Matchings under distance constraints I.

Péter Madarasi¹

Accepted: 15 May 2021 / Published online: 28 June 2021
© The Author(s) 2021

Abstract

This paper introduces the *d-distance matching problem*, in which we are given a bipartite graph $G = (S, T; E)$ with $S = \{s_1, \dots, s_n\}$, a weight function on the edges and an integer $d \in \mathbb{Z}_+$. The goal is to find a maximum-weight subset $M \subseteq E$ of the edges satisfying the following two conditions: (i) the degree of every node of S is at most one in M , (ii) if $s_i t, s_j t \in M$, then $|j - i| \geq d$. This question arises naturally, for example, in various scheduling problems. We show that the problem is NP-complete in general and admits a simple 3-approximation. We give an FPT algorithm parameterized by d and also show that the case when the size of T is constant can be solved in polynomial time. From an approximability point of view, we show that the integrality gap of the natural integer programming model is at most $2 - \frac{1}{2d-1}$, and give an LP-based approximation algorithm for the weighted case with the same guarantee. A combinatorial $(2 - \frac{1}{d})$ -approximation algorithm is also presented. Several greedy approaches are considered, and a local search algorithm is described that achieves an approximation ratio of $3/2 + \epsilon$ for any constant $\epsilon > 0$ in the unweighted case. The novel approaches used in the analysis of the integrality gap and the approximation ratio of locally optimal solutions might be of independent combinatorial interest.

Keywords Distance matching · Restricted *b*-matching · Constrained matching · Scheduling · Parameterized algorithms · Approximation algorithms · Integrality gap

1 Introduction

In the *perfect d-distance matching problem*, one is given a bipartite graph $G = (S, T; E)$ with $S = \{s_1, \dots, s_n\}$, $T = \{t_1, \dots, t_k\}$, a weight function on the edges $w : E \rightarrow \mathbb{R}_+$ and an integer $d \in \mathbb{Z}_+$. The goal is to find a maximum-weight subset $M \subseteq E$ of the edges such that the degree of every node of S is one in M and if $s_i t, s_j t \in M$, then $|j - i| \geq d$. In the (non-perfect) *d-distance matching problem*, some of the nodes of S might remain uncovered. Note

Supported by the ÚNKP-20-3 New National Excellence Program of the Ministry for Innovation and Technology from the source of the National Research, Development and Innovation Fund. The present paper is an extended and revised version of a conference proceeding presented at ISCO 2020.

✉ Péter Madarasi
madarasi@cs.elte.hu

¹ Department of Operations Research, Eötvös Loránd University, Budapest, Hungary

Fig. 1 **a** A feasible perfect 3-distance matching. **b** An infeasible 3-distance matching

that the order of nodes in $S = \{s_1, \dots, s_n\}$ affects the set of feasible d -distance matchings, but the order of $T = \{t_1, \dots, t_k\}$ is indifferent. For example, Fig. 1a is a feasible perfect 3-distance matching, but the example shown in Fig. 1b is not, because edges s_1t_2 and s_3t_2 violate the 3-distance condition.

An application of this problem for $w \equiv 1$ is as follows. Imagine n consecutive all-day events $s_1, \dots, s_n$ each of which must be assigned one of k watchmen $t_1, \dots, t_k$. For each event s_i , a set of possible watchmen is given — those who are qualified to be on guard at event s_i . Assign exactly one watchman to each of the events such that no watchman is assigned to more than one of any d consecutive events, where $d \in \mathbb{Z}_+$ is given. In the weighted version of the problem, let $w_{s_it_j}$ denote the level of safety of event s_i if watchman t_j is on watch, and the objective is to maximize the level of overall safety.

As another application of the above question, consider n items $s_1, \dots, s_n$ one after another on a conveyor belt, and k machines $t_1, \dots, t_k$. Each item s_i is to be processed on the conveyor belt by one of the qualified machines $N(s_i) \subseteq \{t_1, \dots, t_n\}$ such that if a machine processes item s_i , then it can not process any of the next $d - 1$ items — because the conveyor belt is running.

Motivated by the first application, in the *cyclic d -distance matching problem* the nodes of S are considered to be in cyclic order. The focus of this paper is on the above (perfect) d -distance matching problem, but some of the proposed approaches also apply for the cyclic case. In particular, the 3-approximation greedy algorithm achieves the same guarantee for the weighted cyclic case (see Sect. 3.3), and the $(3/2 + \epsilon)$ -approximation algorithm for the unweighted case (see Sect. 4.2).

Previous work Observe that in the special case $d = |S|$, one gets the classic (perfect) bipartite matching problem. For $d = 1$, the problem reduces to the b -matching problem, and one can show that it is a special case of the circulation problem for $d = 2$. This implies that the problem is solvable in strongly polynomial time for $d = 1, 2$, since the circulation problem can be solved in strongly polynomial time Tardos (1985) and the b -matching problem is a special case of it.

A feasible d -distance matching M can be thought of as a b -matching that contains none of the subgraphs $\{(s_i, s_j), \{t\}; \{s_it, s_jt\} : s_it, s_jt \in E \text{ and } |i - j| \leq d\}$, where $b_s = 1$ for $s \in S$ and $b_t = |S|$ for $t \in T$. A similar problem is the $K_{p,p}$ -free p -matching problem Makai (2007). Here one is given an arbitrary family $\mathcal{T}$ of the subgraphs of G isomorphic with $K_{p,p}$, and the goal is to find a maximum-cardinality b -matching which induces no subgraph of $\mathcal{T}$, where $b : S \cup T \rightarrow \{0, \dots, p\}$. This problem can be solved in polynomial time. Note that in the distance matching problem, b is different, and the type of the forbidden subgraphs is $K_{2,1}$.

Another similar problem is the following. Given a partition $E_1, \dots, E_k$ of E and positive integers $r_1, \dots, r_k$, find a perfect matching M for which $|M \cap E_i| \leq r_i$. The problem is

introduced and shown to be NP-complete in Itai et al. (1978). Note that the side constraints in the distance matching problem are similar, but the degree constraints are different and our edge sets do not form a partition of E .

Several other versions of the “restricted” (b)-matching problem have been introduced, for example in Baste et al. (2019); Bérczi and Végő (2010); Fürst and Rautenbach (2019); Pap (2005).

The perfect d -distance matching problem is a special case of the list-coloring problem on interval graphs Zeithofer and Wess (2003). Here a proper vertex coloring must be found for which the color of each node v is chosen from a predefined list of colors C_v . Given an instance of the d -distance matching problem, we construct an interval graph $H = (V, F)$ such that there is a one-to-one correspondence between perfect d -distance matchings of $G = (S, T; E)$ and proper list colorings of H . Let the nodes of H be the intervals $\{R_d(s) : s \in S\}$, where $R_d(s_i) = \{s_i, \dots, s_{\min(i+d-1, |S|)}\}$ for $i = 1, \dots, n$, and let two distinct nodes $R_{s_i}, R_{s_j} \in V$ be connected by an edge if and only if $R_{s_i} \cap R_{s_j} \neq \emptyset$. Finally, let the list of possible colors of node $R_d(s_i)$ be the neighbors of s_i in G . Observe that two nodes $s_i, s_j \in S$ can be assigned to the same node of T in a distance matching M if and only if there is no edge between $R_d(s_i)$ and $R_d(s_j)$ in F . The latter holds, however, if and only if nodes $R_d(s_i) \in V$ and $R_d(s_j) \in V$ can be assigned color t simultaneously in a proper list coloring of H (note that t is in both lists of colors). Hence, there is a one-to-one correspondence between the perfect d -distance matchings of G and the proper list colorings of H .

The perfect d -distance matching problem is also a special case of the frequency assignment problem Aardal et al. (2007). Let $S = \{s_1, \dots, s_n\}$ be a set of antennas and let $T = \{t_1, \dots, t_k\}$ be a set of frequencies. There is an edge between $s \in S$ and $t \in T$ if antenna s can be set to frequency t . We are also given an interference graph of the antennas, in which two antennas are connected if and only if they may interfere with each other. The goal is to assign a frequency to each antenna such that no two interfering antennas are assigned the same (or in a certain sense similar, see Aardal et al. (2007)) frequency. To reduce the d -distance matching problem to the frequency assignment problem, let two antennas s_i, s_j interfere if and only if $|i - j| \leq d$. This corresponds to the setting on the plane when antennas $s_1, \dots, s_n$ are located along a straight line in this order such that the Euclidean distance between s_i and s_{i+1} is one for $i = 1, \dots, n - 1$, and two antennas may interfere if and only if their Euclidean distance is at most d . By this construction, there exists a feasible frequency assignment (in which no two interfering antennas are assigned the same frequency) if and only if there exists a perfect d -distance matching.

Our results This paper settles the complexity of the distance matching problem and gives an FPT algorithm parameterized by d . An efficient algorithm for constant T is also given. We present an LP-based $(2 - \frac{1}{2d-1})$ -approximation algorithm for the weighted distance matching problem, which implies that the integrality gap of the natural IP model is at most $2 - \frac{1}{2d-1}$. An interesting alternative proof for the integrality gap is also given. We also describe a combinatorial $(2 - \frac{1}{d})$ -approximation algorithm for the weighted case. One of the main contributions of the paper is a $(3/2 + \epsilon)$ -approximation algorithm for the unweighted case for any constant $\epsilon > 0$ in the unweighted case. The proof is based on revealing the structure of locally optimal solutions recursively. A generalization of König’s edge-coloring theorem Frank (2011, p. 74) is given to the distance matching problem, as well.

Notation Throughout the paper, assume that $G = (S, T; E)$ contains no loops or parallel edges, unless stated otherwise. Let $\Delta(v)$ and $N(v)$ denote the set of incident edges

to node v and the neighbors of v , respectively. For a subset $X \subseteq E$ of the edges, $N_X(v)$ denotes the neighbors of v for edge set X . We use $\deg(v)$ to denote the degree of node v . Let $L_d(s_i) = \{s_{\max(i-d+1, 1)}, \dots, s_i\}$ and $R_d(s_i) = \{s_i, \dots, s_{\min(i+d-1, |S|)}\}$. The maximum of the empty set is $-\infty$ by definition. Given a function $f : A \rightarrow B$, both $f(a)$ and f_a denote the value f assigns to $a \in A$, and let $f(X) = \sum_{a \in X} f(a)$ for $X \subseteq A$. Let χ_Z denote the characteristic vector of set Z , i.e. $\chi_Z(y) = 1$ if $y \in Z$, and 0 otherwise. Occasionally, the braces around sets consisting of a single element are omitted, e.g. $\chi_e = \chi_{\{e\}}$ for $e \in E$.

2 Complexity

This section settles the complexity of the d -distance matching problem. First, we introduce the following NP-complete problem.

