

A PTAS FOR THE STEINER FOREST PROBLEM IN DOUBLING METRICS*

T-H. HUBERT CHAN[†], SHUGUANG HU[†], AND SHAOFENG H.-C. JIANG[†]

Abstract. We achieve a (randomized) polynomial-time approximation scheme (PTAS) for the Steiner forest problem in doubling metrics. Before our work, a PTAS was given only for the Euclidean plane in [G. Borradaile, P. N. Klein, and C. Mathieu, in FOCS, IEEE Computer Society, 2008, pp. 115–124]. Our PTAS also shares similarities with the dynamic programming for sparse instances used in [Y. Bartal, L. Gottlieb, and R. Krauthgamer, in STOC, ACM, 2012, pp. 663–672] and [T-H. H. Chan and S.-H. Jiang, in SODA, SIAM, 2016, pp. 754–765]. However, extending previous approaches requires overcoming several nontrivial hurdles, and we make the following technical contributions. (1) We prove a technical lemma showing that Steiner points have to be “near” the terminals in an optimal Steiner tree. This enables us to define a heuristic to estimate the local behavior of the optimal solution, even though the Steiner points are unknown in advance. This lemma also generalizes previous results in the Euclidean plane and may be of independent interest for related problems involving Steiner points. (2) We develop a novel algorithmic technique known as “adaptive cells” to overcome the difficulty of keeping track of multiple components in a solution. Our idea is based on but significantly different from the previously proposed “uniform cells” in [G. Borradaile, P. N. Klein, and C. Mathieu, in FOCS, IEEE Computer Society, 2008, pp. 115–124], where techniques cannot be readily applied to doubling metrics.

Key words. approximation algorithms, doubling metrics, Steiner forest problem

AMS subject classifications. 68W25, 68W20

DOI. 10.1137/16M1107206

1. Introduction. We consider the Steiner forest problem (SFP) in a metric space (X, d) . An instance of the problem is given by a collection W of n terminal pairs $\{(a_i, b_i) : i \in [n]\}$ in X , and the objective is to find a minimum weight graph $F = (V, E)$ (where V is a subset of X and the edge weights are induced by the metric space) such that every pair in W is connected in F .

1.1. Problem background. The problem is well known in the computer science community. In general metrics, Chlebík and Chlebíková [13] showed that SFP is **NP**-hard to approximate with a ratio better than $\frac{96}{95}$. The best known approximation ratio achievable in polynomial time is 2 [19, 2]. Recently, Gupta and Kumar [21] gave a purely combinatorial greedy-based algorithm that also achieves a constant ratio. However, it is still an open problem to break the 2-approximation barrier in general metrics for SFP.

SFP in Euclidean plane and planar graphs. In light of the aforementioned hardness result [13], restrictions are placed on the metric space to achieve $(1 + \epsilon)$ approximation in polynomial time. In the Euclidean plane, a randomized polynomial-time approximation scheme (PTAS) was obtained in [10] using the dynamic programming framework proposed by Arora [3]. Later, a simpler analysis was presented in [8], in which a new structural property was proved and additional information was incor-

*Received by the editors December 8, 2016; accepted for publication (in revised form) June 11, 2018; published electronically August 29, 2018. The conference version of this paper appeared in FOCS, IEEE Computer Society, 2016.

<http://www.siam.org/journals/sicomp/47-4/M110720.html>

Funding: This work was partially supported by a grant from the Hong Kong RGC under contract 17217716.

[†]Department of Computer Science, The University of Hong Kong, Hong Kong (hubert@cs.hku.hk, sghu@cs.hku.hk, sfjiang@cs.hku.hk).

porated in the dynamic programming algorithm. It was only suggested that similar techniques might be applicable to higher-dimensional Euclidean space.

Going beyond the Euclidean plane, a PTAS for planar graphs was obtained in [9] and, more generally, on bounded genus graphs. As a building block, they also obtained a PTAS for graphs with bounded treewidth.

Steiner tree problems. A notable special case of SFP is the Steiner tree problem (STP), in which all terminals are required to be connected. In general metrics, the minimum spanning tree (MST) on the terminal points simply gives a 2-approximation. There is a long line of research to improve the 2-approximation, and the state-of-the-art approximation ratio 1.39 was presented in [11] via an LP rounding approach. On the other hand, it is **NP**-hard to approximate STP with a ratio better than $\frac{96}{95}$ [13].

For the group STP¹ in general metrics, it is **NP**-hard to approximate within $\log^{2-\epsilon} n$ [22] unless $\text{NP} \subseteq \text{ZTIME}(n^{\text{polylog}(n)})$. On the other hand, it is possible to approximate within $O(\log^3 n)$, as shown in [18]. Restricting ourselves to planar graphs, the group STP can be approximated within $O(\log n \text{ poly log log } n)$ [14], and very recently this result was improved to a PTAS [7].

For more related works, we refer the reader to a survey by Hauptmann and Karpiński [23], who gave a comprehensive literature review of STP and its variations.

PTASs for other problems in doubling metrics. Doubling dimension captures the local growth rate of a metric space. A k -dimensional Euclidean dimension has doubling dimension $O(k)$. A challenge in extending algorithms for low-dimensional Euclidean space to doubling metrics is the lack of geometric properties in doubling metrics. Although quasi-PTASs (QPTASs) for various approximation problems in doubling metrics, such as the traveling salesman problem (TSP) and STP, were presented in [25], a PTAS was only recently achieved for TSP [6]. Subsequently, a PTAS was also achieved for group TSP in doubling metrics [12]. Before this work, the existence of a PTAS for SFP (or even the special case STP) in doubling metrics remained an open problem.

1.2. Our contribution and techniques. Although PTASs for TSP (and its group variant) were known, as we shall explain later, the nature of SFP- and TSP-related problems is quite different. Hence, it is interesting to investigate what new techniques are required for SFP. Fundamentally, it is an important question whether the notion of doubling dimension captures sufficient properties of a metric space to design a PTAS for SFP, even without the geometric properties that are crucially used in obtaining approximation schemes for SFP in the Euclidean plane [10].

In this paper, we settle this open problem by giving a (randomized) PTAS for SFP in doubling metrics. We remark that previously even a PTAS for SFP in higher-dimensional Euclidean space was not totally certain, as the authors of [10] only gave a PTAS for the Euclidean plane, and whether their approach may be generalized to higher dimensions is unknown.

THEOREM 1.1 (PTAS for SFP in doubling metrics). *For any $0 < \epsilon < 1$, there is a (randomized) algorithm that takes an instance of SFP with n terminal pairs in a metric space with doubling dimension at most k , and returns a $(1 + \epsilon)$ -approximate solution with constant probability, running in time $O(n^{O(1)^k}) \cdot \exp(\sqrt{\log n} \cdot O(\frac{k}{\epsilon})^{O(k)})$.*

We next give an overview of our techniques. On a high level, we use the divide and conquer framework that was originally used by Arora [3] to achieve a PTAS for

¹The problem generalizes the STP, in such a way that each terminal point is replaced by a group of points, and the goal is to find the minimum graph that visits at least one point from each group.

TSP in Euclidean space and was extended recently to doubling metrics [6].

However, we shall explain that it is nontrivial to adapt this framework to SFP and explain how we overcome the difficulties encountered. Moreover, we shall provide some insights regarding the relationship between Euclidean and doubling metrics and discuss the implications of our technical lemmas.

Summary of framework. As in [6], a PTAS is designed for a class of special instances known as *sparse* instances. Then, it can be shown that the general instances can be decomposed into sparse instances. Roughly speaking, an instance is sparse if there is an optimal solution such that for any ball B with radius r , the portion of the solution in B has weight that is small with respect to r .

The PTAS for the sparse instances is usually based on a dynamic program, which is based on a randomized hierarchical decomposition, as in [25, 6]. This framework has also been successfully applied to achieve a PTAS for group TSP in doubling metrics [12]. Intuitively, sparsity is used to establish the property that with high enough probability, a cluster in the randomized decomposition cuts a (near) optimal tour only a few times [6, Lemma 3.1]. However, SFP brings new significant challenges when such a framework is applied. We next describe the difficulties and give an overview of our technical contributions.

Challenge 1: It is difficult to detect a sparse instance, because which Steiner points are used by the optimal solution are unknown. Let us first consider STP, which is a special case of SFP in which all (pairs of) terminals are required to be connected. In other words, the optimal Steiner tree is the minimum weight graph that connects all terminals. Unlike TSP, in which the points visited by a tour are clearly known in advance, it is not known which points will be included in the optimal Steiner tree.

In [6], a crucial step was to estimate the sparsity of a ball B , which measures the weight of the portion of the optimal solution restricted to B . For the TSP tour, this can be estimated from the points inside B that have to be visited. However, for solution involving Steiner points, it is difficult to analyze the solution inside some ball B , because it is possible that there are few (or even no) terminals inside B , but the optimal solution could potentially have lots of Steiner points and a large weight inside B .

Our solution: Analyzing the distribution of Steiner points in an optimal Steiner tree in doubling metrics. We resolve this issue by showing a technical characterization of Steiner points in an optimal Steiner tree for doubling metrics. This technical lemma is used crucially in our proofs, and we remark that it could be of interest for other problems involving Steiner points in doubling metrics.

LEMMA 1.2 (formal version in Lemma 3.1). *For a terminal set S with diameter D , if an optimal Steiner tree spanning S has no edge longer than γD , then every Steiner point in the solution is within $O(\sqrt{\gamma}) \cdot D$ distance to some terminal in S , where the big O hides the dependence on the doubling dimension.*

We observe that variants of Lemma 1.2 have been considered on the Euclidean plane. In [16, 17], it was shown that if the terminal set consists of n evenly distributed points on a unit circle, then for large enough n , there are no Steiner points in an optimal Steiner tree. To see how this relates to our lemma, when n is sufficiently large, it follows that adjacent points in the circle are very close to each other. Hence, any long edge in a Steiner tree could be replaced by some short edge between adjacent terminals in the circle. Our lemma then implies that all Steiner points must be near the terminals, which is a weaker statement than the conclusion in [16] but is enough

for our purposes. We emphasize that the results in [16, 17] rely on the geometric properties of the Euclidean plane. However, in our lemma, we only use that the doubling dimension is bounded.

Implication of Lemma 1.2 on sparsity heuristic. We next demonstrate an example of how we use this technical lemma. In Lemma 3.3, we argue that our sparsity heuristic provides an upper bound on the weight of the portion of an optimal solution F within some ball B .

The idea is that we remove the edges in F within B and add back some edges of small total weight to maintain connectivity. We first add an MST H on some net-points N within B of an appropriate scale $\gamma \cdot D$. Using the property of doubling dimension, we argue that the number of points in H is bounded and so is its weight. In one of our case analyses, there are two sets S and T of terminals that are far apart $d(S, T) \geq D$, and we wish to argue that in the optimal Steiner tree F connecting S and T , there is an edge $\{u, v\}$ of length at least $\Omega(\gamma) \cdot D$. If this is the case, we could remove this edge and connect u and v to their corresponding net-points directly. For contradiction's sake, we assume there is no such edge, but Lemma 1.2 implies that every Steiner point must be close to either S or T . Since S and T are far apart, this means that there is a long edge after all.

Conversely, in Lemma 3.5, we also use this technical lemma to show that if the sparsity heuristic for some ball B is large, then the portion of the optimal solution F inside B is also large.

Challenge 2: In doubling metrics, the number of cells for keeping track of connectivity in each cluster could be too large. Unlike the case for TSP variants [6, 12], the solution for SFP need not be connected. Hence, in the dynamic programming algorithm for SFP, in addition to keeping track of what *portals* are used to connect a cluster to points outside, we need to keep information on which portals the terminals inside a cluster are connected. In previous works [10], the notion of *cells* was used for this purpose.

Previous technique: Cell property. The idea of *cell property* was first introduced in [10], which gave a PTAS for SFP in the Euclidean plane using dynamic programming. Since there would have been an exponential number of dynamic program entries if we kept information on which portal was used by every terminal to connect to its partner outside the cluster, the high level idea is to partition a cluster into smaller clusters (already provided by the hierarchical decomposition) known as *cells*. Loosely speaking, the *cell property* ensures that every terminal inside the same cell must be connected to points outside the cluster in the same way. More precisely, a solution F satisfies the cell property if for every cluster C and every cell e inside C , there is only one component in the portion of F restricted to C that connects e to points outside C .

A great amount of work was actually needed in [10] and a subsequent work [8] to show that it is enough to consider cells whose diameters are constant times smaller than that of its cluster. This allows the number of dynamic program entries to be bounded, which is necessary for a PTAS.

Difficulty encountered for doubling metrics. When the notion of cell is applied to the dynamic program for SFP in doubling metrics, an important issue is that the diameters of cells need to be about $\Theta(\log n)$ times smaller than that of its cluster, because there are around $\Theta(\log n)$ levels in the hierarchical decomposition. Hence, the number of cells in a cluster is $\Omega(\text{poly } \log n)$, which would eventually lead to a QPTAS only. A similar situation is observed when dynamic programming was first used for TSP on doubling metrics [25]. However, the idea of using sparsity as in [6] does not seem to immediately provide a solution.

Our solution: Adaptive cells. Since there are around $\Theta(\log n)$ levels in the hierarchical decomposition, it seems very difficult to increase the diameter of cells in a cluster. Our key observation is that the cells are needed only for covering the portion of a solution inside a cluster that touches the cluster boundary. Hence, we use the idea of *adaptive cells*. Specifically, for each connected component A in the solution crossing a cluster C , we define the corresponding *basic cells* such that if the component A has larger weight, then its corresponding basic cells (with respect to cluster C) will have larger diameters. Combining with the notion of sparsity and bounded doubling dimension, we can show that we only need to pay attention to a small number of cells.

Further cells for refinement. Since the dynamic program entries are defined in terms of the hierarchical decomposition and the entries for a cluster are filled recursively with respect to those of its child clusters, we would like the cells to have a *refinement* property; i.e., if a cluster C has some cell e (which itself is some descendant cluster of C), then the child C' containing e has either e or all children of e as its cells.

At first glance, a quick fix may be to push down each basic cell in C to its child clusters. Although we could still bound the number of relevant cells, it would be difficult to bound the cost to achieve the cell property. The reason is that the basic cells from higher levels are too large for the descendant clusters. When more than one relevant component intersects such a large cell, we need to add edges to connect the components. However, if the diameter of the cell is too large compared to the cluster, these extra edges would be too costly.

We resolve this issue by introducing *nonbasic* cells for a cluster: *promoted* cells and *virtual* cells. These cells are introduced to ensure that every sibling of a basic cell is present. Moreover, only nonbasic cells of a cluster will be passed to its children. We show in Lemma 5.16 that the total number of *effective* cells for a cluster is not too large. Moreover, Lemma 5.13 shows that the refinement property still holds even if we only pass the nonbasic cells down to the child clusters. More importantly, we show that as long as we enforce the cell property for the basic cells, the cell property for all cells are automatically ensured. This means that it is sufficient to bound the cost to achieve the cell property with respect to only the basic cells.

