

Deterministic Min-Cost Matching with Delays

Yossi Azar*

Amit Jacob-Fanani†

Abstract

We consider the online Minimum-Cost Perfect Matching with Delays (MPMD) problem introduced by Emek et al. (STOC 2016), in which a general metric space is given, and requests are submitted in different times in this space by an adversary. The goal is to match requests, while minimizing the sum of distances between matched pairs in addition to the time intervals passed from the moment each request appeared until it is matched.

In the online Minimum-Cost Bipartite Perfect Matching with Delays (MBPMD) problem introduced by Ashlagi et al. (APPROX/RANDOM 2017), each request is also associated with one of two classes, and requests can only be matched with requests of the other class.

Previous algorithms for the problems mentioned above, include randomized $O(\log n)$ -competitive algorithms for known and finite metric spaces, n being the size of the metric space, and a deterministic $O(m)$ -competitive algorithm, m being the number of requests.

We introduce $O\left(m^{\log(\frac{3}{2}+\epsilon)}\right)$ -competitive deterministic algorithms for both problems and for any fixed $\epsilon > 0$. In particular, for a small enough ϵ the competitive ratio becomes $O(m^{0.59})$. These are the first deterministic algorithms for the mentioned online matching problems, achieving a sub-linear competitive ratio. Our algorithms do not need to know the metric space in advance.

1 Introduction

In the algorithmic graph theory, a *Perfect Matching* is a subset of graph edges, in which each vertex of the graph is incident on exactly one edge of the subset, and the weight of the matching is the sum of the weights of the edges of the subset. In the well known *Minimum-Cost Perfect Matching* problem a weighted graph is given, and a *Perfect Matching* of minimum weight is to be found. The *Blossom Algorithm* due to Edmonds [9] is the first algorithm to solve this problem in polynomial time.

Many versions of the *Minimum-Cost Perfect Matching* problem have been studied over the last few decades, some of the noticeable variants are online versions of the problem (e.g. *Minimum-Cost Perfect Matchings with Online Vertex Arrival* due to Kalyanasundaram and Pruhs [14]).

In this paper we suggest a deterministic algorithm for the *Minimum-Cost Perfect Matching with Delays* (MPMD) variant, which was introduced by Emek et al. [10], and a similar deterministic algorithm for another variation of the problem - the *Minimum-Cost Bipartite Perfect Matching with Delays* (MBPMD) problem, which was introduced by Ashlagi et al. [2].

To illustrate the MPMD problem, imagine players logging in through a server to an online game at different times, unknown a priori to the server they have connected through. The server then needs to match between the players while maximizing their satisfaction from playing the game. Players feel satisfied when they play against players at a level similar to their own. Therefore, when

*School of Computer Science, Tel Aviv University, Israel; azar@tau.ac.il

†School of Computer Science, Tel Aviv University, Israel; amitj@mail.tau.ac.il

pairing players, the server needs to consider the difference in levels between the players, called the *connection cost*.

Once logged in, a player doesn't necessarily start playing instantly, as the server can postpone the decision regarding with whom to match the player, until a good match is found (i.e. another player at a similar level logs in to the game). This is a poor strategy since players are unhappy when forced to wait too long until they start playing. The time a player has to wait until the game starts is called the *delay cost*.

More formally, an adversary presents requests at points in a general metric space, in an online manner. The goal is to produce a minimum-cost perfect matching when the cost of an edge is the sum of its *connection cost* (the distance between the two points in the metric space) and the *delay cost* of the two requests matched by the edge. All requests have to be matched by the server after a finite time from the moment they have arrived.

The MBPMD problem is an extension of the MPMD problem (due to Ashlagi et al. [2]), in which each of the requests may take one of two colors, and each edge of the matching, must be incident on one request from each color. The MBPMD problem has many applications, such as matching drivers to passengers (Uber, Lyft), job finding platforms, etc.

Background. The standard method used to measure an online algorithm's performance is its competitive ratio. We use this method when comparing the performance of matching algorithms for both MPMD and MBPMD. An algorithm is α -competitive if the maximum ratio between the cost of the algorithm to the cost of the optimum solution, over all inputs, is bounded by α .

The first algorithm for MPMD was developed by Emek et al. [10] with an expected competitive ratio $O(\log^2 n + \log \Delta)$ on a finite metric space of size n , where Δ is the aspect-ratio of the metric space (the ratio of the maximum distance to the minimum distance between any two points in the metric space). Azar et al. [3] improved the competitive ratio to $O(\log n)$, and showed a lower bound of $\Omega(\sqrt{\log n})$ (both deterministic and randomized). Ashlagi et al. [2] improved this lower bound to $\Omega\left(\frac{\log n}{\log \log n}\right)$ (both deterministic and randomized). They also gave an $O(\log n)$ -competitive randomized algorithm for MBPMD.

All mentioned above algorithms are randomized (on a general finite metric). In online algorithms where one cannot repeat the algorithm in case the cost is high, a deterministic algorithm is preferable. Bienkowski et al. [7] provided the first deterministic algorithm for MPMD on general metrics, with a competitive-ratio of $O(m^{2.46})$, m being the number of requests. While the previous algorithms require the metric space to be known a priori, their algorithm does not, and is also applicable when the metric space is revealed in an online manner. Bienkowski et al. also noted that the algorithm of [3] can be used to provide an $O(n)$ -competitive deterministic algorithm for a general known metric space. Recently, Bienkowski et al. [6] provided a new primal-dual deterministic algorithm for MPMD on general metrics, with a competitive-ratio of $O(m)$, m being the number of requests.

Prior to our result there was no deterministic sub-linear competitive algorithm, neither in n nor in m .

Our Contribution. In this paper we introduce deterministic algorithms for both versions of the problem, both with a competitive ratio $O\left(\frac{1}{\epsilon} m^{\log(\frac{3}{2} + \epsilon)}\right)$. When the constant ϵ is small enough, this becomes $O(m^{0.59})$. Our algorithms do not need to know the metric space in advance.

We present a simple algorithm, which is an adaptation of the greedy algorithm for the *Minimum-Cost Perfect Matching* problem by Reingold and Tarjan [21] to an online environment. In our

algorithm, requests grow hemispheres around them in a metric that is the Cartesian product of the original metric and the time axis (also called the *time-augmented metric space*). The hemispheres radii grow slowly in the negative direction of the time axis. Once a request is found on the boundary of another request's hemisphere, they are matched by the algorithm. Our analysis is inspired by the analysis of the original greedy algorithm by Reingold and Tarjan.

In the bipartite case, the algorithm is essentially the same, but requests are matched only if they are of different colors.

