A Bicriteria Approximation Scheme for the Minimum Well-Separated Pair Decomposition

Guilherme D. da Fonseca
Aix-Marseille Université and LIS, France

Abstract

A well-separated pair decomposition (WSPD) of a set P of n points covers all $\binom{n}{2}$ pairs of points by a family of *dumbbells* $\{X, Y\}$, each covering the pairs of $X \times Y$, such that the diameters of X and Y are at most an ε-fraction of their distance. In $\mathbb{R}^d$, WSPDs with $O(n/\varepsilon^d)$ dumbbells always exist, but many inputs admit much smaller ones. However, computing a WSPD of minimum size is NP-hard in the plane. For $d \geq 2$, the only known approximation that preserves the separation has a factor of $2^{O(d)}$, and Buchin et al. (ESA 2026) asked whether a $(1+\delta)$-approximation exists.

We give a bicriteria approximation scheme: for fixed d and $\delta \in (0, 1]$, we compute in $f(d, \delta) \cdot n^6$ time a $(1+\delta)\varepsilon$-WSPD whose size is at most $(1+\delta)$ times the minimum size of an ε-WSPD. The same holds in metric spaces of bounded doubling dimension. The algorithm combines shifting over scales with the local-growth technique of Nieberg, Hurink, and Kern, applied to a graph whose vertices are the pairs of points. In contrast, we show that in general metric spaces, relaxing the separation by a factor below 2 keeps the problem hard: for every fixed rational $\varepsilon \leq 1/4$, it is NP-hard to compute such a relaxed WSPD whose size is within a factor n^c of the minimum, for every $c < 2/3$.

2012 ACM Subject Classification Theory of computation → Computational geometry; Theory of computation → Approximation algorithms analysis

Keywords and phrases Well-separated pair decomposition, approximation scheme, set cover, local growth, doubling metrics

Acknowledgements Use of AI: This paper was prepared with substantial help from the AI assistant Claude (Anthropic). It took part in developing the algorithms and proofs, checking the references, and writing the text. The author checked all results and is responsible for the content.

1 Introduction

Let P be a set of n points in a metric space (M, dist) and let $\varepsilon > 0$. For $X \subseteq P$ let $\mathrm{diam}(X)$ be its diameter, and for $X, Y \subseteq P$ let $\mathrm{dist}(X, Y) = \min_{x \in X, y \in Y} \mathrm{dist}(x, y)$. A *dumbbell* is a set $\{X, Y\}$ of two disjoint nonempty subsets of P, called its *sides*. It is *ε-separated* if

$$\max\{\mathrm{diam}(X), \mathrm{diam}(Y)\} \leq \varepsilon \cdot \mathrm{dist}(X, Y).$$

Throughout the paper, a *pair* is an unordered pair $\{p, q\}$ of distinct points of P, and the dumbbell $\{X, Y\}$ *covers* the pairs $\{x, y\}$ with $x \in X$ and $y \in Y$. A family $\mathcal{W}$ of ε-separated dumbbells is an *ε-WSPD* of P if it covers all $\binom{n}{2}$ pairs of P. Its *size* is $|\mathcal{W}|$. It is a *partition* WSPD if every pair is covered exactly once. Unless stated otherwise, WSPDs in this paper are *cover* WSPDs, that is, pairs may be covered several times.

Since their introduction by Callahan and Kosaraju [5], WSPDs have become a basic tool for proximity problems, including spanners, approximate minimum spanning trees, nearest neighbors, and distance oracles; see [24, 29]. In $\mathbb{R}^d$, a partition ε-WSPD of size $O(n/\varepsilon^d)$ can be computed in $O(n \log n + n/\varepsilon^d)$ time [5], and similar bounds hold in doubling metrics [17]. On many inputs, however, much smaller WSPDs exist. For instance, the exponentially spread points of Section 2.1 admit a WSPD of size $n - 1$, which is optimal by Lemma 3. Since the size of a WSPD determines the size of the resulting spanners and the cost of many downstream algorithms, it is natural to ask for a smallest one. In MINWSPD, we are given

Table 1 Approximation results for MINWSPD. In every row, the size is compared with OPT_ε, the minimum size of a cover ε-WSPD. The row of [19] requires $\varepsilon < 1/33$. The last row is a hardness result for every fixed rational $\varepsilon \le 1/4$, which holds for every $\tau < 2$ and every $c < 2/3$, unless P = NP.

Setting	Problem	Size	Separation	Time	Reference
$\mathbb{R}^1$	cover	$(1+\delta)\,\mathrm{OPT}_\varepsilon$	ε	$n^{O(1/\delta^2)}$	[3]
$\mathbb{R}^1$	cover	$3\,\mathrm{OPT}_\varepsilon$	ε	$O(n \log n)$	[3]
$\mathbb{R}^1$	both	$10\,\mathrm{OPT}_\varepsilon$	ε	polynomial	[3]
Doubling dim. κ	both	$2^{O(\kappa)}\mathrm{OPT}_\varepsilon$	ε	$2^{O(\kappa)}(n \log n + \mathrm{OPT})$ exp.	[3, 17]
Any metric	cover	$7\,\mathrm{OPT}_\varepsilon$	33ε	polynomial	[19]
$\mathbb{R}^d$, fixed d	cover	$(1+\delta)\,\mathrm{OPT}_\varepsilon$	$(1+\delta)\varepsilon$	$f(d,\delta)\,n^6$	Theorem 4
Doubling dim. κ	cover	$(1+\delta)\,\mathrm{OPT}_\varepsilon$	$(1+\delta)\varepsilon$	$f(\kappa,\delta)\,n^6$	Theorem 11
Any metric	cover	no $n^c\mathrm{OPT}_\varepsilon$	$\tau\varepsilon$, $\tau < 2$	—	Theorem 17

P and $\varepsilon \in (0,1)$ and we must compute an ε-WSPD of minimum size. We write OPT_ε, or OPT when ε is clear, for this minimum size of a cover WSPD.

Previous work. Buchin et al. [3] initiated the study of MINWSPD. They showed that the problem is NP-hard in the plane when ε depends on n, and that in general metric spaces, already for the fixed value $\varepsilon = 1/3$, it cannot be approximated within a factor n^c for some constant $c > 0$ unless P = NP, by padding a reduction from clique cover with many copies of the input graph at a large distance. On the real line, they gave a PTAS based on local search for pseudodisks [23], an $O(n \log n)$-time 3-approximation based on the local-ratio technique [2], and a 10-approximation for the partition version. They also showed that the net-tree WSPD of Har-Peled and Mendel [17] is a $2^{O(\kappa)}$-approximation in metric spaces of doubling dimension κ, computable in $2^{O(\kappa)}(n \log n + \mathrm{OPT})$ expected time. Since the net-tree WSPD is a partition WSPD and OPT_ε is at most the minimum size of a partition ε-WSPD, this is also a $2^{O(d)}$-approximation for the partition version in $\mathbb{R}^d$. They asked whether a $(1+\delta)$-approximation, or a simple $O(1)$-approximation, exists in $\mathbb{R}^d$ for $d \ge 2$.

Har-Peled, Raichel, and Robson [19, Thm. 2.14] construct, in polynomial time and in every finite metric space, an ε-WSPD whose size is at most $7\,\mathrm{OPT}_{\varepsilon/33}$ for every $\varepsilon \in (0,1)$. Their theorem is stated with unspecified constants, and the values 7 and 33 come from its proof. This bicriteria guarantee loses constant factors in both the size and the separation.

Our results. Table 1 summarizes the known and new results.

1. **A bicriteria approximation scheme** (Theorem 4). For fixed d and $\delta \in (0,1]$, we compute a $(1+\delta)\varepsilon$-WSPD of $P \subset \mathbb{R}^d$ of size at most $(1+\delta)\mathrm{OPT}_\varepsilon$ in $f(d,\delta) \cdot n^6$ time, where $f(d,\delta) = 2^{2^{O(d(1/\delta+\log(d/\delta)))}}$. The running time does not depend on ε. Equivalently, we compute an ε-WSPD of size at most $(1+\delta)\mathrm{OPT}_{\varepsilon/(1+\delta)}$. In fixed-dimensional Euclidean space, this improves the bicriteria factors $(7, 33)$ of Har-Peled et al. [19] to $(1+\delta, 1+\delta)$, at the cost of a much larger running time, whereas their result holds in every metric space. The question of Buchin et al., a $(1+\delta)$-approximation that keeps the separation ε, remains open (see Remark 10). The scheme extends to metric spaces of fixed doubling dimension κ (Theorem 11). There, relaxing the separation by a factor $1+\delta$ brings the size factor from the $2^{O(\kappa)}$ of the net-tree WSPD down to $1+\delta$, with the same running time bound as in $\mathbb{R}^d$, with κ in place of d.
2. **Limits in general metrics** (Theorem 17). By a direct reduction from clique cover, for every rational ε and every $\tau < 2$ with $\tau\varepsilon < 1/2$ (in particular, every $\tau < 2$ when $\varepsilon \le 1/4$),

unless P = NP, no polynomial-time algorithm computes a $\tau\varepsilon$-WSPD of size OPT_ε in general metric spaces, and, using the hardness of approximating the chromatic number, none computes one of size at most $n^c\mathrm{OPT}_\varepsilon$ for any $c < 2/3$. Hence, unless P = NP, no bicriteria approximation scheme exists there. This extends the hardness result of Buchin et al. [3], which is for $\varepsilon = 1/3$ and exact separation and whose padding with many copies we reuse, with a self-contained proof and an explicit exponent. For $\varepsilon < 1/33$, Har-Peled et al. [19] achieve a constant size factor with the separation factor 33, so for small ε the smallest separation factor that allows a constant size factor lies between 2 and 33.