Further techniques: Global cell property. We note that the cell property in [10] is localized. In particular, for each cluster C , we restrict the solution inside C , which could have components disconnected within C but are actually connected globally. In order to enforce the localized cell property as in [10], extra edges would need to be added for these locally disconnected components. Instead, we enforce a *global* cell property, in which for every cell e in a cluster C , there is only one (global) connected component in the solution that intersects e and crosses the boundary of cluster C . A consequence of this is that if there are m components in the solution, then at most $m - 1$ extra edges are needed to maintain the global cell property. This implication is crucially used in our charging argument to bound the cost for enforcing the cell property for the basic cells. However, this would imply that in the dynamic program entries, we need to keep additional information on how the portals of a cluster are connected outside the cluster.

Combining the ideas: A more sophisticated dynamic program. Even though our approaches to tackling the encountered issues are intuitive, it is a nontrivial task to balance between different tradeoffs and keep just enough information in the dynamic program entries but still ensure that the entries can be filled in polynomial time.

2. Preliminaries. We consider a metric space $M = (X, d)$ (see [15, 24] for more details on metric spaces). For $x \in X$ and $\rho \geq 0$, a *ball* $B(x, \rho)$ is the set $\{y \in X \mid d(x, y) \leq \rho\}$. The *diameter* $\text{Diam}(Z)$ of a set $Z \subset X$ is the maximum distance between points in Z . For $S, T \subset X$, we denote $d(S, T) := \min\{d(x, y) : x \in S, y \in T\}$, and for $u \in X$, $d(u, T) := d(\{u\}, T)$. Given a positive integer m , we denote $[m] := \{1, 2, \dots, m\}$.

A set $S \subset X$ is a ρ -packing if any two distinct points in S are at a distance more than ρ away from each other. A set S is a ρ -cover for $Z \subseteq V$ if for any $z \in Z$, there exists $x \in S$ such that $d(x, z) \leq \rho$. A set S is a ρ -net for Z if S is a ρ -packing and a ρ -cover for Z . We assume that a ρ -net for any ball in X can be constructed efficiently.

We consider metric spaces with *doubling dimension* [5, 20] at most k ; this means that for all $x \in X$, for all $\rho > 0$, every ball $B(x, 2\rho)$ can be covered by the union of at most 2^k balls of the form $B(z, \rho)$, where $z \in X$. The following fact captures a standard property of doubling metrics.

FACT 2.1 (packing in doubling metrics [20]). *Suppose that in a metric space with doubling dimension at most k , a ρ -packing S has diameter at most R . Then, $|S| \leq (\frac{2R}{\rho})^k$.*

We consider undirected graphs $G = (V, E)$ in the metric space, where $V \subset X$, $E \subseteq \binom{V}{2}$, and an edge $e = \{x, y\} \in E$ receives weight $d(x, y)$ from the metric space M . The weight $w(G)$ or cost of a graph is the sum of its edge weights. Let $V(G)$ denote the vertex set of a graph G .

We consider the Steiner forest problem (SFP). Given a collection $W = \{(a_i, b_i) \mid i \in [n]\}$ of terminal pairs in X , the goal is to find an undirected graph F (having vertex set in X) with minimum cost such that each pair of terminals is connected in F . The nonterminal vertices in $V(F)$ are called Steiner points.

Rescaling instance. Fix constant $\epsilon > 0$. Since we consider asymptotic running time to obtain $(1 + \epsilon)$ -approximation, we consider sufficiently large $n > \frac{1}{\epsilon}$. Suppose $R > 0$ is the maximum distance between a pair of terminals. Then R is a lower bound on the cost of an optimal solution. Moreover, the optimal solution F has cost at most nR , and hence we do not need to consider distances larger than nR . Observe that each Steiner point in F has degree at least 3, so the number of Steiner points is at most the number of terminal points. Moreover, there are at most $2n$ terminal points, since $|W| = n$. Hence, F contains at most $4n$ vertices. Because of this, if we consider an $\frac{\epsilon R}{32n^2}$ -net S for X and replace every point in F with its closest net-point in S , the cost increases by at most $\epsilon \cdot \text{OPT}$. Hence, after rescaling, we can assume that the interpoint distance is at least 1 and we consider distances up to $O(\frac{n^3}{\epsilon}) = \text{poly}(n)$. By the property of doubling dimension (Fact 2.1), we can hence assume $|X| \leq O(\frac{n}{\epsilon})^{O(k)} \leq O(n)^{O(k)}$.

Hierarchical nets. As in [6], we consider some parameter $s = (\log n)^{\frac{c}{k}} \geq 4$, where $0 < c < 1$ is a universal constant that is sufficiently small (as required in Lemma 5.31). Set $L := O(\log_s n) = O(\frac{k \log n}{\log \log n})$. A greedy algorithm can construct $N_L \subseteq N_{L-1} \subseteq \dots \subseteq N_1 \subseteq N_0 = N_{-1} = N_{-2} = N_{-3} = \dots = X$ such that for each i , N_i is an s^i -net for X , where we say *distance scale* s^i is of *height* i .

Net-respecting solution. As defined in [6], a graph F is net-respecting with respect to $\{N_i\}_{i \in [L]}$ and $\epsilon > 0$ if for every edge $\{x, y\}$ in F , both x and y belong to N_i , where $s^i \leq \epsilon \cdot d(x, y) < s^{i+1}$.

The following lemma from [6] shows that any graph may be converted to a net-respecting graph edge-by-edge, and the weight of the graph is not increased a lot. In all our arguments, we shall use Lemma 2.2 implicitly when we need to convert a

solution to be net-respecting.

LEMMA 2.2 (conversion to net-respecting solution [6]). *For any edge $e := \{u, v\}$, there is a net-respecting path P such that $w(P) \leq (1 + O(\epsilon)) \cdot w(e)$.*

Given an instance W of a problem, let $\text{OPT}(W)$ be an optimal solution; when the context is clear, we also use $\text{OPT}(W)$ to denote the cost $w(\text{OPT}(W))$ as well; similarly, $\text{OPT}^{nr}(W)$ refers to an optimal net-respecting solution. By Lemma 2.2, $w(\text{OPT}^{nr}(W)) \leq (1 + O(\epsilon)) \cdot w(\text{OPT}(W))$.

2.1. Overview. As in [6, 12], we achieve a PTAS for SFP by the framework of sparse instance decomposition.

Sparse solution and dynamic program. Given a graph F and a subset $S \subseteq X$, $F|_X$ is the subgraph induced by the vertices in $V(F) \cap X$. A graph F is called q -sparse if for all $i \in [L]$ and all $u \in N_i$, $w(F|_{B(u, 3s^i)}) \leq q \cdot s^i$.

We show that for SFP (in section 5) there is a dynamic program DP that runs in polynomial time such that if an instance W has an optimal net-respecting solution that is q -sparse for some small enough q , $\text{DP}(W)$ returns a $(1 + \epsilon)$ -approximation with high probability (at least $1 - \frac{1}{\text{poly}(n)}$).

Sparsity heuristic. Since one does not know the optimal solution in advance, we estimate the local sparsity with a heuristic. For $i \in [L]$ and $u \in N_i$, given an instance W , the heuristic $H_u^{(i)}(W)$ is supposed to estimate the sparsity of an optimal net-respecting solution in the ball $B' := B(u, O(s^i))$.² We shall see in section 3 that the heuristic actually gives a constant approximation to some appropriately defined subinstance W' in the ball B' .

Generic algorithm. We describe a generic framework that applies to SFP. A similar framework is also used in [12, 6] to obtain PTASs for TSP-related problems. Given an instance W , we describe the recursive algorithm $\text{ALG}(W)$ as follows:

1. **Base case.** If $|W| = n$ is smaller than some constant threshold, solve the problem by brute force, recalling that $|X| \leq O(\frac{n}{\epsilon})^{O(k)}$.
2. **Sparse instance.** If for all $i \in [L]$, for all $u \in N_i$, $H_u^{(i)}(W)$ is at most $q_0 \cdot s^i$, for some appropriate threshold q_0 , call the subroutine $\text{DP}(W)$ to return a solution, and terminate.
3. **Identify critical instance.** Otherwise, let i be the smallest height such that there exists $u \in N_i$ with *critical* $H_u^{(i)}(W) > q_0 \cdot s^i$; in this case, choose $u \in N_i$ such that $H_u^{(i)}(W)$ is maximized.
4. **Decomposition into sparse instances.** Decompose the instance W into appropriate subinstances W_1 and W_2 (possibly using randomness). Loosely speaking, W_1 is a sparse enough subinstance induced in the region around u at distance scale s^i , and W_2 captures the rest. We note that $H_u^{(i)}(W_2) \leq q_0 \cdot s^i$ such that the recursion will terminate. We also require the the subinstances contain some common terminal pairs, so that the union of the solutions to the subinstances will be a solution to W . Moreover, the following property holds:

$$(2.1) \quad \mathbf{E}[\text{OPT}(W_1)] \leq \frac{1}{1 - \epsilon} \cdot (\text{OPT}^{nr}(W) - \mathbf{E}[\text{OPT}^{nr}(W_2)]),$$

where the expectation is over the randomness of the decomposition.

²Depending on the context, we may use $H_u^{(i)}(W)$ to denote the value or the solution.

5. **Recursion.** Call the subroutine $F_1 := \text{DP}(W_1)$, and solve $F_2 := \text{ALG}(W_2)$ recursively; return the union $F_1 \cup F_2$.

Analysis of approximation ratio. We follow the inductive proof as in [6] to show that with constant probability (where the randomness comes from DP), $\text{ALG}(W)$ returns a tour with expected length at most $\frac{1+\epsilon}{1-\epsilon} \cdot \text{OPT}^{nr}(W)$, where expectation is over the randomness of decomposition into sparse instances in step 4.

As we shall see, in $\text{ALG}(W)$, the subroutine DP is called at most $\text{poly}(n)$ times (either explicitly in the recursion or the heuristic $H^{(i)}$). Hence, with constant probability, all solutions returned by all instances of DP have appropriate approximation guarantees.

Suppose F_1 and F_2 are solutions returned by $\text{DP}(W_1)$ and $\text{ALG}(W_2)$, respectively. Since we assume that W_1 is sparse enough and DP behaves correctly, $w(F_1) \leq (1 + \epsilon) \cdot \text{OPT}(W_1)$. The induction hypothesis states that $\mathbf{E}[w(F_2)|W_2] \leq \frac{1+\epsilon}{1-\epsilon} \cdot \text{OPT}^{nr}(W_2)$.

In step 4, (2.1) guarantees that $\mathbf{E}[\text{OPT}(W_1)] \leq \frac{1}{1-\epsilon} \cdot (\text{OPT}^{nr}(W) - \mathbf{E}[\text{OPT}^{nr}(W_2)])$. Hence, it follows that $\mathbf{E}[w(F_1) + w(F_2)] \leq \frac{1+\epsilon}{1-\epsilon} \cdot \text{OPT}^{nr}(W) = (1 + O(\epsilon)) \cdot \text{OPT}(W)$, achieving the desired ratio.

Analysis of running time. As mentioned above, if $H_u^{(i)}(W)$ is found to be critical, then the instance is decomposed subinstances W_1 and W_2 , and $H_u^{(i)}(W_2)$ should be small. This implies $H_u^{(i)}(W_2)$ should not be critical again, and hence it follows that there will be at most $|X| \cdot L = \text{poly}(n)$ recursive calls to ALG . Therefore, as far as obtaining polynomial running times, it suffices to analyze the running time of the dynamic program DP. The details are in section 5.4.

2.2. Paper organization. In order to apply the above framework to obtain a PTAS for SFP, we shall describe in detail the following components:

1. (section 3) Design a heuristic H such that for each $i \in [L]$ and $u \in N_i$, the heuristic $H_u^{(i)}(W)$ gives an upper bound for $\text{OPT}^{nr}(W)|_{B(u, 3s^i)}$.
2. (section 4) When a critical $H_u^{(i)}(W)$ is found, decompose W into instances W_1 and W_2 such that (2.1) holds.
3. (section 5) Design a dynamic program DP that gives $(1 + \epsilon)$ -approximation to sparse instances in polynomial time.

3. Sparsity heuristic for SFP. Suppose a collection W of terminal pairs is an instance of SFP. For $i \in [L]$ and $u \in N_i$, recall that we wish to estimate $\text{OPT}^{nr}(W)|_{B(u, 3s^i)}$ with some heuristic $H_u^{(i)}(W)$. We consider a more general heuristic $T_u^{(i,t)}$ associated with the ball $B(u, ts^i)$ for $t \geq 1$. The following auxiliary subinstance deals with terminal pairs that are separated by the ball.

Auxiliary subinstance. Fix $\delta := \Theta(\frac{\epsilon}{k})$, where the constant depends on the proof of Lemma 4.3. For $i \in [L]$, $u \in N_i$, and $t \geq 1$, the subinstance $W_u^{(i,t)}$ is induced by each pair $\{a, b\} \in W$ as follows:

- (a) If both $a, b \in B(u, ts^i)$, or if exactly one of them is in $B(u, ts^i)$ and the other is in $B(u, (t + \delta)s^i)$, then $\{a, b\}$ is also included in $W_u^{(i,t)}$.
- (b) Suppose j is the index such that $s^j < \delta s^i \leq s^{j+1}$. If $a \in B(u, ts^i)$ and $b \notin B(u, (t + \delta)s^i)$, then $\{a, a'\}$ is included in $W_u^{(i,t)}$, where a' is the nearest point to a in N_j .
- (c) If both a and b are not in $B(u, ts^i)$, then the pair is excluded.

Defining heuristic. We define $H_u^{(i)}(W) := T_u^{(i,4)}(W)$ in terms of a more general heuristic, where $T_u^{(i,t)}(W)$ is the cost of a constant approximate net-respecting solu-

tion of SFP on the instance $W_u^{(i,t)}$. For example, we can first apply the primal-dual algorithm in [19] that gives a 2-approximation of SFP and then make it net-respecting, and we have $T_u^{(i,t)}(W) \leq 2(1 + \Theta(\epsilon)) \cdot \text{OPT}(W_u^{(i,t)})$.

One potential issue is that $\text{OPT}^{nr}(W)$ might use Steiner points in $B(u, ts^i)$, even if $W_u^{(i,t)}$ is empty. We shall prove a structural property of the optimal Steiner tree³ in Lemma 3.1, and Lemma 3.1 implies Lemma 3.2, which helps us to resolve this issue.

LEMMA 3.1 (distribution of Steiner points in the optimal Steiner tree). *Suppose S is a terminal set with $\text{Diam}(S) \leq D$, and suppose F is an optimal Steiner tree with terminal set S . If the longest edge in F has weight at most γD ($0 < \gamma \leq 1$), then for any Steiner point r in F , $d(r, S) \leq 4k\gamma \log_2 \frac{4}{\gamma} \cdot D$.*

Proof. Since F is an optimal solution, all Steiner points in F have degree at least 3. Fix any Steiner point r in F .

