

A QPTAS for TSP with Fat Weakly Disjoint Neighborhoods in Doubling Metrics

T-H. Hubert Chan*

Khaled Elbassioni†

Abstract

We consider the Traveling Salesman Problem with Neighborhoods (TSPN) in doubling metrics. The goal is to find a shortest tour that visits each of a collection of n subsets (*regions* or *neighborhoods*) in the underlying metric space. We give a QPTAS when the regions are what we call α -fat weakly disjoint. This notion combines the existing notions of diameter variation, fatness and disjointness for geometric objects and generalizes these notions to any arbitrary metric space. Intuitively, the regions can be grouped into a bounded number of types, where in each type, the regions have similar upper bounds for their diameters, and each such region can designate a point such that these points are far away from one another.

Our result generalizes the PTAS for TSPN on the Euclidean plane by Mitchell [27] and the QPTAS for TSP on doubling metrics by Talwar [30]. We also observe that our techniques directly extend to a QPTAS for the Group Steiner Tree Problem on doubling metrics, with the same assumption on the groups.

1 Introduction

We consider the Traveling Salesman Problem with Neighborhoods (TSPN) in a metric space (V, d) . An instance of the problem is given by a collection W of n subsets $\{P_1, P_2, \dots, P_n\}$ in V . Each subset $P_j \subset V$ is known as a *neighborhood* or *region*. The objective is to find a minimum length tour that visits at least one point from each region.

This problem generalizes the well-known Traveling Salesman Problem (TSP), for which there are PTAS's for low-dimensional Euclidean metrics [26, 3, 28], and a QPTAS for doubling metrics [30]. The neighborhood version of the problem was first introduced by Arkin and Hassin [1], who gave constant approximation for the case when the regions are in the plane

and “well-behaved” (e.g., disks, parallel and similar length segments, bounded ratio between the largest and smallest diameters). The general version of the problem was shown to have an inapproximability threshold of $\Omega(\log^{2-\epsilon} n)$ for any $\epsilon > 0$ by Halperin and Krauthgamer [20]. There is an almost matching upper bound of $O(\log N \log k \log n)$ -approximation, using the results of Garg et al. [17] and Fakcharoenphol et al. [15], where N is the total number of points in V and k is the maximum number of points in each region.

Special cases are considered where (V, d) is taken to be the Euclidean plane. However, if the regions are allowed to be intersecting connected subsets, the problem remains APX-hard [9, 29]. Further restrictions are placed on the regions. For connected polygonal regions, Mata and Mitchell [24] gave an $O(\log n)$ -approximation, and Gudmundsson and Levkopoulos [18] reduced the running time to $O(N^2 \log N)$, where N is the total number of vertices of the polygons.

Regions are often assumed to be “fat”¹ and disjoint. In fact, no constant factor approximation algorithm is known for the case of intersecting non-fat regions. Dumitrescu and Mitchell [12] considered connected regions that are all about the same size, fat and disjoint, and gave a PTAS in this case, using the “guillotine” method.

Berg et al. [9] gave constant approximation for slightly more general regions of varying size, but are still disjoint, fat and convex. Elbassioni et al. [14] generalized this to the discrete case where each neighborhood consists of a discrete set of points in a fat though not necessarily convex region, and gave a constant approximation. This constant approximation was further generalized in [13], where the neighborhoods are intersecting, connected and have comparable diameters.

The best previously known result for getting a $(1 + \epsilon)$ -approximation is by Mitchell [27], who obtained a PTAS for the Euclidean plane, where the regions are fat and almost disjoint. This result is obtained by the “guillotine subdivision” technique, which unfortunately

*Department of Computer Science, University of Hong Kong, Hong Kong (hubert@cs.hku.hk). This research was done while the author was at Max-Planck-Institut für Informatik, 66123 Saarbrücken, Germany.

†Max-Planck-Institut für Informatik, 66123 Saarbrücken, Germany (elbassio@mpi-inf.mpg.de).

¹Intuitively, the fatness of a region measures the ratio between the smallest circumscribing radius and the largest inscribing radius. A disk is fat, while a line segment is not.

only works for 2 dimensions. On the other hand, the hierarchical decomposition technique by Arora [3] and Talwar [30] is applicable to more general metrics. However, as pointed out by Mitchell [27], previous attempts in applying this technique have led to only limited success.

Feremans and Grigoriev [16] suggested a PTAS by using Arora’s framework [3] in the case where regions are of similar size and bounded perimeter. They proposed that the regions are never divided in the decomposition. However, some technical issues would arise,² but unfortunately they were not addressed in the short paper. We instead allow regions to be divided, but have to overcome a major technical hurdle and successfully use Arora’s framework to design an approximation scheme for TSPN.

Our Contribution. We give a $(1 + \varepsilon)$ -approximation for instances on metrics with bounded doubling dimension [4, 7, 19].³ This includes low-dimensional Euclidean metrics, and hence is a generalization of Mitchell’s result [27] for 3 or more dimensions. Moreover, since the doubling dimension is well defined for any metric, our framework covers metrics that do not have any geometric structure, and the regions need not be convex or even connected, where such notions might not even be applicable in the first place. For more applications of doubling metrics, the reader is referred to [6, 22, 23, 21].

Nevertheless, we still need to place some restrictions on the regions, because the problem is APX-hard in general on the plane [11], which has bounded doubling dimension. We combine the notions of diameter variation, fatness and disjointness for geometric spaces, and define for regions in general metrics the notion of α -fat weak disjointness. We assume that the regions have Δ types of radii. For the regions within the same type, there is some $\rho > 0$ such that there is a ρ -packing⁴ consisting of one point from each region, and all the regions have diameters at most $O(\alpha\rho)$.

Our definition allows very general regions. Intuitively, all we require is that regions of similar diameters should each designate a point within, such that these points are far away from one another; the regions can otherwise intersect arbitrarily. The assumption that there are only a bounded number Δ of types of region diameters is also necessary, as we show in Appendix B

²In particular, keeping regions intact would destroy the so-called “padding property” of the decomposition, which is essential in Arora’s argument [3] for the existence of a good *portal respecting* tour.

³Intuitively, a set has bounded doubling dimension if any set can be covered by a bounded number of sets with half its diameter.

⁴A ρ -packing is a set of points with inter-point distance larger than ρ .

that otherwise the problem remains APX-hard⁵. Of course, the catch with working on such weak assumptions is that the running time of our algorithm is only quasi-polynomial, which is not surprising, because there is only a QPTAS known even for TSP on doubling metrics by Talwar [30].

Main Result. We augment the hierarchical decomposition method [3, 30] for TSP to give a randomized algorithm that approximates TSPN.

THEOREM 1.1. *Suppose that we are given an instance of TSPN, where the underlying metric space has doubling dimension at most k , and the regions are α -fat weakly disjoint with at most Δ types of radii. Then, there is a QPTAS that, with constant probability, gives a TSP tour of length at most $(1 + \varepsilon)OPT$ in time $\exp\{O(\frac{\Delta}{\varepsilon})^k O(\alpha)^{k^2} \log^k n\}$.*

For the case of Euclidean metrics, we can remove the dependence on Δ if we use a stronger notion of fatness as in [9, 31].