Related Work. First we consider related work **with delays**. Since Emek et al. [10] introduced the notion of online problems with delayed service, there has been a growing number of works studying such problems (e.g. *Online Service with Delays* [4], *Minimum-Cost Bipartite Perfect Matching with Delays* [2], *Minimum-Cost Perfect Matching with Delays for Two Sources* [11]). Works dealing with the *Minimum-Cost Perfect Matching with Delays* and *Minimum-Cost Bipartite Perfect Matching with Delays* problems, such as the papers by Emek et al. [10], Azar et al. [3], Ashlagi et al. [2] and Bienkowski et al. [7], are the most closely related to this work. As mentioned above, Emek et al. [10] provided a randomized $O(\log^2 n + \log \Delta)$ -competitive algorithm for MPMD on general metrics, in which n is the size of the metric space and Δ is the aspect ratio. They consider the randomized embeddings of the general metric space into a distribution over metrics given by hierarchically separated full binary trees, with distortion $O(\log n)$, and give a randomized algorithm for the hierarchically separated trees metrics.

Subsequently, Azar et al. [3] provided a randomized $O(\log n)$ -competitive algorithm for the same problem, thus improving the original upper bound. They used randomized embedding of the general metric space into a distribution over metrics given by hierarchically separated trees of height $O(\log n)$, with distortion $O(\log n)$. Then they give a deterministic $O(1)$ -space-competitive (that is the competitive ratio associated with the *connection cost*) and $O(h)$ -time-competitive (that is the competitive ratio associated with the *delay cost*) algorithm over tree metrics, where h is the height of the tree. This yields a competitive ratio of $O(\log n)$. Moreover, they provided a randomized $\Omega(\sqrt{\log n})$ lower bound, confirming a conjecture made by Emek et al. [10] that the competitive ratio of any online algorithm for the problem must depend on n .

Ashlagi et al. [2] improved the lower bound on the competitive ratio to $\Omega\left(\frac{\log n}{\log \log n}\right)$, almost matching the upper bound of Azar et al. of $O(\log n)$. The rest of the paper focuses on the bipartite version of the problem, providing an $O(\log n)$ -competitive ratio by the adaptation of the algorithm of Azar et al. [3] to the bipartite case.

In order to provide a *deterministic* algorithm, Bienkowski et al. [7] used a different approach for the problem - they used a semi-greedy scheme of a ball-growing algorithm. In their analysis, they fix an optimal matching, and charge the cost of each matching-edge generated by their algorithm against the cost of an existing matching-edge of the optimal matching. As mentioned above, their algorithm achieves a competitive ratio of $O(m^{2.46})$, where m is the number of requests.

Bienkowski et al. improved this result in [6] by providing a new $O(m)$ -competitive LP-based algorithm. Briefly, their algorithm maintains a primal relaxation of the matching problem and its dual (the programs evolve in time as more requests arrive). Dual variables are increased along time, until a dual constraint (corresponding to a pair of requests) becomes tight, which results in the algorithm connecting the pair. They also proved that their analysis is tight (the competitive-ratio of their algorithm is $\Omega(m)$). Recall that our algorithm achieves a sub-linear competitive-ratio (in m).

Next we consider related work **without delays**. The *Online Minimum Weighted Bipartite Matching* (OMM) problem due to [14, 16] is another important online version of the *Minimum-*

Cost Perfect Matching problem, in which k vertices are given a priori, and k additional vertices are revealed at different times, together with the distances from the first k vertices. The algorithm then needs to match the later k vertices to the first k vertices, while trying to minimize the total weight of the produced matching. In this version, delay of the algorithm's decision is not available. Kalyanasundaram and Pruhs [14] and Khuller et al. [16] showed independently a tight upper and lower bounds of $2k - 1$ on the deterministic competitive ratio of the problem.

The first sub-linear competitive randomized algorithm for the problem, was given by Meyerson et al. [18] using randomized embeddings into trees, with a competitive ratio of $O(\log^3 k)$. Consequently, Bansal et al. [5] improved this upper bound by providing a $O(\log^2 k)$ -competitive randomized algorithm. In addition, they showed an $\Omega(\log k)$ lower bound on the competitive ratio for randomized algorithms.

The special case of line-metrics is argued to be the most interesting instance of OMM (e.g. [17]). Kalyanasundaram and Pruhs conjectured in 1998 [15] that there exists a 9-competitive deterministic algorithm for OMM on line-metrics, but in 2003 Fuchs et al. [12] disproved the conjecture, proving a lower bound of 9.001 for deterministic algorithms. This is the best known lower bound thus far.

Antoniadis et al. [1] presented the first sub-linear deterministic algorithm for line-metrics, with a competitive ratio of $O\left(\frac{1}{\epsilon} k^{\log(\frac{3}{2} + \epsilon)}\right)$. Recently, Nayyar and Raghvendra [19] improved this upper bound to $O(\log^2 k)$ by careful analysis of the deterministic algorithm present in [20]. Gupta and Lewi [13] provided a randomized $O(\log k)$ -competitive algorithm for doubling metrics, hence for line-metrics as well.

To summarize, the best known deterministic upper bound on the competitive ratio for line-metrics is $O(\log^2 k)$, and best known lower bound is 9.001. For randomized algorithms the best known upper bound is $O(\log k)$.

Paper Organization. We describe the algorithm for *Minimum-Cost Perfect Matching with Delays* in Section 3 and analyze its performance in Section 3.1. Through an example in Appendix A we show that our analysis is tight, and prove that the competitive ratio of our algorithm indeed depends on the number of requests, and not on the size of the metric space. In addition, we show in Appendix B that minor natural changes to the algorithm, do not transform the competitive ratio into a function of the size of the metric space (in the case of a finite metric space) instead of the number of requests. In Section 4 we present the algorithm for *Minimum-Cost Bipartite Perfect Matching with Delays* and analyze its performance.

2 Preliminaries

A *metric space* $\mathcal{M} = (S, d)$ is a set S and a distance function $d : S \times S \rightarrow \mathbb{R}^+$ that meets the following conditions: non-negativity, symmetry, the triangle-inequality, and that $d(x, y) = 0$ if and only if $x = y$. When S is finite, we refer to \mathcal{M} as a *finite metric space*, and an *infinite metric space* otherwise.

2.1 Model

In the online *Minimum-Cost Perfect Matching with Delays* problem on a metric space $\mathcal{M} = (S, d)$ (known a priori to the algorithm), an input instance $\mathcal{I} = \langle r_i \rangle_{i=1}^m$ is presented to the algorithm in an online fashion, so that each request r_i is revealed to the algorithm at time $t(r_i)$ at the location $x(r_i) \in S$. The number of requests m is even and unknown a priori to the algorithm.

The online algorithm should produce a perfect matching in real time. Formally, two requests p, q can be matched by the algorithm at any time $t \geq \max(t(p), t(q))$, if they have not been matched yet by the algorithm.