Denote $K := \lceil \log_2(\gamma D) \rceil$. Suppose we consider r as the root of the tree F . We shall show that there is a path of small weight from r to some terminal. Without loss of generality, we can assume that all terminals are leaves, because once we reach a terminal, there is no need to visit its descendants. For simplicity, we can assume that each internal node (Steiner point) has exactly two children, because we can ignore extra branches if an internal has more than two children.

For $i \leq K$, let E_i be the set of edges in F that have weights in the range $(2^{i-1}, 2^i]$, and we say that such an edge is of *type i* . For each node u in F , denote F_u as the subtree rooted at u . Suppose we consider F_u and remove all edges in $\cup_{j \geq i} E_j$ from F_u ; in the resulting forest, let $M_u^{(i)}$ be the number of connected components that contain at least one terminal. We shall prove the following statement by structural induction on the tree F .

For each node $u \in F$, there exists a leaf $x \in F_u$ such that

$$d(x, u) \leq \sum_{i \leq K} 2^i \log_2 M_u^{(i)}.$$

Base case. If u is a leaf, then the statement is true.

Inductive step. Suppose u has children u_1 and u_2 such that $\{u, u_1\} \in E_i$ and $\{u, u_2\} \in E_{i'}$, where $i \geq i'$. Suppose x_1 and x_2 are the leaves in F_{u_1} and F_{u_2} , respectively, from the induction hypothesis. Observe that $M_u^{(i)} = M_{u_1}^{(i)} + M_{u_2}^{(i)}$. We consider two cases:

(1) Suppose $M_{u_1}^{(i)} \leq M_{u_2}^{(i)}$. Then, we can pick x_1 to be the desired leaf, because the extra distance $d(u_1, u) \leq 2^i$ can be accounted for, as $2M_{u_1}^{(i)} \leq M_u^{(i)}$, and $M_{u_1}^{(j)} \leq M_u^{(j)}$ for $j \neq i$. More precisely, $d(x_1, u) \leq d(x_1, u_1) + d(u_1, u) \leq 2^i \cdot (1 + \log_2 M_{u_1}^{(i)}) + \sum_{j \leq K, j \neq i} 2^j \log_2 M_{u_1}^{(j)} \leq \sum_{j \leq K} 2^j \log_2 M_u^{(j)}$, where the second inequality follows from the induction hypothesis for u_1 .

(2) Suppose $M_{u_2}^{(i)} < M_{u_1}^{(i)}$. Then, similarly we pick x_2 to be the desired leaf, because the extra distance is $d(u_2, u) \leq 2^{i'} \leq 2^i$. This completes the inductive step.

Next, it suffices to give an upper bound for each $M^{(i)} := M_r^{(i)}$ for root r . Suppose after removing all tree edges in $\cup_{j \geq i} E_j$, P and Q are two connected components each containing at least one terminal. Then, observe that the path in F connecting P and Q must contain an edge e with weight at least 2^{i-1} . It follows that $d(P, Q) \geq 2^{i-1}$;

³Recall that the Steiner tree problem is a special case of SFP where the goal is to return a minimum cost tree that connects all terminals.

otherwise, we can replace e in F with another edge of length less than 2^{i-1} to obtain a Steiner tree with strictly less weight. It follows that each cluster has a terminal representative that form a 2^{i-1} -packing. Hence, we have $M^{(i)} \leq (\frac{4D}{2^i})^k$ by the packing property of doubling metrics (Fact 2.1).

Therefore, every Steiner point r in F has a terminal within distance $\sum_{i \leq K} k \cdot 2^i \log_2 \frac{4D}{2^i} \leq 4k\gamma D \log_2 \frac{4}{\gamma}$. \square

Given a graph F , a *chain* in F is specified by a sequence of points (p_1, p_2, \dots, p_l) such that there is an edge $\{p_i, p_{i+1}\}$ in F between adjacent points, and the degree of an internal point p_i (where $2 \leq i \leq l-1$) in F is exactly 2.

The following lemma is crucially used in Lemmas 3.3 and 3.5. Since both lemmas are dealing with optimal *net-respecting* solutions, we need to consider net-respecting solutions in Lemma 3.2 as well. We note that in an optimal net-respecting solution, a Steiner point can have a Steiner point of degree 2, and that is the reason we consider long *chains* instead of long edges.

LEMMA 3.2 (Steiner tree of well-separated terminals contains a long chain). *Suppose S and T are terminal sets in a metric space with doubling dimension at most k such that $\text{Diam}(S \cup T) \leq D$, and $d(S, T) \geq \tau D$, where $0 < \tau < 1$. Suppose F is an optimal net-respecting Steiner tree connecting the points in $S \cup T$. Then, there is a chain in F with weight at least $\frac{\tau^2}{4096k^2} \cdot D$ such that any internal point in the chain is a Steiner point.*

Proof. Denote $\gamma := \frac{\tau^2}{4096k^2}$. Suppose for contradiction's sake that all chains in F have weight less than γD . We consider a minor \widehat{F} that is obtained from F by merging Steiner points of degree 2 with adjacent points. Hence, the vertex set of \widehat{F} is comprised of the terminals together with Steiner points in F with degree at least 3. Moreover, an edge in \widehat{F} corresponds to a chain in F , and its weight is defined to be the weight of the corresponding chain.

Then, by using the argument in Lemma 3.1, we can prove that every point u in \widehat{F} is within distance at most $4k\gamma \log_2 \frac{8}{\gamma} \cdot D$ to a terminal. Precisely, we shall replace the F in the argument of Lemma 3.1 with \widehat{F} . We observe that the only difference caused by this replacement is when we use the optimality of the solutions. Specifically, in Lemma 3.1, we use the fact that when an edge e connects point sets P and Q , both contain at least one terminal (i.e., removing e results in the disconnectivity of P and Q); it has to be $d(P, Q) \geq w(e)$, while the corresponding fact for \widehat{F} is $d(P, Q) \geq \frac{w(e)}{1+\Theta(\epsilon)} \geq \frac{w(e)}{2}$, where the first inequality follows from the optimality of F and the fact that there exists a path connecting P and Q with weight at most $(1 + O(\epsilon)) \cdot d(P, Q)$ by Lemma 2.2.

Obtaining contradiction. Recall that the terminal sets S and T are well separated: $d(S, T) \geq \tau D$. Since all Steiner points in \widehat{F} are at distance at most $4k\gamma \log_2 \frac{8}{\gamma} \cdot D$ from the terminals, it follows that there must be an edge in \widehat{F} with length at least $\tau D - 8k\gamma \log_2 \frac{8}{\gamma} \cdot D > \tau D - 32k\sqrt{\gamma}D > \gamma D$. \square

LEMMA 3.3. *Suppose F is an optimal net-respecting solution for an SFP instance W . Then, for any i and $u \in N_i$ and $t \geq 1$, $w(F|_{B(u, ts^i)}) \leq \mathbb{T}_u^{(i, t+1)}(W) + O(\frac{skt}{\epsilon})^{O(k)} s^i$.*

Proof. Given an optimal net-respecting solution F , we shall construct another net-respecting solution in the following steps:

1. Remove edges in $F|_{B(u, ts^i)}$.
2. Add edges corresponding to the heuristic $\mathbb{T}_u^{(i, t+1)}(W)$.

3. Add edges in an MST H of $N_j \cap B(u, (t + 2)s^i)$, where $s^j \leq \Theta(\frac{\epsilon}{(t+1)k^2}) \cdot s^i < s^{j+1}$, where the constant in Θ depends on Lemma 3.2; convert each added edge into a net-respecting path if necessary.

By the choice of j and the packing property, $|H| \leq O(\frac{stk}{\epsilon})^{O(k)}$, and hence the weight of edges added in this step is $O(\frac{stk}{\epsilon})^{O(k)} \cdot s^i$.

4. To ensure feasibility, replace some edges without increasing the weight.

If we can show that the resulting solution is feasible for W , then the optimality of F implies the result. We denote $B := B(u, ts^i)$ and $\widehat{B} := B(u, (t + 1)s^i)$.

Feasibility. Define

$$\widehat{V}_1 := \{x : x \in B \mid \exists\{x, y\} \in F$$

$$\text{s.t. } y \notin B \text{ and } y \text{ is connected in } F|_{X \setminus B} \text{ to some point outside } \widehat{B}\}$$

and

$$\widehat{V}_2 := \{x : x \in \widehat{B} \setminus B \mid x \text{ is connected in } F|_{\widehat{B}} \text{ to some point in } \widehat{V}_1,$$

$$\exists\{x, y\} \in F \text{ s.t. } y \notin \widehat{B}\}.$$

In step 4, we will ensure that all points in $\widehat{V}_1 \cup \widehat{V}_2$ are connected to the MST H .

If a pair $\{a, b\} \in W$ has both terminals in \widehat{B} , then they will be connected by the edges corresponding to $\mathsf{T}_u^{(i,t+1)}(W)$. If $a \in \widehat{B}$ and $b \notin \widehat{B}$, then edges for the heuristic $\mathsf{T}_u^{(i,t+1)}(W)$ ensure that a is connected to H ; moreover, in the original tree F , if the path from a to b does not meet any node in \widehat{V}_2 , then this path is preserved; otherwise, there is a portion of the path from a point in \widehat{V}_2 to b that is still preserved. If both a and b are outside \widehat{B} , then they might be connected in F via points in \widehat{V}_2 ; however, since all points in \widehat{V}_2 are connected to H , feasibility is ensured.

We next elaborate on how step 4 is performed. Consider a connected component U in $F|_{\widehat{V}_1 \cup (\widehat{B} \setminus B)}$ that contains a point in \widehat{V}_1 . Let $S_1 := U \cap \widehat{V}_1$ and $S_2 := U \cap \widehat{V}_2$. If $S_2 = \emptyset$, then there is an edge connecting S_1 directly to a point outside \widehat{B} . This means that both its end-points are in N_j by the net-respecting property, and hence S_1 is already connected to H .

Next, if there is a point $z \notin \widehat{B}$ connected directly to some point $y \in S_2$ such that $d(y, z) \geq \frac{s^i}{2}$, then by the net-respecting property, $y \in N_j$, and so again U is connected to H . Otherwise, we have $d(S_1, S_2) \geq \frac{s^i}{2}$. We next replace U with an optimal net-respecting Steiner tree \widehat{U} connecting $S_1 \cup S_2$. Since U itself is net-respecting, this does not increase the cost.

Observing that $\text{Diam}(S_1 \cup S_2) \leq 2(t + 1)s^i$, we can use Lemma 3.2 to conclude that there exists a chain in \widehat{U} from some point u to v such that its length is at least $\Theta(\frac{1}{k^2(t+1)}) \cdot s^i$. Hence, we can remove this chain and use its weight to add a net-respecting path from each of u and v to its nearest point in N_j . This does not increase the cost and ensures that both S_1 and S_2 are connected to H .

Therefore, we have shown that step 4 ensures that all points in \widehat{V}_1 and \widehat{V}_2 are connected to H . \square

It is because of Lemma 3.3 that we choose $\mathsf{H}_u^{(i)}(W) := \mathsf{T}_u^{(i,4)}(W)$ to be the heuristic.

COROLLARY 3.4 (threshold for critical instance). *Suppose F is an optimal net-respecting solution for an SFP instance W , and $q \geq \Theta(\frac{sk}{\epsilon})^{\Theta(k)}$. If for all $i \in [L]$ and $u \in N_i$, $\mathsf{H}_u^{(i)}(W) \leq qs^i$, then F is $2q$ -sparse.*

LEMMA 3.5. *Suppose W is an SFP instance. Consider $i \in [L]$, $u \in N_i$, and $t \geq t' \geq 1$. Suppose F is a net-respecting solution for $W_u^{(i,t)}$. Then, $\mathsf{T}_u^{(i,t')}(W) \leq 4(1 + \epsilon) \cdot w(F) + O(\frac{skt'}{\epsilon})^{O(k)} s^i$.*

Proof. We first show that there is a feasible solution for $W_u^{(i,t')}$ with weight at most $2 \cdot w(F) + O(\frac{skt'}{\epsilon})^{O(k)} s^i$. Then, the heuristic $\mathsf{T}_u^{(i,t')}(W)$ gives the weight of a net-respecting solution with cost at most $4(1 + \epsilon) \cdot w(F) + O(\frac{skt'}{\epsilon})^{O(k)} s^i$.

We first include F in the solution. It suffices to handle the terminal pairs in $W_u^{(i,t')} \setminus W_u^{(i,t)}$. Such a pair $\{a, a'\}$ must be induced from $\{a, b\} \in W_u^{(i,t)}$ such that $a \in B(u, t' s^i)$ and $b \in B(u, (t + \delta)s^i) \setminus B(u, (t' + \delta)s^i)$. We next add more edges such that a is connected to a' , which lies in N_j , where $s^j \leq \Theta(\frac{\delta^2}{t'k^2}) \cdot s^i < s^{j+1}$.

We add an MST H on the points in $N_j \cap B(u, (t' + \delta)s^i)$. This has cost at most $O(\frac{skt'}{\epsilon})^{O(k)} \cdot s^i$.

Consider a connected component U of F . Consider the terminal pairs $\{a, b\} \in W_u^{(i,t)}$ connected by U such that $a \in B(u, t' s^i)$ and $b \in B(u, (t + \delta)s^i)$; let S_1 be those terminals a , and let S_2 be those terminals b . Suppose \widehat{U} is an optimal net-respecting Steiner tree connecting $S_1 \cup S_2$. Since U is also net-respecting, it follows that the weight of \widehat{U} is at most that of U .

Since $d(S_1, S_2) \geq \delta s^i$ and $\text{Diam}(S_1 \cup S_2) \leq \Theta(t')s^i$, it follows from Lemma 3.2 that there exists a chain from p to q in \widehat{U} with weight at least $\Theta(\frac{\delta^2}{t'k^2}) \cdot s^i$. Hence, we can remove this chain and use this weight to connect p and q to each of their closest points in N_j . This ensures that each point $a \in S_1$ is connected to its closest point in N_j via the MST H .

If we perform this operation on each connected component U of F , the weight of edges added is at most $w(F)$. Hence, we have shown that there is a feasible solution to $W_u^{(i,t')}$ with cost at most $2 \cdot w(F) + O(\frac{skt'}{\epsilon})^{O(k)} s^i$, as required. \square

4. Decomposition into sparse instances. In section 3, we defined a heuristic $\mathsf{H}_u^{(i)}(W)$ to detect a critical instance around some point $u \in N_i$ at distance scale s^i . We next describe how the instance W can be decomposed into W_1 and W_2 such that (2.1) in section 2.1 is satisfied.

Since the ball centered at u with radius around s^i could potentially separate terminal pairs in W , we use the idea in section 3 for defining the heuristic to decompose the instance.

