure of the completeness case. This is formalized in Lemma 3.13.

LEMMA 3.13. *There exists an optimal Steiner tree in which the length of each edge is $1/2$ and the number of Steiner points is exactly $\chi(G)$.*

PROOF. Let $T(\tilde{P} \cup X, E)$ be an optimal Steiner tree satisfying properties (P1), (P2), (P3), (P4), (P5), (P6), and (P7).

By (P3), (P7), (P2), and (P4), each terminal is either connected to a Steiner point by an edge of length $1/2$ or connected to $\mathbf{0}$ by an edge of length 1 . Each edge from a Steiner point to $\mathbf{0}$ is length $1/2$.

For each edge of length 1 from a terminal, we can add an intermediate Steiner point halfway along that edge (corresponding to the terminal) such that the edge is split into two edges of length $1/2$. Then, the cost is fixed and each edge is of length $1/2$.

Now, by Corollary 3.6, for each Steiner point s , $P(s)$ corresponds to an independent set in G . Since each edge in the graph is length $1/2$, the cost of the tree is $(n + |X|)/2$, where X is the set of Steiner points in the tree. Since each nonzero terminal is adjacent to a Steiner point, by Corollary 3.6, T must have at least $\chi(G)$ Steiner points. But, the construction in the completeness case shows a tree with $\chi(G)$ Steiner points and this structure is possible (adding the intermediate Steiner points for edges directly to $\mathbf{0}$), implying the result. ■

This completes the proof of the soundness case of Theorem 1.1.

4. The Metric Steiner Problem on Graphs

In this section, we show a gap preserving reduction from SP_3 to DST in general metrics. In Section 4.1 we describe how we associate a metric space to a set system. The resultant metric spaces will be called *set system spaces*. When the distances in these metric spaces satisfy an additional collection of constraints (see Definition 4.4), these set system spaces will be sufficient to yield a gap preserving reduction from SP_3 to DST. The details of this reduction are provided in Subsection 4.2.

4.1 From Set Systems to Metric Spaces

In this subsection, we describe how to associate a metric space to a set system⁷. We also introduce more restricted metric spaces that facilitate gap preserving reductions from SP_3 to DST.

DEFINITION 4.1 (Set System Space). Let $([n], \mathcal{S})$ be a set system such that for all $S \in \mathcal{S}$, we have $|S| = 3$, and $|\mathcal{S}| := m$. Let $\alpha_X, \alpha_P, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \gamma_2, \tau \in \mathbb{R}^+$. Then, we define the $(\alpha_X, \alpha_P, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \gamma_2, \tau)$ -set system space corresponding to $([n], \mathcal{S})$ to be the space $(\tilde{P} \cup X, \Delta)$ containing $n + m + 1$ points and a distance function Δ . We define $P := \{t_i : i \in [n]\}$, $\tilde{P} = P \cup \{r\}$, and $X := \{s_j : S_j \in \mathcal{S}\}$. We refer to the points in P as *universe elements*, r as the *root*, and the points in X as *facilities*. For convenience, for $s_j \in X$, we define $\Gamma(s_j) = S_j$.

⁷ For the purposes of our application to DST, we specify this association to metric spaces only for set systems where each set in the collection is of size exactly 3.

Interpreting the points of the space as vertices in a graph, we define the weight function Δ by defining every pairwise edge weight as follows (the weight of the edge between each point and itself is 0).

- $\Delta(r, s) = \alpha_X$ for all $s \in X$.
- $\Delta(r, t) = \alpha_P$ for all $t \in P$.
- $\Delta(t_i, s) = \beta_{\text{in}}$ for all $i \in [n]$ such that $i \in \Gamma(s)$.
- $\Delta(t_i, s) = \beta_{\text{out}}$ for all $i \in [n]$ such that $i \notin \Gamma(s)$.
- $\Delta(s, s') = \gamma_i$ for all $s, s' \in X$ such that $|\Gamma(s) \cap \Gamma(s')| = i$, for $i = 0, 1$, and 2 .
- $\Delta(t, t') = \tau$ for all distinct $t, t' \in P$.

To elucidate our choice of notation, we also provide intuitive definitions of each parameter.

- α_X is the distance from the root terminal to a Steiner point in X .
- α_P is the distance from the root terminal to a terminal in P .
- β_{in} is the distance between a Steiner point s and non-root terminal t_i such that the set corresponding to s contains the universe element corresponding to t_i .
- β_{out} is the distance between a Steiner point s and non-root terminal t_i such that the set corresponding to s *does not* contain the universe element corresponding to t_i .
- γ_i is the distance between Steiner points s and s' such that their corresponding sets have an intersection of size i .
- τ is the distance between non-root terminals.

Note that, although we say “distance” in the above intuitive descriptions, the weight function Δ may not define a metric, so set system spaces may not be metric spaces. The following definition resolves this issue.

DEFINITION 4.2 (Metric Compatibility). Suppose that we have $\alpha_X, \alpha_P, \beta_{\text{in}}, \beta_{\text{out}}, \gamma_0, \gamma_1, \gamma_2, \tau \in \mathbb{R}^+$, such that all of the following inequalities hold.

1. $\alpha_X \leq \min(\alpha_P + \beta_{\text{in}}, \alpha_P + \beta_{\text{out}})$.
2. $\alpha_P \leq \min(\alpha_X + \beta_{\text{in}}, \alpha_X + \beta_{\text{out}})$.
3. $\beta_{\text{in}} \leq \min(\alpha_X + \alpha_P, \beta_{\text{out}} + \tau, \gamma_i + \beta_{\text{out}} : i \in \{0, 1, 2\})$.
4. $\beta_{\text{out}} \leq \min(\alpha_X + \alpha_P, \beta_{\text{in}} + \tau, \gamma_i + \beta_{\text{in}} : i \in \{0, 1, 2\})$.
5. $\gamma_i \leq \min(2\alpha_X, \gamma_j + \gamma_k, 2\beta_{\text{in}}, 2\beta_{\text{out}} : i, j, k \in \{0, 1, 2\})$.
6. $\tau \leq \min(2\alpha_P, 2\beta_{\text{in}}, 2\beta_{\text{out}})$.

In this case, we call the tuple $(\alpha_X, \alpha_P, \beta_{\text{in}}, \beta_{\text{out}}, \gamma_0, \gamma_1, \gamma_2, \tau)$ *metric compatible*. Note that the constraints $\gamma_i \leq \beta_{\text{in}} + \beta_{\text{out}}$ and $\tau \leq \beta_{\text{in}} + \beta_{\text{out}}$ are implied by $\gamma_i \leq 2\beta_{\text{in}}, 2\beta_{\text{out}}$ and $\tau \leq 2\beta_{\text{in}}, 2\beta_{\text{out}}$.

A general diagram of an $(\alpha_X, \alpha_P, \beta_{\text{in}}, \beta_{\text{out}}, \gamma_0, \gamma_1, \gamma_2, \tau)$ -set system space is given in Figure 4. Note that not all pairwise distances in the diagram are labeled.

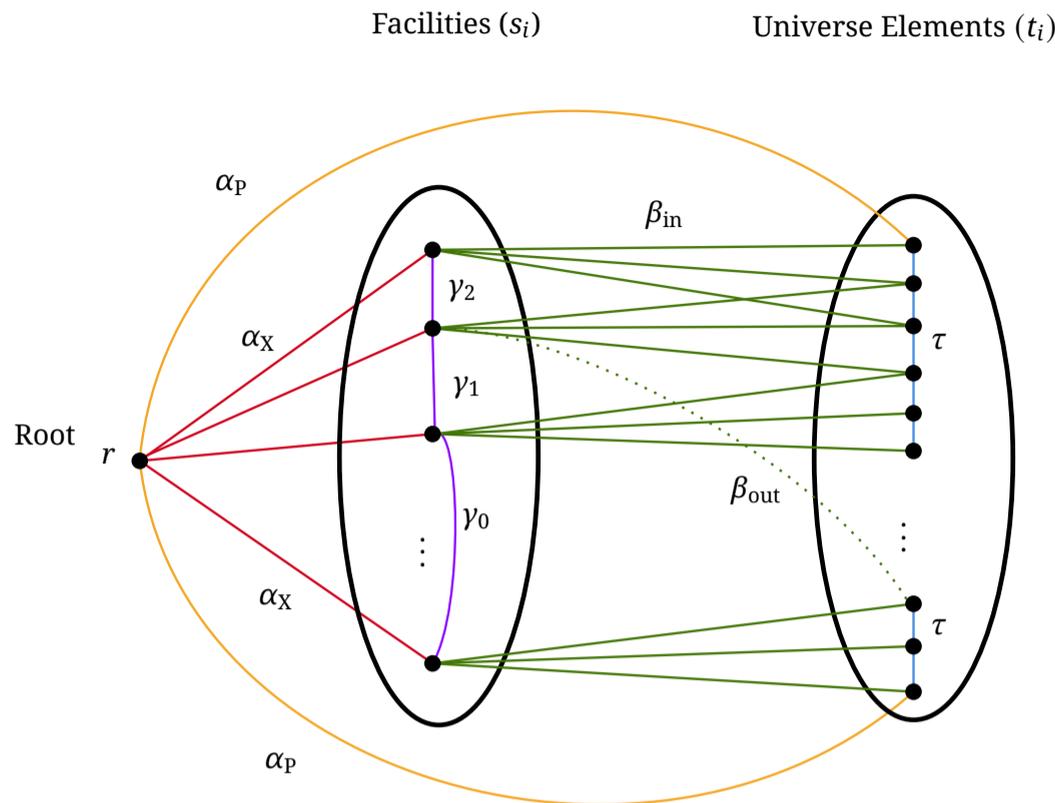


Figure 4. A general $(\alpha_x, \alpha_p, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \gamma_2, \tau)$ -set system space with edges between different groups of nodes color-coded

PROPOSITION 4.3. *If $(\alpha_x, \alpha_p, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \gamma_2, \tau)$ is metric compatible, then any $(\alpha_x, \alpha_p, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \gamma_2, \tau)$ -set system space is a metric space.*

PROOF. Inequalities 1–6 in Definition 4.2 are precisely those necessary for the triangle inequality to hold with distance function Δ . The other metric constraints follow trivially from Definition 4.1. ■

Hence, if $(\alpha_x, \alpha_p, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \gamma_2, \tau)$ is metric compatible, associated set system space weight functions Δ are distance functions. Now we introduce the notion of *Steiner embeddability*, a further restriction on metric compatible tuples. Tuples with these restrictions induce a notion of distance on metric spaces related to set systems that facilitates the construction of a gap preserving reduction from SP_3 to DST.

DEFINITION 4.4 (Steiner embeddability and Steiner spaces). We call a metric compatible tuple

$(\alpha_x, \alpha_p, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \gamma_2, \tau)$ *Steiner embeddable* if

(P1) *Steiner proximity:* $\alpha_x \leq 3\gamma_2/2, \alpha_p, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \tau$.

(P2) *Root proximity:* $\alpha_p \leq \beta_{out}$.

(P3) *Steiner utility:* $\beta_{in} + \frac{\alpha_x}{3} < \min(\alpha_p, \tau)$.

(P4) *Steiner diameter*: $\min(\alpha_P, \tau) \leq \beta_{in} + \gamma_2$.

