Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Online Duet between Metric Embeddings and Minimum-Weight Perfect Matchings*

Sujoy Bhore† Arnold Filtser‡ Csaba D. Tóth§

Abstract

Low-distortional metric embeddings are a crucial component in the modern algorithmic toolkit. In an online metric embedding, points arrive sequentially and the goal is to embed them into a simple space irrevocably, while minimizing the distortion. Our first result is a deterministic online embedding of a general metric into Euclidean space with distortion $O(\log n) \cdot \min\{\sqrt{\log \Phi}, \sqrt{n}\}$ (or, $O(d) \cdot \min\{\sqrt{\log \Phi}, \sqrt{n}\}$ if the metric has doubling dimension d), solving affirmatively a conjecture by Newman and Rabinovich (2020), and quadratically improving the dependence on the aspect ratio Φ from Indyk et al. (2010). Our second result is a stochastic embedding of a metric space into trees with expected distortion $O(d \cdot \log \Phi)$, generalizing previous results (Indyk et al. (2010), Bartal et al. (2020)).

Next, we study the problem of *online minimum-weight perfect matching* (*MWPM*). Here a sequence of $2n$ points $s_1, \ldots s_{2n}$ in a metric space arrive in pairs, and one has to maintain a perfect matching on the first $2i$ points $S_i = \{s_1, \ldots s_{2i}\}$. We allow recourse (as otherwise the order of arrival determines the matching). The goal is to return a perfect matching that approximates the *minimum-weight* perfect matching on S_i, while minimizing the recourse. Online matchings are among the most studied online problems, however, there is no previous work on online MWPM. One potential reason for this is that online MWPM is drastically non-monotone, which makes online optimization highly challenging. Our third result is a randomized algorithm with competitive ratio $O(d \cdot \log \Phi)$ and recourse $O(\log \Phi)$ against an oblivious adversary, this result is obtained via our new stochastic online embedding. Our fourth result is a deterministic algorithm that works against an adaptive adversary, using $O(\log^2 n)$ recourse, and maintains a matching of total weight at most $O(\log n)$ times the weight of the MST, i.e., a matching of lightness $O(\log n)$. We complement our upper bounds with a strategy for an oblivious adversary that, with recourse r, establishes a lower bound of $\Omega(\frac{\log n}{r \log r})$ for both competitive ratio as well as lightness.

1 Introduction

The traditional model of algorithms design solves problems where the entire input is given in advance. In contrast, online algorithms work under conditions of uncertainty, gradually receiving an input sequence $\sigma = \sigma_1, \sigma_2, \ldots, \sigma_n$ (where σ_i is presented at step i). The algorithm has to serve them in the order of occurrence, where the decisions are irrevocable, and without prior knowledge of subsequent terms of the input. The objective is to optimize the total cost paid on the entire sequence σ. The performance of an online algorithm ALG is measured using competitive analysis, where ALG is compared to an optimal offline algorithm that knows the entire sequence in advance and can provide the solution with optimum cost [BE98, Ch. 1]. The two most central adversarial models in online algorithms are adaptive and oblivious. In the *adaptive adversary* model, the sequence of arriving points is determined "on the fly", and may depend on the previous decisions made by the algorithm. This is a restrictive model, and in particular, randomization is not helpful in this model. An algorithm is *k-competitive against an adaptive adversary* if, for every sequence σ of requests, the cost of the algorithm is at most k times the optimal offline solution.[1] An *oblivious adversary* assumes that the input sequence is determined

*The full version of the paper can be accessed at http://arxiv.org/abs/2310.14078

†Department of Computer Science & Engineering, Indian Institute of Technology Bombay, Mumbai, India. Email: sujoy@cse.iitb.ac.in

‡Bar Ilan University, Ramat Gan, Israel. Email: arnold.filtser@biu.ac.il. This research was supported by the Israel Science Foundation (grant No. 1042/22).