We also observe that in every metric space every ε-WSPD has at least $n - 1$ dumbbells, which is tight already on the line (Lemma 3). For partition WSPDs, this follows from the Graham–Pollak theorem [14], as noted by Buchin et al. [3]. This argument does not extend to covers, since K_n can be covered by $\lceil \log_2 n \rceil$ bicliques; for $\varepsilon = 1$, these bicliques are 1-separated dumbbells of the uniform metric, so the bound needs $\varepsilon < 1$. The proof adapts Lemma 3.1 of Callahan and Kosaraju [4], who show for partition WSPDs in $\mathbb{R}^d$ that replacing every edge of a minimum spanning tree by any pair of the same dumbbell yields a spanning tree.

Overview of techniques. Call two pairs *compatible* if some ε-separated dumbbell covers both. Then minWSPD is a clique cover problem in the compatibility graph on the $\binom{n}{2}$ pairs, restricted to the cliques that are realized by dumbbells. We solve it with the local-growth technique of Nieberg, Hurink, and Kern [26, 25], in the form used by Pirwani and Salavatipour [27]. Around an uncovered pair, we grow hop-balls in the compatibility graph until increasing the radius by one multiplies the optimal cover size of the ball by at most $1 + \delta/3$. No dumbbell of an optimal solution covers pairs of two different resulting cores, which gives the lower bound. Two further ingredients make this work. First, a separated dumbbell only covers pairs whose lengths differ by a factor at most $1 + 2\varepsilon$, so a shifting over scales in the spirit of Hochbaum and Maass [20] reduces the problem to bands of bounded scale ratio, in which the optimal cover size of a hop-ball grows polynomially with its radius. Second, inside a band, rounding the sides of a dumbbell to a grid of cell size $\Theta(\delta\varepsilon)$ times the scale relaxes the separation by a factor at most $1 + \delta$ and turns every local subproblem into a finite problem of size depending only on d and δ, which we solve by brute force.

Further related work. Beyond doubling metrics, WSPDs of near-linear size exist for unit-disk graph metrics [13, 19], low-density graphs [15], and c-packed graphs [8]. Relaxations of WSPDs include semi-separated pair decompositions, which require only one side to be small [30, 1], and the approximate biclique covers of Buchin et al. [3], which drop the condition on the diameters. Since every dumbbell is a biclique of the complete graph on P, minWSPD is a restricted biclique cover problem. For general bipartite graphs on n vertices, minimum biclique cover is hard to approximate within $n^{1-\gamma}$ for every $\gamma > 0$ [7]. For weighted biclique covers of δ-disk bipartite graphs, Esmailpour, Le, and Sintos [11, Thm. D.3] relax the distance threshold by a factor $1 + \varphi$, as we relax the separation, and lose a factor $O((\varphi^{-d} + \log n) \log n)$ in the cost. Similarly, Har-Peled and Lee [16] give a PTAS for covering points by weighted fat objects that may expand by a δ-fraction of their diameter. Local search gives approximation schemes for many geometric set cover and hitting set problems with pseudodisks [23], and shifting and local growth give approximation schemes for covering and packing problems in geometric intersection graphs [20, 10, 26, 27, 9].

Organization. Section 2 introduces notation and proves the lower bound. Section 3 presents the approximation scheme in $\mathbb{R}^d$, and Section 4 its extension to doubling metrics. Section 5 shows the hardness in general metrics. We conclude with open problems in Section 6.

2 Preliminaries

This section fixes notation and collects the simple properties of separated dumbbells that the rest of the paper relies on. A separated dumbbell stays separated when its sides shrink, and all the pairs it covers have almost the same length, with the endpoints on each side close to each other. We then define the compatibility graph, in which MINWSPD becomes a clique cover problem, and prove a lower bound on the size of every WSPD.

For a pair $e = \{x_e, y_e\}$, its *scale* is $s(e) = \mathrm{dist}(x_e, y_e)$. We write $\mathbb{B}(z, R)$ for the closed ball of radius R centered at z. For a set S of pairs, $\theta_\varepsilon(S)$ denotes the minimum number of ε-separated dumbbells that cover all pairs of S, so $\mathrm{OPT}_\varepsilon = \theta_\varepsilon(\binom{P}{2})$. We use the following two basic facts, which follow directly from the definitions.

▶ **Observation 1** (Heredity). *If $\{X, Y\}$ is ε-separated and $\emptyset \neq X' \subseteq X$, $\emptyset \neq Y' \subseteq Y$, then $\{X', Y'\}$ is ε-separated.*

▶ **Observation 2** (Locality). *Let $\{X, Y\}$ be an ε-separated dumbbell with $\mathrm{dist}(X, Y) = \ell$ that covers a pair e. Then $\ell \leq s(e) \leq (1 + 2\varepsilon)\ell$. Moreover, $X \subseteq \mathbb{B}(x, \varepsilon\ell)$ for every $x \in X$, and similarly for Y.*

Proof. We have $s(e) \geq \ell$ by definition, and $s(e) \leq \mathrm{diam}(X) + \ell + \mathrm{diam}(Y) \leq (1 + 2\varepsilon)\ell$ by the triangle inequality. The second claim follows from $\mathrm{diam}(X) \leq \varepsilon\ell$. ◀

Two pairs a and b are *compatible* if some ε-separated dumbbell covers both. By Observation 1, a and b are compatible if and only if at least one of the dumbbells $\{\{x_a, x_b\}, \{y_a, y_b\}\}$ and $\{\{x_a, y_b\}, \{y_a, x_b\}\}$ has disjoint sides and is ε-separated. Hence, compatibility can be tested with $O(1)$ distance computations. The *compatibility graph* H has the pairs as vertices, and an edge between any two compatible pairs. The pairs covered by an ε-separated dumbbell form a clique in H.

2.1 A lower bound for cover WSPDs

A minimum spanning tree of P has $n - 1$ edges, and we show that no dumbbell covers two of them. Intuitively, if a dumbbell covered two tree edges, then the tree paths between the endpoints on each side would be short, because the sides are small. These two paths and the shorter edge would connect the endpoints of the longer edge without using it, which is impossible in a tree (Figure 1).

▶ **Lemma 3.** *Let (M, dist) be a metric space, $P \subseteq M$ with $|P| = n$, and $\varepsilon \in (0, 1)$. Every ε-WSPD of P has at least $n - 1$ dumbbells.*

Proof. Let T be a minimum spanning tree of the complete graph on P with edge weights dist. Recall the *bottleneck property* of minimum spanning trees: for all $u, v \in P$, every edge on the path between u and v in T has weight at most $\mathrm{dist}(u, v)$.

Let $\mathcal{W}$ be an ε-WSPD. Map every edge pq of T to a dumbbell $\{X, Y\} \in \mathcal{W}$ with $p \in X$ and $q \in Y$. Since T has $n - 1$ edges, it suffices to show that this map is injective.

Suppose that two distinct edges pq and $p'q'$ of T, with $\mathrm{dist}(p, q) \leq \mathrm{dist}(p', q')$, are mapped to the same dumbbell $\{X, Y\}$. After relabeling, $p, p' \in X$ and $q, q' \in Y$. Then

$$\mathrm{dist}(p, p') \leq \mathrm{diam}(X) \leq \varepsilon\ \mathrm{dist}(X, Y) \leq \varepsilon\ \mathrm{dist}(p, q) < \mathrm{dist}(p', q'),$$

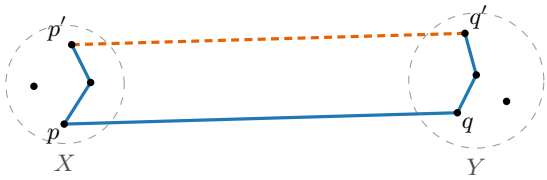

Figure 1 The proof of Lemma 3. If one dumbbell $\{X, Y\}$ covered two edges pq and $p'q'$ of a minimum spanning tree, with $\mathrm{dist}(p, q) \leq \mathrm{dist}(p', q')$, then the tree paths from p' to p and from q to q' would use only edges shorter than $p'q'$. Together with pq, they would connect p' and q' (blue, solid) without using $p'q'$ (orange, dashed).

and similarly $\mathrm{dist}(q, q') < \mathrm{dist}(p', q')$. By the bottleneck property, all edges on the tree paths from p' to p and from q to q' are shorter than $p'q'$, so neither path contains $p'q'$. Together with the edge $pq \neq p'q'$, these paths form a walk from p' to q' in $T - p'q'$, contradicting the fact that every edge of a tree is a bridge. ◀

The bound of Lemma 3 is attained already on the line: for the points $p_i = (4/\varepsilon)^i$, $i = 1, \ldots, n$, the $n - 1$ dumbbells $\{\{p_1, \ldots, p_i\}, \{p_{i+1}\}\}$, $i = 1, \ldots, n - 1$, form an ε-WSPD, as for the planar points $(i, (4/\varepsilon)^i)$ of Buchin et al. [3].