Lemma 1 *Given a bipartite graph $G = (S, T; E)$ and $S_1, S_2 \subseteq S$ such that $S_1 \cup S_2 = S$, it is NP-complete to decide if there exists $M \subseteq E$ for which $|M| = |S|$ and both $M \cap E_1$ and $M \cap E_2$ are matchings, where E_i denotes the edges induced by T and S_i for $i = 1, 2$. The problem remains NP-complete even if the maximum degree of the graph is at most 4.*

Proof We reduce the 3-Dimensional Matching problem to the problem defined in the lemma statement. Here, one is given three finite disjoint sets X, Y, Z and a set of hyperedges $\mathcal{H} \subseteq X \times Y \times Z$. A subset of the hyperedges $F \subseteq \mathcal{H}$ is called 3-dimensional matching if $x_1 \neq x_2$, $y_1 \neq y_2$ and $z_1 \neq z_2$ for any two distinct triples $(x_1, y_1, z_1), (x_2, y_2, z_2) \in F$. Being one of Karp's 21 NP-complete problems Karp (1972), it is NP-complete to decide whether there exists a 3-dimensional matching $F \subseteq \mathcal{H}$ of size $|Z|$. In fact, the problem remains NP-complete even if no element of $X \cup Y \cup Z$ occurs in more than three triples in $\mathcal{H}$ (Garey and Johnson (1979, p. 221)). Without loss of generality, one might assume that $|X| = |Y| = |Z|$. Let $\mathcal{H}_z = \{e_1^z, \dots, e_{k_z}^z\}$ denote the set of hyperedges incident to $z \in Z$, i.e. $\mathcal{H}_z = \mathcal{H} \cap (X \times Y \times \{z\})$ for each $z \in Z$. To reduce the 3-dimensional matching problem to the above problem, consider the following construction.

First define a bipartite graph $G = (S, T; E)$ where $S = X \cup (\mathcal{H} \setminus \{e_1^z : z \in Z\}) \cup Y$, $T = \mathcal{H}$ and E is as follows. For each $s \in S \cap (X \cup Y)$, add an edge between s and all the hyperedges $e \in T$ incident to s ; and connect each $e_i^z \in S \cap \mathcal{H}$ to hyperedges $e_{i-1}^z, e_i^z \in T$ for each $z \in Z$ and $i = 2, \dots, k_z$. Let $S_1 = S \setminus Y$ and $S_2 = S \setminus X$. Figure 2a and 2b show an instance of the 3-dimensional matching problem and the corresponding construction, respectively. Each hyperedge is represented by a unique line style, e.g. the dotted lines represent hyperedge $e_1 = (x_2, y_1, z_1)$ in Fig. 2a, and the dotted lines correspond to the same hyperedge e_1 in Fig. 2b. Note that the edges represented by a straight line in Fig. 2b do not represent hyperedges, but the edges between hyperedges. The highlighted edges in Fig. 2a and 2b correspond to the same feasible 3-dimensional matching.

Observe that there exists a 3-dimensional matching F of size $|Z|$ if and only if there exists $M \subseteq E$ for which $|M| = |S|$ and both $M \cap E_1$ and $M \cap E_2$ are matchings, where E_i denotes the edges incident to S_i ($i = 1, 2$). Indeed, such an $M \subseteq E$ matches all nodes of S into T , therefore there exists a unique hyperedge $e_z^* \in T \cap \mathcal{H}_z$ for each $z \in Z$ that is not matched to $S \cap \mathcal{H}$, but to exactly one element of $x \in X$ and exactly one element of $y \in Y$ (because all hyperedges in $S \cap \mathcal{H}_z$ are matched into $T \cap \mathcal{H}_z$, and these edges of M cover all but one hyperedge of $T \cap \mathcal{H}_z$). These three edges correspond to the inclusion of hyperedge (x, y, z) . This way one obtains a 3-dimensional matching F of size $|Z|$. On the other hand, if a 3-dimensional matching F is given for which $|F| = |Z|$, then one might easily construct

Fig. 2 Illustration of the proof of Lemma 1. Each hyperedge is represented by a unique line style. The highlighted hyperedges on (a) and the highlighted edges on (b) correspond to the same feasible solution

the desired $M \subseteq E$ as follows. For each $(x, y, z) \in F$, let i be the unique index such that $e_i^z = (x, y, z) \in F$ and extend M with edges $xe_i^z, ye_i^z \in E$. Let us also include a perfect matching between $S \cap \mathcal{H}_z$ and $(T \cap \mathcal{H}_z) \setminus \{e_i^z\}$ (such a perfect matching exists, because the induced subgraph consists of at most two disjoint paths of odd length). It is easy to see that $|M| = |S|$ and both $M \cap E_1$ and $M \cap E_2$ are matchings, and hence M is feasible.

To complete the proof, observe that the maximum degree in G is at most four if one starts with an instance of the 3-dimensional matching for which no element of $X \cup Y \cup Z$ occurs in more than three triples. Hence, the problem indeed remains NP-complete even if the maximum degree is 4.

□

In what follows, the previous problem is reduced to the d -distance matching problem, hence the hardness of the latter.

Theorem 1 *It is NP-complete to decide if a graph has a perfect d -distance matching, even if the maximum degree of the graph is at most 4.*

Proof It suffices to reduce the problem from Lemma 1 to the perfect d -distance matching problem. Let $G = (S, T; E)$; $S_1, S_2 \subseteq S$, $S_1 \cup S_2 = S$ be an instance of the above problem. Without loss of generality, one might assume that $S_1 \not\subseteq S_2$ and $S_2 \not\subseteq S_1$.

To construct an instance $G' = (S', T'; E')$, $d \in \mathbb{N}$ of the perfect d -distance matching problem, let $G' = G$ and modify G' as follows. Order the nodes of S' such that the nodes of $S_1 \setminus S_2$, $S_1 \cap S_2$ and $S_2 \setminus S_1$ appear in this order (the order of the elements inside the three sets is arbitrary). Insert $|S_1 \setminus S_2|$ and $|S_2 \setminus S_1|$ new nodes to S' right after the last node of S_1 and right before the first node covered by S_2 , respectively. Finally, add $|S_1 \setminus S_2| + |S_2 \setminus S_1|$ new nodes to T' , extend E' with the edges of a perfect matching between the newly added nodes and let $d = |S|$. Figure 3 illustrates the construction. The blank nodes on the figure are the newly inserted ones, and S'_i is the union of S_i and the i^{th} set of new nodes added to S' . The highlighted edges correspond to those in Figure 2b.

To complete the proof, observe that there exists a perfect $|S|$ -distance matching in G' if and only if there exists $M \subseteq E$ for which $|M| = |S|$ and both $M \cap E_1$ and $M \cap E_2$ are matchings. Indeed, from M one obtains a perfect $|S|$ -distance matching in G' by simply adding the perfect matching between the new nodes. To see the other direction, one has to

Fig. 3 Illustration of the construction in the proof of Theorem 1 for the problem instance presented in Fig. 2b. There exists a perfect 9-distance matching if and only if the problem given in Fig. 2b has a feasible solution of size 9

remove this perfect matching from the perfect S -distance matching. Note that the maximum degree in G' is not larger than in G , hence the problem remains hard even if the maximum degree is at most 4. $\square$

3 Weighted d -distance matching problem

This section presents various approaches to the weighted d -distance matching problem. Section 3.1 presents an FPT algorithm Downey and Fellows (2013) parameterized by d , while Sect. 3.2 settles the case when the size of T is constant. A simple greedy approach is presented in Sect. 3.3. Finally, Sects. 3.4.1 and 3.4.2 are devoted to the investigation of the natural linear programming model.

3.1 FPT algorithm parameterized by d

In what follows, an FPT algorithm parameterized by d is presented for the weighted (perfect) d -distance matching problem. Observe that the weighted d -distance matching problem easily reduces to the perfect case by adding a new node t_s to T and a new edge st_s of weight zero for each $s \in S$, therefore the algorithm is given only for the weighted perfect d -distance matching problem. The next claim gives a way to reduce the problem so that it admits an efficient dynamic programming solution.

Claim 1 Suppose that $s \in S$ is such that $\deg(s) \geq 2d$. Then, we can remove an arbitrary minimum-weight edge of $\Delta(s)$ from the edge set without changing the weight of the optimal perfect d -distance matching.

Proof Let st be a minimum-weight edge incident to node s . In order to prove that st can be removed, it suffices to show that there is a maximum-weight d -distance matching that does not use edge st . Given a d -distance matching M that contains edge st , let $Z \subseteq T$ denote the nodes that M assigns to $L_d(s) \cup R_d(s)$.

Since $|Z| \leq 2d - 1$, there exists a node $t' \in N(s) \setminus Z$ for which $w_{st} \leq w_{st'}$. Observe that $M' = (M \cup \{st'\}) \setminus \{st\}$ is a perfect d -distance matching of weight at least $w(M)$, which

does not contain edge st . Indeed, the degree of s remains one, and the only edge M' contains between nodes $L_d(s) \cup R_d(s)$ and t' is st' itself (by contradiction, if there were another edge $s't' \in M' \setminus \{st'\}$ for some $s' \in L_d(s) \cup R_d(s)$, then $s't'$ would be in M , and hence $t' \in Z$ would hold). $\square$

Based on this claim, the problem can be reduced so that the degree of each node $s \in S$ is at most $2d - 1$. The reduction can be performed in $\mathcal{O}(m + n)$ steps by removing all but the $2d - 1$ heaviest edges incident to each node $s \in S$. To this end, let us assume that each edge weight occurs only once (otherwise fix an arbitrary order between the ties), and for each $s \in S$, find the $(2d - 1)$ th lightest edge e_s in $\Delta(s)$ with the linear time selection algorithm, and then eliminate all edges of $\Delta(s)$ which are lighter than e_s .

In what follows, a dynamic programming approach is presented to solve the reduced problem in $\mathcal{O}((2d - 1)^{d+1}n)$ steps.

For $i \geq d$, let $f(s_i, z_1, \dots, z_d)$ denote the weight of the maximum-weight perfect d -distance matching if the problem is restricted to the first i nodes of S and s_{i-j+1} is assigned to its neighbor z_j for $j = 1, \dots, d$. Formally, let

$$f(s_i, z_1, \dots, z_d) = -\infty$$

if $z_1, \dots, z_d$ are not distinct, otherwise, $f(s_i, z_1, \dots, z_d)$ can be defined by the following recursive formula.

$$f(s_i, z_1, \dots, z_d) = \begin{cases} \sum_{j=1}^d w_{s_{i-j+1}z_j} & \text{if } i = d \\ w_{s_i z_1} + \max_{t \in N(s_{i-d})} f(s_{i-1}, z_2, \dots, z_d, t) & \text{if } i > d, \end{cases} \quad (1)$$

where $i \geq d$, $s_i \in S$, $z_j \in N(s_{i-j+1})$ for $j = 1, \dots, d$ and $z_1, \dots, z_d$ are distinct. To see that recursion (1) holds, observe that in its first case, by definition, $f(s_d, z_1, \dots, z_d)$ is the weight of matching $\{s_j z_{d-j+1} : j = 1, \dots, d\}$. In the second case of (1), s_i must be mapped to z_1 , and we want to find the maximum-weight perfect d -distance matching on the first $i - 1$ nodes of S which maps s_{i-j+1} to its neighbor z_j for $j = 2, \dots, d$. To this end, we want to find a node $t \in N(s_{i-d})$ (to be assigned to node s_{i-d}) which maximizes $f(s_{i-1}, z_2, \dots, z_d, t)$.