We call any $(\alpha_X, \alpha_P, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \gamma_2, \tau)$ -set system space with $(\alpha_X, \alpha_P, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \gamma_2, \tau)$ Steiner embeddable a *Steiner space*. The constraints in Definition 4.2 ensure that the parameters can be realized as distances in a metric space, and the constraints in Definition 4.4 ensure sufficient properties for proving hardness of approximation of DST via a reduction from (ε, δ) -SP₃.

4.2 Hardness of DST from Set Packing

In this subsection we describe a general reduction from (ε, δ) -SP₃ to DST using the language of Steiner spaces.

THEOREM 4.5. *Let $(\alpha_X, \alpha_P, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \gamma_2, \tau)$ be a Steiner embeddable tuple. For an instance $([n], \mathcal{S})$ of (ε, δ) -SP₃ where $|\mathcal{S}| = m$, let $(\tilde{P} \cup X, \Delta)$ be the Steiner space guaranteed by Definitions 4.2 and 4.4. Then, we have the following guarantees on an instance $(\tilde{P} \cup X)$ of DST over $(\tilde{P} \cup X, \Delta)$:*

Completeness: *If $([n], \mathcal{S})$ admits a set packing of size $n/3$, then there is a Steiner tree for \tilde{P} of cost*

$$n(\alpha_X/3 + \beta_{in}).$$

Soundness: *Every Steiner tree of \tilde{P} must be of cost at least*

$$n(\alpha_X/3 + \beta_{in}) \cdot \left(1 - \varepsilon + \frac{(4\varepsilon - 2\delta)}{3} \cdot \frac{\min(\alpha_P, \tau, \beta_{in} + \alpha_X/2)}{\alpha_X/3 + \beta_{in}} + \frac{(2\delta - \varepsilon)}{3} \cdot \frac{\min(\alpha_P, \tau)}{\alpha_X/3 + \beta_{in}} \right).$$

Run Time: *Given $\alpha_X, \alpha_P, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \gamma_2, \tau$ and $([n], \mathcal{S})$, the above instance of DST can be constructed in $\text{poly}(n + m)$ time.*

REMARK 4.6. The completeness case bound comes from Steiner trees of the form shown in Figure 5. In the soundness case, we prove that the optimal structure of the Steiner tree involves packing as many terminals as possible into a tree structure analogous to the completeness case: groups of three terminals each at distance β_{in} from a Steiner point connected to the root r (we conclude this by Lemma 4.17). This is the $(1 - \varepsilon)$ term. Then, whether its better to try and group the remaining terminals into pairs connected to common Steiner points and handle them individually depends precisely on which term minimizes $\min(\alpha_P, \tau, \beta_{in} + \alpha_X/2)$. The first term is the per-terminal cost of connecting the remaining terminals each individually to the root, the second term is the per-terminal cost of connecting the remaining terminals each to some terminal already connected (via a Steiner point) to the root, and the final term in the per-terminal cost of connecting paired up terminals to the root via a common Steiner point. If $\beta_{in} + \alpha_X/2$ minimizes this expression, then it is optimal to pair up as many terminals as possible. However, the number of achievable pairs is limited by the minimum size set cover restriction

in the soundness case of SP_3 (this is where the $(4\varepsilon - 2\delta)/3$ comes from). Then, the remaining terminals are connected to the root either directly or via edges to already connected terminals.

Nonetheless, the expression for the inapproximability factor in the soundness case is quite intricate. One may wonder if this theorem implies a stronger inapproximability result than the simple reduction from Set Cover mentioned in Section 1.2 for DST in general metrics.

COROLLARY 4.7. *Suppose (ε, δ) - SP_3 is NP-hard and there exists a Steiner embeddable tuple $(\alpha_X, \alpha_P, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \gamma_2, \tau)$ that also satisfies*

1. $\alpha_X = \beta_{in}$,
2. $\min(\alpha_P, \tau, \beta_{in} + \alpha_X/2) = 3\alpha_X/2$, and
3. $\min(\alpha_P, \tau) = 2\alpha_X$.

Then, it is NP-hard to approximate DST in general metric spaces within a factor of $1 + \delta/4$.

PROOF. The hardness of approximation factor is achieved by dividing the costs in the soundness and completeness cases of Theorem 4.5. Then, we wish to maximize both

$$\frac{\min(\alpha_P, \tau, \beta_{in} + \alpha_X/2)}{\alpha_X/3 + \beta_{in}}$$

and

$$\frac{\min(\alpha_P, \tau)}{\alpha_X/3 + \beta_{in}}.$$

The numerator of the first fraction is at most $\beta_{in} + \alpha_X/2$. The numerator of the second fraction is at most $\alpha_P \leq \alpha_X + \beta_{in}$ by the triangle inequality. Supposing that both inequalities were in fact tight, our problem would reduce to maximizing $\frac{\alpha_X/2 + \beta_{in}}{\alpha_X/3 + \beta_{in}}$ and $\frac{\alpha_X + \beta_{in}}{\alpha_X/3 + \beta_{in}}$. Both are maximized for minimum β_{in} and, since $\alpha_X \leq \beta_{in}$ by (P1), the best possible is $\beta_{in} = \alpha_X$. Assuming there is some tuple $(\alpha_X, \alpha_P, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \gamma_2, \tau)$ that both satisfies these additional constraints and is Steiner embeddable, we would then have that DST is NP-hard to approximate within a factor of

$$1 - \varepsilon + \frac{(4\varepsilon - 2\delta)}{3} \cdot \frac{9}{8} + \frac{(2\delta - \varepsilon)}{3} \cdot \frac{3}{2} = 1 + \delta/4. \quad \blacksquare$$

In fact, DST in the ℓ_∞ -metric can satisfy these constraints and achieve this hardness of approximation factor. See Theorem 5.5.

REMARK 4.8. Interestingly, as shown in Section 1.2, there exists a simple reduction from Set Cover with sets of size 3 to DST which yields this same hardness in general metric spaces. Nonetheless, ε , the set packing parameter in (ε, δ) - SP_3 is highly relevant in certain other metric spaces. See, for example Theorem 5.3 and Corollary 5.7 for applications to ℓ_p metric spaces.

We now prove Theorem 4.5. Given any hard instance of SP_3 , using the language of Subsection 4.1, we can consider the corresponding $(\alpha_X, \alpha_P, \beta_{in}, \beta_{out}, \gamma_0, \gamma_1, \gamma_2, \tau)$ -set system space $((P \cup \{r\}) \cup X, \Delta)$. This space may be interpreted as an instance of DST with terminals $\tilde{P} = P \cup \{r\}$, candidate Steiner points X , and pairwise distances given by Δ determining edge weights.

Completeness. In the completeness case, the SP_3 instance has a solution that partitions $[n]$ into sets from \mathcal{S} . Let C be such a partition. For each $\{i, j, k\} \in C$, connect t_i , t_j , and t_k to $s = \Gamma^{-1}(\{i, j, k\})$. Then, connect each such s to r . This results in a Steiner tree of total cost $n(\alpha_X/3 + \beta_{in})$. See Figure 5.

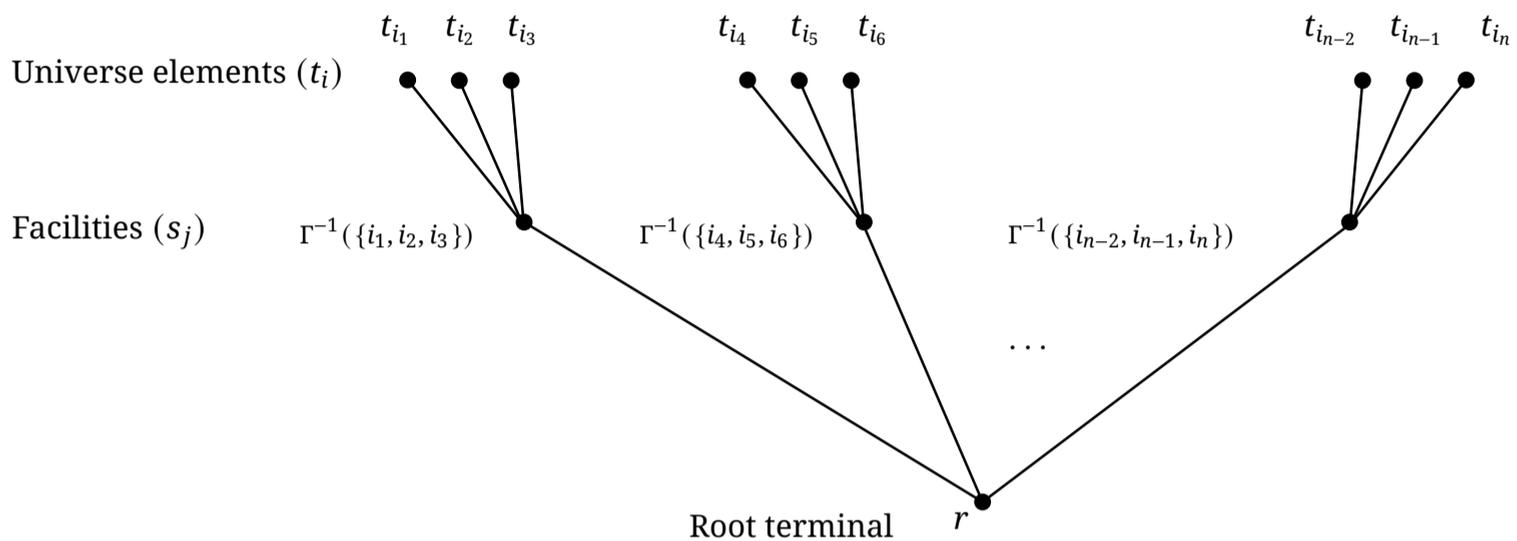


Figure 5. Steiner tree construction in the completeness case of Theorem 4.5.

Soundness. In the soundness case, we prove a series of consistent claims about the adjacencies of universe elements (non- r terminals) and Steiner points. The claims rely intimately on the properties of Steiner embeddable tuples. Ultimately, the claims will show that there is some minimum length Steiner tree mirroring much of the structure of the completeness case. This then yields insight into the size of a set packing in the SP_3 instance, thereby lower-bounding the length of a minimum Steiner tree.

We want r to serve as a root for the Steiner tree with its subtrees inducing groups of related Steiner points. To achieve this goal, we need to restrict the adjacencies of both universe elements and Steiner points. In Proposition 4.9 we show that there exists a minimum length Steiner tree such that, for each universe element t_i , t_i is adjacent to only Steiner points s such that $i \in \Gamma(s)$. We hope to interpret the Steiner points in the Steiner tree as a choice of sets from \mathcal{S} . This property is important since it means the adjacency of a universe element to a Steiner point can be interpreted as the set corresponding to the Steiner point covering the universe element.

Lemma 4.10 further restricts the adjacencies of universe elements. It shows that we may assume that edges between universe elements cannot connect substantial portions of the tree—one of the endpoints of such edges must be a leaf node. Indeed, the leaf node endpoint will ultimately become a universe element not covered by the set packing induced by the Steiner tree. Notably, Lemma 4.10 leaves open the option that a universe element could serve as a connecting hub for the tree, avoiding the need for Steiner points and limiting the role of r . Lemma 4.12,

proved using Lemma 4.11, discounts that possibility, enforcing that non-leaf node universe elements must be connected to Steiner points. However, even with Lemma 4.10, terminals could conceivably connect to many Steiner points adjacent to other terminals. Lemma 4.13 shows that we may assume that this is not the case. Proposition 4.14, combined with Lemma 4.12, shows that we may assume that universe elements of degree greater than 1 are not adjacent to r . This yields some insight into how groups of terminals connected to common Steiner points are connected to r .