Decomposing a critical instance. We define a threshold $q_0 := \Theta(\frac{sk}{\epsilon})^{\Theta(k)}$ according to Corollary 3.4. As stated in section 2.1, a critical instance is detected by the heuristic when a smallest $i \in [L]$ is found for which there exists some $u \in N_i$ such that $\mathsf{H}_u^{(i)}(W) = \mathsf{T}_u^{(i,4)}(W) > q_0 s^i$. Moreover, in this case, $u \in N_i$ is chosen to maximize $\mathsf{H}_u^{(i)}(W)$. To achieve a running time with an $\exp(O(1)^k \log(k))$ dependence on the doubling dimension k , we also apply the technique in [12] to choose the cutting radius carefully.

CLAIM 4.1 (choosing radius of cutting ball). *Denote $\mathsf{T}(\lambda) := \mathsf{T}_u^{(i,4+2\lambda)}(W)$. Then, there exists $0 \leq \lambda < k$ such that $\mathsf{T}(\lambda + 1) \leq 30k \cdot \mathsf{T}(\lambda)$.*

Proof. Suppose the contrary is true. Then, it follows that $\mathsf{T}(k) > (30k)^k \cdot \mathsf{T}(0)$. We shall obtain a contradiction by showing that there is a solution for the instance $W_u^{(i,4+2k)}$ corresponding to $\mathsf{T}(k) = \mathsf{T}_u^{(i,4+2k)}(W)$ with small weight.

Define N'_i to be the set of points in N_i that cover $B(u, (2k + 5)s^i)$, and similarly

define N'_j , where $s^j \leq \delta \cdot s^i \leq s^{j+1}$.

Define edge set F to be the union of an MST on N'_j together with the union of the edge sets $H_v^{(i)}$ over $v \in N'_j$. It follows that F is a feasible solution for the instance $W_u^{(i, 4+2k)}$. By the choice of u and q_0 , we have $w(F) \leq |N'_j| \cdot 2(2k+5) \cdot s^i + |N'_j| \cdot T(0) \leq q_0 s^i + (4k+10)^k \cdot T(0) \leq (15k)^k \cdot T(0)$.

Hence, we have an upper bound for the heuristic $T(k) \leq 2(1 + \Theta(\epsilon)) \cdot w(F) \leq (30k)^k \cdot T(0)$, which gives us the desired contradiction. \square

Cutting ball and subinstances. Suppose $\lambda \geq 0$ is picked as in Claim 4.1, and sample $h \in [0, \frac{1}{2}]$ uniformly at random. Recall that $\delta := \Theta(\frac{\epsilon}{k})$. Define $B := B(u, (4 + 2\lambda + h)s^i)$ and $\widehat{B} := B(u, (4 + 2\lambda + h + \delta)s^i)$. The instances W_1 and W_2 are induced by each pair $\{a, b\} \in W$ as follows:

- (a) If $a \in B$ and $b \in \widehat{B}$, then include $\{a, b\}$ in W_1 .
- (b) If $a \in B$ and $b \notin \widehat{B}$, then include $\{a, a'\}$ in W_1 and $\{a', b\}$ in W_2 , where a' is the closest point in N_j to a and $s^j \leq \delta \cdot s^i < s^{j+1}$.
- (c) If both a and b are not in B , then include $\{a, b\}$ in W_2 .

LEMMA 4.2 (subinstances are sparse). *The subinstances W_1 and W_2 satisfy the following:*

- (i) *If F_1 is feasible for W_1 and F_2 is feasible for W_2 , then the union $F_1 \cup F_2$ is feasible for W .*
- (ii) *The subinstance W_2 does not have a critical instance with height less than i , and $H_u^{(i)}(W_2) = 0$.*
- (iii) $H_u^{(i)}(W_1) \leq O(s)^{O(k)} \cdot q_0 \cdot s^i$.

Proof. The first two statements follow immediately from the construction (to see (ii), recall the definition of auxiliary subinstance defined in section 3). For the third statement, we use the fact that there is no critical instance at height $i - 1$ to show that there is a solution to W_1 with small cost.

Specifically, we consider an MST H on $N_j \cap B(u, 5s^i)$, where $s^j \leq \delta \cdot s^{i-1} < s^{j+1}$. Then, we have $w(H) \leq q_0 \cdot s^i$.

Moreover, we consider the union of solutions corresponding to $H_v^{(i-1)}(W)$ over $v \in N_{i-1} \cap B(u, 5s^i)$. The cost is $O(s)^{O(k)} \cdot q_0 \cdot s^i$.

Hence, the union of H together with the edges for the $H_v^{(i-1)}(W)$'s is feasible for W_1 , and this implies that $H_u^{(i)}(W_1) \leq O(s)^{O(k)} \cdot q_0 \cdot s^i$. \square

LEMMA 4.3 (combining costs of subinstances). *Suppose F is an optimal net-respecting solution for W . Then, for any realization of the decomposed subinstances W_1 and W_2 as described above, there exist net-respecting solutions F_1 and F_2 for W_1 and W_2 , respectively, such that $(1 - \epsilon) \cdot \mathbf{E}[w(F_1)] + \mathbf{E}[w(F_2)] \leq w(F)$, where the expectation is over the randomness to generate W_1 and W_2 .*

Proof. Let B and \widehat{B} be defined as above, and denote $\overline{B} := B(u, (4 + 2\lambda + 1) \cdot s^i)$. Since δ is sufficiently small, we have $B \subset \widehat{B} \subset \overline{B}$.

We start by including $F|_B$ in T_1 and including the remaining edges in F in F_2 . We will then show how to add extra edges with expected weight at most $\epsilon \cdot \mathbf{E}[w(F_1)]$ to make F_1 and F_2 feasible. This will imply the lemma.

Define N to be the subset of N_j that covers the points in \overline{B} , where $s^j < \delta s^i \leq s^{j+1}$. We include a copy of an MST H of N in each of F_1 and F_2 and make it net-respecting. This costs at most $|N| \cdot O(k) \cdot s^i \leq O(\frac{ks}{\epsilon})^{O(k)} \cdot s^i$.

We next include the edges of F in the annulus $\widehat{B} \setminus B$ (of width δ) into F_1 . This

has expected cost at most $\delta \cdot w(F|_{\overline{B}})$.

Connecting crossing points. To ensure the feasibility of F_1 , we connect the following sets of points to N . We denote $V_1 := \{x \in B \mid \exists y \in \widehat{B} \setminus B, \{x, y\} \in F\}$, $V_2 := \{y \in \widehat{B} \setminus B \mid \exists x \in B, \{x, y\} \in F\}$, and $V_3 := \{x \in \widehat{B} \mid \exists y \notin \widehat{B}, \{x, y\} \in F\}$.

We shall connect each point in $V_1 \cup V_2 \cup V_3$ to its closest point in N . Note that if such a point x is incident to some edge in F with weight at least $\frac{s^i}{4}$, then the net-respecting property of F implies that x is already in N . Otherwise, this is because some edge $\{x, y\}$ in F is cut by either B or \widehat{B} , which happens with probability at most $O(\frac{d(x,y)}{s^i})$. Hence, each edge $\{x, y\} \in F|_{\overline{B}}$ has an expected contribution of $\delta s^i \cdot O(\frac{d(x,y)}{s^i}) = O(\delta) \cdot d(x, y)$.

Similarly, to ensure the feasibility of F_2 , we ensure each point in the following set is connected to N . Denote $\widehat{V}_1 := \{x \in B \mid \exists y \notin B, \{x, y\} \in F\}$. By the same argument, the expected cost to connect each point to N is also at most $O(\delta) \cdot w(F|_{\overline{B}})$.

Charging the extra costs to F_1 . Apart from using edges in F , the extra edges come from two copies of the MST H and other edges with cost $O(\delta) \cdot w(F|_{\overline{B}})$. We charge these extra costs to F_1 .

Since $T_u^{(i,4)}(W) > q_0 \cdot s^i$ and F_1 is a net-respecting solution for $W_u^{(i,4+2\lambda+h)}$, by Lemma 3.5, $w(F_1) \geq \frac{1}{4(1+\epsilon)}(T^{(i)}(u, 4) - O(\frac{sk}{\epsilon})^{O(k)} \cdot s^i) > \frac{q_0}{8} \cdot s^i$ by choosing large enough q_0 .

Therefore, the cost for the two copies of the MST H is at most $O(\frac{ks}{\epsilon})^{O(k)} \cdot s^i \leq \frac{\epsilon}{2} \cdot w(F_1)$.

We next give an upper bound on $w(F|_{\overline{B}})$, which is at most $T_u^{(i,4+2(\lambda+1))}(W) + O(\frac{sk}{\epsilon})^{O(k)} \cdot s^i$, by Lemma 3.3. By the choice of λ , we have $T_u^{(i,4+2(\lambda+1))}(W) \leq 30k \cdot T_u^{(i,4+2\lambda)}(W)$. Moreover, by Lemma 3.5, $T_u^{(i,4+2\lambda)}(W) \leq 4(1+\epsilon) \cdot w(F_1) + O(\frac{sk}{\epsilon})^{O(k)} \cdot s^i$. Hence, we can conclude that $w(F|_{\overline{B}}) \leq O(k) \cdot w(F_1)$.

Hence, by choosing small enough $\delta = \Theta(\frac{\epsilon}{k})$, we can conclude that the extra costs $O(\delta) \cdot w(F|_{\overline{B}}) \leq \frac{\epsilon}{2} \cdot w(F_1)$.

Therefore, we have shown that $\mathbf{E}[w(F_1)] + \mathbf{E}[w(F_2)] \leq w(F) + \epsilon \cdot w(F_1)$, where the right-hand side is a random variable. Taking expectation on both sides and rearranging gives the required result. \square

5. A PTAS for sparse SFP instances. Our dynamic program follows the divide and conquer strategy as in previous works on TSP [3, 25, 6] that are based on hierarchical decomposition. However, to apply the framework to SFP, we need a version of the *cell property* that is more sophisticated than previous works [10, 8].

We shall first give a review of the hierarchical decomposition techniques in section 5.1. In section 5.2, we shall give a high level overview of our new ideas. In section 5.3, we shall elaborate on the precise implementations of the ideas. Finally, we shall show the details of DP in section 5.4 and conclude a PTAS for sparse SFP instances (in Corollary 5.35).

5.1. Review on hierarchical decomposition.

DEFINITION 5.1 (single-scale decomposition [1]). *At height i , an arbitrary ordering π_i is imposed on the net N_i . Each net-point $u \in N_i$ corresponds to a cluster center and samples random h_u from a truncated exponential distribution Exp_i having density function $t \mapsto \frac{\chi}{\chi-1} \cdot \frac{\ln \chi}{s^i} \cdot e^{-\frac{t \ln \chi}{s^i}}$ for $t \in [0, s^i]$, where $\chi = O(1)^k$. Then, the cluster at u has random radius $r_u := s^i + h_u$.*

The clusters induced by N_i and the random radii form a decomposition Π_i , where a point $p \in X$ belongs to the cluster with center $u \in N_i$ such that u is the first point

in π_i to satisfy $p \in B(u, r_u)$. We say that the partition Π_i cuts a set P if P is not totally contained within a single cluster.

The results in [1] imply that the probability that a set P is cut by Π_i is at most $\frac{\beta \cdot \text{Diam}(P)}{s^i}$, where $\beta = O(k)$.

DEFINITION 5.2 (hierarchical decomposition). *Given a configuration of random radii for $\{N_i\}_{i \in [L]}$, decompositions $\{\Pi_i\}_{i \in [L]}$ are induced as in Definition 5.1. At the top height $L-1$, the whole space is partitioned by Π_{L-1} to form height- $(L-1)$ clusters. Inductively, each cluster at height $i+1$ is partitioned by Π_i to form height- i clusters, until height 0 is reached. Observe that a cluster has $K := O(s)^k$ child clusters. Hence, a set P is cut at height i if and only if the set P is cut by some partition Π_j such that $j \geq i$; this happens with probability at most $\sum_{j \geq i} \frac{\beta \cdot \text{Diam}(P)}{s^j} = \frac{O(k) \cdot \text{Diam}(P)}{s^i}$.*

Portals. As in [4, 25, 6], each height- i cluster U is equipped with portals such that a solution F is *portal-respecting* if for every edge $\{x, y\}$ in F between a point x in U and some point y outside U , at least one of x and y must be a portal of cluster U . As mentioned in [6], the portals of a cluster need not be points of the cluster itself but are just used as connection points. For a height- i cluster C , its portals are comprised of the subset of net-points in $N_{i'}$ that cover C , where i' is the maximum index such that $s^{i'} \leq \max\{1, \frac{\epsilon}{4\beta L} \cdot s^i\}$. As noted in [25, 6, 12], any solution can be made to be portal-respecting with a multiplicative factor of $1 + O(\epsilon)$ in cost.


Since a height- i cluster has diameter $O(s^i)$, by Fact 2.1, the cluster has at most $m := O(\frac{\beta L s}{\epsilon})^k$ portals.

(m, r) -light solution. A solution F is called (m, r) -light if it is portal-respecting for a hierarchical decomposition in which each cluster has at most m portals, and for each cluster, at most r of its portals are used in F to connect points in the cluster to the points outside.

5.2. Proof overview. As discussed in section 1, the DP entry has to keep track of the portals to which each terminal is connected, and a naive implementation can lead to an exponential run time. In the Euclidean plane, the notion of cells was introduced in [10] to resolve the issue, and the idea is to group several terminals in a cell (which is a subset of points) such that all terminals in a cell have similar connectivity and to argue that enforcing the connectivity is not very costly. Moreover, the number of cells is small for each cluster, which is a constant for the Euclidean plane, and it is feasible to keep track of the portals to which each cell is connected.

Indeed, it is natural to define the cells in a uniform way such that all child clusters of diameter γD are cells of a cluster of diameter D , where $0 < \gamma < 1$. However, in doubling metrics, it turns out γ has to be set to as small as $\frac{1}{L}$, which leads to poly $\log n$ number of cells for each cluster. This only gives a QPTAS.

Our idea is to use the sparsity to reduce the number of cells. One implication of the sparsity is that the solution may be cheaply modified such that there are only a small number of components crossing a cluster. Moreover, in the cell idea of [10], it is actually sufficient that cells only cover the crossing components. Hence, we define a cell for each crossing component, and we make the size of the cells depend on the weight of the components. That is, we use cells of a coarse gratitude to cover larger cells and use cells of a fine gratitude to cover small cells. Figure 5.1 illustrates how our cells differ from the uniform cells used in [10]. We implement this idea by defining basic cells in Definition 5.6.

FIG. 5.1. *Uniform cells versus adaptive cells.*

Observe that our cell set is not uniform, and one consequence is that the cell set for a child cluster may not be a refinement (see Definition 5.12) of that of the parent cluster. This is a severe issue for the dynamic programming algorithm, as it may not deduce the correct connectivity information of a subproblem from its child subproblems. To resolve this issue, we introduce nonbasic cells (see Definition 5.10) which together with the basic cells form the final cell set which we call the effective cells (see Definition 5.11).

We show that our cell is indeed of a small number in Lemma 5.16. Furthermore, even though the cells are defined with respect to the optimal solution which is unknown in advance, we can afford to “guess” the cell set, and the possibilities of all possible cell sets is bounded, which follows from Lemma 5.15.

