

An $O(n \log n)$ -Time Approximation Scheme for Geometric Many-To-Many Matching

Sayan Bandyapadhyay  

Department of Computer Science, Portland State University, OR, USA

Jie Xue  

Department of Computer Science, New York University Shanghai, China

Abstract

Geometric matching is an important topic in computational geometry and has been extensively studied over decades. In this paper, we study a geometric-matching problem, known as *geometric many-to-many matching*. In this problem, the input is a set S of n colored points in \mathbb{R}^d , which implicitly defines a graph $G = (S, E(S))$ where $E(S) = \{(p, q) : p, q \in S \text{ have different colors}\}$, and the goal is to compute a minimum-cost subset $E^* \subseteq E(S)$ of edges that cover all points in S . Here the cost of E^* is the sum of the costs of all edges in E^* , where the cost of a single edge e is the Euclidean distance (or more generally, the L_p -distance) between the two endpoints of e . Our main result is a $(1 + \varepsilon)$ -approximation algorithm with an optimal running time $O_\varepsilon(n \log n)$ for geometric many-to-many matching in any fixed dimension, which works under any L_p -norm. This is the first near-linear approximation scheme for the problem in any $d \geq 2$. Prior to this work, only the bipartite case of geometric many-to-many matching was considered in \mathbb{R}^1 and \mathbb{R}^2 , and the best known approximation scheme in \mathbb{R}^2 takes $O_\varepsilon(n^{1.5} \cdot \text{poly}(\log n))$ time.

2012 ACM Subject Classification Theory of computation \rightarrow Computational geometry; Theory of computation \rightarrow Design and analysis of algorithms

Keywords and phrases many-to-many matching, geometric optimization, approximation algorithms

Digital Object Identifier 10.4230/LIPIcs.SoCG.2024.12

Related Version *Full Version*: <https://arxiv.org/abs/2402.15837>

1 Introduction

A central topic in computational geometry is the study of optimization problems on edge-weighted graphs that are defined geometrically (sometimes known as *geometric graphs*). Typically, geometric graphs use points in \mathbb{R}^d as their vertices, and the Euclidean distance (or distance under other norms) between two points naturally defines the weight of an edge. Many fundamental graph optimization problems have been investigated on geometric graphs, including minimum spanning tree [2, 30, 38], Steiner trees [8, 12, 39], traveling salesman problem [4, 5, 17], spanners [6, 24, 29], matching [1, 44, 45], etc. By exploiting the underlying geometry, these problems can usually be solved much more efficiently on geometric graphs, compared to general edge-weighted graphs.

In this paper, we focus on a particularly important class of problems on geometric graphs, the matching-related problems. These problems have applications in a wide range of areas, e.g., computational biology [11], data mining [10], computational music [40, 41], machine learning [37], etc. When studying matching problems on geometric graphs, there are two settings commonly used in the literature. The first one, called the *bipartite* setting, requires the input points to be bichromatic and the graph considered is the complete bipartite graph consisting of edges between points with different colors. The second one, called the *complete* setting, simply considers the complete graph induced by the input points.



© Sayan Bandyapadhyay and Jie Xue;


licensed under Creative Commons License CC-BY 4.0

40th International Symposium on Computational Geometry (SoCG 2024).

Editors: Wolfgang Mulzer and Jeff M. Phillips; Article No. 12; pp. 12:1–12:15

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



12:2 Approximation for Geometric Many-To-Many Matching

One of the matching problems that has received most attention on geometric graphs is the classical *minimum-weight perfect matching* problem, where the goal is to compute a set of vertex-disjoint edges with minimum total weight that cover all vertices. The problem can be solved in $O(nm + n^2 \log n)$ time on any graph with n vertices and m edges, by the seminal work of Fredman and Tarjan [21]. A popular line of research [1, 25, 26, 33, 42, 44, 45] in computational geometry investigated minimum-weight perfect matching on geometric graphs, leading to much faster exact and approximation algorithms. The exact algorithms were designed for the problem in \mathbb{R}^2 . The best algorithm for the bipartite setting [26] runs in $O(n^2 \cdot \text{poly}(\log n))$ time where n is the number of input points, while the best algorithm for the complete setting [44] runs in $O(n^{1.5} \cdot \text{poly}(\log n))$ time. The approximation algorithms are much more general and efficient. It was known that in both bipartite and complete settings, the problem admits $(1 + \varepsilon)$ -approximation algorithms with running time $O_\varepsilon(n \cdot \text{poly}(\log n))$ in \mathbb{R}^d for any fixed d [1, 33, 45]. Here $O_\varepsilon(\cdot)$ hides factors depending only on ε .

The minimum-weight perfect matching problem has an interesting variant, which is also a classical problem known as *many-to-many matching* or *edge cover* [20, 23, 27, 31, 32, 46]. In this variant, the only difference is that the edges in the solution need not to be vertex-disjoint. In other words, it simply asks for a set of edges with minimum total weight that cover all vertices. Many-to-many matching can be reduced to minimum-weight perfect matching [27], and is thus polynomial-time solvable. On geometric graphs, many-to-many matching, while having received less attention than minimum-weight perfect matching, also has a long history. Eiter and Mannila [18] introduced the problem for the first time in 1997, under the name of *link distance*, in order to measure the similarity between two sets of points. Colannino and Toussaint [15] showed that bipartite geometric many-to-many matching in \mathbb{R}^1 can be solved in $O(n^2)$ time. Later, Colannino et al. [14] improved this result and obtained an optimal $O(n \log n)$ -time algorithm, which completely settles the complexity of the problem in \mathbb{R}^1 (for the bipartite setting). Several variants of the problem in \mathbb{R}^1 have also been considered [34, 35, 36]. Recently, Bandyopadhyay et al. [9] studied bipartite geometric many-to-many matching in \mathbb{R}^2 and designed two algorithms. The first algorithm solves the problem in $O(n^2 \cdot \text{poly}(\log n))$ time. This algorithm is based on the general reduction from many-to-many matching to perfect matching, and exploits various geometric data structures to implement the reduction and the Hungarian algorithm [28] for perfect matching in an efficient way. The second algorithm is a $(1 + \varepsilon)$ -approximation algorithm which runs in $O_\varepsilon(n^{1.5} \cdot \text{poly}(\log n))$ time. The basic idea of this algorithm is similar to the first one, but it uses the multi-scale algorithm of Gabow and Tarjan [22] for perfect matching instead of the Hungarian algorithm, which can be implemented more efficiently in the geometric setting by losing a factor of at most $1 + \varepsilon$ in cost. In higher dimensions, no nontrivial results for geometric many-to-many matching were known, to the best of our knowledge.

As one can see in the above discussion, in terms of exact algorithms, the best known bounds for geometric many-to-many matching are similar to the bounds for geometric (minimum-weight) perfect matching, at least in the bipartite setting – both problems can be solved in near-quadratic time in \mathbb{R}^2 [9, 26] and are open in higher dimensions. However, in terms of approximation, geometric many-to-many matching is much less well-understood than geometric perfect matching. Even in \mathbb{R}^2 , no approximation scheme for geometric many-to-many matching with near-linear running time was known, while geometric perfect matching admits near-linear approximation schemes in any fixed dimension [1, 33, 45]. This motivates the following natural question, which is the subject of this paper.

Question: Does geometric many-to-many matching admit a $(1 + \varepsilon)$ -approximation algorithm with running time $O_\varepsilon(n \cdot \text{poly}(\log n))$ in \mathbb{R}^d for any fixed d ?

We answer the above question affirmatively by giving such an approximation scheme for geometric many-to-many matching. In fact, our results are much stronger and more general. Our algorithm has an optimal $O_\varepsilon(n \log n)$ running time, works under any L_p -norm, and applies to both bipartite and complete settings (and beyond). We shall discuss our results in detail in the next section.

Our results

We study geometric many-to-many matching in a *colored* setting which simultaneously generalizes the aforementioned bipartite and complete settings for geometric graphs. Here, each input point is associated with a color (the total number of colors can be unbounded) and the graph considered has edges between every pair of points with different colors. Clearly, the colored setting is equivalent to the bipartite setting when there are only two colors and is equivalent to the complete setting when all points have distinct colors. Let S be a set of colored points in \mathbb{R}^d . We write $E(S) = \{(p, q) : p, q \in S \text{ have different colors}\}$ as the edge set of the geometric graph induced by S . The L_p -cost of an edge $e \in E(S)$ is the L_p -distance between its two endpoints, and the L_p -cost of a subset $E \subseteq E(S)$ is the sum of the L_p -costs of all edges in E . We say $E \subseteq E(S)$ covers a point in S if the point is an endpoint of an edge in E . The *geometric many-to-many matching* problem is formally defined as follows.

GEOMETRIC MANY-TO-MANY MATCHING

Input: A set S of n colored points in \mathbb{R}^d .

Goal: Compute $E^* \subseteq E(S)$ with minimum L_p -cost which covers all points in S .

When studying the problem, we assume that the input points in S are already clustered by their colors so that we do not need extra time to compute the color-partition of S . Alternatively, one can assume the colors belong to $[n] = \{1, \dots, n\}$ and thus the color-partition of S can be computed in linear time. Our main result is the following theorem.

