





Deterministic Primal-Dual Algorithms for Online k -Way Matching with Delays

Naonori Kakimura¹  and Tomohiro Nakayoshi² 

¹ Keio University, 3-14-1 Hiyoshi, Kohoku-ku, Yokohama, Kanagawa 223-8522, Japan
kakimura@math.keio.ac.jp

² The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8654, Japan
nakayoshi-tomohiro@g.ecc.u-tokyo.ac.jp

Abstract. In this paper, we study the Min-cost Perfect k -way Matching with Delays (k -MPMD), recently introduced by Melnyk et al. In the problem, m requests arrive one-by-one over time in a metric space. At any time, we can irrevocably make a group of k requests who arrived so far, that incurs the distance cost among the k requests in addition to the sum of the waiting cost for the k requests. The goal is to partition all the requests into groups of k requests, minimizing the total cost. The problem is a generalization of the min-cost perfect matching with delays (corresponding to 2-MPMD). It is known that no online algorithm for k -MPMD can achieve a bounded competitive ratio in general, where the competitive ratio is the worst-case ratio between its performance and the offline optimal value. On the other hand, k -MPMD is known to admit a randomized online algorithm with competitive ratio $O(k^5 \log n)$ for a certain class of k -point metrics called the H -metric, where n is the size of the metric space. In this paper, we propose a deterministic online algorithm with a competitive ratio of $O(mk^2)$ for the k -MPMD in H -metric space. Furthermore, we show that the competitive ratio can be improved to $O(m + k^2)$ if the metric is given as a diameter on a line.

Keywords: Online Matching · Online Algorithm · Competitive Analysis

1 Introduction

Consider an online gaming platform supporting two-player games such as Chess. In such a platform, players arrive one-by-one over time, and stay in a queue to participate in a match. The platform then tries to suggest a suitable opponent for each player from the queue. In order to satisfy the players, the platform aims to maximize the quality of the matched games. Specifically, we aim to minimize the distance of the matched players (e.g., the difference of their ratings) as well as the sum of the players' waiting time.

This work was supported by JSPS KAKENHI Grant Numbers JP20H05795, JP22H05001, and JP21H03397.

The above situation can be modeled as the problem called Online Matching with Delays, introduced by Emek et al. [11]. In the setting, arriving requests (or players) are embedded in a metric space so that the distance of each pair is determined. For the Online Matching with Delays, Emek et al. [11] proposed a randomized algorithm with a competitive ratio of $O(\log^2 n + \log \Delta)$, where n is the number of points in a metric space, and Δ is the ratio of the maximum to minimum distance between two points. The competitive ratio was later improved to $O(\log n)$ by Azar et al. [5]. We remark that both algorithms require that a metric space is finite and all the points in the metric space are known in advance (we note that arriving requests may be embedded into the same point more than once). Bienkowski et al. [9] presented a primal-dual algorithm with a competitive ratio of $O(m)$, where m is the number of requests. Another algorithm with a better competitive ratio of $O(m^{0.59})$ was proposed by Azar et al. [6].

In this paper, we consider a generalization of Online Matching with Delays, called the Min-cost Perfect k -way Matching with Delays (k -MPMD) [22]. In the problem, requests arrive one-by-one over time. At any time, instead of choosing a pair of requests, we make a group of k requests. This corresponds to an online gaming platform that allows more than two players to participate, such as mahjong ($k = 4$), Splatoon ($k = 8$), Apex Legends ($k = 60$), and Fortnite ($k = 100$). Then we aim to partition all the requests into groups of size- k subsets, minimizing the sum of the distance of the requests in the same group and the total waiting time.

To generalize to k -MPMD, it is necessary to measure the distance of a group of $k > 2$ requests. That is, we need to introduce a metric space that defines distances for any subset of k points. Although there are many ways of generalizing a standard distance between two points to $k > 2$ points in the literature [4, 16], Melnyk et al. [22] showed that most known generalized metrics on k points cannot achieve a bounded competitive ratio for the k -MPMD. Melnyk et al. [22] then introduced a new interesting class of generalized metric, called H -metric, and proposed a randomized algorithm for the k -MPMD on H -metric with a competitive ratio of $O(k^5 \log n)$, extending Azar et al. [5].

The main contribution of this paper is to propose a deterministic algorithm for the k -MPMD on H -metric with a competitive ratio of $O(mk^2)$, where m is the number of requests. The proposed algorithm adopts a primal-dual algorithm based on a linear programming relaxation of the k -MPMD.