THEOREM 1.2. *Suppose that we are given an instance of TSPN, where the underlying metric is the k -dimensional Euclidean space, and the regions are disjoint and α -fat in the sense defined in [9]. Then, there is a QPTAS that, with constant probability, gives a TSPN tour of length at most $(1 + \varepsilon)OPT$ in time $\exp\{O(\frac{1}{\varepsilon})^{O(k)} O(\alpha)^{O(k^2)} \log^{O(k)} n\}$.*

Our Techniques. Our approximation scheme is built on top of the hierarchical decomposition method used for TSP by Arora [3] and Talwar [30]. The main technical hurdle is that a cluster can partially intersect many regions, causing an exponential number of dynamic program entries for that cluster. We resolve this issue via the following approaches.

1. When a region is separated by clusters, we charge the extra cost incurred to the radius of the cluster. The sum of the radii of the clusters can be charged to the length of the optimal tour. This is done by extending a lemma appearing in [14, 27] to doubling metrics, which gives a lower bound on the length of any tour that hits all weakly disjoint regions with similar diameters.
2. By considering the probability that a region is separated by the clusters, we carefully prune the search in lower levels of the dynamic program. The number of partially intersecting regions that a cluster needs to explicitly consider is greatly

⁵However, as we shall see, this assumption is not necessary in the case of Euclidean metric

reduced to poly-logarithmic, and hence this allows the running time of the approximation scheme to be quasi-polynomial.

Extension to Group Steiner Tree Problem (GSTP). Observing that the optimal length for GSTP is at least half of that for TSPN, we have the corresponding version of Corollary 3.1 for GSTP, which leads to a QPTAS using the same techniques.

2 Notation and Preliminaries

We denote a metric space by $M = (V, d)$.⁶ (For basic properties of metric spaces, please refer to standard texts [10, 25].) A ball $B(x, \rho)$ is the set $\{y \in V \mid d(x, y) \leq \rho\}$. The diameter $\text{Diam}(Z)$ of a set Z is the maximum distance between points in Z . A set Z of points is a ρ -packing, if any two distinct points in Z are at a distance more than ρ away from each other.

Problem Definition. An instance of the metric TSP with neighborhoods (TSPN) is given by a metric space $M = (V, d)$ and a collection of n neighborhoods or regions $W := \{P_j \mid j \in [n]\}$, where each P_j is a subset of V . The objective is to find a minimum TSP tour that visits at least one point from each region. We require that the regions satisfy a weak disjointness condition.

DEFINITION 2.1. (α -FAT WEAKLY DISJOINT REGIONS) The regions $\{P_j\}_j$ are α -fat weakly disjoint with Δ types of radii if the regions can be partitioned into Δ sets $\{W_l\}_{l \in [\Delta]}$ such that for each set W_l , the following conditions hold.

1. There exists $\rho_l > 0$ such that for each region P_j in W_l , there exists some point $z_j \in P_j$ such that the set $\{z_j\}_j$ is a ρ_l -packing. We say that the region P_j has center z_j and core radius $r_j := \rho_l$.
2. Every region P_j is contained in the ball $B(z_j, \alpha r_j)$, and we denote $P_j = P_j(z_j, \alpha r_j)$.

Observe that regions from different W_l 's can intersect arbitrarily.

The assumption that there are only a bounded number Δ of types of region radii is necessary, as we show in Appendix B that otherwise the problem remains APX-hard, even if the regions are disjoint balls.

Examples.

(1) Suppose the problem is defined in the Euclidean space, and the regions are continuous disjoint balls, i.e., for each P_j , there exist $z_j \in V$ and $\rho_j \geq 0$ such that $P_j = B(z_j, \rho_j)$. Suppose further that the regions are partitioned into Δ sets $\{W_l\}_{l \in [\Delta]}$ such that any two regions in the same W_l have their radii differ by a multiplicative factor of at most 2. Note that in this case,

$\Delta \leq 1 + \log_2 \max\{\frac{\rho_i}{\rho_j} \mid \rho_j > 0\}$. Suppose that in some W_l , all the regions $P_j = B(z_j, \rho_j)$ satisfy $\rho \leq \rho_j \leq 2\rho$. Then, it follows by the disjointness of the balls that the corresponding $\{z_j\}_j$ forms a 2ρ -packing, and obviously, P_j is contained in $B(z_j, 2\rho)$. Hence, these regions are 1-fat weakly disjoint with Δ types of radii. One may consider, more generally, fat regions in the sense defined by Mitchell [27], and note that his definition is included in ours. Hence, our results apply to the class of fat regions considered in [27].

(2) Suppose the problem is defined on a finite metric, and for each region P_j , there is some $\rho_j > 0$ and $z_j \in V$ such that $B(z_j, \rho_j) \subseteq P_j \subseteq B(z_j, \alpha\rho_j)$. Suppose further that two regions are disjoint if their corresponding ρ_j 's are within a factor of 2 from each other. We can partition the regions into Δ sets $\{W_l\}_{l \in [\Delta]}$ such that any two regions in the same W_l have their ρ_j 's differ by a multiplicative factor of at most 2. One can check that we also have $\Delta \leq 1 + \log_2 \max\{\frac{\rho_i}{\rho_j} \mid \rho_j > 0\}$, and with respect to such $\{W_l\}_l$ the regions are 2α -fat weakly disjoint.

REMARK 2.1. The parameter α in Definition 2.1 depends on how the regions are grouped into the sets $\{W_l\}_l$, and also on how the centers of regions are picked within each set. By decreasing the number Δ of sets W_l , one might possibly increase α . However, we are not concerned about the optimal way to form the sets $\{W_l\}_l$ to obtain the best α and Δ . We just assume that we are given a partition $\{W_l\}_l$ of regions (together with the corresponding core radius and the centers of regions in each such W_l) such that the regions are α -fat weakly disjoint with respect to this $\{W_l\}_l$, for some $\alpha \geq 1$. The only requirement that we need in order to avoid too many W_l 's is that, the ratio of the core radii from two different W_l 's should be at least some constant at least 2.

REMARK 2.2. Observe that the α -fat weak disjointness condition implies that if all the regions in some W_l have diameters at least δ , then the corresponding centers form a $\frac{\delta}{2\alpha}$ -packing. Moreover, the diameters of the regions in W_l are within a factor of 2α from one another.

Restricting the Tour inside B_0 . Without loss of generality, we can assume that there is a region P_0 which contains only one point p_0 . For finite metrics, we can try each p_0 in P_0 , and consider those TSPN tours that pass through p_0 ; for the special case of Euclidean metrics, see Appendix 5. We let R to be the minimum radius of a ball centering at p_0 that intersects all regions. Suppose OPT is the length of the optimal tour. Then, it follows that $2R \leq \text{OPT} \leq 2nR$. Hence, the optimal tour must be contained in the ball $B_0 := B(p_0, nR)$. Therefore,

⁶Observe that V could be an infinite set.

without loss of generality, we only need to consider the points in B_0 .