► **Theorem 1.** *For any fixed $d \in \mathbb{N}$ and $p \geq 1$, geometric many-to-many matching in \mathbb{R}^d under the L_p -norm admits a $(1 + \varepsilon)$ -approximation algorithm with running time $O_\varepsilon(n \log n)$.*

Note that the running time in Theorem 1 is *optimal* up to a factor depending on ε . Indeed, as observed in [14], any approximation algorithm for geometric many-to-many matching in \mathbb{R}^1 requires $\Omega(n \log n)$ time, due to a reduction from set equality.

Interestingly, our algorithm in Theorem 1 completely bypasses the reduction from many-to-many matching to minimum-weight perfect matching. This allows us to avoid the techniques for perfect matching such as augmenting paths, which were commonly used in the previous algorithms for geometric matching problems [1, 9, 33]. Instead, our algorithm exploits the nice structures of the many-to-many matching problem itself, and solves the problem by nontrivially combining Baker's shifting technique [7], grid techniques, approximate nearest-neighbor search [13], and the FPT algorithm for integer linear programming [16].

2 Preliminaries

Basic notations. We use \mathbb{N} to denote the set of natural numbers including 0. For a number $n \in \mathbb{N}$, we write $[n] = \{1, \dots, n\}$. A *colored point* in \mathbb{R}^d is a point $p \in \mathbb{R}^d$ with a color, which we denote by $\text{cl}(p)$. Let S be a set of colored points in \mathbb{R}^d . We define $E(S) = \{(p, q) : p, q \in S \text{ and } \text{cl}(p) \neq \text{cl}(q)\}$ as the *edge set* on S . Here, the pairs in $E(S)$ are *unordered*, i.e., (p, q) and (q, p) are viewed as one element in $E(S)$. For a subset $E \subseteq E(S)$, we denote by $V(E) \subseteq S$ the subset consisting of the endpoints of the edges in E .

Foreign neighbors. Let S be a set of colored points in \mathbb{R}^d . For a point $p \in S$, a *foreign neighbor* of p in S refers to another point $q \in S$ satisfying $\text{cl}(q) \neq \text{cl}(p)$. We say q is a *c-approximate nearest foreign neighbor* of p in S (with respect to a metric $\text{dist} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$) if $\text{dist}(p, q) \leq c \cdot \text{dist}(p, q')$ for all foreign neighbors q' of p in S . The following lemma is a direct consequence of the dynamic approximate nearest neighbor data structure of Chan et al. [13]. We omit its proof here, which can be found in the full version of the paper.

► **Lemma 2.** *Given a set S of n colored points in \mathbb{R}^d , one can compute in $O_\varepsilon(n \log n)$ time a function $\text{ann} : S \rightarrow S$ which maps each point in S to a $(1 + \varepsilon)$ -approximate nearest foreign neighbor (with respect to the Euclidean distance) of that point in S . The algorithm generalizes to the L_p -norm for any fixed $p \geq 1$.*

Grids. A *d-dimensional grid* refers to a (infinite) set of axis-parallel hyperplanes that partition \mathbb{R}^d into same-sized axis-parallel hypercubes (called *grid cells* or simply *cells*). A d -dimensional grid can be characterized by a number $w > 0$ called the *cell-size* and a vector $(k_1, \dots, k_d) \in \mathbb{R}^d$ called the *offset*. Specifically, the grid with cell-size w and offset (k_1, \dots, k_d) , denoted by $\Gamma_w(k_1, \dots, k_d)$, consists of the hyperplanes whose equations are of the form $x_i = wt + k_i$ for $i \in [d]$ and $t \in \mathbb{Z}$. In other words, $\Gamma_w(k_1, \dots, k_d)$ is the (unique) grid in which the cells are hypercubes of side-length w and (k_1, \dots, k_d) is a grid point (i.e., a vertex of a cell). Figure 1 presents the 2-dimensional grid $\Gamma_3(1, 2)$.

	(1, 8)	(4, 8)	(7, 8)	(10, 8)
	(1, 5)	(4, 5)	(7, 5)	(10, 5)
	(1, 2)	(4, 2)	(7, 2)	(10, 2)

■ **Figure 1** The 2-dimensional grid $\Gamma_3(1, 2)$. The numbers are the coordinates of the grid points.

For a set S of points in \mathbb{R}^d , a (d -dimensional) grid naturally induces a partition of S in which each part consists of the points in one grid cell. To guarantee that every point in \mathbb{R}^d belongs to exactly one grid cell, we define the cells as hypercubes that are closed on the lower side and open on the higher side. Formally, every grid cell of $\Gamma_w(k_1, \dots, k_d)$ is a hypercube $\prod_{i=1}^d [t_i w + k_i, (t_i + 1)w + k_i)$ where $(t_1, \dots, t_d) \in \mathbb{Z}^d$.

3 The approximation scheme

We present our approximation scheme under the Euclidean norm. Its extension to any L_p -norm is straightforward, and will be briefly discussed in Section 3.4.

Let S be a set of n colored points in \mathbb{R}^d . For each point p , let $\text{nn}(p) \in S$ be the nearest foreign neighbor of p in S (with respect to the Euclidean distance). For simplicity of exposition, we shall first present our algorithm under the assumption that $\text{nn}(p)$ for every $p \in S$ is known to us. It is unlikely to compute all nearest foreign neighbors in near-linear time for $d \geq 3$ [19]. However, as we only want an approximation algorithm, it turns out that knowing *approximate* nearest foreign neighbors (which can be computed efficiently using Lemma 2) is already sufficient. We shall discuss this in Section 3.4.

We define $\phi(p)$ as the distance between p and $\text{nn}(p)$ for all $p \in S$. The feasible solutions for geometric many-to-many matching on S are subsets $E \subseteq E(S)$ satisfying $V(E) = S$. It is more convenient to consider an equivalent formulation of the problem where all subsets of $E(S)$ are feasible solutions, which we call the *penalized* formulation. In this formulation, we allow the solution to not cover all points in S , but for every uncovered point $p \in S$, there is a penalty of $\phi(p)$ added to the cost of the solution. Formally, the cost of a solution $E \subseteq E(S)$ is $\sum_{e \in E} |e| + \sum_{p \in S \setminus V(E)} \phi(p)$, where $|e|$ denotes the *length* of e , i.e., the Euclidean distance between the endpoints of e . To see this formulation is equivalent to the original one, let opt (resp., opt') be the optimum of the original (resp., penalized) formulation of the problem. Clearly, a solution of the original formulation is also a solution of the penalized formulation with the same cost, which implies $\text{opt} \geq \text{opt}'$. The following lemma shows $\text{opt} \leq \text{opt}'$.

► **Lemma 3.** *Given a subset $E \subseteq E(S)$, one can compute in $O(n + |E|)$ time another subset $E' \subseteq E(S)$ such that $V(E') = S$ and $\sum_{e \in E'} |e| \leq \sum_{e \in E} |e| + \sum_{p \in S \setminus V(E)} \phi(p)$.*

Proof. We simply set $E' = E \cup \{(p, \text{nn}(p)) : p \in S \setminus V(E)\}$, which can be computed in $O(n + |E|)$ time (given the nearest foreign neighbors). Clearly, $V(E') = S$. For an edge $e = (p, \text{nn}(p))$, we have $|e| = \phi(p)$. Thus, $\sum_{e \in E'} |e| \leq \sum_{e \in E} |e| + \sum_{p \in S \setminus V(E)} \phi(p)$. ◀

As $\text{opt} = \text{opt}'$, a c -approximation solution for the original formulation is also a c -approximation solution for the penalized formulation. On the other hand, Lemma 3 shows that given a c -approximation solution for the penalized formulation, one can compute in linear time a c -approximation solution for the original formulation. Thus, the two formulations are equivalent, and it suffices to solve the penalized formulation.

We sort the points in S by their coordinates in every dimension, and also by their ϕ -values. The benefit of considering the penalized formulation is that it allows us to properly define *subproblems*. Specifically, for $R \subseteq S$ and $E \subseteq E(R)$, we write

$$\text{cost}_R(E) = \sum_{e \in E} |e| + \sum_{p \in R \setminus V(E)} \phi(p).$$

We define a subproblem $\mathbf{Prob}(R)$ for every $R \subseteq S$, which aims to compute $E \subseteq E(R)$ that minimizes $\text{cost}_R(E)$. Note that $\mathbf{Prob}(R)$ is not exactly equivalent to the (penalized) geometric many-to-many matching problem on R , as the penalty $\phi(p)$ for $p \in R$ is defined by the nearest foreign neighbor of p in S (rather than R). We observe the following simple fact.

► **Fact 4.** *Let $R \subseteq S$ and $E^* \subseteq E(R)$ be an optimal solution of $\mathbf{Prob}(R)$. Then any edge $e = (p, q) \in E^*$ satisfies $|e| \leq \phi(p) + \phi(q)$.*

Proof. Assume there exists $e = (p, q) \in E^*$ with $|e| > \phi(p) + \phi(q)$. Then $\text{cost}_R(E^* \setminus \{e\}) \leq \text{cost}_R(E^*) - |e| + \phi(p) + \phi(q) < \text{cost}_R(E^*)$, contradicting the optimality of E^* . ◀

Clearly, our final goal is to compute a $(1+\varepsilon)$ -approximation solution for $\mathbf{Prob}(S)$. Without loss of generality, assume $\varepsilon \leq 1$. For $R \subseteq S$, we denote by $\text{opt}(R) = \min_{E \subseteq E(R)} \text{cost}_R(E)$ as the optimum of subproblem $\mathbf{Prob}(R)$. Our algorithm first applies two reductions, which eventually reduce $\mathbf{Prob}(S)$ to certain well-structured subproblems. These reductions are presented in Sections 3.1 and 3.2. Then we use grid technique together with the FPT algorithm for integer linear programming to solve these subproblems, which is discussed in Section 3.3. Finally, we put everything together and prove Theorem 1 in Section 3.4.