To design a primal-dual algorithm, we first formulate a linear programming relaxation of the offline problem, that is, when a sequence of requests is given in advance. We remark that even the offline setting is NP-hard when $k \geq 3$, as it includes the triangle packing problem. We first show that H -metric can be approximated by a standard metric (Theorem 1). This allows us to construct a linear programming problem with variables for each pair of requests such that the optimal value gives a lower bound on the offline version of the k -MPMD. Using the linear programming problem, we can design a primal-dual algorithm by extending the one by Bienkowski et al. [9] for Online Matching with Delays. We show that, by the observation on H -metric (Theorem 1) again, the cost

of the output can be upper-bounded by the dual objective value of our linear programming problem.

An interesting special case of the H -metric is the diameter on a line. That is, points are given on a 1-dimensional line, and the distance of k points is defined to be the maximum difference in the k points. In the context of an online gaming platform, the diameter on a line can be interpreted as the difference of players' ratings. In this case, we show that the competitive ratio of our algorithm can be improved to $O(m + k^2)$. Moreover, we construct an instance such that our algorithm achieves the competitive ratio of $\Omega(m/k)$.

Related Work. An online algorithm for the matching problem was first introduced by Karp et al. [15]. They considered the online bipartite matching problem where arriving requests are required to match upon their arrival. Since then, the problem has been studied extensively in theory and practice. For example, motivated by internet advertising, Mehta et al. [20] generalized the problem to the AdWords problem. See also Mehta [21] and Goel and Mehta [13]. The weighted variant of the online bipartite matching problem is considered in the literature. It includes the vertex-weighted online bipartite matching [1], the problem with metric costs [8, 23, 25], and the problem with line metric cost [2, 12, 14, 17]. We remark that the edge-weighted online bipartite matching in general has no online algorithm with bounded competitive ratio [1].

This paper deals with a variant of the online matching problem with delays, in which arriving requests are allowed to make decision later with waiting costs. Besides the related work [5, 6, 9, 11] mentioned before, Liu et al. [18] extended the problem to the one with non-linear waiting costs. Other delay costs are studied in [7, 10, 19]. Ashlagi et al. [3] studied the online matching problem with deadlines, where each arriving request has to make a decision by her deadline. Pavone et al. [24] considered online hypergraph matching with deadlines.

Paper Organization. This paper is organized as follows. In Sect. 2, we formally define the minimum-cost perfect k -way matching problem and H -metric. We also discuss useful properties of H -metrics which will be used in our analysis. In Sect. 3, we present our main algorithm for the k -MPMD on H -metric. In Sect. 4, we show that there exists an instance such that our algorithm admits an almost tight competitive ratio. Due to the space limitation, the proofs of lemmas and theorems are omitted, which may be found in the full version of this paper.

2 Preliminaries

2.1 Minimum-Cost Perfect k -Way Matching with Delays

In this section, we formally define the problem k -MPMD. Let (χ, d) be a generalized metric space where χ is a set and $d : \chi^k \rightarrow [0, \infty)$ represents a distance among k elements.

In the problem, m requests u_1, u_2, \dots, u_m arrive one-by-one in this order. The arrival time of u_i is denoted by $\text{atime}(u_i)$. When u_i arrives, the location $\text{pos}(u_i)$ of u_i in the metric space χ is revealed. Thus, an instance of the problem is given as a tuple $\sigma = (V, \text{atime}, \text{pos})$, where $V = \{u_1, \dots, u_m\}$, $\text{atime} : V \rightarrow \mathbb{R}_+$, and $\text{pos} : V \rightarrow \chi$ such that $\text{atime}(u_1) \leq \dots \leq \text{atime}(u_m)$. We note that m may be unknown in advance, but we assume that m is a multiple of k .

At any time τ , with the only information for requests arrived so far, an online algorithm can make a set of k requests v_1, \dots, v_k in V , where we say that v_1, \dots, v_k are *matched*, if they satisfy the following two conditions: (a) The requests v_1, \dots, v_k have already arrived, that is, $\text{atime}(v_i) \leq \tau$ for any $i = 1, \dots, k$; (b) None of v_1, \dots, v_k has been matched to other requests yet. The cost to match v_1, \dots, v_k at time τ is defined to be

$$d(\text{pos}(v_1), \text{pos}(v_2), \dots, \text{pos}(v_k)) + \sum_{i=1}^k (\tau - \text{atime}(v_i)).$$

The first term means the distance cost among the k requests and the second term is the total waiting cost of the k requests.

The objective of the problem is to design an online algorithm that matches all the requests, minimizing the total cost. In other words, an online algorithm finds a family of disjoint subsets of size k that covers all the requests. We call a family of disjoint subsets of size k a k -way *matching*, and a k -way matching is called *perfect* if it covers all the requests.