The combination of these results severely limits the adjacency of universe elements. They are either adjacent to the root r , adjacent to some other universe element connected to a Steiner point, or are connected to some Steiner point whose corresponding set contains their corresponding element of $[n]$. In the former two cases, the universe element is a leaf node in T . In the latter case, while the universe element may not be a leaf node, if it is not, its only other neighbors are leaf node terminals. Treating those leaf node terminals as “uncovered” universe elements allows us to treat all terminals as leaf nodes.

The most important and most technical part of the proof is Lemma 4.15. Lemma 4.15 shows that we may assume that there are no edges between Steiner points in the Steiner tree. In combination with the severely limited adjacencies of universe elements, this shows that we may assume that removing r from the tree (and ignoring terminals not adjacent to Steiner points) divides the tree into groups of universe elements connected to Steiner points. By Proposition 4.9, these universe elements correspond to elements of $[n]$ contained in the sets corresponding to the Steiner points. So, the Steiner points adjacent to three terminals actually correspond precisely to a packing of sets from \mathcal{S} .

We want this induced packing to be large to yield insight into the size of maximum packing in the SP_3 instance. Indeed, the induced packing corresponds to a maximum packing, as shown in Lemma 4.17. Lemma 4.16 merely facilitates the proof of Lemma 4.17. The rest of the proof follows from this correspondence.

Now we proceed with the proof. Let $T = (T_V, T_E)$ be a minimum Steiner tree of \tilde{P} . We wish to show that we may assume a nearly identical structure to the completeness case, particularly in that it is optimal to pack as many terminals into sets as possible. As we will show, these packings arise in the Steiner tree as groups of three terminals t_i, t_j, t_k connected to a single Steiner point $s = \Gamma^{-1}(\{i, j, k\})$ which is in turn connected to r . In each successive claim, we assume that the properties of the previous claims initially hold.

PROPOSITION 4.9. *We may assume that $(t_i, s) \in T_E$ for $t_i \in P$ and $s \in X$ only if $i \in \Gamma(s)$.*

PROOF. Suppose $(t_i, s) \in T_E$ for $t_i \in P$, $s \in X$, and $i \notin \Gamma(s)$. Then, $\Delta(t_i, s) = \beta_{\text{out}} \geq \alpha_P, \alpha_X$ from (P1) and (P2) of Definition 4.4. Remove (t_i, s) from T_E . Then, one of the resulting connecting components contains r so we may add edge (r, s) or (r, t_i) to reconnect the graph. The change in cost is then either $\alpha_X - \beta_{\text{out}}$ or $\alpha_P - \beta_{\text{out}}$ which, in either case, is at most 0. This operation

reduces the number of edges of the form $(t_i, s) \in T_E$ for $t_i \in P$, $s \in X$, and $i \notin \Gamma(s)$, so repeating this process completes the proof since it yields a minimum cost T satisfying the claim. ■

LEMMA 4.10. *We may assume that if $(t, t') \in T_E$ for some $t, t' \in P$, then $\deg(t) = 1$ or $\deg(t') = 1$ in T .*

PROOF. Suppose otherwise. Let $(t, t') \in T_E$ with $t, t' \in P$ and $\deg(t), \deg(t') > 1$. Remove (t, t') . If the resultant connected component not containing r contains a Steiner point s , we may add an edge (r, s) to reconnect the tree. Since $\alpha_X \leq \tau$ by (P1), the cost of the tree does not increase.

Otherwise, without loss of generality assume that the connected component containing t' does not contain a Steiner point and does not contain r . Hence it merely contains terminals $\{t_{i_1}, t_{i_2}, \dots, t_{i_k}\} \subseteq P$. Remove all $(k-1)$ edges in this connected component. Each was of length τ . Now, add an edge (t, t_{i_j}) for $1 \leq j \leq k$. This reconnects the tree and fixes its original cost.

Note that in either case, the degree of each terminal is non-increasing. Hence, we reduced the number of edges between terminals with both terminals having degree > 1 , so repeating this process yields the desired result. Also note that we do not add any edges of the form considered in Proposition 4.9, so the two claims are consistent. ■

LEMMA 4.11. *T includes some $s \in X$. In particular, there exists $(s, t) \in T_E$ such that $s \in X$ and $t \in P$.*

PROOF. If there are no Steiner points in T , then we may assume that either every $t \in P$ is a leaf node connected to r (in the case of $\alpha_P < \tau$) or there exists $\tilde{t} \in P$ such that \tilde{t} is connected to r and every other terminal is a leaf node with an edge to \tilde{t} . This is because, since T is a tree, it has a fixed number of edges. Additionally, each edge is of length α_P or τ . Then, in the former case, since $\alpha_P < \tau$, there is a unique tree with all edges of length α_P (namely connecting every universe element to r), and the tree is minimum, that must be the tree. In the latter case, the minimum possible length of the tree is having one edge of length α_P (to connect some universe element to r) and the rest of length τ . The described tree is such a tree.

Now, let $t_i, t_j, t_k \in P$ and $s \in X$ such that $\Gamma(s) = \{i, j, k\}$. We may assume w.l.o.g. that none of the terminals are \tilde{t} in the latter case (using $m > 3$). Then, remove the edges incident to t_i, t_j, t_k , disconnecting the graph into four connected components. We may reconnect the graph by introducing s as a Steiner point, adding edges from t_i, t_j , and t_k to s and then adding an edge from s to r . The change in cost is $3\beta_{in} + \alpha_X - 3 \min(\alpha_P, \tau)$. In either case the change in cost is negative, using (P3), contradicting minimality of T . ■

LEMMA 4.12. *We may assume that $t \in P$ has $\deg(t) > 1$ in T only if t is adjacent to a Steiner point.*

PROOF. Suppose that $t \in P$ has degree greater than 1 but is not adjacent to any Steiner points. By Lemma 4.10, then all of the neighbors of t are either r or $t' \in P$ with $\deg(t') = 1$.

Now, if $\tau > \alpha_p$, removing all of t 's incident edges and connecting t and all of its neighbors $t' \in P$ to r reduces the cost of the tree, preserves connectivity, and decreases the degree of t to 1. Note that the number of edges removed equals the number of edges added. This is because t must be adjacent to r since all of its neighbors are leaf nodes from Lemma 4.10 and T is connected.

If $\tau \leq \alpha_p$, since there must be some $\tilde{t} \in P$ connected to a Steiner point s by Lemma 4.11, we can drop all of the edges incident to t and add edges from t and all its previously neighboring terminals to \tilde{t} . This does not increase the cost of the tree since $\tau \leq \alpha_p$ in this case. It also preserves connectivity of the tree.

Repeating this process then yields the claim. Note that this process is consistent with Proposition 4.9 since it does not add any edges from Steiner points to terminals. This process is also consistent with Lemma 4.10 since one endpoint of every edge added between elements of P in the second case is a leaf node. ■

LEMMA 4.13. *We may assume that each $t \in P$ is adjacent to at most one Steiner point.*

PROOF. Suppose otherwise, let $t_i \in P$ be adjacent to more than one Steiner point. From Proposition 4.9, t_i is adjacent only to $s \in S$ such that $i \in \Gamma(s)$. Drop edges from t_i to Steiner points such that t_i remains in the connected component with r . Then, add the same number of edges from r to a Steiner point in each other connected component. This does not increase the cost of T since $\alpha_X \leq \beta_{in}$ from (P1). Additionally, this process reduces the number of elements of P adjacent to more than one Steiner point, so repeating this completes the claim.

This process is consistent with Proposition 4.9 since none of the edges considered in that claim are added. It is consistent with Lemmas 4.10 and 4.12 since the degree of non root terminals only decreases in this process. ■

PROPOSITION 4.14. *We may assume that if $t \in P$ such that t is adjacent to a Steiner point s , then t is not adjacent to r .*

PROOF. Suppose that $(t, r) \in T_E$. Drop (t, r) and add edge (s, r) . This reconnects the tree and does not increase its cost since $\alpha_X \leq \alpha_p$ by (P1). As usual, repeating this process yields the desired claim. Moreover, this process is consistent with the above claims. We do not add an edges of the form to violate Proposition 4.9 or Lemma 4.13, and we only decrease the degree of terminals so this process is consistent with Lemmas 4.10 and 4.12. ■

LEMMA 4.15. *We may assume that $(s, s') \notin T_E$ for $s, s' \in X$.*

PROOF. First, if $\gamma_0, \gamma_1, \gamma_2 \geq \alpha_X$, then the result follows easily. Suppose $(s, s') \in T_E$ for $s, s' \in X$. Then, remove the edge between them. Assume w.l.o.g. that s is in the resultant connected component with r . Then, add edge (s', r) to reconnect the graph. Since $\gamma_i = \Delta(s, s') \geq \Delta(s', r) = \alpha_X$, the cost of the tree does not increase under this operation.

If instead $\gamma_0, \gamma_1 \geq \alpha_X$ but $\gamma_2 < \alpha_X$ (the only other case by (P1)), we proceed carefully. By applying the procedure above we may assume that $(s, s') \in T_E$ only if $|\Gamma(s) \cap \Gamma(s')| = 2$. Fix $s \in X$ with edges to other Steiner points. We may assume that s has at most 3 neighboring Steiner points. This is because each of s 's neighbors overlap $\Gamma(s) = \{i, j, k\}$ on two elements by Proposition 4.9, and there are three choices of pairs of elements to overlap on ($\{i, j\}$, $\{i, k\}$, or $\{j, k\}$). Hence, for all but one neighbor per pair of overlapping set elements, the edges to s can be adjusted to be edges to one another. For example, if s' and s'' are Steiner points adjacent to s such that $\{i, j\} \subset \Gamma(s'), \Gamma(s'')$, then $\Delta(s', s'') = \Delta(s, s') = \Delta(s, s'') = \beta_{\text{in}}$ so the edge (s, s'') may be replaced by the edge (s', s'') without changing the cost or connectivity of the tree.

Similarly, we may assume that s is not adjacent to r by moving an edge from s to r to be an edge from r to one of s 's neighboring Steiner points. Now we consider two cases based on the number of neighboring Steiner points to s .

Case 1: s is adjacent to two or three Steiner points. If s is adjacent to at least two Steiner points s' and s'' , then $\Gamma(s) \subseteq \Gamma(s') \cup \Gamma(s'')$ since $|\Gamma(s) \cap \Gamma(s')| = 2$, $|\Gamma(s) \cap \Gamma(s'')| = 2$, and $\Gamma(s), \Gamma(s'), \Gamma(s'')$ are distinct sets of size 3. Hence, using Proposition 4.9, for each non-root terminal adjacent to s , we may remove its edge to s and add a new edge to s' or s'' (depending on which induces less cost) without changing the connectivity or cost of T .

At this point, s has degree 2 or 3. If s has degree 2, s 's only neighbors are two Steiner points s' and s'' such that $|\Gamma(s') \cap \Gamma(s'')| = 1$. Drop s and its incident edges. Add an edge from one of s 's former neighbors to r to reconnect T . This induces an additive cost of $\alpha_X - 2\gamma_2 \leq 0$. This inequality holds since $2\gamma_2 \geq \gamma_1$ by the triangle inequality and $\alpha_X \leq \gamma_1$ by (P1).