Finally, we introduce a new cell property (see Definition 5.18) which is a weaker version of that considered in [10]. Our cell property requires that all components that connect a cell to outside its cluster be connected in the eventual solution, while in [10] it is required to connect all components inside the *cluster*. The new cell property enables a charging argument in Lemma 5.19 to show that the optimal solution may be modified to satisfy the cell property cheaply.

5.3. Structural property. In this section, we shall define the cell property (Definition 5.18) with respect to the effective cells (Definition 5.11), where the effective cells are carefully chosen to implement our adaptive cells idea. Specifically, the effective cells are defined by the union of the basic cells (Definition 5.6) and the nonbasic cells (Definition 5.10). Moreover, the virtual cells and the promoted cells (Definition 5.8) are introduced in order to define the nonbasic cells. Finally, we shall prove the structural property in Lemma 5.19.

Notations and parameters. Let $\text{ht}(C)$ denote the height of a cluster C , $\text{des}(C)$ denote the collection of all descendant clusters of C (including C), and $\text{par}(C)$ denote the parent cluster of C . For $x \in \mathbb{R}_+$, let $\lfloor x \rfloor_s$ denote the largest power of s that is at most x , and let $\lceil x \rceil_s$ denote the smallest power of s that is at least x . Define $\hat{\gamma}_0 := \Theta(\frac{\epsilon}{ks^2L})$, and define $\hat{\gamma}_1 := \Theta(\frac{\epsilon}{s^2})$. Define γ_0 such that $\frac{1}{\gamma_0} := \lceil \frac{1}{\hat{\gamma}_0} \rceil_s$, and define γ_1 such that $\frac{1}{\gamma_1} := \lfloor \frac{1}{\hat{\gamma}_1} \rfloor_s$. We note that $\gamma_0 < \gamma_1$.

DEFINITION 5.3 (cell). *Suppose C is a cluster of height i . A p -cell of C is a height- $\log_s p$ child cluster of C .*

DEFINITION 5.4 (crossing component). *Suppose C is some cluster and F is a solution for SFP. We say that a subset A crosses C if there exist points $x, y \in A$ such that $x \in C$ and $y \notin C$. A component A in F is called a crossing component of C if A*

crosses C .

In the following, we shall introduce the notions of basic cells, owner of basic cells, promoted cells, virtual cells, nonbasic cells, and effective cells. All of these are defined with respect to some feasible solution to SFP. We assume there is an underlying feasible solution F when talking about these definitions.

The basic cells are defined with respect to crossing components. Also, the size of the cell depends on the size of the crossing component, and in the following definition, we specify its precise size by function $h(i, l)$.

Adaptive cells. For each cluster C , we shall define its *basic cells* whose heights depend on the weights l of the crossing components of C in the solution F . We consider three cases.

Define $I_1(l) := \{i \mid \lfloor l \rfloor_s \geq s^i\}$, $I_2(l) := \{i \mid \frac{\gamma_0}{\gamma_1} s^i \leq \lfloor l \rfloor_s < s^i\}$, and $I_3(l) := \{i \mid i \leq L, \lfloor l \rfloor_s < \frac{\gamma_0}{\gamma_1} s^i\}$. Define a function $h : [L] \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that

$$h(i, l) = \begin{cases} \gamma_1 s^i & \text{for } i \in I_1(l), \\ \gamma_1 \lfloor l \rfloor_s & \text{for } i \in I_2(l), \\ \gamma_0 s^i & \text{for } i \in I_3(l). \end{cases}$$

LEMMA 5.5. $\frac{h(i+1, l)}{s} \leq h(i, l) \leq h(i+1, l)$.

Proof. If both i and $i+1$ lie in the same $I_j(l)$ ($j \in \{1, 2, 3\}$), then it holds immediately.

Otherwise, it is either $i \in I_2(l)$ but $i+1 \in I_3(l)$, or $i \in I_1(l)$ but $i+1 \in I_2(l)$, noting that it is not possible that $i \in I_1(l)$ and $i+1 \in I_3(l)$:

- If $i \in I_1(l)$ and $i+1 \in I_2(l)$. This implies $s^i = \lfloor l \rfloor_s$. Hence, $h(i, l) = \gamma_1 s^i = \gamma_1 \lfloor l \rfloor_s = h(i+1, l)$.
- If $i \in I_2(l)$ and $i+1 \in I_3(l)$. This implies $s^i = \frac{\gamma_1}{\gamma_0} \lfloor l \rfloor_s$. Hence, $s \cdot h(i, l) = s \cdot \gamma_1 \lfloor l \rfloor_s = \gamma_0 s^{i+1} = h(i+1, l)$.

This implies the inequality. □

DEFINITION 5.6 (basic cell). *Suppose C is a cluster of height i and A is a crossing component of C . Define $l := w(A)$. Define the basic cells of A in C , $\text{Bas}_A(C)$, to be the collection of the $h(i, l)$ -cells of C that intersect A . Define the basic cells of C , $\text{Bas}(C)$, to be the union of $\text{Bas}_A(C)$ for all crossing components A of C .*

DEFINITION 5.7 (owner of a basic cell). *For some cluster C , define the owner of $e \in \text{Bas}(C)$ to be the minimum weight crossing component A such that $e \in \text{Bas}_A(C)$.*

As noted in section 5.2, the nonbasic cells are used make sure the cells satisfy the refinement property. The promoted cells and virtual cells are used to define the nonbasic cells, and the detailed definition is as follows.

DEFINITION 5.8 (promoted cell and virtual cell). *Suppose C is a cluster of height i . Let S be the set of child clusters of C that is not in $\text{Bas}(C)$ but has a sibling in $\text{Bas}(C)$.*

Consider each $e \in S$:

- *If there exists a child cluster C' of C such that $e \in \text{Bas}(C')$, then define $\text{Pro}_e(C) := \text{des}(e) \cap \text{Bas}(C')$, and define $\text{Vir}_e(C) := \emptyset$, where $C' \subset C$ is any one that satisfies $e \in \text{Bas}(C')$.*
- *Otherwise, define $\text{Pro}_e(C) := \emptyset$, and define $\text{Vir}_e(C) := e$.*

Finally, $\text{Pro}(C) := \bigcup_{e \in S} \text{Pro}_e(C)$, and $\text{Vir}(C) := \bigcup_{e \in S} \text{Vir}_e(C)$, and elements in $\text{Pro}(C)$ and $\text{Vir}(C)$ are called promoted cells and virtual cells, respectively.

LEMMA 5.9. *For any cluster C , if $e \in \text{Vir}(C)$, then for any cluster $C' \subset C$ (C' may equal C), $e \setminus \{e' \in \text{Bas}(C') \mid e' \subsetneq e\}$ has no intersection with any crossing component of C' .*

Proof. Suppose not. Then, there exist a cluster $C' \subset C$ and a crossing component A of C' such that A intersects $u := e \setminus \{e' \in \text{Bas}(C') \mid e' \subsetneq e\}$. This implies that there exists $u' \in \text{Bas}_A(C')$ such that $e \subset u'$. By Lemma 5.5 and the fact that $h(\text{ht}(C'), w(A)) \geq s^{\text{ht}(e)}$ and that $h(0, w(A)) < s^{\text{ht}(e)}$, we know that there exists a cluster $C'' \subset C' \subset C$ such that $e \in \text{Bas}_A(C'')$. This contradicts the definition of virtual cells. \square

DEFINITION 5.10 (nonbasic cell). *We define the nonbasic cells $\text{NBas}(C)$ for a cluster C . If C is the root cluster, then $\text{NBas}(C) := \text{Pro}(C) \cup \text{Vir}(C) \setminus \text{Bas}(C)$. For any other cluster C , let $\text{NBas}(C) := \{e \cap C \mid e \in \text{Pro}(C) \cup \text{Vir}(C) \cup \text{NBas}(\text{par}(C)) \setminus \text{Bas}(C)\}$.*

DEFINITION 5.11 (effective cell). *For a cluster C , define the effective cells of C as $\text{Eff}(C) := \text{Bas}(C) \cup \text{NBas}(C)$.*

DEFINITION 5.12 (refinement). *Suppose S_1 and S_2 are collections of clusters. We say S_1 is a refinement of S_2 if for any $e \in S_2$, either $e \in S_1$, or all child clusters of e are in S_1 .*

LEMMA 5.13. *Suppose C is a cluster that is not a leaf. Define $\{C_i\}_i$ to be the collection of all the child clusters of C . Then, $\cup_i \text{Eff}(C_i)$ is a refinement of $\text{Eff}(C)$.*

Proof. Define $S := \cup_i \text{Eff}(C_i)$. It is sufficient to prove that for any $e \in \text{Eff}(C)$, either $e \in S$, or all child clusters of e are in S .

If $e \in \text{NBas}(C)$ and $e \neq C$, then $e \in S$ follows from Definitions 5.10 and 5.11. If $e \in \text{NBas}(C)$ but $e = C$, then also by Definitions 5.10 and 5.11, $C \cap C_i = C_i \subset \text{Eff}(C_i)$, and this implies that all child clusters of e are in S .

Otherwise, $e \in \text{Bas}(C)$; then by Lemma 5.5, we know that either $e \in S$, or there exists $e' \subset e$ such that $\text{ht}(e') = \text{ht}(e) - 1$ and $e' \in S$. Then, all siblings of e' are in S by the definition of promoted cells and virtual cells. This implies that all child clusters of e are in S . \square

The following is to show that the number of cells in each cell set is small, and the number of possible cell sets for each cluster is also bounded.

DEFINITION 5.14 (candidate center). *Suppose C is a cluster of height i . The set of candidate centers of C , denoted as $\text{Can}(C)$, is the subset of $\cup_{j=\log_s \gamma_0^2 s^i}^i N_j$ that may become a center of C 's child cluster in the hierarchical decomposition.*

LEMMA 5.15. *For any cluster C , the centers of clusters in $\text{Eff}(C)$ are chosen from $\text{Can}(C)$, and $|\text{Can}(C)| \leq \kappa$, where $\kappa := O(\frac{1}{\gamma_0})^{O(k)}$.*

Proof. We first prove that centers of cluster in $\text{Eff}(C)$ are chosen from $\text{Can}(C)$:

- For $e \in \text{Bas}(C)$, by the definition of the basic cells, we have $\text{ht}(e) \geq \log_s \gamma_0 s^i$.
- For $e \in \text{Pro}(C)$, we have that e is a basic cell of some cluster C' , and hence $\text{ht}(e) \geq \log_s \gamma_0^2 s^i$.
- For $e \in \text{Vir}(C)$, since it is a sibling of a basic cell, $\text{ht}(e) \geq \log_s \gamma_0 s^i$.
- For $e \in \text{NBas}(C)$, there is a cluster C'' such that $C \subset C''$ and $e \in \text{Pro}(C'') \cup \text{Vir}(C'')$.

Hence, $\text{ht}(e) \geq \log_s \gamma_0^2 s^i$. Therefore, centers of clusters in $\text{Eff}(C)$ are in $\text{Can}(C)$.

We then bound $|\text{Can}(C)|$. Suppose $i := \text{ht}(C)$. Observe that a center of height $j \leq i$ that may become a center of a child cluster of C is contained in a ball of diameter $O(s^j)$. Moreover, N_j is an s^j packing. Hence, by the packing property,

$$|\text{Can}(C)| \leq O\left(\frac{1}{\gamma_0}\right)^{O(k)}. \quad \square$$

LEMMA 5.16. *Suppose Eff is defined in terms of a solution that is (m, r) -light. Then, for each cluster C , $|\text{Eff}(C)| \leq \rho$, where $\rho := O(\log_s \frac{1}{\gamma_0}) \cdot r^2 \cdot O(\frac{s}{\gamma_1})^{O(k)}$.*

Proof. Suppose C is of height i . We give upper bounds for $|\text{Bas}(C)|$, $|\text{Pro}(C)|$, $|\text{Vir}(C)|$, and $|\text{NBas}(C)|$, respectively.

Bounding $|\text{Bas}(C)|$. Fix a crossing component A of C , and suppose $l := w(A)$. We upper bound $|\text{Bas}_A(C)|$:

- If $i \in I_1(l)$, then $\text{Bas}_A(C)$ is a subset of $\gamma_1 s^i$ -cells of C . By the packing property, $|\text{Bas}_A(C)| \leq O\left(\frac{1}{\gamma_1}\right)^k$.
- If $i \in I_2(l)$, then $\text{Bas}_A(C)$ is a subset of $\gamma_1 \lfloor l \rfloor_s$ -cells of C . Since all the $\gamma_1 \lfloor l \rfloor_s$ -cells that intersect A are inside a ball of diameter $O(l)$, by the packing property, $|\text{Bas}_A(C)| \leq O\left(\frac{s}{\gamma_1}\right)^{O(k)}$.
- If $i \in I_3(l)$, then $\text{Bas}_A(C)$ is a subset of $\gamma_0 s^i$ -cells of C . Since all the $\gamma_0 s^i$ -cells that intersect A are inside a ball of diameter $O(\frac{2\gamma_0}{\gamma_1} s^i)$ (which follows from $i \in I_3(l)$ so that A is small), by the packing property, $|\text{Bas}_A(C)| \leq O\left(\frac{1}{\gamma_1}\right)^k$.

Since the solution is r -light, there are at most r crossing components. Therefore,

$$|\text{Bas}(C)| \leq r \cdot O\left(\frac{s}{\gamma_1}\right)^{O(k)}.$$

Bounding $|\text{Pro}(C)|$ and $|\text{Vir}(C)|$. Recall that for $e \in \text{Bas}(C)$, and for $e' \notin \text{Bas}(C)$ that is a sibling of e , we either include e to $\text{Vir}(C)$, or include $\text{des}(e) \cap \text{Bas}(C')$ to $\text{Pro}(C)$, for some child cluster C' of C . In either case, the number of added elements is at most $r \cdot O\left(\frac{s}{\gamma_1}\right)^{O(k)}$, and we charge this to e .

We observe that for each $e \in \text{Bas}(C)$, it has at most $O(s)^k$ siblings by the packing property. Therefore, each e is charged at most $O(s)^k$ times. We conclude that

$$|\text{Pro}(C) \cup \text{Vir}(C)| \leq O(s)^k \cdot r^2 \cdot O\left(\frac{s}{\gamma_1}\right)^{O(k)}.$$

Bounding $|\text{NBas}(C)|$. Suppose P is the set consisting of C and all its ancestor clusters. Recalling the definition, $\text{NBas}(C)$ is a subset of the inside C clusters of $\cup_{p \in P} (\text{Pro}(p) \cup \text{Vir}(p))$.