Let $\langle p_i, q_i, t_i \rangle_{i=1}^{\frac{m}{2}}$ be the set of pairs of requests matched by the algorithm, and their matching times (p_i and q_i were matched by the algorithm at t_i), then the cost of the matching produced by the algorithm is

$$\sum_{i=1}^{\frac{m}{2}} d(x(p_i), x(q_i)) + |t_i - t(p_i)| + |t_i - t(q_i)|$$

In other words, the cost is the sum of the connection cost of all matched pairs in addition to the sum of the delay cost of all requests. The goal of the algorithm is to minimize this cost.

The *Minimum-Cost Bipartite Perfect Matching with Delays* is virtually the same problem as the *Minimum-Cost Perfect Matching with Delays* problem, except that each request r_i is associated with one of two classes, so that each request r_i can be matched to a request r_j if and only if $\text{class}(r_i) \neq \text{class}(r_j)$.

2.2 The time-augmented metric space

Given a metric space $\mathcal{M} = (S, d)$ define the *time-augmented metric space* as $\mathcal{M}_T = (S \times \mathbb{R}, D)$ where D is a distance function defined as

$$D((l_1, t_1), (l_2, t_2)) = d(l_1, l_2) + |t_1 - t_2|$$

assuming $(l_1, t_1), (l_2, t_2) \in S \times \mathbb{R}$. That is, the time axis was added as another dimension in the metric space. One can easily verify that D indeed defines a metric.

The following Lemma shows that for offline algorithms, solving the Minimum-Cost Perfect Matching **with Delays** problem in the metric space \mathcal{M} is equivalent to solving the Minimum-Cost Perfect Matching problem in \mathcal{M}_T .

Lemma 1. Assume $\mathcal{I} = \langle r_i \rangle_{i=1}^m$ is an instance of MPMD then OPT can be computed as the weight of an optimal solution for the Minimum Metric Perfect Matching problem on the instance \mathcal{I} as points in the time-augmented metric space \mathcal{M}_T .

Proof. Let OPT^* be an optimal solution for *Minimum Metric Perfect Matching* over the instance \mathcal{I} . We show that $\text{OPT} = \text{OPT}^*$.

Let A be the solution for *Minimum Metric Perfect Matching* over the instance \mathcal{I} , which matches the pairs corresponding to those matched by OPT . The cost of A is at most the cost of OPT , since for a given pair (u, v) matched by OPT at time $t_{uv} \geq \max(t(u), t(v))$, OPT would pay $t_{uv} - t(u) + t_{uv} - t(v) + d(x(u), x(v))$, while A would pay $D(u, v) = |t(u) - t(v)| + d(x(u), x(v))$ which cannot be larger. Therefore $\text{OPT}^* \leq A \leq \text{OPT}$.

For the other direction we define an online algorithm B which matches the pairs corresponding to those matched by OPT^* , as soon as the two end-points arrive. For a given pair of requests (p, q) matched by B , it pays

$$\max(t(p), t(q)) - t(p) + \max(t(p), t(q)) - t(q) + d(x(p), x(q)) = |t(p) - t(q)| + d(x(p), x(q))$$

Therefore the cost paid by B is the same as the cost paid by OPT^* .

Hence $\text{OPT} \leq B = \text{OPT}^*$. □

3 A Deterministic Algorithm for MPMD on General Metrics

Our algorithm ($\text{ALG}(\epsilon)$) is parametrized with a constant $\epsilon \in \mathbb{R}$. Upon the arrival of a request $p \in S \times \mathbb{R}$, the algorithm begins to grow a hemisphere surrounding p in the negative direction of the time axis, such that the radius growth rate is ϵ . Therefore, at time t , a request $q \in S \times \mathbb{R}$ is on the hemisphere's boundary if and only if $\epsilon(t - t(p)) = D(p, q)$ and $t(q) \leq t(p)$, where D is the distance function defined by the time-augmented metric space \mathcal{M}_T . The algorithm matches a request q to a request p as soon as q is found on the boundary of p 's hemisphere.

Note that the algorithm does not need to know the metric space in advance, but it only requires that together with any arriving request p , it learns the distances from p to all previous requests.

Algorithm 1 A Deterministic Algorithm for MPMD on General Metrics

```

1: procedure  $\text{ALG}(\epsilon)$ 
2:   At every moment  $t$ :
3:     Add the new requests that arrive at time  $t$ 
4:     for each unmatched request  $p$  do
5:       for each unmatched request  $q \neq p$  do
6:         if  $t(p) \geq t(q)$  and  $t = t(p) + \frac{D(p,q)}{\epsilon}$  then
7:           match $(p, q)$ 
8:         end if
9:       end for
10:    end for
11: end procedure

```

The algorithm is described as a continuous process but can be easily discretized using priority queues over anticipated matching events for each pair.

The algorithm breaks ties arbitrarily (i.e. a request that is on multiple hemispheres at the same time, or multiple requests that are on the same hemisphere). Note that for the analysis of the algorithm we may assume that there are no ties, as an adversary might slightly perturb the points so that the algorithm would choose the worse option.

3.1 Analysis

Theorem 1. $\text{ALG}(\epsilon)$ is $O\left(\frac{1}{\epsilon} m^{\log(\frac{3+\epsilon}{2})}\right)$ -competitive.

Given $\epsilon \in \mathbb{R}$ we run $\text{ALG}(\epsilon)$ over the instance $\mathcal{I} = \langle r_i \rangle_{i=1}^m$, that is with a hemisphere growth rate of ϵ . For the analysis, we denote ALG_{ON} to be the cost paid by $\text{ALG}(\epsilon)$, and ALG_{OFF} to be the weight of the matching produced by $\text{ALG}(\epsilon)$, when viewing \mathcal{I} as points in the time-augmented metric space \mathcal{M}_T . OPT is the cost of an optimal solution for MPMD over the instance \mathcal{I} .

Consider the last two pairs of requests to be matched by ALG . They consist of four requests, name them a, b, c, d , such that (a, b) is one pair, and (c, d) is the second pair. Assume w.l.o.g that (a, b) were matched at time t_{ab} , and (c, d) at $t_{cd} \geq t_{ab}$. Also, assume w.l.o.g that $t(a) \leq t(b)$.

Lemma 2.

- (1) $D(a, b) \leq (1 + \epsilon)D(a, c)$ and $D(a, b) \leq (1 + \epsilon)D(a, d)$
- (2) $D(a, b) \leq (1 + \epsilon)D(b, c)$ and $D(a, b) \leq (1 + \epsilon)D(b, d)$

Proof. We only prove $D(a, b) \leq (1+\epsilon)D(a, c)$ and $D(a, b) \leq (1+\epsilon)D(b, c)$ since there is no difference between c and d .

To prove (1), we look at two cases, that are $t(c) \geq t(a)$, and $t(c) < t(a)$.

