

Approximation Algorithms for Bipartite Matching with Metric and Geometric Costs

Pankaj K. Agarwal
Duke University

R. Sharathkumar
Stanford University

ABSTRACT

Let $G = G(A \cup B, A \times B)$, with $|A| = |B| = n$, be a weighted bipartite graph, and let $d(\cdot, \cdot)$ be the cost function on the edges. Let $w(M)$ denote the weight of a matching in G , and M^* a minimum-cost perfect matching in G . We call a perfect matching M c -approximate, for $c \geq 1$, if $w(M) \leq c \cdot w(M^*)$. We present three approximation algorithms for computing minimum-cost perfect matchings in G .

First, we consider the case when $d(\cdot, \cdot)$ is a metric. For any $\delta > 0$, we present an algorithm that, in $O(n^{2+\delta} \log n \log^2(1/\delta))$ time, computes a $O(1/\delta^\alpha)$ -approximate matching of G , where $\alpha = \log_3 2 \approx 0.631$. Next, we assume the existence of a dynamic data structure for answering approximate nearest neighbor (ANN) queries under $d(\cdot, \cdot)$. Given two parameters $\varepsilon, \delta \in (0, 1)$, we present an algorithm that, in $O(\varepsilon^{-2} n^{1+\delta} \tau(n, \varepsilon) \log^2(n/\varepsilon) \log(1/\delta))$ time, computes a $O(1/\delta^\alpha)$ -approximate matching of G , where $\alpha = 1 + \log_2(1 + \varepsilon)$ and $\tau(n, \varepsilon)$ is the query and update time of an $(\varepsilon/2)$ -ANN data structure.

Finally, we present an algorithm that works even if $d(\cdot, \cdot)$ is not a metric but admits an ANN data structure for $d(\cdot, \cdot)$. In particular, we present an algorithm that computes, in $O(\varepsilon^{-1} n^{3/2} \tau(n, \varepsilon) \log^4(n/\varepsilon) \log \Delta)$ time, a $(1 + \varepsilon)$ -approximate matching of A and B ; here Δ is the ratio of the largest to the smallest-cost edge in G , and $\tau(n, \varepsilon)$ is the query and update time of an (ε/c) -ANN data structure for some constant $c > 1$.

We show that our results lead to faster matching algorithms for many geometric settings.

Categories and Subject Descriptors

F.2.2 [Theory of Computation]: Nonnumerical Algorithms and Problems—*geometrical problems and computations*; G.2.2 [Discrete Mathematics]: Graph Theory—*graph algorithms*

General Terms

Algorithms, Theory

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

STOC'14, May 31– June 03, 2014, New York, New York, USA.

Copyright 2014 ACM 978-1-4503-2710-7/14/05 ...\$15.00 .

<http://dx.doi.org/10.1145/2591796.2591844>

Keywords

Matching, approximation algorithms

1. INTRODUCTION

Problem Statement. Let $G := G(A, B) = G(A \cup B, A \times B)$ be a bipartite graph with $|A| = |B| = n$, and let $d : A \times B \rightarrow \mathbb{R}_{\geq 0}$ be the edge-cost function. Throughout this paper we assume that $d(\cdot, \cdot)$ can be evaluated in $O(1)$ time. A *perfect matching* M is a set of n vertex-disjoint edges in $G(A, B)$. The cost of M , $w(M)$, is the sum of cost of its edges, i.e., $w(M) = \sum_{(a,b) \in M} d(a,b)$. When the cost function $d(\cdot, \cdot)$ is not obvious from the context, we will use the notation $w(M, d)$ to represent the cost of M under d . The *minimum-cost bipartite matching* (or *optimal matching* for brevity) is a perfect matching with the smallest cost; let M^* be such a matching. A perfect matching M for G is called c -approximate, for $c \geq 1$, if $w(M) \leq c \cdot w(M^*)$. In this paper, we develop approximation algorithms for computing optimal matchings when A and B are points in a metric space, and when A and B are point sets in \mathbb{R}^d and $d(\cdot, \cdot)$ is not necessarily a metric. The problem of computing an optimal matching in metric and geometric settings arises in several applications such as computer vision, shape analysis, computer graphics. For example, the problem in computer vision of tracking multiple objects can be formulated as computing an optimal matching between two sets of objects, and a metric is defined to measure the distance between a pair of objects [6].

Previous work. Computing a maximum cardinality matching in a graph (both bipartite and general graphs) with n vertices and m edges takes $O(m\sqrt{n})$ time [20, 23]. The bound was improved for sparse graphs by Madry [22] (see also [17, 28]). For weighted bipartite graphs, an optimal matching with n vertices and m edges can be computed, in $O(mn + n^2 \log n)$ time, using the Hungarian algorithm [12]. If the edge costs are positive integers bounded by $n^{O(1)}$, Gabow and Tarjan show that an optimal matching can be computed in $O(m\sqrt{n} \log n)$ time for bipartite graphs [14], and $O(m\sqrt{n} \log^{3/2} n)$ time for general graphs [15]—almost matching the time bounds for that of the unweighted case. There are algorithms that use algebraic methods to compute optimal matching in matrix multiplication time [8, 27]. Bipartite matching is a special case of general graph matching, and the known algorithms for the latter are more complex.

If A and B are points in a metric space, computing an optimal bipartite matching of A and B seems more challenging than computing an optimal matching on a complete graph

of $2n$ points in the same metric space.¹ For instance, an optimal matching in a set of $2n$ points in \mathbb{E}^2 , the 2D Euclidean plane, can be computed in $O(n^{3/2} \log^5 n)$ time [33], while the best known algorithm for computing an optimal matching between two point sets of size n each in \mathbb{E}^2 takes $O(n^{2+\delta})$ time, for any constant $\delta > 0$ [1]. Recently, Sharathkumar [29] presented a $O(n^{3/2+\delta} \log(n\Delta))$ -time algorithm to compute an optimal matching between two point sets in \mathbb{E}^2 with positive integer coordinates bounded by Δ . A sub-quadratic algorithm for computing an optimal bipartite matching of two arbitrary point sets in \mathbb{E}^2 has remained elusive so far.

The same trend holds for approximation algorithms. Several approximation algorithms are known for computing a (perfect) matching in a set of points (i.e., non-bipartite case) in arbitrary metric space [25, 26]. The best among those—a 2-approximation algorithm by Goemans and Williamson [18]—runs in near-quadratic time [7, 13]. However, it does not extend to the bipartite setting. The only known near-quadratic time algorithm for the bipartite case computes a matching whose expected cost is within $O(\log n)$ of the optimal matching [11].

A near-linear time algorithm for computing a $(1 + \varepsilon)$ -approximate matching in a set of n points in \mathbb{E}^d , for $d = O(1)$, was first presented by Arora [4], and later improved by Varadarajan and Agarwal [34]. Following this, considerable effort was made to design near-linear-time approximation algorithm for bipartite matching in low-dimensional Euclidean spaces [2, 21, 31, 34]. Only recently Sharathkumar and Agarwal [31] presented a near-linear time Monte-Carlo algorithm for the bipartite case. The best-known deterministic approximation algorithm known for low-dimensional Euclidean spaces is by Varadarajan and Agarwal [34] that runs in time $O(n^{3/2}(\log(n)/\varepsilon)^{O(d)})$.

In high-dimensional Euclidean space, when d is not assumed to be a constant, Goel *et al.* [16] presented an algorithm for computing a $2(1 + \varepsilon)$ -approximate matching in a set of $2n$ points in \mathbb{E}^d that runs in $d^{O(1)}n^{1+1/(1+\varepsilon)}$ time, for any constant $\varepsilon > 0$. For the bipartite case, besides the cubic-time exact algorithm, the only known approximation algorithms are an $O(dn)$ time algorithm that computes an $O(d \log n)$ -approximate matching [11], and an $O(nd \log^{O(1)} n)$ -time algorithm that computes an $O(\log d \log n)$ -approximate matching [3].

A closely related problem is the *maximum-cost matching* problem — compute a matching with the maximum cost. While exact algorithms for maximum-cost matching and minimum-cost matching have the same running time, approximating minimum-cost matching is considerably harder than approximating maximum-cost matching. For example, a simple greedy algorithm produces a 0.5-approximation to maximum weight matching for any weighted graph. A similar greedy algorithm gives only an $O(n^{0.58})$ approximation [26] even when the edge costs are assumed to satisfy triangle inequality. In fact, as observed by Avis [5], computing an $O(n)$ -approximation of minimum-cost matching on a graph with arbitrary edge costs takes the same time as computing the maximum cardinality matching of a dense unweighted graph, i.e., $\Omega(m\sqrt{n})$ time. In contrast, Duan and Pettie [9] present an $O(m/\varepsilon \log(1/\varepsilon))$ -time algorithm to compute a $(1 - \varepsilon)$ -approximate maximum-cost matching.

¹Note that bipartite matching is not a special case of the non-bipartite matching in this setting.

Our results. We present three approximation algorithms for computing a minimum-cost matching in G . First, given a constant $\delta \in (0, 1/3)$, we present, in Section 3, a deterministic algorithm that computes, in $O(n^{2+\delta} \log n \log^2 1/\delta)$ time, an $O((1/\delta)^\alpha)$ -approximate matching, where $\alpha = \log_3 2 \approx 0.631$, for the case when $d(\cdot, \cdot)$ is an arbitrary metric that can be evaluated in $O(1)$ time. This is the first $O(1)$ -approximation algorithm that runs in $o(n^{2.5})$ time. Moreover, our algorithm computes, in $O(n^2 \log^{O(1)} n)$ time, an $O((\log n / \log \log n)^\alpha)$ -approximate matching; this compares favorably to the randomized quadratic-time algorithm for computing a $O(\log n)$ -approximate matching. Our algorithm is a variant of the scaling algorithms [14, 30]. Roughly speaking, our algorithm computes a partial matching at each scale, commits this matching, and solves the problem recursively for unmatched points at the next scale. We use the metric property of $d(\cdot, \cdot)$ to bound the cost of an optimal matching of the unmatched points, and show that we compute an approximate matching at each scale.