To measure the performance of an online algorithm, we define the competitive ratio. For an instance σ , let $\mathcal{ALG}(\sigma)$ be the cost incurred by the online algorithm, and let $\mathcal{OPT}(\sigma)$ be the optimal cost when we know in advance a sequence of requests V as well as $\text{atime}(u_i)$ and $\text{pos}(u_i)$ for each request u_i . The *competitive ratio* of the online algorithm is defined as $\sup_{\sigma} \frac{\mathcal{ALG}(\sigma)}{\mathcal{OPT}(\sigma)}$.

2.2 H -Metric

In this section, we define H -metric, introduced by Melnyk et al. [22]. Recall that a function $d : \chi^2 \rightarrow [0, \infty)$ is called a *distance function* (or a *metric*) if d satisfies the following three axioms:

- (**Symmetry**) $d(p_1, p_2) = d(p_2, p_1)$ for any $p_1, p_2 \in \chi$.
- (**Positive definiteness**) $d(p_1, p_2) \geq 0$ for any $p_1, p_2 \in \chi$, and $d(p_1, p_2) = 0$ if and only if $p_1 = p_2$.
- (**Triangle inequality**) $d(p_1, p_3) \leq d(p_1, p_2) + d(p_2, p_3)$ for any $p_1, p_2, p_3 \in \chi$.

We first define a k -point metric as a k -variable function satisfying generalizations of the symmetry axiom and the positive definiteness axiom.

Definition 1. We call a function $d : \chi^k \rightarrow [0, \infty)$ a k -point metric if it satisfies the following two axioms.

Π : For any permutation π of $\{p_1, \dots, p_k\}$, we have $d(p_1, \dots, p_k) = d(\pi(p_1), \dots, \pi(p_k))$.

O_D : It holds that $d(p_1, \dots, p_k) \geq 0$. Moreover, $d(p_1, \dots, p_k) = 0$ if and only if $p_1 = p_2 = \dots = p_k$.

There are several ways of generalizing the triangle inequality to k -variable functions. One possibility is the following axiom: for any $p_1, \dots, p_k, a \in \chi$ and any $i \in \{1, \dots, k\}$, it holds that

$$\Delta_H : d(p_1, \dots, p_k) \leq d(p_1, \dots, p_i, \underbrace{a, \dots, a}_{k-i}) + d(\underbrace{a, \dots, a}_i, p_{i+1}, \dots, p_k).$$

We note that it is identical to the triangle inequality when $k = 2$.

For a multiset S on χ , we denote by $elem(S)$ the set of all distinct elements contained in S . In addition to the generalized triangle inequality, we consider the relationship between $d(p_1, \dots, p_k)$ and $d(p'_1, \dots, p'_k)$ when $elem(\{p_1, \dots, p_k\}) \subseteq elem(\{p'_1, \dots, p'_k\})$. The *separation axiom* \mathcal{S}_H says that, for some nonnegative integer $\gamma \leq k - 1$,

$$\begin{aligned} d(p_1, \dots, p_k) &\leq d(p'_1, \dots, p'_k) \quad \text{if } elem(\{p_1, \dots, p_k\}) \subset elem(\{p'_1, \dots, p'_k\}), \\ d(p_1, \dots, p_k) &\leq \gamma \cdot d(p'_1, \dots, p'_k) \quad \text{if } elem(\{p_1, \dots, p_k\}) = elem(\{p'_1, \dots, p'_k\}). \end{aligned}$$

The H -metric is a k -point metric that satisfies all the above axioms.

Definition 2 (Melnik et al. [22]). A k -point metric $d_H : \chi^k \rightarrow [0, \infty)$ is an H -metric with parameter $\gamma \leq k - 1$ if it satisfies Π , O_D , Δ_H and \mathcal{S}_H with parameter γ .

We remark that there are weaker conditions than Δ_H and \mathcal{S}_H , generalizing the triangle inequality, which yields other classes of k -point metrics such as the n -metric [4] and the K -metric [16]. See [22] for the formal definition. Melnyk et al. [22], however, showed that the k -MPMD cannot be solved for such more general metrics. Specifically, they proved that there exists no randomized algorithm for the k -MPMD ($k \geq 5$) problem on n -metric or K -metric against an oblivious adversary that has a competitive ratio which is bounded by a function of the number of points n .

2.3 Properties of H -Metric

In this section, we discuss approximating H -metric by a standard metric, and present specific examples of H -metric.