By definition, the weight of the optimal d -distance matching is

$$\max\{f(s_n, z_1, \dots, z_d) : z_j \in N(s_{n-j+1}) \text{ for } j = 1, \dots, d\}. \quad (2)$$

Observe that the number of subproblems is $\mathcal{O}(n(2d - 1)^d)$, since the degree of each $s \in S$ is at most $2d - 1$. Recursion (1) gives a way to compute $f(s_i, z_1, \dots, z_d)$ in $\mathcal{O}(d)$ steps if the subproblems are computed in appropriate order, i.e. the value $f(s_{i-1}, z'_1, \dots, z'_d)$ is available for all necessary $z'_1, \dots, z'_d \in T$. Therefore the number of steps to compute all the subproblems is $\mathcal{O}(dn(2d - 1)^d)$. Furthermore, the optimum value can be computed in $\mathcal{O}((2d - 1)^d)$ steps by (2). The overall running time of the algorithm is $\mathcal{O}(dn(2d - 1)^d + \text{poly}(|S| + |T|))$.

Remark 1 To improve the running time to $\mathcal{O}(nd^{d+1} + \text{poly}(|S| + |T|))$, observe that when $\max_{t \in N(s_{i-d})} f(s_{i-1}, z_2, \dots, z_d, t)$ is computed in (1), we need to consider only the (at most) d heaviest edges of $\Delta(s_{i-d})$ which are not incident to any of $z_2, \dots, z_d \in T$, since we only need to make sure that there is no conflict with the $d - 1$ nodes on the left of s_{i-d} . This way the number of subproblems is $\mathcal{O}(nd^d)$, and the overall number of steps is $\mathcal{O}(nd^{d+1} + \text{poly}(|S| + |T|))$. Similarly, in (2) one needs to consider only the d heaviest edges incident to z_j which are not incident to any of $z_{j+1}, \dots, z_d$, therefore there are at most d^d different configurations to be taken into account in (2).

3.2 Efficient algorithm for constant $|T|$

If the size of T is constant, then one can solve the problem efficiently as well. First, consider the following subproblems. Let $f(s_i, d_1, \dots, d_{|T|})$ denote the weight of the optimal perfect d -distance matching when the problem is restricted to $s_1, \dots, s_i$, and t_j cannot be matched to nodes $s_{i-d_j+1}, \dots, s_i$ for $j = 1, \dots, |T|$ (here $d_j = 0$ means that t_j can be matched to any node). Formally, $f(s_i, d_1, \dots, d_{|T|})$ can be defined as follows. If $i \geq 2$, then let

$$\begin{aligned} f(s_i, d_1, \dots, d_{|T|}) \\ = \max_{t_j \in N(s_i): d_j=0} \{w_{s_i t_j} + f(s_{i-1}, d'_1, \dots, d'_{j-1}, d-1, d'_{j+1}, \dots, d'_{|T|})\}, \end{aligned} \quad (3)$$

where $d'_k = \max(d_k - 1, 0)$ for $k = 1, \dots, |T|$. If $i = 1$, then let

$$f(s_1, d_1, \dots, d_{|T|}) = \max_{t_j \in N(s_1): d_j=0} w_{s_1 t_j}. \quad (4)$$

By definition, the weight of the optimal d -distance matching is given by

$$\max_{t_i \in N(s_n)} f(s_{n-1}, 0, \dots, 0, \underbrace{d-1}_{i^{\text{th}}}, 0, \dots, 0). \quad (5)$$

The number of subproblems to be solved is $\mathcal{O}(nd^{|T|})$, each of which can be computed in $\mathcal{O}(|T|)$ steps by (3) and (4). Once all the subproblems are computed, it takes additional $\mathcal{O}(|T|)$ steps to compute the optimal value by (5). Hence the overall number of steps is $\mathcal{O}(n|T|d^{|T|})$.

A similar approach settles the non-perfect case for constant $|T|$, the details of which are left to the reader.

3.3 A greedy algorithm

This section describes a greedy method for the weighted (not-necessarily-perfect) d -distance matching problem, and proves that it is a 3-approximation algorithm.

Algorithm 1 GREEDY

Let $e_1, \dots, e_m$ be the edges in non-increasing order by their weights.

$M := \emptyset$

for $i = 1, \dots, m$ **do**

if $M \cup \{e_i\}$ is a feasible d -distance matching **then**

$M := M \cup \{e_i\}$

end if

end for

output M

Theorem 2 GREEDY is a 3-approximation algorithm for the weighted d -distance matching problem.

Proof Assume that GREEDY returns edges $f_1, \dots, f_p$, and it selects them in this order. Let M_i denote a maximum-weight d -distance matching that contains $f_1, \dots, f_i$, where $0 \leq i \leq p$, i.e.

$$M_i = \arg \max \{w(M) : f_1, \dots, f_i \in M \text{ and } M \text{ is a } d\text{-distance matching}\}.$$

Fig. 4 Tight examples for Theorems 2, 8 and 9. **a** For $d = 2$ and unit weights, GREEDY might select edge s_2t_2 only, while the largest 2-distance matching is of cardinality 3. **b** For $d = 2$ and unit weights, both S -GREEDY and T -GREEDY select edge s_1t_1 only, while the largest 2-distance matching is of cardinality 2

Furthermore, let θ_i denote the weight of M_i for $i = 0, \dots, p$. Note that θ_0 is the weight of the optimal d -distance matching and θ_p is the weight of the matching GREEDY returns. Observe that there exist edges $e, e', e'' \in M_i \setminus \{f_1, \dots, f_i\}$ such that $(M_i \setminus \{e, e', e''\}) \cup \{f_{i+1}\}$ is a feasible d -distance matching, which contains edges $f_1, \dots, f_{i+1}$. By the greedy selection rule, $w_e, w_{e'}, w_{e''} \leq w_{f_{i+1}}$, one gets

$$\theta_{i+1} \geq \theta_i + w_{f_{i+1}} - w_e - w_{e'} - w_{e''} \geq \theta_i - 2w_{f_{i+1}} \quad (6)$$

holds for all $i = 0, \dots, p-1$. A simple inductive argument shows that (6) implies $\theta_p \geq \theta_0 - 2 \sum_{i=1}^p w_{f_i}$, therefore $3\theta_p \geq \theta_0$ follows, which completes the proof.

The analysis is tight even for $d = 2$ and $w \equiv 1$ in the sense that GREEDY might return only one edge, while the largest 2-distance matching consists of 3 edges, see Fig. 4a for an example. $\square$

Remark 2 The above proof shows that GREEDY is a 3-approximation algorithm for the more general *cyclic d -distance matching problem*, in which the nodes of S are considered in cyclic order.

3.4 Linear programming

The following two sections prove that the integrality gap of the natural integer programming model is at most $2 - \frac{1}{2d-1}$, and present an LP-based $(2 - \frac{1}{2d-1})$ -approximation algorithm for the weighted d -distance matching problem. First consider the relaxation of the natural 0–1 integer programming formulation of the weighted d -distance matching problem.

$$\max \sum_{st \in E} w_{st} x_{st} \quad (\text{LP1})$$

s.t.

$$x \in \mathbb{R}_+^E \quad (7a)$$

$$\sum_{st \in \Delta(s)} x_{st} \leq 1 \quad \forall s \in S \quad (7b)$$

$$\sum_{s't \in E: s' \in R_d(s)} x_{s't} \leq 1 \quad \forall s \in S, t \in T \quad (7c)$$

One gets the relaxation of the 0–1 integer programming formulation (LP2) of the weighted perfect d -distance matching problem by tightening (7b) to equality in LP1.

3.4.1 Integrality gap

This section proves that the integrality gap of **LP1** is at most $2 - \frac{1}{2d-1}$, and proves the integrality of **LP1** and **LP2** in special cases. The former result also follows from the LP-based approximation algorithm described in Sect. 3.4.2. The following definition plays a central role both in the analysis of the integrality gap and in the LP-based approximation algorithm presented in the next section.

Definition 1 Given a feasible solution x of **LP1**, an order of the edges $e_1 = s^1 t^1, \dots, e_m = s^m t^m$ is **θ -flat with respect to x** if

$$\xi_i + \bar{\xi}_i \leq \theta - x_{e_i} \quad (8)$$

holds for each $i = 1, \dots, m$, where $\xi_i = \sum \{x_{e_j} : j > i, e_j \in \Delta(s^i)\}$ and $\bar{\xi}_i = \sum \{x_{e_j} : j > i, e_j \in \Delta(t^i), s^j \in L_d(s^i) \cup R_d(s^i)\}$.

That is, an order of the edges is θ -flat if the sum of x on those edges among $e_{i+1}, \dots, e_m$ that are hit by an edge e_i is at most $\theta - x_{e_i}$ for every i . Note that any order of the edges is 3-flat by definition, since for any edge $e = st$, the sum of variables on all edges incident to s is at most 1 by (7b), whereas the sum on the edges induced by $L_d(s) \cup R_d(s)$ and $\{t\}$ is at most 2 by (7c). The following lemma further improves this bound to $2 - \frac{1}{2d-1}$.

Lemma 2 *There exists an optimal solution $x \in \mathbb{Q}^E$ of **LP1** and an order $e_1 = s^1 t^1, \dots, e_m = s^m t^m$ of the edges that is $(2 - \frac{1}{2d-1})$ -flat with respect to x .*

Proof Let $E_s \subseteq \Delta(s)$ denote the first $\min(2d-1, \deg(s))$ largest weight edges incident to node s for each $s \in S$. Let x be an optimal solution to **LP1** for which $\gamma(x) = \sum \{x_e : e \in E \setminus \bigcup_{s \in S} E_s\}$ is minimal. Towards a contradiction, suppose that $\gamma(x) > 0$. By definition, $\gamma(x) > 0$ implies that there exists an edge $st \in E \setminus \bigcup_{k=1}^n E_k$ for which $x_{st} > 0$. There exists an edge $st' \in E_s$ such that $x' = x - \epsilon x_{st} + \epsilon x_{st'}$ is feasible for sufficiently small $\epsilon > 0$, otherwise $x(\bigcup \{\Delta(s') : s' \in L_d(s) \cup R_d(s)\}) \geq 2d-1 + \epsilon$ would hold, which is not possible because of the constraints (7b). But then $w x \leq w x'$ and $\gamma(x') < \gamma(x)$, contradicting the minimality of $\gamma(x)$. Therefore $\gamma(x) = 0$ follows, meaning that $x_e = 0$ holds for each $e \in E \setminus \bigcup_{s \in S} E_s$. Hence one can restrict the edge set to $\bigcup_{s \in S} E_s$ without change in the optimal objective value, which implies that there exists a rational optimal solution $x \in \mathbb{Q}^E$ of **LP1** with $\gamma(x) = 0$.