In Section 4, we improve the above algorithm if $d(\cdot, \cdot)$ admits an efficient dynamic data structure for answering ε -approximate nearest-neighbor (ANN) queries. We present an algorithm that given two parameters $\varepsilon, \delta \in (0, 1)$, computes an $O(1/\delta^\alpha)$ -approximate matching of A and B in $O(\varepsilon^{-2} n^{1+\delta} \tau(n, \varepsilon) \log^2(n/\varepsilon) \log(1/\delta))$ time. Here $\alpha = 1 + \log_2(1 + \varepsilon)$ and $\tau(n, \varepsilon)$ is the query and update time of the $\varepsilon/2$ -ANN data structure. Note that we sacrifice the approximation factor a little but gain substantially in the running time, provided $\tau(n, \varepsilon)$ is small. If we set $\varepsilon = 1/\log_2(1/\delta)$, the algorithm computes a $O(1/\delta)$ -approximate matching.

Finally, by combining the ideas of the second algorithm with the Hungarian algorithm, we present an algorithm that works even when $d(\cdot, \cdot)$ is not necessarily a metric but admits an efficient ANN data structure (e.g. squared Euclidean distance). More precisely, we present in Section 5 an algorithm that computes a $(1 + \varepsilon)$ -approximate matching of A and B in time $O(\varepsilon^{-1} n^{3/2} \tau(n, \varepsilon) \log^4(n/\varepsilon) \log \Delta)$; here Δ is the ratio of the largest to the smallest-cost edge in G . Previously, data structures for answering (additive) weighted nearest neighbor (WNN) queries were used to speed up the Hungarian algorithm for geometric settings [30]. It is tempting to attain an approximation algorithm by replacing the WNN data structure in these algorithms with the one that answers approximate WNN queries, with the hope that the latter can be answered more quickly. However, it is unlikely that approximate WNN queries can be answered quickly—in time similar to unweighted ANN queries, because exact NN queries can be reduced to answering approximate WNN queries. We overcome this difficulty by using a different approach that allows us to use a weaker notion of approximate WNN queries which can be answered using ANN data structure.

These results lead to faster matching algorithms for many geometric settings. We just mention two of them here:

(i) For $A, B \subseteq \mathbb{E}^d$, with $d = \Omega(\log n)$, and for a constant $\delta > 0$, an $O(1/\delta^{\log_2(1/\delta)})$ -approximate matching can be computed in $O(n^{1+\delta} \log n)$ time. This is the first sub-quadratic algorithm for computing an $O(1)$ -approximate matching of high-dimensional point sets.

(ii) For $A, B \subseteq \mathbb{R}^d$, with $d = O(1)$, and for a constant $\delta > 0$, we obtain a *deterministic* algorithm that computes an $O(\frac{1}{\delta})$ -approximate matching of A and B in $O(n^{1+\delta} \log n)$ time. The previously best known deterministic algorithm computes a $\frac{1}{\delta}$ -approximate matching in $O(n^{3/2} \log^{O(d)} n)$

time [34] (using our second algorithm). Alternatively, we also obtain a *deterministic* algorithm that computes a $(1+\varepsilon)$ -approximate matching of A and B in $O(\frac{1}{\varepsilon^{O(d)}} n^{3/2} \log^5(n/\varepsilon))$ time (using our third algorithm), while the algorithm in [34] takes $O(n^{3/2}(\log(n)/\varepsilon)^{O(d)})$ time.

2. PRELIMINARIES

In this section, we present a few concepts that will be useful for our algorithms. Let A, B, G , and $d(\cdot, \cdot)$ be as defined earlier. Given a matching M on $G(A, B)$, an *alternating path* (or cycle) is a simple path (resp. cycle) whose edges are alternately in and not in M . A *free vertex* is a vertex of $A \cup B$ that has not been matched, i.e., not incident on any edge of M . An *augmenting path* P is an alternating path between two free vertices. M can be *augmented* by one edge along P , by setting $M = M \oplus P$, where \oplus is the symmetric difference, i.e., remove $P \cap M$ from M and add $P \setminus M$ to M . We refer to this step as an *augmentation* step.

An *alternating tree* is a tree rooted at a free vertex of B in which every path to the leaf node is an alternating path. While describing our algorithms, sometimes it is useful to view G as a directed graph \vec{G}_M , defined as follows: Every edge $(a, b) \in M$ with $a \in A$ and $b \in B$, is directed from a to b , and every edge $(a, b) \notin M$ is directed from b to a . The in-degree (resp. out-degree) of every vertex of B (resp. A) is at most 1. Any alternating path is a directed path in \vec{G}_M . An alternating tree \mathcal{T} rooted at a free vertex $b \in B$ is a directed tree in \vec{G}_M , in the sense that \mathcal{T} has a directed path from the root b to every vertex in \mathcal{T} .

Recall that Hungarian and many other algorithms assign dual weights to each vertex of G . Let $d(\cdot, \cdot)$ be a distance function. For a vertex $v \in A \cup B$, let $y(v)$ be its *dual weight*. A matching M and a set of dual weights are *dual feasible* if

$$\begin{aligned} y(a) + y(b) &\leq d(a, b) & \forall (a, b) \in A \times B, \\ y(a) + y(b) &= d(a, b) & \forall (a, b) \in M. \end{aligned} \quad (1)$$

It is well-known that any dual-feasible perfect matching is optimal with respect to $d(\cdot, \cdot)$; see [24]. Introduced by Gabow and Tarjan [14], a *1-feasible matching* is a matching M and set of dual weights $y(\cdot)$ such that

$$\begin{aligned} y(a) + y(b) &\leq d(a, b) + 1 & \forall (a, b) \in A \times B, \\ y(a) + y(b) &= d(a, b) & \forall (a, b) \in M. \end{aligned} \quad (2)$$

A *1-optimal matching* is a perfect matching that is 1-feasible. The notion of 1-feasible matching will be used in Section 3. An edge $(a, b) \in A \times B$ is called *eligible* if

$$\begin{aligned} y(a) + y(b) &= d(a, b) & (a, b) \in M, \\ y(a) + y(b) &= d(a, b) + 1 & (a, b) \notin M. \end{aligned} \quad (3)$$

The subgraph of G induced by the set of eligible edges is called the *eligible graph* (with respect to M).

Relaxing the notion of feasible edges differently, a matching M is called ε -relaxed if the dual weights satisfy the following condition:

$$\begin{aligned} y(a) + y(b) &\leq d(a, b) + \lceil \varepsilon d(a, b) \rceil, & \forall (a, b) \in A \times B, & (4) \\ y(a) + y(b) &= d(a, b) & \forall (a, b) \in M. & (5) \end{aligned}$$

The following lemma shows the significance of ε -relaxed matchings.

LEMMA 2.1. *Let M be a perfect $(\varepsilon/2)$ -relaxed matching of A and B , and let M^* be an optimal matching of them. If $w(M^*) \geq 2n/\varepsilon$, then*

$$w(M) \leq (1 + \varepsilon)w(M^*).$$

PROOF. Since M is a perfect matching, from (5), we have

$$w(M) = \sum_{(u,v) \in M} d(u, v) = \sum_{(u,v) \in M} y(u) + y(v) = \sum_{v \in A \cup B} y(v).$$

Every edge $(a, b) \in M^*$ satisfies (4) and $w(M^*, d) \geq \frac{2n}{\varepsilon}$, therefore

$$\begin{aligned} \sum_{a \in A \cup B} y(a) &\leq \sum_{(u,v) \in M^*} d(u, v) + \left\lceil \frac{\varepsilon}{2} d(u, v) \right\rceil \\ &\leq \left(1 + \frac{\varepsilon}{2}\right)w(M^*) + n \leq (1 + \varepsilon)w(M^*, d). \end{aligned}$$

□

3. ALGORITHM FOR ARBITRARY METRIC SPACE

Let A, B be two point sets of size n each in a metric space, and let $d(\cdot, \cdot)$ be the underlying metric such that the distance between any two points of $A \cup B$ can be computed in $O(1)$ time. We describe an approximation algorithm for computing an optimal matching of A and B under $d(\cdot, \cdot)$. Our algorithm is a variant of the *scaling algorithms* for matching [14, 30]. Before presenting the algorithm, we briefly describe the main idea behind it. This description, although somewhat incorrect, conveys the main intuition.

Main idea. Suppose we fix a scale of the Gabow-Tarjan algorithm [14]. At a fixed scale, their algorithm runs in stages, each of which computes, in $O(n^2)$ time, a maximal set of vertex-disjoint augmenting paths, augments the matching along these paths, and updates the dual weights of the vertices. If we run their algorithm for n^α stages, then it computes a matching M_1 that matches, in $O(n^{2+\alpha})$ time, all but $O(n^{1-\alpha})$ points of A and B , and $w(M_1) \leq 2w(M^*)$, where M^* is the optimal matching of A, B . Next, we compute an optimal matching M_2 of the free points of A and B , using the Hungarian algorithm, which takes $O(n^{3(1-\alpha)})$ time, and return $M_1 \cup M_2$ as a matching of A, B . Using the fact that $d(\cdot, \cdot)$ is a metric, we can argue that $w(M_2) \leq w(M_1) + w(M^*) \leq 3w(M^*)$. By setting $\alpha = 1/4$, we obtain a 5-approximate matching of A, B in $O(n^{9/4})$ time. We can repeat this process once more, but this time use the $O(n^{9/4})$ algorithm, instead of Hungarian algorithm, to compute M_2 . We obtain a faster algorithm that returns an $O(1)$ -approximate matching, though the constant is now larger than 5. By iterating this process multiple times, we can obtain a trade-off between the approximation factor and the running time. How we do this correctly and obtain a good trade-off requires some care, as described below.