So, we may assume that s has degree exactly 3. But consider removing s and its incident edges. This reduces the cost of T by $3\gamma_2$. To reconnect T , we add edges from r to a Steiner point in each of the connected components of T not containing r . Both components may be clearly contain Steiner points. This costs $2\alpha_X$. Hence, since $3\gamma_2/2 \geq \alpha_X$ by (P1), this does not increase the cost of the tree.

Case 2: s is adjacent to exactly one Steiner point. Suppose s is only adjacent to one other Steiner point, s' . Then, since $|\Gamma(s) \cap \Gamma(s')| = 2$, we can remove edges from s to all but one of its neighboring universe elements and replace those with edges from the terminals to s' without changing the connectivity or cost of T . Now, we have that s is not adjacent to r , s is adjacent to at most one Steiner point (our present case), and s is adjacent to at most one universe element. Then, remove s and its incident edges from T . If s was not adjacent to any terminals, the tree remains connected and its cost decreases (and note that s is not adjacent to r).

Otherwise, suppose that s was adjacent to a universe element t . If t had degree greater than 1, then, by Lemmas 4.12 and 4.13 and Proposition 4.14, t 's other neighbors are exactly leaf node universe elements. Remove the edges from those terminals to t and replace those

edges with edges to some other terminal adjacent to a Steiner point (for example, some terminal adjacent to s'). This fixes the cost of the tree.

Now, add an edge from t to either a terminal adjacent to a Steiner point or r (whichever edge is cheaper). This changes the cost of the tree by an additive factor of

$$\min(\alpha_p, \tau) - \gamma_2 - \beta_{\text{in}} \leq 0$$

and reconnects the tree. The inequality holds by (P4).

Repeating this process removes all edges between Steiner points since it reduces the number of edges between Steiner points by at least one in each iteration.

Now we check that this process is consistent with the above claims. We do not add any edges to violate Proposition 4.9. This process is consistent with Lemmas 4.10 and 4.12 since the only time we add an edge between non- r terminals in this process is in Case 2 and all such added edges are between a leaf node and a terminal adjacent to a Steiner point. This process is consistent with Lemma 4.13 since we only replace edges from non- r terminals to Steiner points with different Steiner point endpoints than adding additional edges to Steiner points. Finally, the only time we add an edge from a non- r terminal to r is in Case 2, and we only do this when the non- r terminal is a leaf node. Hence, this process is consistent with Proposition 4.14. ■

One important consequence of the above claims is that any path between two Steiner points in T must pass through r . That facilitates computing the cost of the tree. Before completing the proof, we prove a lemma that facilitates our analysis.

LEMMA 4.16. *We may assume that each Steiner point has degree at least 3 in T .*

PROOF. If there is a Steiner point of degree 1, dropping the Steiner point and its incident edge yields a lower cost tree, violating minimality of T . Now suppose that s is a Steiner point of degree 2 in T . By Lemma 4.15 and Proposition 4.9, s 's only possible neighbors are r and $t_i \in P$ such that $i \in \Gamma(t_i)$. One of the two neighbors must be r since T is connected and, by Lemma 4.12 and Proposition 4.14, none of s 's neighboring universe elements nor their neighboring universe elements may be adjacent to r . Then, s 's neighbors are precisely r and some universe element.

If $\alpha_p \leq \tau$, then we may assume that s 's neighboring universe element is a leaf node. In that case, drop s and its incident edges and connect s 's neighboring universe element to r . This does not increase the cost of the tree by the triangle inequality and maintains connectivity.

If $\tau < \alpha_p$, let t be s 's universe element neighbor. Remove s and its incident edges and add an edge from r to t . As above, this retains connectivity and does not increase the cost of the tree by the triangle inequality. Now, the tree must have an additional Steiner point by Lemma 4.11 (otherwise we contradict minimality). That Steiner point must be adjacent to a universe element \tilde{t} since otherwise it is degree 1. But replacing the edge (t, r) with an edge (t, \tilde{t}) then reduces the cost of the tree while preserving connectivity, contradicting minimality of the tree.

Hence, degree 2 Steiner points are only even possible when $\alpha_p \leq \tau$. In that case, this simple modification is evidently consistent with all of the above claims. ■

LEMMA 4.17. *The Steiner points in T adjacent to 3 terminals correspond exactly to a maximum packing of sets in \mathcal{S} .*

PROOF. By Proposition 4.9 and Lemma 4.13, we know that a given terminal t_i is adjacent to s only if $i \in \Gamma(s)$ and is adjacent to at most one such s . Then, since every $S \in \mathcal{S}$ has $|S| = 3$, the Steiner points adjacent to 3 terminals correspond exactly to a disjoint choice of sets from \mathcal{S} .

It remains to prove that this choice of sets is of maximum size. To show this, we need to understand the cost of T .

Partition the edges of T into three sets E_1, E_2 , and E_3 defined as follows: edges incident to universe elements not adjacent to Steiner points, edges incident to Steiner points adjacent to two universe elements, and edges incident to Steiner points adjacent to three universe elements. Lemma 4.10 and Proposition 4.14 imply that every edge without a leaf universe element involves a Steiner point. Then, Proposition 4.9, Lemma 4.16, and the fact that the sets in \mathcal{S} are each cardinality 3 shows that this accounts for all for all edges in T .

Now, let x_1 be the number of universe elements not adjacent to Steiner points, x_2 the number of universe elements adjacent to Steiner points adjacent to two universe elements, and x_3 the number of universe elements adjacent to Steiner points adjacent to three universe elements. By Lemma 4.10 and Proposition 4.14 combined with the restrictions on the adjacency of Steiner points from Proposition 4.9 and Lemma 4.15, every Steiner point must be adjacent to r since T is connected. Then, the total cost contributed by edges in E_1 is $\min(\alpha_p, \tau)x_1$, $(\alpha_x/2 + \beta_{in})x_2$ for E_2 , and $(\alpha_x/3 + \beta_{in})x_3$ for E_3 .

Then, the cost of T is precisely

$$\min(\alpha_p, \tau)x_1 + (\beta_{in} + \alpha_x/2)x_2 + (\beta_{in} + \alpha_x/3)x_3.$$

But, since $\beta_{in} + \alpha_x/3 < \min(\alpha_p, \tau)$ by (P3) and $\beta_{in} + \alpha_x/3 < \beta_{in} + \alpha_x/2$, this cost is minimized for maximizing x_3 . This means that the set of Steiner points in T adjacent to 3 terminals must correspond to a maximum packing of sets in \mathcal{S} (by minimality of the cost of T). ■

As a result of Lemma 4.17, we have the following.

COROLLARY 4.18. *In the soundness case we have*

$$\text{cost}_\Delta(T) \geq n(\alpha_x/3 + \beta_{in}) \cdot \left(1 - \varepsilon + \frac{(4\varepsilon - 2\delta)}{3} \cdot \frac{\min(\alpha_p, \tau, \beta_{in} + \alpha_x/2)}{\alpha_x/3 + \beta_{in}} + \frac{(2\delta - \varepsilon)}{3} \cdot \frac{\min(\alpha_p, \tau)}{\alpha_x/3 + \beta_{in}} \right).$$

PROOF. Define x_1, x_2 , and x_3 as in the proof of Lemma 4.17 and let P_1, P_2 , and P_3 partition P into the universe elements counted by x_1, x_2 , and x_3 . From the proof of Lemma 4.17, we have that the cost of the T is given by

$$\text{cost}_\Delta(T) = \min(\alpha_p, \tau)x_1 + (\beta_{in} + \alpha_x/2)x_2 + (\beta_{in} + \alpha_x/3)x_3.$$

From Lemma 4.17, this is minimized for maximum x_3 . Now, as stated in Lemma 4.17, the sets in \mathcal{S} corresponding to each Steiner point adjacent to three universe elements correspond to a packing of sets. Since we are in the soundness case, we then have $x_3 \leq (1 - \epsilon)n$. Since we are lower bounding the cost of T and maximizing x_3 minimizes the cost of T , suppose $x_3 = (1 - \epsilon)n$.

Now we make use of δ , describing how to construct a set cover \mathcal{S}' from T . For each Steiner point $s \in T$, add $\Delta(s)$ to \mathcal{S}' . For each remaining uncovered universe element, add some set covering that element. This yields a set cover of size at most $x_1 + x_2/2 + x_3/3$. Since we are in the soundness case (and using $x_3 = (1 - \epsilon)n$) we have that

$$x_1 + x_2/2 + (1 - \epsilon)(n/3) \geq (1 + \delta)(n/3)$$

and

$$x_1 + x_2 + (1 - \epsilon)n = n.$$

Then, this implies $(4\epsilon - 2\delta)/3 \cdot n \geq x_2$. Our bound on $\text{cost}_\Delta(T)$ then follows by maximizing x_2 when $\beta_{\text{in}} + \alpha_X/2 < \min(\alpha_p, \tau)$. ■

5. APX-hardness of DST in ℓ_p -metrics

In this section, we describe the embedding of the metric Steiner tree instances (described in the previous section), into \mathbb{R}^n equipped with the ℓ_p -metric for $p \in (1, \infty]$. The APX-hardness when $p = 1$ follows easily from the known APX-hardness of CST in ℓ_1 -metric or ℓ_0 -metric given in [60] (or see Theorem A.6 for a simplified reduction).

5.1 An Embedding into ℓ_p -metric Spaces

For a given p , our corresponding instance of DST in ℓ_p -metric will be as follows. For an instance $([n], \mathcal{S})$ of (ϵ, δ) -SP₃ where $|\mathcal{S}| = m$, let $P = \{\mathbf{e}_i : i \in [n]\}$. Then, the set of terminals are defined to be $\tilde{P} = P \cup \{\mathbf{0}\}$ (with $r = \mathbf{0}$ the distinguished “root” terminal). Let θ_p be a scalar defined for each p . Then, the set of facilities is defined to be

$$X = \{\theta_p \cdot (\mathbf{e}_i + \mathbf{e}_j + \mathbf{e}_k) : \{i, j, k\} \in \mathcal{S}\}.$$

We choose θ_p such that the following properties hold:

1. $0 < \theta_p \leq 1/2$ and
2. $3((1 - \theta_p)^p + 2\theta_p^p)^{1/p} + (3\theta_p^p)^{1/p} < 3$.

Note that this is possible since, for

$$f_p(x) = 3((1 - x)^p + 2x^p)^{1/p} + (3x^p)^{1/p} - 3$$

we have $f_p(0) = 0$ and $\frac{df_p(0)}{dx} = -3 + 3^{1/p} < 0$ for $p > 1$. This point configuration induces a complete graph $G = (\tilde{P} \cup X, E)$ with edge costs given by distances in the ℓ_p -metric.

Now, we check that the conditions of Theorem 4.5 hold. Clearly our choices of distances are metric compatible since they arise from a metric space.