Case $t(c) \geq t(a)$: Upon the arrival of c and b , the algorithm begins to grow hemispheres surrounding them, and in particular a might be on their boundaries. Since (a, b) was the first pair to be matched, a was on b 's hemisphere before it was on c 's hemisphere (otherwise (a, c) should have been matched first). Therefore $t(b) + \frac{D(a, b)}{\epsilon} \leq t(c) + \frac{D(a, c)}{\epsilon}$, and we conclude

$$D(a, b) \leq D(a, c) + \epsilon(t(c) - t(b)) \leq D(a, c) + \epsilon(t(c) - t(a)) \leq (1 + \epsilon)D(a, c)$$

Case $t(c) < t(a)$: Upon the arrival of a and b , the algorithm begins to grow hemispheres surrounding them. In particular, a might be on the boundary of b 's hemisphere, and c might be on the boundary of a 's hemisphere. Since (a, b) was the first pair to be matched, a was on b 's hemisphere before c was on a 's hemisphere (otherwise (a, c) should have been matched first). Therefore $t(b) + \frac{D(a, b)}{\epsilon} \leq t(a) + \frac{D(a, c)}{\epsilon}$. Thus, we conclude that

$$D(a, b) \leq D(a, c) + \epsilon(t(a) - t(b)) = D(a, c) - \epsilon(t(b) - t(a)) \leq D(a, c) \leq (1 + \epsilon)D(a, c)$$

To prove (2), we look at the two cases $t(c) \geq t(b)$, and $t(c) < t(b)$.

Case $t(c) \geq t(b)$: Upon the arrival of c and b , the algorithm begins to grow hemispheres surrounding them. In particular, a might be on the boundary of b 's hemisphere, and b might be on the boundary of c 's hemisphere. Since (a, b) was the first pair to be matched, a was on b 's hemisphere before b was on c 's hemisphere (otherwise (b, c) should have been matched first). Therefore $t(b) + \frac{D(a, b)}{\epsilon} \leq t(c) + \frac{D(b, c)}{\epsilon}$. Thus, we conclude that

$$D(a, b) \leq D(b, c) + \epsilon(t(c) - t(b)) \leq D(b, c) + \epsilon D(b, c) = (1 + \epsilon)D(b, c)$$

Case $t(c) < t(b)$: Upon b 's arrival, the algorithm begins to grow a hemisphere surrounding it, and in particular a and c might be on its boundary. Since (a, b) was the first pair to be matched, a was on b 's hemisphere before c was (otherwise (b, c) should have been matched first). Therefore $t(b) + \frac{D(a, b)}{\epsilon} \leq t(b) + \frac{D(b, c)}{\epsilon}$. Thus, we conclude that

$$D(a, b) \leq D(b, c) \leq (1 + \epsilon)D(b, c)$$

□

We use the following well known observation.

Observation 1. *The union of any two matchings is a set of vertex-disjoint cycles. In every such cycle, the edges alternate between the two matchings. Note that two parallel edges are considered a cycle.*

Let $\mathcal{C} = \{C_1, \dots, C_k\}$ be the set of cycles (vertices and edges) generated from taking the union of the matchings produced by ALG and OPT. Define $l_1, \dots, l_k \in \mathbb{R}$ such that l_i is the total length of edges of ALG in C_i . Define similarly $l_1^*, \dots, l_k^* \in \mathbb{R}$ for edges of OPT.

Lemma 3. $\frac{\text{ALG}_{\text{OFF}}}{\text{OPT}} \leq \max_i \frac{l_i}{l_i^*}$

Proof.

$$\frac{\text{ALG}_{\text{OFF}}}{\text{OPT}} = \frac{\sum_{i=1}^k l_i}{\sum_{i=1}^k l_i^*} = \sum_{j=1}^k \frac{l_j^*}{\sum_{i=1}^k l_i^*} \frac{l_j}{l_j^*} \leq \sum_{j=1}^k \frac{l_j^*}{\sum_{i=1}^k l_i^*} \max_r \frac{l_r}{l_r^*} = \max_r \frac{l_r}{l_r^*} \sum_{j=1}^k \frac{l_j^*}{\sum_{i=1}^k l_i^*} = \max_r \frac{l_r}{l_r^*}$$

□

Lemma 4. Denote \hat{l}_i^* the cost paid by an optimal algorithm for *Minimum Metric Perfect Matching* on the instance constructed from the vertices of C_i , and \hat{l}_i the cost of running ALG over the vertices of C_i . Then $\hat{l}_i^* = l_i^*$ and $\hat{l}_i = l_i$.

Proof. To prove $\hat{l}_i^* = l_i^*$ assume by contradiction that $l_i^* < \hat{l}_i^*$. Notice that the subset of edges of OPT contained in C_i is a legal solution for *Minimum Metric Perfect Matching* with cost l_i^* . Clearly l_i^* is less than \hat{l}_i^* , contradicting the definition of \hat{l}_i^* . For the other direction, let E be the edges matched by OPT, and \hat{E} be the edges matched by an optimal algorithm for *Minimum Metric Perfect Matching* on the instance constructed from the vertices of C_i . Define $\bar{E} = (E \setminus C_i) \cup \hat{E}$. Notice that \bar{E} is a legal solution for *Minimum Metric Perfect Matching* on the instance \mathcal{I} with cost $\sum_{i=1}^k l_i^* - l_i^* + \hat{l}_i^* < \text{OPT}$ contradicting the definition of OPT. Therefore $l_i^* = \hat{l}_i^*$.

To prove $\hat{l}_i = l_i$ we show that K - the matching produced by ALG when running over the vertices of C_i , is the same as E_i - the subset of edges matched by ALG and contained in C_i , when running on the instance \mathcal{I} . Let $r = \frac{|C_i|}{2}$ where $|C_i|$ is the number of edges in C_i , and note that $|E_i| = r = |K|$, since both E_i and K are matchings over C_i . Sort the edges of E_i by the time they are formed from first to last: $e_1 = (u_1, v_1), \dots, e_r = (u_r, v_r)$, and the same for the edges of K : $k_1 = (p_1, q_1), \dots, k_r = (p_r, q_r)$.

Assume by contradiction that $E_i \neq K$, and let j be the lowest index with $e_j \neq k_j$. Let t_e be the time that e_j was formed and t_k be the time that k_j was formed. At $\min(t_e, t_k)$, just before e_j and k_j were formed, E_i and K contained the same set of edges. Therefore the points that were not matched by ALG until $\min(t_e, t_k)$, are the same in the two cases, and obviously the radii of the hemispheres at $\min(t_e, t_k)$ are the same in both cases as well. Thus, if v_j and u_j still exist in ALG's run on \mathcal{I} at that time, and v_j is on u_j 's hemisphere, then at the same time both v_j and u_j exist in ALG's run on C_i , and v_j is on u_j 's hemisphere. Thus ALG would match the pair (u_j, v_j) when running on C_i at $t_e = t_k$, concluding $e_j = k_j$ and contradicting the assumption. \square

Corollary 1. By virtue of Lemma 3 and Lemma 4 it suffices to consider $\frac{\text{ALG}_{\text{OFF}}}{\text{OPT}}$ when the union of the matchings produced by ALG and OPT forms a single cycle.