REMARK 2.3. *Suppose the optimal tour visits p_j in each P_j . If we replace each p_j by $p'_j \in P_j$ such that $d(p_j, p'_j) \leq \frac{\varepsilon R}{2n}$, then we change the length of the tour by at most εOPT . Hence, we can assume that each region has radius of either 0 or at least $\frac{\varepsilon R}{2n}$. However, we can have two regions of large radii that almost touch each other.*

We measure the complexity of the given metric by its doubling dimension.

DEFINITION 2.2. (DOUBLING DIMENSION [4, 19]) *The doubling dimension of a metric space (V, d) is at most k if for all $x \in V$, 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 $z \in V$.*

Observe that a set of points in k -dimensional Euclidean space induces a metric space with doubling dimension at most $O(k)$. Unless otherwise stated, we use only the doubling property of Euclidean metrics, and we give explicit emphasis when the geometric properties of Euclidean metrics are used.

Given $\rho > 0$, recall that a ρ -net for a set U of points is a subset S such that every point in U is within a distance of ρ from some point in S and any two points in S are at a distance of more than ρ away from each other. The following fact states that for a doubling metric, one cannot pack too many points in some fixed ball such that the points are far away from one another.

FACT 2.1. (PACKING IN DOUBLING METRICS [19]) *Suppose Z is a set of points in a metric space with doubling dimension at most k . If Z is contained in some ball of radius $2^s \rho$ and for all $y, z \in Z$ such that $y \neq z$, $d(y, z) > \rho$, then $|Z| \leq 2^{(s+1)k}$.*

On a high level, we use a divide and conquer paradigm. Hence, we would need a desirable scheme for dividing up the metric space. The following decomposition schemes are widely used in the metric embedding literature [5, 15].

DEFINITION 2.3. (PADDED DECOMPOSITION) *Given a finite metric space (V, d) , a positive parameter $D > 0$ and $\beta > 1$, a D -bounded β -padded decomposition is a distribution Π over partitions of V such that the following conditions hold.*

- (a) *For each partition \mathcal{P} in the support of Π , the diameter of every cluster in \mathcal{P} is at most D .*
- (b) *Suppose $S \subseteq V$ is a set with diameter δ . If \mathcal{P} is sampled from Π , then the set S is partitioned by \mathcal{P} with probability at most $\beta \cdot \frac{\delta}{D}$.*

We consider distances of geometrically decreasing scales. Recall the relevant distances are between $\frac{\varepsilon R}{2n}$ and $2nR$. We consider powers of 2, and have $L := \lceil \log_2 \frac{4n^2}{\varepsilon} \rceil$ distance scales. We let $D_L := 4nR$ and $D_{i-1} := \frac{D_i}{2}$, for $1 \leq i \leq L$.

DEFINITION 2.4. (PADDED HIERARCHICAL DECOMPOSITION) *Given a metric space (V, d) , a β -padded hierarchical decomposition is a family $\{\Pi_i\}_i$ of distributions of partitions of (V, d) such that:*

- (a) *Each Π_i is a D_i -bounded β -padded decomposition of (V, d) , and*
- (b) *Suppose a hierarchical partition $\{\mathcal{P}_i\}_i$ is in the support of $\{\Pi_i\}_i$. Then, for $0 \leq i < L$, each cluster in \mathcal{P}_i is completely contained in some parent cluster in \mathcal{P}_{i+1} .*

FACT 2.2. (PADDED HIERARCHICAL DECOMPOSITION FOR k -DIMENSIONAL EUCLIDEAN METRICS [3]) *Suppose a metric space resides in k -dimensional Euclidean space. Then, the randomly shifted quadtree construction in [3] gives a k -padded hierarchical decomposition. Moreover, for any hierarchical partition sampled from it, any height- $(i+1)$ cluster contains at most $K := 2^k$ height- i children clusters.*

FACT 2.3. (PADDED HIERARCHICAL DECOMPOSITION FOR DOUBLING METRICS [30]) *Suppose a metric has doubling dimension at most k . Then, it admits an $O(k)$ -padded hierarchical decomposition. Moreover, for any hierarchical partition sampled from it, any height- $(i+1)$ cluster contains at most $K := 2^{O(k)}$ height- i children clusters.*

2.1 Arora's and Talwar's Approximation Schemes for TSP

We give a very brief review of the hierarchical decomposition method used by Arora [3] (for low-dimensional Euclidean metrics) and Talwar [30] (for doubling metrics) to design approximation schemes for TSP. A more complete description is given in Appendix A.

1. A hierarchical partition $\{\mathcal{P}_i\}_i$ is sampled as in Definition 2.4. Each cluster C in each level contains a set $U(C)$ of points called *portals*, which can, for instance, be a fine enough net of C . The search space is restricted to *portal respecting tours*, i.e., those that enter or leave a cluster only through its portals. Given positive integers m and r , a TSP tour is (m, r) -light with respect to some hierarchical partitioning and portaling scheme if every cluster in every height of the partition contains at most m portals and the tour enters and leaves each cluster only through its portals for at most r times. It

is shown in [3, 30] that for appropriate values of m and r , with constant probability, there is some (m, r) -light tour that has length at most $(1 + \varepsilon)$ times the optimal length.

2. A dynamic program is used to find the shortest (m, r) -light tour with respect to some hierarchical partition and portaling scheme. Each cluster C has entries, each of which is indexed by a configuration consisting of a collection I of entry/exit portal pairs. Since only tours that enter and exit each cluster at most r times are considered, each such I contains at most r entry/exit pairs of portals. Moreover, an entry stores the minimum length of the internal segments consistent with its configuration. The running time of the dynamic program depends on the values m and r , as well as the maximum number K of children clusters that a parent cluster can have.

3 Augmenting the Hierarchical Decomposition Method for TSPN

The main difficulty in applying the hierarchical decomposition method (or other similar divide and conquer method) is that when a sub-problem contains partial regions, the corresponding dynamic program would possibly need to try all combinations of whether the sub-problem is responsible for those intersecting partial regions. This can potentially increase the number of dynamic program entries by a factor of $2^{\Omega(n)}$. We prove a structure theorem that can reduce the number of intersecting regions that a cluster needs to explicitly consider. In particular, if a region P is first divided up at a certain height in the hierarchical partition, then it is only necessary for descendant clusters down to certain height to explicitly consider the region P . These descendant clusters each has a *potential site*, which when *activated*, can be the point responsible for the divided region. Descendant clusters further down need not be concerned about that the divided region P any more. We first look at what exactly happens when a region is divided up in the hierarchical decomposition method.