REMARK 5.1. Note that, for $p \in (1, \infty)$,

$$\begin{aligned} \text{— } \alpha_X &= 3^{1/p}\theta_p, & \text{— } \gamma_0 &= 6^{1/p}\theta_p, \\ \text{— } \alpha_P &= 1, & \text{— } \gamma_1 &= 4^{1/p}\theta_p, \\ \text{— } \beta_{\text{in}} &= ((1 - \theta_p)^p + 2\theta_p^p)^{\frac{1}{p}}, & \text{— } \gamma_2 &= 2^{1/p}\theta_p, \text{ and} \\ \text{— } \beta_{\text{out}} &= (1 + 3\theta_p^p)^{\frac{1}{p}}, & \text{— } \tau &= 2^{1/p}. \end{aligned}$$

When $p = \infty$, we instead have

$$\begin{aligned} \text{— } \alpha_X &= \gamma_0 = \gamma_1 = \gamma_2 = \theta_p, \\ \text{— } \alpha_P &= \beta_{\text{out}} = \tau = 1, \text{ and} \\ \text{— } \beta_{\text{in}} &= 1 - \theta_p. \end{aligned}$$

PROPOSITION 5.2. For each $p \in (1, \infty]$, the corresponding tuple $(\alpha_X, \alpha_P, \beta_{\text{in}}, \beta_{\text{out}}, \gamma_0, \gamma_1, \gamma_2, \tau)$ from the above is Steiner embeddable.

PROOF. For all p , the tuple is clearly metric compatible since it arises from a metric space. It clear that (P1), (P2), (P3), and (P4) hold when $p = \infty$. Now we prove that these constraints hold when p is finite. Note first that

$$\begin{aligned} 3^{1/p}\theta_p &\leq ((1 - \theta_p)^p + 2\theta_p^p)^{1/p} \\ &< ((1 - \theta_p)^p + 2\theta_p^p)^{1/p} + 3^{1/p}\theta_p/3 < 1 \leq 2^{1/p}. \end{aligned}$$

The first inequality follows from $0 < \theta_p \leq 1/2$ and the last inequality follows from $3((1 - \theta_p)^p + 2\theta_p^p)^{1/p} + (3\theta_p^p)^{1/p} < 3$. This immediately implies $\alpha_X \leq \alpha_P, \beta_{\text{in}}, \tau$. Since $0 < \theta_p \leq 1/2$, $\beta_{\text{in}} < \beta_{\text{out}}$, and $\alpha_X \leq \gamma_1, \gamma_2$ is clear. Also observe that $st \leq 3\gamma_2/2$ if and only if $(3/2)^{\frac{1}{p}} \leq 3/2$ which holds for all $p \geq 1$. This proves (P1) holds. The inequality $\alpha_P \leq \beta_{\text{out}}$ is immediate (showing (P2)) and $\beta_{\text{in}} + \frac{\alpha_X}{3} < \min(\alpha_P, \tau)$ holds by our choice of θ_p (showing (P3)).

It then remains to show $\min(\alpha_P, \tau) \leq \beta_{\text{in}} + \gamma_2$, (P4). But note that

$$1 = (1 - \theta_p) + \theta_p \leq \beta_{\text{in}} + \gamma_2,$$

completing the proof. ■

Then, applying Theorem 4.5 to Proposition 5.2, we have the following.

THEOREM 5.3. Let $p \in (1, \infty)$. Suppose that

1. $0 < \theta_p \leq 1/2$,
2. $3[(1 - \theta_p)^p + 2\theta_p^p]^{\frac{1}{p}} + 3^{\frac{1}{p}}\theta_p < 3$,
3. (ε, δ) -SP₃ is NP-hard.

Then, DST in the ℓ_p -metric is NP-hard to approximate within a factor of

$$(1-\varepsilon) + \frac{(4\varepsilon - 2\delta)}{3} \cdot \frac{\min(1, [(1 - \theta_p)^p + 2\theta_p^p]^{\frac{1}{p}} + \frac{3^{1/p}}{2}\theta_p)}{[(1 - \theta_p)^p + 2\theta_p^p]^{\frac{1}{p}} + 3^{\frac{1-p}{p}}\theta_p} + \frac{(2\delta - \varepsilon)}{3} \cdot \frac{1}{[(1 - \theta_p)^p + 2\theta_p^p]^{\frac{1}{p}} + 3^{\frac{1-p}{p}}\theta_p} > 1.$$

In particular, for all $p \in (1, \infty)$, there exists θ_p such that the first two conditions above hold and Theorem 2.2 implies $(\varepsilon_0, \varepsilon_0/2)$ -SP₃ is NP-hard for some $\varepsilon_0 > 0$.

REMARK 5.4. Curiously, we observe that the importance of ε and δ depends dramatically on p . For $p \in (1, 1/\log_3(4/3))$, δ is sometimes irrelevant. This occurs when θ_p is chosen such that $1 < [(1 - \theta_p)^p + 2\theta_p^p]^{\frac{1}{p}}$. Intuitively, this means that in the soundness case, the optimal Steiner tree not only corresponds to a maximal set packing (as proved in Lemma 4.17), but also that the terminals (universe elements) that are not yet covered by the packing, would like to connect by an edge directly with the root in the optimal Steiner tree, instead of having two terminals (universe elements) connecting to a common Steiner point (which is then connected to the root). This happens for $\theta_2 = 1/6$, for example. See Figure 1.

By applying the Fréchet embedding to the hard instance in Theorem 4.5, we have that only δ is relevant to the hardness of DST in ℓ_∞ -metric (and not ε). In many other settings, it appears that both ε and δ play a role with larger δ permitting better hardness of approximation. Below we prove the non-dependency of ε for DST in ℓ_∞ -metric, without invoking the Fréchet embedding.

THEOREM 5.5. *If (ε, δ) -SP₃ is NP-hard, then DST in ℓ_∞ -metric is NP-hard to approximate within a factor of $1 + \delta/4$.*

PROOF. Note that $\theta_\infty = 1/2$ satisfies the required properties for Proposition 5.2. Then, from Corollary 4.7, it suffices to simply check that, for $p = \infty$, $\beta_{\text{in}} + \alpha_X/2 = \min(\alpha_p, \tau, \beta_{\text{in}} + \alpha_X/2)$, $\min(\alpha_p, \tau) = \alpha_X + \beta_{\text{in}}$, and $\beta_{\text{in}} = \alpha_X$. This is clear from observing the distances in Remark 5.1. ■

COROLLARY 5.6. *For all $p \in [1, \infty]$, DST in ℓ_p -metric is APX-hard.*

PROOF. The case of DST in ℓ_1 -metric follows from the known APX-hardness of CST in ℓ_1 -metric or ℓ_0 -metric given in [60] (or see Theorem A.6). The cases of $p \in (1, \infty)$ and $p = \infty$ follow from Theorems 5.3 and 5.5, respectively. ■

5.2 Hardness of Approximating DST in ℓ_p -metrics

Optimizing for the choice of θ_p in the construction from Proposition 5.2, we have the following.

COROLLARY 5.7. *If (ε, δ) -SP₃ is NP-hard, then for $p \in [\frac{1}{\log_3(4/3)}, \infty]$, DST in ℓ_p -metric is NP-hard to approximate within a factor*

$$1 + \varepsilon \left(\frac{1}{2} - \frac{1}{2 \cdot 3^{1/p}} \right) + 2\delta \left(\frac{1}{2 \cdot 3^{1/p}} - \frac{3}{8} \right).$$

In particular, we have DST in ℓ_p -metric is NP-hard to approximate within a factor of $1 + \varepsilon/8$.

PROOF. The second part of the claim occurs when $\delta = \varepsilon/2$. In this case, the second fractional term in Theorem 5.3 is nullified. Then, observe that, for a fixed p , while

$$[(1 - \theta_p)^p + 2\theta_p^p]^{\frac{1}{p}} + \frac{3^{1/p}}{2}\theta_p \leq 1,$$

the hardness of approximation factor is maximized for maximal θ_p . By our constraints on θ_p , $\theta_p \leq 1/2$. When $\theta_p = 1/2$ is valid we have

$$3[(1 - 1/2)^p + 2(1/2)^p]^{\frac{1}{p}} + 3^{\frac{1}{p}}(1/2) = 2 \cdot 3^{1/p} < 3.$$

This then holds for $\frac{1}{\log_3(3/2)} < p$. Additionally,

$$[(1 - \theta_p)^p + 2\theta_p^p]^{\frac{1}{p}} + \frac{3^{1/p}}{2}\theta_p \leq 1,$$

holds for $\theta_p = 1/2$ if and only if $3^{1/p} \cdot \frac{3}{4} \leq 1$ so $p \geq \frac{1}{\log_3(4/3)}$. Hence, for $p \in [\frac{1}{\log_3(4/3)}, \infty]$, the embedding with $\theta_p = 1/2$ yields that DST in ℓ_p -metric is hard to approximate within a factor of $1 + \varepsilon/8$.

If (ε, δ) -SP₃ is NP-hard for $\delta > \varepsilon/2$, we get improved hardness for $p \in [\frac{1}{\log_3(4/3)}, \infty]$ by using $\theta_p = 1/2$. Computing yields the claimed expression. ■

When $\delta > \varepsilon/2$, optimizing the hardness exactly is more challenging. This is because, for p finite, the maximum value of

$$\frac{1}{[(1 - \theta_p)^p + 2\theta_p^p]^{\frac{1}{p}}}$$

arises for $\theta_p < 1/2$, so the optimal value depends intimately on δ . Likewise, when $p \in (1, \frac{1}{\log_3(4/3)})$, we need $\theta_p < 1/2$. At that point even finding a permissible choice of θ_p requires an explicit computation.

We summarize the dependence of the hardness of DST in ℓ_p -metric on ε and δ (and therefore the dependence on the hardness of Set Packing and Set Cover) in Figure 1. The purple region, $p \in \{1, \infty\}$, is the range where there is no dependency on ε , and hence the reduction only uses the hardness of set packing (see Corollary 4.7 and Theorem A.6). In the red range, $[1/(\log_3(4/3)), \infty)$, we show in Corollary 5.7 a dependency on both ε and δ .

The behavior in the yellow range is less clear. For p such that the optimal choice of θ_p satisfies

$$[(1 - \theta_p)^p + 2\theta_p^p]^{\frac{1}{p}} + \frac{3^{1/p}}{2}\theta_p < 1,$$

the hardness depends both on δ and ε ; larger δ will yield increased hardness of approximation. This is, however, not the case for all values of p in the yellow region. The choice of $p = 2$ is a notable exception (see Corollary A.2). In the green region there is no dependence on δ . This is

because the optimal choice of θ_p satisfies

$$[(1 - \theta_p)^p + 2\theta_p^p]^{\frac{1}{p}} + \frac{3^{1/p}}{2}\theta_p \geq 1.$$

6. Reduction from CST to DST

In this section, we link CST to DST, proving that DST is computationally harder than CST. The sense in which this holds is formalized in Theorem 6.1.

THEOREM 6.1. *Let (X, d) be a metric space such that there exists an algorithm for CST in (X, d) that runs in $O(f(n))$ time for some computable function f and n the number of input terminals. Let P be an instance of CST in (X, d) , and let T^* be an optimal Steiner tree of P . Then, for any $\varepsilon > 0$, there exists a $\text{poly}(n, \varepsilon)$ -time algorithm outputting $X \subset X$ with the following properties.*

1. $|X| = \text{poly}(n, \varepsilon)$.
2. If T is the optimal Steiner tree on DST instance (P, X) , then

$$\text{cost}(T^*) \leq \text{cost}(T) \leq (1 + \varepsilon) \text{cost}(T^*).$$

To prove this, we appeal to a powerful structural result about optimal Steiner trees.