Let x be as above, and let $e_1 = s^1 t^1, \dots, e_m = s^m t^m$ be the order of the edges given by Algorithm 2 for input x . To prove that this order is $(2 - \frac{1}{2d-1})$ -flat with respect to x , let ξ_i and $\bar{\xi}_i$ ($i = 1, \dots, n$) be as in Definition 1. First observe that $\bar{\xi}_i \leq 1 - x_i$ holds for each $i = 1, \dots, n$, because the algorithm places each edge $\bigcup_{j=1}^{i-1} \Delta(s^j)$ before e_i . Hence, to obtain (8), it suffices to prove that $\xi_i \leq 1 - \frac{1}{2d-1}$. For any node $s \in S$, if there exists an edge $st \in \Delta(s)$ for which $x_{st} \geq \frac{1}{2d-1}$, then $\xi_i \leq 1 - \frac{1}{2d-1}$ follows for each $e_i \in \Delta(s)$, since $x_e \geq \frac{1}{2d-1}$ holds for the first edge $e \in \Delta(s)$ selected by Algorithm 2. Otherwise, there exists no edge $st \in \Delta(s)$ for which $x_{st} \geq \frac{1}{2d-1}$. Therefore $\xi_i \leq x(\Delta(s)) < |E_s| \frac{1}{2d-1} \leq 1$ follows for $e_i \in \Delta(s)$, which completes the proof if $|E_i| < 2d-1$. Hence one can assume that $|E_i| = 2d-1$. Next we argue that $x' = x + \epsilon x_{st'}$ is feasible for some $st' \in E_s$ and sufficiently small $\epsilon > 0$. By contradiction, if there existed no such edge st' , then it is one of the constraints (7c) that prevents us from increasing $x_{st'}$ for each $st' \in \Delta(s)$. However, these tight constraints imply that $x(\bigcup \{\Delta(s') : s' \in L_d(s) \cup R_d(s)\}) = 2d-1$, but this can not

be the case, because $x(\Delta(s)) < 1$. Hence x' is a feasible solution for some $st' \in \Delta(s)$ and sufficiently small $\epsilon > 0$ — contradicting the optimality of x .

Therefore $\xi_i \leq 1 - \frac{1}{2d-1}$ follows for $i = 1, \dots, n$, which means that the order of the edges is $(2 - \frac{1}{2d-1})$ -flat. $\square$

Theorem 3 *The integrality gap of LP1 is at most $2 - \frac{1}{2d-1}$.*

Proof Let $\theta = 2 - \frac{1}{2d-1}$. By Lemma 2, there exists a solution $x \in \mathbb{Q}^E$ to LP1 and an order of the edges $e_1 = s^1 t^1, \dots, e_m = s^m t^m$ that is θ -flat with respect to x . First, it will be shown that there exist d -distance matchings $M_1, \dots, M_q$ and coefficients $\lambda_1, \dots, \lambda_q \in \mathbb{Q}_+$ such that $\sum_{i=1}^q \lambda_i \chi_{M_i} = x$ and $\lambda := \sum_{i=1}^q \lambda_i \leq \theta$.

Let $K \in \mathbb{N}$ be the lowest common denominator of $\{x_e : e \in E\}$, and let $q = \lfloor K\theta \rfloor$. The main observation is that each edge $e \in E$ can be assigned a set of colors $C_e \subseteq \{1, \dots, q\}$ such that each color class corresponds to a feasible d -distance matching and $|C_e| = Kx_e$. To prove this, the edges are greedily colored one by one in order $e_m, \dots, e_1$. By induction, assume that edges $e_m, \dots, e_{i+1}$ already have their color sets. It suffices to assign a color set C_{e_i} to edge e_i which is of size Kx_{e_i} and distinct from both $A := \bigcup \{C_{e_j} : j > i, e_j \in \Delta(s^i)\}$ and $B := \bigcup \{C_{e_j} : j > i, e_j \in \Delta(t^i), s^j \in R_d(s^i) \cup L_d(s^i)\}$. Without loss of generality, assume that $x_{e_i} > 0$ (otherwise $C_{e_i} = \emptyset$). By (8), one gets $|A \cup B| \leq |A| + |B| = K(\xi_i + \bar{\xi}_i) \leq \lfloor K(\theta - x_{e_i}) \rfloor = \lfloor K\theta \rfloor - Kx_{e_i} = q - Kx_{e_i}$, thus $|A \cup B| + Kx_{e_i} \leq q$. That is, the number of free colors is at least Kx_{e_i} , so let C_{e_i} be any Kx_{e_i} colors in $\{1, \dots, q\} \setminus (A \cup B)$.

Let the desired d -distance matching M_i consist of the edges with color i for $i = 1, \dots, q$. Set $\lambda_i = \frac{1}{K}$ for all $i = 1, \dots, q$, and observe that both $\sum_{i=1}^q \lambda_i \chi_{M_i} = x$ and $\sum_{i=1}^q \lambda_i = \sum_{i=1}^q \frac{1}{K} = \frac{q}{K} \leq \theta$ hold.

Now, we are ready to argue that there exists a λ -approximate solution among $M_1, \dots, M_q$. By contradiction, suppose that $\lambda w(M_i) < w(M^*)$ for each $i = 1, \dots, q$, where M^* is an optimal distance matching. Observe that $w \sum_{i=1}^q \lambda_i \chi_{M_i} = \sum_{i=1}^q \lambda_i w(M_i) < \frac{1}{\lambda} w(M^*) \sum_{i=1}^q \lambda_i = w(M^*)$, that is, the LP optimum is strictly smaller than the IP optimum, which is a contradiction. Therefore, the largest weight d -distance matchings among $M_1, \dots, M_q$ are indeed λ -approximate. Since $\lambda = \sum_{i=1}^q \lambda_i \leq 2 - \frac{1}{2d-1}$, the proof is complete. $\square$

Note that the above approach is algorithmic, but it does not necessarily run in polynomial time — since q may be exponential in the size of the graph. The next section presents a polynomial-time method and re-proves that the integrality gap is at most $2 - \frac{1}{2d-1}$.

Algorithm 2 The ordering procedure for Lemma 2

Let x be a given fractional solution to LP1 and let $G = (S, T; E)$ be a copy of the graph.

$j := 1$

for $i = 1, \dots, n$ **do**

while $\deg(s_i) \neq 0$ **do**

 Choose an edge $s_i t \in \Delta(s_i)$ for which $x_{s_i t}$ is as large as possible.

$e_j := s_i t$

$j := j + 1$

$E := E \setminus \{s_i t\}$

end while

end for

output $e_1, \dots, e_m$

Fig. 5 For $w \equiv 1$ and $d = 5$, $x \equiv 1/2$ is an optimal solution to **LP1**, and the highlighted edges form an optimal 5-distance matching, hence the integrality gap is $6/5$

Remark 3 Figure 5 provides an example with (the largest known) integrality gap $6/5$. Using this instance, one might easily derive an example (by adding two new nodes t_5 and t_6 to T , and two new edges s_3t_5, s_6t_6) for which no perfect 5-distance matching exists, but there is a fractional perfect 5-distance matching — meaning that the integrality gap of LP2 is unbounded as it was expected due to the complexity of the problem.

In what follows, the integrality of **LP1** and LP2 are shown in special cases.

Theorem 4 *If $d = 1$ or $d = 2$, then both **LP1** and LP2 are integral.*

Proof For $d = 1$, the matrix of **LP1** and LP2 is the incidence matrix of a bipartite graph, which is a well-known network matrix Frank (2011, Page 149). For $d = 2$, one can easily construct a directed graph and a spanning tree (in this case a directed caterpillar) for which the corresponding network matrix is the matrix of **LP1** and LP2. As the right-hand side of both programs are integral and their matrices are network matrices (and hence totally unimodular) for $d = 1, 2$, the proof is complete. Note that the statement for $d = 1$ also follows from Theorem 3. $\square$

Note that the matrix of **LP1** and LP2 is not totally unimodular for $d \geq 3$ if the input graph is the complete bipartite graph — the technical proof is omitted here. Therefore, the proof of Theorem 4 can not work for $d \geq 3$ and one can not expect that **LP1** and LP2 remain integral. Having said that, LP2 still describes the convex hull of the integral solutions for $d = |T|$, but not because of total unimodularity:

Theorem 5 *If $d = |T|$, then LP2 is integral.*

Proof Let A denote the matrix of LP2, and let $\tilde{x}$ be an optimal integral solution. If $\tilde{x}$ is not an optimal LP solution, then there is no complementary dual solution y , therefore — by Farkas' lemma — there exists $z \in \mathbb{R}^E$ for which

$$wz > 0 \quad (9a)$$

$$Az = 0 \quad (9b)$$

$$\tilde{x}_e = 1 \implies z_e \leq 0 \quad \forall e \in E \quad (9c)$$

$$\tilde{x}_e = 0 \implies z_e \geq 0 \quad \forall e \in E. \quad (9d)$$

Let $z^j = (z_{s_1t_j}, z_{s_2t_j}, \dots, z_{s_{|S|}t_j})$. Observe that $z_i^j = z_k^j$ for all $j = 1, \dots, d$ whenever $i \equiv k \pmod d$, which allows the simplification of (9a–9d). For all $i = 1, \dots, d$ and $j = 1, \dots, |T|$, let $\hat{z}_i^j$ be a new variable representing all variables $\{z_{i'}^j : i' \equiv i \pmod d\}$, and consider the following formulation.

$$\max \hat{w}\hat{z}$$

$$\sum_{i=1}^d \hat{z}_i^j = 0 \forall j = 1, \dots, |T| \quad (10a)$$

$$\sum_{j=1}^{|T|} \hat{z}_i^j = 0 \forall i = 1, \dots, d \quad (10b)$$

$$\hat{z}_i^j \leq 0 \forall s_i t_j \in E : i \in \{1, \dots, d\} \text{ and } \tilde{x}_{s_i t_j} = 1 \quad (10c)$$

$$\hat{z}_i^j \geq 0 \forall s_i t_j \in E : i \in \{1, \dots, d\} \text{ and } \tilde{x}_{s_i t_j} = 0 \quad (10d)$$

$$-1 \leq z \leq 1 \quad (10e)$$

where $\hat{w}_i^j = \sum \{w_{i'}^j : i' \in \{1, \dots, |S|\} \text{ and } i' \equiv i \pmod{d}\}$. Note that system (9a–9d) has a feasible solution if and only if (10a–10e) has one with positive objective value. As the optimal value of (10a–10e) is finite and its matrix is totally unimodular (the incidence matrix of a bipartite graph and identity matrices under it), there is an integral solution $\hat{z}^*$ to (10a–10e) with a positive objective value. This particular solution corresponds to an integral solution z^* to (9a–9d) with the same positive weight. But this means that $\tilde{x} + z^*$ is an integral solution of LP2, for which $w\tilde{x} < w(\tilde{x} + z^*)$ holds, contradicting the fact that $\tilde{x}$ was an optimal integral solution. $\square$

Note that the analogous statement for LP1 does not hold.

3.4.2 $(2 - \frac{1}{2d-1})$ -approximation algorithm for the weighted d -distance matching

This section presents an “almost greedy” LP-based $(2 - \frac{1}{2d-1})$ -approximation algorithm and re-proves that the integrality gap is at most $\theta := 2 - \frac{1}{2d-1}$.