3 Bicriteria Approximation Scheme

The scheme follows the local-growth approach for covering problems in geometric intersection graphs. Starting from an uncovered pair, it grows a neighborhood in the compatibility graph until one more step of growth barely increases the cost of covering it. It then covers the enlarged neighborhood optimally, removes the covered pairs, and repeats. The few extra dumbbells spent on the outer layer buy a useful property: every pair compatible with the inner neighborhood is removed, so no dumbbell covers pairs of two different inner neighborhoods, and an optimal WSPD pays separately for each of them. For this to give a factor $1 + \delta$ in polynomial time, the growth must stop after a bounded number of steps, and each neighborhood must be covered optimally in bounded time. The first property requires that all pairs have comparable lengths, and the second requires a finite set of candidate dumbbells, which is where we relax the separation.

▶ **Theorem 4.** *There is a function f such that the following holds. Let $d \geq 1$ and $\delta \in (0, 1]$. Given a set P of n points in $\mathbb{R}^d$ and $\varepsilon \in (0, 1)$, one can compute in $f(d, \delta) \cdot n^6$ time a $(1 + 2\delta/3)\varepsilon$-WSPD $\mathcal{W}$ of P with $|\mathcal{W}| \leq (1 + \delta)\mathrm{OPT}_\varepsilon$. Moreover, $f(d, \delta) = 2^{2^{O(d(1/\delta + \log(d/\delta)))}}$.*

Throughout this section, let $\zeta = \delta/3 \leq 1/3$, $m = \lceil 1/\zeta \rceil$, $\varepsilon' = (1 + 2\zeta)\varepsilon$, and $\Lambda = (1 + 2\varepsilon)^m \leq 3^m$ (as $\varepsilon < 1$). The algorithm has three layers. It partitions the pairs into bands of bounded scale ratio by shifting over scales (Section 3.1). Inside each band, the dumbbells are rounded to a grid, which makes local subproblems finite and is the only step that relaxes the separation from ε to ε' (Section 3.2). Finally, each band is covered by local growth in the compatibility graph (Section 3.3).

3.1 Shifting over scales

By Observation 2, a dumbbell covers only pairs of almost the same length. We therefore group the pairs by length into bands, each spanning a bounded ratio of lengths, and cover each band separately. A dumbbell of an optimal WSPD that covers pairs of two bands is

then paid twice, but this happens only for dumbbells whose pairs straddle a band boundary. As Hochbaum and Maass [20] do with strips, we try m shifted positions of the boundaries. For some of them, only a $1/m$ fraction of the optimal dumbbells is paid twice.

If $n \leq 1$, the empty family is an ε-WSPD, so we assume $n \geq 2$. Let $s_0 = \min_e s(e)$ and define the *level* of a pair e as $\lambda(e) = \lfloor \log_{1+2\varepsilon}(s(e)/s_0) \rfloor$. For an offset $o \in \{0, \ldots, m-1\}$, the pairs with $\lfloor (\lambda(e) - o)/m \rfloor = j$ form the *band* E_j^o. Each band consists of m consecutive levels, so the scales of its pairs lie in $[\sigma_j, \Lambda\sigma_j)$ with $\sigma_j = s_0(1+2\varepsilon)^{jm+o}$. We assume the real RAM model with the floor function and logarithms. This model is too powerful in general [28], but we use these operations only to compute levels and to locate grid cells (Section 3.2), on quantities computed from the input and never to encode or extract bits. The algorithm processes only the nonempty bands, of which there are at most $\binom{n}{2}$ per offset, so its running time does not depend on the spread of P.

For each offset o and each nonempty band j, the algorithm computes a family $\mathcal{W}_j^o$ of ε'-separated dumbbells covering E_j^o (Section 3.3). It outputs $\mathcal{W} = \bigcup_j \mathcal{W}_j^o$ for an offset o that minimizes $\sum_j |\mathcal{W}_j^o|$.

▶ **Lemma 5.** *There is an offset o with $\sum_j \theta_\varepsilon(E_j^o) \leq (1 + 1/m)\mathrm{OPT}_\varepsilon$.*

Proof. Fix a minimum ε-WSPD $\mathcal{W}^*$ and an offset o. For every band E_j^o, the dumbbells of $\mathcal{W}^*$ that cover at least one pair of E_j^o cover all of E_j^o. Therefore $\sum_j \theta_\varepsilon(E_j^o) \leq \sum_{Z \in \mathcal{W}^*} \beta^o(Z)$, where $\beta^o(Z)$ is the number of bands (for offset o) that contain a pair covered by Z.

By Observation 2, the pairs covered by Z have scales in $[\ell, (1+2\varepsilon)\ell]$, so their levels take at most two consecutive values λ_0 and $\lambda_0 + 1$. Hence $\beta^o(Z) \leq 2$, and $\beta^o(Z) = 2$ only for the single offset $o \equiv \lambda_0 + 1 \pmod{m}$ (Figure 2). Averaging over the m offsets gives $\frac{1}{m}\sum_o \sum_Z \beta^o(Z) \leq (1+1/m)|\mathcal{W}^*|$, so some offset o satisfies $\sum_Z \beta^o(Z) \leq (1+1/m)\mathrm{OPT}_\varepsilon$. ◀

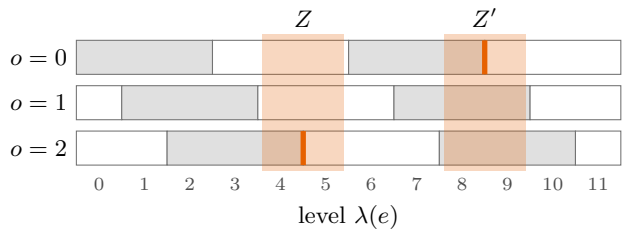

Figure 2 Bands of $m = 3$ consecutive levels for the offsets $o = 0, 1, 2$. The pairs covered by a dumbbell have at most two consecutive levels, so the dumbbell meets two bands for at most one offset (thick boundary): $o = 2$ for Z and $o = 0$ for Z'.

3.2 Rounding

The neighborhoods grown in Section 3.3 are geometrically small, but they may contain arbitrarily many points, so the number of dumbbells that can cover them is unbounded. To solve them by brute force, we round every dumbbell to a grid whose cells are much smaller than ε times the scale of the band: each side is replaced by the union of the cells it meets. Points in the same cell then become interchangeable, and a neighborhood is described by the bounded number of cells it meets. Rounding enlarges the sides by a tiny amount, so it increases their diameters and decreases their distance only slightly. This relaxes the separation from ε to ε', but it never increases the number of dumbbells.

Fix a band with scales in $[\sigma, \Lambda\sigma)$ and consider the axis-parallel grid of half-open cubes $\prod_{i=1}^{d}[k_i\eta, (k_i+1)\eta)$, $k_i \in \mathbb{Z}$, of side

$$\eta = \frac{\zeta\,\varepsilon\,\sigma}{12\sqrt{d}}.$$

Every point of $\mathbb{R}^d$ lies in exactly one cell. A *cell dumbbell* is a set $\{A, B\}$ of two nonempty finite sets of grid cells such that the regions $\bigcup A$ and $\bigcup B$ satisfy

$$\max\{\operatorname{diam}(\textstyle\bigcup A), \operatorname{diam}(\bigcup B)\} \le \varepsilon' \operatorname{dist}(\bigcup A, \bigcup B).$$

Here the diameter of a region and the distance between two regions are those of their closures. It covers a pair e if one endpoint of e lies in $\bigcup A$ and the other in $\bigcup B$. If both sides contain points of P, the dumbbell $\{P \cap \bigcup A,\ P \cap \bigcup B\}$ is ε'-separated, because taking subsets does not increase diameters and does not decrease distances. We identify a cell dumbbell with this dumbbell (Figure 3). For a set S of pairs of the band, $\theta_\square(S)$ denotes the minimum number of cell dumbbells covering S.