Extra Cost due to Divided Regions. Suppose a region P with diameter δ is first divided in the hierarchical partition at diameter scale D_i . By the property of β -padded decomposition, this happens with probability at most $\beta \cdot \frac{\delta}{D_i}$. (Recall that for metrics with doubling dimension at most k , $\beta = O(k)$.) We do not know exactly the point $p \in P$ that the optimal tour visits. However, suppose we can *somehow* ensure that the tour visits a point u (not necessarily in P) instead of p that satisfies $d(p, u) \leq \gamma D_i$ (for some small $\gamma < 1$); and then the tour makes a further detour at u and visits

a point q in P such that $d(u, q) \leq \gamma D_i$. Observing that there are at most L values of i , the expected extra cost incurred is at most $\sum_i \beta \cdot \frac{\delta}{D_i} \cdot 4\gamma D_i = 4L\beta\gamma\delta$. (It would be soon apparent why we perform such a convoluted detour.) This intuition suggests that it is useful to obtain a lower bound on OPT in terms of the diameters of the regions.

The following lemma is an extension of the packing lemmas in [14, 27] to doubling metrics.

LEMMA 3.1. (EXISTENCE OF A PACKING AMONG FAT WEAKLY DISJOINT REGIONS) *Suppose W_l is a set of α -fat weakly disjoint regions of the same type, all with core radius ρ . Let Q be a set of points that intersect every region in W_l . Suppose that the underlying metric has doubling dimension at most k and $|W_l| > (8\alpha)^k$. Then, there exists $T := \lceil \frac{|W_l|}{(8\alpha)^k} \rceil$ points in Q that form an $\alpha\rho$ -packing.*

Proof. Let Q be the set of points that intersects every region in W_l , i.e., for each $P \in W_l$, the intersection $Q \cap P$ is non-empty. Let $T := \lceil \frac{|W_l|}{(8\alpha)^k} \rceil \geq 2$. It suffices to show, by induction on t , that for $1 \leq t \leq T$, there exists a set of points $Q_t := \{p_j \mid 1 \leq j \leq t\}$ in Q such that any two points in Q_t are at distance more than $\alpha\rho$ from each other.

For $t = 1$, pick any $p_1 \in Q$ and set $Q_1 := \{p_1\}$. Then, the result is trivially true. Suppose for some $1 \leq t < T$, there exists $Q_t := \{p_\lambda\}_{\lambda=1}^t$ in Q such that any two points in Q_t are at least $\alpha\rho$ apart.

Let $Z := \{z_j \mid P_j(z_j, \alpha\rho_j) \in W_l\}$ be the set of centers of regions in W_l , which all have core radii ρ . From Definition 2.1, the set Z is a ρ -packing.

For each $1 \leq \lambda \leq t$, let $Z_\lambda := \{z \in Z \mid d(z, p_\lambda) \leq 2\alpha\rho\}$. Observe that since the doubling dimension of the underlying metric is at most k , by Fact 2.1, we have $|Z_\lambda| \leq (8\alpha)^k$. It follows that $|\cup_{\lambda=1}^t Z_\lambda| \leq t \cdot (8\alpha)^k < |Z| = |W_l|$. Hence, there exists some center $z \in Z$ that is not in any of the existing Z_λ 's. Suppose p_{t+1} is a point in the region centering at z that the tour visits, and hence $p_{t+1} \in B(z, \alpha\rho)$. Now, for each $1 \leq \lambda \leq t$, by the triangle inequality, $d(p_\lambda, p_{t+1}) \geq d(p_\lambda, z) - d(z, p_{t+1}) > \alpha\rho$, since $d(p_\lambda, z) > 2\alpha\rho$ and $d(z, p_{t+1}) \leq \alpha\rho$. Setting $Q_{t+1} := Q_t \cup \{p_{t+1}\}$ completes the inductive step. \square

By taking Q to be the set of points in the set W_l of regions that a TSP tour visits, we have the following corollary.

COROLLARY 3.1. (LOWER BOUND ON OPT VIA DIAMETERS OF REGIONS) *The length of any TSP tour visiting all regions of the same type in W_l as in Lemma 3.1 is at least $\frac{|W_l|}{(8\alpha)^k} \cdot \alpha\rho$; moreover, we have $\sum_{P \in W_l} \text{Diam}(P) \leq 2(8\alpha)^k \text{OPT}$.*

Distinguishing between Common and Rare Types of Core Radii. Recall that the set W of regions are grouped into sets $\{W_l\}_{l \in [\Delta]}$, where the regions in each group have their diameters within a factor of 2 from one another. Let $W_c := \cup_{l: |W_l| > (8\alpha)^k} W_l$ be the regions with common types of radii, and $W_r := W \setminus W_c$ be those with rare types of radii. By Corollary 3.1, $\sum_{P \in W_c} \text{Diam}(P) \leq 2\Delta \cdot (8\alpha)^k \text{OPT}$, and observe that $|W_r| \leq \Delta \cdot (8\alpha)^k$.

LEMMA 3.2. (APPROXIMATE POINT LOCATION FOR DIVIDED REGIONS) *Suppose a hierarchical partition is sampled as in Fact 2.3, and a region P has diameter δ . Consider the following operation of modifying a given TSP tour.*

1. *Suppose that p is the point in P for which the TSP tour visits. Consider the height- i partition (with diameter scale D_i) for which the region P is first divided.*
2. *Let $0 < \gamma < 1$ and suppose that u is an arbitrary point (not necessarily in the region P) such that $d(u, p) \leq \gamma D_i$.*
3. *In the given tour, replace p with u .*
4. *Suppose $q \in P$ is a point such that $d(q, u) \leq \gamma D_i$. (The points p and q could be the same.) Then, make a detour at point u : visit point q and then back again at u .*

Then, the expected increase in the length of the tour is at most $4L\beta\gamma\delta$.

Proof. First, observe that the probability that a region P with diameter δ is first divided at the height- i partition is at most $\beta \cdot \frac{\delta}{D_i}$, by the property of β -padded decomposition. Note that the increase in length after the modification procedure is at most $4\gamma D_i$. Finally, observing that there can be L possible values of i for which this can happen, the expected increase in the tour length is as required. \square

Combining Corollary 3.1 and Lemma 3.2, we have the following structure theorem for TSPN.

THEOREM 3.1. (STRUCTURE LEMMA FOR TSPN) *Consider a TSPN instance on an underlying metric with doubling dimension at most k , and suppose that a hierarchical partition is sampled as in Fact 2.3. Moreover, for each region P in W_c , the approximate point location modification is performed as in Lemma 3.2 on any given tour. Then, the expected increase in length is at most $8(8\alpha)^k L\beta\gamma\Delta \text{OPT}$, which is at most $\frac{\varepsilon}{2} \cdot \text{OPT}$, if we set $\gamma = O(\frac{\varepsilon}{(8\alpha)^k L\beta\Delta})$. In particular, if the given tour is an (m, r) -light TSPN tour, whose length has an expected difference from the optimal length of at most $\frac{\varepsilon}{2} \cdot \text{OPT}$, then the resulting TSPN tour has an expected increase from the optimal length of at most $\varepsilon \cdot \text{OPT}$.*

We next give the details of the approximate point location procedure in Lemma 3.2.