Algorithm 3 θ -approximation algorithm for the weighted distance matching problem

Let $e_1, \dots, e_m$ be a θ -flat order with respect to a solution x of LP1 (see Lemma 2).

procedure WDMLPAPX(E, w)

$E := E \setminus \{e \in E : w_e \leq 0\}$

if $E = \emptyset$ **then**

return $\emptyset$

end if

Let st be the first edge according to the above order that appears in E .

$M' := \text{WDMLPAPX}(E \setminus \{st\}, w')$, where $w' := w - w_{st} \chi_{\Delta(s) \cup \{s' : t \in \Delta(t) : s' \in R_d(s)\}}$

if $M' \cup \{st\}$ is a feasible d -distance matching **then**

return $M' \cup \{st\}$

else

return M'

end if

Theorem 6 Algorithm 3 is a θ -approximation algorithm for the weighted d -distance matching problem if a θ -flat order of the edges is given in the first step of the algorithm.

Proof The proof is by induction on the number of edges. Let M denote the distance matching found by WDMLPAPX(E, w), and let x be as defined in Algorithm 3. In the base case, if $E = \emptyset$, then $\theta w(M) \geq wx$ holds. Let $st \in E$ be the first edge with respect to the order of the edges

used by Algorithm 3. By induction, $\theta w'(M') \geq w'x$ holds for $M' = \text{WDMLPAPX}(E \setminus \{st\}, w')$, where $w' = w - w_{st} \chi_{\Delta(s) \cup \{s't \in \Delta(t): s' \in R_d(s)\}}$. The key observation is that

$$\theta(w - w')(M) \geq \theta w_{st} \geq (w - w')x \quad (11)$$

follows from the definition of w' and the order of the edges. Hence, one gets

$$\theta w(M) = \theta(w - w')(M) + \theta w'(M) \geq (w - w')x + w'x = wx, \quad (12)$$

where $w'(M) = w'(M')$ because $w'_{st} = 0$. Therefore, M is indeed a θ -approximate solution, which completes the proof. $\square$

Theorem 6 also implies that the integrality gap of LPI is at most θ . Note that if we have a θ' -flat order of the edges in the first step of Algorithm 3, then it outputs a θ' -approximate solution. We believe that there always exists a θ' -flat order of the edges for some $\theta' < \theta$, i.e. it is possible to improve Lemma 2, which would automatically improve both the integrality gap and the approximation guarantee of the algorithm to θ' .

3.5 A combinatorial $(2 - \frac{1}{d})$ -approximation algorithm

This section presents a $(2 - \frac{1}{d})$ -approximation algorithm for the weighted distance matching problem. Let $k \in \{d - 1, \dots, 3d - 3\}$ be such that $2d - 1$ divides $|S| + k$, and add k new dummy nodes $s_{n+1}, \dots, s_{n+k}$ to the end of S in this order. Let us consider the extended node set in cyclic order. Observe that the new cyclic problem is equivalent to the original one. Let H_j denote the subgraph of G induced by $R_d(s_j) \cup T$, where $R_d(s_j)$ is the set consisting of node s_j and the next $d - 1$ nodes on its right in the new cyclic problem. For each such subgraph H_j , let F_j denote a maximum-weight matching of it with respect to w . Let

$$G_i = (S_i, T; E_i) = \bigcup_{j=0}^{\frac{n+k}{2d-1}-1} H_{i+j(2d-1)}$$

and

$$M_i = \bigcup_{j=0}^{\frac{n+k}{2d-1}-1} F_{i+j(2d-1)}$$

for $i = 1, \dots, 2d - 1$, where $S_i \subseteq S$. Let $i^* = \arg \max\{w(M_i) : i = 1, \dots, 2d - 1\}$. For example, consider the graph in Fig. 6 with $d = 3$. The nodes of G_4 are highlighted on the figure and the edges of M_4 are the wavy ones. Nodes $s_6, \dots, s_{10}$ are the five dummy nodes.

Since M_{i^*} can be computed in strongly polynomial time, we obtain a strongly-polynomial-time $(2 - \frac{1}{d})$ -approximation algorithm by the following theorem.

Theorem 7 M_{i^*} is a $(2 - \frac{1}{d})$ -approximate d -distance matching.

Proof Each node of S is covered by at most one edge of M_i , as M_i is the union of matchings no two of which cover the same node of S . If $s_it, s_jt \in M_i$, then s_it and s_jt belong to two distinct matchings $F_k, F_l \subseteq M_i$ for some k, l , hence $|j - i| \geq d$. From this, the feasibility of M_i follows for all $i = 1, \dots, 2d - 1$, and M_{i^*} being one of them, it is feasible as well.

To show the approximation guarantee, let M^* be an optimal d -distance matching. For each node $s \in S$, let $\mu_s \in \mathbb{R}_+$ denote the weight of the edge covering s in M^* and zero if

Fig. 6 Tight example for Theorem 7 in the case $d = 3$. The wavy edges form a possible output of the algorithm. (Recall that the nodes of S are in cyclic order.)

M^* does not cover s . Note that $\sum_{s \in S} \mu_s = w(M^*)$ by definition, and

$$\sum_{s \in S_i} \mu_s \leq w(M_i) \quad (13)$$

follows because $\sum_{s \in S_i} \mu_s$ is the weight of a d -distance matching which covers no nodes outside G_i . Observe that

$$dw(M^*) = d \sum_{s \in S} \mu_s = \sum_{i=1}^{2d-1} \sum_{s \in S_i} \mu_s \leq \sum_{i=1}^{2d-1} w(M_i) \leq (2d-1)w(M_{i^*}) \quad (14)$$

holds, where the second equation holds because μ_s occurs exactly d times as a summand in $\sum_{i=1}^{2d-1} \sum_{s \in S_i} \mu_s$ for all $s \in S$, the first inequality follows from (13), while the last one holds because M_{i^*} is a largest-weight d -distance matching among $M_1, \dots, M_{2d-1}$. By (14), one gets $w(M^*) \leq (2 - \frac{1}{d})w(M_{i^*})$, which completes the proof of the theorem. $\square$

The analysis is tight in the sense that, for every $d \in \mathbb{Z}_+$, there exists a graph G for which the algorithm returns a d -distance matching M for which $w(M^*) = (2 - \frac{1}{d})w(M)$, where M^* is an optimal d -distance matching. Let S and T consist of $2d-1$ and d nodes, respectively. Add edge $s_i t_i$ for $i = 1, \dots, d$, and edge $s_{i+d} t_i$ for $i = 1, \dots, d-1$. Note that the edge set is a feasible d -distance matching itself, and the above algorithm returns a matching that covers exactly d nodes of S . Hence the approximation ratio of the found solution is $\frac{2d-1}{d}$. Figure 6 shows the construction for $d = 3$, where $s_6, \dots, s_{10}$ are the dummy nodes.

4 Unweighted d -distance matching

First, two refined greedy approaches are considered for the unweighted case, then the analysis of the approximation ratio of locally optimal solutions follows.

4.1 Greedy algorithms

This section describes two refined greedy algorithms for the unweighted d -distance matching problem, and proves that both of them achieve an approximation guarantee of 2.

Theorem 8 *S-GREEDY is a 2-approximation algorithm for the unweighted d -distance matching problem.*

Proof Assume that S -GREEDY returns edges $f_1, \dots, f_p$, and it selects them in this order. Let M_i and θ_i be as above in the proof of Theorem 2, i.e. let $M_i = \arg \max\{w(M) : f_1, \dots, f_i \in M \text{ and } M \text{ is a } d\text{-distance matching}\}$ and let θ_i denote the weight of M_i for

Algorithm 4 *S*-GREEDY

```

Let  $s_1, \dots, s_n$  be the nodes of  $S$  in the given order
 $M := \emptyset$ 
for  $i = 1, \dots, n$  do
  if  $M \cup \{s_i t\}$  is feasible for some  $s_i t \in \Delta(s_i)$  then
     $j := \arg \min\{j : s_i t_j \in \Delta(s_i) \text{ and } M \cup \{s_i t_j\} \text{ is feasible}\}$ 
     $M := M \cup \{s_i t_j\}$ 
  end if
end for
output  $M$ 

```

$i = 0, \dots, p$. Observe that, as opposed to the proof of Theorem 2, there exist two edges $e, e' \in M_i \setminus \{f_1, \dots, f_i\}$ such that $(M_i \setminus \{e, e'\}) \cup \{f_{i+1}\}$ is a feasible d -distance matching containing edges $f_1, \dots, f_{i+1}$. Indeed, if there were three edges to leave out, then one of them would be incident to $\{s_{i-d+1}, \dots, s_{i-1}\}$, but then Algorithm 4 would have picked this edge instead of f_{i+1} . By the greedy selection rule, one gets

$$\theta_{i+1} \geq \theta_i + 1 - 1 - 1 = \theta_i - 1 \quad (15)$$

holds for all $i = 0, \dots, p-1$. A straightforward inductive argument shows that (15) implies $\theta_p \geq \theta_0 - p$, therefore $2\theta_p \geq \theta_0$ follows, which completes the proof.

The analysis is tight in the sense that *S*-GREEDY might return only one edge, while the largest 2-distance matching consists of 2 edges, see Fig. 4b. $\square$

Algorithm 5 *T*-GREEDY

```

Let  $s_1, \dots, s_n$  be the nodes of  $S$  in the given order
 $M := \emptyset$ 
 $C := \emptyset$ 
for  $j = 1, \dots, k$  do
   $M_j := \emptyset$ 
   $C_j := \emptyset$ 
   $i := 1$ 
  while  $i \leq n$  do
    if  $s_i t_j \in E$  and  $s_i \notin C$  then
       $M_j := M_j \cup \{s_i t_j\}$ 
       $C_j := C_j \cup \{s_i\}$ 
       $i := i + d$ 
    else
       $i := i + 1$ 
    end if
  end while
   $M := M \cup M_j$ 
   $C := C \cup C_j$ 
end for
output  $M$ 

```

Theorem 9 *T*-GREEDY is a 2-approximation algorithm for the unweighted d -distance matching problem.

Proof Let M_S and M_T denote the edge sets *T*-GREEDY (Algorithm 5) and *S*-GREEDY (Algorithm 4) outputs, respectively. It suffices to prove that $M_S = M_T$. By contradiction,

suppose that $M_S \neq M_T$. Let s_i be the first node in S for which $\Delta(s_i) \cap M_S \neq \Delta(s_i) \cap M_T$, and choose the edge $s_i t_j \in \Delta(s_i) \cap (M_S \Delta M_T)$ such that j is the smallest possible.

Case 1: $s_i t_j \in M_S \setminus M_T$. First, observe that M_T covers node s_i , otherwise it would have included $s_i t_i$. Therefore, T -GREEDY assigns node s_i to $t_{j'}$, where $j' \neq j$. If $j' < j$, then S -GREEDY would have chosen edge $s_i t_{j'}$ instead of $s_i t_j$. If $j' > j$, then T -GREEDY would have included $s_i t_j$ instead of $s_i t_{j'}$ to M_T .