§Department of Mathematics, California State University Northridge, Los Angeles, CA; and Department of Computer Science, Tufts University, Medford, MA, USA. Email: csaba.toth@csun.edu. Research on this paper was supported, in part, by the NSF award DMS-2154347.

[1]This paper is focused on minimization problems. In a maximization problem, an algorithm is k-competitive if for every sequence σ, the cost of the algorithm is at least a k fraction of the optimal offline solution.

Copyright © 2024
Copyright for this paper is retained by authors

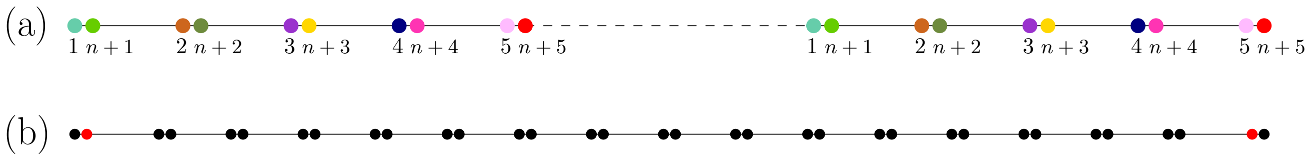

Figure 1: (a) Example where the weight of an online matching is arbitrarily far form optimum assuming irrevocable decisions (i.e., no recourse). The metric is the real line. We first receive the pairs $\{i, W+i\}_{i=1}^{n}$, and then the pairs $\{i+\varepsilon, W+i+\varepsilon\}_{i=1}^{n}$, for sufficiently small ε and large W. The weight of the online perfect matching (specified in the illustration) is $2n \cdot W$, while the cost of the optimal perfect matching is $2\varepsilon \cdot n$.
(b) Example of the drastic non-monotonicity of minimum-weight perfect matching. The metric is the real line, where each point $\{1, 2 \ldots, 2n\}$ appears twice, while the points $\{0, 2n+1\}$ appear once. Then the weight of a perfect matching is $2n+1$. After introducing the pair $\{0, 2n+1\}$ (red in the figure), the weight of the perfect matching drops to 0.

in advance (however, unknown to the algorithm). Here randomization can be useful. A randomized online algorithm is *k-competitive against an oblivious adversary* if, for every sequence σ of requests, the expected cost of the algorithm is at most k times the optimal offline solution.[1]

Online problems arise in various areas of computer science, such as scheduling, network optimization, data structures, resource management in operating systems, etc.; see [BN09, KVV90, Alb03, BFK+17]. Some preeminent examples of online problems are k-server [BCR23], job scheduling [LLMV20], routing [AAF+97], load balancing [Aza05], among many others.

One of the most fundamental and well-studied problems in the online algorithms world is online matching. Starting with the seminal paper by Karp, Vazirani, and Vazirani [KVV90], a large body of work on "online matchings" is devoted to the *online maximum bipartite matching (server-client model)* problem, where one side of the bipartite graph (*servers*) is fixed and the vertices of the other side (*clients*) are revealed one at a time: The objective is to maintain a maximum matching (not necessarily a perfect matching). Since then, numerous variants of this problem have been studied, see e.g. [AS22, BN09, DH09, GKM+19, GM08, KP98, MSVV07, SWW95].

Online Minimum-Weight Perfect Matching. In this paper, we study online *minimum-weight* perfect matchings (MWPM) in metric spaces. Points $s_1, \ldots s_{2n}$ arrive sequentially from a metric space (X, d_X) (unknown in advance). For each new point s_i, we are given the distances to all previous points: $\{d_X(s_i, s_j)\}_{j=1}^{i-1}$. [2] Denote by $S_i = \{s_1, \ldots, s_{2i}\}$ the set of the first $2i$ points. The goal is to maintain a perfect matching M_i on S_i such that the difference between M_i and M_{i+1} is bounded by a constant or a polylogarithmic function of n, and the weight of M_i is as small as possible.