Melnik et al. proved that H -metric can be approximated by the sum of distances between all pairs [22, Theorem 6]. We refine their results as in the theorem below, which will be used in the next section.

Theorem 1. Let d_H be an H -metric on χ with parameter γ . Define a metric $d : \chi^2 \rightarrow [0, \infty)$ as

$$d(p_1, p_2) := d_H(p_1, p_2, \dots, p_2) + d_H(p_2, p_1, \dots, p_1)$$

for any $p_1, p_2 \in \chi$. Then it holds that

$$\frac{1}{\gamma k^2} \cdot \sum_{i=1}^{k-1} \sum_{j=i+1}^k d(p_i, p_j) \leq d_H(p_1, \dots, p_k) \leq \sum_{i=1}^k d(v, p_i), \tag{1}$$

for all $v \in \{p_1, \dots, p_k\}$.

We conclude this section with providing specific examples of H -metric. We note that the examples below satisfy that $\gamma = 1$, and thus the approximation factor in Theorem 1 becomes small.

Let $d : \chi^2 \rightarrow [0, \infty)$ be a distance function. We define a k -point metric d_{\max} by $d_{\max}(p_1, \dots, p_k) = \max_{i,j \in \{1, \dots, k\}} d(p_i, p_j)$. Then it turns out to be an H -metric.

Proposition 1. *Let $d : \chi^2 \rightarrow [0, \infty)$ be a distance function. Then the k -point metric d_{\max} is an H -metric with $\gamma = 1$.*

For real numbers $p_1, \dots, p_k \in \mathbb{R}$, we define the *diameter on a line* as $d_D(p_1, \dots, p_k) = \max_{i,j \in \{1, \dots, k\}} |p_i - p_j|$. By Proposition 1, d_D is an H -metric.

For a distance function $d : \chi^2 \rightarrow [0, \infty)$, we define another H -metric d_{HC} by

$$d_{HC}(p_1, \dots, p_k) = \min \left\{ \sum_{e \in C} d(e) \mid C \subseteq \binom{\chi}{2}, C \text{ forms a Hamiltonian circuit in } \{p_1, \dots, p_k\} \right\}$$

where $\binom{\chi}{2} = \{(p, q) \mid p, q \in \chi, p \neq q\}$. This means that $d_{HC}(p_1, \dots, p_k)$ equals to the minimum cost of a Hamiltonian circuit contained in $\{p_1, \dots, p_k\}$ with respect to cost d .

Proposition 2. *Let $d : \chi^2 \rightarrow [0, \infty)$ be a distance function. Then the k -point metric d_{HC} is an H -metric with parameter $\gamma = 1$.*

3 k -MPMD on H -Metric Space

This section proposes a primal-dual algorithm for k -MPMD on H -metric space. Let (χ, d_H) be an H -metric space with parameter γ .

3.1 Linear Programming Relaxation

This subsection introduces a linear programming relaxation for computing the offline optimal value $OPT(\sigma)$ for a given instance σ .

We first give some notation. Let $\mathcal{E} = \{F \subseteq V \mid |F| = k\}$. For any subset $S \subseteq V$, we denote $\text{sur}(S) = |S| \bmod k$, which is the number of remaining requests when we make a k -way matching of size $\lfloor |S|/k \rfloor$ among S . We denote $\Delta(S) = \{F \in \mathcal{E} \mid F \cap S \neq \emptyset, F \setminus S \neq \emptyset\}$, which is the family of k request sets that intersect both S and $V \setminus S$.

Preparing a variable x_F for any subset $F \in \mathcal{E}$, we define a linear programming problem:

$$(\mathcal{P}) \quad \left\{ \begin{array}{ll} \min. & \sum_{F \in \mathcal{E}} \text{opt-cost}(F) \cdot x_F \\ \text{s.t.} & \sum_{F \in \Delta(S)} x_F \geq \left\lceil \frac{\text{sur}(S)}{k} \right\rceil, \quad \forall S \subseteq V \\ & x_F \geq 0, \quad \forall F \in \mathcal{E} \end{array} \right. \quad (2)$$

where, for any $F = (v_1, \dots, v_k) \in \mathcal{E}$, we define

$$\text{opt-cost}(F) := d_H(\text{pos}(v_1), \dots, \text{pos}(v_k)) + \sum_{i=1}^k \left(\max_j \text{atime}(v_j) - \text{atime}(v_i) \right).$$

Notice that $\text{opt-cost}(F)$ is the cost of choosing F at the moment when all the requests in F have arrived.