Case 2: $s_i t_j \in M_T \setminus M_S$. Observe that M_S covers node s_i , otherwise S -GREEDY could have included edge $s_i t_i$. Therefore, S -GREEDY assigns node s_i to $t_{j'}$, where $j' \neq j$. Similarly to the argument in Case 1, it is easy to see that neither $j' < j$ nor $j' > j$ is possible.

Figure 4b shows that the approximation ratio is tight. $\square$

4.2 Local search

This section investigates the approximation ratio of the so-called locally optimal solutions. First, consider the following notion, which plays a central role throughout the section.

Definition 2 Given an edge $e^* \in E$, let $\mathcal{H}(e^*, M) \subseteq M$ denote the inclusion-wise minimal subset of M for which $M \setminus \mathcal{H}(e^*, M) \cup \{e^*\}$ is a feasible d -distance matching.

We say that an edge e^* *hits* the edges of $\mathcal{H}(e^*, M)$, or that $\mathcal{H}(e^*, M)$ is the *hit set* of edge e^* with respect to M . Similar notation and terminology are used for a subset of the edges as follows.

Definition 3 Given an edge set $X \subseteq E$, let $\mathcal{H}(X, M) \subseteq M$ denote the set of edges hit by at least one edge in X , i.e. let $\mathcal{H}(X, M) = \bigcup_{e^* \in X} \mathcal{H}(e^*, M)$.

Definition 4 A d -distance matching M is l -**locally optimal** if there exists no d -distance matching $X \subseteq E \setminus M$ such that $l \geq |X| > |\mathcal{H}(X, M)|$. Similarly, M is l -**locally optimal with respect to** M^* if there exists no $X \subseteq M^* \setminus M$ such that $l \geq |X| > |\mathcal{H}(X, M)|$, where M^* is an d -distance matching.

For unit weights, the possible outputs of GREEDY (Algorithm 1) are exactly the 1-locally optimal solutions.

Claim 2 A d -distance matching M is 1-locally optimal if and only if there exists a permutation of E such that GREEDY outputs M for $w \equiv 1$.

Proof If M is the output of GREEDY, then there exists no edge e outside M which can be added to M (otherwise GREEDY could have added e when it tried to), hence M is 1-locally optimal by definition. On the other hand, if M is 1-locally optimal, then permute E such that the edges of M come first. As $w \equiv 1$, one can choose this particular permutation in the first line of Algorithm 1. To complete the proof, observe that its output is M itself, since it includes all edges of M as M is feasible, and may not include any other edges, because M is 1-locally optimal. $\square$

In what follows, an upper bound q_l is shown on the approximation ratio of l -locally optimal solutions for each $l \geq 1$, where q_l is defined by the following recursion.

$$q_l = \begin{cases} 3, & \text{if } l = 1 \\ 2, & \text{if } l = 2 \\ \frac{4q_{l-2} - 3}{2q_{l-2} - 1}, & \text{if } l \geq 3. \end{cases} \quad (16)$$

For $l = 1, 2, 3, 4$, the statement can be proved by a simple argument, given below. However, this approach does not seem to work in the general case. The proof of the general case, which is much more involved and quite esoteric, is given after the following theorem.

Theorem 10 *If M, M^* are d -distance matchings such that M is l -locally optimal with respect to M^* , then the approximation ratio $|M^*|/|M|$ is at most q_l , where $l = 1, \dots, 4$ and q_l is as defined above.*

Proof Let $M_i^* = \{e^* \in M^* : |\mathcal{H}(e^*, M)| = i\}$ for $i = 0, \dots, 3$. Note that $M_0^*, M_1^*, M_2^*, M_3^*$ is a partition of M^* , and $M_0^* = \emptyset$ since each edge of M^* hits at least one edge of M if $l \geq 1$. Since each edge $e \in M$ can be hit by at most three edges of M^* , one gets

$$3|M| \geq \sum_{e^* \in M^*} |\mathcal{H}_+(e^*, M)| = |M_1^*| + 2|M_2^*| + 3|M_3^*|. \quad (17)$$

Case $l = 1$.

It easily follows from (17) that

$$|M^*| = |M_1^*| + |M_2^*| + |M_3^*| \leq |M_1^*| + 2|M_2^*| + 3|M_3^*| \leq 3|M|. \quad (18)$$

Case $l = 2$. Similarly,

$$\begin{aligned} 2|M^*| &= 2(|M_1^*| + |M_2^*| + |M_3^*|) \leq |M_1^*| + |M_1^*| + 2|M_2^*| + 3|M_3^*| \\ &\leq |M_1^*| + 3|M| \leq 4|M|, \end{aligned} \quad (19)$$

where the second inequality follows from (17) and the third one holds because M is 2-locally optimal with respect to M^* .

Case $l = 3$. For $l = 3$, one has to show that $|M^*|/|M| \leq 9/5$. In the following computation, inequality (17) is forced with an appropriate coefficient so that the rest admits the application of case $l = 1$ to a derived problem instance.

$$\begin{aligned} 5|M^*| &= 5(|M_1^*| + |M_2^*| + |M_3^*|) = 2(|M_1^*| + 2|M_2^*| + 3|M_3^*|) \\ &\quad + 3|M_1^*| + |M_2^*| - |M_3^*| \leq 6|M| + 3|M_1^*| + |M_2^*| - |M_3^*| \leq 9|M|, \end{aligned} \quad (20)$$

where the first inequality holds by (17), while the last one by the following claim.

Claim 3 *If M is 3-locally optimal with respect to M^* , then*

$$|M_2^*| - |M_3^*| \leq 3(|M| - |M_1^*|). \quad (21)$$

Proof It suffices to show that there exist d -distance matchings $\tilde{M}, \tilde{M}^*$ such that

- (1) $|\tilde{M}| = |M| - |M_1^*|$,
- (2) $|\tilde{M}^*| = |M_2^*|$,
- (3) $\tilde{M}$ is 1-locally optimal with respect to $\tilde{M}^*$.

Then, condition 3) implies that $|\tilde{M}^*| \leq 3|\tilde{M}|$ holds, from which the inequality to be proved follows by substituting 1) and 2).

Let $\tilde{M} = M \setminus \mathcal{H}(M_1^*, M)$ and $\tilde{M}^* = M_2^*$. Clearly, both 1) and 2) hold. By contradiction, suppose that 3) does not hold, that is, there exists $e_1^* \in \tilde{M}^*$ such that $\tilde{M} \cup \{e_1^*\}$ is a feasible d -distance matching. By definition, $e_1^* \in M_2^*$, therefore e_1^* hits exactly two edges e_1, e_2 in M . Neither e_1 , nor e_2 are in $\tilde{M}$, thus $e_1, e_2 \in \mathcal{H}(M_1^*, M)$, that is e_j is hit by an edge $e_{j+1}^* \in M_1^*$ for $j = 1, 2$. Note that e_1^*, e_2^*, e_3^* are pairwise distinct edges, and $\mathcal{H}(\{e_1^*, e_2^*, e_3^*\}, \tilde{M}) = \{e_1, e_2\}$, contradicting that M is 3-locally optimal. $\square$

Case $l = 4$. One has to show that $|M^*|/|M| \leq 5/3$. As in the previous case, inequality (17) will be applied with an appropriate multiplier so that the rest admits the application of case $l = 2$ to a derived problem instance.

$$\begin{aligned} 6|M^*| &= 6(|M_1^*| + |M_2^*| + |M_3^*|) = 2(|M_1^*| + 2|M_2^*| + 3|M_3^*|) + 4|M_1^*| + 2|M_2^*| \\ &\leq 6|M| + 4|M_1^*| + 2|M_2^*| \leq 10|M|, \end{aligned} \quad (22)$$

where the first inequality holds by (17), the last one by the following claim.

Claim 4 *If M is 4-locally optimal with respect to M^* , then*

$$2|M_2^*| \leq 4(|M| - |M_1^*|). \quad (23)$$

Proof It suffices to show that there exist d -distance matchings $\tilde{M}, \tilde{M}^*$ such that

- (1) $|\tilde{M}| = |M| - |M_1^*|$,
- (2) $|\tilde{M}^*| = |M_2^*|$,
- (3) $\tilde{M}$ is 2-locally optimal with respect to $\tilde{M}^*$.

Then, condition 3) implies that $|\tilde{M}^*| \leq 2|\tilde{M}|$ holds, from which one obtains the inequality to be proved by substituting 1) and 2).

Let $\tilde{M} = M \setminus \mathcal{H}(M_1^*, M)$ and $\tilde{M}^* = M_2^*$. As in the proof of the claim in case $l = 3$, one might show that 3) holds, hence the desired inequality follows. $\square$

This concludes the proof of the theorem. $\square$

It is worth noting that the proof for $l = 3, 4$ refers inductively to the case $l-2$, which is quite unexpected. The same idea does not seem to work for $l = 5$. Based on cases $l = 1, 2, 3, 4$, one gains the following analogous computation.

$$\begin{aligned} 13|M^*| &= 13(|M_1^*| + |M_2^*| + |M_3^*|) = 4(|M_1^*| + 2|M_2^*| + 3|M_3^*|) + 9|M_1^*| \\ &\quad + 5|M_2^*| + |M_3^*| \leq 12|M| + 9|M_1^*| + 5|M_2^*| + |M_3^*| \leq 21|M|, \end{aligned} \quad (24)$$

where the last inequality requires that $5|M_2^*| + |M_3^*| \leq 9(|M| - |M_1^*|)$. However, the latter inequality does not admit a constructive argument similar to the cases $l = 3, 4$ (see the proof of Theorem 10). To overcome this complication, consider the following extended problem setting, which surprisingly does admit a constructive argument.

Definition 5 Let R be a set of (parallel) loops on the nodes of S . A subset $M \subseteq E \cup R$ is **(R,d)-distance matching** if it is the union of a d -distance matching and R .

Consider the following extension of Definition 2.

Definition 6 Given an (R, d) -distance matching M and an edge $sv \in (S \times T) \cup R$, let

$$\mathcal{H}_+(sv, M) = \begin{cases} \mathcal{H}(sv, M \setminus R) \cup \{e \in R : e \text{ is incident to nodes } s\}, & \text{if } sv \in S \times T \\ sv, & \text{if } sv \in R. \end{cases}$$

In other words, each $st \in E$ hits the edges hit by $\mathcal{H}(st, M)$ and all the loops incident to node s , while each loop hits only itself. A natural way to define the hit set of multiple edges is as follows.

Definition 7 Given an edge set $X \subseteq E$, let $\mathcal{H}_+(X, M) = \bigcup_{e \in X} \mathcal{H}_+(e, M)$.

Using $\mathcal{H}_+$, the definition of l -locally optimal d -distance matchings can be naturally extended to (R, d) -distance matchings.

Definition 8 An (R, d) -distance matching M is **l -locally optimal** if there exists no d -distance matching $X \subseteq E \setminus M$ such that $l \geq |X| > |\mathcal{H}_+(X, M)|$. Similarly, M is **l -locally optimal with respect to M^*** if there exists no $X \subseteq M^* \setminus M$ such that $l \geq |X| > |\mathcal{H}_+(X, M)|$, where M^* is an (R, d) -distance matching.

Note that each of these definitions reduces to its original counterpart if $R = \emptyset$. Therefore, it suffices to show that ϱ_l is an upper bound on the approximation ratio of (R, l) -locally optimal solutions.