A standard *online* algorithm can add edges to the matching, but the decisions are irrevocable, and therefore no edge is ever deleted. In this setting, the matching is completely determined by the order of points: $M = \{\{s_{2i-1}, s_{2i}\} : i = 1, \ldots, n\}$, and the weight of M may be arbitrarily far from the optimum (see an example in Figure 1 (a)). For this reason, we allow *recourse* r: the online algorithm has to maintain a perfect matching on S_i, and in each step, it can delete up to r edges. Our primary focus is the trade-off between recourse and the weight of the matching M_i.

Surprisingly, even though matchings are one of the most meticulously studied online problems, essentially no previous results were known for the online minimum-weight perfect matching problem with recourse.[3] Note that the server-client model based online algorithms are significantly different from ours, and in general are not helpful for our problem. One natural difficulty is that in contrast to other classical optimization problems, e.g., MST or TSP, the minimum weight of perfect matching is drastically non-monotone:[4] it can decrease from a large weight to 0 after introducing a single new pair! (See Figure 1 (b) for an illustration.) This non-monotonicity is

[2]Equivalently, there is an underlying complete weighted graph $G = (V, E, w)$ with weight respecting triangle inequality. For each new vertex, we receive the edges to all previously arrived vertices.

[3]The related problem of online minimum-weight perfect matching with delays [ACK17, DU23, EKW16] has been studied previously; see Section 1.3.

[4]Due to the triangle inequality, the cost of an optimal TSP tour can only increase as new points arrive. The cost of an MST could decrease after additional points arrive, but this could happen up to at most a factor of 2 (due to the fact that the MST is a 2 approximation of the minimum Steiner tree).

Copyright © 2024
Copyright for this paper is retained by authors

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

the major bottleneck for maintaining a good approximation with limited recourse. We further discuss related classical online optimization problems in Section 1.3.

For online MST, for example, Gu et al. [GGK16] achieve a competitive ratio of $2^{O(k)}$ with a single recourse for every k new points (that is fractional recourse, on average). However, for MWPM we show (Proposition 1.1) that there is no competitive online algorithm if we are allowed to use a single recourse per vertex pair, which already holds for a sequence of 8 points on a real line, even if the sequence is known in advance. That is to say, online MWPM is a much more challenging problem than online MST.

Online Metric Embeddings. Low-distortion metric embeddings are a crucial component in the modern algorithmic toolkit with applications in approximation algorithms [LLR95], distributed algorithms [KKM+12], online algorithms [BBMN15], and many more. A metric embedding is a map $f: X \to Y$ between the points of two metric spaces (X, d_X) and (Y, d_Y). The *contraction* and *expansion* of the map f are the smallest $\rho, t \geq 1$, respectively, such that for every pair $x, y \in X$,

$$\rho^{-1} \cdot d_X(x, y) \leq d_Y(f(x), f(y)) \leq t \cdot d_X(x, y) .$$

The *distortion* of the map is then $\rho \cdot t$. The embedding is non-contractive (non-expansive) if $\rho = 1$ ($t = 1$). Such an embedding is also called *dominating*. Arguably, in the TCS community, the two most celebrated metric embeddings are the following: (1) every n-point metric space embeds into Euclidean space ℓ_2 with distortion $O(\log n)$ [Bou85], and (2) every n-point metric space stochastically embeds into a distribution over tree metrics (in fact ultrametrics[5]) with expected distortion $O(\log n)$ [FRT04] (see also [Bar96, Bar98, Bar04]). Specifically, there is a distribution $\mathcal{D}$ over pairs (f, U), where U is an ultrametric, and $f: X \to U$ is a dominating embedding, such that for all $x, y \in X$, we have $\mathbb{E}_{(f,U)\sim\mathcal{D}}[d_U(f(x), f(y)] \leq O(\log n) \cdot d_X(x, y)$.