We shall first prove that if $\text{ht}(p) - \text{ht}(C) > 2 \log_s \frac{1}{\gamma_0}$, then there is no element in $\text{Pro}(p) \cup \text{Vir}(p)$ that can appear in $\text{NBas}(C)$ for any $p \in P$. Suppose not, and consider some p such that $\text{ht}(p) - \text{ht}(C) > 2 \log_s \frac{1}{\gamma_0}$. Let $j := \text{ht}(p)$. We observe that all elements in $\text{Pro}(p) \cup \text{Vir}(p)$ have height at least $\log_s \gamma_0^2 s^j = j - 2 \log_s \frac{1}{\gamma_0}$ by Definitions 5.8 and 5.6. However, if some element in $\text{Pro}(p) \cup \text{Vir}(p)$ appears in $\text{NBas}(C)$, then it has height less than $j - 2 \log_s \frac{1}{\gamma_0}$ by $\text{ht}(C) < \text{ht}(p) - 2 \log_s \frac{1}{\gamma_0}$. This is a contradiction. Therefore,

$$|\text{NBas}(C)| \leq O\left(\log_s \frac{1}{\gamma_0}\right) \cdot r^2 \cdot O\left(\frac{s}{\gamma_1}\right)^{O(k)}.$$

Hence, $|\text{Eff}(C)| = |\text{Bas}(C)| + |\text{NBas}(C)| \leq O(\log_s \frac{1}{\gamma_0}) \cdot r^2 \cdot O\left(\frac{s}{\gamma_1}\right)^{O(k)}. \quad \square$

DEFINITION 5.17 (disjointness). *For a collection of clusters S , define $\text{Dis}(S) := \{e \setminus \cup_{e' \in S: e' \subsetneq e} e'\}_{e \in S}$. We say e is induced by u in S if $u \in S$ and $e = u \setminus \cup_{e' \in S: e' \subsetneq u} e'$. Define the height of e as the height of u .*

DEFINITION 5.18 (cell property). *Suppose F is an SFP solution, and suppose f maps a cluster C to a collection of child clusters of C . We say that f satisfies the cell property in terms of F if for all clusters C , for all $e \in \text{Dis}(f(C))$, there is at most one crossing component of C in F that intersects e .*

LEMMA 5.19 (structural property). *Suppose an instance has a q -sparse optimal net-respecting solution F . Moreover, for each $i \in [L]$, for each $u \in N_i$, point u samples $O(k \log n)$ independent random radii as in Definition 5.1. Then, with constant probability, there exists a configuration from the sampled radii that defines a hierarchical decomposition, under which there exists an (m, r) -light solution F' that includes all the points in F , and Eff defined in terms of F' satisfies the cell property, where the following hold:*

- $\mathbf{E}[w(F')] \leq (1 + O(\epsilon)) \cdot w(F)$.
- $m := O(\frac{skL}{\epsilon})^k$ and $r := O(1)^k \cdot q \log_s \log n + O(\frac{k}{\epsilon})^k + O(\frac{s}{\epsilon})^k$.

Proof. We observe that the argument in [6, Lemma 3.1] readily gives an (m, r) -light solution \widehat{F} with the desired m and r and also satisfies $\mathbf{E}[w(\widehat{F})] \leq (1 + \epsilon) \cdot w(F)$.

We shall first show additional steps with additional cost at most $\epsilon w(F)$ in expectation, so that Bas defined in terms of the resultant solution satisfies the cell property. Then, we shall show that this implies that Eff defined in terms of the resultant solution also satisfies the cell property (hence no more additional cost caused).

Maintaining cell property: Basic cells. For $i := L, L - 1, L - 2, \dots, 0$, for each height- i cluster C , we examine $e \in \text{Dis}(\text{Bas}(C))$ in the nondecreasing order of its height. If there are at least two crossing components that intersect e , we add edges in e to connect all crossing components that intersect e . We note that each added edge connects two components in F , and edges added are of length at most $\text{Diam}(e)$. At the end of the procedure, we define the solution as F' . We observe that Bas defined in terms of F' satisfies the cell property.

Recall that each added edge connects two components. We charge the cost of the edge to one of the components to which it connects. Moreover, after a rearrangement (at the end of the procedure), we can make sure each edge is charged to one of the components to which it connects and each component is charged at most once.

Bounding the cost. We shall show that for a fixed component A , the expected cost it takes charge of is at most $\epsilon \cdot w(A)$. Define $l := w(A)$. The expected cost that A is charged for is at most the following (up to a constant):

$$\sum_{i=1}^L \Pr[A \text{ is charged for an edge in a cell of height } i] \cdot s^{i+1}.$$

Define $p_i := \Pr[A \text{ is charged for an edge in a cell of height } i]$. Then,

$$\begin{aligned} \sum_{i=0}^L p_i \cdot s^{i+1} &\leq \sum_{i:s^i \leq 2\gamma_1 l} s^{i+1} + \sum_{i:s^i > 2\gamma_1 l} p_i s^{i+1} \\ &\leq O(\gamma_1 s)l + \sum_{i:s^i > 2\gamma_1 l} p_i s^{i+1} \\ &\leq O(\epsilon)l + \sum_{i:s^i > 2\gamma_1 l} p_i s^{i+1}. \end{aligned}$$

Fix an i such that $s^i > 2\gamma_1 l$, and we shall upper bound p_i . Suppose in the event corresponding to p_i that A takes charge of an edge inside a cell e that is a basic cell

of some height- h cluster. Note that h and e are random, and recall that the edge is inside a cell of height i . We shall give a lower bound of h .

CLAIM 5.20. $s^h \geq \frac{s^i}{2\gamma_0}$.

Proof. Define the weight of the owner of e to be l' . We first show that h must be in $I_3(l')$. By the procedure of maintaining cell property, we know that $l' \leq l$.

If $h \in I_1(l')$, then $\lfloor l' \rfloor_s \geq s^h$, and $s^i \leq 2\gamma_1 s^h$ by e is of height i and the choice of radius in the single-scale decomposition. This implies that $s^i \leq 2\gamma_1 s^h \leq 2\gamma_1 l$, which cannot happen, since we assume $s^i > 2\gamma_1 l$.

If $h \in I_2(l')$, then $s^i \leq 2\gamma_1 \lfloor l' \rfloor_s$. This implies that $s^i \leq 2\gamma_1 l$, which cannot happen as well.

Therefore, $h \in I_3(l')$. This implies that $2\gamma_0 s^h \geq s^i$. □

Since the event that the edge is taken by A automatically implies that A is cut by a height- h cluster, and the probability that A is cut at a height- j cluster is at most $O(k) \cdot \frac{l}{s^j}$ for $j \in [L]$, we conclude that

$$p_i \leq \sum_{j: s^j \geq \frac{s^i}{2\gamma_0}} \Pr[A \text{ is cut at height } j] \leq O(k) \cdot \sum_{j: s^j \geq \frac{s^i}{2\gamma_0}} \frac{l}{s^j} \leq O(\gamma_0 k) \cdot \frac{l}{s^i}.$$

Hence $\sum_{i: s^i > 2\gamma_1 l} p_i s^{i+1} \leq O(\gamma_0 k s L) \cdot l \leq O(\epsilon)l$.

Maintaining cell property: Effective cells. Next we show that **Bas** defined in terms of F' satisfies that the cell property implies that **Eff** defined in terms of F' also satisfies the cell property.

Fix a cluster C , and fix $e \in \text{Dis}(\text{Eff}(C))$. We shall prove that there is at most one crossing component of C that intersects e in F' . Suppose e is induced by u in $\text{Eff}(C)$.

LEMMA 5.21. *If there is no cluster \widehat{C} such that $C \subset \widehat{C}$ and $u \in \text{Vir}(\widehat{C})$, then there exists cluster C' such that $u \in \text{Bas}(C')$, $\text{ht}(C') \leq \text{ht}(C)$, and $\text{Eff}(C)$ is a refinement of $\text{des}(u) \cap \text{Bas}(C')$.*

Proof. If $u \in \text{Bas}(C)$, then we define $C' = C$, and the lemma follows.

If $u \in \text{NBas}(C)$, then there exists C'' such that $C \subset C''$ and $u \in \text{Pro}(C'')$. This is by the definition of nonbasic cells, and by the assumption that there is not cluster \widehat{C} such that $C \subset \widehat{C}$ and $u \in \text{Vir}(\widehat{C})$. Then, by the definition of the promoted cells, there exists cluster C' such that $u \in \text{Bas}(C')$, $\text{ht}(C') < \text{ht}(C'')$, and $\text{des}(u) \cap \text{Bas}(C') \subset \text{Eff}(C'')$. Since $u \in \text{NBas}(C) \subset \text{Eff}(C)$ and by Lemma 5.13, we know that $\text{Eff}(C)$ is a refinement of $\text{des}(u) \cap \text{Bas}(C')$. Hence, it remains to show $\text{ht}(C') \leq \text{ht}(C)$.

Suppose for contradiction that $\text{ht}(C') > \text{ht}(C)$, so $\text{ht}(C) < \text{ht}(C') < \text{ht}(C'')$. By the definition of nonbasic cells, we have that $\text{NBas}(C') \cap \text{Bas}(C') = \emptyset$. Since $u \in \text{Bas}(C')$, we know that $u \notin \text{NBas}(C')$. However, this implies that $u \notin \text{NBas}(C)$, which contradicts the assumption that $u \in \text{NBas}(C)$. □

If there exists cluster \widehat{C} such that $u \in \text{Vir}(\widehat{C})$ and $C \subset \widehat{C}$, then by Lemma 5.9, there is no crossing component of C in F' that intersects e .

Otherwise, there is no cluster \widehat{C} such that $u \in \text{Vir}(\widehat{C})$ and $C \subset \widehat{C}$. By Lemma 5.21, there exists a cluster C' such that $u \in \text{Bas}(C')$, $\text{ht}(C') \leq \text{ht}(C)$, and $\text{Eff}(C)$ is a refinement of $\text{des}(u) \cap \text{Bas}(C')$. We pick any one of such C' . Define $e' \in \text{Dis}(\text{Bas}(C'))$ as the one induced by u in $\text{Bas}(C')$. Since **Bas** defined in terms of F' satisfies the cell property, there is at most one crossing component of C' that intersects e' .

LEMMA 5.22. $e \subset e'$.

Proof. Recall that $e \in \text{Dis}(\text{Eff}(C))$ is induced by u in $\text{Eff}(C)$ and $e' \in \text{Dis}(\text{Bas}(C'))$ is induced by u in $\text{Bas}(C')$. Then, we can write $e = u \setminus P$ and $e' = u \setminus P'$ such that $P \subset \text{Eff}(C)$ and $P' \subset \text{Bas}(C')$. Since $P' = \text{des}(u) \cap \text{Bas}(C')$ and $\text{Eff}(C)$ is a refinement of $\text{des}(u) \cap \text{Bas}(C')$, we know that $P' \subset P$. This implies that $e \subset e'$. \square

Since $\text{ht}(C) \geq \text{ht}(C')$, any crossing component of C is also a crossing component of C' . Moreover, Lemma 5.22 implies that $e \subset e'$. Hence, if there are two crossing components A_1, A_2 of C that intersect e , then A_1 and A_2 are also crossing components of C' and both of them intersect e' . However, this cannot happen, since Bas satisfies the cell property and there is at most one crossing component in C' that intersects e' . Therefore, there is at most one crossing component of C that intersects e . \square

5.4. Dynamic program. Recall that the input of DP is an instance that has a q -sparse optimal net-respecting solution, where $q \leq O(s)^{O(k)} \cdot q_0$, by Lemma 4.2 and Corollary 3.4. In the DP algorithm, $O(k \log n)$ random radii are independently sampled for each $u \in N_i, i \in [L]$, and then a dynamic programming based algorithm is used to find a near optimal SFP solution over all hierarchical decompositions defined by the radii. In this section, we shall describe in detail the dynamic program and an algorithm that solves the dynamic program efficiently. For completeness, we shall also analyze the correctness of the dynamic program.

We first describe the information needed to identify each cluster at each height.

Information to identify a cluster. Each cluster is identified by the following information:

1. Height i and cluster center $u \in N_i$. This has $L \cdot O(n^k)$ combinations, recalling that $|N_i| \leq O(n^k)$.
2. For each $j \geq i$, and $v \in N_j$ such that $d(u, v) \leq O(s^j)$, the random radius chosen by (v, j) . Observe that the space around $B(u, O(s^i))$ can be cut by net-points in the same or higher heights that are nearby with respect to their distance scales. As argued in [6], the number of configurations that are relevant to (u, i) is at most $O(k \log n)^{L \cdot O(1)^k} = n^{O(1)^k}$, where $L = O(\log_s n)$ and $s = (\log n)^{\Theta(\frac{1}{k})}$.
3. For each $j > i$, which cluster at height j (specified by the cluster center $v_j \in N_j$) contains the current cluster at height i . This has $O(1)^{kL} = n^{O(\frac{k^2}{\log \log n})}$ combinations.

To define the dynamic program, we start by defining the entries.

Entries of DP. We define entries as $(C, (R, Y), (\text{BAS}, \text{NBAS}), (g, P))$. Define $U := \text{Dis}(\text{BAS} \cup \text{NBAS})$. We define the following *internal constraints* for entries, where the parameters m, r are as defined in Lemma 5.19, and ρ is as defined in Lemma 5.16:

- C is a cluster.
- R is a subset of the m predefined portals such that $|R| \leq r$. This will denote the active portals.
- $Y \subset 2^R$ is a partition of R . We intend to use it to record the subsets of portals that are connected inside C .
- BAS and NBAS are collections of subclusters of C such that $\text{BAS} \cap \text{NBAS} = \emptyset$ and $|\text{BAS} \cup \text{NBAS}| \leq \rho$, and the centers of the clusters in $\text{BAS} \cup \text{NBAS}$ are chosen from $\text{Can}(C)$. Moreover, $e \in \text{BAS}$ implies that any sibling cluster of e is in $\text{BAS} \cup \text{NBAS}$. We intend to use this to record the basic cells and nonbasic cells.
- g is a mapping from U to 2^Y . For some $e \in U$, we intend to use $g(e)$ to

denote the portals to which e connects inside C .

- $P \subset 2^Y$ is a partition of Y such that for all $e \in U$, $g(e) = Q$ implies that Q is a subset of a part in P . The intended use of P is to denote the portals that are to be connected outside C .

We only consider the entries that satisfy the internal constraints. We capture the intended use of an entry formally as follows.

DEFINITION 5.23 (compatibility). *Suppose F is a graph on the metric space and E is an entry. Let $E := (C, (R, Y), (\text{BAS}, \text{NBAS}), (g, P))$. Define $F' := F|_{C \cup R}$. We say F is compatible to E if F' satisfies the following:*

1. A part y is in Y if and only if F' connects all the portals in y .
2. BAS covers all components of F' that intersect R .
3. For $e \in U$, $g(e)$ is exactly the collection of subsets of Y to which e is connected by F' .
4. Every terminal in C is visited by F' .
5. Every isolated terminal of C is connected to at least one portal in R by F' .
6. Every terminal pair that both lie in C is either in the same component of F' , or they are connected to y_1 and y_2 in Y by F' and $\{y_1, y_2\}$ is a subset of a part in P .

We bound the number of entries in the following lemma.