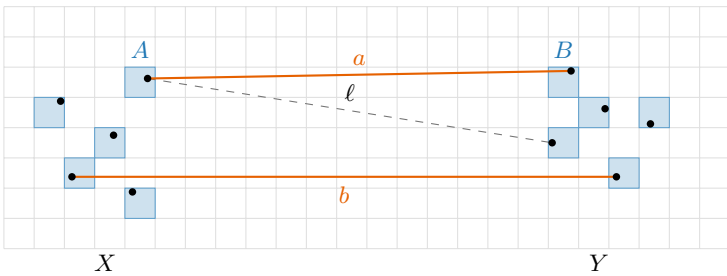

Figure 3 A dumbbell $\{X, Y\}$ covering the pairs a and b, and its cell dumbbell $\{A, B\}$.

▶ **Lemma 6.** *Let $\{X, Y\}$ be an ε-separated dumbbell covering a pair e with $s(e) \ge \sigma$, and let A and B be the sets of cells meeting X and Y. Then $\{A, B\}$ is a cell dumbbell that covers every pair covered by $\{X, Y\}$. In particular, $\theta_\square(S) \le \theta_\varepsilon(S)$ for every set S of pairs of the band.*

Proof. Let $\ell = \operatorname{dist}(X, Y)$. By Observation 2, $\ell \ge \sigma/(1+2\varepsilon) \ge \sigma/3$, so $\eta\sqrt{d} \le \zeta\varepsilon\ell/4$. Every point of $\bigcup A$ is within distance $\eta\sqrt{d}$ of a point of X. Therefore

$$\operatorname{diam}(\textstyle\bigcup A) \le \varepsilon\ell + 2\eta\sqrt{d} \le \varepsilon\ell(1+\zeta/2) \quad \text{and} \quad \operatorname{dist}(\bigcup A, \bigcup B) \ge \ell - 2\eta\sqrt{d} \ge \ell(1-\zeta/2),$$

and the same holds for B. Since $1+\zeta/2 \le (1+2\zeta)(1-\zeta/2)$ for $\zeta \le 1/2$, the dumbbell $\{A, B\}$ is ε'-separated. Finally, $X \subseteq \bigcup A$ and $Y \subseteq \bigcup B$, so $\{A, B\}$ covers every pair covered by $\{X, Y\}$. ◀

▶ **Lemma 7.** *Let $\rho \ge 1$, let $u, w \in \mathbb{R}^d$, and let S be a set of pairs of the band whose endpoints lie in $\mathbb{B}(u, R) \cup \mathbb{B}(w, R)$, where $R = \rho\,\varepsilon\Lambda\sigma$. Let Q be the set of cells that contain an endpoint of a pair of S. Then $|Q| \le K = 2(24\sqrt{d}\,\rho\Lambda/\zeta + 2)^d$, and:*

1. *$\theta_\square(S) \le K^2$;*
2. *$\theta_\square(S)$ and an optimal family of cell dumbbells can be computed in $g(d, \zeta, \rho) + O(|S| \log |S|)$ time, for some function g that does not depend on n or ε.*

Proof. A ball of radius R meets at most $(2R/\eta + 2)^d$ cells, which gives the bound on $|Q|$.

(1) For a pair $e \in S$ with endpoints in cells c and c', the set $\{\{c\},\{c'\}\}$ covers e. It is a cell dumbbell because $\operatorname{diam}(c) = \eta\sqrt{d} \le \zeta\varepsilon\sigma/12$ and $\operatorname{dist}(c,c') \ge \sigma - 2\eta\sqrt{d} \ge \sigma/2$ (in particular, $c \ne c'$).

(2) A cell dumbbell that covers no pair of S can be dropped. If a cell dumbbell $\{A,B\}$ covers a pair of S, then $A \cap Q$ and $B \cap Q$ are nonempty, and $\{A \cap Q, B \cap Q\}$ is a cell dumbbell that covers the same pairs of S, because shrinking the regions does not increase diameters and does not decrease distances. So we may assume $A, B \subseteq Q$. Every cell dumbbell covering a pair of S satisfies $\operatorname{dist}(\bigcup A, \bigcup B) \le \Lambda\sigma$, hence $\operatorname{diam}(\bigcup A) \le \varepsilon'\Lambda\sigma \le 2\varepsilon\Lambda\sigma$. Therefore, if $c \in A$, then every cell of A lies in the cube of $(2D/\eta+3)^d$ cells centered at c, where $D = 2\varepsilon\Lambda\sigma$, and so A is contained in a set of at most $K_0 = (48\sqrt{d}\,\Lambda/\zeta + 3)^d$ cells determined by any one of its cells. There are at most $K \cdot 2^{K_0}$ choices for each side. The diameter of a union of cells and the distance between two unions of cells are attained by pairs of cells and can be computed from their corners, so testing a candidate takes time polynomial in K_0. Whether a cell dumbbell covers a pair of S depends only on the cells that contain the endpoints of the pair, so by sorting in $O(|S|\log|S|)$ time we replace S by its set of at most K^2 distinct pairs of cells. By (1), an optimal solution uses at most K^2 cell dumbbells, and exhaustive search over families of at most K^2 candidates takes time bounded by a function of d, ζ, and ρ only. Finally, the scale $\varepsilon\sigma$ cancels in R/η and D/η, and $\Lambda \le 3^m$, so after replacing Λ by 3^m none of these bounds depend on ε. ◀

3.3 Local growth inside a band

We now cover a single band. The algorithm grows a ball around an uncovered pair in the compatibility graph, one hop at a time, as long as each hop multiplies the optimal cover size of the ball by more than $1+\zeta$. The endpoints of the pairs in a ball of radius r lie within distance $r\varepsilon\Lambda\sigma$ of the endpoints of the starting pair, where σ is the scale of the band, so after rounding, the optimal cover size grows only polynomially with r, and the exponential growth must stop after a bounded number of hops. The last ball, the core, is then covered together with one more hop, which costs at most a factor $1+\zeta$ more than covering the core alone, and the pairs covered are removed. In particular, every pair compatible with the core is removed, so no dumbbell covers pairs of two different cores, and an optimal cover pays for each core separately.

Fix a band $E = E_j^o$ with scales in $[\sigma, \Lambda\sigma)$. The algorithm maintains the set $U \subseteq E$ of uncovered pairs, initially $U = E$. Let $H[U]$ be the compatibility graph induced by U, and let $N^r(v)$ be the set of pairs at hop distance at most r from v in $H[U]$. While $U \ne \emptyset$, the algorithm performs a *round* consisting of the following steps:

1. pick any $v \in U$;
2. let $r \ge 0$ be the smallest integer with $\theta_\square(N^{r+1}(v)) \le (1+\zeta)\,\theta_\square(N^r(v))$;
3. add to $\mathcal{W}_j^o$ a minimum family of cell dumbbells covering $N^{r+1}(v)$, and remove from U all pairs that it covers.

The integer r in step 2 exists by Lemma 9 below. We call $N^r(v)$ the *core* of the round. In round i, we write U_i for the set U at the start of the round, and v_i, r_i, and $C_i = N^{r_i}(v_i)$ for the objects chosen in the round. All pairs of $N^{r_i+1}(v_i)$ are removed from U in round i.

▶ **Lemma 8.** *Let $C_1, C_2, \ldots$ be the cores in the order of the rounds. Then $\sum_i \theta_\square(C_i) \le \theta_\varepsilon(E)$, and consequently $|\mathcal{W}_j^o| \le (1+\zeta)\,\theta_\varepsilon(E)$.*

Proof. Let $\mathcal{W}^*$ be a minimum family of ε-separated dumbbells covering E. We claim that no dumbbell $Z \in \mathcal{W}^*$ covers pairs $a \in C_i$ and $b \in C_k$ with $i < k$. Suppose it does (Figure 4). Then a and b are compatible. Since $b \in U_k \subseteq U_i$ and $a \in C_i \subseteq U_i$, the pairs a and b

are adjacent in $H[U_i]$. Hence $b \in N^{r_i+1}(v_i)$, so b was removed from U in round $i < k$, a contradiction.

Therefore, the sets $\mathcal{W}_i^* = \{Z \in \mathcal{W}^* : Z \text{ covers a pair of } C_i\}$ are pairwise disjoint. Since $\mathcal{W}_i^*$ covers C_i, Lemma 6 gives $|\mathcal{W}_i^*| \geq \theta_\varepsilon(C_i) \geq \theta_\square(C_i)$. Summing over i gives the first claim. For the second claim, round i adds $\theta_\square(N^{r_i+1}(v_i)) \leq (1+\zeta)\,\theta_\square(C_i)$ dumbbells to $\mathcal{W}_j^o$, and summing over the rounds gives the bound. ◀

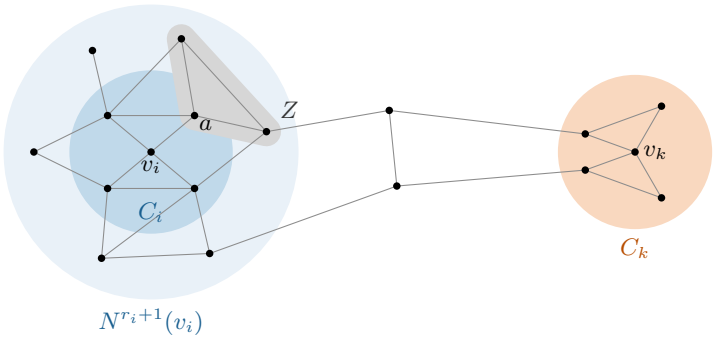

Figure 4 Round i in $H[U_i]$ (schematic, $r_i = 1$). The pairs of the core C_i (dark) and of $N^{r_i+1}(v_i)$ (light) are removed. The pairs covered by a dumbbell Z form a clique of H, so if Z covers a pair $a \in C_i$, then all its pairs in U_i lie in $N^{r_i+1}(v_i)$. Hence no dumbbell covers pairs of C_i and of a later core C_k.