To elaborate on the intuition behind these technical definitions and to understand how R influences locally optimality, suppose that we are given a feasible d -distance matching M , which we want to make l -locally optimal. To this end, one needs to find a d -distance matching $X \subseteq E \setminus M$ of cardinality at most l that hits strictly fewer edges of M than its cardinality. In an (R, d) -distance matching, however, the number of edges hit by such a subset X can be larger because of the loops (as $\mathcal{H}_+(X, M)$ also counts those), meaning that the requirements for l -locally optimality are relaxed. Intuitively, the loops incident to a node $s \in S$ can be thought of as the “resistances” of s : the more loops s has, the less we want to replace the edge of M incident to s with some other edge of $\Delta(s)$. Note, however, that the loops also contribute to the size of the matching, which will be crucial in the proof of the next theorem.

Theorem 11 If M, M^* are (R, d) -distance matchings such that M is l -locally optimal with respect to M^* , then the approximation ratio $|M^*|/|M|$ is at most ϱ_l , where $l \geq 1$ and ϱ_l is as defined above.

Proof As in the proof of Theorem 10, let $M_i^* = \{e^* \in M^* : |\mathcal{H}_+(e^*, M)| = i\}$ for $i \in \mathbb{N}$, and let $M_{i+}^* = \bigcup_{k=i}^{\infty} M_k^*$. Note that $M_0^*, M_1^*, \dots$ is a partition of M^* , for which $R \subseteq M_1^*$ by definition, and $M_0^* = \emptyset$ since each edge of M^* hits at least one edge of M if $l \geq 1$. Similar to (17), observe that each edge $e \in M$ can be hit by at most three edges of M^* , therefore

$$3|M| \geq \sum_{e^* \in M^*} |\mathcal{H}_+(e^*, M)| = \sum_{k=1}^{\infty} k|M_k^*|. \quad (25)$$

The proof is by induction on l . The argument for $l = 1, 2$ is analogous to that in the proof of Theorem 10.

Case 1: $l = 1$.

It easily follows from (25) that

$$|M^*| = \sum_{k=1}^{\infty} |M_k^*| \leq \sum_{k=1}^{\infty} k|M_k^*| \leq 3|M|. \quad (26)$$

Case 2: $l = 2$. Similarly,

$$2|M^*| = 2 \sum_{k=1}^{\infty} |M_k^*| \leq |M_1^*| + \sum_{k=1}^{\infty} k|M_k^*| \leq |M_1^*| + 3|M| \leq 4|M|, \quad (27)$$

where the second inequality follows from (25) and the third one holds because M is 2-locally optimal with respect to M^* .

Case 3: $l \geq 3$. One has to show that $|M^*|/|M| \leq 4\varrho_{l-2} - 3/2\varrho_{l-2} - 1$.

First, introduce the notation $\alpha(M, M^*) = \sum_{k=3}^{\infty} (k-2)|M_k^*|$.

In the following computation, inequality (25) is forced with an appropriate multiplier so that the rest admits the application of case $l-2$ to a derived problem instance (see Lemma 3).

Note that the approach is similar to computations (20) and (22).

$$\begin{aligned}
 (2\varrho_{l-2} - 1)|M^*| &= (2\varrho_{l-2} - 1) \sum_{k=1}^{\infty} |M_k^*| = (\varrho_{l-2} - 1) \sum_{k=1}^{\infty} k|M_k^*| \\
 &\quad + \sum_{k=1}^{\infty} ((k-1) - (k-2)\varrho_{l-2})|M_k^*| \\
 &\leq 3(\varrho_{l-2} - 1)|M| + \sum_{k=1}^{\infty} ((k-1) - (k-2)\varrho_{l-2})|M_k^*| \\
 &= 3(\varrho_{l-2} - 1)|M| + \varrho_{l-2}|M_1^*| + |M_2^*| + \sum_{k=3}^{\infty} (k-1)|M_k^*| - \varrho_{l-2} \sum_{k=3}^{\infty} (k-2)|M_k^*| \\
 &= 3(\varrho_{l-2} - 1)|M| + \varrho_{l-2}|M_1^*| + |M_2^*| + \sum_{k=3}^{\infty} |M_k^*| + \alpha(M, M^*) - \varrho_{l-2}\alpha(M, M^*) \\
 &= 3(\varrho_{l-2} - 1)|M| + \varrho_{l-2}|M_1^*| + |M_2^*| + |M_{3+}^*| + \alpha(M, M^*) - \varrho_{l-2}\alpha(M, M^*) \\
 &\leq 3(\varrho_{l-2} - 1)|M| + \varrho_{l-2}|M| = (4\varrho_{l-2} - 3)|M|, \tag{28}
 \end{aligned}$$

where the first inequality holds by (25) and the last one by the following lemma. Note that if $R = \emptyset$, then $\alpha(M, M^*) = |M_3^*|$ and $|M_k| = 0$ for $k \geq 4$, hence (28) gives back (20) and (22) for $l = 3, 4$, respectively. The following lemma completes the proof of (28).

Lemma 3 *If $l \geq 3$ and $M, M^*, \alpha(M, M^*)$ are as above, then*

$$|M_{2+}^*| + \alpha(M, M^*) \leq \varrho_{l-2}(|M| - |M_1^*| + \alpha(M, M^*)) \tag{29}$$

Proof It suffices to show that if M is l -locally optimal with respect to M^* , then there exist $\tilde{M}$, $\tilde{M}^*$ and $\tilde{R}$ such that

- (1) $\tilde{M}$ and $\tilde{M}^*$ are $(\tilde{R}, d)$ -distance matchings,
- (2) $|\tilde{M}| = |M| - |M_1^*| + \alpha(M, M^*)$,
- (3) $|\tilde{M}^*| = |M_{2+}^*| + \alpha(M, M^*)$,
- (4) $|\tilde{R}| = \alpha(M, M^*)$,
- (5) $\tilde{M}$ is $(l-2)$ -locally optimal with respect to $\tilde{M}^*$.

Then, condition 5) implies that $|\tilde{M}^*| \leq \varrho_{l-2}|\tilde{M}|$ holds by induction, from which one obtains (29) by substituting 2) and 3). We define $\tilde{R}$, $\tilde{M}$ and $\tilde{M}^*$ such that

$$\begin{aligned}
 \tilde{R} &= \bigcup_{s^*t^* \in M_{3+}^*} \{|\mathcal{H}_+(s^*t^*, M)| - 2 \text{ parallel loops incident to } s^*\}, \\
 \tilde{M} &= M \setminus \mathcal{H}_+(M_1^*, M) \cup \tilde{R} \text{ and} \\
 \tilde{M}^* &= M_{2+}^* \cup \tilde{R}.
 \end{aligned}$$

It is easy to see that $\tilde{M}$, $\tilde{M}^*$ and $\tilde{R}$ fulfill 1)-4). In the rest of the proof, we argue that 5) holds, as well. By contradiction, suppose that 5) does not hold, that is, there exists $Z \subseteq \tilde{M}^* : l-2 \geq |Z| > |\mathcal{H}_+(Z, \tilde{M})|$. Assume that the instance of the problem at hand is minimal in the sense that $|M| + |M^*| + |\tilde{M}| + |\tilde{M}^*| + |Z|$ is minimal. First, various useful properties of minimal problem instances are derived. Note that $|Z| = |\mathcal{H}_+(Z, M)| + 1$ can be assumed, otherwise $|Z| > |\mathcal{H}_+(Z, M)| + 1$ and therefore one could have removed an arbitrary edge from Z .

Observe that if an edge $e \in \mathcal{H}_+(Z, \tilde{M})$ were hit by a sole edge $e^* \in Z$, then $l-2 \geq |Z \setminus \{e^*\}| > |\mathcal{H}_+(Z \setminus \{e^*\}, \tilde{M})|$ would hold, i.e. one could have left out e^* from Z . Therefore,

each edge $e \in \mathcal{H}_+(Z, \tilde{M})$ is hit by at least two edges of Z . This also implies that $s^*t^* \in Z$ if and only if $\{e \in R : e \text{ is incident to } s^*\} \subseteq Z$. Using this, $Z = \tilde{M}^*$ follows, because removing all edges $\tilde{M}^* \setminus Z$ from M^* and all those loops from R that are incident to the removed edges, one obtains a smaller instance (where $\alpha(M, M^*)$, $\tilde{R}$, $\tilde{M}$ and $\tilde{M}^*$ need to be adjusted appropriately after the edge-removal), which satisfies 1)-4) but not 5). Clearly, M remains l -locally optimal with respect to M^* after the edge-removal. So, one can assume that $Z = \tilde{M}^*$.

A minimal instance also fulfills that there exist no edges $e \in M$ and $e^* \in M_1^*$ such that e is not hit by any edge of M_{2+}^* , (that is, $\mathcal{H}_+(e^*, M) \setminus \mathcal{H}_+(M_{2+}^*, M) = \emptyset$), otherwise the removal of e and e^* results in a smaller instance satisfying 1)-4) but not 5). Observe that after the removal, M remains l -locally optimal with respect to M^* , because there exists no $X \subseteq M^* \setminus \{e^*\}$ such that $e \in \mathcal{H}_+(X, M)$ (since e is not hit by any edge of M_{2+}^*), therefore if the new instance were not l -locally optimal, then the original instance would not have been either. So, one can assume that $\mathcal{H}_+(M_1^*, M) \setminus \mathcal{H}_+(M_{2+}^*, M) = \emptyset$.