There are three main differences between our and traditional scaling algorithms [14, 30]: (i) We commit the edges computed in the matching computed in one scale, and the next scale computes a matching of only unmatched points. (ii) We compute a partial matching at each scale. By choosing the number of steps for which we run the algorithm at each scale, we keep both the running time and the cost of the matching under control. (iii) In successive scales, we scale down the edge costs instead of scaling up, so that the cost

of the optimal matching of unmatched points is linear in the number of the points.

Algorithm. We now describe the algorithm in detail. We choose a parameter $\delta \in (0, 1/3]$ and set $\varepsilon = \frac{1}{2 \log_3 1/\delta}$. Let $\omega > 0$ be a real value such that $\mathfrak{w}(M^*)/2 \leq \omega \leq \mathfrak{w}(M^*)$. For now, we assume that we have such a value of ω at our disposal. At the end, we describe how to choose ω . The algorithm works in phases. In the beginning of phase i , for $i \geq 0$, we have sets $A_i \subseteq A$ and $B_i \subseteq B$ of points and a matching \mathbb{M}_i of the points in $A \setminus A_i$ and $B \setminus B_i$; set $|A_i| = |B_i| = n_i$. For $i = 0$, $A_0 = A$, $B_0 = B$, $n_0 = n$, and $\mathbb{M}_0 = \emptyset$. If $n_i \leq n^{2/3}$, we compute an optimal matching of A_i and B_i using the Hungarian algorithm, which takes $O(n_i^3) = O(n^2)$ time. So, assume $n_i \geq n^{2/3}$. In phase i , we set the distance function to be $\mathfrak{d}_i(\cdot, \cdot)$, defined as follows:

$$\mathfrak{d}_i(a, b) = \begin{cases} \left\lceil \frac{2\mathfrak{d}(a, b) \cdot n}{\varepsilon \omega} \right\rceil & \text{if } i = 0, \\ \left\lceil \frac{\mathfrak{d}_{i-1}(a, b)}{2(1 + \varepsilon)^2 n^{\varphi_{i-1}}} \right\rceil & \text{if } i > 0, \end{cases} \quad (6)$$

where $\varphi_i = 3^i \delta$.

We compute a (possibly partial) matching M_i of A_i and B_i under $\mathfrak{d}_i(\cdot, \cdot)$ using the algorithm PHASE-MATCH described below. We set $\mathbb{M}_{i+1} \leftarrow \mathbb{M}_i \cup M_i$. If M_i is a perfect matching between A_i and B_i , we stop and return \mathbb{M}_{i+1} . Otherwise, we set A_{i+1} (resp. B_{i+1}) to be the set of free vertices of A_i (resp. B_i) and move to phase $i + 1$.

We now describe the PHASE-MATCH algorithm, which is similar to Gabow and Tarjan's [14] algorithm but with some additional twists. We set $\kappa_i = \frac{30}{\varepsilon} n^{\varphi_i}$. The procedure works in stages. In the beginning of each stage, we have a set $A_i^F \subseteq A_i$ and $B_i^F \subseteq B_i$ of free vertices, dual weights $y(v)$ for each point in $A_i \cup B_i$ and a 1-feasible matching M computed so far. Initially, $A_i^F = A_i$, $B_i^F = B_i$, and $M = \emptyset$, and $y(v) = 0$ for all $v \in A_i \cup B_i$. At the end of the stage, if we have computed a perfect matching of A_i, B_i or run the procedure for κ_i stages, we stop and return M as M_i .

Each stage of PHASE-MATCH consists of the following three steps:

(i) **Finding augmenting paths.** Let $\mathcal{E} \subseteq A_i \times B_i$ be the set of *eligible* edges. We compute, in $O(n_i^2)$ time, a maximal set \mathbb{P} of vertex-disjoint augmenting paths in \mathcal{E} , as described in [14, 30].

(ii) **Augmentation step.** We augment M along \mathbb{P} , i.e., set $M = M \oplus (\bigcup_{P \in \mathbb{P}} P)$. For each vertex $b \in B_i$ that lies on a path in \mathbb{P} , set $y(b) \leftarrow y(b) - 1$, to ensure that M continues to be 1-feasible after the augmentation step. This step takes $O(n_i)$ time because the paths in \mathbb{P} are vertex disjoint.

(iii) **Dual-adjustment step.** Consider the directed subgraph G_i formed by \mathcal{E} , the set of eligible edges, as defined in Section 2. Let S (resp. T) be the set of vertices of B_i (resp. A_i) that are reachable from a vertex of B_i^F in G_i . Set $y(b) = y(b) + 1$ for all $b \in S$ and $y(a) = y(a) - 1$ for all $a \in T$. This step takes $O(n_i^2)$ time.

This completes the description of the PHASE-MATCH procedure. Since each stage takes $O(n_i^2)$ time, the total time taken by PHASE-MATCH is $O(n_i^2 \kappa_i)$.

Analysis. Before analyzing the performance of the algorithm, we make a few observations about the PHASE-MATCH procedure, which are well-known; see [14, 30].

LEMMA 3.1. (i) *The matching M is always 1-feasible.* (ii) *At the end of the augmentation step in each stage, there is no augmenting path composed solely of eligible edges.* (iii) *For any vertex $v \in A_i \cup B_i$ that is free after k stages, $y(v) = 0$ if $v \in A_i$ and $y(v) = k$ if $v \in B_i$.*

We claim that our (overall) algorithm maintains the following invariants after i phases. Let M_i^* be the optimal perfect matching of A_i and B_i under $\mathfrak{d}(\cdot, \cdot)$; note that $M_0^* = M^*$.

(I1) $\mathfrak{w}(M_i) \leq (1 + \varepsilon)\mathfrak{w}(M_i^*)$,

(I2) $\mathfrak{w}(M_{i+1}^*) \leq (2 + \varepsilon)\mathfrak{w}(M_i^*)$, and

(I3) $n_i \leq n^{1 - \Phi_i}$, where $\Phi_i = \sum_{j=0}^{i-1} \varphi_j = \frac{3^i - 1}{2} \delta$.

We note that (I2) holds because $\mathfrak{d}(\cdot, \cdot)$ is a metric, and this is the only place where the metric property is used. Assuming that these invariants hold, we analyze the performance of the algorithm.

(I3) implies that the algorithm runs as long as $\Phi_i < 1/3$. In other words, the algorithm stops after $\mu \leq \lceil \log_3(\frac{2}{3\delta} + 1) \rceil$ phases.

Next, we analyze the running time of the algorithm. In phase i , PHASE-MATCH takes $O(n_i^2 \kappa_i) = O(\frac{n^{2+\delta}}{\varepsilon} n^{2-2\Phi_i})$ time. Since

$$\varphi_i - 2\Phi_i = 3^i \delta - 2 \frac{3^i - 1}{2} \delta = \delta,$$

the PHASE-MATCH procedure in phase i runs in $O(n^{2+\delta}/\varepsilon)$ time. Plugging $\mu = O(\log(1/\delta))$ and $\varepsilon = \frac{1}{2 \log_3(1/\delta)}$, the overall running time of the algorithm is $O(n^{2+\delta} \log^2(1/\delta))$.

The following lemma, which bounds the cost of the matching computed by the algorithm, follows from (I1), (I2), and the value of μ .

LEMMA 3.2. *Let \mathbb{M} be the final matching computed by the algorithm, then $\mathfrak{w}(\mathbb{M}) = O(1/\delta^\alpha) \cdot \mathfrak{w}(M^*)$, where $\alpha = \log_3 2 \approx 0.631$.*

PROOF. Using the invariant (I1),

$$\mathfrak{w}(\mathbb{M}) = \sum_{i=0}^{\mu} \mathfrak{w}(M_i) \leq (1 + \varepsilon) \sum_{i=0}^{\mu} \mathfrak{w}(M_i^*).$$

By applying (I2) repeatedly, we obtain

$$\mathfrak{w}(M_i^*) \leq 2^i (1 + \varepsilon)^i \mathfrak{w}(M_0^*) = 2^i (1 + \varepsilon)^i \mathfrak{w}(M^*).$$

Hence,

$$\mathfrak{w}(\mathbb{M}) \leq (1 + \varepsilon) \mathfrak{w}(M^*) \sum_{i=0}^{\mu-1} 2^i (1 + \varepsilon)^i \leq (1 + \varepsilon) (2(1 + \varepsilon))^\mu \mathfrak{w}(M^*).$$

Since $\mu \leq \lceil \log_3(\frac{2}{3\delta} + 1) \rceil$ and $\varepsilon = \frac{1}{2 \log_3(1/\delta)}$, $(1 + \varepsilon)^\mu = O(1)$ and $2^\mu = O(1/\delta^\alpha)$, where $\alpha = \log_3 2$. Hence, $\mathfrak{w}(\mathbb{M}) = O(1/\delta^\alpha) \mathfrak{w}(M^*)$. \square

We now prove invariants (I1)–(I3). We first note that $\mathfrak{d}_i(\cdot, \cdot)$ is not an exact scaled version of $\mathfrak{d}(\cdot, \cdot)$ because of the ceiling operator in (6). Nevertheless, we can prove the following properties for $\mathfrak{d}(\cdot, \cdot)$.

LEMMA 3.3. (i) *For any $i \geq 0$, $\mathfrak{d}_i(\cdot, \cdot)$ is a metric.*

(ii) *For $i \geq 1$, $\mathfrak{d}_i(a, b) \geq \frac{\delta}{\varepsilon}$ for any $(a, b) \in A_i \times B_i$.*

(iii) For any $i \geq 1$, there is a scaling factor σ_i such that

$$(1 - \varepsilon/3)\sigma_i \mathbf{d}_i(a, b) \leq \mathbf{d}(a, b) \leq \sigma_i \mathbf{d}_i(a, b).$$

PROOF. (i) Suppose $d(\cdot, \cdot)$ is a distance function that satisfies triangular inequality, i.e., for any three points a, b and c , $d(a, b) + d(b, c) \geq d(a, c)$. This inequality continues to hold if we multiply it by a parameter $k \in \mathbb{R}$ and take the ceiling, so we get

$$\lceil kd(a, b) \rceil + \lceil kd(b, c) \rceil \geq \lceil kd(a, c) \rceil. \quad (7)$$

Part (i) can be proved by induction on i . For $i = 0$, set $k = \frac{n}{\varepsilon\omega}$ and $d(\cdot, \cdot)$ to $\mathbf{d}(\cdot, \cdot)$, and we have $\mathbf{d}_0(\cdot, \cdot)$ satisfying triangle inequality. Assume $\mathbf{d}_{i-1}(\cdot, \cdot)$ satisfies triangle inequality. Setting $k = \frac{1}{2(1+\varepsilon)^2 n^{\varphi_{i-1}}}$ and $d(\cdot, \cdot) = \mathbf{d}_{i-1}(\cdot, \cdot)$, we have $\mathbf{d}_i(\cdot, \cdot)$ satisfying triangle inequality.