While [Bou85, FRT04] enjoined tremendous success and have numerous applications, they require to know the metric space (X, d_X) in advance, and hence cannot be used in an online fashion where the points are revealed one by one. In this paper, we first focus on online metric embeddings:

Definition 1. (Online Embedding) *An* online embedding *of a sequence of points $x_1, \ldots, x_k$ from a metric space (X, d_X) into a metric space (Y, d_Y) is a sequence of embeddings $f_1, \ldots, f_k$ such that for every i, f_i is a map from $\{x_1, \ldots, x_i\}$ to Y, and f_{i+1} extends f_i (i.e., $f_i(x_j) = f_{i+1}(x_j)$ for $j \leq i$). The embedding has* expansion $\alpha = \max_{i,j\leq k} \frac{d_Y(f_k(x_i), f_k(x_j))}{d_X(x_i, x_j)}$, *and* contraction $\beta = \max_{i,j\leq k} \frac{d_X(x_i, x_j)}{d_Y(f_k(x_i), f_k(x_j))}$. *The* distortion *of the online embedding is $\alpha \cdot \beta$. If $\beta \leq 1$, we say that the embedding is* dominating.

A stochastic online embedding *is a distribution $\mathcal{D}$ over dominating online embeddings. A stochastic online embedding $\mathcal{D}$ has expected distortion t if for every x_i, x_j, $\mathbb{E}_{f_k\sim\mathcal{D}}[d_M(f_k(x_i), f_k(x_j))] \leq t \cdot d_X(x_i, x_j)$. In an online embedding algorithm, the embedding f_i can depend only on $\{x_1, \ldots, x_i\}$.*

For stochastic online embeddings, the sequence of points should be fixed in advance (but unknown to the algorithm). That is, a deterministic online embedding can be used against an adaptive adversary, while stochastic online embedding can be used only against an oblivious adversary.

Indyk, Magen, Sidiropoulos, and Zouzias [IMSZ10] (see also [ERW10]) observed that Bartal's original embedding [Bar96] can be used in an online fashion to produce a stochastic embedding into trees (ultrametrics) with expected distortion $O(\log n \cdot \log \Phi)$, where $\Phi = \frac{\max_{x,y\in X} d_X(x,y)}{\min_{x,y\in X} d_X(x,y)}$ is the *aspect ratio* (a.k.a. *spread*). Their original embedding had the caveat that the number of metric points n and the aspect ratio Φ must be known in advance. Later, Bartal, Fandina, and Umboh [BFU20] removed these restrictions. They also provided an $\Omega\left(\frac{\log n \cdot \log \Phi}{\log \log n}\right)$ lower bound, showing this distortion to be tight up to second order terms. For the case where the input metric space (X, d_X) has doubling dimension[6] ddim (known in advance), Indyk et al. [IMSZ10] constructed a stochastic online embedding into ultrametrics with expected distortion $2^{O(\mathrm{ddim})} \cdot \log \Phi$.

In an attempt to construct an online version of Bourgain's embedding [Bou85], Indyk et al. [IMSZ10] constructed online stochastic embedding of an arbitrary n-point metric space into the Euclidean space ℓ_2 with expected distortion $O(\log n \cdot \sqrt{\log \Phi})$ (again, n and Φ must be known in advance), or with expected distortion $2^{O(\mathrm{ddim})} \cdot \log \Phi$ for the case where the metric space has doubling dimension ddim (known in advance).

[5] An ultrametric is a metric space with the strong triangle inequality: $\forall x, y, z,\ d_X(x, y) \leq \max\{d_X(x, z), d_X(z, y)\}$. In particular, ultrametric can be represented as the shortest path metric of a tree graph. See Defintion 2.

[6] A metric (X, d) has doubling dimension ddim if every ball of radius $2r$ can be covered by 2^{ddim} balls of radius r.