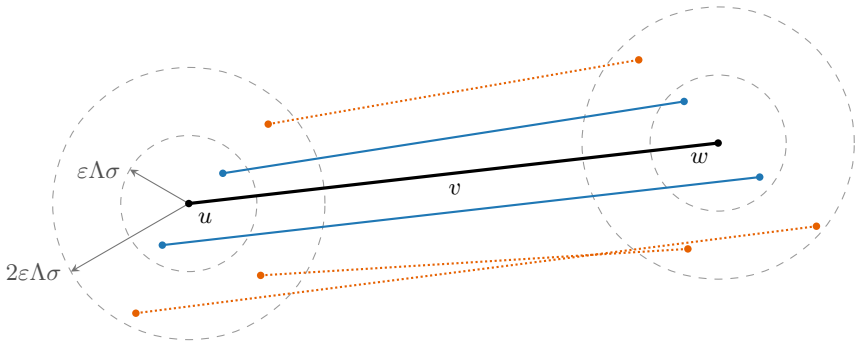

Figure 5 Pairs at hop distance 1 (blue, solid) and 2 (orange, dotted) from the pair $v = \{u, w\}$ in H (schematic). By the proof of Lemma 9, their endpoints lie in the dashed balls of radius $\varepsilon\Lambda\sigma$ and $2\varepsilon\Lambda\sigma$ around u and w.

▶ **Lemma 9.** *In every round, an integer r as in step 2 exists and satisfies $r \leq r_{\max}$ for some $r_{\max} = O\big((d/\zeta)(1/\zeta + \log(d/\zeta))\big)$.*

Proof. Let $a, b \in E$ be compatible, and let Z be an ε-separated dumbbell covering both. Each side of Z contains one endpoint of a and one endpoint of b. By Observation 2, these two endpoints are at distance at most $\varepsilon\, s(a) < \varepsilon\Lambda\sigma$. By induction, all endpoints of pairs of $N^r(v)$ lie within distance $r\varepsilon\Lambda\sigma$ of an endpoint of v (Figure 5).

By Lemma 7(1) with $\rho = r+1$, and since $2 \leq 2\sqrt{d}\,(r+1)\Lambda/\zeta$, this implies $\theta_\square(N^r(v)) \leq \mu_r := 4\big(26\sqrt{d}\,(r+1)\Lambda/\zeta\big)^{2d}$. Now let $r \geq 1$ be such that the condition of step 2 fails for every $r' < r$, that is, $\theta_\square(N^{r'+1}(v)) > (1+\zeta)\,\theta_\square(N^{r'}(v))$. Since $\theta_\square(N^0(v)) = \theta_\square(\{v\}) = 1$, we get $(1+\zeta)^r < \theta_\square(N^r(v)) \leq \mu_r$. Taking logarithms and using $\ln(1+\zeta) \geq \zeta/2$, $\ln\Lambda \leq$

$m \ln 3$, and $m \le 1/\zeta + 1$, we get $r \le \xi\,(L + \ln(r+1))$ with $\xi = O(d/\zeta)$, $\xi \ge 1$, and $L = 1/\zeta + \ln(d/\zeta) \ge 1$. This implies $r = O(\xi(L + \ln \xi))$, and $\ln \xi = O(L)$, so $r \le r_{\max}$ for some $r_{\max} = O\big((d/\zeta)(1/\zeta + \log(d/\zeta))\big)$. Hence the condition of step 2 holds for some $r \le r_{\max}$. ◀

Proof of Theorem 4. All dumbbells in the output are cell dumbbells whose cells contain points of P (by the proof of Lemma 7(2)), hence ε'-separated, and $\varepsilon' = (1 + 2\delta/3)\varepsilon$. Let o^* be the offset of Lemma 5. Since the algorithm chooses an offset that minimizes $\sum_j |\mathcal{W}_j^o|$, Lemmas 5 and 8 give

$$|\mathcal{W}| \le \sum_j |\mathcal{W}_j^{o^*}| \le (1+\zeta) \sum_j \theta_\varepsilon(E_j^{o^*}) \le (1+\zeta)^2 \mathrm{OPT}_\varepsilon \le (1+3\zeta)\mathrm{OPT}_\varepsilon = (1+\delta)\mathrm{OPT}_\varepsilon,$$

using $1/m \le \zeta$ and $\zeta \le 1$.

By Lemma 9, every value $\theta_\square(N^r(v))$ evaluated by the algorithm has $r \le r_{\max} + 1$, and $N^r(v)$ satisfies the hypothesis of Lemma 7 with $\rho = r_{\max} + 2$ and u, w the endpoints of v. Hence each such value, and the corresponding optimal family, is computed in $g'(d,\delta) + O(n^2 \log n)$ time. Here, since $\log \Lambda \le m \log 3$ and $\log r_{\max} = O(\log(d/\zeta))$, the numbers K and K_0 of the proof of Lemma 7 are $2^{O(d(1/\zeta + \log(d/\zeta)))}$. The exhaustive search examines at most $(K \cdot 2^{K_0})^{2K^2} = 2^{O(K_0 K^2)}$ families, so $g'(d,\delta) = 2^{2^{O(d(1/\delta + \log(d/\delta)))}}$. Building H takes $O(n^4)$ compatibility tests. For each of the m offsets, there are $O(n^2)$ rounds in total over all bands, and each round computes its hop-balls by one breadth-first search in $O(n^4)$ time and evaluates $O(r_{\max})$ values of $\theta_\square$. The total running time is therefore $O(m\, n^6) + O(m\, r_{\max}\, n^2 (g'(d,\delta) + n^2 \log n))$, which is $f(d,\delta) \cdot n^6$, since $\zeta = \delta/3$. ◀

▶ Remark 10. The only place where the separation is relaxed is the rounding. Suppose that, for a set S of pairs as in Lemma 7, one could decide in $|S|^{h(k)}$ time whether S can be covered by k ε-separated dumbbells, and find such a cover. Then Lemmas 5, 8, and 9 hold with θ_ε in place of $\theta_\square$ throughout. The only additional step is a bound on $\theta_\varepsilon(S)$ for such sets S, which replaces Lemma 7(1). For this, consider the grid of side $\eta' = \varepsilon\sigma/(4\sqrt{d})$. If a pair of the band has endpoints in cells c and c', then $\mathrm{diam}(c) = \varepsilon\sigma/4$ and $\mathrm{dist}(c, c') \ge \sigma/2$, so $\{P \cap c, P \cap c'\}$ is an ε-separated dumbbell that covers the pair. As in Lemma 7, this gives $\theta_\varepsilon(S) \le K'^2$ with $K' = 2(8\sqrt{d}\,\rho\Lambda + 2)^d$. Hence the local-growth radius is bounded by a function of d and δ, every local instance has optimum at most K'^2 for the corresponding ρ, and we would obtain an ε-WSPD of size at most $(1+\delta)\mathrm{OPT}_\varepsilon$ in $n^{O(h(K'^2))}$ time, that is, a PTAS that does not relax ε. For clique partition in unit disk graphs, such local solvers rely on the fact that optimal clusterings in the plane can be separated by lines [6]; see [27, 9]. We do not know an analogue of this property for WSPDs, even in the plane.

4 Doubling Metrics

The scheme of Section 3 uses the geometry of $\mathbb{R}^d$ only in the rounding step. The grid needs coordinates, but the analysis uses only two of its properties: every cell has small diameter, and a ball meets few cells. Both properties also hold for a net of the points in a metric space of bounded doubling dimension. We therefore replace the grid cells by clusters of points around the net points, and the rest of the algorithm and of its analysis carries over.

A metric space has *doubling dimension at most* κ if every ball of radius $2R$ can be covered by 2^κ balls of radius R, and its doubling dimension is the smallest such κ. We use the doubling dimension of P itself, viewed as a metric space, so all balls below are centered at points of P. The doubling dimension of $\mathbb{R}^d$ is $\Theta(d)$, and a subset has at most twice the

doubling dimension of the space containing it. We assume that distances can be evaluated in $O(1)$ time.