(ii) If the algorithm does not terminate in phase i , then phase i runs for $\frac{30}{\varepsilon} n^{\varphi_i}$ iterations, so for any $(a, b) \in A_i^F \times B_i^F$, Lemma 3.1(iii) implies that

$$y_i(a) + y_i(b) = \frac{30}{\varepsilon} n^{\varphi_i} \leq \mathbf{d}_i(a, b) + 1$$

or $\mathbf{d}_i(a, b) \geq \frac{30}{\varepsilon} n^{\varphi_i} - 1$. But

$$\mathbf{d}_{i+1}(a, b) \geq \frac{\mathbf{d}_i(a, b)}{2(1+\varepsilon)^2 n^{\varphi_i}} \geq \frac{30}{2(1+\varepsilon)^2 \varepsilon} - 1 \geq \frac{6}{\varepsilon}.$$

The last inequality follows since $\varepsilon = \frac{1}{2 \log_3(1/\delta)} \leq 1/2$.

(iii) By ignoring the ceiling operator in (6), we only decrease the value of $\mathbf{d}_i(\cdot, \cdot)$, so using (6) repeatedly, we obtain

$$1/\sigma_i \cdot \mathbf{d}(a, b) \leq \mathbf{d}_i(a, b) \quad (8)$$

where, $\sigma_i = \frac{\omega \varepsilon (2(1+\varepsilon)^2)^i}{n^{1-\Phi_i}}$. On the other hand,

$$\mathbf{d}_i(a, b) \leq \frac{1}{2(1+\varepsilon)^2 n^{\varphi_{i-1}}} \mathbf{d}_{i-1}(a, b) + 1 \quad (9)$$

Unraveling the recurrence and following a straight-forward, albeit somewhat tedious, calculation, we obtain

$$\mathbf{d}_i(a, b) \leq \frac{\mathbf{d}(a, b)}{\sigma_i} + 2 \quad (10)$$

Now using part (ii) of the lemma, we obtain

$$(1 - \varepsilon/3)\sigma_i \mathbf{d}_i(a, b) \leq \mathbf{d}(a, b) \leq \sigma_i \mathbf{d}_i(a, b). \quad (11)$$

□

COROLLARY 3.4. For $i \geq 1$, Let M and M' be two (possibly partial) matchings of A_i, B_i .

(i) If M is a perfect matching, then $\mathbf{w}(M, \mathbf{d}_i) \geq \frac{6}{\varepsilon} n_i$.

(ii) $(1 - \varepsilon/3) \frac{\mathbf{w}(M, \mathbf{d}_i)}{\mathbf{w}(M', \mathbf{d}_i)} \leq \frac{\mathbf{w}(M, \mathbf{d})}{\mathbf{w}(M', \mathbf{d})} \leq \frac{1}{1 - \varepsilon/3} \frac{\mathbf{w}(M, \mathbf{d}_i)}{\mathbf{w}(M', \mathbf{d}_i)}$.

Now, we are ready to prove the invariants (I1)–(I3).

Proof of (I1): For every edge $(a, b) \in M_i$, $y(a) + y(b) = \mathbf{d}_i(a, b)$ and for every free vertex v of $A_i \cup B_i$, $y(v) \geq 0$, therefore,

$$\mathbf{w}(M_i, \mathbf{d}_i) \leq \sum_{v \in A_i \cup B_i} y(v).$$

For every edge $(a, b) \in M_i^*$, $y(a) + y(b) \leq \mathbf{d}_i(a, b) + 1$. Every vertex of $A_i \cup B_i$ is incident on exactly one edge of M_i^* , so

$$\sum_{v \in A_i \cup B_i} y(v) \leq \sum_{(a, b) \in M_i^*} \mathbf{d}_i(a, b) + n_i \leq \mathbf{w}(M_i^*, \mathbf{d}_i) + n_i. \quad (12)$$

Consequently,

$$\mathbf{w}(M_i, \mathbf{d}_i) \leq \mathbf{w}(M_i^*, \mathbf{d}_i) + n_i. \quad (13)$$

We prove (I1) for cases $i = 0$ and $i > 0$ separately:

(i) **Case $i = 0$:** From (13), we have $\mathbf{w}(M_0, \mathbf{d}_0) \leq \mathbf{w}(M^*, \mathbf{d}_0) + n$. Since $\mathbf{d}_0(a, b) = \lceil \frac{2\mathbf{d}(a, b) \cdot n}{\varepsilon\omega} \rceil$ and $\omega \leq \mathbf{w}(M^*)$, we obtain the following inequalities:

$$\begin{aligned} \mathbf{w}(M_0, \mathbf{d}_0) &\leq \sum_{(a, b) \in M^*} \left(\frac{2n}{\varepsilon\omega} \mathbf{d}(a, b) + 1 \right) + n \\ &\leq \frac{2n}{\varepsilon\omega} \mathbf{w}(M^*) + 2n, \end{aligned} \quad (14)$$

$$\mathbf{w}(M_0, \mathbf{d}_0) \geq \frac{2n}{\varepsilon\omega} \sum_{(a, b) \in M_0} \mathbf{d}(a, b) = \frac{2n}{\varepsilon\omega} \mathbf{w}(M_0). \quad (15)$$

Combining (14) and (15) and using $\omega \leq \mathbf{w}(M^*)$,

$$\mathbf{w}(M_0) \leq \mathbf{w}(M^*) + \varepsilon\omega \leq (1 + \varepsilon)\mathbf{w}(M^*).$$

(ii) **Case $i > 0$:** From (13) and Corollary 3.4 (i), we get

$$\mathbf{w}(M, \mathbf{d}_i) \leq \mathbf{w}(M_i^*, \mathbf{d}_i) + n_i \leq (1 + \varepsilon/6)\mathbf{w}(M_i^*, \mathbf{d}_i).$$

Finally, by Corollary 3.4(ii),

$$\mathbf{w}(M_i) \leq \frac{1 + \varepsilon/6}{1 - \varepsilon/3} \mathbf{w}(M_i^*) \leq (1 + \varepsilon)\mathbf{w}(M_i^*, \mathbf{d}).$$

Proof of (I2): $M_i \oplus M_i^*$ is a set of alternating cycles and a set \mathbb{P} of vertex-disjoint augmenting paths. Each path in \mathbb{P} connects a free vertex of B_i to a free vertex of A_i , i.e., $|\mathbb{P}| = n_{i+1}$. Using (I1),

$$\sum_{P \in \mathbb{P}} \sum_{(a, b) \in P} \mathbf{d}(a, b) = \mathbf{w}(M_i) + \mathbf{w}(M_i^*) \leq (2 + \varepsilon)\mathbf{w}(M_i^*).$$

Recall that $\mathbf{d}(\cdot, \cdot)$ satisfies the triangle inequality, so if the endpoints of a path $P_j \in \mathbb{P}$ are $(a_j, b_j) \in A_i \times B_i$, then $\mathbf{d}(a_j, b_j) \leq \sum_{(a, b) \in P_j} \mathbf{d}(a, b)$. Hence,

$$\mathbf{w}(M_{i+1}^*, \mathbf{d}) \leq \sum_{P \in \mathbb{P}} \sum_{(a, b) \in P} \mathbf{d}(a, b) \leq (2 + \varepsilon)\mathbf{w}(M_i^*, \mathbf{d}).$$

Proof of (I3): We claim that

$$\mathbf{w}(M_i^*, \mathbf{d}_i) \leq \frac{5}{\varepsilon} n^{1-\Phi_i}. \quad (16)$$

Suppose this claim is true. Then $n_i \leq n^{1-\Phi_i}$ because $\mathbf{w}(M_i^*, \mathbf{d}_i) \geq \frac{5}{\varepsilon} n_i$, by Corollary 3.4 (i).

It thus suffices to prove the above claim, which we prove by induction on i . For $i = 0$, $\mathbf{d}_0(a, b) - 1 \leq \frac{2n}{\varepsilon\omega} \mathbf{d}(a, b)$. Since $\omega \geq \mathbf{w}(M^*)/2$ we have,

$$\mathbf{w}(M_0^*, \mathbf{d}_0) - n \leq \frac{2n}{\varepsilon\omega} \mathbf{w}(M_0^*) \leq \frac{2n}{\varepsilon\omega} \mathbf{w}(M^*, \mathbf{d}) \leq \frac{4n}{\varepsilon}.$$

Therefore, $\mathbf{w}(M_0^*, \mathbf{d}_0) \leq \frac{5}{\varepsilon} n$.