Assigning Anchor Points for a Divided Region in W_c . We describe how the point u is picked for a region P (that is first divided at height- i), as in Step 2 of Lemma 3.2. Observe that P is totally contained in some height- $(i-1)$ cluster C_{i-1} . For the special case when the diameter of P is at most γD_i , then we pick an arbitrary point $p \in P$ and replace the region P with the singleton $\{p\}$; we emphasize that in this case p is NOT an anchor point for the region P . Otherwise, consider the descendant clusters of C_{i-1} that intersect with P , in decreasing height. As soon as the diameter of an intersecting cluster C drops below γD_i , or if we have reached the lowest height where C is a height-0 cluster (which has diameter at most $\frac{\varepsilon R}{n}$, see Remark 2.3), we pick u to be any arbitrary point inside C , and we say u is an anchor point at height- i for the region P ; in this case, it is not necessary to consider further the descendant clusters of C for assigning anchor points. Note that we do not know which point in the region the optimal tour would visit, but we can ensure that the correct point would have an anchor point within a distance of γD_i .

Potential Site in a Cluster. Observe that in the above description, an anchor point u for some region is an arbitrary point in some cluster C . Hence, we pick an arbitrary point $u(C)$ in each cluster as a *potential site*, which when *activated*, can be an anchor point for regions partially intersecting C . We require that if u is a potential site for a cluster C , then it must also be one for one of its children clusters.

Ambiguous Regions for a Cluster. Recall ultimately, we want to limit the number of regions that intersect a cluster for which the dynamic program has to explicitly consider. Given a cluster C at height- i , its *ambiguous regions* are those regions P partially intersecting C that satisfy one of the following properties.

1. The region P is in W_r , i.e., its type of core radius is rare.
2. The cluster C or any of its descendant clusters contain potential sites that can be anchor points (at the corresponding heights) for the region P .

Technical Issues Involving Padded Hierarchical Decomposition. Before we can bound the number of ambiguous regions for a cluster, there are some issues concerning the padded hierarchical decomposition that need to be clarified.

1. We know that the optimal solution is contained in the ball $B(p_0, nR)$. However, if we simply take this ball as the height- L cluster, then some regions

would be divided at height- L with probability 1, thereby violating the padded-property. It suffices to pick $\eta \in [\frac{1}{2}, 1]$ uniformly at random, and take the height- L cluster to be $B(p_0, \frac{1}{2}\eta D_L)$ (assuming β is sufficiently large, say $\beta \geq 2$).

2. In the argument that follows, for $i \in [L]$, we would need the existence of β -padded decomposition at height- $(i + \Gamma)$, where $\Gamma := \lceil \log_2 \frac{1}{\gamma} \rceil$. In particular, we need the property that at height- $(i + \Gamma)$, the cluster containing p_0 has diameter at most $D_{i+\Gamma}$, and the probability that a region with diameter δ is divided at this height is at most $\beta \cdot \frac{\delta}{D_{i+\Gamma}}$. For $1 \leq s \leq \Gamma$, we can simply set the “imaginary cluster” at height- $(L + s)$ to be $B(p_0, \frac{1}{2}\eta D_{L+s})$, where $\eta \in [\frac{1}{2}, 1]$ is the same as in 1. The imaginary clusters are only for the sake of the proof and do not play any role in the actual algorithm.

LEMMA 3.3. (BOUNDING THE NUMBER OF AMBIGUOUS REGIONS) *The number of ambiguous regions for a cluster is at most $H := \Delta \cdot (8\alpha)^k + O(\frac{\alpha}{\gamma})^k$, where k is the doubling dimension.*

Proof. Suppose cluster C is at height- i , where $i \leq L$. The number of its ambiguous regions in W_r is at most $|W_r| \leq \Delta \cdot (8\alpha)^k$. We next count the number of its ambiguous regions in W_c .

We first bound the number of ambiguous regions in W_c having diameter at least D_i . If a region P has diameter larger than $\frac{D_i}{\gamma}$, then it cannot have its anchor point in cluster C . The reason is that such a region P would be first divided at a height of at least $i + \Gamma$, recall $\Gamma := \lceil \log_2 \frac{1}{\gamma} \rceil$. Hence, it follows that region P 's anchor points must be at a height larger than i . So we may assume that $D_i \leq \text{Diam}(P) \leq \frac{D_i}{\gamma}$. Observe that such a region can be of at most Γ types of core radii. The centers of the α -fat weakly disjoint regions from each such type form some packing, and by Remark 2.2 and Fact 2.1, there are at most $O(\alpha)^k$ regions from each such type. Hence, there can be totally at most $\Gamma \cdot O(\alpha)^k = O(\frac{\alpha}{\gamma})^k$ such ambiguous regions.

It remains to bound the number of ambiguous regions in W_c having diameter less than D_i . Note that if a region P has diameter less than γD_i , then there would be no anchor points for the region P (at any height), and so P cannot be ambiguous. The reason is that the region must be first divided at a height $i' \geq i$, and hence the diameter of region P is at most $\gamma D_{i'}$. In this case, an arbitrary point p in P (which is NOT an anchor point) is picked and the region P is reduced to $\{p\}$. Again, note that if the diameter of a region P is at least δ , then its core radius is at least $\frac{\delta}{2\alpha}$. By Remark 2.2 and Fact 2.1, there can be at most $O(\frac{D_i + \delta}{\delta/\alpha})^k$ regions with diameter around δ that intersect cluster C . Hence,

the total number of ambiguous regions having diameter less than D_i is dominated by the term corresponding to $\delta = \gamma D_i$, which is $O(\frac{\alpha}{\gamma})^k$.

Summing up the number of ambiguous regions in all the cases gives the required bound. \square

4 Dynamic Program for TSPN

We describe details of the augmented dynamic program for finding the shortest TSPN tour after the approximate point location modification as in Theorem 3.1, in addition to ensuring the (m, r) -lightness property in the original dynamic program. In Lemma 3.3, we bound the number H of ambiguous regions for each cluster. Hence, in the dynamic program, the number of configurations for a cluster increases by a factor of at most 2^H . Observing that there are at most $K = O(1)^k$ children clusters for any parent cluster, we show that the running time of the dynamic program increases by a factor of at most $2^{O(HK)}$.

THEOREM 4.1. *Suppose that we are given an instance of TSPN, where the underlying metric space has doubling dimension at most k , and the regions are α -fat weakly disjoint with at most Δ types of radii. Then, with constant probability, the augmented hierarchical decomposition method gives a TSPN tour of length at most $(1 + \varepsilon)OPT$ in time $\text{TIME}(TSP) \cdot 2^{O(HK)} = \exp\{O(\frac{\Delta}{\varepsilon})^k O(\alpha)^{k^2} \log^k n\}$, where $\text{TIME}(TSP)$ is the time for approximating TSP with the hierarchical decomposition method used by Arora [3] or Talwar [30].*

Proof. In view of Theorem 3.1, we prove the theorem by giving the construction of the augmented dynamic program for approximating TSPN.

Outline of the TSPN Algorithm. A hierarchical partition is first sampled. Then, the portals for each cluster are assigned as in [3] or [30]. As described in Section 3, the potential sites are chosen, the anchor points for the regions are assigned, and the ambiguous regions for each cluster are determined. Then, the following dynamic program finds a desirable TSPN tour.