3.1 First reduction

In the first step, we reduce the problem $\mathbf{Prob}(S)$ to subproblems $\mathbf{Prob}(R)$ where the points in R have similar ϕ -values. For $R \subseteq S$, let $\Delta_R = \max_{p \in R} \phi(p) / \min_{p \in R} \phi(p)$. Our goal in this section is to prove the following lemma.

► **Lemma 5.** *Suppose for every subset $R \subseteq S$ satisfying $\Delta_R \leq (\frac{3}{\varepsilon})^{\lceil \frac{3}{\varepsilon} \rceil}$, one can compute in $O_\varepsilon(|R|)$ time a $(1 + \frac{\varepsilon}{2})$ -approximation solution for $\mathbf{Prob}(R)$. Then one can compute in $O_\varepsilon(n)$ time a $(1 + \varepsilon)$ -approximation solution for $\mathbf{Prob}(S)$.*

The basic idea of this reduction is the following. We distinguish the edges in $E(S)$ as balanced edges and unbalanced edges. The balanced edges are those whose two endpoints have similar ϕ -values. It turns out that the unbalanced edges can be “ignored” almost for free. Regarding only the balanced edges, we can then apply Baker’s shifting technique [7] on the ϕ -values to decompose the problem. Below we discuss the reduction in detail.

For each point $p \in S$, we write $\phi'(p) = \log_{3/\varepsilon} \phi(p)$. We say an edge $e = (p, q) \in E(S)$ is *balanced* if $|\phi'(p) - \phi'(q)| \leq 1$, and *unbalanced* otherwise.

► **Observation 6.** *For any unbalanced edge $e = (p, q) \in E(S)$, $\phi(p) + \phi(q) \leq (1 + \frac{\varepsilon}{3}) \cdot |e|$.*

Proof. Without loss of generality, assume $\phi(p) \geq \phi(q)$. As e is unbalanced, $|\phi'(p) - \phi'(q)| > 1$ and thus $\phi(p) > \frac{3}{\varepsilon} \cdot \phi(q)$. Note that $|e| \geq \phi(p)$ by the definition of $\phi(p)$. Therefore, we have $\phi(p) + \phi(q) \leq (1 + \frac{\varepsilon}{3}) \cdot \phi(p) \leq (1 + \frac{\varepsilon}{3}) \cdot |e|$. ◀

Set $w = \lceil \frac{3}{\varepsilon} \rceil$. For each $i \in [w]$, we construct the (1-dimensional) grid $\Gamma_w(i)$ – recall that $\Gamma_w(i)$ is the grid with cell-size w and offset i . For an edge $e = (p, q) \in E(S)$, we define $I_e = \{i \in [w] : \phi'(p) \text{ and } \phi'(q) \text{ lie in the same cell of } \Gamma_w(i)\}$.

► **Observation 7.** *For any balanced edge $e \in E(S)$, $|I_e| \geq w - 1$.*

Proof. Suppose $e = (p, q)$. Since e is balanced, $|\phi'(p) - \phi'(q)| \leq 1$ and thus there exists at most one integer $i^* \in [\phi'(p), \phi'(q)]$. For any $i \in [w]$, $i \notin I_e$ iff i^* exists and i is congruent with i^* modulo w . Thus, $|I_e| = w$ if i^* does not exist and $|I_e| = w - 1$ if i^* exists. ◀

For each $i \in [w]$, the grid $\Gamma_w(i)$ induces a partition of the values in $\{\phi'(p) : p \in S\}$, which in turn induces a partition \mathcal{R}_i of S . In other words, \mathcal{R}_i partitions S into subsets each of which contains the points in S whose ϕ' -values lying in one cell of $\Gamma_w(i)$.

► **Observation 8.** *There exists $i \in [w]$ such that $\sum_{R \in \mathcal{R}_i} \text{opt}(R) \leq (1 + \frac{\varepsilon}{3}) \cdot \text{opt}(S)$.*

Proof. Let $E^* \subseteq E(S)$ be an optimal solution of $\mathbf{Prob}(S)$. Also, let $B^* \subseteq E^*$ and $U^* \subseteq E^*$ be the subsets consisting of balanced and unbalanced edges in E^* , respectively. For $i \in [w]$, define $N_i^* = \{e \in B^* : i \notin I_e\}$. By Observation 7, each balanced edge $e \in B^*$ is contained in at most one set N_i^* . Thus, we have $\sum_{i=1}^w \sum_{e \in N_i^*} |e| \leq \sum_{e \in B^*} |e|$, which implies the existence of an index $i \in [w]$ such that $\sum_{e \in N_i^*} |e| \leq (\sum_{e \in B^*} |e|) / w \leq \frac{\varepsilon}{3} \cdot (\sum_{e \in B^*} |e|)$. We show that $\sum_{R \in \mathcal{R}_i} \text{opt}(R) \leq (1 + \frac{\varepsilon}{3}) \cdot \text{opt}(S)$. Note that N_i^* contains exactly the balanced edges in E^* whose two endpoints belong to different sets in \mathcal{R}_i . Therefore, we have

$$\sum_{R \in \mathcal{R}_i} \text{opt}(R) \leq \sum_{R \in \mathcal{R}_i} \text{cost}_R(B^* \cap E(R)) \leq \text{cost}_S(E^*) - \sum_{e \in U^* \cup N_i^*} |e| + \sum_{p \in V(U^* \cup N_i^*)} \phi(p).$$

By Observation 6, $\sum_{p \in V(U^*)} \phi(p) - \sum_{e \in U^*} |e| \leq \frac{\varepsilon}{3} \sum_{e \in U^*} |e|$. Furthermore, every $e = (p, q) \in E(S)$ satisfies $|e| \geq \max\{\phi(p), \phi(q)\} \geq \frac{1}{2}(\phi(p) + \phi(q))$ by the definition of ϕ . Therefore,

$\sum_{e \in N_i^*} |e| \geq \frac{1}{2} \sum_{p \in V(N_i^*)} \phi(p)$. It follows that $\sum_{p \in V(N_i^*)} \phi(p) - \sum_{e \in N_i^*} |e| \leq \sum_{e \in N_i^*} |e| \leq \frac{2}{3} \sum_{e \in B^*} |e|$, where the second inequality follows from the choice of i . Now,

$$\begin{aligned} \sum_{R \in \mathcal{R}_i} \text{opt}(R) &\leq \text{cost}_S(E^*) - \sum_{e \in U^* \cup N_i^*} |e| + \sum_{p \in V(U^* \cup N_i^*)} \phi(p) \\ &\leq \text{cost}_S(E^*) + \left(\sum_{p \in V(U^*)} \phi(p) - \sum_{e \in U^*} |e| \right) + \left(\sum_{p \in V(N_i^*)} \phi(p) - \sum_{e \in N_i^*} |e| \right) \\ &\leq \text{cost}_S(E^*) + \frac{\varepsilon}{3} \sum_{e \in U^*} |e| + \frac{\varepsilon}{3} \sum_{e \in B^*} |e| \\ &\leq \left(1 + \frac{\varepsilon}{3} \right) \cdot \text{cost}_S(E^*), \end{aligned}$$

where the last inequality follows from the fact $\sum_{e \in U^*} |e| + \sum_{e \in B^*} |e| = \sum_{e \in E^*} |e| \leq \text{cost}_S(E^*)$. As $\text{cost}_S(E^*) = \text{opt}(S)$, we conclude that $\sum_{R \in \mathcal{R}_i} \text{opt}(R) \leq \left(1 + \frac{\varepsilon}{3} \right) \cdot \text{opt}(S)$. ◀

Using the above observation, we can now prove Lemma 5. We consider every $i \in [w]$ and compute the partition \mathcal{R}_i of S . Note that \mathcal{R}_i can be computed in $O(n)$ time as we sorted the points in S by their ϕ -values. For every $R \in \mathcal{R}_i$, the ϕ -values of the points in R differ by at most w and thus $\Delta_R \leq \left(\frac{3}{\varepsilon}\right)^w = \left(\frac{3}{\varepsilon}\right)^{\lceil \frac{3}{\varepsilon} \rceil}$. Therefore, by our assumption, for every $R \in \mathcal{R}_i$, we can compute in $O_\varepsilon(|R|)$ time a $\left(1 + \frac{\varepsilon}{2}\right)$ -approximation solution $E_R^* \subseteq E(R)$ for $\mathbf{Prob}(R)$. The union $E_i^* = \bigcup_{R \in \mathcal{R}_i} E_R^*$ is a solution of $\mathbf{Prob}(S)$ and $\text{cost}_S(E_i^*) = \sum_{R \in \mathcal{R}_i} \text{cost}_R(E_R^*) \leq \left(1 + \frac{\varepsilon}{2}\right) \sum_{R \in \mathcal{R}_i} \text{opt}(R)$. The total time for constructing E_i^* is $O_\varepsilon(n)$, because $\sum_{R \in \mathcal{R}_i} |R| = n$. We construct the solution E_i^* for all $i \in [w]$ and finally output the best one among them. Observation 8 guarantees the existence of $i \in [w]$ such that

$$\text{cost}_S(E_i^*) \leq \left(1 + \frac{\varepsilon}{2}\right) \sum_{R \in \mathcal{R}_i} \text{opt}(E_R^*) \leq \left(1 + \frac{\varepsilon}{2}\right) \left(1 + \frac{\varepsilon}{3}\right) \cdot \text{opt}(S) \leq (1 + \varepsilon) \cdot \text{opt}(S).$$

Therefore, our algorithm gives a $(1 + \varepsilon)$ -approximation solution for $\text{opt}(S)$. Since $w = O_\varepsilon(1)$, the total running time is still $O_\varepsilon(n)$. This completes the proof of Lemma 5.