By induction hypothesis, assume that

$$\mathbf{w}(M_{i-1}^*, \mathbf{d}_{i-1}) \leq \frac{5}{\varepsilon} n^{1-\Phi_{i-1}}.$$

Using (II), (I2), (6) and Corollary 3.4(ii), we can write

$$\mathfrak{w}(M_i^*, \mathbf{d}_{i-1}) \leq \frac{2+\varepsilon}{1-\varepsilon/3} \mathfrak{w}(M_{i-1}^*, \mathbf{d}_{i-1}) \leq 2(1+\varepsilon) \frac{5}{\varepsilon} n^{1-\Phi_{i-1}}; \quad (17)$$

the last inequality follows because $\varepsilon \leq 1/2$. On the other hand, by Lemma 3.3(ii),

$$(1-\varepsilon/6)\mathbf{d}_i(a, b) \leq \mathbf{d}_i(a, b) - 1 \leq \frac{\mathbf{d}_{i-1}(a, b)}{2(1+\varepsilon)^2 n^{\varphi_{i-1}}}.$$

Therefore

$$\mathbf{d}_i(a, b) \leq \frac{\mathbf{d}_{i-1}(a, b)}{2(1+\varepsilon)n^{\varphi_{i-1}}}. \quad (18)$$

Combining (17) and (18),

$$\mathfrak{w}(M_i^*, \mathbf{d}_i) \leq \frac{5}{\varepsilon} n^{1-\Phi_{i-1}-\varphi_{i-1}} = \frac{5}{\varepsilon} n^{1-\Phi_i},$$

as desired. Finally, to guess the value of ω , we compute, for a constant $c > 1$, an n^c -approximate matching of A and B in $O(n^2)$ time [30]. Let Θ be the cost of the matching, then we run the above algorithm with $\omega = \Theta/2^i$, for all $1 \leq i \leq \lceil c \log_2 n \rceil$, and choose the minimum-cost matching among them. Since one of these values of ω lies in the range $[\mathfrak{w}(M^*)/2, \mathfrak{w}(M^*)]$, we obtain the following.

THEOREM 3.5. *Let A, B be two point sets of size n each in a metric space such that the distance between any two points of A and B can be computed in $O(1)$ time, and let $\delta \in (0, 1/3)$ be a constant. Then an $O(1/\delta^\alpha)$ -approximate matching of A, B can be computed in $O(n^{2+\delta} \log n \log^2(1/\delta))$ time, where $\alpha = \log_3 2 \approx 0.631$.*

4. ANN DATA STRUCTURE BASED ALGORITHM

In this section, we present a more efficient algorithm by adapting the algorithm described in Section 3 and using a dynamic data structure for answering approximate nearest-neighbor (ANN) queries. Given two parameters, $0 \leq \delta \leq 1/3$ and $0 < \varepsilon < 1$, it computes a $O(1/\delta^\alpha)$ -approximate matching, for $\alpha = 1 + \log_2(1+\varepsilon)$, in $O(n^{1+\delta} \tau(n, \varepsilon) \log^2(n/\varepsilon))$ time, where $\tau(n, \varepsilon)$ is the update and query time of an $\varepsilon/2$ -approximate nearest neighbor ($\varepsilon/2$ -ANN) data structure.

We will now use the notion of ε -relaxed matching described in Section 2, instead of 1-feasible matching. Given an ε -relaxed matching M , we call an edge $(a, b) \in A \times B$ *strongly eligible* if $(a, b) \in M$ (in which case $y(a) + y(b) = \mathbf{d}(a, b)$) or $(a, b) \notin M$ and $y(a) + y(b) = \mathbf{d}(a, b) + \lceil \varepsilon \mathbf{d}(a, b) \rceil$. We call $(a, b) \notin M$ *weakly eligible* if $\mathbf{d}(a, b) < y(a) + y(b) < \mathbf{d}(a, b) + \lceil \varepsilon \mathbf{d}(a, b) \rceil$. Given M , let Σ (resp. Ω) be the set of all strongly (resp. weakly) eligible edges of $A \times B$.

We call a graph $G = (A \cup B, \mathcal{E})$ an ε -eligible graph if $\Sigma \subseteq \mathcal{E}$ and $\mathcal{E} \subseteq \Sigma \cup \Omega$, i.e., \mathcal{E} contains all strongly eligible edges and possibly some of the weakly eligible edges; an edge that is not weakly or strongly eligible is not in \mathcal{E} . Note that G is not unique as we do not specify which of the weakly eligible edges should be in G . We also note that we do not construct G explicitly. Instead, we will be relying on a data structure to explore an ε -eligible graph.

Algorithm. The structure of the overall algorithm is the same as in Section 3, we just change some of the parameters. In particular, in the beginning of phase i , for $i \geq 0$, we have sets A_i and B_i of n_i points each and a matching M_i

of $A \setminus A_i$ and $B \setminus B_i$. We set the threshold of n_i to $n^{1/3}$ instead of $n^{2/3}$, i.e., if $n_i < n^{1/3}$ we execute the Hungarian algorithm to compute an optimal matching of A_i, B_i , which takes $O(n_i^3) = O(n)$ time. So assume $n_i \geq n^{1/3}$. We set the distance function in phase i to be $\mathbf{d}_i(\cdot, \cdot)$ as defined in (6) except that we set $\varphi_i = 2^i \delta$ (instead of $3^i \delta$)

Suppose we have a data structure, which we describe later, that given an ε -relaxed matching M and weights $y(v)$ for all $v \in A_i \cup B_i$, maintains a set $K \subset A_i$ and supports the following two operations:

- **DELETE** (a): Deletes a from K .
- **ELIGIBLE** (b): Given a point $b \in B_i$, if there is a strongly eligible edge incident on b whose other end-point is in K , it returns an edge (a, b) with $a \in K$ and $(a, b) \in \Sigma \cup \Omega$, i.e., (a, b) is strongly or weakly eligible edge. If no strongly eligible edge is incident upon b , the procedure either returns NULL or an edge $(a, b) \in \Omega$ with a in K .

Let $\tilde{\tau}(n_i, \varepsilon)$ be the time taken by each of these two operations. The data structure can be constructed in $O(n_i \tilde{\tau}(n_i, \varepsilon))$ time.

The PHASE-MATCH procedure also has the same structure as in Section 3, except that we use the above data structure to compute augmenting paths. PHASE-MATCH works in stages and computes a $(\varepsilon/2)$ -relaxed matching M of A_i and B_i . We set $\kappa_i = \frac{30}{\varepsilon} n^{\varphi_i}$, where $\varphi_i = 2^i \delta$. In the beginning of each stage, we have sets $A_i^F \subseteq A_i$, $B_i^F \subseteq B_i$ of free vertices, and we build the above data structure on A_i . Each stage again consists of three steps. We describe the modifications to the other two steps.

(i) Finding augmenting paths. As before, we construct a maximal set of vertex-disjoint augmenting paths \mathbb{P} of edges that are strongly or weakly eligible – the paths are maximal in the sense that no additional augmenting path composed of only strongly eligible edges exists. We compute \mathbb{P} by doing a depth first search from the free vertices of B_i . Suppose we are currently at a vertex $b \in B_i$, then we call **ELIGIBLE**(b) to find a vertex of $a \in A_i$ that can be reached from b using a strongly or weakly eligible edge. If we visit a vertex $a \in A_i$, we delete a from the data structure. We omit the details from here.

The total time spent by the procedure is $O(n \tilde{\tau}(n_i, \varepsilon))$. Let $T \subset A_i$ and $S \subset B_i$ be all the points that are visited by this step.

(ii) Augmentation step. We augment along every path of \mathbb{P} . The only difference is that, for every path $P \in \mathbb{P}$, we decrease the dual weights of points in $B \cap P$ so that every new edge (a, b) added to the matching satisfies $y(a) + y(b) = \mathbf{d}(a, b)$. Since any edge $(a, b) \in P$ that is not in the matching is either weakly or strongly eligible, it satisfies $y(a) + y(b) > \mathbf{d}(a, b)$. Therefore, the reduction of the dual weight is at least 1. Total time taken by this step is $O(n_i)$.

(iii) Dual-adjustment step. Increase the dual weight of every point in S by 1, and reduce the dual weight of every point in T by 1.

This completes the description of PHASE-MATCH. Each stage of the algorithm runs in $O(n_i \tilde{\tau}(n_i, \varepsilon))$ time, so the total time for each phase is $O(n_i \kappa_i \tilde{\tau}(n_i, \varepsilon))$.

Analysis. The following variant of Lemma 3.1 now holds:

LEMMA 4.1. (i) *The matching M is always $(\varepsilon/2)$ -relaxed.*
(ii) *At the end of the augmentation step in each stage, there*

is no augmenting paths composed of only strongly eligible edges. (iii) For any vertex $v \in A_i \cup B_i$ that is free after k stages, $y(v) = 0$ if $v \in A_i$ and $y(v) = k$ if $v \in B_i$.

We now analyze the performance of the algorithm. The invariants (I1)–(I3) of Section 3 still hold but $\Phi_i = \sum_{j=0}^i \varphi_j = (2^i - 1)\delta$. Following the same analysis as in Section 3, we obtain that the algorithm stops within $\mu = \lceil \log_2(2/3\delta + 2) \rceil = O(\log 1/\delta)$ phases, and the phase i takes $O(n_i \tilde{\tau}(n_i, \varepsilon) \kappa_i) = O(\varepsilon^{-1} n^{1+\delta} \tilde{\tau}(n_i, \varepsilon))$ steps. Summing over all phases, the total running time is $O(\varepsilon^{-1} n^{1+\delta} \tilde{\tau}(n_i, \varepsilon) \log(1/\delta))$. The same argument as in Lemma 3.2 shows that if \mathbb{M} is the final matching computed by the algorithm, then $w(\mathbb{M}) \leq O(1/\delta^\alpha) w(M^*)$ where $\alpha = 1 + \log_2(1 + \varepsilon)$. Again, we guess the value of ω as in Section 3 by using an $n^{O(1)}$ -approximate matching cost. To obtain an $n^{O(1)}$ -approximate matching, we use the algorithm described in [30] in conjunction with approximate nearest neighbor data structure in a straight-forward way.

Data structure. We now describe the data structure used by PHASE-MATCH algorithm. It uses, as a black-box, a data structure to answer approximate nearest-neighbor (ANN) queries for $d(\cdot, \cdot)$. As proved in [30], an ANN data structure for $d(\cdot, \cdot)$ can also answer queries under $\mathbf{d}_i(\cdot, \cdot)$ — i.e., the ceiling operator can be ignored. So, we assume that we have an ANN data structure on $\mathbf{d}_i(\cdot, \cdot)$ and that it can answer a query and delete a point in $\tau(n, \varepsilon)$ time.