Configuration of a Cluster. The configuration of a cluster C includes the following.

1. A collection I of portal entry/exit points as before. (Recall that (m, r) -lightness implies that $|I| \leq r$.)
2. A bit vector of length equal to the number of ambiguous regions that cluster C has. Each such bit indicates whether the cluster is responsible for the corresponding ambiguous region.
3. A bit indicating whether the potential site $u(C)$ is activated.

Since a cluster has at most H ambiguous regions, the number of configurations for the cluster increases by a multiplicative factor of at most 2^{H+1} . Each entry also stores the minimum length of the segments consistent with the configuration, and also the segments themselves if a tour needs to be constructed eventually. We focus mainly on the new features of the dynamic program.

Base Case: Filling Entries of Height-0 Clusters.

Suppose C is a height-0 cluster. We describe the possible configurations for such a cluster and the corresponding partial lengths stored under its entries.

1. If the cluster C contains singleton regions P , then $u(C)$ must be activated in the configuration to be responsible for all those P . Otherwise, the cluster C has both configurations in which $u(C)$ is activated and also those in which $u(C)$ is not. For those configurations in which $u(C)$ is activated, the collection I of entry/exit points contains only a singleton $\{x\}$, where x is the closest portal to $u(C)$ in C ; ⁷ moreover, twice the distance from x to $u(C)$ needs to be included to the length of the partial tour within C . If $u(C)$ is not activated in a configuration, then the set I is empty; and such a configuration needs not take part any further in the dynamic program.
2. If $u(C)$ is activated in a configuration, there is a choice for each ambiguous region P of C of whether C is responsible for P .
3. For each region P for which the cluster C is responsible (either by necessity as in 1 or by choice as in 2), add twice the distance between $u(C)$ and its closest point in P to the length of partial tour for the entry under the corresponding configuration.

Inductive Step: Determining the Configuration and Combining the Solutions from the Children Clusters. The entries of a height- i cluster are computed from those of its children clusters. We consider each combination of configurations of the children clusters. A configuration for C is formed by interleaving different portal entry/exit pairs as indicated by the children clusters' configurations. (Recall that an interleaving is valid, only if it results in a collection I of entry/exit points that has size at most r , since we look for (m, r) -light tours.) We now concentrate on how the new parameters involved in the configuration operate. Here are the steps to be performed to determine the configuration of the cluster C , and compute the corresponding entries.

⁷Note that we can always choose a potential site $u(C)$ that is also a portal, in which case x would also be $u(C)$.

1. For each region P that is totally contained in C , but is divided by the children clusters, the configurations of the children clusters should reflect that at least one of them is responsible for the region P . Otherwise, such a combination of the children clusters' configurations is invalid, and we move on to the next combination.
2. The potential site $u = u(C)$ is activated in the configuration of C , if and only if the site u is also activated in the corresponding child cluster that contains u .
3. Suppose the region P is ambiguous for cluster C . If possible, determine whether C is responsible for P from the configurations of the children clusters.
4. If the configurations of the children clusters give no information of whether C is responsible for an ambiguous region P , then it must be the case that the cluster C has a potential site $u(C)$ that can be an anchor point at height- i for the region P . If the potential site $u(C)$ is not activated, then the cluster C is not responsible for the ambiguous region P in the configuration of C . If the potential site $u(C)$ is activated, then there is a choice of whether C is responsible for region P . In the case where C is responsible for the region P , twice the distance of $u(C)$ to the closest point in $P \cap C$ is added to the length stored in the corresponding entry of C . Note that considering choices for each ambiguous region can increase the running time by a factor of at most 2^H .
5. For each configuration of C formed in the manner described above, we update the corresponding entry for C if the new partial length is less than that of the existing one.

Increase in the Running Time. We analyze the running time of the augmented dynamic program. Note that the number of configurations for a cluster increases by a factor of at most 2^{H+1} , and the number of children cluster is at most K . Hence the total number of combinations of the children clusters' configurations increases by a factor of at most $2^{K(H+1)}$. For each such combination, the time to combine them increases by a factor of at most 2^H , as described in Step 4 above. Hence, it follows the total time of the augmented dynamic program increases by a factor of at most $2^{O(HK)}$. Observing that $H = \Delta \cdot (8\alpha)^k + O(\frac{\alpha}{\gamma})^k$, $K = O(1)^k$, $\gamma = O(\frac{\epsilon}{(8\alpha)^k L \beta \Delta})$, $L = O(\log \frac{n}{\epsilon})$ and $\beta = O(k)$ gives the required running time. \square

5 Special Case: Euclidean Metrics

We consider the special case when the underlying metric is the k -dimensional Euclidean metric, with disjoint

regions. If we use a stronger notion of fatness, then we can remove the running time's dependence on the number Δ of types of region diameters.

DEFINITION 5.1. (α -FATNESS [9, 31]) *A region $P \subseteq \mathbb{R}^k$ is said to be α -fat if for any k -dimensional ball B which does not fully contain P and whose center lies in P , the volume of the intersection of P and B is at least $1/\alpha^k$ times the volume of B .*

Restricting the Tour in B_0 . Recall that for finite metrics, we can try each point p_0 in the region P_0 to form some bounding ball $B(p_0, nR)$. We need to use a different approach if every region is continuous and contains an infinite number of points. Using the method outlined in [9], one can approximate a minimum box intersecting all regions with center c and radius R . Note that we can assume that there is some region P_0 with diameter at most nR . Otherwise, if all regions have diameters at least nR , then by the definition of α -fatness and disjointness, there can be at most $O(\alpha)^k$ regions; and hence there are only a constant number of regions and the problem becomes trivial. Note that for any point p_0 in P_0 , the ball $B(p_0, (n+1)R)$ intersects all regions. Hence, instead of trying each p_0 in P_0 , we can just conclude that the optimal tour must be inside the ball $B(c, 2n^2R)$. We still have only $L = O(\log \frac{n}{\epsilon})$ length scales to consider.

Bounding the Number Δ of Types of Region Diameters. Note that we have shown that the optimal tour must be inside the ball $B(c, 2n^2R)$, and as before we can assume that a region is either a singleton or has diameter at least $\frac{\epsilon R}{2n}$. For regions having diameter between $\frac{\epsilon R}{2n}$ and $2n^2R$, we can group them between powers of 2, so that there are at most $O(\log \frac{n}{\epsilon})$ types. Note that by α -fatness and disjointness, there can be at most $O(\alpha)^k$ regions having diameter larger than $2n^2R$. Hence, it follows that the number of types Δ of region diameters is at most $O(\log \frac{n}{\epsilon}) + O(\alpha)^k$.

Hence, Theorem 1.2 follows immediately from Theorem 1.1.