Now we are ready to derive that $|\mathcal{H}_+(M^*, M)| < |M^*| \leq l$ holds — contradicting that M is l -locally optimal. The first inequality is shown by

$$\begin{aligned} |\mathcal{H}_+(M^*, M)| &= |\mathcal{H}_+(M_{2+}^*, M)| = |\mathcal{H}_+(M_{2+}^*, M) \cap \mathcal{H}_+(M_1^*, M)| \\ &\quad + |\mathcal{H}_+(M_{2+}^*, M) \setminus \mathcal{H}_+(M_1^*, M)| = |\mathcal{H}_+(M_1^*, M)| + |\mathcal{H}_+(M_{2+}^*, M) \setminus \mathcal{H}_+(M_1^*, M)| \\ &= |M_1^*| + |\mathcal{H}_+(M_{2+}^*, M) \setminus \mathcal{H}_+(M_1^*, M)| = |M_1^*| + |\mathcal{H}_+(M_{2+}^*, M \setminus \mathcal{H}_+(M_1^*, M))| \\ &= |M_1^*| + |\mathcal{H}_+(M_{2+}^*, \tilde{M} \setminus \tilde{R})| = |M_1^*| + |\mathcal{H}_+(Z, \tilde{M}) \setminus \tilde{R}| = |M_1^*| + |Z| - 1 - |\tilde{R}| \\ &= |M_1^*| + |\tilde{M}^*| - 1 - |\tilde{R}| = |M_1^*| + |M_{2+}^*| + |\tilde{R}| - 1 - |\tilde{R}| = |M^*| - 1. \end{aligned} \quad (30)$$

Next, we show that $|M^*| \leq l$.

$$\begin{aligned} |M^*| &= |M_{2+}^*| + |M_1^*| = |M_{2+}^*| + |\mathcal{H}_+(M_1^*, M)| \\ &= |M_{2+}^*| + |\mathcal{H}_+(M_{2+}^*, M) \cap \mathcal{H}_+(M_1^*, M)| \\ &= |M_{2+}^*| + \left| \bigcup_{e^* \in M_{2+}^*} \mathcal{H}_+(e^*, M) \cap \mathcal{H}_+(M_1^*, M) \right| \\ &\leq |M_{2+}^*| + \sum_{e^* \in M_{2+}^*} |\mathcal{H}_+(e^*, M) \cap \mathcal{H}_+(M_1^*, M)| \\ &= |M_{2+}^*| + \sum_{e^* \in M_{2+}^*} (|\mathcal{H}_+(e^*, M)| - |\mathcal{H}_+(e^*, M) \setminus \mathcal{H}_+(M_1^*, M)|) \\ &\leq |M_{2+}^*| + \sum_{e^* \in M_{2+}^*} |\mathcal{H}_+(e^*, M)| - 2|\mathcal{H}_+(M_{2+}^*, M) \setminus \mathcal{H}_+(M_1^*, M)| \\ &= |M_{2+}^*| + \sum_{e^* \in M_{2+}^*} |\mathcal{H}_+(e^*, M)| - 2(|\mathcal{H}_+(Z, \tilde{M}) \setminus \tilde{R}|) \\ &= |M_{2+}^*| + \sum_{e^* \in M_{2+}^*} |\mathcal{H}_+(e^*, M)| - 2(|M_{2+}^*| - 1) \\ &= |M_{2+}^*| + 2|M_{2+}^*| + |\tilde{R}| - 2(|M_{2+}^*| - 1) \\ &= |M_{2+}^*| + |\tilde{R}| + 2 = |\tilde{M}^*| + 2 = |Z| + 2 \leq l, \end{aligned} \quad (31)$$

where the second inequality holds by the following computation.

$$2|\mathcal{H}_+(M_{2+}^*, M) \setminus \mathcal{H}_+(M_1^*, M)| = 2|\mathcal{H}_+(M_{2+}^*, M \setminus \mathcal{H}_+(M_1^*, M))|$$

$$\begin{aligned}
&= 2|\mathcal{H}_+(M_{2+}^*, \tilde{M} \setminus \tilde{R})| = 2|\mathcal{H}_+(\tilde{M}^*, \tilde{M} \setminus \tilde{R})| \\
&= 2|\mathcal{H}_+(Z, \tilde{M} \setminus \tilde{R})| \leq \sum_{e^* \in Z} |\mathcal{H}_+(e^*, \tilde{M} \setminus \tilde{R})| \\
&= \sum_{e^* \in \tilde{M}_{2+}^*} |\mathcal{H}_+(e^*, M \setminus \mathcal{H}_+(M_1^*, M))| \\
&= \sum_{e^* \in \tilde{M}_{2+}^*} |\mathcal{H}_+(e^*, M) \setminus \mathcal{H}_+(M_1^*, M)|, \quad (32)
\end{aligned}$$

where the inequality holds because each edge of $\mathcal{H}_+(Z, \tilde{M})$ is hit at least twice by Z .

Combining (30) and (31), one obtains that $|\mathcal{H}_+(M^*, M)| < |M^*| \leq l$, which contradicts that M is l -locally optimal with respect to M^* . Hence — in contrast to the indirect assumption — condition 5) holds, and this proves the lemma. $\square$

Note that if $R = \emptyset$, then Lemma 3 gives back (20) and (22) for $l = 3, 4$, respectively. By Lemma 3, inequality (28) follows, meaning that the desired recursion (16) gives a valid upper bound on the approximation ratio of the l -locally optimal solutions. $\square$

Corollary 1 *The approximation ratio of l -locally optimal d -distance matchings is at most q_l , where q_l is as defined above.*

Proof Let M^* denote an optimal d -distance matching. By definition, M is l -locally optimal with respect to M^* , therefore M is $(\emptyset, l)$ -locally optimal with respect to M^* . By Theorem 11, one gets $|M^*|/|M| \leq q_l$, which completes the proof. $\square$

Corollary 2 *For any constant $\epsilon > 0$, there exist a polynomial-time algorithm for the unweighted d -distance matching problem that achieves an approximation guarantee of $3/2 + \epsilon$.*

Proof By Corollary 1, the approximation ratio of l -locally optimal solutions is at most q_l . One might easily show that $\lim_{l \rightarrow \infty} q_l = 3/2$. Hence for any $\epsilon > 0$, there exists $l_0 \in \mathbb{N}$ such that $q_l \leq 3/2 + \epsilon$. To complete the proof, observe that l_0 is independent from the problem size, therefore one can compute an l_0 -locally optimal solution in polynomial time. Note that the number of improvements is at most the size of the matching and hence it is polynomial as well. $\square$

Remark 4 Figures 4a, 7 and 8 show that the upper bound on the approximation ratio of l -locally optimal solutions given by Theorem 11 is tight for $l = 1, 2$ and 3, respectively. It remains open whether the analysis is tight for $l \geq 4$.

Remark 5 Similar proof shows that for any constant $\epsilon > 0$, the above local-search algorithm is a $(3/2 + \epsilon)$ -approximation algorithm for the unweighted cyclic d -distance matching problem.

Fig. 7 The wavy edges form a 2-locally optimal 2-distance matching M , and $M^* = E \setminus M$ is the optimal 2-distance matching. The approximation ratio $|M^*|/|M|$ is $q_2 = 2$

Fig. 8 The wavy edges form a 3-locally optimal 5-distance matching M , and $M^* = E \setminus M$ is the optimal 5-distance matching. The approximation ratio $|M^*|/|M|$ is $\varrho_3 = 9/5$

5 Regular distance matching

The following theorem is a straightforward generalization of the well-known result that every regular bipartite graph has a perfect matching.

Definition 9 An instance of the d -distance matching problem is r -regular if $\deg(s) = r$ for each $s \in S$ and the number of edges between t and $R_d(s_i)$ is r for each $t \in T$ and $i = 1, \dots, n - d + 1$.

Theorem 12 *If a problem instance is r -regular, then there exists a perfect d -distance matching.*

Proof There exists a perfect matching between $\{s_1, \dots, s_d\}$ and T , because the induced graph is r -regular. By induction, assume that the degrees of $\{s_1, \dots, s_{i-1}\}$ in M are one, where $i - 1 \geq d$. Let t denote the node that M assigns to s_{i-d+1} . If $s_i t \notin E$, then the number of edges between t and $L_d(s_i)$ is $r - 1$, meaning that the instance at hand is not r -regular, hence $s_i t \in E$. Therefore, $M \cup \{s_i t\}$ is feasible for the first i nodes of S , hence the claim follows. $\square$

If we leave out a perfect d -distance matching from an r -regular problem instance, then it becomes an $r - 1$ regular instance, hence one gets the following generalization of König's edge-coloring theorem Frank (2011, p. 74).

Corollary 3 *If a problem instance is r -regular, then the edge set of the graph partitions into r perfect d -distance matchings.*

6 Conclusion

This paper introduced the d -distance matching problem. We proved that the problem is NP-complete in general and admits a 3-approximation. We gave an FPT algorithm parameterized by d and also settled the case when the size of T is constant. The integrality gap of the natural integer programming model is shown to be at most $2 - \frac{1}{2d-1}$, and an LP-based approximation algorithm for the weighted case is given with the same guarantee. Using a different approach, a combinatorial $(2 - \frac{1}{d})$ -approximation algorithm was also described. Several greedy approaches, including a local search algorithm, were presented. The latter method achieves an approximation ratio of $3/2 + \epsilon$ for any constant $\epsilon > 0$ in the unweighted case.

The problem itself has several generalizations (e.g. pose degree bounds on the nodes of both S and T , cyclic version of the problem, distance-constraints on both node classes, etc.), which are subjects for further research.

Acknowledgements The project has been supported by the ÚNKP-20-3 New National Excellence Program of the Ministry for Innovation and Technology. The author is grateful to Kristóf Bérczi and Alpár Jüttner for their comments that greatly improved an earlier version of the manuscript. The author is indebted to an anonymous reviewer of an earlier version of the manuscript for suggesting an approach which led to the algorithm described in Sect. 3.5. The author is also grateful to both anonymous referees of this manuscript for their valuable comments and suggestions.

Funding Open access funding provided by Eötvös Loránd University.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aardal, K. I., van Hoesel, S. P. M., Koster, A. M. C. A., Mannino, C., & Sassano, A. (2007). Models and solution techniques for frequency assignment problems. *Annals of Operations Research*, 153(1), 79–129. <https://doi.org/10.1007/s10479-007-0178-0>.
- Baste, J., Rautenbach, D., & Sau, I. (2019). Approximating maximum uniquely restricted matchings in bipartite graphs. *Discrete Applied Mathematics*, 267, 30–40.
- Bérczi, K., & Végh, L. A. (2010). Restricted b-matchings in degree-bounded graphs. In F. Eisenbrand & F. B. Shepherd (Eds.), *Integer programming and combinatorial optimization* (pp. 43–56). Berlin: Springer.
- Downey, R. G., & Fellows, M. R. (2013). *Parameterized complexity*. Berlin: Springer Publishing Company.
- Frank, A. (2011). *Connections in combinatorial optimization*. Oxford: Oxford University Press.
- Fürst, M., & Rautenbach, D. (2019). On some hard and some tractable cases of the maximum acyclic matching problem. *Annals of Operations Research*, 279(1–2), 291–300.
- Garey, M. R., & Johnson, D. S. (1979). *Computers and intractability: A guide to the theory of NP-completeness (Series of books in the mathematical sciences), first edition*. W. H: Freeman.
- Itai, A., Rodeh, M., & Tanimoto, S. (1978). Some matching problems for bipartite graphs. *Journal of the ACM*, 25, 517–525. <https://doi.org/10.1145/322092.322093>.
- Karp, R. (1972). Reducibility among combinatorial problems. In R. Miller & J. Thatcher (Eds.), *Complexity of computer computations* (pp. 85–103). New York: Plenum Press.
- Madarasi, P. (2020). The distance matching problem. In M. Baiou, B. Gendron, O. Günlük, & A. R. Mahjoub (Eds.), *Combinatorial optimization* (pp. 202–213). Cham: Springer International Publishing.
- Makai, M. (2007). On maximum cost $k_{t,t}$ -free t -matchings of bipartite graphs. *SIAM Journal on Discrete Mathematics*, 21(2), 349–360. <https://doi.org/10.1137/060652282>.
- Pap, G.: Alternating paths revisited ii: Restricted b-matchings in bipartite graphs. EGRES Technical Report TR-2005-13 (2005)
- Tardos, E. (1985). A strongly polynomial minimum cost circulation algorithm. *Combinatorica*, 5, 247–255.
- Zeithofer, T., & Wess, B. (2003). List-coloring of interval graphs with application to register assignment for heterogeneous register-set architectures. *Signal Processing*, 83, 1411–1425.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.