For two integers $i, j, i \leq j$, we use $[i : j]$ to represent all integers between i and j ; we also view $[i : j]$ as in interval. Two intervals $X = [i_1 : j_1]$ and $Y = [i_2 : j_2]$ with $i_2 \geq i_1$ are called ρ -well separated (or simply well-separated) if $[i_1 : j_1] \cap [i_2 : j_2] = \emptyset$ and $j_1 - i_1, j_2 - i_2 \leq \rho(i_2 - j_1)$. We refer to $j_1 - i_1$ as the length of X . A family $\mathbb{I} = \{(X_1, Y_1), (X_2, Y_2), \dots, (X_u, Y_u)\}$ of pairs of intervals, with X_i lying to the left of Y_i , where $X_i = [x_i : x'_i]$ and $Y_i = [y_i : y'_i]$ is called a ρ -well separated pair decomposition of $[0 : N]$ if

- (i) Each X_i and Y_i is a ρ -well separated pair, and
- (ii) for any pair $x, y \in [0 : N]$, with $x < y$, there is an index i such that $x \in X_i$ and $y \in Y_i$.

This is the 1-dimensional version of the general well-separated pair decomposition of a point set; see [19, Chapter III]. A standard binary decomposition can generate \mathbb{I} such that (a) $u = O(N/\rho)$, and (b) for any integer $x \in [0 : N]$, there are $O(\rho^{-1} \log N)$ pairs in \mathbb{I} such that x is contained in one of the two intervals in the pair.

Suppose $K \subseteq A_i$ is the set of m points on which we wish to build the data structure. We set $\rho = \varepsilon/10$ and $N = \frac{5}{\varepsilon}n$ (it can be shown that the maximum absolute value of any dual weight is at most $\frac{5}{\varepsilon}n$), and construct, in $O(n/\varepsilon^2 \log(n/\varepsilon))$ time, a WSPD \mathbb{I} of $[0 : N]$. For every pair in $(X_i, Y_i) \in \mathbb{I}$, let $K_i = \{a \mid a \in A_i, |y(a)| \in X_i\}$; recall that dual weights of all points of A are non-positive but we consider their absolute values. For every i , we maintain a dynamic $(\varepsilon/2)$ -ANN data structure on K_i . Let $\tau(m, \varepsilon)$ be the query and update time of this data structure. DELETE operation is straight-forward and takes $\tilde{\tau}(m, \varepsilon) = O(\varepsilon^{-1} \log(n/\varepsilon) \tau(m, \varepsilon))$ time. ELIGIBLE(b) works as follows: Identify all pairs $(X_i, Y_i) \in \mathbb{I}$ such that $y(b) \in Y_i$.² Let $Q \subseteq \mathbb{I}$ be this set. For each pair in $(X_i, Y_i) \in Q$, retrieve an $(\varepsilon/2)$ -ANN a_i of b from the corresponding ANN data structure. If the edge (a_i, b) , for

²Note that, for every eligible edge, $d(\cdot, \cdot)$ is positive, so $y(b) > |y(a)|$ and no need to consider the case where $y(b) \in X_i$.

any of these retrieved a_i 's is weakly or strongly eligible, the procedure returns a_i and NULL otherwise. The time taken by the procedure is $O(\varepsilon^{-1} \log(n/\varepsilon) \tau(m, \varepsilon))$.

The following lemma establishes the correctness of ELIGIBLE procedure.

LEMMA 4.2. ELIGIBLE(b) returns NULL only when there are no strongly eligible edges of the form (a, b) where $a \in K$.

PROOF. Suppose ELIGIBLE(b) returns NULL and there is a strongly eligible edge incident on b . Let (a, b) be this edge, i.e., $y(a) + y(b) = \mathbf{d}_i(a, b) + \lceil \varepsilon \mathbf{d}_i(a, b) \rceil$. Since $y(a) \leq 0$, we can write this as

$$|y(b)| - |y(a)| \geq (1 + \varepsilon) \mathbf{d}_i(a, b). \quad (19)$$

By property (i) of \mathbb{I} and since $|y(a)|, |y(b)|$ are integers at most N , there is a pair $(X_i, Y_i) \in \mathbb{I}$ such that $|y(a)| \in X_i, |y(b)| \in Y_i$. Since X_i and Y_i are $(\varepsilon/10)$ -well separated, the length of X_i (resp. Y_i) is at most $(\varepsilon/10)(|y(b)| - |y(a)|) \leq (1 + \varepsilon)(\varepsilon/10) \mathbf{d}_i(a, b) \leq (\varepsilon/5) \mathbf{d}_i(a, b)$. Consider the corresponding ANN data structure for K_i and observe that a is in this data structure. Let a' be the $\varepsilon/2$ -approximate nearest neighbor of b returned by the ANN data structure. Then $\mathbf{d}_i(a', b) \leq (1 + \varepsilon/2) \mathbf{d}_i(a, b)$. Since our procedure returns NULL, (a', b) is not eligible and therefore

$$\begin{aligned} |y(b)| - |y(a')| &\leq \mathbf{d}_i(a', b) \leq (1 + \varepsilon/2) \mathbf{d}_i(a, b) \\ &\leq |y(b)| - |y(a)| - (\varepsilon/2) \mathbf{d}_i(a, b) \end{aligned}$$

or $|y(a')| - |y(a)| \geq (\varepsilon/2) \mathbf{d}_i(a, b)$ which is a contradiction of the fact that the length of interval X_i that contains the $|y(a)|, |y(a')|$ is $(\varepsilon/10) \mathbf{d}_i(a, b)$. \square

Putting everything together, we obtain the following.

THEOREM 4.3. Given two point sets A and B of size n each, a metric $d(\cdot, \cdot)$ that can be evaluated in $O(1)$ time, two parameters $\varepsilon, \delta \in (0, 1)$, and an ANN data structure under $d(\cdot, \cdot)$, an $O(1/\delta^\alpha)$ -approximate matching of A and B can be computed in $O(\varepsilon^{-2} n^{1+\delta} \tau(n, \varepsilon) \log^2(n/\varepsilon) \log(1/\delta))$ time, where $\alpha = 1 + \log_2(1 + \varepsilon)$ and $\tau(n, \varepsilon)$ is the query and update time of an $\varepsilon/2$ -ANN data structure.

Using existing dynamic data structures for ANN queries, we obtain the results mentioned in Introduction for geometric settings. Because of lack of space, we omit the details from here.

5. MATCHING WITH ANN UNDER AN ARBITRARY COST FUNCTION

We present an algorithm to compute a $(1 + \varepsilon)$ -approximate matching for any cost function $d(\cdot, \cdot)$, provided we have a data structure that can answer ε/c -ANN queries under $d(\cdot, \cdot)$ in $\tau(n, \varepsilon)$ time, for some $c > 1$. While this algorithm is slower than the one presented in Section 4, it does not require the metric assumption on $d(\cdot, \cdot)$ and the approximation factor is much smaller.

At a high-level, this algorithm can be viewed as the execution of the phase 0 of the algorithm of Section 4, but instead of running it for at most $\frac{30}{\varepsilon} n^\delta$ stages, we let it run for \sqrt{n} stages. When the algorithm stops, it has computed, in $O(\varepsilon^{-1} n^{3/2} \tau(n, \varepsilon) \log n)$ time, a partial $(\varepsilon/2)$ -relaxed matching M_1 of A and B so that only $O(\sqrt{n})$ points of A and B are free — this is similar to a number of previous algorithms; see

e.g. [14, 20, 30]. For the remaining $O(\sqrt{n})$ free points, unlike Section 4, we do not commit M_1 . Instead we iteratively compute augmenting paths using Hungarian search, starting with M_1 , and compute a perfect $(\varepsilon/2)$ -relaxed matching M of A and B . By Lemma 2.1, M is an ε -approximate matching of A, B .

In the rest of the section, we focus on the Hungarian search procedure. At each stage, it takes a partial $(\varepsilon/2)$ -relaxed matching and its dual weights as input, adjusts the dual weights and computes an augmenting path consisting only of strongly and weakly eligible edges. The search procedure maintains an alternating tree consisting of eligible edges rooted at the free vertices of B . Let S be the set of points of B that are in the alternating tree and T be the set of points of A that are not in any alternating tree constructed so far. Initially $S = B^F$ and $T = A$ (or $A \setminus T = \emptyset$). For simplicity, with a slight abuse of notation, we use $d(\cdot, \cdot)$ to denote $d_0(\cdot, \cdot)$.

For a pair $(a, b) \in A \times B$, let $\xi(a, b) = d(a, b) - y(a) - y(b)$. Due to $\varepsilon/2$ -relaxed conditions, $\xi(a, b) \geq -(\varepsilon/2)d(a, b)$. Let $a^*, b^* = \operatorname{argmin}_{a \in T, b \in S} \xi(a, b)$ be the *weighted bi-chromatic closet pair* (BCP) of S and T . If $S \times T$ contains a strongly eligible edge, then $\xi(a^*, b^*) = -(\varepsilon/2)d(a, b)$.

We define a weaker notion of δ -approximate BCP (or δ -BCP for brevity), which is sufficient for our algorithm. A pair $\tilde{a} \in T$ and $\tilde{b} \in S$ is called a δ -BCP if $\xi(\tilde{a}, \tilde{b}) \leq \xi(a, b) + \delta d(a, b)$ for any pair $(a, b) \in S \times T$. We describe a data structure which using an $(\delta/20)$ -ANN data structure maintains an δ -approximate BCP under the following three operations, each of which takes $\tilde{\tau}(n, \varepsilon)$ time.

- (i) INSERT(b): Insert a weighted point b with weight $y(b)$ to S .
- (ii) DELETE(a): Delete a point a from T .
- (iii) WT-UPDATE(Δ): For each point $b \in S$, set its weight $y(b) = y(b) + \Delta$.