Lemma 5. Let $\gamma \in \mathbb{R}$ s.t. $\gamma > 2$ and let $f : \mathbb{N} \rightarrow \mathbb{R}$ satisfy the recurrence relation

$$f(2k) = \min_{1 \leq i \leq k-1} \left\{ f(2i), \frac{1}{\gamma} (f(2i) + f(2k - 2i)) \right\}, \quad f(2) = 1$$

Then,

$$f(n) = \Omega\left(\frac{1}{n^{\log(\frac{\gamma}{2})}}\right)$$

Proof. We prove by induction on k that $f(2k) \geq \left(\frac{2}{\gamma}\right)^{\log k}$.

Base Case ($k = 1$): $f(2) = 1$, and $\left(\frac{2}{\gamma}\right)^{\log 1} = \left(\frac{2}{\gamma}\right)^0 = 1$.

Inductive step: Assume the claim holds for all $j < k$.

By the induction hypothesis for every $j < k$ it holds that $f(2j) \geq \left(\frac{2}{\gamma}\right)^{\log j} > \left(\frac{2}{\gamma}\right)^{\log k}$. Therefore, from the definition of f

$$f(2k) \geq \min \left(\left(\frac{2}{\gamma}\right)^{\log k}, \frac{1}{\gamma} (f(2) + f(2k - 2)), \frac{1}{\gamma} (f(4) + f(2k - 4)), \dots \right)$$

Define $h(j) = \frac{1}{\gamma} (f(2j) + f(2k - 2j))$, so

$$f(2k) \geq \min \left(\left(\frac{2}{\gamma} \right)^{\log k}, \min_{1 \leq j \leq k-1} \{h(j)\} \right)$$

By the induction hypothesis,

$$h(j) \geq \frac{1}{\gamma} \left(\left(\frac{2}{\gamma} \right)^{\log j} + \left(\frac{2}{\gamma} \right)^{\log k-j} \right) \geq \min_{x \in \mathbb{R}} \frac{1}{\gamma} \left\{ \left(\frac{2}{\gamma} \right)^{\log x} + \left(\frac{2}{\gamma} \right)^{\log k-x} \right\}$$

$\left(\frac{2}{\gamma} \right)^{\log x} + \left(\frac{2}{\gamma} \right)^{\log k-x}$ is symmetric about $x = \frac{k}{2}$. Moreover, it is a concave function as it is the sum of two concave functions, thus the minimum point occurs at $x = \frac{k}{2}$.

We found that $h(j) \geq \frac{1}{\gamma} \left(\left(\frac{2}{\gamma} \right)^{\log \frac{k}{2}} + \left(\frac{2}{\gamma} \right)^{\log \frac{k}{2}} \right) = \left(\frac{2}{\gamma} \right)^{\log \frac{k}{2}+1} = \left(\frac{2}{\gamma} \right)^{\log k}$

Hence, we conclude

$$f(2k) \geq \min \left(\left(\frac{2}{\gamma} \right)^{\log k}, \left(\frac{2}{\gamma} \right)^{\log k} \right) = \left(\frac{2}{\gamma} \right)^{\log k} = \frac{1}{k^{\log \frac{\gamma}{2}}}$$

□

Lemma 6. $\text{ALG}_{\text{OFF}} \leq O \left(m^{\log(\frac{3+\epsilon}{2})} \right) \text{OPT}$

Proof. We view the requests as if they were in the time-augmented metric space \mathcal{M}_T , and analyze the performance of ALG in an offline manner. By Corollary 1 we analyze the performance of ALG when $G = (\mathcal{I}, E)$, the union of the matchings produced by ALG and OPT, forms a single cycle.

Denote E_O the subset of edges matched by OPT, and E_A the subset of edges matched by ALG. Consider again the last two pairs of requests to be matched by ALG, that is (a, b) and (c, d) , and assume that $t_{ab} \leq t_{cd}$ and $t(b) \geq t(a)$ (t_{ab} is the time that ALG matched (a, b) , and t_{cd} is the time that ALG matched (c, d)). Denote $T = \sum_{e \in E \setminus \{(c, d)\}} D(e)$, and let $O = \sum_{e \in E_O} D(e)$. From the triangle inequality we have that $D(c, d)$ is smaller than T , therefore

$$\frac{\text{ALG}_{\text{OFF}}}{\text{OPT}} = \frac{D(c, d) + T - O}{O} \leq \frac{2T - O}{O} = 2\frac{T}{O} - 1 \quad (1)$$

We will bound $\frac{T}{O}$ from below, by developing and solving a recurrence relation similar to the one developed in [21], thus giving an upper bound on $\frac{\text{ALG}_{\text{OFF}}}{\text{OPT}}$.

Scale the distances so that $T = 1$. Of course, $\frac{T}{O}$ stays the same. Let $f(m)$ be the minimal value of $\frac{T}{O}$ over all possible inputs of size m ($|\mathcal{I}| = m$), when the union of the matchings produced by ALG and OPT forms a single cycle.

For the sake of this analysis consider Figure 1.

Let P_{ca} be the alternating path from c to a , and P_{db} be the alternating path from d to b . Denote $\alpha = \sum_{e \in P_{ca}} D(e)$, and $\beta = \sum_{e \in P_{db}} D(e)$. Then, by the triangle inequality

$$\alpha \geq D(a, c) \quad (2)$$

From Lemma 2 we have

$$(1 + \epsilon)D(a, c) \geq D(a, b) \quad (3)$$




Figure 1: The cycle formed by the union of the matchings produced by ALG and OPT. The length of P_{ca} is α , and the length of P_{db} is β .

It follows from Equations (2) and (3) that

$$1 - \alpha - \beta = D(a, b) \leq (1 + \epsilon)\alpha \quad (4)$$

Similarly $1 - \alpha - \beta \leq (1 + \epsilon)\beta$.

Let $2i$ be the number of points on P_{ca} , then $f(m)$ satisfies the recurrence relation

$$f(m) = \min_{\substack{1 \leq i < \frac{m}{2} - 1 \\ 0 < 1 - \alpha - \beta \leq (1 + \epsilon)\alpha \\ 0 < 1 - \alpha - \beta \leq (1 + \epsilon)\beta}} \{ \alpha f(2i) + \beta f(m - 2i) \} \quad (5)$$

Conditioning on t , $f(t)$ and $f(m - t)$ are constant, therefore $\alpha f(t) + \beta f(m - t)$ becomes a linear function in α and β , so its minimum must occur at a vertex of the polyhedron defined by the minimization constraints (see for example [8]).