LEMMA 5.24 (number of entries). *There are at most $O(n^{O(1)^k}) \cdot O(\kappa mr)^{O(k)^k \cdot \rho r}$ number of entries. Moreover, for any fixed cluster C , the number of entries with C as the cluster is at most $O(\kappa mr)^{O(k)^k \cdot \rho r}$. (κ is defined as in Lemma 5.15.)*

Proof. Since R is a set of at most r portals chosen from m predefined portals, there are at most $O(m^r)$ possibilities of R . Then, after R is fixed, there are $O(r^r)$ possibilities of Y , since Y is a partition of Y and $|R| \leq r$.

To count the number of BAS and NBAS, we count the union $S := \text{BAS} \cup \text{NBAS}$ of them, and then for any fixed S , we count the number of ways to assign elements in S to BAS and NBAS. Since it is required that the centers of clusters in S be chosen from $\text{Can}(C)$, to form S , we first choose at most ρ centers from $\text{Can}(C)$. There are at most $O(\kappa^\rho)$ possibilities for this by Lemma 5.15. For each chosen center u that is of height i_u , we count the number of configurations of the cluster C_u centered at u . Since C is already fixed, we only need to consider relevant radii for clusters of height less than $\text{ht}(C)$ and at least i_u . Since $u \in \text{Can}(C)$, and for $j \geq i_u$ there are $O(1)^k$ clusters of height- j can affect u , we conclude that there are at most $O(k \log n)^{O(1)^k \cdot \log_s(\frac{1}{\gamma_0})} \leq O(k \log n)^{O(k)^k}$ configurations for C_u . Since $|S| \leq \rho$, there are at most $O(k \log n)^{O(k)^k \cdot \rho}$ configurations for all clusters in S , for any given centers. Therefore, there are $O(\kappa)^{O(k)^k \cdot \rho}$ possibilities for S in total. Then, we assign elements in S to one of BAS, NBAS, and there are at most $2^{|S|} \leq 2^\rho$ numbers of them. In conclusion, the number of possibilities for BAS and NBAS is at most $O(\kappa)^{O(k)^k \cdot \rho}$.

With S fixed, we count the number of possibilities of g . Since g is a mapping from U to 2^Y , the number of such a mapping is at most $O((2^{|Y|})^{|U|}) \leq O(2^{\rho \cdot r})$. Finally, observe that P is a partition of Y , and $|Y| \leq r$. This implies that P has at most $O(r^r)$ possibilities.

Therefore, after fixing C , there are at most $O(\kappa mr)^{O(k)^k \cdot \rho r}$ possibilities for tuples $((R, Y), (\text{BAS}, \text{NBAS}), (g, P))$.

We then count the number of possibilities of C . Observe that there are $O(n)^{O(k)}$ centers for C . For a fixed center, since the number of configurations is at most $n^{O(1)^k}$,

we conclude that there are at most $O(n^{O(1)^k}) \cdot O(\kappa m r)^{O(k)^k \cdot \rho r}$ entries in total. \square

After we define the entries, we shall (recursively) define the value that is associated with each entry. The intended value of an entry E is the weight of the minimum graph that is *recursively compatible* with E (see Definition 5.28).

DEFINITION 5.25 (child entry collection). *Consider an entry E with cluster C . We say a collection of entries $\{(C_i, (R_i, Y_i), (\text{BAS}_i, \text{NBAS}_i), (g_i, P_i))\}_i$ is a child entry collection of E if $\{C_i\}_i$ is a partition of C with $\text{ht}(C_i) = \text{ht}(C) - 1$ for all i .*

DEFINITION 5.26 (portal graph). *We say a graph G is a portal graph of a collection of entries $I := \{(C_i, (R_i, Y_i), (\text{BAS}_i, \text{NBAS}_i), (g_i, P_i))\}_i$ if the vertex set of G is $\cup_i R_i$.*

DEFINITION 5.27 (consistency checking). *Consider an entry E denoted as $E := (C, (R, Y), (\text{BAS}, \text{NBAS}), (g, P))$, a child entry collection I of E denoted as $I := \{(C_i, (R_i, Y_i), (\text{BAS}_i, \text{NBAS}_i), (g_i, P_i))\}_i$, and a portal graph G of I . We say G and I are consistent with E if all checks in the following procedure are passed:*

1. Check whether $\cup_i (C_i \cup R_i) = C \cup R$.
2. We shall define Y' to be a partition of $R' := \cup_i R_i$. Initialize $Y' := \cup_i Y_i$, and whenever there are $y_1, y_2 \in Y'$ connected by G or $y_1 \cap y_2 \neq \emptyset$, replace them by the union of them. Check whether Y' restricted to R is exactly Y .
3. For each $e \in \text{BAS}$, check whether there exists i and $e' \in \text{BAS}_i$ such that $e' = e$ or e' is a child cluster of e .
4. For each $e \in \text{NBAS}$, check whether either there exists i and $e' \in \text{BAS}_i \cup \text{NBAS}_i$ such that $e = e'$, or all child clusters of e are in $\cup_i (\text{BAS}_i \cup \text{NBAS}_i)$.
5. Define g'_i to be a mapping from U_i to 2^Y , where $g'_i(e) := \{y \cap R \mid y \in Y' \wedge \exists y' : (y' \in g_i(e) \wedge y \cap y' \neq \emptyset)\}$ for $e \in U_i$. Here $g'_i(e)$ intends to mean the parts in Y to which e connects, which is defined by “extending” $g_i(e)$ with respect to G . For each i and $u \in \text{BAS}_i$, if there exists $e \in U_i$ such that $e \subset u$ and $g'_i(e) \neq \emptyset$, then check whether there exists $u' \in \text{BAS}$ such that $u = u'$ or u is a child cluster of u' .
6. Define a mapping g' from U to 2^Y , where $g'(e) := \cup_i \cup_{e' \in U_i : e' \subset e} g'_i(e')$ for $e \in U$. Check whether g' is exactly g . We observe that here we consider $e' \subset e$ only, and we shall see later why this is sufficient.
7. For each i , for each $y_1, y_2 \in Y_i$ ($y_1 \neq y_2$) such that y_1, y_2 are in the same part of P_i , check whether either there exists $y \in Y'$ such that $y_1 \cup y_2 \subset y$, or there exist $y'_1, y'_2 \in Y'$ such that $y'_1 \neq y'_2$, $y_1 \subset y'_1$, $y_2 \subset y'_2$, $y'_1 \cap R \neq \emptyset$, $y'_2 \cap R \neq \emptyset$, and $\{y'_1 \cap R, y'_2 \cap R\}$ is a subset of a part in P . This intends to check whether the parts in P_i are connected by G or the information in P_i 's is passed to P .
8. For each terminal pair (a, b) such that $a \in C_i$ and $b \in C_j$ for $i \neq j$, suppose $a \in e_i$ and $b \in e_j$ for $e_i \in U_i$ and $e_j \in U_j$. Check whether $g_i(e_i)$ is connected by G to $g_j(e_j)$ or whether $g'_i(e_i) \neq \emptyset$, $g'_j(e_j) \neq \emptyset$, and $g'_i(e_i) \cup g'_j(e_j)$ is a subset of a part in P . This intends to check whether (a, b) are already connected by G ; otherwise, they will be connected outside C .
9. For each isolated terminal a in C , check whether there exists i and $e \in U_i$ such that $a \in e$ and $g'_i(e)$ is nonempty.

DEFINITION 5.28 (recursive compatibility). *Consider some graph F on the metric space and an entry $E := (C, (R, Y), (\text{BAS}, \text{NBAS}), (g, P))$. F is recursively compatible with E if there exists a set S of entries with $E \in S$ and with a unique entry in S that corresponds to each descendant cluster of C such that the following requirements hold:*

- For each $E' := (C', (R', Y'), (BAS', NBAS'), (g', P'))$ in S , we require that $F' := F|_{C' \cup R'}$ be compatible with E' .
- For each $E' := (C', (R', Y'), (BAS', NBAS'), (g', P'))$ in S , suppose that the child entry collection consisting of elements in S is I' , and define $I' := \{(C_t, (R_t, Y_t), (BAS_t, NBAS_t), (g_t, P_t))\}_t$. Define $G' := F|_{\cup_t R_t}$. (Note that G' is a portal graph of I' .) We require that I' and G' be consistent with E' .

Value of entries. For any entry $E := (C, (R, Y), (BAS, NBAS), (g, P))$, we shall define its value $\text{val}(E)$. The height-0 clusters are corresponding to the base cases. In particular, for any $C := \{x\}$ that is a height-0 cluster, we define entries with such C and with $BAS := \{C\}$, $NBAS := \emptyset$, $R := C$, $Y := \{R\}$, $g(C) := Y$, $P := \{Y\}$ to be the base entries. All base entries have value 0. All other (nonbase) entries with height-0 clusters have value ∞ .

We then define $\text{val}(E)$ when $\text{ht}(C) \neq 0$. Define \mathcal{I}_E to be the set of tuples (I, G) such that I is a child entry collection of E and G is a portal graph of I , and I, G are consistent. The value of E is defined as $\text{val}(E) := \min_{(I, G) \in \mathcal{I}_E} \{w(G) + \text{val}(I)\}$, where $\text{val}(I) = \sum_{E' \in I} \text{val}(E')$. As we shall see in Lemma 5.32, for any entry E , if $\text{val}(E) \neq \infty$, then there actually exists a graph that is recursively compatible with E with weight $\text{val}(E)$.

LEMMA 5.29 (counting \mathcal{I}_E). *For any entry E , the number of possibilities of \mathcal{I}_E is at most $O(k \log n)^{O(s)^k} \cdot O(\kappa m r)^{O(sk)^{O(k)} \cdot pr^2}$, where κ is defined as in Lemma 5.15.*

Proof. Define $E := (C, (R, Y), (BAS, NBAS), (g, P))$. We first bound the number of possibilities of child entry collections $I := \{(C_i, (R_i, Y_i), (BAS_i, NBAS_i), (g_i, P_i))\}_i$ of C . To define I , we start by defining $\{C_i\}_i$. By the packing property, there are at most $O(s)^k$ centers for the child clusters of C . For each center u of the child cluster, there are at most $O(k \log n)$ possible radii. Hence, there are at most $O(k \log n)^{O(s)^k}$ possibilities for $\{C_i\}_i$.

By Lemma 5.24, for any fixed C_i , there are at most Z possibilities for tuples $((R_i, Y_i), (BAS_i, NBAS_i), (g_i, P_i))$, where $Z := O(\kappa m r)^{O(k)^k \cdot pr}$. Therefore, there are at most $Z^{O(s)^k} \cdot O(k \log n)^{O(s)^k}$ possibilities of I .

For a fixed I , the vertex set of the portal graph G of I is fixed, and there are at most $O(s)^k \cdot r$ vertices in G . Then, the number of possibilities of G for a fixed I is at most the number of edge sets, and it is at most $2^{O(s)^{O(k)} \cdot r^2}$, since there are at most $(O(s)^k \cdot r)^2$ edges.

In conclusion, there are at most $O(k \log n)^{O(s)^k} \cdot Z^{O(s)^k} \cdot 2^{O(s)^{O(k)} \cdot r^2}$ possibilities of \mathcal{I}_E , which is at most $O(k \log n)^{O(s)^k} \cdot O(\kappa m r)^{O(sk)^{O(k)} \cdot pr^2}$. \square

Final entry. The final entry is the entry with C being the root cluster, $R, BAS, NBAS$ being \emptyset , and Y, g, P being uniquely defined from $R, BAS, NBAS = \emptyset$. We use the value of the final entry as the output of DP.

Evaluating the final entry Although we only care about the value of the final entry, it may be necessary to evaluate the value of other entries. We shall define a (recursive) algorithm in Definition 5.30 that takes an entry and returns the value of the input. To get the value of the final entry which is the output of DP, we invoke the algorithm with the final entry as the input.

We note that the counting arguments in Lemmas 5.24 and 5.29 can both be naturally implemented as algorithms, with additional $O(n^{O(k)})$ factors in the running time compared with the corresponding counting bounds. We will make use of these implementations as subroutines in Definition 5.30. Moreover, the natural implementation

of the consistency checking procedure in Definition 5.27 runs in time $O(n^{O(k)})$.

DEFINITION 5.30 (algorithm for evaluating value of entries). *We define a recursive procedure that evaluates the value of an input entry E with cluster C :*

- If $\text{ht}(C) = 0$, then the value of it is already defined, and we return its value.
- If $\text{ht}(C) > 0$ and $\text{val}(E)$ is already calculated, then we return the calculated value.
- Otherwise, $\text{ht}(C) > 0$ and $\text{val}(E)$ has not yet calculated. The following procedure is executed:
 1. Set the default value for $\text{val}(E) := \infty$.
 2. Calculate \mathcal{I}_E .
 3. For each element $(I, G) \in \mathcal{I}_E$, use the consistency checking procedure defined in Definition 5.27 to check whether I and G are consistent with E . If they are consistent, then recursively use this procedure to calculate $\text{val}(I) + w(G)$, and update $\text{val}(E)$ if $\text{val}(E) + w(G)$ is smaller than $\text{val}(E)$.
 4. Finally, return $\text{val}(E)$ as the output.

LEMMA 5.31 (running time). *The running time for the algorithm defined in Definition 5.30 is at most $O(n^{O(1)^k}) \cdot \exp(\sqrt{\log n}) \cdot O(\frac{k}{\epsilon})^{O(k)}$.*

Proof. Suppose the input is $E := (C, (R, Y), (\text{BAS}, \text{NBAS}), (g, P))$. We observe that once the value for some entry is calculated, it would not be calculated again, and recalling the value takes constant time. Then, we shall bound the time when $\text{val}(E)$ is not yet calculated and $\text{ht}(C) \neq 0$.

Observe that for any given I with $\text{val}(E')$ for all $E' \in I$ known and a graph G such that $(I, G) \in \mathcal{I}_E$, evaluating $\text{val}(I) + w(G)$ takes $O(n)^{O(k)}$ time. Therefore, combining with Lemmas 5.24 and 5.29, there are at most $O(n^{O(1)^k}) \cdot Z$ entries, and it takes $O(n)^{O(k)} \cdot O(k \log n)^{O(s)^k} \cdot O(\kappa m r)^{O(sk)^{O(k)} \cdot \rho r^2}$ time to evaluate each. In conclusion, the time for evaluating all the entries is at most $O(n^{O(1)^k}) \cdot O(k \log n)^{O(s)^k} \cdot O(\kappa m r)^{O(sk)^{O(k)} \cdot \rho r^2}$.