References

- [1] E. M. Arkin and R. Hassin. Approximation algorithms for the geometric covering salesman problem. *Discrete Applied Mathematics*, 55(3):197–218, 1994.
- [2] S. Arnborg and A. Proskurowski. Linear time algorithms for NP-hard problems restricted to partial k -trees. *Discrete Appl. Math.*, 23(1):11–24, 1989.
- [3] S. Arora. Approximation algorithms for geometric TSP. In *The traveling salesman problem and its variations*, volume 12 of *Comb. Optim.*, pages 207–221. Kluwer Acad. Publ., Dordrecht, 2002.
- [4] P. Assouad. Plongements lipschitziens dans \mathbf{R}^n . *Bull. Soc. Math. France*, 111(4):429–448, 1983.
- [5] Y. Bartal. Probabilistic approximations of metric spaces and its algorithmic applications. In *37th FOCS*, pages 184–193, 1996.
- [6] H. T-H. Chan, A. Gupta, B. M. Maggs, and S. Zhou. On hierarchical routing in DOubling metrics. In *16th SODA*, pages 762–771, 2005.
- [7] K. L. Clarkson. Nearest neighbor queries in metric spaces. *Discrete Comput. Geom.*, 22(1):63–93, 1999.
- [8] A. E. F. Clementi, P. Crescenzi, and G. Rossi. On the complexity of approximating colored-graph problems. In *Proc. 5th Internat. Computing and Combinatorics Conference*, volume 1627 of *Lecture Notes in Computer Science*, pages 281–290. Springer, 1999.
- [9] M. de Berg, J. Gudmundsson, M.J. Katz, C. Levcopoulos, M.H. Overmars, and A. F. van der Stappen. TSP with Neighborhoods of varying size. *J. of Algorithms*, 57:22–36, 2005.
- [10] M. M. Deza and M. Laurent. *Geometry of cuts and metrics*, volume 15 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1997.
- [11] M. Dror and J. B. Orlin. Combinatorial optimization with explicit delineation of the ground set by a collection of subsets. In *SIAM J. Discrete Math.* 21(4), pages 1019–1034, 2008.
- [12] A. Dumitrescu and J. S. B. Mitchell. Approximation algorithms for TSP with neighborhoods in the plane. *J. Algorithms*, 48(1):135–159, 2003.
- [13] K. Elbassioni, A. V. Fishkin, and R. Sitters. Approximation algorithms for the Euclidean traveling salesman problem with discrete and continuous neighborhoods. *Int. J. Comput. Geometry Appl.*, 19(2):173–193, 2009.
- [14] K. Elbassioni, A.V. Fishkin, N. Mustafa, and R. Sitters. Approximation algorithms for Euclidean group TSP. In *ICALP*, volume 3580 of *Lecture Notes in Computer Science*, pages 1115–1126. Springer, 2005.
- [15] J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. System Sci.*, 69(3):485–497, 2004.
- [16] C. Feremans and A. Grigoriev. Approximation schemes for the generalized geometric problems with geographic clustering. In *Proceedings of the 21st European Workshop on Computational Geometry EWCG*, pages 101–102, 2005.
- [17] N. Garg, G. Konjevod, and R. Ravi. A polylogarithmic approximation algorithm for the group Steiner tree problem. *J. Algorithms*, 37(1):66–84, 2000. (Preliminary version in *9th SODA*, pages 253–259, 1998).
- [18] J. Gudmundsson and C. Levcopoulos. A fast approximation algorithm for TSP with neighborhoods. *Nordic J. Computing*, 6(4):469–488, 1999.
- [19] A. Gupta, R. Krauthgamer, and J. R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *44th FOCS*, pages 534–543, 2003.
- [20] E. Halperin and R. Krauthgamer. Polylogarithmic inapproximability. In *Proceedings of the thirty-fifth ACM symposium on Theory of computing*, pages 585–

594. ACM Press, 2003.
- [21] S. Har-Peled and M. Mendel. Fast constructions of nets in low dimensional metrics, and their applications. In *21st SOCG*, pages 150–158, 2005.
- [22] R. Krauthgamer and J. R. Lee. Navigating nets: simple algorithms for proximity search. In *15th SODA*, pages 798–807, 2004.
- [23] R. Krauthgamer and J. R. Lee. The black-box complexity of nearest-neighbor search. *Theoret. Comput. Sci.*, 348(2-3):262–276, 2005.
- [24] C. S. Mata and J. S. B. Mitchell. Approximation algorithms for geometric tour and network design problems (extended abstract). In *Proc. 11th Annual ACM Symposium on Computational Geometry*, pages 360–369, 1995.
- [25] J. Matoušek. *Lectures on discrete geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002.
- [26] J. S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: a simple polynomial-time approximation scheme for geometric TSP, k -MST, and related problems. *SIAM J. Comput.*, 28(4):1298–1309, 1999.
- [27] J. S. B. Mitchell. A PTAS for TSP with neighborhoods among fat regions in the plane. In *SODA*, pages 11–18, 2007.
- [28] S. B. Rao and W. D. Smith. Approximating geometrical graphs via “spanners” and “banyans”. In *STOC '98 (Dallas, TX)*, pages 540–550. ACM, New York, 1999.
- [29] S. Safra and O. Schwartz. On the complexity of approximating TSP with Neighborhoods and related problems. In *Proc. 11th Annual European Symposium on Algorithms*, volume 2832 of *Lecture Notes in Computer Science*, pages 446–458. Springer, 2003.
- [30] K. Talwar. Bypassing the embedding: Algorithms for low-dimensional metrics. In *36th STOC*, pages 281–290, 2004.
- [31] A. F. van der Stappen. Motion planning amidst fat obstacles. PhD Dissertation, 1994, Department of Computer Science, Utrecht University, Utrecht, Netherlands.

Appendix A: Review: Approximating TSP via the Hierarchical Decomposition and Portaling Method on Bounded Growth Metrics

Our techniques are based on the approximation schemes by Arora [3] (for low-dimensional Euclidean metrics) and Talwar [30] (for doubling metrics), to which we refer as the *Hierarchical Decomposition Method*. We give a brief review of the construction and highlight the relevant properties that are crucial to our augmented scheme. For the moment, consider the case where each region P_j contains only one point.

On a high level, the method divides the metric space hierarchically into smaller clusters. The partial

solutions for smaller clusters are solved and combined together to form the global solution through dynamic programming. We next give the main ingredients of the method.

(1) Padded Hierarchical Decomposition & Portaling Scheme

Portal Assignment and (m, r) -Light Tours. Suppose a hierarchical partition $\{\mathcal{P}_i\}$ is sampled as in Definition 2.4. For each $0 \leq i < L$, each height- i cluster C has a set $U(C)$ of points called *portals*. We consider *portal respecting* tours, i.e., those that enter or leave a cluster only through its portals. This would limit the size of the search space for TSP tours. However, to ensure that a tour of good quality is still possible, the set $U(C)$ is chosen to be a fine enough θD_i -net of C , where $\theta = O(\frac{\varepsilon}{\beta^L})$ is suitably small. Given positive integers m and r , a TSP tour is (m, r) -light with respect to some hierarchical partitioning and portaling scheme if every cluster in every height of the partition contains at most m portals and the tour enters and leaves each cluster only through its portals for at most r times.