The vertices of this polyhedron are $(1, 0)$, $(0, 1)$, $(\frac{1}{3 + \epsilon}, \frac{1}{3 + \epsilon})$, so

$$f(m) = \min_{1 \leq i \leq \frac{m}{2} - 1} \left\{ f(2i), \frac{1}{3 + \epsilon} (f(2i) + f(m - 2i)) \right\} \quad (6)$$

Also note that $f(2) = 1$, since there is only one way to match two points, so $T = O$. The conditions of Lemma 5 are met with $\gamma = 3 + \epsilon$, thus

$$f(m) = \Omega \left(\frac{1}{m^{\log(\frac{3 + \epsilon}{2})}} \right)$$

Finally, from 1 we conclude

$$\frac{\text{ALG}_{\text{OFF}}}{\text{OPT}} \leq 2 \frac{T}{O} - 1 \leq \frac{2}{f(m)} = O \left(m^{\log(\frac{3 + \epsilon}{2})} \right)$$

□

Lemma 7. $\text{ALG}_{\text{ON}} = \Theta \left(\frac{1}{\epsilon} \right) \text{ALG}_{\text{OFF}}$

Proof. Assume two requests p and q were matched by ALG at time t . Assume w.l.o.g that $t(p) \geq t(q)$. The contribution of this pair to ALG_{ON} , is

$$\begin{aligned} t - t(p) + t - t(q) + d(x(p), x(q)) = \\ t - t(p) + t - t(p) + t(p) - t(q) + d(x(p), x(q)) = 2(t - t(p)) + D(p, q) \end{aligned}$$

On the contrary, the contribution of this pair to ALG_{OFF} , is just $D(p, q)$.

Note that t is the time that q was on p 's hemisphere, so $t = t(p) + \frac{D(p, q)}{\epsilon}$, hence the ratio between ALG_{ON} and ALG_{OFF} for this pair is

$$\frac{2\frac{D(p, q)}{\epsilon} + D(p, q)}{D(p, q)} = 1 + \frac{2}{\epsilon}$$

Summing over all matched pairs we get $\frac{\text{ALG}_{\text{ON}}}{\text{ALG}_{\text{OFF}}} = 1 + \frac{2}{\epsilon} = \Theta\left(\frac{1}{\epsilon}\right)$. \square

Finally we prove Theorem 1 using the inequalities proven in the previous lemmas.

Proof of Theorem 1. Combining Lemma 1, Lemma 6 and Lemma 7 we have

$$\text{ALG}_{\text{ON}} \leq O\left(\frac{1}{\epsilon}\right) \text{ALG}_{\text{OFF}} \leq O\left(\frac{1}{\epsilon} m^{\log\left(\frac{3+\epsilon}{2}\right)}\right) \text{OPT}$$

Hence, $\text{ALG}(\epsilon)$ is $O\left(\frac{1}{\epsilon} m^{\log\left(\frac{3+\epsilon}{2}\right)}\right)$ -competitive. \square

In Appendix A we show that the analysis is tight, and that the competitive ratio is indeed a function of m , and not of n (the size of the metric space). In Appendix B we show that growing hemispheres in space while ignoring the time axis, and other similar hacks, only worsen the competitive ratio.

4 The Bipartite Case

For the bipartite case, we suggest the same algorithm as in the monochromatic case. The only difference is that we match a request q to a request p as soon as q is found on the boundary of p 's hemisphere, **and that q and p do not belong to the same class**.

Algorithm 2 A Deterministic Algorithm for MBPMD on General Metrics

```

1: procedure ALG-B( $\epsilon$ )
2:   At every moment  $t$ :
3:     Add the new requests that arrive at time  $t$ 
4:     for each unmatched request  $p$  do
5:       for each unmatched request  $q \neq p$  do
6:         if  $t(p) \geq t(q)$  and  $t = t(p) + \frac{D(x(p), x(q))}{\epsilon}$  and  $\text{class}(q) \neq \text{class}(p)$  then
7:           match( $p, q$ )
8:         end if
9:       end for
10:    end for
11: end procedure

```

4.1 Analysis

We prove the following theorem:

Theorem 2. $\text{ALG-B}(\epsilon)$ is $O\left(\frac{1}{\epsilon} m^{\log\left(\frac{3+\epsilon}{2}\right)}\right)$ -competitive.

Observation 1, Lemma 3 and Lemma 4 hold for the bipartite case as well, therefore using Corollary 1 we may assume that the union of ALG-B and OPT forms a single cycle.

The key difference in the analysis for this case, is that when we consider the last four requests to be matched, not every two of them could have been matched to each other. Therefore Lemma 2 does not hold, but a weaker yet similar result does.

Consider the last two pairs of requests to be matched by ALG-B. Name them (a, b) and (c, d) , and assume w.l.o.g that (a, b) were matched at time t_{ab} , and (c, d) at $t_{cd} \geq t_{ab}$. Also, assume w.l.o.g that $t(a) \leq t(b)$.

Lemma 8. *If $\text{class}(a) = \text{class}(d) \neq \text{class}(b) = \text{class}(c)$ then*

$$(1) \ D(a, b) \leq (1 + \epsilon)D(a, c)$$

$$(2) \ D(a, b) \leq (1 + \epsilon)D(b, d)$$

We omit the proof of this lemma as it is the same as the proof of Lemma 2 for the relevant cases.

Considering Figure 1 we have the following lemma.

Lemma 9. *$\text{class}(a) = \text{class}(d) \neq \text{class}(b) = \text{class}(c)$*

Proof. From the alternation property of Observation 1 we have that the number of edges along P_{ca} must be odd (since the number of OPT edges along P_{ca} must be one more than ALG-B edges along P_{ca}). Moreover, the classes of the requests along P_{ca} alternate as well (since every edge must match requests of different classes). Since there are odd number of edges along P_{ca} , there are odd number of class alternations along P_{ca} , so the class of the last request along P_{ca} (that is $\text{class}(c)$) must be different from the class of the first request along P_{ca} (that is $\text{class}(a)$). Thus $\text{class}(c) \neq \text{class}(a)$ and of course $\text{class}(a) \neq \text{class}(b)$, $\text{class}(c) \neq \text{class}(d)$, so $\text{class}(a) = \text{class}(d) \neq \text{class}(b) = \text{class}(c)$. \square

Using Lemma 9 and Lemma 8 we repeat the proof of Lemma 6 and achieve the following result:

Lemma 10. $\text{ALG-B}_{\text{OFF}} \leq O\left(m^{\log(\frac{3+\epsilon}{2})}\right) \text{OPT}$

The main theorem for the bipartite case now follows:

Proof of Theorem 2. Lemma 7 and Lemma 1 hold for ALG-B as well, thus from Lemma 10 we have

$$\text{ALG-B}_{\text{ON}} \leq O\left(\frac{1}{\epsilon}\right) \text{ALG-B}_{\text{OFF}} \leq O\left(\frac{1}{\epsilon} m^{\log(\frac{3+\epsilon}{2})}\right) \text{OPT}$$

Hence, $\text{ALG-B}(\epsilon)$ is $O\left(\frac{1}{\epsilon} m^{\log(\frac{3+\epsilon}{2})}\right)$ -competitive. \square

5 Concluding Remarks and Open Problems

In this paper we presented the first sub-linear competitive deterministic algorithm for *Minimum-Cost Perfect Matching with Delays* as a function of m , the number of requests. We also provided a similar algorithm for the problem of *Minimum-Cost Bipartite Perfect Matching with Delays* achieving the same competitive ratio.