▶ **Theorem 11.** *There is a function f such that the following holds. Let $\kappa \geq 1$ and $\delta \in (0,1]$. Given a set P of n points of doubling dimension at most κ in a metric space and $\varepsilon \in (0,1)$, one can compute in $f(\kappa,\delta) \cdot n^6$ time a $(1+2\delta/3)\varepsilon$-WSPD $\mathcal{W}$ of P with $|\mathcal{W}| \leq (1+\delta)\mathrm{OPT}_\varepsilon$. Moreover, $f(\kappa,\delta) = 2^{2^{O(\kappa(1/\delta+\log(\kappa/\delta)))}}$, and the algorithm does not need to know κ.*

The algorithm is that of Section 3, with the same $\zeta = \delta/3$, m, ε', and Λ, and with the grid replaced by a net. Observations 1 and 2 and Lemma 5 and the first paragraph of the proof of Lemma 9 use only the triangle inequality, so they hold verbatim. Lemma 8 also holds once $\theta_\square$ is redefined as below, with Lemma 12 in place of Lemma 6. The only geometric part of Section 3 is Section 3.2, and we now replace Lemmas 6 and 7.

Nets and cells. Fix a band with scales in $[\sigma, \Lambda\sigma)$ and let $\eta = \zeta\varepsilon\sigma/48$. Let $P_\eta \subseteq P$ be a maximal subset with pairwise distances greater than η, computed greedily in $O(n^2)$ time. By maximality, every point of P is within distance η of P_η. Assign every point of P to a nearest point of P_η. The *cell* $\Gamma(q)$ of $q \in P_\eta$ is the set of points assigned to q, so $\Gamma(q) \subseteq \mathbb{B}(q,\eta)$ and $\mathrm{diam}(\Gamma(q)) \leq 2\eta$. For $A \subseteq P_\eta$, let $\Gamma(A) = \bigcup_{q\in A}\Gamma(q)$. A *cell dumbbell* is now a set $\{A,B\}$ of two nonempty subsets of P_η such that $\{\Gamma(A),\Gamma(B)\}$ is an ε'-separated dumbbell, which we identify with $\{A,B\}$. As before, $\theta_\square(S)$ denotes the minimum number of cell dumbbells covering a set S of pairs. Unlike grid cells, these cells are sets of points of P, so separation is tested directly on the points.

▶ **Lemma 12.** *Let $\{X,Y\}$ be an ε-separated dumbbell covering a pair e with $s(e) \geq \sigma$, and let $A,B \subseteq P_\eta$ be the net points whose cells meet X and Y. Then $\{A,B\}$ is a cell dumbbell that covers every pair covered by $\{X,Y\}$. In particular, $\theta_\square(S) \leq \theta_\varepsilon(S)$ for every set S of pairs of the band.*

Proof. Let $\ell = \mathrm{dist}(X,Y) \geq \sigma/(1+2\varepsilon) \geq \sigma/3$. Every $z \in \Gamma(A)$ lies in a cell $\Gamma(q)$ that contains a point $x \in X$, so $\mathrm{dist}(z,x) \leq \mathrm{dist}(z,q) + \mathrm{dist}(q,x) \leq 2\eta$. Hence $\mathrm{diam}(\Gamma(A)) \leq \varepsilon\ell + 4\eta$, similarly for B, and $\mathrm{dist}(\Gamma(A),\Gamma(B)) \geq \ell - 4\eta$. Since $4\eta = \zeta\varepsilon\sigma/12 \leq \zeta\varepsilon\ell/4$, the dumbbell $\{\Gamma(A),\Gamma(B)\}$ is $\varepsilon(1+\zeta/4)/(1-\zeta/4)$-separated, and $1+\zeta/4 \leq (1+2\zeta)(1-\zeta/4)$ for $\zeta \leq 1/2$. Finally, $X \subseteq \Gamma(A)$ and $Y \subseteq \Gamma(B)$. ◀

▶ **Lemma 13.** *For every $z \in P$ and $R \geq \eta$, we have $|P_\eta \cap \mathbb{B}(z,R)| \leq (4R/\eta)^\kappa$.*

Proof. Let $t = \lfloor \log_2(2R/\eta) \rfloor + 1$. Applying the doubling property t times covers $\mathbb{B}(z,R) \cap P$ by $2^{\kappa t} \leq (4R/\eta)^\kappa$ balls of radius $R/2^t < \eta/2$. Each of them contains at most one point of P_η. ◀

▶ **Lemma 14.** *Let $\rho \geq 1$, let $u,w \in P$, and let S be a set of pairs of the band whose endpoints lie in $\mathbb{B}(u,R) \cup \mathbb{B}(w,R)$, where $R = \rho\varepsilon\Lambda\sigma$. Let $Q \subseteq P_\eta$ be the net points whose cells contain an endpoint of a pair of S. Then $|Q| \leq K = 2(384\rho\Lambda/\zeta)^\kappa$, and:*

1. *$\theta_\square(S) \leq K^2$;*
2. *$\theta_\square(S)$ and an optimal family of cell dumbbells can be computed in $2^{O(K^3)} + O(4^K n^2)$ time.*

Proof. Every $q \in Q$ is within distance $\eta \leq R$ of an endpoint, so $Q \subseteq \mathbb{B}(u,2R) \cup \mathbb{B}(w,2R)$, and Lemma 13 gives $|Q| \leq 2(8R/\eta)^\kappa = K$, since $R/\eta = 48\rho\Lambda/\zeta$.

(1) Let $e \in S$ have endpoints in the cells of q and q'. The set $\{\{q\}, \{q'\}\}$ is a cell dumbbell covering e, because $\operatorname{diam}(\Gamma(q)) \le 2\eta = \zeta\varepsilon\sigma/24$ and $\operatorname{dist}(\Gamma(q), \Gamma(q')) \ge \sigma - 4\eta \ge \sigma/2$ (in particular, $q \ne q'$).

(2) A cell dumbbell that covers no pair of S can be dropped. If a cell dumbbell $\{A, B\}$ covers a pair of S, then $A \cap Q$ and $B \cap Q$ are nonempty, and replacing $\{A, B\}$ by $\{A \cap Q, B \cap Q\}$ keeps it ε'-separated by Observation 1 and keeps the pairs of S that it covers, so we may assume $A, B \subseteq Q$. There are at most 4^K such candidates, and each is tested on the points in $O(n^2)$ time. Whether a candidate covers a pair of S depends only on the cells of its endpoints, which gives at most K^2 types of pairs. By (1), an optimal solution uses at most K^2 candidates, and exhaustive search over families of at most K^2 of the 4^K candidates takes $2^{O(K^3)}$ time. ◀

Proof of Theorem 11. All output dumbbells are cell dumbbells, hence ε'-separated, and $\varepsilon' = (1 + 2\delta/3)\varepsilon$. The size bound follows from Lemmas 5 and 8, with Lemma 12 in place of Lemma 6, exactly as in the proof of Theorem 4. For the running time, the proof of Lemma 9, with Lemma 14(1) in place of Lemma 7, gives

$$(1+\zeta)^r < \theta_\square(N^r(v)) \le 4\big(384(r+1)\Lambda/\zeta\big)^{2\kappa},$$

so $r \le r_{\max}$ for some $r_{\max} = O\big((\kappa/\zeta)(1/\zeta + \log(\kappa/\zeta))\big)$. Every set $N^r(v)$ that the algorithm evaluates has $r \le r_{\max} + 1$ and satisfies the hypothesis of Lemma 14 with $\rho = r_{\max} + 2$, and $\Lambda \le 3^m$ gives $K = 2^{O(\kappa(1/\zeta + \log(\kappa/\zeta)))}$. As in Lemma 7, the scale $\varepsilon\sigma$ cancels in R/η, so with $\Lambda \le 3^m$ the bound on K does not depend on ε. The remaining steps are as in the proof of Theorem 4, and computing one net per nonempty band takes $O(n^4)$ time per offset, so the total time is $f(\kappa, \delta) \cdot n^6$. The algorithm never uses κ: the growth stops by itself, and the local solver enumerates the subsets of whichever set Q it encounters. ◀

5 Hardness in General Metrics

Our proof does not extend to general metrics, because the growth bound of Lemma 9 requires bounded doubling dimension. In fact, relaxing the separation by a factor $\tau < 2$ keeps the problem hard there. We show this with a direct reduction from clique cover, which works for every ε and $\tau \ge 1$ with $\tau < 2$ and $\tau\varepsilon < 1/2$, in particular for every $\tau < 2$ when $\varepsilon \le 1/4$. The metric consists of copies of a graph: two vertices of a copy are at distance 1 if they are adjacent and 2 otherwise, and two vertices of different copies are at a large distance. A separated dumbbell between two copies then has cliques as sides, even if the separation is relaxed by a factor below 2, so the minimum WSPD encodes clique covers of the graph. For $\tau = 1$ and $\varepsilon = 1/3$, this is the setting of the hardness results of Buchin et al. [3, Thm. 28, Lemma 29], and our construction pads the instance with many copies of the input graph as in their Lemma 29. The new ingredients are the slack $\tau < 2$, all $\varepsilon < 1/2$, and the explicit exponent 2/3, which follows from the product bound (1) and the hardness of the chromatic number.