Copyright © 2024
Copyright for this paper is retained by authors

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

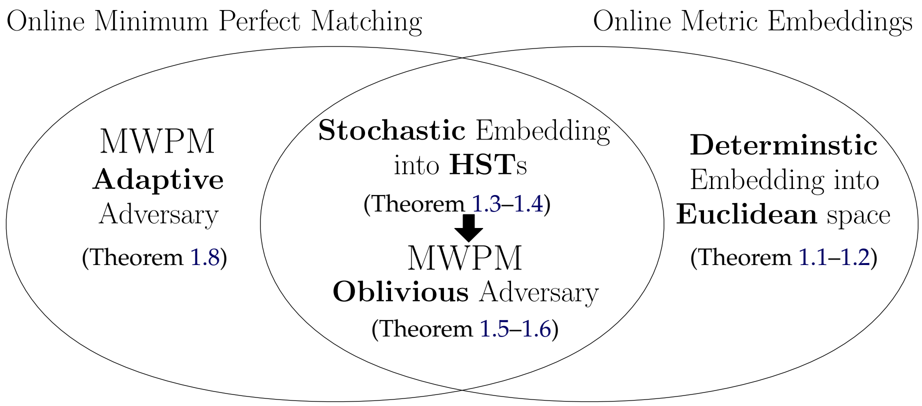

Figure 2: A "duet" between online metric embeddings and minimum-weight perfect matchings: a Venn diagram of the relationship between the various results in this paper.

Newman and Rabinovich [NR20] showed that every deterministic embedding into ℓ_2 must have distortion $\Omega(\min\{\sqrt{n},\sqrt{\log\Phi}\})$. They conjectured[7] that a similar upper bound holds:

CONJECTURE 1. ([NR20]) *Every sequence of n points in a metric space received in an online fashion can be deterministically embedded into Euclidean space ℓ_2 with distortion* $\mathrm{poly}(n)$.

1.1 Our Results The results in this paper are twofold: We study both online minimum-weight perfect matchings and online metric embeddings. The connections between them are illustrated in Figure 2. The proofs of the lemmas and theorems in this extended abstract can be found in the full version of this paper (available on arXiv).

1.1.1 Metric Embeddings Our results on online metric embeddings are summarized in Table 1. Our first result is a deterministic online embedding with distortion $O(\mathrm{ddim})\cdot\min\{\sqrt{\log\Phi},\sqrt{n}\}$ into Euclidean space ℓ_2 where the input metric has doubling dimension ddim. No prior knowledge (of any of n, Φ, and ddim) is required. As every n-point metric space has doubling dimension $O(\log n)$, this result simultaneously: (1) proves the conjecture by Newman and Rabinovich [NR20] (with an upper bound of $O(\sqrt{n}\log n)$); (2) matches the lower bound up to second order terms; (3) exponentially improves the dependence on ddim compared to [IMSZ10]; (4) quadratically improves the dependence on Φ; (5) gives a deterministic distortion guarantee instead of expected distortion; and (6) removes the requirement to know ddim and Φ in advance. In fact, the quadratic improvement in the dependence on Φ answers an open question in [IMSZ10] (see Remark 2 in [IMSZ10]).

THEOREM 1.1. *For a sequence of metric points $x_1,\dots,x_n$ arriving in an online fashion, there is a deterministic online embedding into Euclidean space ℓ_2 with distortion $O(\mathrm{ddim})\cdot\min\{\sqrt{\log\Phi},\sqrt{n}\}$. Here Φ is the aspect ratio, and* ddim *is the doubling dimension of $\{x_1,\dots,x_n\}$. No prior knowledge is required.*

Our second result is a lower bound showing that for constant doubling dimension, Theorem 1.1 is tight. Our lower bound holds even for stochastic online embeddings with expected distortion guarantees. This generalizes [NR20], where it was shown that there is a family of n-point metric spaces with aspect ratio $\Phi = 2^{\Omega(n)}$ such that every deterministic embedding into ℓ_2 requires distortion $\Omega(\sqrt{\log\Phi}) = \Omega(\sqrt{n})$ (however their family does not have bounded doubling dimension).