One open problem is to decide if a deterministic algorithm with a better competitive ratio exists, in particular a $\text{polylog}(m)$ -competitive one, by showing a lower bound or providing an algorithm for the problem. In addition, the problem of finding a sub-linear in n competitive deterministic algorithm is still open.

References

- [1] Antonios Antoniadis, Neal Barcelo, Michael Nugent, Kirk Pruhs, and Michele Scquizzato. A $o(n)$ -competitive deterministic algorithm for online matching on a line. In *Approximation and Online Algorithms - 12th International Workshop*, pages 11–22, 2014.
- [2] Itai Ashlagi, Yossi Azar, Moses Charikar, Ashish Chiplunkar, Ofir Geri, Haim Kaplan, Rahul M. Makhijani, Yuyi Wang, and Roger Wattenhofer. Min-cost bipartite perfect matching with delays. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 1:1–1:20, 2017.
- [3] Yossi Azar, Ashish Chiplunkar, and Haim Kaplan. 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*, pages 1051–1061, 2017.
- [4] Yossi Azar, Arun Ganesh, Rong Ge, and Debmalya Panigrahi. Online service with delay. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 551–563, 2017.
- [5] Nikhil Bansal, Niv Buchbinder, Anupam Gupta, and Joseph Naor. A randomized $O(\log^2 k)$ -competitive algorithm for metric bipartite matching. *Algorithmica*, 68(2):390–403, 2014.
- [6] Marcin Bienkowski, Artur Kraska, Hsiang-Hsuan Liu, and Pawel Schmidt. A primal-dual online deterministic algorithm for matching with delays. *CoRR*, abs/1804.08097, 2018.
- [7] Marcin Bienkowski, Artur Kraska, and Pawel Schmidt. A match in time saves nine: Deterministic online matching with delays. In *Approximation and Online Algorithms - 15th International Workshop*, pages 132–146, 2017.
- [8] George B. Dantzig. *Linear programming and extensions*. Princeton University Press, 1963.
- [9] Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965.
- [10] Yuval Emek, Shay Kutten, and Roger Wattenhofer. Online matching: haste makes waste! In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 333–344, 2016.
- [11] Yuval Emek, Yaacov Shapiro, and Yuyi Wang. Minimum cost perfect matching with delays for two sources. In *Algorithms and Complexity - 10th International Conference*, pages 209–221, 2017.
- [12] Bernhard Fuchs, Winfried Hochstättler, and Walter Kern. Online matching on a line. *Electronic Notes in Discrete Mathematics*, 13:49–51, 2003.
- [13] Anupam Gupta and Kevin Lewi. The online metric matching problem for doubling metrics. In *Automata, Languages, and Programming - 39th International Colloquium*, pages 424–435, 2012.
- [14] Bala Kalyanasundaram and Kirk Pruhs. Online weighted matching. *J. Algorithms*, 14(3):478–488, 1993.
- [15] Bala Kalyanasundaram and Kirk Pruhs. On-line network optimization problems. In *Online Algorithms, The State of the Art (the book grow out of a Dagstuhl Seminar)*, pages 268–280, 1996.

- [16] Samir Khuller, Stephen G. Mitchell, and Vijay V. Vazirani. On-line algorithms for weighted bipartite matching and stable marriages. In *Automata, Languages and Programming, 18th International Colloquium*, pages 728–738, 1991.
- [17] Elias Koutsoupias and Akash Nanavati. The online matching problem on a line. In *Approximation and Online Algorithms, First International Workshop*, pages 179–191, 2003.
- [18] Adam Meyerson, Akash Nanavati, and Laura J. Poplawski. Randomized online algorithms for minimum metric bipartite matching. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 954–959, 2006.
- [19] Krati Nayyar and Sharath Raghvendra. An input sensitive online algorithm for the metric bipartite matching problem. In *58th IEEE Annual Symposium on Foundations of Computer Science*, pages 505–515, 2017.
- [20] Sharath Raghvendra. A robust and optimal online algorithm for minimum metric bipartite matching. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 18:1–18:16, 2016.
- [21] Edward M. Reingold and Robert Endre Tarjan. On a greedy heuristic for complete matching. *SIAM J. Comput.*, 10(4):676–681, 1981.

A The competitive ratio is a function of m

Following Section 3.1, a question arises - whether Theorem 1 can be modified to prove $\frac{\text{ALG}}{\text{OPT}} \leq O\left(\frac{1}{\epsilon} n^{\log(\frac{3}{2}+\epsilon)}\right)$ for a finite metric space of size n .

We show that for every $\text{ALG}(\epsilon)$ there is an instance with $n = 1$ for which $\frac{\text{ALG}}{\text{OPT}} \geq \Omega\left(\frac{1}{\epsilon} m^{\log(\frac{3+2\epsilon}{2+2\epsilon})}\right)$. The instance we give is essentially the example given by [21], over the time axis, and with distances scaled to consider the progress of time. Let $k = \log(m)$ and consider Figure 2 which describes a series of requests with the recurrence relation

$$a_i = \frac{b_i}{1 + \epsilon}, \quad b_i = 2b_{i-1} + a_{i-1}, \quad b_1 = 1 \quad (7)$$

Lemma 11. *ALG matches $(r_2, r_3), \dots, (r_{m-2}, r_{m-1})$ and (r_1, r_m) .*

Proof. We prove the lemma by induction on k .

Base Case ($k = 1$): The only point that $r_2 = r_m$ can be matched to is r_1 .

Inductive step: Assume the claim holds for $k - 1$. We start by showing that ALG will match the pairs $(r_2, r_3), \dots, (r_{\frac{m}{2}-2}, r_{\frac{m}{2}-1})$.