For a graph G with vertex set V, a *clique* is a nonempty set of pairwise adjacent vertices (a single vertex is a clique). Let $\bar{\chi}(G)$ be the minimum number of cliques whose union is V, and let $\omega(G)$ be the maximum size of a clique. For two graphs G_1 and G_2 with vertex sets V_1 and V_2, let $\operatorname{cc}(G_1, G_2)$ be the minimum number of products $A \times B$ of a clique A of G_1 and a clique B of G_2 whose union is $V_1 \times V_2$. If G_2 has a single vertex, then $\operatorname{cc}(G_1, G_2) = \bar{\chi}(G_1)$. Moreover,

$$\frac{|V_1|\,|V_2|}{\omega(G_1)\,\omega(G_2)} \;\le\; \operatorname{cc}(G_1, G_2) \;\le\; \bar{\chi}(G_1)\,\bar{\chi}(G_2), \tag{1}$$

since each product covers at most $\omega(G_1)\,\omega(G_2)$ elements, and the products of two clique covers cover $V_1 \times V_2$.

Block metrics. Let $G_1, \ldots, G_k$ be graphs with pairwise disjoint vertex sets $V_1, \ldots, V_k$, and let $\Delta \geq 2$. The *block metric* on $V_1 \cup \cdots \cup V_k$ is defined as follows. Two distinct vertices of the same block V_i are at distance 1 if they are adjacent in G_i and at distance 2 otherwise, and two vertices of different blocks are at distance Δ (Figure 6). This is a metric: if x and z lie in different blocks, then every y lies in a block different from that of x or from that of z, so $\mathrm{dist}(x,y) + \mathrm{dist}(y,z) \geq \Delta$; if x and z lie in the same block, then $\mathrm{dist}(x,z) \leq 2$, while $\mathrm{dist}(x,y) + \mathrm{dist}(y,z) \geq 2$ for $y \notin \{x,z\}$.

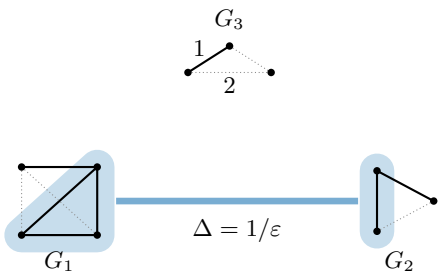

Figure 6 A block metric and a dumbbell whose sides are cliques of G_1 and G_2.

▶ **Lemma 15.** *Let $\varepsilon > 0$ and $\tau \in [1,2)$ with $\tau\varepsilon < 1/2$, and let P be a block metric with parameter $\Delta = 1/\varepsilon$. For a dumbbell $\{X,Y\}$ of P, the following are equivalent: (i) $\{X,Y\}$ is $\tau\varepsilon$-separated; (ii) $\{X,Y\}$ is ε-separated; (iii) either X and Y are singletons contained in the same block, or there are $i \neq j$ such that X is a clique of G_i and Y is a clique of G_j.*

Proof. Note that $\Delta = 1/\varepsilon > 2\tau \geq 2$. (ii) implies (i) because $\tau \geq 1$. Assume (i), and let $\ell = \mathrm{dist}(X,Y) \in \{1,2,\Delta\}$. If $\ell \leq 2$, then both sides have diameter at most $2\tau\varepsilon < 1$, so they are singletons, and since $\ell \leq 2$ they lie in the same block. If $\ell = \Delta$, then both sides have diameter at most $\tau\varepsilon\Delta = \tau < 2$. Two points at distance less than 2 lie in the same block and are adjacent, so each side is a clique of one graph, and the two sides lie in different blocks because $\ell = \Delta$. Hence (iii) holds. Finally, assume (iii). Two singletons are ε-separated. If X and Y are cliques of G_i and G_j with $i \neq j$, then all distances between X and Y are Δ and both diameters are at most $1 = \varepsilon\Delta$, so (ii) holds. ◀

▶ **Lemma 16.** *Under the assumptions of Lemma 15, every $\tau\varepsilon$-WSPD of P is an ε-WSPD, and*

$$\mathrm{OPT}_\varepsilon(P) \;=\; \sum_i \binom{|V_i|}{2} + \sum_{i<j} \mathrm{cc}(G_i, G_j).$$

Proof. The first claim follows from Lemma 15. By the same lemma, a pair of two points of V_i is covered only by the dumbbell consisting of these two singletons, and a dumbbell that covers a pair with one point in V_i and one in V_j, $i < j$, has a clique A of G_i as one side and a clique B of G_j as the other, and covers exactly the pairs corresponding to $A \times B$. Such dumbbells cover all pairs between V_i and V_j if and only if the corresponding products cover $V_i \times V_j$. Since every dumbbell covers pairs of only one of these parts, the minimum is the sum of the minima of the parts. ◀

Our main hardness result follows from the hardness of approximating the chromatic number. We use it in the following form: for every $\gamma \in (0, 1/2)$, it is NP-hard, given a graph G on n vertices such that $\bar{\chi}(G) \le n^\gamma$ or $\omega(G) \le n^\gamma$, to decide which of the two cases holds. For the complement $\bar{G}$, the proof of Theorem 1.2 of Zuckerman [31], which derandomizes a reduction of Feige and Kilian [12], shows that it is NP-hard to distinguish $\chi_f(\bar{G}) \le n^{\gamma/2}$ from $\alpha(\bar{G}) \le n^{\gamma/2}$, where χ_f is the fractional chromatic number and α the independence number. Since $\chi(\bar{G}) \le (1 + \ln n)\,\chi_f(\bar{G})$ (Lovász [22, Cor. 2 and Thm. 6]), and $\bar{\chi}(G) = \chi(\bar{G})$ and $\omega(G) = \alpha(\bar{G})$, the claim follows for n large enough.

▶ **Theorem 17.** *Let $\varepsilon > 0$ be rational, $\tau \in [1, 2)$ with $\tau\varepsilon < 1/2$, and $c < 2/3$. Unless $P = NP$, no polynomial-time algorithm computes, for every finite metric space P on N points, a $\tau\varepsilon$-WSPD of P of size at most $N^c \mathrm{OPT}_\varepsilon(P)$.*

Proof. Since $c < 2/3$, there is $\gamma \in (0, 1/2)$ with $3c + 4\gamma < 2$. Given a graph G on $n \ge 2$ vertices, let P be the block metric with parameter $\Delta = 1/\varepsilon$ and $\beta = n^2$ blocks, each a copy of G. It has $N = n^3$ points, so it can be built, and the algorithm run on it, in time polynomial in n. By Lemma 16,

$$\mathrm{OPT}_\varepsilon(P) = \beta\binom{n}{2} + \binom{\beta}{2}\,\mathrm{cc}(G, G).$$

If $\bar{\chi}(G) \le n^\gamma$, then $\mathrm{cc}(G, G) \le n^{2\gamma}$ by (1), and $\mathrm{OPT}_\varepsilon(P) \le n^4/2 + n^{4+2\gamma}/2 \le n^{4+2\gamma}$. If $\omega(G) \le n^\gamma$, then $\mathrm{cc}(G, G) \ge n^{2-2\gamma}$ by (1), and since $\binom{\beta}{2} \ge n^4/4$, every $\tau\varepsilon$-WSPD of P, which is an ε-WSPD by Lemma 16, has at least $n^{6-2\gamma}/4$ dumbbells. In the first case, the algorithm outputs at most $N^c n^{4+2\gamma} = n^{3c+4+2\gamma}$ dumbbells, which is less than $n^{6-2\gamma}/4$ as soon as $n^{2-3c-4\gamma} > 4$. Hence, for n above a constant, comparing the output size with $n^{6-2\gamma}/4$ decides which case holds, and smaller instances are solved by brute force. ◀

For $c = 0$, Theorem 17 states that no polynomial-time algorithm computes a $\tau\varepsilon$-WSPD of size $\mathrm{OPT}_\varepsilon(P)$. This case needs only the NP-completeness of clique cover [21]. Given a graph G on n vertices, let P be the block metric with parameter $\Delta = 1/\varepsilon$ and two blocks, G and a graph with a single vertex. By Lemma 16, a $\tau\varepsilon$-WSPD of P of size $\mathrm{OPT}_\varepsilon(P)$ has exactly $\binom{n}{2} + \bar{\chi}(G)$ dumbbells, so such an algorithm would compute $\bar{\chi}(G)$. For $\tau = 1$ and $\varepsilon = 1/3$, this is the reduction of Buchin et al. [3, Thm. 28].