Let \mathcal{M} be a perfect k -way matching with optimal cost $\mathcal{OPT}(\sigma)$. Define a 0-1 vector $(x_F)_{F \in \mathcal{E}}$ such that $x_F = 1$ if and only if $F \in \mathcal{M}$. Then the vector satisfies the constraint (2). Moreover, the cost incurred by $F \in \mathcal{M}$ is equal to $\text{opt-cost}(F)$. This is because the optimal algorithm that returns \mathcal{M} chooses F at the moment when all the requests in F have arrived. Thus the objective value for the vector $(x_F)_{F \in \mathcal{E}}$ is equal to $\mathcal{OPT}(\sigma)$, and hence the optimal value of (\mathcal{P}) gives a lower bound of $\mathcal{OPT}(\sigma)$.

We further relax the above LP (\mathcal{P}) by replacing x_F 's with variables for all pairs of requests. Let $E = \{(u, v) \mid u, v \in V, u \neq v\}$, and we prepare a variable x_e for any $e \in E$. We often call an element in E an *edge*.

We denote by $\delta(S)$ the set of pairs between S and $V \setminus S$. Define the following linear programming problem:

$$(\mathcal{P}') \quad \left\{ \begin{array}{ll} \min. & \sum_{e \in E} \frac{1}{\gamma k^2} \cdot \text{opt-cost}(e) \cdot x_e \\ \text{s.t.} & \sum_{e \in \delta(S)} x_e \geq \text{sur}(S) \cdot (k - \text{sur}(S)), \quad \forall S \subseteq V \\ & x_e \geq 0, \quad \forall e \in E \end{array} \right. \quad (3)$$

where, for any $e = (v_1, v_2) \in E$ with $p_1 = \text{pos}(v_1)$ and $p_2 = \text{pos}(v_2)$, we define

$$\begin{aligned} d(p_1, p_2) &:= d_H(p_1, p_2, \dots, p_2) + d_H(p_2, p_1, \dots, p_1), \text{ and} \\ \text{opt-cost}(e) &:= d(p_1, p_2) + |\text{atime}(v_1) - \text{atime}(v_2)|. \end{aligned}$$

The following lemma follows from Theorem 1.

Lemma 1. *It holds that, for any $F = (v_1, \dots, v_k) \in \mathcal{E}$,*

$$\frac{1}{\gamma k^2} \cdot \sum_{i=1}^{k-1} \sum_{j=i+1}^k \text{opt-cost}(v_i, v_j) \leq \text{opt-cost}(F) \leq \sum_{i=1}^k \text{opt-cost}(v, v_i), \quad (4)$$

where $v = \arg \max_{u \in F} \text{atime}(u)$.

For any perfect k -way matching \mathcal{M} , define an edge subset M such that $e \in M$ if and only if the pair e is contained in some set F of \mathcal{M} . Thus we represent each set in \mathcal{M} with a complete graph of k vertices. We will show below that the characteristic vector for M is feasible to (\mathcal{P}') . Here, for a subset $X \subseteq E$, the characteristic vector $\mathbb{1}_X \in \{0, 1\}^E$ is defined to be

$$\mathbb{1}_X(x) = \begin{cases} 1 & (x \in X) \\ 0 & (x \notin X) \end{cases}.$$

Moreover, this implies that the optimal value of (\mathcal{P}') , denoted by $\mathcal{P}'(\sigma)$, is a lower bound of $\text{OPT}(\sigma)$ for any instance σ .

Lemma 2. *Let \mathcal{M} be a perfect k -way matching. Define an edge subset $M = \{(u, v) \in E \mid \exists F \in \mathcal{M} \text{ s.t. } u, v \in F\}$. Then $x = \mathbb{1}_M$ is a feasible solution to \mathcal{P}' . Furthermore, $\mathcal{P}'(\sigma) \leq \text{OPT}(\sigma)$ holds.*

The dual linear programming problem of (\mathcal{P}') is

$$(\mathcal{D}') \quad \left\{ \begin{array}{ll} \max. & \sum_{S \subseteq V} \text{sur}(S) \cdot (k - \text{sur}(S)) \cdot y_S \\ \text{s.t.} & \sum_{S: e \in \delta(S)} y_S \leq \frac{1}{\gamma k^2} \cdot \text{opt-cost}(e), \quad \forall e \in E \\ & y_S \geq 0, \quad \forall S \subseteq V \end{array} \right. \quad (5)$$

The weak duality of LP implies that $\mathcal{D}'(\sigma) \leq \mathcal{P}'(\sigma)$, where $\mathcal{D}'(\sigma)$ is the dual optimal value.