THEOREM 1.2. *For every $n\in\mathbb{N}$, there is a family $\mathcal{M}_n$ of metric spaces with $O(n)$ points, aspect ratio $\Phi = 4^n$, and uniformly constant doubling dimension, where each metric $(X,d_X)\in\mathcal{M}_n$ constitutes a shortest path metric of a series parallel (in particular, planar) graph, such that every stochastic online embedding into ℓ_2 has expected distortion $\Omega(\sqrt{\log\Phi}) = \Omega(\sqrt{n})$.*

[7] The conjecture appears in the full arXiv version (in the conference version, it is stated as an open problem).

Copyright © 2024
Copyright for this paper is retained by authors

	Input space	Host Space	Distortion	Reference	Deter?	Pri.Kno.
1.	General	ℓ_2	$O(\log n \cdot \sqrt{\log \Phi})$	[IMSZ10]		n, Φ
2.	General		$\Omega(\min\{\sqrt{n}, \sqrt{\log \Phi}\})$	[NR20]	yes	
3.	Doubling		$2^{O(\mathrm{ddim})} \cdot \log \Phi$	[IMSZ10]		ddim, Φ
4.	Doubling		$O(\mathrm{ddim} \cdot \sqrt{\log \Phi})$	Theorem 1.1	yes	none
5.	Constant ddim		$\Omega(\sqrt{\log \Phi})$	Theorem 1.2		
6.	General	HST (ultrametric)	$O(\log n \cdot \log \Phi)$	[IMSZ10, BFU20]		none
7.	General		$\tilde{\Omega}(\log n \cdot \log \Phi)$	[BFU20]		
8.	Doubling		$2^{O(\mathrm{ddim})} \cdot \log \Phi$	[IMSZ10]		ddim, Φ
9.	Doubling		$O(\mathrm{ddim} \cdot \log \Phi)$	Theorem 1.3		none
10.	$(\mathbb{R}^d, \|\cdot\|_2)$		$O(\sqrt{d} \cdot \log \Phi)$	Theorem 1.4		none
11.	$\mathbb{R}$		$\Omega(\min\{n, \log \Phi\})$	[IMSZ10]		
12.	General	Tree	2^{n-1}	[NR20]	yes	none
13.	General		$\Omega(2^{\frac{n}{2}})$	[NR20]	yes	

Table 1: Summary of new and previous result on online metric embeddings. Doubling stands for metric space with doubling dimension ddim. The column "Deter?" indicates whether the embedding is deterministic (in particular works against adaptive adversary), or the guarantee is only in expectation (in particular works only against oblivious adversary). The column "Pri.Kno" indicates what prior knowledge is required by the embedding (applicable only for the upper bounds).

Our third result is a stochastic embedding into ultrametrics with expected distortion $O(\mathrm{ddim} \cdot \log \Phi)$. This is a generalization of [BFU20] (as every metric space has doubling dimension $O(\log n)$), and an exponential improvement in the dependence on ddim compared with [IMSZ10].

THEOREM 1.3. *Given a sequence of metric points $x_1, x_2, \dots$ arriving in an online fashion, there is a stochastic metric embedding into an ultrametric (a 2-HST) with expected distortion $O(\mathrm{ddim} \cdot \log \Phi)$, where ddim, Φ are the doubling dimension and the aspect ratio of the metric space. No prior knowledge is required.*

REMARK 1. In fact, for a pair of points $\{x_j, x_k\}$ where $j < k$, the expected distortion guarantee provided by Theroem 1.3 is $O(\mathrm{ddim}_j) \cdot \log \Phi_j$, where ddim_j and Φ_j are the doubling dimension and aspect ratio of the metric space induced by the prefix $\{x_1, \dots, x_j\}$. This is also known as prioritized distortion. See [EFN18, BFN19, FGK20, EN22] for further details on prioritized distortion.