For $\tau < 1$, a $\tau\varepsilon$-WSPD is an ε-WSPD, so for $\varepsilon < 1/2$ Theorem 17 also holds for every $\tau < 1$. The factor 2 is inherent to block metrics: for $\tau \ge 2$, a side at distance $\Delta = 1/\varepsilon$ may contain two nonadjacent vertices of a block, and the cliques no longer play a role. Metrics of this type cannot spread the two distances inside a block further apart either, because the triangle inequality forces the distance of two nonadjacent vertices with a common neighbor to be at most twice the edge length. The exponent $2/3$ is also the limit of this construction: with n^b blocks instead of n^2, the same argument gives every $c < \min\{b, 2\}/(b + 1)$, which is largest for $b = 2$.

In particular, for every $\delta \in (0, 1]$ and every fixed rational $\varepsilon \le 1/4$, no polynomial-time algorithm achieves the separation factor $1 + 2\delta/3 < 2$ of Theorem 11 in general metrics with the size factor N^c for any $c < 2/3$, unless P = NP, let alone with the size factor $1 + \delta$. Some polynomial size factor is always achievable: the $\binom{N}{2}$ dumbbells with singleton sides form a WSPD, and $\mathrm{OPT}_\varepsilon \ge N - 1$ by Lemma 3, so the best exponent lies between $2/3$ and 1. On the other hand, for $\varepsilon < 1/33$, Har-Peled et al. [19] achieve the size factor 7 with the separation factor 33 in every metric space.

6 Conclusion

We gave a bicriteria approximation scheme for MINWSPD in fixed dimension and in doubling metrics, and showed that in general metrics no such scheme exists, even for a fixed $\varepsilon \leq 1/4$, unless P = NP.

Several questions remain open. Is there a PTAS that does not relax ε, for instance via a separability lemma for optimal WSPDs in the plane? Can the approximation scheme be made to run in near-linear time? Does local search that exchanges up to k dumbbells give a $(1 + O(1/\sqrt{k}))$-approximation in the plane, as it does for covering problems with pseudodisks [23] and low-density objects [18], and for MINWSPD on the line [3]? Is MINWSPD NP-hard on the line, as asked by Buchin et al., or in the plane for a fixed ε? Does an approximation scheme exist for the partition version? Finally, in general metrics and for small ε, what is the smallest τ such that a $\tau\varepsilon$-WSPD of size $O(\mathrm{OPT}_\varepsilon)$ can be computed in polynomial time? By Theorem 17 and the result of Har-Peled et al. [19], it lies between 2 and 33. Going beyond 2 requires metrics other than the block metrics of Lemma 15.

References

1. Mohammad A. Abam and Sariel Har-Peled. New constructions of SSPDs and their applications. *Comput. Geom.*, 45(5–6):200–214, 2012. doi:10.1016/j.comgeo.2011.12.003.
2. Reuven Bar-Yehuda. One for the price of two: A unified approach for approximating covering problems. *Algorithmica*, 27(2):131–144, 2000. doi:10.1007/s004530010009.
3. Kevin Buchin, Jacobus Conradi, Sariel Har-Peled, Antonia Kalb, Abhiruk Lahiri, Lukas Plätz, Carolin Rehs, and Sampson Wong. On small pair decompositions for point sets. In *Proc. 34th Annual European Symposium on Algorithms (ESA)*, volume 388 of *LIPIcs*, pages 102:1–102:22, 2026. Full version: arXiv:2601.22728. doi:10.4230/LIPIcs.ESA.2026.102.
4. Paul B. Callahan and S. Rao Kosaraju. Faster algorithms for some geometric graph problems in higher dimensions. In *Proc. 4th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 291–300, 1993.
5. Paul B. Callahan and S. Rao Kosaraju. A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields. *J. ACM*, 42(1):67–90, 1995. doi:10.1145/200836.200853.
6. Vasilis Capoyleas, Günter Rote, and Gerhard J. Woeginger. Geometric clusterings. *J. Algorithms*, 12(2):341–356, 1991. doi:10.1016/0196-6774(91)90007-L.
7. Parinya Chalermsook, Sandy Heydrich, Eugenia Holm, and Andreas Karrenbauer. Nearly tight approximability results for minimum biclique cover and partition. In *Proc. 22nd Annual European Symposium on Algorithms (ESA)*, volume 8737 of *LNCS*, pages 235–246, 2014. doi:10.1007/978-3-662-44777-2_20.
8. Lindsey Deryckere, Joachim Gudmundsson, André van Renssen, Yuan Sha, and Sampson Wong. A WSPD, separator and small tree cover for c-packed graphs. In *Proc. 19th International Symposium on Algorithms and Data Structures (WADS)*, volume 349 of *LIPIcs*, pages 21:1–21:15, 2025. doi:10.4230/LIPIcs.WADS.2025.21.
9. Adrian Dumitrescu and János Pach. Minimum clique partition in unit disk graphs. *Graphs Combin.*, 27(3):399–411, 2011. doi:10.1007/s00373-011-1026-1.
10. Thomas Erlebach, Klaus Jansen, and Eike Seidel. Polynomial-time approximation schemes for geometric intersection graphs. *SIAM J. Comput.*, 34(6):1302–1323, 2005. doi:10.1137/S0097539702402676.
11. Aryan Esmailpour, Khoi Le, and Stavros Sintos. Compact representations of geometric bipartite graphs via weighted biclique covers, 2026. arXiv:2608.20767. doi:10.48550/arXiv.2608.20767.
12. Uriel Feige and Joe Kilian. Zero knowledge and the chromatic number. *Journal of Computer and System Sciences*, 57(2):187–199, 1998. doi:10.1006/jcss.1998.1587.

13 Jie Gao and Li Zhang. Well-separated pair decomposition for the unit-disk graph metric and its applications. *SIAM J. Comput.*, 35(1):151–169, 2005. doi:10.1137/S0097539703436357.

14 Ronald L. Graham and Henry O. Pollak. On the addressing problem for loop switching. *Bell System Tech. J.*, 50(8):2495–2519, 1971. doi:10.1002/j.1538-7305.1971.tb02618.x.

15 Joachim Gudmundsson and Sampson Wong. A well-separated pair decomposition for low density graphs. In *Proc. 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3688–3723, 2026. doi:10.1137/1.9781611978971.135.

16 Sariel Har-Peled and Mira Lee. Weighted geometric set cover problems revisited. *J. Comput. Geom.*, 3(1):65–85, 2012. doi:10.20382/jocg.v3i1a4.

17 Sariel Har-Peled and Manor Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM J. Comput.*, 35(5):1148–1184, 2006. doi:10.1137/S0097539704446281.

18 Sariel Har-Peled and Kent Quanrud. Approximation algorithms for polynomial-expansion and low-density graphs. *SIAM J. Comput.*, 46(6):1712–1744, 2017. doi:10.1137/16M1079336.

19 Sariel Har-Peled, Benjamin Raichel, and Eliot W. Robson. Well-separated pairs decomposition revisited, 2025. arXiv:2509.05997.

20 Dorit S. Hochbaum and Wolfgang Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *J. ACM*, 32(1):130–136, 1985. doi:10.1145/2455.214106.

21 Richard M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972. doi:10.1007/978-1-4684-2001-2_9.

22 László Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13(4):383–390, 1975. doi:10.1016/0012-365X(75)90058-8.

23 Nabil H. Mustafa and Saurabh Ray. Improved results on geometric hitting set problems. *Discrete Comput. Geom.*, 44(4):883–895, 2010. doi:10.1007/s00454-010-9285-9.

24 Giri Narasimhan and Michiel H. M. Smid. *Geometric Spanner Networks*. Cambridge University Press, 2007. doi:10.1017/CBO9780511546884.

25 Tim Nieberg and Johann Hurink. A PTAS for the minimum dominating set problem in unit disk graphs. In *Proc. 3rd Workshop on Approximation and Online Algorithms (WAOA)*, volume 3879 of *LNCS*, pages 296–306, 2006. doi:10.1007/11671411_23.

26 Tim Nieberg, Johann Hurink, and Walter Kern. Approximation schemes for wireless networks. *ACM Trans. Algorithms*, 4(4):49:1–49:17, 2008. doi:10.1145/1383369.1383380.

27 Imran A. Pirwani and Mohammad R. Salavatipour. A weakly robust PTAS for minimum clique partition in unit disk graphs. *Algorithmica*, 62(3–4):1050–1072, 2012. doi:10.1007/s00453-011-9503-8.

28 Arnold Schönhage. On the power of random access machines. In *Proc. 6th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 71 of *LNCS*, pages 520–529. Springer, 1979. doi:10.1007/3-540-09510-1_42.

29 Michiel H. M. Smid. The well-separated pair decomposition and its applications. In Teofilo F. Gonzalez, editor, *Handbook of Approximation Algorithms and Metaheuristics, Volume 2: Contemporary and Emerging Applications*, chapter 4, pages 71–84. Chapman and Hall/CRC, 2nd edition, 2018. doi:10.1201/9781351235426.

30 Kasturi R. Varadarajan. A divide-and-conquer algorithm for min-cost perfect matching in the plane. In *Proc. 39th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 320–331, 1998. doi:10.1109/SFCS.1998.743466.

31 David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3(6):103–128, 2007. doi:10.4086/toc.2007.v003a006.