3.2 Greedy Dual for k -MPMD (GD- k)

We present our proposed algorithm, called *Greedy Dual for k -MPMD* (GD- k). The proposed algorithm extends the one by Bienkowski et al. [9] for 2-MPMD using the LP (\mathcal{P}') .

In the algorithm GD- k , we maintain a family of subsets of requests, called *active sets*. At any time, any request v arrived so far belongs to exactly one active set, denoted by $A(v)$. We also maintain a k -way matching \mathcal{M} . A request not in $\bigcup_{F \in \mathcal{M}} F$ is called *free*, and, for a subset $S \subseteq V$ of requests, $\text{free}(S)$ is the set of free requests in S .

When request v arrives, we initialize $A(v) = \{v\}$ and $y_S = 0$ for any subset $S \subseteq V$ such that $v \in S$. At any time, for an active set S such that $\text{free}(S)$ is nonempty, we increase y_S with rate r , where r is set to be $1/(\gamma k^2)$. Then, at some point, there exists an edge $e = (u, v) \in E$ such that $\sum_{S: e \in \delta(S)} y_S = \frac{1}{\gamma k^2} \cdot \text{opt-cost}(e)$, which we call a *tight* edge. When it happens, we merge the active sets $A(u)$ and $A(v)$ to a large subset $S = A(u) \cup A(v)$, that is, we update

$A(w) = S$ for all $w \in S$. We also mark the tight edge e . If $|\text{free}(S)| \geq k$, we partition $\text{free}(S)$ arbitrarily into subsets of size k with $\text{sur}(S)$ free requests, and add these size- k subsets to \mathcal{M} .

The pseudo-code of the algorithm is given as in Algorithm 1.

Let T be the time when all requests are matched in the algorithm. For any subset S , we denote the value of y_S at time τ in the algorithm by $y_S(\tau)$.

Algorithm 1 Greedy Dual for k -MPMD

```

1: procedure GD- $k(\sigma)$ 
2:    $\mathcal{M} \leftarrow \emptyset$ 
3:   for all moments  $t$  do
4:     if a request  $v$  arrives then
5:        $A(v) \leftarrow \{v\}$ 
6:       for all subsets  $S \ni v$  do
7:          $y_S \leftarrow 0$ 
8:       end for
9:       modify constraints of  $(\mathcal{D}')$ .
10:    end if
11:    if there exists  $e = (u, v) \in E$  such that  $\sum_{S:e \in \delta(S)} y_S = \frac{1}{\gamma k^2} \cdot \text{opt-cost}(e)$  and
     $A(u) \neq A(v)$  then
12:       $S \leftarrow A(u) \sqcup A(v)$ 
13:      for all  $v \in S$  do
14:         $A(v) \leftarrow S$ 
15:      end for
16:      mark  $e$ 
17:      while  $|\text{free}(S)| \geq k$  do
18:        choose arbitrarily a set  $F$  of  $k$  requests from  $S$ 
19:         $\mathcal{M} \leftarrow \mathcal{M} \cup \{F\}$ 
20:      end while
21:    end if
22:    for all sets  $S$  which are active and  $\text{free}(S) \neq \emptyset$  do
23:      increase continuously  $y_S$  at the rate of  $r$  per unit time
24:    end for
25:  end for
26: end procedure

```

We show that y_S 's maintained in Algorithm 1 are always dual feasible.

Lemma 3. *For any request v , it holds that*

$$\sum_{S:v \in S} y_S(\tau) \leq r \cdot (\tau - \text{atime}(v)) \quad (6)$$

at any time $\tau \geq \text{atime}(v)$. This holds with equality while v is not matched.

Lemma 4. *Let $r = \frac{1}{\gamma k^2}$. Then, at any time τ , $y_S(\tau)$ maintained in Algorithm 1 is a feasible solution to (\mathcal{D}') .*

3.3 Competitive Ratio of GD- k

To bound the competitive ratio of GD- k , we evaluate the distance cost and the waiting cost separately. We will show that each cost is upper-bounded by the dual optimal value of $\mathcal{D}'(\sigma)$.

Waiting Cost. We can upper-bound the waiting cost of the output as follows.

Lemma 5. *Let $\mathcal{M} = \{M_1, \dots, M_p\}$ be a perfect k -way matching returned by Algorithm 1, and let τ_ℓ be the time when we match M_ℓ . Then it holds that*

$$\sum_{\ell=1}^p \sum_{i=1}^k (\tau_\ell - \text{atime}(v_{\ell,i})) = \frac{1}{r} \cdot \sum_{S \subseteq V} \text{sur}(S) \cdot y_S(T) \leq \frac{1}{r} \cdot \mathcal{D}'(\sigma),$$

where we denote $M_\ell = \{v_{\ell,1}, \dots, v_{\ell,k}\}$.