THEOREM 6.2 ([24, 7, 6]). *Let P be an instance of CST in metric space (X, d) and let T^* be an optimal Steiner tree of P . For any $\varepsilon > 0$, there exists constant C_ε and Steiner tree T' of P with the following properties.*

1. $\text{cost}(T^*) \leq \text{cost}(T') \leq (1 + \varepsilon) \text{cost}(T^*)$.
2. T' is formed by partitioning P into parts of size at most C_ε , finding optimal Steiner trees of each part, and connecting those Steiner trees via edges between their respective terminals.

This follows from Theorem 3.1 of [24] with $C_\varepsilon \leq 2^{1+\varepsilon^{-1}}$. The best possible C_ε can be derived from Theorem 3.2 of [7]. This result can also be derived from the proof of Theorem 3.2 of [6] (with somewhat weaker $C_\varepsilon \leq 2^{32/\varepsilon \ln(8/\varepsilon)}$).

PROOF OF THEOREM 6.1. Consider every subset of P of size at most $C_{\varepsilon/2}$ as an instance of CST in (X, d) . There are $\text{poly}(n, \varepsilon)$ such instances and optimal Steiner trees (optimal up to an arbitrarily small factor $1 + \varepsilon'$ due to numerical precision) can be computed in constant time (in ε and ε'). Let X be the union of the Steiner points used in the Steiner tree solutions for these instances. Note that $|X| = \text{poly}(n)$ since Steiner trees on n terminals may be assumed to use at most $n - 2$ Steiner points (this follows by the triangle inequality and a degree counting

argument—see Section 3.4 of [33], for example). Let T' be the near-optimal Steiner tree from Theorem 6.2. Let Q be some set of terminals in the partition of terminals induced by T' . Then, the optimal Steiner tree of DST instance (Q, X) costs at most a $(1 + \varepsilon')$ factor more than the cost of the optimal Steiner tree for CST instance Q . Then, since this holds for each part in the partition of terminals induced by T' from Theorem 6.2, if T is the optimal Steiner tree on DST instance (P, X) , then $\text{cost}(T) \leq (1 + \varepsilon') \text{cost}(T') \leq (1 + \varepsilon')(1 + \varepsilon/2) \text{cost}(T^*)$. Since ε' can be made arbitrarily small, this implies the desired result. ■

In particular, Theorem 6.1 holds for all ℓ_p -metric spaces and in any dimension. This shows that CST is essentially no harder than DST in those spaces.

COROLLARY 6.3. *For all $\varepsilon > 0$ and $p \in [1, \infty]$, if there exists a polynomial time $(1 + \alpha)$ -approximation algorithm for DST in the ℓ_p -metric, then there exists a polynomial time $(1 + \alpha + \varepsilon)$ -approximation algorithm for CST in the ℓ_p -metric.*

PROOF. From Theorem 6.1, it suffices to show that CST in the ℓ_p -metric admits an algorithm that runs in $O(f(n))$ time for some computable function f and n the number of input terminals. Given a set of terminals P with $|P| = n$, consider the subspace of dimension at most n formed by those points. Then, the approximation scheme of Arora (see 1.1.1 of [4]) runs in $O(f(n))$ time for a computable function f , yielding the result. ■

Theorem 6.1 and Corollary 6.3 give a construction for efficient approximation algorithms of CST using the literature on DST.

COROLLARY 6.4. *For all $\varepsilon > 0$ and $p \in [1, \infty]$, CST in the ℓ_p -metric admits a polynomial time $(\ln(4) + \varepsilon)$ -approximation algorithm.*

PROOF. Run the algorithm implicit in 6.1 and Corollary 6.3 to compute a DST instance. Then, apply the $(\ln(4) + \varepsilon')$ -approximation algorithm for DST from [10] and choose ε' to be appropriately small. ■

Acknowledgement

We would like to thank Fabrizio Grandoni for discussion about DST.

References

- [1] Amir Abboud, MohammadHossein Bateni, Vincent Cohen-Addad, Karthik C. S., and Saeed Seddighin. On Complexity of 1-Center in Various Metrics. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2023)*, volume 275 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 1:1–1:19, Dagstuhl, Germany. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. DOI (7, 50)
- [2] Sara Ahmadian, Ashkan Norouzi-Fard, Ola Svensson, and Justin Ward. Better guarantees for k-means and Euclidean k-median by primal-dual algorithms. *SIAM Journal on Computing*, 49(4), 2020. DOI (4)
- [3] Ernst Althaus, Joschka Kupilas, and Rouven Naujoks. On the low-dimensional Steiner minimum tree problem in Hamming metric. *Theoretical Computer Science*, 505:2–10, 2013. DOI (4)
- [4] Sanjeev Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998. DOI (3, 4, 7, 8, 44)
- [5] Christina Backes, Alexander Rurainski, Gunnar W. Klau, Oliver Müller, Daniel Stöckel, Andreas Gerasch, Jan Küntzer, Daniela Maisel, Nicole Ludwig, Matthias Hein, et al. An integer linear programming approach for finding deregulated subgraphs in regulatory networks. *Nucleic acids research*, 40(6):e43–e43, 2012. DOI (2)
- [6] Yair Bartal and Lee-Ad Gottlieb. Near-linear time approximation schemes for Steiner tree and forest in low-dimensional spaces. *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, STOC '21*, pages 1028–1041, Virtual Italy. ACM, 2021. DOI (4, 7, 15, 43)
- [7] Al Borchers and Ding-Zhu Du. The k-Steiner ratio in graphs. *SIAM Journal on Computing*, 26(3):857–869, 1997. DOI (7, 15, 43)
- [8] Marcus Brazil, Ronald L. Graham, Doreen A. Thomas, and Martin Zachariasen. On the history of the Euclidean Steiner tree problem. *Archive for history of exact sciences*, 68(3):327–354, 2014. DOI (3)
- [9] Jarosław Byrka, Fabrizio Grandoni, Thomas Rothvoß, and Laura Sanità. Steiner tree approximation via iterative randomized rounding. *Journal of the ACM (JACM)*, 60(1):1–33, 2013. DOI (3)
- [10] Jarosław Byrka, Thomas Pensyl, Bartosz Rybicki, Aravind Srinivasan, and Khoa Trinh. An improved approximation for k-median and positive correlation in budgeted optimization. *ACM Transactions on Algorithms*, 13(2), March 2017. DOI (4, 44)
- [11] Márcia R. Cerioli, Luerbio Faria, Talita O. Ferreira, Carlos A. J. Martinhon, Fábio Protti, and Bruce Reed. Partition into cliques for cubic graphs: planar case, complexity and approximation. *Discrete Applied Mathematics*, 156(12):2270–2278, 2008. DOI (5, 18)
- [12] Deeparnab Chakrabarty, Nikhil R. Devanur, and Vijay V. Vazirani. New geometry-inspired relaxations and algorithms for the metric steiner tree problem. *Mathematical Programming*, 130(1):1–32, 2011. DOI (3)
- [13] Deeparnab Chakrabarty, Maryam Negahbani, and Ankita Sarkar. Approximation algorithms for continuous clustering and facility location problems. *30th Annual European Symposium on Algorithms, ESA 2022, September 5–9, 2022, Berlin/Potsdam, Germany*, volume 244 of *LIPIcs*, 33:1–33:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. DOI (4)
- [14] Xiuzhen Cheng and Ding-Zhu Du. Steiner trees in industry, volume 11. Springer Science & Business Media, 2013. DOI (2)
- [15] Miroslav Chlebík and Janka Chlebíková. Complexity of approximating bounded variants of optimization problems. *Theoretical Computer Science*, 354(3):320–338, 2006. DOI (49)
- [16] Miroslav Chlebík and Janka Chlebíková. The Steiner tree problem on graphs: inapproximability results. *Theoretical Computer Science*, 406(3):207–214, 2008. DOI (3)
- [17] Jun-Dong Cho. Steiner tree problems in VLSI layout designs, *Steiner Trees in Industry*, pages 101–173. Springer, 2001. DOI (2)
- [18] F. R. K. Chung and R. L. Graham. A new bound for Euclidean Steiner minimal trees. *Annals of the New York Academy of Sciences*, 440(1):328–346, 1985. DOI (9)
- [19] Fan R. K. Chung and Edgar N. Gilbert. Steiner trees for the regular simplex. *Steiner Trees for the Regular Simplex*, 4(2):313–325, 1976. (9)
- [20] Vincent Cohen-Addad and Karthik C. S. Inapproximability of clustering in L_p metrics. *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9–12, 2019*, pages 519–539. IEEE Computer Society, 2019. DOI (4, 7, 51)
- [21] Vincent Cohen-Addad, Karthik C. S., and Euiwoong Lee. Johnson coverage hypothesis: inapproximability of k-means and k-median in ℓ_p -metrics. *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 – 12, 2022*, pages 1493–1530. SIAM, 2022. DOI (4, 7, 51)