For our algorithm, we set up $\delta = \varepsilon/2$. Using an $(\varepsilon/2)$ -BCP data structure, we execute Hungarian search as follows: Let (\tilde{a}, \tilde{b}) be the pair maintained by the data structure. There are two cases: First, if $\xi(\tilde{a}, \tilde{b}) \leq 0$, then (\tilde{a}, \tilde{b}) is weakly eligible, and we add (\tilde{b}, \tilde{a}) to the alternating tree and remove \tilde{a} from T . If \tilde{a} is unmatched, we have found an augmenting path and the search terminates. Otherwise, let (\tilde{a}, b') be the edge in the current matching. We add (\tilde{a}, b') to the alternating tree, and add b' to S .

Next, if $\xi(\tilde{a}, \tilde{b}) > 0$, i.e., (\tilde{a}, \tilde{b}) is not eligible, we set the dual weight of every point $b \in S$ to $y(b) = y(b) + \xi(\tilde{a}, \tilde{b})$ and the dual weight of every point $a \in A \setminus T$ to $y(a) = y(a) - \xi(\tilde{a}, \tilde{b})$. We update the weights in the data structure using the WT-UPDATE procedure.³ After this step (\tilde{a}, \tilde{b}) is weakly eligible, and we proceed as above.

Following the proofs of [1, 30, 32], it can be shown that the search grows alternating tree without violating the $(\varepsilon/2)$ -relaxed constraints and terminates when an augmenting path is found. Omitting the details, the total time spent in computing all $O(\sqrt{n})$ augmenting path is $O(n^{3/2}\tilde{\tau}(n, \varepsilon))$.

Recall that the phase 0 of the algorithm described in Section 4 assumes that we have a value $\omega \in [\underline{w}(M^*)/2, \overline{w}(M^*)]$.

³Since points in $A \setminus T$ do not participate in computing \tilde{a} and \tilde{b} , we can update their dual weights at the end of the search procedure. The total time taken for this step would be $O(n)$.

Let $\Delta = \frac{\max d(a, b)}{\min d(a, b)}$, where the maximum and minimum are taken over all distinct pairs of $A \times B$, is the *spread* of $A \times B$. Then we run the above algorithm on $O(\log \Delta)$ different choices of ω and choose the minimum-cost matching among them. Putting everything together, the total running time of the algorithm is $O(n^{3/2}\tilde{\tau}(n, \varepsilon) \log \Delta)$.

BCP data structure. We now describe the data structure for maintaining a δ -BCP of S and T under the operations (i)–(iii). For a weighted point $q \in A \cup B$ and for a subset $P \subset A \cup B$, we define a δ -approximate weighted nearest neighbor (or δ -WNN) in P to be a point $\tilde{p} \in P$ such that $\xi(q, \tilde{p}) \leq \min_{p \in P} \xi(p, q) + \delta d(p, q)$. The heart of our data structure is a dynamic data structure for answering δ -WNN queries on a subset of S or of T . For now suppose we have such a data structure. Eppstein [10] has described a data structure for maintaining the BCP of two point sets under insertion and deletion of points, for any given distance function provided that there is a data structure for answering the nearest-neighbor queries under that distance function. It basically maintains a set of $O(n \log n)$ candidate pairs of points for bi-chromatic closet pair (BCP), using a dynamic data structure for nearest-neighbor queries. It can be shown that if we plug our data structure for answering δ -WNN queries, then his data structure can be adapted in a straight-forward manner to maintain a δ -BCP of $S \times T$. In particular, if the query and update time of the δ -WNN data structure on a set of m points is $O(\varphi(m, \delta))$, then a δ -BCP of $S \times T$ can be maintained in $O(\varphi(n, \delta) \log^2 n)$ time. We omit the details from this version.

δ -WNN data structure. Let $P \subseteq S$ be a set of m weighted points, and let $\delta > 0$ be a parameter. We describe how to answer a δ -WNN query for a point $q \in T$. The case when $P \subseteq T$ is symmetric. Let $w(p)$ denote the weight of a point $p \in P$. Since the weight of the query point q does not affect the choice of WNN of q in P , we assume $w(q) = 0$ and set $\xi(p, q) = d(p, q) - w(p)$. We focus on insertion/deletion of points in P . The weight update operation can be handled using the technique by Vaidya [32] so that the weights of points in P need not be updated over time. We refer to [32] for details.

Data structure. The weights of points in P lie in the range $[-N, N]$ for $N = 5n^2/\varepsilon$. For simplicity assume that N is of the form 2^k for some integer k . Let \mathcal{J} be a 1D quad-tree on the interval $[-2N, 2N]$, whose leaves correspond to intervals of the form $[i, i + 1]$. Each node $v \in \mathcal{J}$ is associated with a *canonical interval* J_v . For each node $v \in \mathcal{J}$, we associate a *canonical subset* $P_v \subseteq P$ where $P_v = \{p \in P \mid w(p) \in J_v\}$. We do not construct the entire \mathcal{J} but only those nodes v for which $P_v \neq \emptyset$. Set $\delta' = \delta/20$. We maintain P_v in a dynamic δ' -ANN data structure that can handle insertion/deletion of points.

This completes the description of the data structure. A point can be inserted into and deleted from the data structure in a straight-forward manner.

Query procedure. Let q be a query point, and let $\varphi(q) = \min_{p \in P} \xi(p, q)$. Before describing the query procedure, we briefly sketch the main idea. Suppose we know $\varphi(q)$. It can be shown that for any point $p \in P$, $|\varphi(q) + w(p)| \leq d(p, q)$, so we can cluster the points of P whose weights differ by at most $\delta'|\varphi(q) + w(p)|$ into a single cluster, compute a δ' -ANN for each cluster, and return the nearest one among

them. But we do not know $\varphi(q)$, so we guess its value and repeatedly refine our guess. We show that the procedure terminates within $O(\log N)$ steps, and at the end we have a good estimate of $\varphi(q)$ as well as a δ -WNN of q in P . We now describe the procedure in detail.

We maintain a value $\chi = \xi(q, \tilde{p})$ for some point $\tilde{p} \in P$ as an estimate of $\varphi(q)$. Initially, we choose \tilde{p} to be an arbitrary point of P . By identifying the leaf of \mathcal{T} that contains $-\chi$, we compute a partition of P into a family of $u = O(\frac{1}{\delta} \log N)$ canonical subsets $\mathcal{P} = \{P_1, \dots, P_u\}$ with corresponding canonical intervals J_1, \dots, J_u such that for any $i \leq u$,

$$|J_u| \leq \delta' |\chi + w(p)| \quad \forall p \in P_i. \quad (20)$$

See e.g. [19]. For each $i \leq u$, we compute a δ' -ANN p_i of q in P_i . Set $K = \{p_i \mid 1 \leq i \leq u\}$ and $p^* = \arg \min_{p \in K \cup \{\tilde{p}\}} \xi(q, p)$. If

$$\chi \leq \min_{p \in K} \xi(q, p) + \frac{\delta}{10} \mathbf{d}(q, a), \quad (21)$$

then the algorithm returns p^* ; otherwise ($p^* \neq \tilde{p}$), we set $\tilde{p} = p^*$ and $\chi = \xi(q, p^*)$, and repeat the procedure. This completes the description of the query procedure.

LEMMA 5.1. *The point p^* returned by the procedure is a δ -WNN(q).*

PROOF. Suppose on the contrary, there is a point $p' \in P$ such that $\xi(p, q) > \xi(p', q) + \delta \mathbf{d}(p', q)$. Let χ^* be the final value of χ when the algorithm terminates. Then

$$\xi(p^*, q) \leq \chi^* \leq \xi(p^*, q) + \frac{\delta}{10} \mathbf{d}(p', q). \quad (22)$$

Therefore

$$\chi^* + w(p') > (1 + \delta) \mathbf{d}(p', q). \quad (23)$$

Let $P_i \in \mathcal{P}$ be the canonical subset that contains p' . Since p_i is a δ' -ANN of q in P_i ,

$$\mathbf{d}(p_i, q) \leq (1 + \delta') \mathbf{d}(p', q). \quad (24)$$

Furthermore, by (21),

$$\chi^* + w(p_i) \leq (1 + \delta/10) \mathbf{d}(p_i, q) \leq (1 + \delta/4) \mathbf{d}(p', q). \quad (25)$$

Since both $p', p_i \in P_i$, by (20),

$$|w(p_i) - w(p')| \leq \delta' \cdot \min\{|\chi^* + w(p')|, |\chi^* + w(p_i)|\}. \quad (26)$$

Using (23)–(26), one can prove that

$$|w(p') - w(p_i)| \leq \frac{\delta}{4} \mathbf{d}(q, p').$$

Therefore,

$$\begin{aligned} \xi(p^*, q) &\leq \xi(p_i, q) = \mathbf{d}(p_i, q) - w(p_i) \\ &\leq (1 + \delta') \mathbf{d}(p', q) - w(p') + \frac{\delta}{4} \mathbf{d}(q, p') \\ &< \xi(p', q) + \delta \mathbf{d}(p', q), \end{aligned}$$

a contradiction. Hence, p^* is a δ -WNN of q in P_i . \square

LEMMA 5.2. *The query procedure terminates in $O(\log N)$ iterations.*

PROOF. Let $\hat{p} = \arg \min_{p \in P} \xi(q, a) - \delta' \mathbf{d}(q, a)$. Suppose χ_{curr} is the value of χ in an iteration of the query procedure at which it does not terminate. Let \mathcal{P} be the family of canonical subsets computed in this iteration, and let $P_i \in \mathcal{P}$

be the canonical subset that contains \hat{p} . Then $\mathbf{d}(q, p_i) \leq (1 + \delta') \mathbf{d}(q, \hat{p})$. Since both $p_i, \hat{p} \in P_i$, by (20), $|w(p_i) - w(\hat{p})| \leq \delta' |\chi_{\text{curr}} + w(\hat{p})|$. This implies that the new value of χ , denoted by χ_{new} , satisfies

$$\chi_{\text{new}} \leq \xi(\hat{p}, q) + \delta' \mathbf{d}(\hat{p}, q) + \delta' |\chi_{\text{curr}} + w(\hat{p})|.$$

If $|\chi_{\text{curr}} + w(\hat{p})| \leq 2\mathbf{d}(q, \hat{p})$, then it can be verified that χ_{new} satisfies (21), and the algorithm stops after the next iteration.