3.2 Second reduction

In this section, we further reduce a subproblem $\mathbf{Prob}(R)$ with bounded Δ_R to subproblems $\mathbf{Prob}(Q)$ where Q has a small bounding box compared to the values $\phi(p)$ for $p \in Q$. For $Q \subseteq S$, let W_Q be the side-length of the smallest axis-parallel hypercube containing Q . Our goal in this section is to prove the following lemma.

► **Lemma 9.** *Suppose for every subset $Q \subseteq S$ satisfying $W_Q \leq 2 \lceil \frac{4d}{\varepsilon} \rceil \left(\frac{3}{\varepsilon}\right)^{\lceil \frac{3}{\varepsilon} \rceil} \cdot \min_{p \in Q} \phi(p)$, one can compute in $O_\varepsilon(|Q|)$ time a $\left(1 + \frac{\varepsilon}{5}\right)$ -approximation solution for $\mathbf{Prob}(Q)$. Then for every subset $R \subseteq S$ satisfying $\Delta_R \leq \left(\frac{3}{\varepsilon}\right)^{\lceil \frac{3}{\varepsilon} \rceil}$, one can compute in $O_\varepsilon(|R|)$ time a $\left(1 + \frac{\varepsilon}{2}\right)$ -approximation solution for $\mathbf{Prob}(R)$.*

This reduction is done again by a shifting technique. But this time, we apply grid shifting to the space \mathbb{R}^d . As the points in R have similar ϕ -values, the edges in an optimal solution of $\mathbf{Prob}(R)$ also have similar lengths by Fact 4. This nice property allows us to use a shifted grid to decompose the problem by losing a factor of $1 + O(\varepsilon)$ in cost.

Consider a subset $R \subseteq S$ satisfying $\Delta_R \leq \left(\frac{3}{\varepsilon}\right)^{\lceil \frac{3}{\varepsilon} \rceil}$. Let $E^* \subseteq E(R)$ be an optimal solution of $\mathbf{Prob}(R)$. Set $r = \lceil \frac{4d}{\varepsilon} \rceil$, $\phi = \max_{p \in R} \phi(p)$, and $w = r \cdot 2\phi$. We say an edge $e \in E^*$ is compatible with a tuple $(k_1, \dots, k_d) \in [r]^d$ if the two endpoints of e lie in the same cell of the d -dimensional grid $\Gamma_w(k_1 \cdot 2\phi, \dots, k_d \cdot 2\phi)$.

► **Observation 10.** Every $e \in E^*$ is compatible with at least $(r-1)^d$ tuples in $[r]^d$.

Proof. Let $e \in E^*$. As $\phi = \max_{p \in R} \phi(p)$, by Fact 4, we have $|e| \leq 2\phi$. Thus, for every $i \in [d]$, there exists at most one integer $k_i^* \in \mathbb{N}$ such that the hyperplane $x_i = k_i^* \cdot 2\phi$ separates the two endpoints of e . Note that e is compatible with a tuple $(k_1, \dots, k_d) \in [r]^d$ iff for every $i \in [d]$ such that k_i^* exists, k_i and k_i^* are not congruent modulo r . For $i \in [d]$, define $K_i = [r]$ if k_i^* does not exist and $K_i = \{k \in [r] : k \text{ is not congruent with } k_i^* \text{ modulo } r\}$ if k_i^* exists. Then the tuples which e is compatible with are exactly those in $\prod_{i=1}^d K_i$. We have $|K_i| \geq r-1$ for all $i \in [d]$. Therefore, $|\prod_{i=1}^d K_i| \geq (r-1)^d$. ◀

For a tuple $\sigma = (k_1, \dots, k_d) \in [r]^d$, we denote by \mathcal{Q}_σ the partition of R induced by the grid $\Gamma_w(k_1 \cdot 2\phi, \dots, k_d \cdot 2\phi)$.

► **Observation 11.** There exists a tuple $\sigma \in [r]^d$ such that $\sum_{Q \in \mathcal{Q}_\sigma} \text{opt}(Q) \leq (1 + \frac{\varepsilon}{4}) \cdot \text{opt}(R)$.

Proof. For $\sigma \in [r]^d$, denote by $N_\sigma^* = \{e \in E^* : e \text{ is not compatible with } \sigma\}$. By Observation 10, each $e \in E^*$ is contained in at most $r^d - (r-1)^d$ sets N_σ^* . Thus, we have

$$\sum_{\sigma \in [r]^d} \sum_{e \in N_\sigma^*} |e| \leq (r^d - (r-1)^d) \cdot \sum_{e \in E^*} |e| \leq (r^d - (r-1)^d) \cdot \text{cost}_R(E^*).$$

So there exists some $\sigma \in [r]^d$ such that $\sum_{e \in N_\sigma^*} |e| \leq (r^d - (r-1)^d)/r^d \cdot \text{cost}_R(E^*)$. Note that $(r^d - (r-1)^d)/r^d \leq dr^{d-1}/r^d = d/r \leq \frac{\varepsilon}{4}$, which implies $\sum_{e \in N_\sigma^*} |e| \leq \frac{\varepsilon}{4} \cdot \text{cost}_R(E^*)$. We show that σ satisfies the condition $\sum_{Q \in \mathcal{Q}_\sigma} \text{opt}(Q) \leq (1 + \frac{\varepsilon}{4}) \cdot \text{opt}(R)$. Clearly,

$$\sum_{Q \in \mathcal{Q}_\sigma} \text{opt}(Q) \leq \sum_{Q \in \mathcal{Q}_\sigma} \text{cost}_Q(E^* \cap E(Q)) \leq \text{cost}_R(E^*) - \sum_{e \in N_\sigma^*} |e| + \sum_{p \in V(N_\sigma^*)} \phi(p).$$

Every edge $e = (p, q) \in E(S)$ satisfies $|e| \geq \max\{\phi(p), \phi(q)\} \geq \frac{1}{2}(\phi(p) + \phi(q))$ by the definition of ϕ . Therefore, $\sum_{e \in N_\sigma^*} |e| \geq \frac{1}{2} \sum_{p \in V(N_\sigma^*)} \phi(p)$. It follows that

$$\text{cost}_R(E^*) - \sum_{e \in N_\sigma^*} |e| + \sum_{p \in V(N_\sigma^*)} \phi(p) \leq \text{cost}_R(E^*) + \sum_{e \in N_\sigma^*} |e|,$$

which implies $\sum_{Q \in \mathcal{Q}_\sigma} \text{opt}(Q) \leq \text{cost}_R(E^*) + \sum_{e \in N_\sigma^*} |e|$. Since $\sum_{e \in N_\sigma^*} |e| \leq \frac{\varepsilon}{4} \cdot \text{cost}_R(E^*)$, we finally have $\sum_{Q \in \mathcal{Q}_\sigma} \text{opt}(Q) \leq (1 + \frac{\varepsilon}{4}) \cdot \text{cost}_R(E^*) = (1 + \frac{\varepsilon}{4}) \cdot \text{opt}(R)$. ◀

Using the above observation, we can now prove Lemma 9. We consider every tuple $\sigma \in [r]^d$. For each $\sigma \in [r]^d$, we first compute the partition \mathcal{Q}_σ of R , which can be done in $O(|R|)$ time as the points in R are sorted in every dimension. For every $Q \in \mathcal{Q}_\sigma$,

$$W_Q \leq w = 2r\phi \leq 2 \left\lceil \frac{4d}{\varepsilon} \right\rceil \cdot \max_{p \in R} \phi(p) = 2 \left\lceil \frac{4d}{\varepsilon} \right\rceil \cdot \Delta_R \min_{p \in R} \phi(p),$$

and thus $W_Q \leq 2 \lceil \frac{4d}{\varepsilon} \rceil (\frac{3}{\varepsilon})^{\lceil \frac{3}{\varepsilon} \rceil} \cdot \min_{p \in Q} \phi(p)$ as $\Delta_R \leq (\frac{3}{\varepsilon})^{\lceil \frac{3}{\varepsilon} \rceil}$ and $\min_{p \in R} \phi(p) \leq \min_{p \in Q} \phi(p)$. Therefore, by our assumption, for every $Q \in \mathcal{Q}_\sigma$, we can compute in $O_\varepsilon(|Q|)$ time a $(1 + \frac{\varepsilon}{5})$ -approximation solution $E_Q^* \subseteq E(Q)$ for $\mathbf{Prob}(Q)$. The union $E_\sigma^* = \bigcup_{Q \in \mathcal{Q}_\sigma} E_Q^*$ is a solution of $\mathbf{Prob}(R)$ and $\text{cost}_R(E_\sigma^*) = \sum_{Q \in \mathcal{Q}_\sigma} \text{cost}_Q(E_Q^*) \leq (1 + \frac{\varepsilon}{5}) \sum_{Q \in \mathcal{Q}_\sigma} \text{opt}(Q)$. The total time for constructing E_σ^* is $O_\varepsilon(|R|)$, because $\sum_{Q \in \mathcal{Q}_\sigma} |Q| = |R|$. We construct the solution E_σ^* for all $\sigma \in [r]^d$ and finally output the best one among them. Observation 11 guarantees the existence of $\sigma \in [r]^d$ such that

$$\text{cost}_R(E_\sigma^*) \leq \left(1 + \frac{\varepsilon}{5}\right) \sum_{Q \in \mathcal{Q}_\sigma} \text{opt}(Q) \leq \left(1 + \frac{\varepsilon}{5}\right) \left(1 + \frac{\varepsilon}{4}\right) \cdot \text{opt}(R) \leq \left(1 + \frac{\varepsilon}{2}\right) \cdot \text{opt}(R).$$

Therefore, our algorithm gives a $(1 + \frac{\varepsilon}{2})$ -approximation solution for $\text{opt}(R)$. Since $r^d = O_\varepsilon(1)$, the total running time is still $O_\varepsilon(|R|)$. This completes the proof of Lemma 9.