Distance Cost. We say that a set $S \subseteq V$ is *formerly-active* at time τ if S is not active at time τ , but has been active before time τ .

Lemma 6. *Let S be an active or formerly-active set at time τ . Then, marked edges both of whose endpoints are contained in S form a spanning tree in S .*

We now evaluate the distance cost.

Lemma 7. *Let $\mathcal{M} = \{M_1, \dots, M_p\}$ be a perfect k -way matching returned by Algorithm 1. Then it holds that*

$$\sum_{\ell=1}^p d(\text{pos}(v_{\ell,1}), \dots, \text{pos}(v_{\ell,k})) \leq 4\gamma mk \cdot \sum_S \text{sur}(S) \cdot (k - \text{sur}(S)) \cdot y_S(T) \leq 4\gamma mk \mathcal{D}'(\sigma),$$

where we denote $M_\ell = \{v_{\ell,1}, \dots, v_{\ell,k}\}$.

Competitive Ratio. Summarizing the above discussion, we obtain Theorem 2.

Theorem 2. *Let d_H be an H -metric with parameter γ . Setting $r = 1/(\gamma k^2)$, Greedy Dual for k -MPMD achieves a competitive ratio $(4mk + k^2)\gamma$ for k -MPMD.*

Proof. Let σ be an instance of k -MPMD. It follows from Lemmas 5 and 7 that the cost of the returned perfect k -way matching is upper-bounded by $(4mk + k^2)\gamma \cdot \mathcal{D}'(\sigma)$. By the weak duality and Lemma 4, we observe that $\mathcal{D}'(\sigma) \leq \mathcal{P}'(\sigma) \leq \text{OPT}(\sigma)$. Thus the theorem holds. \square

Finally, we consider applying our algorithm to the problem with specific H -metrics such as d_{\max} and d_{HC} given in Sect. 2.3. Since they have parameter $\gamma = 1$, it follows from Theorem 2 that GD- k achieves a competitive ratio $O(mk + k^2)$. In the case of d_{\max} , we can further improve the competitive ratio.

Theorem 3. *For the k -MPMD on a metric space (χ, d_{\max}) , GD- k achieves a competitive ratio $O(m + k^2)$.*

4 Lower Bound of GD- k for a Diameter on a Line

In this section, we show a lower bound on the competitive ratio for GD- k for the metric d_D . Recall that $d_D(p_1, \dots, p_k) = \max_{i,j \in \{1, \dots, k\}} |p_i - p_j|$ for $p_1, \dots, p_k \in \mathbb{R}$.

We define an instance $\sigma_l = (V, \text{pos}, \text{atime})$ where $V = \{u_1, u_2, \dots, u_m\}$ as follows. Suppose that the number m of requests is equal to $m = sk^2$ for some integer s . Let p_1, \dots, p_k be k points in \mathbb{R} such that $d(p_i, p_{i+1}) = 2$ for any $i = 1, 2, \dots, k - 1$.

For $i = 1, 2, \dots, sk$ and $j = 1, 2, \dots, k$, define $\text{atime}(u_{k(i-1)+j}) = t_i$ and $\text{pos}(u_{k(i-1)+j}) = p_j$ for $j = 1, 2, \dots, k$, where we define $t_1 = 0$ and $t_i = 1 + (2i - 3)\varepsilon$ for $i \geq 2$. Thus, at any time t_i ($i = 1, \dots, sk$), the k requests $u_{k(i-1)+1}, \dots, u_{k(i-1)+k}$ arrive at every point in p_1, \dots, p_k , respectively.

Then it holds that $\mathcal{OPT}(\sigma_l) \leq k + k\varepsilon + k^3\varepsilon + mk\varepsilon$, while the output of GD- k has cost at least $m + k + (m - k)\varepsilon$.