If $|\chi_{\text{curr}} + w(\hat{p})| > 2\mathbf{d}(q, \hat{p})$, then

$$|\chi_{\text{new}} + w(\hat{p})| \leq (1 + \delta') \mathbf{d}(q, \hat{p}) + \delta' |\chi_{\text{curr}} + w(\hat{p})|.$$

Therefore

$$\frac{|\chi_{\text{new}} + w(\hat{p})|}{|\chi_{\text{curr}} + w(\hat{p})|} - \delta' \leq \frac{1}{2} (1 + \delta').$$

Since $\delta' = \delta/20 \leq 1/20$, $|\chi_{\text{new}} + w(\hat{p})| \leq \frac{3}{4} |\chi_{\text{curr}} + w(\hat{p})|$. The query procedure thus stops in $O(\log N)$ iterations. \square

Hence, the query procedure takes $O(\delta^{-1} \tau(m, \delta) \log^2 N)$ time, assuming that a δ' -ANN query can be answered in $\tau(m, \delta)$ time. Putting everything together we obtain the following.

THEOREM 5.3. *Given $G(A \cup B, A \times B)$ and a distance function $\mathbf{d}(\cdot, \cdot)$ on the edges, for a parameter $0 < \varepsilon \leq 1$, a $(1 + \varepsilon)$ -approximate matching of A and B can be computed in $O(\varepsilon^{-1} n^{3/2} \tau(n, \varepsilon) \log^4(n/\varepsilon) \log \Delta)$ time, where $\tau(n, \varepsilon)$ is the query and update time of an (ε/c) -approximate nearest neighbor data structure, for some constant $c > 1$.*

6. CONCLUSION

In this paper, we presented an $O(n^{2+\delta})$ time algorithm to compute an $O(1/\delta^{0.631})$ -approximate matching of A, B , under an arbitrary metric. This leads to a fast approximation algorithm for computing a bipartite matching under several useful distance functions such as edit distance, hamming distance etc. We also presented faster approximation algorithms for computing bipartite matchings under cost functions that support approximate nearest-neighbor data structure. As a consequence, we obtained improved approximation algorithms for bipartite matching of points in fixed and high-dimensional Euclidean spaces. We conclude by mentioning a few open problems:

- Is there an $O(n^2)$ -time algorithm for computing a $(1+\varepsilon)$ -approximate matching, for any constant $\varepsilon > 0$, in an arbitrary metric space?
- Is there a $o(n^{3/2})$ -time algorithm for computing a $O(1)$ -approximate bipartite matching of points under the squared Euclidean distance function?
- Can the optimal Euclidean bipartite matching between two points sets in \mathbb{R}^2 be computed in subquadratic time?

Acknowledgments. Work by Pankaj Agarwal is supported by NSF under grants CCF-09-40671, CCF-10-12254, and CCF-11-61359, by an ARO grant, by Grant 2012/229 from the U.S.-Israel Binational Science Foundation, and by an ERDC contract W9132V-11-C-0003. Work by R. Sharathkumar is supported by NSF grants FODAVA 808515, CCF 808515, and 1161480, and by the grant FA9550-12-1-XXXX from the U.S. Air Force Office of Scientific Research (AFOSR) and the Defense Advanced Research Projects Agency (DARPA).

7. REFERENCES

- [1] P. K. Agarwal, A. Efrat, and M. Sharir, Vertical decomposition of shallow levels in 3-dimensional arrangements and its applications, *SIAM J. Comput.*, 29 (1999), 39–50.
- [2] P. K. Agarwal and K. R. Varadarajan, A near-linear constant-factor approximation for euclidean bipartite matching?, *Proc. 12th Annual Sympos. Comput. Geom.*, 2004, pp. 247–252.
- [3] A. Andoni, P. Indyk, and R. Krauthgamer, Earth mover distance over high-dimensional spaces, *Proc. 19th Annual ACM-SIAM Sympos. Discrete Algo.*, 2008, pp. 343–352.
- [4] S. Arora, Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems, *J. ACM*, 45 (1998), 753–782.
- [5] D. Avis, Two greedy heuristics for the weighted matching problem, *Proc. of 9th S. E. Conf. on Comb., Graph Theory, and Computing*, 1978, pp. 65–76.
- [6] H.-T. Chen, H.-H. Lin, and T.-L. Liu, Multi-object tracking using dynamical graph matching, *IEEE Conference Comp. Vision Pattern Recog.*, 2 (2001), 210.
- [7] R. Cole, R. Hariharan, M. Lewenstein, and E. Porat, A faster implementation of the Goemans-Williamson clustering algorithm., *Proc. 11th Annual Sympos. Discrete Algo.*, 2001, pp. 17–25.
- [8] M. Cygan, H. N. Gabow, and P. Sankowski, Algorithmic applications of Baur-Strassen’s theorem: Shortest cycles, diameter and matchings, *Proc. 54th Annual IEEE Sympos. Found. Comp. Sci.*, 2012, pp. 531–540.
- [9] R. Duan and S. Pettie, Linear-time approximation for maximum weight matching, *J. ACM*, 61 (2014), 1:1–1:23.
- [10] D. Eppstein, Dynamic euclidean minimum spanning trees and extrema of binary functions, *Discrete Comput. Geom.*, 13 (1995), 111–122.
- [11] J. Fakcharoenphol, S. Rao, and K. Talwar, A tight bound on approximating arbitrary metrics by tree metrics, *Proc. 35th Annual ACM Sympos. Theory Comput.*, 2003, pp. 448–455.
- [12] M. L. Fredman and R. E. Tarjan, Fibonacci heaps and their uses in improved network optimization algorithms, *J. ACM*, 34 (1987), 596–615.
- [13] H. Gabow and S. Pettie, The dynamic vertex minimum problem and its application to clustering-type approximation algorithms, in: *Proc. 8th Scandinavian Workshop Algo. Theory*, 2002, pp. 190–199.
- [14] H. N. Gabow and R. Tarjan, Faster scaling algorithms for network problems, *SIAM J. Comput.*, 18 (1989), 1013–1036.
- [15] H. N. Gabow and R. E. Tarjan, Faster scaling algorithms for general graph-matching problems, *J. ACM*, 38 (1991), 815–853.
- [16] A. Goel, P. Indyk, and K. Varadarajan, Reductions among high dimensional proximity problems, *Proc. 12th Annual ACM-SIAM Sympos. Discrete Algo.*, 2001, pp. 769–778.
- [17] A. Goel, M. Kapralov, and S. Khanna, Perfect matchings in $O(n \log n)$ time in regular bipartite graphs, *Proc. 42nd Annual Sympos. Theory Comput.*, 2010, pp. 39–46.
- [18] M. Goemans and D. Williamson, A general approximation technique for constrained forest problems, *SIAM J. Comput.*, 24 (1995), 296–317.
- [19] S. Har-Peled, *Geometric Approximation Algorithms*, American Mathematical Society, 2011.
- [20] J. Hopcroft and R. Karp, An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs, *SIAM J. Comput.*, 2 (1973), 225–231.
- [21] P. Indyk, A near linear time constant factor approximation for Euclidean bichromatic matching (cost), *Proc. 18th Annual ACM-SIAM Sympos. Discrete Algo.*, 2007, pp. 39–42.
- [22] A. Madry, Navigating central path with electrical flows: from flows to matchings, and back point clouds, *Proc. 54th Annual IEEE Sympos Foundat. Comp. Sc.*, 2013, pp. 253–262.
- [23] S. Micali and V. V. Vazirani, An $o(\sqrt{v})e$ algorithm for finding maximum matching in general graphs, *Proc. 21st Annual IEEE Sympos. Foundat. Comp. Sc.*, 1980, pp. 17–27.
- [24] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*, Prentice-Hall, Inc., 1982.
- [25] D. A. Plaisted, Heuristic matching for graphs satisfying the triangle inequality, *J. Algorithms*, 5 (1984), 163 – 179.
- [26] E. M. Reingold and R. E. Tarjan, On a greedy heuristic for complete matching, *SIAM J. Comput.*, (1981), 676–681.
- [27] P. Sankowski, Weighted bipartite matching in matrix multiplication time, in: *Proc. 33rd Intn. Conf. Automata, Languages and Programming*, Vol. 4051, 2006, pp. 274–285.
- [28] P. Sankowski, Maximum weight bipartite matching in matrix multiplication time, *Theo. Comput. Sci.*, 410 (2009), 4480 – 4488.
- [29] R. Sharathkumar, A sub-quadratic algorithm for bipartite matching of planar points with bounded integer coordinates, *Proc. 29th Annual Sympos. Comput. Geom.*, 2013, pp. 9–16.
- [30] R. Sharathkumar and P. K. Agarwal, Algorithms for transportation problem in geometric settings, *Proc. 23rd Annual ACM-SIAM Sympos. Discrete Algo.*, 2012, pp. 306–317.
- [31] R. Sharathkumar and P. K. Agarwal, A near-linear time approximation algorithm for geometric bipartite matching., *Proc. 44th Annual ACM Annual Sympos. Theory Comput.*, 2012, pp. 385–394.
- [32] P. M. Vaidya, Geometry helps in matching, *SIAM J. Comput.*, 18 (1989), 1201–1225.
- [33] K. R. Varadarajan, A divide-and-conquer algorithm for min-cost perfect matching in the plane, *Proc. 39th Annual IEEE Sympos. Foundat. Comp. Sc.*, 1998, pp. 320–331.
- [34] K. R. Varadarajan and P. K. Agarwal, Approximation algorithms for bipartite and non-bipartite matching in the plane, *Proc. 10th Annual ACM-SIAM Sympos. Discrete Algo.*, 1999, pp. 805–814.