THEOREM A.1. (STRUCTURE THEOREM FOR TSP) *Given an instance of TSP in some underlying metric M , there exists a padded hierarchical decomposition scheme such that with probability at least $\frac{1}{2}$, the hierarchical partition sampled from it admits an (m, r) -light TSP tour of length at most $(1 + \varepsilon)OPT$, where*

- (a) *if M has doubling dimension at most k , then $m = (\frac{kL}{\varepsilon})^{O(k)}$ and $r = (\frac{kL}{\varepsilon})^k$; [30]*
 (b) *if M is in k -dimensional Euclidean space, then $m = 2k \cdot (O(\frac{\sqrt{k}L}{\varepsilon}))^{k-1}$ and $r = 2k \cdot (O(\frac{\sqrt{k}}{\varepsilon}))^{k-1}$. [3]⁸*

(2) Dynamic Programming for Finding (m, r) -Light Tours

We outline a dynamic program to find the shortest (m, r) -light tour with respect to some hierarchical partition and portaling scheme. Similar constructions are used by Arnbourg and Proskurowski [2], Arora [3] and Talwar [30], and our construction for TSPN is built upon this construction.

Configuration of a Cluster. For each cluster C with its portals $U(C)$, there are entries, each of which is indexed by a configuration that consists of a collection I of pairwise disjoint subsets of $U(C)$ of size 1 or 2.

An entry for cluster C indexed by I represents the scenario in which a tour visits cluster C via portals

⁸The original definition of (m, r) -lightness in [3] counts crossings on only one facet of the bounding box, and hence there is an extra factor of $2k$ here.

described by subsets in I . A 2-subset $\{u, v\}$ in I means there is a portion of the tour that enters and exits via portals u and v . A 1-subset $\{x\}$ in I means the tour enters and leaves cluster C through the portal x . We keep track of the length of the portion of the tour that is within the cluster C . The entry indexed by I stores the length of the shortest possible internal segments, for tours consistent with the scenario imposed by I . Note that if we have to construct the tour, under each entry we have to store the internal segments of the tour as well. Note that for a tour to be (m, r) -light, we must have $|U(C)| \leq m$ and $|I| \leq r$. Hence, each cluster has at most m^{2r} entries.

Time for Dynamic Program. As mentioned in [30], if each parent cluster C has at most K children clusters, then the time to fill up all the entries of C is at most the product of the number of configurations for all the children clusters and $(Kr)!$. This product is at most $(mKr)^{2Kr}$. Since there are at most n clusters from each of the L levels, the total time is $nL(mKr)^{2Kr}$.

Appendix B: APX-Hardness for Unbounded Types of Region Diameters

We next motivate why we need to make the assumption that the number of types of diameters for the regions is bounded. It was shown in [13] that the TSPN problem is APX-hard for the case where all objects are line-segments in the plane of almost equal length. We can modify this reduction to get the following result.

THEOREM B.1. *The TSPN problem for doubling metrics with regions being disjoint balls of arbitrary types of radii is APX-hard.*

Proof. We reduce the TSPN problem from VERTEX-COVER for 3-partite graphs, which cannot be approximated within a factor $34/33$, unless $P = NP$ [8].

For completeness, we first describe the construction in [13] again. Given a 3-partite graph G on n vertices, we define an instance of TSPN as shown in Figure 1. The vertices of the graph correspond to points on the plane, and the edges correspond to neighborhoods (of size 2) of the TSPN instance in the obvious way: two points form a neighborhood if and only if the corresponding vertices in the graph are adjacent. Furthermore, we define a large number of singleton neighborhoods which together form a polygon with perimeter L . The small equilateral triangle in the closeup has side-length d . If d is small enough, then an optimal tour follows the polygon and jumps up and down to some of the vertices. The extra cost of the detour for each

such vertex is $2d - d = d$. Consider an optimal tour and let S be the set of vertices of G that are visited, then $\text{OPT} = L + |S|d$. Now we let $d = 1/n$ and choose the distance between any two vertices substantially larger, say $4/n$. We let the perimeter of the polygon be 10. If there is a vertex cover of size $n/2$, then there exists a tour of length $L + nd/2 = 10.5$. On the other hand, if there exists a tour of length at most $10 + \beta$, then there must be a vertex cover of size at most βn . Taking $\beta = 34/66$ shows that TSPN cannot be approximated within a factor $(10 + 34/66)/10.5 \approx 1.0014$.

We modify the above construction in the following way.

1. **Disjoint Neighborhoods.** Observe that in the above construction, the neighborhoods are not disjoint. In particular, if a vertex v has degree d in the given graph, then the corresponding point u would be contained in d neighborhoods. For such a vertex, we have d copies $\{u_1, u_2, \dots, u_d\}$ of the point u . Each neighborhood can now take their unique copy of the point u . These locations of these d points are just tiny perturbations from the original location of the point u . This perturbation is so tiny that the following is true. Given any tour, the tour can be modified such that if one copy u_i is visited, then every copy would be visited, with the increase in tour length being an arbitrarily small fraction of the optimal length. Note that we now have disjoint neighborhoods in the Euclidean plane and we still preserve the APX-hardness of the reduction. We denote the points in the metric space we have constructed so far by X .
2. **Neighborhoods as Disjoint Balls.** Suppose W is the set of size-2 neighborhoods in the construction. We are going to augment the metric and the neighborhoods so that each neighborhood in W is contained in some ball in the new metric. Observe that X are points in the Euclidean plane, and can be represented by a weighted complete graph G_X , where the length of each edge is the Euclidean distance between the corresponding points. Since the Euclidean plane has constant doubling dimension, it follows that the metric induced by G_X also has bounded doubling dimension. We augment the metric and the neighborhoods in the following way. Suppose there are $w = |W|$ neighborhoods of size 2. Let $\Lambda > 4$ be a large enough parameter, for instance, Λ is at least 100 times the optimal length. For each $1 \leq i \leq w$, for the neighborhood $P_i = \{x_i, y_i\}$ in W , we create a new point z_i in the graph G_X and add edges $\{z_i, x_i\}$ and $\{z_i, y_i\}$ to the graph with length Λ^i , and we also




Figure 1: The reduction for a 3-partite graph on 12 vertices.

define a new neighborhood $P'_i := \{x_i, y_i, z_i\}$. We let the augmented graph be G_Z with the augmented set of points Z in the metric induced by G_Z . Observe that the new set of neighborhoods are in the form $B_Z(z_i, \Lambda^i)$. More importantly, they are now disjoint balls.

Notice that we do not need to consider any tour that visits any z_i , because such a tour would have length at least 100 times that of the optimal tour. Hence, the APX-hardness reduction is preserved. It remains to see if the augmented metric still has constant doubling dimension. Now, observe that the augmented points z_i are at geometrically increasing distances from the original graph G_X (which itself induces a metric with bounded doubling dimension), and hence the metric induced by the augmented graph G_Z also has constant doubling dimension.

It follows that we have constructed an instance of the TSPN in a metric with constant doubling dimension, whose regions are either singletons or disjoint balls of the form $B_Z(z_i, \Lambda^i) = \{x_i, y_i, z_i\}$, as required. \square