3.3 Solving a well-structured subproblem

By the reductions of Lemmas 5 and 9, it now suffices to show that for every $Q \subseteq S$ satisfying $W_Q \leq 2 \lceil \frac{4d}{\varepsilon} \rceil (\frac{3}{\varepsilon})^{\lceil \frac{3}{\varepsilon} \rceil} \cdot \min_{p \in Q} \phi(p)$, one can compute in $O_\varepsilon(|Q|)$ time a $(1 + \frac{\varepsilon}{5})$ -approximation solution for **Prob**(Q). In other words, our goal is to prove the following lemma.

► **Lemma 12.** *For every subset $Q \subseteq S$ satisfying $W_Q \leq 2 \lceil \frac{4d}{\varepsilon} \rceil (\frac{3}{\varepsilon})^{\lceil \frac{3}{\varepsilon} \rceil} \cdot \min_{p \in Q} \phi(p)$, one can compute in $O_\varepsilon(|Q|)$ time a $(1 + \frac{\varepsilon}{5})$ -approximation solution for **Prob**(Q).*

The basic idea to solve such a subproblem **Prob**(Q) is the following. Using the fact $\phi(p) = \Omega_\varepsilon(|W_Q|)$ for all $p \in Q$, we can partition the bounding box of Q into $O_\varepsilon(1)$ small hypercubes such that the points in each hypercube have the same color and very similar ϕ -values. This allows us to view the points in a small hypercube as a single “point” and formulate an integer linear program with $O_\varepsilon(1)$ variables. We then apply the FPT algorithm for integer linear programming [16] to solve the problem. Below we discuss this in detail.

If all points in Q have the same color, then $E(Q) = \emptyset$ and the subproblem **Prob**(Q) is trivial. So assume Q contains points of at least two colors. Let \square be a hypercube containing Q with side-length $w = 2 \lceil \frac{4d}{\varepsilon} \rceil (\frac{3}{\varepsilon})^{\lceil \frac{3}{\varepsilon} \rceil} \cdot \min_{p \in Q} \phi(p)$. We set $r = 2d \lceil \frac{4d}{\varepsilon} \rceil (\frac{3}{\varepsilon})^{\lceil \frac{3}{\varepsilon} \rceil} \lceil \frac{4d}{\varepsilon} \rceil$ and evenly partition \square into r^d smaller hypercubes with side-length $\frac{w}{r}$. Let \mathcal{C} be the set of the smaller hypercubes which contain at least one point in Q . We have the following observation.

► **Observation 13.** *For every $C \in \mathcal{C}$, all points in $Q \cap C$ have the same color.*

Proof. Equivalently, we show no edge in $E(Q)$ has both endpoints in C . Assume $e = (p, q)$ with $p, q \in Q \cap C$. Then $|e| \leq \frac{dw}{r} < \phi(p)$, contradicting the fact that $|e| \geq \phi(p)$. ◀

By the above observation, for each $C \in \mathcal{C}$, we can define the *color* of C , denoted by $\text{cl}(C)$, as the color of the points in $Q \cap C$. For a subset $E \subseteq E(Q)$, we define a function $f_E : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{N}$, where $f_E(C, C')$ is equal to the number of edges in E whose one endpoint is in C and the other endpoint is in C' . Also, we define a function $g_E : \mathcal{C} \rightarrow \mathbb{N}$, where $g_E(C) = |Q \cap C| - |V(E) \cap C|$ is the number of points in C not matched by E .

► **Observation 14.** *Let $f : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{N}$ and $g : \mathcal{C} \rightarrow \mathbb{N}$ be two functions. If there exists $E \subseteq E(Q)$ such that $f_E = f$ and $g_E = g$, then the following conditions hold.*

1. *For any $C, C' \in \mathcal{C}$, we have $f(C, C') = f(C', C) = 0$ if $\text{cl}(C) = \text{cl}(C')$, and $f(C, C') = f(C', C) \leq |Q \cap C| \cdot |Q \cap C'|$ if $\text{cl}(C) \neq \text{cl}(C')$.*
2. *For any $C \in \mathcal{C}$, we have $g(C) + \sum_{C' \in \mathcal{C}} f(C, C') \geq |Q \cap C|$.*

Conversely, if f and g satisfy the conditions, then one can compute $E \subseteq E(Q)$ in $O_\varepsilon(|Q|)$ time such that $f_E(C, C') \leq f(C, C')$ for all $(C, C') \in \mathcal{C} \times \mathcal{C}$ and $g_E(C) \leq g(C)$ for all $C \in \mathcal{C}$.

Proof. Suppose $f_E = f$ and $g_E = g$ for some $E \subseteq E(Q)$. By the definition of f_E , it is clear that $f(C, C') = f(C', C)$ for any $C, C' \in \mathcal{C}$. For any $C, C' \in \mathcal{C}$ with $\text{cl}(C) = \text{cl}(C')$, there cannot be any edge in E with one endpoint is in C and the other endpoint is in C' , and thus $f(C, C') = 0$. For any $C, C' \in \mathcal{C}$ with $\text{cl}(C) \neq \text{cl}(C')$, there can be at most $|Q \cap C| \cdot |Q \cap C'|$ edges in E with one endpoint is in C and the other endpoint is in C' , and thus $f(C, C') \leq |Q \cap C| \cdot |Q \cap C'|$. So condition 1 holds. To see condition 2, observe that for any $C \in \mathcal{C}$, $|V(E) \cap C| \leq \sum_{C' \in \mathcal{C}} f_E(C, C')$. Thus, by the definition of g_E , we directly have $g(C) + \sum_{C' \in \mathcal{C}} f(C, C') \geq g(C) + |V(E) \cap C| = |Q \cap C|$.

Now suppose f and g satisfy the two conditions. For $(C, C') \in \mathcal{C} \times \mathcal{C}$ with $f(C, C') > |Q|$, we set $f(C, C') = |Q|$. After this change, f and g still satisfy the conditions. We construct the desired subset $E \subseteq E(Q)$ as follows. Initially, set $E = \emptyset$. For every $C, C' \in \mathcal{C}$ with

12:10 Approximation for Geometric Many-To-Many Matching

$\text{cl}(C) \neq \text{cl}(C')$, we shall pick $f(C, C')$ edges in $E(Q)$ with one endpoint is in C and the other endpoint is in C' , and add them to E ; this is possible since $f(C, C') \leq |Q \cap C| \cdot |Q \cap C'|$ by condition 1. The resulting E guarantees $f_E = f$ (and thus the f_E -values are smaller than or equal to the original f -values). To further guarantee $g_E(C) \leq g(C)$ for all $C \in \mathcal{C}$, we use the following rule to pick edges. When we want to pick an edge e with one endpoint in C and the other endpoint in C' (and add it to E), we always choose the endpoints of e among the points in $Q \cap C$ and $Q \cap C'$ that are not the endpoints of the edges in the current E ; if no such points exist, we then choose the endpoints of e arbitrarily. In this way, we can guarantee that at the end, for every $C \in \mathcal{C}$, either $|V(E) \cap C| = \sum_{C' \in \mathcal{C}} f(C, C')$ or $Q \cap C = V(E) \cap C$. In the former case, $g_E(C) = |Q \cap C| - |V(E) \cap C| = |Q \cap C| - \sum_{C' \in \mathcal{C}} f(C, C') \leq g(C)$ by condition 2. In the latter case, $g_E(C) = 0 \leq g(C)$. This simple construction of E can be done in $O_\varepsilon(|Q|)$ time since $f(C, C') \leq |Q|$ for all $(C, C') \in \mathcal{C} \times \mathcal{C}$. \blacktriangleleft

For convenience, we write $\phi(C) = \min_{p \in Q \cap C} \phi(p)$ for $C \in \mathcal{C}$. For $C, C' \in \mathcal{C}$, let $\text{dist}(C, C')$ denote the Euclidean distance between the centers of the hypercubes C and C' . Now for functions $f : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{N}$ and $g : \mathcal{C} \rightarrow \mathbb{N}$, we define

$$\text{cost}(f, g) = \frac{1}{2} \sum_{(C, C') \in \mathcal{C} \times \mathcal{C}} (f(C, C') \cdot \text{dist}(C, C')) + \sum_{C \in \mathcal{C}} (g(C) \cdot \phi(C)).$$

\blacktriangleright **Observation 15.** *For every subset $E \subseteq E(Q)$, we have*

$$\left(1 - \frac{\varepsilon}{21}\right) \text{cost}(f_E, g_E) \leq \text{cost}_Q(E) \leq \left(1 + \frac{\varepsilon}{21}\right) \text{cost}(f_E, g_E).$$

In particular, for any two subsets $E, E' \subseteq E(Q)$ satisfying $\text{cost}(f_E, g_E) \leq c \cdot \text{cost}(f_{E'}, g_{E'})$, we have $\text{cost}_Q(E) \leq (1 + \frac{\varepsilon}{10})c \cdot \text{cost}_Q(E')$.