If the points arrive from Euclidean d-dimensional space, we obtain expected distortion $O(\sqrt{d} \cdot \log \Phi)$, which is a quadratic improvement in the dependence on the dimension.

THEOREM 1.4. *Given a sequence of points $x_1, x_2, \dots$ in Euclidean d-space $(\mathbb{R}^d, \|\cdot\|_2)$ arriving in an online fashion, there is a stochastic metric embedding into an ultrametric (a 2-HST) with expected distortion $O(\sqrt{d} \cdot \log \Phi)$, where Φ is the aspect ratio (unknown in advance).*

Bartal et al. [BFU20] used their stochastic online embedding to design competitive online algorithms for certain network design problems. Surprisingly, they showed that in many cases the dependence on the aspect ratio can be avoided. One can improve some parameters by pluging in our Theorem 1.3–1.4 into their framework. One example is the Subadditive Constrained Forest problem [GW95], where we can improve the competitive ratio from $O(\log^2 k)$ to $O(\mathrm{ddim} \cdot \log k)$, (or $O(\sqrt{d} \cdot \log k)$ for points in Euclidean d-space). See the full paper for further discussion.

1.1.2 Minimum Weight Perfect Matchings Our results on online minimum-weight perfect matchings are summarized in Table 2.

In the full paper, we design a randomized algorithm against an oblivious adversary that maintains a perfect matching with competitive ratio $O(\mathrm{ddim} \cdot \log \Phi)$ and recourse $O(\log \Phi)$.

Copyright © 2024
Copyright for this paper is retained by authors

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Adversary	Metric	Recourse	Approx. ratio	Approx. type	Reference
Adaptive	General	$O(\log^2 n)$	$O(\log n)$	lightness	Theorem 1.8
Oblivious	General	$O(\min\{\log^3 n, \log \Phi\})$	$O(\mathrm{ddim} \cdot \log \Phi)$	comp. ratio	Theorem 1.5
Oblivious	$(\mathbb{R}^d, \|.\|_2)$	$O(\min\{\log^3 n, \log \Phi\})$	$O(\sqrt{d} \cdot \log \Phi)$	comp. ratio	Theorem 1.6
Oblivious	2-HST	$O(\min\{\log^3 n, \log \Phi\})$	$O(1)$	comp. ratio	Lemma 1.1
Oblivious	General	r	$\Omega(\frac{\log n}{r \log r})$	comp. ratio & lightness	Theorem 1.7
Oblivious	General	1	∞	comp. ratio	Proposition 1.1

Table 2: Summary of our online algorithms and lower bounds for online minimum weight perfect matching with recourse. The input is n points in a metric space with aspect ratio Φ.

THEOREM 1.5. *There is a randomized algorithm that, for any sequence of metric points $x_1, \ldots, x_{2n}$ revealed by an oblivious adversary in an online fashion with aspect ratio Φ and doubling dimension* ddim *(both unknown in advance), maintains a perfect matching of expected competitive ratio $O(\mathrm{ddim} \cdot \log \Phi)$ with recourse $O(\log \Phi)$. Alternatively, the recourse can be bounded by $O(\log^3 n)$.*

Moreover, we show that the competitive ratio can be further improved to $O(\sqrt{d} \cdot \log \Phi)$ if the input points are from Euclidean d-space. Following the exact same lines, we conclude:

THEOREM 1.6. *There is a randomized algorithm such that given points $x_1, \ldots, x_{2n} \in \mathbb{R}^d$ revealed by an oblivious adversary in an online fashion with aspect ratio $\Phi = \frac{\max_{i,j} \|x_i - x_j\|_2}{\min_{i,j} \|x_i - x_j\|_2}$ (unknown in advance), maintains a perfect matching of expected competitive ratio $O(\sqrt{d} \cdot \log \Phi)$ with recourse $O(\log \Phi)$. Alternatively, the recourse can be bounded by $O(\log^3 n)$.*

Note that, every n-point metric space has doubling dimension $O(\log n)$. For example, for the shortest path metric of unweighted graphs, Theorem 1.5 provides a competitive ratio of $O(\log^2 n)$.