- [22] Vincent Cohen-Addad, Karthik C. S., and Euiwoong Lee. On approximability of clustering problems without candidate centers. *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 2635–2648. SIAM, 2021. DOI (15)
- [23] William H. E. Day, David S. Johnson, and David Sankoff. The computational complexity of inferring rooted phylogenies by parsimony. *Mathematical Biosciences*, 81(1):33–42, 1986. DOI (3)
- [24] D.-Z. Du, Y. Zhang, and Q. Feng. On better heuristic for euclidean steiner minimum trees. *Proceedings of 32nd Annual Symposium of Foundations of Computer Science, FOCS '91*, pages 431–439, 1991. DOI (7, 15, 43)
- [25] Ding-Zhu Du and Warren D. Smith. Disproofs of Generalized Gilbert–Pollak Conjecture on the Steiner Ratio in Three or More Dimensions. *Journal of Combinatorial Theory, Series A*, 74(1):115–130, 1996. DOI (9)
- [26] Hai Du, Weili Wu, Zaixin Lu, and Yinfeng Xu. On the Steiner ratio in \mathcal{R}_n . *Discrete Mathematics, Algorithms and Applications*, 03(04):473–489, 2011. DOI (9)
- [27] David Fernández-Baca and Jens Lagergren. On the approximability of the steiner tree problem in phylogeny. *Discrete Applied Mathematics*, 88(1-3):129–145, 1998. DOI (4)
- [28] Henry L. Fleischmann, Surya Teja Gavva, and Karthik C. S. On approximability of steiner tree in ℓ_p -metrics. *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024, Alexandria, VA, USA, January 7-10, 2024*, pages 1669–1703. SIAM, 2024. DOI (1)
- [29] Les R Foulds and Ronald L Graham. The steiner problem in phylogeny is NP-complete. *Advances in Applied mathematics*, 3(1):43–49, 1982. DOI (3, 4)
- [30] Arnaldo Garcia and Henning Stichtenoth. On the asymptotic behaviour of some towers of function fields over finite fields. *Journal of Number Theory*, 61(2):248–273, 1996. DOI (52)
- [31] Michael R. Garey, Ronald L. Graham, and David S. Johnson. The complexity of computing Steiner minimal trees. *SIAM Journal on Applied Mathematics*, 32(4):835–859, 1977. DOI (3, 8)
- [32] Michael R. Garey and David S. Johnson. The rectilinear Steiner tree problem is NP-complete. *SIAM Journal on Applied Mathematics*, 32(4):826–834, 1977. DOI (3)
- [33] Edgar N. Gilbert and Henry O. Pollak. Steiner minimal trees. *SIAM Journal on Applied Mathematics*, 16(1):1–29, 1968. DOI (3, 8, 9, 44)
- [34] Maurice Hanan. On Steiner’s problem with rectilinear distance. *SIAM Journal on Applied Mathematics*, 14(2):255–265, 1966. DOI (3, 8)
- [35] Frank K. Hwang and Dana S. Richards. Steiner tree problems. *Networks*, 22(1):55–89, 1992. DOI (2)
- [36] Trey Ideker, Owen Ozier, Benno Schwikowski, and Andrew F Siegel. Discovering regulatory and signalling circuits in molecular interaction networks. *Bioinformatics*, 18(suppl_1):S233–S240, 2002. DOI (2)
- [37] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *Journal of Computer and System Sciences*, 62(2):367–375, 2001. DOI (7)
- [38] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001. DOI (7)
- [39] Alexandr O Ivanov and Alexey A Tuzhilin. The Steiner ratio Gilbert–Pollak conjecture is still open: clarification statement. *Algorithmica*, 62:630–632, 2012. DOI (9)
- [40] William B. Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Contemporary Mathematics*. Volume 26. American Mathematical Society, Providence, Rhode Island, 1984., pages 189–206. DOI (7)
- [41] Richard M. Karp. Reducibility among combinatorial problems, *Complexity of computer computations*, pages 85–103. Springer, 1972. DOI (3, 11)
- [42] Marek Karpinski and Alexander Zelikovskiy. New approximation algorithms for the Steiner tree problems. *Journal of Combinatorial Optimization*, 1(1):47–65, 1997. DOI (3)
- [43] Thomas Lengauer. Combinatorial algorithms for integrated circuit layout. Springer Science & Business Media, 2012. DOI (2)
- [44] Vladimir Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics-Doklady*, 10(8), 1966. (4)
- [45] Ivana Ljubić. Solving Steiner trees: recent advances, challenges, and perspectives. *Networks*, 77(2):177–204, 2021. DOI (2)
- [46] Joseph S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems. *SIAM Journal on Computing*, 28(4):1298–1309, 1999. DOI (3, 7)
- [47] Mohammad Noormohammadpour, Cauligi S. Raghavendra, Sriram Rao, and Srikanth Kandula. DCCast: efficient point to multipoint transfers across datacenters. *9th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 17)*, 2017. DOI (2)
- [48] Erez Petrank. The hardness of approximation: gap location. *Computational Complexity*, 4(2):133–157, 1994. DOI (17)
- [49] Hans Jürgen Prömel and Angelika Steger. A new approximation algorithm for the Steiner tree problem with performance ratio $5/3$. *Journal of Algorithms*, 36(1):89–101, 2000. DOI (3)
- [50] Harald Räcke. Lecture 12 of Theory of Metric Embeddings, 2006. URL (5)

- [51] Satish Rao and Warren D. Smith. Approximating geometrical graphs via “spanners” and “banyans”. *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998, STOC '98*, pages 540–550. ACM, 1998. DOI (3, 7)
- [52] Mauricio G. C. Resende and Panos M. Pardalos. *Handbook of optimization in telecommunications*. Springer Science & Business Media, 2008. DOI (2)
- [53] Gabriel Robins and Alexander Zelikovsky. Tighter bounds for graph Steiner tree approximation. *SIAM Journal on Discrete Mathematics*, 19(1):122–134, 2005. DOI (3)
- [54] Aviad Rubinfeld. Hardness of approximate nearest neighbor search. *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 1260–1268. ACM, 2018. DOI (7, 51)
- [55] Olga Russakovsky and Andrew Y. Ng. A Steiner tree approach to efficient object detection. *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1070–1077. IEEE, 2010. DOI (2)
- [56] David Sankoff. Minimal mutation trees of sequences. *SIAM Journal on Applied Mathematics*, 28(1):35–42, 1975. DOI (4)
- [57] Naveed A. Sherwani. *Algorithms for VLSI physical design automation*. Springer Science & Business Media, 2012. DOI (2)
- [58] K.W. Shum, I. Aleshnikov, P.V. Kumar, H. Stichtenoth, and V. Deolalikar. A low-complexity algorithm for the construction of algebraic-geometric codes better than the gilbert-varshamov bound. *IEEE Transactions on Information Theory*, 47(6):2225–2241, 2001. DOI (52)
- [59] Luca Trevisan. Non-approximability results for optimization problems on bounded degree instances. *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing, STOC '01*, pages 453–461, 2001. DOI (12)
- [60] Luca Trevisan. When Hamming meets Euclid: the approximability of geometric TSP and Steiner tree. *SIAM Journal on Computing*, 30(2):475–485, 2000. DOI (3, 8, 9, 39, 41)
- [61] Nurcan Tuncbag, Sara J. C. Gosline, Amanda Kedaigle, Anthony R. Soltis, Anthony Gitter, and Ernest Fraenkel. Network-based interpretation of diverse high-throughput datasets through the omics integrator software package. *PLoS computational biology*, 12(4):e1004879, 2016. DOI (2)
- [62] Vijay V. Vazirani. *Approximation algorithms*. Springer, Berlin; New York, 2001. DOI (3)
- [63] H. Todd Wareham. A Simplified Proof of the NP- and MAX SNP-Hardness of Multiple Sequence Tree Alignment. *Journal of Computational Biology*, 2(4):509–514, 1995. DOI (3, 4, 8, 9)
- [64] Alexander Z. Zelikovsky. An 11/6-approximation algorithm for the network Steiner problem. *Algorithmica*, 9(5):463–470, 1993. DOI (3)

A. Explicit Gaps for DST in Specific ℓ_p -metrics

In this section, we compute explicit hardness of approximation constants by plugging in the best inapproximability results known for SP_3 and VC in literature.

A.1 Hardness of Approximation Bounds for Euclidean DST

As the Euclidean metric is of particular importance, we explicitly state our resultant hardness factor for $p = 2$. In this special case, we can make do with a somewhat weaker completeness case for SP_3 than that in (ε, δ) - SP_3 . Instead we consider the decision problem (a, b) -Gap SP_3

DEFINITION A.1 ((a, b) -Gap SP_3). Given a set system $([n], \mathcal{S})$ where for all $S \in \mathcal{S}$, we have $|S| = 3$, and $|\mathcal{S}| = m$, the (a, b) -Gap SP_3 is the problem of deciding which of the following cases hold.

- *Completeness*: There exists a subcollection of \mathcal{S} which are pairwise disjoint and cover at least $(1 - a)n$ elements in $[n]$.
- *Soundness*: Any subcollection of pairwise disjoint sets of \mathcal{S} cover at most $(1 - b)n$ elements in $[n]$.

Observe that $(\varepsilon, \varepsilon/2)$ - SP_3 is actually $(1, 1 - \varepsilon)$ -Gap SP_3 since δ is unnecessary in this formulation. Then, we have the following.

COROLLARY A.2. *Let $\varepsilon > 0$. If $(1, 1 - \varepsilon)$ -Gap SP_3 is NP-hard then Euclidean DST is NP-hard to approximate within a factor of*

$$1 + \frac{\varepsilon(9\sqrt{3} - 15)}{15} \geq 1 + 0.039\varepsilon.$$

More generally, if (a, b) -Gap SP_3 is NP-hard, then Euclidean DST is NP-hard to approximate within a factor of

$$\frac{b(5\sqrt{3}/9 - 1) + 1}{a(5\sqrt{3}/9 - 1) + 1}.$$

PROOF. For $p = 2$, setting $\theta_2 = 1/6$ turns out to yield the maximum hardness of approximation for embeddings of this kind (this can be verified graphically). Note that $\theta_2 = 1/6$ satisfies the requirements of Theorem 5.3: the first requirement is clear and the second follows from $3\sqrt{3}/2 + \sqrt{3}/6 = \sqrt{3}(5/3) < 3$. Then, substitute into Theorem 5.3 and observe that

$$1 < \beta_{\text{in}} + \alpha_X/2 = \sqrt{3}/2 + \sqrt{3}/12 = 7\sqrt{3}/12,$$

as noted in Remark 5.4. This yields hardness of approximation within a factor of

$$1 - \varepsilon + \frac{\varepsilon}{\sqrt{3}/2 + \sqrt{3}/18} = 1 + \frac{9\sqrt{3} - 15}{15} \cdot \varepsilon,$$

proving the first part of the result.

The second part of the proof follows almost identically. We repeat almost the same reduction as in Theorem 4.5: The soundness case is exactly identical and the resultant Steiner tree then has cost at least

$$b(\sqrt{3}/2 + \sqrt{3}/18) + 1 - b = b(5\sqrt{3}/9 - 1) + 1.$$

In the completeness case, we no longer have that the sets partition the universe. Let C be a maximum packing of sets. For each $\{i, j, k\} \in C$, connect e_i , e_j , and e_k to $s = \Gamma^{-1}(\{i, j, k\})$ (following the notation of Section 4). Then connect s to $\mathbf{0}$. For the terminals corresponding to universe elements that are not covered by the sets in C , we connect those terminals directly to $\mathbf{0}$. This incurs costs $a(5\sqrt{3}/9 - 1) + 1$. Taking the ratio in the two cases then yields the desired result. ■

The below hardness of (a, b) - SP_3 is known.

THEOREM A.3 ([15]). *We have $(0.979, 0.969)$ -Gap SP_3 is NP-hard.*

In [15], they prove the hardness of the Max Independent Set problem on 3-regular graphs, a special case of SP_3 . The fact that this is a special case is clear by considering the set of incident edges to each vertex as a set in the set-system and the set of all edges in the graph as the universe of elements.⁸ Applying this to Corollary A.2 we have the following.

COROLLARY A.4. *Euclidean DST is NP-hard to approximate within a factor of 1.00039.*

A.2 Hardness of Approximation Bounds for Hamming and Rectilinear DST

The general embedding described in Section 5.1 notably fails for $p = 1$. However, a simple reduction from VC yields reasonable hardness of approximation bounds. We simultaneously get the same hardness for Hamming (ℓ_0) DST.

We will make use of the following theorem.

THEOREM A.5 ([15]). *$(0.52025, 0.53036)$ -VC is NP-hard on 4-regular graphs.*

We show the following.

THEOREM A.6. *Suppose that (a, b) -VC is NP-hard on max degree Δ graphs. Then DST in both the ℓ_0 and ℓ_1 metrics is NP-hard to approximate within a factor of $\frac{\Delta/2+b}{\Delta/2+a}$.*

⁸ The explicit bounds here comes from a combination of places in [15]: we use Theorem 16 for the bound on $\mu_{3,k}$, the assumption at the beginning of Theorem 17 for bounding $M(H)/k$ by $\mu_{3,k+o(1)}$, the NP-hard decision problem formulation of Max Independent Set stated in the proof of Theorem 17 for our two cases, and the property $|V(H)| = 2M(H)$ used to prove Corollary 18 to apply the bound on $M(H)/k$.

PROOF. The proof is exactly the same as the reduction from Set Cover to DST given in Section 1.2. Interpret the instance of VC as an instance of Set Cover with universe elements being edges and sets being the sets of incident edges to vertices. We can embed the graph used in the reduction in \mathbb{R}^n under the ℓ_0 and ℓ_1 metrics by mapping the “additional vertex” to $\mathbf{0}$, the edges to their characteristic vectors (of Hamming weight 2), and each set of edges incident to a vertex to the characteristic vector of that vertex (of Hamming weight 1). ■

Then, using Theorem A.5, we have the following result.