Proof. We first show for any $e \in E(Q)$, $(1 - \frac{\varepsilon}{21}) \cdot \text{dist}(C, C') \leq |e| \leq (1 + \frac{\varepsilon}{21}) \cdot \text{dist}(C, C')$, where $C, C' \in \mathcal{C}$ are the hypercubes containing the two endpoints of e . Suppose $e = (p, q)$ where $p \in C$ and $q \in C'$. We have $|e| \geq \max\{\phi(p), \phi(q)\} \geq \frac{w}{r} \cdot d \lceil \frac{44}{\varepsilon} \rceil$. Observe that the distance between any two points in C (or C') is at most $\frac{w}{r} \cdot d$. Thus, the difference between $|e|$ and $\text{dist}(C, C')$ is at most $(2/\lceil \frac{44}{\varepsilon} \rceil) \cdot |e| \leq \frac{\varepsilon}{22} \cdot |e|$, which implies $\text{dist}(C, C') \geq \frac{21}{22} \cdot |e|$. It follows that the difference between $|e|$ and $\text{dist}(C, C')$ is at most $\frac{\varepsilon}{21} \cdot \text{dist}(C, C')$. So we have $(1 - \frac{\varepsilon}{21}) \cdot \text{dist}(C, C') \leq |e| \leq (1 + \frac{\varepsilon}{21}) \cdot \text{dist}(C, C')$.

Next, we show that for any $p \in Q$, $\phi(C) \leq \phi(p) \leq (1 + \frac{\varepsilon}{21}) \cdot \phi(C)$, where $C \in \mathcal{C}$ is the hypercube containing p . The inequality $\phi(C) \leq \phi(p)$ follows from the definition of $\phi(C)$. To see $\phi(p) \leq (1 + \frac{\varepsilon}{21}) \cdot \phi(C)$, suppose $\phi(C) = \phi(q)$ for $q \in Q \cap C$. By Observation 13, $\text{cl}(p) = \text{cl}(q)$. This implies $|\phi(p) - \phi(q)| \leq d \cdot \frac{w}{r}$, since the side-length of C is $\frac{w}{r}$. Note that $d \cdot \frac{w}{r} \leq \frac{\varepsilon}{44} \phi(q) \leq \frac{\varepsilon}{21} \phi(q)$. Therefore, $\phi(p) \leq (1 + \frac{\varepsilon}{21}) \cdot \phi(q)$.

Now we prove the observation. For each $p \in Q$, we denote by $C_p \in \mathcal{C}$ the hypercube containing p . Let $E = \{e_1, \dots, e_m\} \subseteq E(Q)$ and suppose $e_i = (p_i, q_i)$ for $i \in [m]$. Then we have $\text{cost}_Q(E) = \sum_{i=1}^m |e_i| + \sum_{p \in Q \setminus V(E)} \phi(p)$. As shown above, $(1 - \frac{\varepsilon}{21}) \cdot \text{dist}(C_{p_i}, C_{q_i}) \leq |e_i| \leq (1 + \frac{\varepsilon}{21}) \cdot \text{dist}(C_{p_i}, C_{q_i})$ for all $i \in [m]$ and $\phi(C_p) \leq \phi(p) \leq (1 + \frac{\varepsilon}{21}) \cdot \phi(C_p)$ for all $p \in Q \setminus V(E)$. Therefore, if we set $\alpha = \sum_{i=1}^m \text{dist}(C_{p_i}, C_{q_i}) + \sum_{p \in Q \setminus V(E)} \phi(C_p)$, then we have $(1 - \frac{\varepsilon}{21}) \cdot \alpha \leq \text{cost}_Q(E) \leq (1 + \frac{\varepsilon}{21}) \cdot \alpha$. It suffices to show $\alpha = \text{cost}(f_E, g_E)$. By the definitions of the functions f_E and g_E , we have

$$\begin{aligned} \alpha &= \frac{1}{2} \sum_{i=1}^m (\text{dist}(C_{p_i}, C_{q_i}) + \text{dist}(C_{q_i}, C_{p_i})) + \sum_{p \in Q \setminus V(E)} \phi(C_p) \\ &= \frac{1}{2} \sum_{(C, C') \in \mathcal{C} \times \mathcal{C}} (f_E(C, C') \cdot \text{dist}(C, C')) + \sum_{C \in \mathcal{C}} (g_E(C) \cdot \phi(C)). \end{aligned}$$

Thus, $\alpha = \text{cost}(f_E, g_E)$. To see the second statement in the observation, let $E, E' \subseteq E(Q)$ satisfying $\text{cost}(f_E, g_E) \leq c \cdot \text{cost}(f_{E'}, g_{E'})$. Write $\alpha = \text{cost}(f_E, g_E)$ and $\alpha' = \text{cost}(f_{E'}, g_{E'})$ for convenience. It follows that

$$\text{cost}_Q(E) \leq \left(1 + \frac{\varepsilon}{21}\right) \cdot \alpha \leq \left(1 + \frac{\varepsilon}{21}\right) c \cdot \alpha' \leq \frac{(1 + \frac{\varepsilon}{21})c}{1 - \frac{\varepsilon}{21}} \cdot \text{cost}_Q(E').$$

As $(1 + \frac{\varepsilon}{21})/(1 - \frac{\varepsilon}{21}) \leq 1 + \frac{\varepsilon}{10}$, we have $\text{cost}_Q(E) \leq (1 + \frac{\varepsilon}{10})c \cdot \text{cost}_Q(E')$. ◀

Thanks to Observation 14 and 15, we can use the following idea to compute a $(1 + \frac{\varepsilon}{5})$ -approximation solution for **Prob**(Q). We say a pair (f, g) of functions $f : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{N}$ and $g : \mathcal{C} \rightarrow \mathbb{N}$ is *valid* if it satisfies the two conditions in Observation 14. First, we compute a valid pair (f, g) satisfying $\text{cost}(f, g) \leq (1 + \frac{\varepsilon}{11}) \cdot \text{cost}(f', g')$ for any valid pair (f', g') . We shall show later how to efficiently compute such a pair (f, g) by solving an integer linear program. Then we use Observation 14 to compute in $O_\varepsilon(|Q|)$ time a subset $E \subseteq E(Q)$ satisfying $f_E(C, C') \leq f(C, C')$ for all $(C, C') \in \mathcal{C} \times \mathcal{C}$ and $g_E(C) \leq g(C)$ for all $C \in \mathcal{C}$. We claim that E is a $(1 + \frac{\varepsilon}{5})$ -approximation solution of **Prob**(Q). To see this, suppose $E^* \subseteq E(Q)$ is an optimal solution of **Prob**(Q). The pair (f_{E^*}, g_{E^*}) is valid by Observation 14. So we have $\text{cost}(f_E, g_E) \leq \text{cost}(f, g) \leq (1 + \frac{\varepsilon}{11}) \cdot \text{cost}(f_{E^*}, g_{E^*})$. Then by Observation 15,

$$\text{cost}_Q(E) \leq \left(1 + \frac{\varepsilon}{10}\right) \left(1 + \frac{\varepsilon}{11}\right) \cdot \text{cost}_Q(E^*) \leq \left(1 + \frac{\varepsilon}{5}\right) \cdot \text{cost}_Q(E^*).$$

Now we show how to compute a valid pair (f, g) satisfying $\text{cost}(f, g) \leq (1 + \frac{\varepsilon}{11}) \cdot \text{cost}(f', g')$ for any valid pair (f', g') . This is done by formulating an integer linear program with $O_\varepsilon(1)$ variables and applying the FPT algorithm for integer linear programming [16]. We view the values $f(C, C')$ for $(C, C') \in \mathcal{C} \times \mathcal{C}$ and $g(C)$ for $C \in \mathcal{C}$ as integer variables. The objective function to be minimized is $\text{cost}(f, g)$, which is a linear function of the variables. To check if (f, g) is valid, it is equivalent to check if (f, g) satisfies the two conditions in Observation 14, which can be described as linear constraints on the variables. Therefore, finding a valid pair (f, g) with minimum $\text{cost}(f, g)$ is equivalent to assigning (non-negative) integer values to the variables to minimize the objective function under the linear constraints. Note that this is not exactly an integer linear program, because the coefficients of the objective function are real numbers (while the coefficients of the linear constraints are all integers). However, as we only need a valid pair (f, g) with *approximately* minimum $\text{cost}(f, g)$, we can round these real coefficients to integers without changing the program too much. Observe that for any distinct $C, C' \in \mathcal{C}$, $\text{dist}(C, C') \geq \frac{w}{r}$. Furthermore, $\phi(p) \geq \frac{w}{r}$ for all $p \in Q$ and thus $\phi(C) \geq \frac{w}{r}$ for all $C \in \mathcal{C}$. We replace every coefficient η in the objective function with a new coefficient $\lfloor \frac{11}{\varepsilon} \rfloor \cdot \frac{\eta}{w/r}$ and obtain a new objective function

$$\text{cost}'(f, g) = \sum_{(C, C') \in \mathcal{C} \times \mathcal{C}} \left(f(C, C') \cdot \left\lfloor \frac{11}{\varepsilon} \right\rfloor \cdot \frac{\text{dist}(C, C')}{w/r} \right) + \sum_{C \in \mathcal{C}} \left(g(C) \cdot \left\lfloor \frac{11}{\varepsilon} \right\rfloor \cdot \frac{\phi(C)}{w/r} \right).$$