Let $t_0 = b_{k-1} + \frac{b_{k-1}}{\epsilon}$, this is the time that the hemisphere of $r_{\frac{m}{2}}$ reaches r_1 unless $r_{\frac{m}{2}}$ is matched by another request at time $t < t_0$. By the induction hypothesis, unless the hemisphere of some r_i with $i > \frac{m}{2}$ reaches past $r_{\frac{m}{2}}$ by $t < t_0$, the hemisphere of $r_{\frac{m}{2}}$ will reach r_1 , after the pairs $(r_2, r_3), \dots, (r_{\frac{m}{2}-2}, r_{\frac{m}{2}-1})$ are matched. Notice that the hemisphere of $r_{\frac{m}{2}}$ may reach r_1 only by time t_0 and the hemisphere of $r_{\frac{m}{2}+1}$ may reach $r_{\frac{m}{2}}$ only by

$$t_1 = b_{k-1} + a_{k-1} + \frac{a_{k-1}}{\epsilon} = b_{k-1} + \frac{b_{k-1}}{1 + \epsilon} \left(1 + \frac{1}{\epsilon}\right) = t_0$$




Figure 2: A series of m requests along the time axis with $n = 1$ and $\frac{\text{ALG}}{\text{OPT}} \geq \Omega\left(\frac{1}{\epsilon} m^{\log\left(\frac{3+2\epsilon}{2+2\epsilon}\right)}\right)$.
In blue is the matching produced by ALG, and in dashed red - a matching of cost $O(m)$.

Therefore, the hemisphere of $r_{\frac{m}{2}+1}$ may reach $r_{\frac{m}{2}}$ only after $(r_2, r_3), \dots, (r_{\frac{m}{2}-2}, r_{\frac{m}{2}-1})$ are matched. Obviously for every $i > \frac{m}{2} + 1$ the hemisphere of r_i would not reach past $r_{\frac{m}{2}}$ by t_0 if the hemisphere of $r_{\frac{m}{2}+1}$ does not, therefore $(r_2, r_3), \dots, (r_{\frac{m}{2}-2}, r_{\frac{m}{2}-1})$ are matched by ALG by time t_0 .

Considering $r_{\frac{m}{2}+1}, \dots, r_m$, again by the induction hypothesis we have that unless r_m is matched by another request before its hemisphere reaches $r_{\frac{m}{2}+1}$, ALG will match the pairs $(r_{\frac{m}{2}+2}, r_{\frac{m}{2}+3}), \dots, (r_{m-2}, r_{m-1})$. Indeed, there is no request after r_m , thus ALG will match these pairs, and we are left to address the requests $r_1, r_{\frac{m}{2}}, r_{\frac{m}{2}+1}, r_m$.

Observe that the hemisphere of r_m reaches $r_{\frac{m}{2}+1}$ at $t = b_k + \frac{b_{k-1}}{\epsilon} > (1 + \frac{1}{\epsilon})b_{k-1} = t_1 = t_0$, hence ALG will match the pair $(r_{\frac{m}{2}+1}, r_{\frac{m}{2}})$. The remaining and last pair to be matched by ALG is (r_1, r_m) of course. \square

The cost of OPT is at most $O(m)$ since $D(u, v) = b_1 = 1$ for every pair (u, v) in the matching $(r_1, r_2) \dots (r_{m-1}, r_m)$. The cost of matching r_1 to r_m is b_k . Out of the pairs $(r_2, r_3) \dots (r_{m-2}, r_{m-1})$ there are 2^i pairs with distance a_{k-i-1} between the two end-points, for $0 \leq i \leq k-2$. Therefore $\text{ALG}_{\text{OFF}} = b_k + \sum_{i=0}^{k-2} 2^i a_{k-i-1}$.

The mutual recurrence relation (7) solves to

$$a_i = \frac{\left(2 + \frac{1}{1+\epsilon}\right)^i}{2\epsilon + 3}, \quad b_i = \left(2 + \frac{1}{1+\epsilon}\right)^{i-1} \quad (8)$$

Therefore

$$\begin{aligned}
\text{ALG}_{\text{OFF}} &= b_k + \sum_{i=1}^{k-1} 2^{k-1-i} a_i > \sum_{i=1}^{k-1} \frac{\left(2 + \frac{1}{1+\epsilon}\right)^i}{2\epsilon + 3} 2^{k-1-i} \\
&= \frac{2^{k-1}}{2\epsilon + 3} \sum_{i=1}^{k-1} \left(1 + \frac{1}{2(1+\epsilon)}\right)^i \\
&= \frac{2^{k-1}}{2\epsilon + 3} (2\epsilon + 3) \left(\left(1 + \frac{1}{2(1+\epsilon)}\right)^{k-1} - 1 \right) \\
&= \left(2 + \frac{1}{1+\epsilon}\right)^{k-1} - 2^{k-1}
\end{aligned}$$

Hence,

$$\frac{\text{ALG}_{\text{OFF}}}{\text{OPT}} \geq \frac{\left(2 + \frac{1}{1+\epsilon}\right)^{k-1} - 2^{k-1}}{2^k} = \Omega\left(m^{\log\left(1 + \frac{1}{2(1+\epsilon)}\right)}\right) = \Omega\left(m^{\log\left(\frac{3+2\epsilon}{2+2\epsilon}\right)}\right)$$

Finally, from Lemma 7 we have

$$\frac{\text{ALG}_{\text{ON}}}{\text{OPT}} = \Omega\left(\frac{1}{\epsilon}\right) \frac{\text{ALG}_{\text{OFF}}}{\text{OPT}} \geq \Omega\left(\frac{1}{\epsilon} m^{\log\left(\frac{3+2\epsilon}{2+2\epsilon}\right)}\right)$$

B Time must be considered

A simple hack that may handle the instance given in Appendix A, is to match immediately two points that are located at the same position in space. Obviously this will not handle some very similar instances, generated by small perturbations of the positions of the requests.

A simple extension of this idea is to ignore the time axis, so that p and q will be matched as soon as $t \geq \min(t(p), t(q)) + \frac{d(x(p), x(q))}{\epsilon}$, i.e. the requests grow spheres only in space, but not in time, and they are matched to each other as soon as one of them is in the sphere of the other.

The instance in Figure 3 shows that the competitive-ratio of this algorithm can be worse as $\Omega(m)$, even though the size of the metric space is $n = 2$.




Figure 3: In blue - the matching produced by the suggested algorithm, of cost $O(m)$.
In dashed red - an alternative matching of cost $O(1 + \delta m)$.

Note that $\Omega(m)$ competitive-ratio will be achieved for this instance, even for similar algorithms which do not consider time, such as matching p to q if $t(p) \geq t(q)$ and $t \geq t(p) + \frac{d(x(p), x(q))}{\epsilon}$, or matching p to q if $t(p) \leq t(q)$ and $t \geq t(p) + \frac{d(x(p), x(q))}{\epsilon}$.