Substituting parameters. Recall that we consider $q \leq O(s)^{O(k)} \cdot q_0$. Observe that $\frac{1}{\gamma_0} := \lceil \frac{1}{\gamma_0} \rceil_s \leq O(\frac{ks^3L}{\epsilon})$, and $\frac{1}{\gamma_1} := \lfloor \frac{1}{\gamma_1} \rfloor_s \leq O(\frac{s^2}{\epsilon})$. Substituting γ_0 and γ_1 , we have $\kappa \leq O(\frac{ksL}{\epsilon})^{O(k)}$ and $\rho \leq O(\frac{sk}{\epsilon})^{O(k)}$. Moreover,

$$r := O(1)^k \cdot q \log_s \log n + O\left(\frac{k}{\epsilon}\right)^k + O\left(\frac{s}{\epsilon}\right)^k \leq O\left(\frac{sk}{\epsilon}\right)^{O(k)}, \quad m \leq O\left(\frac{skL}{\epsilon}\right)^k.$$

By definition, $s := (\log n)^{\frac{\epsilon}{k}}$, $L := O(\log_s n) = O(\frac{k \log n}{c \log \log n})$. Therefore, the running time is at most

$$\begin{aligned} & O(n^{O(1)^k}) \cdot O(k \log n)^{O(s)^k} \cdot O(\kappa m r)^{O(sk)^{O(k)} \cdot \rho r^2} \\ & \leq O(n^{O(1)^k}) \cdot O(k \log n)^{O(s)^k} \cdot O\left(\frac{ksL}{\epsilon}\right)^{O(\frac{sk}{\epsilon})^{O(k)}} \\ & \leq O(n^{O(1)^k}) \cdot \exp\left(O(s)^{O(k)} \cdot O\left(\frac{k}{\epsilon}\right)^{O(k)} \cdot \log \frac{k \log n}{\epsilon}\right) \\ & \leq O(n^{O(1)^k}) \cdot \exp\left(O\left(\frac{k}{\epsilon}\right)^{O(k)} \cdot O(\log n)^{O(c)} \cdot \log \log n\right). \end{aligned}$$

By choosing constant c to be sufficiently small so that $O(\log n)^{O(c)} \cdot \log \log n \leq$

$O(\sqrt{\log n})$, we conclude that the running time is at most $O(n^{O(1)^k}) \cdot \exp(\sqrt{\log n} \cdot O(\frac{k}{\epsilon})^{O(k)})$. \square

LEMMA 5.32 (characterizing the value of entries). *For any entry E such that $\text{val}(E) \neq \infty$, where $E := (C, (R, Y), (\text{BAS}, \text{NBAS}), (g, P))$, $\text{val}(E)$ is the weight of the minimum weight graph that is recursively compatible with E and uses points in $C \cup R$ only.*

Proof. For the clusters of height 0, the lemma holds trivially.

Assuming the lemma holds for all entries with the clusters of height $i - 1$, we prove the lemma for an entry E with C of height i centered at $u \in N_i$, where $i \geq 1$. We shall first show that $\text{val}(E)$ is the weight of some graph that is recursively compatible to the entry and uses points in $C \cup R$ only. Then, we shall show that the value is minimum.

Feasibility. Suppose $(I, G) := \arg \min_{(I', G') \in \mathcal{I}_E} \{\text{val}(I') + w(G')\}$. Define $I = \{E_j\}_j$, where $E_j := (C_j, (R_j, Y_j), (\text{BAS}_j, \text{NBAS}_j), (g_j, P_j))$. Since $\text{val}(E) \neq \infty$, we have $\text{val}(I') \neq \infty$. For $E_j \in I$, by assumption, there exists a graph that is recursively compatible with E_j and uses points in $C_j \cup R_j$ only, and we denote it as F_j . We define a graph F that is the union of F_j for all j , and we define G . Then $w(F) = \text{val}(I) + w(G)$.

We shall show that F is recursively compatible with E . Since $(I, G) \in \mathcal{I}_E$, I and G are consistent with E . Since F_j is recursively compatible to E_j for all j , it remains to verify that F is compatible with E . When we say ‘‘consistency checking procedure,’’ we refer to Definition 5.27:

- F uses points in $C \cap R$ only. This is by definition.
- A part y is in Y if and only if F connects all the portals in the part y . This is by the step 2 of the consistency checking procedure.
- **BAS** covers all components of F that intersect R . This is by step 5 of the consistency checking procedure.
- For $e \in U$, the collection of subsets of Y that e is connected to by F is exactly $g(e)$. We note that steps 3 and 4 of the consistency checking procedure, together with the internal constraint that $e' \in \text{BAS}_j$, imply that any sibling cluster of e' is in $\text{BAS}_j \cup \text{NBAS}_j$ for all j . This implies that $\cup_j (\text{BAS}_j \cup \text{NBAS}_j)$ is a refinement of $\text{BAS} \cup \text{NBAS}$. Then, for each j , for each $e' \in U_j$, and for each $e \in U$, either $e' \subset e$ or $e' \cap e = \emptyset$. Therefore, step 6 is sufficient to ensure this item. (If e' is not a subset of e but $e' \cap e \neq \emptyset$, then the g_j mappings in the subentries do not have sufficient information to determine the portals to which $e' \cap e$ is connected.)
- Every terminal in $C \cup R$ is visited by F . This is by the construction of F , and by F_j it is recursively compatible with E_j for all j .
- Every isolated terminal of C is connected to at least one portal in R by F . This is by step 9 of the consistency checking procedure.
- Every terminal pair that both lie in C is either in the same component of F , or they are connected to y_1 and y_2 in Y by F and $\{y_1, y_2\}$ is a subset of a part in P . This is by steps 7 and 8 of the consistency checking procedure, and by F_j it is recursively compatible with E_j for all j .

This implies that F is recursively compatible with E .

Optimality. Then, we shall show that $\text{val}(E)$ is minimum. Suppose not. Define l as the weight of the minimum weight graph that is recursively compatible with E and uses points in $C \cup R$ only. Define F' to be the corresponding graph recursively compatible with E with weight l . Since F' is recursively compatible with E , there exist $I' := \{E_t := (C_t, (R_t, Y_t), (\text{BAS}_t, \text{NBAS}_t), (g_t, P_t))\}_t$ and a portal graph G' of

I' that are consistent with E . Moreover, there exists a graph F_t that is recursively compatible with E_t for all t .

We note that $(I', G') \subset \mathcal{I}_E$. Therefore, $\sum_t w(F_t) + w(G') = l < \text{val}(E) \leq \text{val}(I') + w(G')$. This implies that $\sum_t w(F_t) < \sum_{E' \in I'} \text{val}(E')$, and hence there exists t such that $w(F_t) < \text{val}(E_t)$. However, we know that F_t is recursively compatible with E_t , and by assumption, $\text{val}(E_t) \leq w(F_t)$. This is a contradiction. \square

COROLLARY 5.33. *There exists a feasible solution to SFP whose weight is the value of the final entry.*

LEMMA 5.34 (good solution is recursively compatible). *Suppose for each $i \in [L]$ and $u \in N_i$ that $O(k \log n)$ radii are fixed. Suppose F is an (m, r) -light solution such that Eff satisfies the cell property in terms of F under one of the hierarchical decompositions defined by the radii. Then, the value of the final entry is at most $w(F)$.*

Proof. We shall show that F is recursively compatible with the final entry, and then Lemma 5.32 implies that the value of the final entry is at most $w(F)$.

Suppose we fix a hierarchical decomposition induced from the given radii such that F is (m, r) -light and Eff satisfies the cell property in terms of F . For each cluster C in the decomposition, we define $F_C := F_{C \cup R}$ and define an entry $E_C := (C, (R_C, Y_C), (\text{BAS}_C, \text{NBAS}_C), (g_C, P_C))$ as follows, where R is the set of active portals for C :

- $R_C := R$.
- Y_C contains a part y if and only if portals in y are connected by F_C .
- $\text{BAS}_C := \text{Bas}(C)$, $\text{NBAS}_C := \text{NBas}(C)$.
- For each $e \in U_C$, let $g_C(e) := Q$, where Q is the collection of parts in Y_C to which e is connected by F_C .
- Define P_C to be any one that satisfies the following:
 1. For each $e \in U_C$, $g_C(e) = Q$ implies Q is a subset of P_C .
 2. For each terminal pair (a, b) that both lie in C , if they are not connected by F_C , then the subsets of portals to which a and b are connected by F_C are in a same part of P_C .

The internal constraints for an entry are satisfied from the definition of the cells, the fact that Eff satisfies the cell property, and Lemmas 5.15 and 5.16.

Then we (uniquely) define $I_C := \{(C_i, (R_i, Y_i), (\text{BAS}_i, \text{NBAS}_i), (g_i, P_i))\}_i$ as the child collection of E_C and define $G_C := F|_{\bigcup_i R_i}$ as the portal graph of I_C .

We then check that I_C and G_C are consistent with E_C . Steps 1, 2, 7, 8, and 9 are immediate. Steps 3, 4, and 5 follow from the definition of the basic cells and nonbasic cells. Inside step 6, we observe that $g'(e)$ is evaluated by looking at $e' \subset e$ only (instead of considering all $e' \in U_i$) for $e' \in U_i$, for some i , and $e \in U$. However, this is indeed sufficient, since Lemma 5.13 asserts that for any $e \in U$, $e' \in U_i$, for any i , either $e' \subset e$ or $e \cap e' = \emptyset$.

It remains to check the following for E_C , for each cluster C .

- A part $y \in Y_C$ if and only if F_C connects all the portals in the part y . This is by definition.
- BAS covers all components of F_C that intersect R_C . This is by definition.
- For $e \in U_C$, the collection of subsets of Y_C to which e is connected by F_C is exactly $g_C(e)$. This is by definition.
- Every terminal in $C \cup R_C$ is visited by F_C . This is by the feasibility of F .
- Every isolated terminal of C is connected to at least one portal in R_C by F_C . This is by the feasibility of F .
- Every terminal pair that both lie in C is either in the same component of F_C ,

or they are connected to y_1 and y_2 in Y_C by F_C and $\{y_1, y_2\}$ is a subset of a part in P_C . This is by definition.

This finishes the proof. \square

Combining Lemmas 5.19 and 5.34, Corollary 5.33, and Lemma 5.31, we conclude a PTAS for sparse SFP instances.

COROLLARY 5.35 (PTAS for sparse SFP instances). *For an instance of SFP that has a q -sparse optimal net-respecting solution, algorithm DP returns a $(1 + \epsilon)$ solution with constant probability, running in time $O(n^{O(1)^k}) \cdot \exp(\sqrt{\log n} \cdot O(\frac{k}{\epsilon})^{O(k)})$ for $q \leq O(s)^{O(k)} \cdot q_0$.*

REFERENCES

- [1] I. ABRAHAM, Y. BARTAL, AND O. NEIMAN, *Advances in metric embedding theory*, in STOC, ACM, 2006, pp. 271–286.
- [2] A. AGRAWAL, P. KLEIN, AND R. RAVI, *When trees collide: An approximation algorithm for the generalized Steiner problem on networks*, SIAM J. Comput., 24 (1995), pp. 440–456, <https://doi.org/10.1137/S0097539792236237>.
- [3] S. ARORA, *Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems*, J. ACM, 45 (1998), pp. 753–782.
- [4] S. ARORA, *Approximation algorithms for geometric TSP*, in The Traveling Salesman Problem and Its Variations, Comb. Optim. 12, Kluwer Academic Publishers, Dordrecht, The Netherlands, 2002, pp. 207–221.
- [5] P. ASSOUD, *Plongements lipschitziens dans \mathbf{R}^n* , Bull. Soc. Math. France, 111 (1983), pp. 429–448.
- [6] Y. BARTAL, L. GOTTLIEB, AND R. KRAUTHGAMER, *The traveling salesman problem: Low-dimensionality implies a polynomial time approximation scheme*, in STOC, ACM, 2012, pp. 663–672.
- [7] M. BATENI, E. D. DEMAINE, M. HAJIAGHAYI, AND D. MARX, *A PTAS for planar group Steiner tree via spanner bootstrapping and prize collecting*, in STOC, ACM, 2016, pp. 570–583.
- [8] M. BATENI AND M. HAJIAGHAYI, *Euclidean prize-collecting Steiner forest*, Algorithmica, 62 (2012), pp. 906–929.
- [9] M. BATENI, M. T. HAJIAGHAYI, AND D. MARX, *Approximation schemes for Steiner forest on planar graphs and graphs of bounded treewidth*, J. ACM, 58 (2011), no. 5, 21.
- [10] G. BORRADAILE, P. N. KLEIN, AND C. MATHIEU, *A polynomial-time approximation scheme for Euclidean Steiner forest*, in FOCS, IEEE Computer Society, 2008, pp. 115–124.
- [11] J. BYRKA, F. GRANDONI, T. ROTHVOß, AND L. SANITÀ, *An improved lp-based approximation for Steiner tree*, in STOC, ACM, 2010, pp. 583–592.
- [12] T.-H. H. CHAN AND S. H.-C. JIANG, *Reducing curse of dimensionality: Improved PTAS for TSP (with neighborhoods) in doubling metrics*, in SODA, SIAM, 2016, pp. 754–765, <https://doi.org/10.1137/1.9781611974331.ch54>.
- [13] M. CHLEBÍK AND J. CHLEBÍKOVÁ, *The Steiner tree problem on graphs: Inapproximability results*, Theoret. Comput. Sci., 406 (2008), pp. 207–214.
- [14] E. D. DEMAINE, M. T. HAJIAGHAYI, AND P. N. KLEIN, *Node-weighted Steiner tree and group Steiner tree in planar graphs*, ACM Trans. Algorithms, 10 (2014), no. 3, 13.
- [15] M. M. DEZA AND M. LAURENT, *Geometry of Cuts and Metrics*, Algorithms Combin. 15, Springer-Verlag, Berlin, 1997.
- [16] D. DU, F. HWANG, AND S. CHAO, *Steiner minimal tree for points on a circle*, Proc. Amer. Math. Soc., 95 (1985), pp. 613–618.
- [17] D.-Z. DU, F. K. HWANG, AND J. WENG, *Steiner minimal trees for regular polygons*, Discrete Comput. Geom., 2 (1987), pp. 65–84.
- [18] N. GARG, G. KONJEVOD, AND R. RAVI, *A polylogarithmic approximation algorithm for the group Steiner tree problem*, J. Algorithms, 37 (2000), pp. 66–84.
- [19] M. X. GOEMANS AND D. P. WILLIAMSON, *A general approximation technique for constrained forest problems*, SIAM J. Comput., 24 (1995), pp. 296–317, <https://doi.org/10.1137/S0097539793242618>.
- [20] A. GUPTA, R. KRAUTHGAMER, AND J. R. LEE, *Bounded geometries, fractals, and low-distortion embeddings*, in FOCS, IEEE Computer Society, 2003, pp. 534–543.
- [21] A. GUPTA AND A. KUMAR, *Greedy algorithms for Steiner forest*, in STOC, ACM, 2015, pp. 871–

- 878.
- [22] E. HALPERIN AND R. KRAUTHGAMER, *Polylogarithmic inapproximability*, in STOC, ACM, 2003, pp. 585–594.
 - [23] M. HAUPTMANN AND M. KARPIŃSKI, *A Compendium on Steiner Tree Problems*, Research Report 85343, Institut für Informatik, University of Bonn, Bonn, Germany, 2013.
 - [24] J. MATOUŠEK, *Lectures on Discrete Geometry*, Grad. Texts in Math. 212, Springer-Verlag, New York, 2002.
 - [25] K. TALWAR, *Bypassing the embedding: Algorithms for low dimensional metrics*, in STOC, ACM, 2004, pp. 281–290.