► **Observation 16.** *Let (f, g) be a valid pair with minimum $\text{cost}'(f, g)$. Then $\text{cost}(f, g) \leq (1 + \frac{\varepsilon}{11}) \cdot \text{cost}(f', g')$ for any valid pair (f', g') .*

Proof. We first claim that $\lfloor \frac{11}{\varepsilon} \rfloor \cdot \frac{\eta}{w/r} \leq (1 + \frac{\varepsilon}{11}) \cdot \lfloor \frac{11}{\varepsilon} \rfloor \cdot \frac{\eta}{w/r}$ for any coefficient η of the cost function. As observed before, $\eta \geq \frac{w}{r}$ and thus $\lfloor \frac{11}{\varepsilon} \rfloor \cdot \frac{\eta}{w/r} \geq \lfloor \frac{11}{\varepsilon} \rfloor$. So we have

$$\left\lfloor \frac{11}{\varepsilon} \right\rfloor \cdot \frac{\eta}{w/r} \leq \left\lfloor \frac{11}{\varepsilon} \right\rfloor \cdot \frac{\eta}{w/r} + 1 \leq \left(1 + \frac{\varepsilon}{11}\right) \cdot \left\lfloor \frac{11}{\varepsilon} \right\rfloor \cdot \frac{\eta}{w/r}.$$

12:12 Approximation for Geometric Many-To-Many Matching

This implies that $\lceil \frac{11}{\varepsilon} \rceil \cdot \frac{\text{cost}(f,g)}{w/r} \leq (1 + \frac{\varepsilon}{11}) \cdot \text{cost}'(f,g)$. By the choice of (f,g) , it then follows that $\lceil \frac{11}{\varepsilon} \rceil \cdot \frac{\text{cost}(f,g)}{w/r} \leq (1 + \frac{\varepsilon}{11}) \cdot \text{cost}'(f',g')$ for any valid pair (f',g') . By the definition of cost' , we directly have $\text{cost}'(f',g') \leq \lceil \frac{11}{\varepsilon} \rceil \cdot \frac{\text{cost}(f',g')}{w/r}$. Therefore, for any valid pair (f',g') , $\lceil \frac{11}{\varepsilon} \rceil \cdot \frac{\text{cost}(f,g)}{w/r} \leq (1 + \frac{\varepsilon}{11}) \cdot \lceil \frac{11}{\varepsilon} \rceil \cdot \frac{\text{cost}(f',g')}{w/r}$ and thus $\text{cost}(f,g) \leq (1 + \frac{\varepsilon}{11}) \cdot \text{cost}(f',g')$. ◀

With the above observation, it suffices to minimize the new objective function $\text{cost}'(f,g)$, which is a linear function of the variables with integer coefficients. Therefore, our task becomes solving an integer linear program. The numbers of variables and constraints are both $O_\varepsilon(1)$. The coefficients in the linear constraints are bounded by $|Q|^{O(1)}$. To bound the coefficients in the objective function, observe that $\text{dist}(C,C') \leq dw$ for any $C,C' \in \mathcal{C}$ (because the side-length of \square is w) and thus $\frac{\text{dist}(C,C')}{w/r} = O_\varepsilon(1)$. Also, since Q contains points of at least two colors, every point $p \in Q$ has a foreign neighbor in Q (which has distance at most dw to p) and thus $\phi(p) \leq dw$. It follows that $\phi(C) \leq dw$ for all $C \in \mathcal{C}$ and thus $\frac{\phi(C)}{w/r} = O_\varepsilon(1)$. So the coefficients in the objective function are bounded by $O_\varepsilon(1)$. The entire integer linear program can be encoded in $O_\varepsilon(\log |Q|)$ bits. It is well-known [16] that an integer linear program encoded in L bits with s variables can be solved in $s^{O(s)} \cdot L^{O(1)}$ time. Therefore, our program can be solved in $O_\varepsilon(\log^{O(1)} |Q|)$ time. Including the time for constructing the program and finding the set E using Observation 14, the total time cost is $O_\varepsilon(|Q|)$, which proves Lemma 12.

3.4 Putting everything together

Combining Lemmas 5, 9, and 12, we can compute a $(1 + \varepsilon)$ -approximation solution for $\text{Prob}(S)$ in $O_\varepsilon(n)$ time. Including the time for pre-sorting, we see that geometric many-to-many matching in any fixed dimension admits a $(1 + \varepsilon)$ -approximation algorithm with running time $O_\varepsilon(n \log n)$, assuming the nearest foreign neighbors of the points are given.

Finally, we provide the last missing piece of our result: how to solve the problem without knowing the nearest foreign neighbors. While it is well-known that the all-nearest-neighbor problem can be solved in $O(n \log n)$ time in any fixed dimension [43], computing all nearest foreign neighbors is much more challenging: it admits an $O(n \log n)$ -time algorithm only for $d = 2$ [3] and has a conjectured $\Omega(n^{4/3})$ lower bound for $d \geq 3$ [19]. However, as we only want to compute an approximation solution, we can actually use $(1 + \varepsilon)$ -approximate nearest foreign neighbors (instead of the exact nearest foreign neighbors) which can be computed in $O_\varepsilon(n \log n)$ time by Lemma 2. For a point $p \in S$, let $\text{ann}(p) \in S$ be a $(1 + \varepsilon)$ -approximate nearest foreign neighbor of p in S . We use $\phi'(p) = \frac{\text{dist}(p, \text{ann}(p))}{1 + \varepsilon}$ as the penalty of a point $p \in S$, where dist denotes the Euclidean distance. Note that $\phi'(p) \leq \phi(p) \leq (1 + \varepsilon) \cdot \phi'(p)$. It is easy to see the following analogy of Lemma 3.

► **Lemma 17.** *Given a subset $E \subseteq E(S)$, one can compute in $O(n + |E|)$ time another subset $E' \subseteq E(S)$ such that $V(E') = S$ and $\sum_{e \in E'} |e| \leq (1 + \varepsilon) \cdot (\sum_{e \in E} |e| + \sum_{p \in S \setminus V(E)} \phi'(p))$.*

Proof. Like what we did in the proof of Lemma 3, we simply set $E' = E \cup \{(p, \text{ann}(p)) : p \in S \setminus V(E)\}$. Clearly, $V(E') = S$. For an edge $e = (p, \text{ann}(p))$, we have $|e| = (1 + \varepsilon) \cdot \phi(p)$. Thus, $\sum_{e \in E'} |e| \leq (1 + \varepsilon) \cdot (\sum_{e \in E} |e| + \sum_{p \in S \setminus V(E)} \phi'(p))$. ◀

Now it suffices to consider the penalized formulation with the new penalty function ϕ' . The same algorithm still works. Indeed, in our algorithm, we only use two properties of the old penalty function ϕ . First, we need $|e| \geq \max\{\phi(p), \phi(q)\}$ for every edge $e = (p, q) \in E(S)$. This also holds for ϕ' . Second, in the proof of Observation 15, we need the inequality that for any $p \in Q$, $\phi(C) \leq \phi(p) \leq (1 + \frac{\varepsilon}{21}) \cdot \phi(C)$, where $C \in \mathcal{C}$ is the hypercube containing p

and $\phi(C) = \min_{p \in Q \cap C} \phi(p)$. When replacing ϕ with ϕ' and setting $\phi'(C) = \min_{p \in Q \cap C} \phi'(p)$, we still have $\phi'(C) \leq \phi(p) \leq (1 + O(\varepsilon)) \cdot \phi'(C)$, because $\phi'(p) \leq \phi(p) \leq (1 + \varepsilon) \cdot \phi'(p)$. Therefore, applying the same algorithm with the new penalty function ϕ' , we can obtain a $(1 + O(\varepsilon))$ -approximation solution, which is sufficient for the purpose of an approximation scheme, since we can always choose ε to be smaller than the required approximation ratio by a constant factor. This completes the proof of Theorem 1 for the Euclidean case.

Generalization to L_p -norms. Our algorithm directly applies to geometric many-to-many matching under the L_p -norm for any $p \geq 1$. The only thing we need to adjust is the parameter r in Section 3.3: for different norms, we need to partition the hypercube \square into different numbers of smaller hypercubes to make Observations 13 and 15 hold.

4 Conclusion and future work

In this paper, we studied the geometric many-to-many matching problem. We give a $(1 + \varepsilon)$ -approximation algorithm with running time $O_\varepsilon(n \log n)$ for geometric many-to-many matching in any fixed dimension under the L_p -norm for any $p \geq 1$. Our result significantly improves and generalizes the previous work on the problem.