Theorem 4. *For a metric space (\mathbb{R}, d_D) , there exists an instance σ_l of m requests such that GD- k admits a competitive ratio $\Omega(\frac{m}{k})$.*

References

1. Aggarwal, G., Goel, G., Karande, C., Mehta, A.: Online vertex-weighted bipartite matching and single-bid budgeted allocations. In: Proceedings of the 2011 Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1253–1264. SODA 2011, SIAM (2011)
2. Antoniadis, A., Barcelo, N., Nugent, M., Pruhs, K., Scquizzato, M.: A $o(n)$ -competitive deterministic algorithm for online matching on a line. *Algorithmica* **81**, 2917–2933 (2019)
3. Ashlagi, I., Burq, M., Dutta, C., Jaillet, P., Saberi, A., Sholley, C.: Edge-weighted online windowed matching. *Math. Oper. Res.* **48**(2), 999–1016 (2023)
4. Assaf, S., Pal, K.: Partial n -metric spaces and fixed point theorems (2015). [arXiv: 1502.05320](https://arxiv.org/abs/1502.05320)
5. Azar, Y., Chiplunkar, A., Kaplan, H.: Polylogarithmic bounds on the competitiveness of min-cost perfect matching with delays. In: Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2017), pp. 1051–1061. SIAM (2017)
6. Azar, Y., Jacob Fanani, A.: Deterministic min-cost matching with delays. *Theor. Comput. Syst.* **64**(4), 572–592 (2020)
7. Azar, Y., Ren, R., Vainstein, D.: The min-cost matching with concave delays problem. In: Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, pp. 301–320. SODA 2021, SIAM (2021)
8. Bansal, N., Buchbinder, N., Gupta, A., Naor, J.S.: An $O(\log^2 k)$ -competitive algorithm for metric bipartite matching. In: Arge, L., Hoffmann, M., Welzl, E. (eds.) ESA 2007. LNCS, vol. 4698, pp. 522–533. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75520-3_47
9. Bienkowski, M., Kraska, A., Liu, H.-H., Schmidt, P.: A primal-dual online deterministic algorithm for matching with delays. In: Epstein, L., Erlebach, T. (eds.) WAOA 2018. LNCS, vol. 11312, pp. 51–68. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-04693-4_4

10. Deryckere, L., Umboh, S.W.: Online matching with set and concave delays (2023). [arXiv: 2211.02394](https://arxiv.org/abs/2211.02394)
11. Emek, Y., Kutten, S., Wattenhofer, R.: Online matching: haste makes waste! In: Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing, pp. 333–344. STOC 2016, ACM (2016)
12. Fuchs, B., Hochstättler, W., Kern, W.: Online matching on a line. Theoret. Comput. Sci. **332**(1), 251–264 (2005)
13. Goel, G., Mehta, A.: Online budgeted matching in random input models with applications to Adwords. In: Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 982–991. SODA 2008, SIAM (2008)
14. Gupta, A., Lewi, K.: The online metric matching problem for doubling metrics. In: Czumaj, A., Mehlhorn, K., Pitts, A., Wattenhofer, R. (eds.) ICALP 2012. LNCS, vol. 7391, pp. 424–435. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31594-7_36
15. Karp, R.M., Vazirani, U.V., Vazirani, V.V.: An optimal algorithm for on-line bipartite matching. In: Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing, pp. 352–358. STOC 1990, ACM (1990)
16. Khan, K.A.: On the possibility of N -topological spaces. Int. J. Math. Arch. **3**(7), 2520–2523 (2012)
17. Koutsoupias, E., Nanavati, A.: The online matching problem on a line. In: Solis-Oba, R., Jansen, K. (eds.) WAOA 2003. LNCS, vol. 2909, pp. 179–191. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24592-6_14
18. Liu, X., Pan, Z., Wang, Y., Wattenhofer, R.: Impatient online matching. In: 29th International Symposium on Algorithms and Computation, ISAAC 2018. LIPIcs, vol. 123, pp. 62:1–62:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2018)
19. Liu, X., Pan, Z., Wang, Y., Wattenhofer, R.: Online matching with convex delay costs (2022). [arXiv:2203.03335](https://arxiv.org/abs/2203.03335)
20. Mehta, A., Saberi, A., Vazirani, U., Vazirani, V.: AdWords and generalized on-line matching. In: 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), pp. 264–273 (2005)
21. Mehta, A.: Online matching and ad allocation. Found. Trends Theor. Comput. Sci. **8**(4), 265–368 (2013)
22. Melnyk, D., Wang, Y., Wattenhofer, R.: Online k -Way Matching with Delays and the H -Metric (2021). [arXiv:2109.06640](https://arxiv.org/abs/2109.06640)
23. Nayyar, K., Raghvendra, S.: An input sensitive online algorithm for the metric bipartite matching problem. In: 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS 2017), pp. 505–515 (2017)
24. Pavone, M., Saberi, A., Schiffer, M., Tsao, M.W.: Technical note-online hypergraph matching with delays. Oper. Res. **70**(4), 2194–2212 (2022)
25. Raghvendra, S.: A robust and optimal online algorithm for minimum metric bipartite matching. In: International Workshop on Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2016). LIPIcs, vol. 60, pp. 18:1–18:16. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2016)