THEOREM A.7. *We have that DST in both the ℓ_0 and ℓ_1 metrics is NP-hard to approximate within a factor of 1.004.*

B. String Metrics

In this section, we show how APX-hardness of DST in the Hamming metric implies APX-hardness of DST in the Ulam and edit distance metrics, using known near-isometric embeddings. This is Theorems 1.5 and 1.6, respectively. Let Σ be an alphabet and Σ^* be the set of strings over the alphabet Σ . Then let $\text{ED} : \Sigma^* \times \Sigma^* \rightarrow \mathbb{R}$, the *edit distance* between strings. That is, the minimum number of single character deletions, insertions, and substitutions required to modify one string into the other. The Ulam metric UD is the same edit distance except with input strings restricted to those without repeated characters.

B.1 Ulam Metric

In [1], they show that there exists an embedding of n -dimensional Hamming space into the set of permutations of $[2n]$ under the Ulam metric such that all pairwise distances are exactly scaled by a factor of 2.

LEMMA B.1 (Lemma 4.5 of [1]). *Let Π_{2n} denote the set of permutations of $[2n]$. There is a function $\eta : \{0, 1\}^n \rightarrow \Pi_{2n}$, such that, for all $x, y \in \{0, 1\}^n$, we have*

$$\text{ed}(\eta(x), \eta(y)) = 2 \|x - y\|_0.$$

Moreover, for all $x \in \{0, 1\}^n$, $\eta(x)$ can be computed in $O(n)$ time.

Then, given a hard instance (P, X) of Hamming DST, consider $(\eta(P), \eta(X))$ as an instance of DST in the Ulam metric. In particular, there is a one-to-one correspondence between Steiner trees in the two instances and their costs are exactly scaled by 2. So, hardness is exactly scaled by 2 (using that $\eta(x)$ can be computed in $O(n)$ time). This completes the prove of Theorem 1.5.

B.2 Edit Distance Metric

In [54], Rubinfeld shows a near-isometric embedding of Hamming space into the space of binary strings under the edit distance metric. We use an explicit formulation given in Lemma A.1 of [20].

LEMMA B.2 ([54]). *For large enough d , there exists a function $\eta : \{0, 1\}^d \rightarrow \{0, 1\}^{d'}$, where $d' = O(d \log d)$ such that for all $a, b \in \{0, 1\}^d$, we have*

$$|\text{ED}(\eta(a), \eta(b)) - c \cdot \log d \cdot \|a - b\|_0| = o(d'),$$

for some constant c . Moreover, for any $a \in \{0, 1\}^d$, $\eta(a)$ can be computed in $2^{o(d)}$ time.

First, we use the fact that Hamming DST is APX-hard in $O(\log n)$ dimensions. This follows from Theorem C.1 in Section C. Then, given an instance (P, X) of Hamming DST in $O(\log n)$ dimensions, for n large enough, consider instance $(\eta(P), \eta(X))$ of ED DST (with η the function from Lemma B.2). This is computable in $\text{poly}(n)$ time by Lemma B.2.

Now, observe that the cost of optimal Steiner trees must be $\Omega(n)$ in both the completeness and soundness cases of Hamming DST (since the minimum distance between any two points is 1). Now, let T be a Steiner tree of (P, X) . By Lemma B.2, the corresponding Steiner tree of $(\eta(P), \eta(X))$ costs at most $c \cdot \log \log n \cdot \text{cost}_0(T) + o(\log n \log \log n)$ and at least $c \cdot \log \log n \cdot \text{cost}_0(T) - o(\log n \log \log n)$. Then, since $\text{cost}_0(T) = \Omega(n)$ the ratio of the bounds on the costs of optimal trees in the soundness and completeness cases of the ED DST are preserved for sufficiently large n since the $\Omega(n)$ term dominates the $o(\log n \log \log n)$ term. Hence, since there is a constant gap for Hamming DST, there is a constant gap for ED DST. This completes the proof of Theorem 1.6.

C. Inapproximability in Low Dimensions

In Section 5.1, we describe gap-preserving reduction from (ε, δ) -SP₃ to DST in n -dimensional ℓ_p -metric spaces. In this section, we show that this reduction can be extended to a configuration of n points in an $O(\log n)$ -dimensional ℓ_p -metric space (for constant p) via a near-isometric embedding from [20, 21]. For the dimensionality reduction for $p = \infty$, refer to [20]. This will show that DST in the ℓ_p -metric is APX-hard even in $O(\log n)$ -dimensions (and essentially preserve the hardness of approximation). In the below, let $\text{supp}(x)$ be the set of nonzero coordinates of x (the *support* of x).

THEOREM C.1 (Implicit in [20, 21]). *Let $\varepsilon > 0$, $B, C \in \mathbb{Z}^+$, and $p \in \mathbb{R}_{\geq 1}$. Let $P \subset \mathbb{R}^n$ such that $|P| = n$ and, for all $x \in P$, $\text{supp}(x) \subseteq [n]$ with $|\text{supp}(x)| \leq C$ and $\|x\|_\infty \leq B$. Then, there exists $\alpha = O_{C,B,\varepsilon}(1/\log n)$ and $\sigma : \mathbb{R}^n \rightarrow \mathbb{R}^{O_{C,B,\varepsilon}(\log n)}$ such that, for all $x, y \in P$,*

$$(1 - \varepsilon) \|x - y\|_p^p \leq \alpha \|\sigma(x) - \sigma(y)\|_p^p \leq (1 + \varepsilon) \|x - y\|_p^p.$$

Additionally, for all $x \in P$, $\sigma(x)$ is computable in $\text{poly}(n)$ time.

Note that all points in instances of DST in Section 5.1 have support of size at most 3. Hence, Theorem C.1 applies. Hence, for a fixed p and any $\varepsilon > 0$, there exists an instance of DST in the ℓ_p -metric in $O(\log n)$ dimensions with all pairwise distances within an arbitrarily small factor of the original distances in n dimensions (after scaling the configuration by $\alpha^{1/p}$). Therefore, the cost of the trees in the completeness and soundness cases change by at most that factor and we get hardness of approximation within an arbitrarily small factor of the original hardness.

To prove Theorem C.1, we use the existence of codes from algebraic geometry.

THEOREM C.2 (Existence of AG codes, [30, 58]). *Let q be a prime squared at least 49. Then, there exists a linear map $C : \mathbb{F}_q^{\log_q n} \rightarrow \mathbb{F}_q^{c \log_q n}$ for some constant c , such that for all $x, y \in \mathbb{F}_q^{\log_q n}$ we have*

$$\|C(x) - C(y)\|_0 \geq (1 - 3/\sqrt{q})(c \log_q n).$$

To define the map $\sigma : \mathbb{R}^n \rightarrow \mathbb{R}^{O(\log n)}$ for Theorem C.1, we first describe the map on $\{0, 1\}_1^n \subset \mathbb{R}^n$, the set of vectors of Hamming weight 1 (the standard basis vectors). Let q be a squared prime at least 49 (we will specify q appropriately later). Given $i \in [n]$, map it bijectively to an element $x \in \mathbb{F}_q^{\log_q n}$ (by enumerating the elements of $\mathbb{F}_q^{\log_q n}$). Call this map f_1 . Let C be the linear map existing from Theorem C.2. Then, apply C to x to yield an element of $\mathbb{F}_q^{c \log_q n}$. Finally, enumerate the elements of \mathbb{F}_q from 1 to q and map $C(f_1(x))$ to $\{0, 1\}^{qc \log_q n}$ by interpreting $C(f_1(x))$ as $c \log_q n$ blocks of q zeroes and ones, $q - 1$ of which being 0 and the 1 indicating the number of the element from \mathbb{F}_q . Call this map f_2 . For example, $(3, 1, 2) \in \mathbb{F}_5^3$ would be mapped to the bitstring $000100100000100 \in \{0, 1\}^{15}$. We interpret this bitstring in $\{0, 1\}^{qc \log_q n}$ as a vector in $\mathbb{R}^{qc \log_q n}$ by the natural inclusion map. Call this map f_3 . Formally, $\sigma = f_3 \circ f_2 \circ C \circ f_1$ for inputs in $\{0, 1\}_1^n$. To extend σ to the rest of \mathbb{R}^n , we define

$$\sigma \left(\sum_{i=1}^n a_i \mathbf{e}_i \right) = \sum_{j=1}^{qc \log_q n} \begin{cases} a_i \cdot \mathbf{e}_j & \text{if } |\{i \in [n] : a_i \sigma(\mathbf{e}_i)_j \neq 0\}| = 1, \\ 0 & \text{otherwise.} \end{cases}$$

Intuitively, given the large minimum distance between codewords from Theorem C.2 and the bounded support of point configurations in Theorem C.1, for most $j \in [qc \log_q n]$, $|\{i \in [n] : \sigma(\mathbf{e}_i)_j = 1\}| \in \{0, 1\}$.

Now, let ε, C, B, p , and P be as in Theorem C.1. Fix $x, y \in P$. Let $x = \sum_{i \in I_1} x_i \mathbf{e}_i$ and $y = \sum_{i \in I_2} y_i \mathbf{e}_i$, where $|I_1|, |I_2| \leq C$. Fix $i \in I_1$. Consider the set of nonzero coordinates on which $\sigma(\mathbf{e}_i)$ does not overlap with $\sigma(\mathbf{e}_j)$ for $i \neq j \in I_1 \cup I_2$. The contribution of each such coordinate k to $\|\sigma(x) - \sigma(y)\|_p^p$ is exactly $|x_i - y_i|^p$.

By Theorem C.2, $\|C(f_1(\mathbf{e}_i)) - C(f_1(\mathbf{e}_j))\|_0 \geq (1 - 3/\sqrt{q})c \log_q n$ for each $i \neq j \in [n]$. Then, since the map f_2 exactly doubles Hamming distance and f_3 preserves Hamming distance, $\|\sigma(\mathbf{e}_i) - \sigma(\mathbf{e}_j)\|_0 \geq (1 - 3/\sqrt{q})2c \log_q n$. In particular, since $\|\sigma(\mathbf{e}_i)\|_0$ is exactly $c \log_q n$, this implies

that $\sigma(\mathbf{e}_i)$ and $\sigma(\mathbf{e}_j)$ can overlap on at most $(3/\sqrt{q})c \log_q n$ nonzero coordinates. Moreover, $\sigma(\mathbf{e}_i)$ can overlap on at most $(6C/\sqrt{q})c \log_q n$ total nonzero coordinates with $\sigma(\mathbf{e}_j)$ for $i \neq j \in I$. The maximum contribution of these overlapping coordinates is B^p (if, for example, $\sigma(\mathbf{e}_i)$ overlaps with $\sigma(\mathbf{e}_{i'})$ for $i' \in I_1$ but doesn't overlap with any mapped standard basis vector from I_2 and $a_i = B$). The minimum contribution is 0. Then, we have

$$(1 - 6C/\sqrt{q})c \log_q n \|x - y\|_p^p \leq \|\sigma(x) - \sigma(y)\|_p^p \leq c \log_q n \|x - y\|_p^p + (6C/\sqrt{q})B^p c \log_q n.$$

Setting q sufficiently large then yields Theorem C.1.