We pose two open questions for future study. First, the running time of our algorithm has an exponential dependency on $\frac{1}{\varepsilon}$, which comes from both the reduction in Section 3.1 and the FPT algorithm for integer linear programming used in Section 3.3. It is interesting to see whether one can improve the bound to $(\frac{1}{\varepsilon})^{O(1)} \cdot n \log n$. Second, as mentioned in the introduction, several variants of geometric many-to-many in \mathbb{R}^1 have been studied in the literature [34, 35, 36]. One can ask whether the techniques in this paper can be applied to obtain nontrivial results for these variants in higher dimensions.

References

- 1 Pankaj K Agarwal, Hsien-Chih Chang, Sharath Raghvendra, and Allen Xiao. Deterministic, near-linear ε -approximation algorithm for geometric bipartite matching. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 1052–1065, 2022.
- 2 Pankaj K Agarwal, Herbert Edelsbrunner, Otfried Schwarzkopf, and Emo Welzl. Euclidean minimum spanning trees and bichromatic closest pairs. *Discrete & Computational Geometry*, 6:407–422, 1991.
- 3 Alok Aggarwal, Herbert Edelsbrunner, Prahakar Raghavan, and Praseon Tiwari. Optimal time bounds for some proximity problems in the plane. *Information Processing Letters*, 42(1):55–60, 1992.
- 4 Sanjeev Arora. Polynomial time approximation schemes for Euclidean TSP and other geometric problems. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 2–11. IEEE, 1996.
- 5 Sanjeev Arora. Nearly linear time approximation schemes for Euclidean TSP and other geometric problems. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 554–563. IEEE, 1997.
- 6 Sunil Arya, Gautam Das, David M Mount, Jeffrey S Salowe, and Michiel Smid. Euclidean spanners: short, thin, and lanky. In *Proceedings of the 27th Annual ACM Symposium on Theory of Computing (STOC)*, pages 489–498, 1995.
- 7 Brenda S Baker. Approximation algorithms for NP-complete problems on planar graphs. *Journal of the ACM*, 41(1):153–180, 1994.

- 8 Sayan Bandyapadhyay, William Lochet, Daniel Lokshtanov, Saket Saurabh, and Jie Xue. Euclidean bottleneck steiner tree is fixed-parameter tractable. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 699–711. SIAM, 2024.
- 9 Sayan Bandyapadhyay, Anil Maheshwari, and Michiel Smid. Exact and approximation algorithms for many-to-many point matching in the plane. In *32nd International Symposium on Algorithms and Computation (ISAAC)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.
- 10 Hiba Belhadi, Karima Akli-Astouati, Youcef Djenouri, and Jerry Chun-Wei Lin. Data mining-based approach for ontology matching problem. *Applied Intelligence*, 50:1204–1221, 2020.
- 11 Amir Ben-Dor, Richard M Karp, Benno Schwikowski, and Ron Shamir. The restriction scaffold problem. In *Proceedings of the 6th Annual International Conference on Computational Biology*, pages 58–66, 2002.
- 12 Marcus Brazil, Ronald L Graham, Doreen A Thomas, and Martin Zachariasen. On the history of the Euclidean steiner tree problem. *Archive for history of exact sciences*, 68:327–354, 2014.
- 13 Timothy M Chan, Sarel Har-Peled, and Mitchell Jones. On locality-sensitive orderings and their applications. *SIAM Journal on Computing*, 49(3):583–600, 2020.
- 14 Justin Colannino, Mirela Damian, Ferran Hurtado, Stefan Langerman, Henk Meijer, Suneeta Ramaswami, Diane L. Souvaine, and Godfried Toussaint. Efficient many-to-many point matching in one dimension. *Graphs and Combinatorics*, 23:169–178, 2007.
- 15 Justin Colannino and Godfried Toussaint. Faster algorithms for computing distances between one-dimensional point sets. *Proceedings of the XI Encuentros de Geometria Computacional*, pages 189–198, 2005.
- 16 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- 17 Mark de Berg, Hans L Bodlaender, Sándor Kisfaludi-Bak, and Sudeshna Kolay. An ETH-tight exact algorithm for Euclidean TSP. In *Proceedings of the 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 450–461. IEEE, 2018.
- 18 Thomas Eiter and Heikki Mannila. Distance measures for point sets and their computation. *Acta informatica*, 34(2):109–133, 1997.
- 19 Jeff Erickson. On the relative complexities of some geometric problems. In *Proceedings of the 7th Canadian Conference on Computational Geometry (CCCG)*, volume 95, pages 85–90, 1995.
- 20 S. M. Ferdous, Alex Pothén, and Arif Khan. New approximation algorithms for minimum weighted edge cover. In Fredrik Manne, Peter Sanders, and Sivan Toledo, editors, *Proceedings of the 8th SIAM Workshop on Combinatorial Scientific Computing, CSC 2018, Bergen, Norway, June 6-8, 2018*, pages 97–108. SIAM, 2018.
- 21 Michael L Fredman and Robert Endre Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *Journal of the ACM*, 34(3):596–615, 1987.
- 22 Harold N Gabow and Robert E Tarjan. Faster scaling algorithms for network problems. *SIAM Journal on Computing*, 18(5):1013–1036, 1989.
- 23 Steven A Gallagher. *A 4/3-approximation for Minimum Weight Edge Cover*. PhD thesis, Purdue University, 2019.
- 24 Sarel Har-Peled, Piotr Indyk, and Anastasios Sidiropoulos. Euclidean spanners in high dimensions. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 804–809. SIAM, 2013.
- 25 Piotr Indyk. A near linear time constant factor approximation for Euclidean bichromatic matching (cost). In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, volume 7, pages 39–42. SIAM, 2007.
- 26 Haim Kaplan, Wolfgang Mulzer, Liam Roditty, Paul Seifert, and Micha Sharir. Dynamic planar voronoi diagrams for general distance functions and their algorithmic applications. *Discrete & Computational Geometry*, 64:838–904, 2020.
- 27 Judith Keijsper and Rudi Pendavingh. An efficient algorithm for minimum-weight bibranching. *Journal of Combinatorial Theory, Series B*, 73(2):130–145, 1998.

- 28 Harold W Kuhn. Variants of the Hungarian method for assignment problems. *Naval Research Logistics Quarterly*, 3(4):253–258, 1956.
- 29 Hung Le and Shay Solomon. Truly optimal Euclidean spanners. *SIAM Journal on Computing*, pages FOCS19–135 – FOCS19–199, 2022.
- 30 William B March, Parikshit Ram, and Alexander G Gray. Fast Euclidean minimum spanning tree: algorithm, analysis, and applications. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 603–612, 2010.
- 31 Yaodong Ni. Fuzzy minimum weight edge covering problem. *Applied mathematical modelling*, 32(7):1327–1337, 2008.
- 32 Robert Z Norman and Michael O Rabin. An algorithm for a minimum cover of a graph. *Proceedings of the American Mathematical Society*, 10(2):315–319, 1959.
- 33 Sharath Raghvendra and Pankaj K. Agarwal. A near-linear time ε -approximation algorithm for geometric bipartite matching. *Journal of the ACM*, 67(3):18:1–18:19, 2020.
- 34 Fatemeh Rajabi-Alni and Alireza Bagheri. An $O(n^2)$ algorithm for the limited-capacity many-to-many point matching in one dimension. *Algorithmica*, 76(2):381–400, 2016.
- 35 Fatemeh Rajabi-Alni and Alireza Bagheri. A fast and efficient algorithm for many-to-many matching of points with demands in one dimension. *arXiv preprint*, 2019. [arXiv:1904.05184](https://arxiv.org/abs/1904.05184).
- 36 Fatemeh Rajabi-Alni and Alireza Bagheri. A faster algorithm for the limited-capacity many-to-many point matching in one dimension. *arXiv preprint*, 2019. [arXiv:1904.03015](https://arxiv.org/abs/1904.03015).
- 37 Petar Ristoski, Petar Petrovski, Peter Mika, and Heiko Paulheim. A machine learning approach for product matching and categorization. *Semantic web*, 9(5):707–728, 2018.
- 38 Michael Ian Shamos and Dan Hoey. Closest-point problems. In *16th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 151–162. IEEE, 1975.
- 39 Warren D Smith. How to find steiner minimal trees in Euclidean d -space. *Algorithmica*, 7:137–177, 1992.
- 40 Godfried Toussaint. The geometry of musical rhythm. In *Japanese Conference on Discrete and Computational Geometry*, pages 198–212. Springer, 2004.
- 41 Godfried T Toussaint et al. A comparison of rhythmic similarity measures. In *ISMIR*, 2004.
- 42 Pravin M Vaidya. Geometry helps in matching. *SIAM Journal on Computing*, 18(6):1201–1225, 1989.
- 43 Pravin M Vaidya. An $O(n \log n)$ algorithm for the all-nearest-neighbors problem. *Discrete & Computational Geometry*, 4:101–115, 1989.
- 44 Kasturi R Varadarajan. A divide-and-conquer algorithm for min-cost perfect matching in the plane. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 320–329. IEEE, 1998.
- 45 Kasturi R Varadarajan and Pankaj K Agarwal. Approximation algorithms for bipartite and non-bipartite matching in the plane. In *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, volume 99, pages 805–814. SIAM, 1999.
- 46 Lee J White. Minimum covers of fixed cardinality in weighted graphs. *SIAM Journal on Applied Mathematics*, 21(1):104–113, 1971.