Theorem 1.5 and Theorem 1.6 are proven by a reduction to *hierarchically well-separated tree* (*HST*, a.k.a. ultrametric, see Definition 2) via Theorem 1.3 and Theorem 1.4, respectively. For an HST of height h, we can maintain a minimum-weight perfect matching (i.e., an optimum matching) using recourse $O(h)$. Using heavy-path decomposition, we can also maintain a $O(1)$-approximate minimum-weight matching with recourse $O(\log^3 n)$ (see Lemma 1.1).

It is clear that without any recourse, the competitive ratio for the online minimum-weight perfect matching is unbounded (see example in Section 1). Can we bound the competitive ratio if we allow one recourse per point pair? That is, for each new point pair, we allow the algorithm to delete one edge when it updates a perfect matching. For online MST, for example, Gu et al. [GGK16] achieve a competitive ratio of $\Omega(\frac{1}{k})$ with one recourse for every k new points.

PROPOSITION 1.1. *Given a single recourse per vertex pair, there is no competitive online algorithm. This already holds for a sequence of 8 points in $\mathbb{R}$, even if the sequence is known in advance.*

Next, we establish a lower bound using points on the real line (with linear aspect ratio) such that the recourse times the competitive ratio must be $\tilde{\Omega}(\log n)$. Note that for metric spaces with polynomial aspect ratio, our Theorem 1.5 and Theorem 1.6 are tight up to a quadratic factor. [8]

THEOREM 1.7. *For every $r \geq 2$, every online algorithm for minimum-weight perfect matching problem with recourse r, even for n points in the real line, has competitive ratio $\Omega\left(\frac{\log n}{r \cdot \log r}\right)$ against an oblivious adversary. Furthermore r can depend on n.*

[8] As we try to optimize simultaneously both the competitive ratio and the recourse, it is natural to define a new parameter called *performance*, which equals competitive ratio times recourse. Thus for metric space with polynomial aspect ratio and constant ddim, our Theorem 1.5 has performance $O(\log^2 n)$, while by Theorem 1.7, the performance is at least $\min_r \left\{r, \frac{\log n}{\log r}\right\} = \Omega(\frac{\log n}{\log \log n})$. Thus, ignoring second order terms, Theorem 1.5 is tight up to a quadratic factor.

Copyright © 2024
Copyright for this paper is retained by authors

Finally, we design a deterministic algorithm against an adaptive adversary that maintains a perfect matching of weight $O(\log n)\cdot\mathrm{Cost}(\mathrm{MST})$ with recourse $O(\log^2 n)$ in any metric space (Theorem 1.8). The *lightness* of weighted graph G on a point set is the ratio $\frac{\mathrm{Cost}(G)}{\mathrm{Cost}(\mathrm{MST})}$ of the weight of G to the weight of an MST, and it is a popular measure in network optimization. The lightness of a perfect matching may be arbitrarily close to zero, and it is always at most one[9]. However, for $2n$ uniformly random points in a unit cube $[0,1]^d$ for constant $d\in\mathbb{N}$, for example, the expected minimum weight of a perfect matching is proportional to the maximum weight of an MST [SRP83, SS89]. Interestingly, in our lower bound (Theorem1.7) the weight of the perfect matching is $\Theta(\mathrm{diam})=\Theta(\mathrm{MST})$. Thus, in particular, it implies that the product of the recourse and the lightness of any oblivious algorithm is $\tilde{\Omega}(\log n)$. Hence, our algorithm against an adaptive adversary is comparable to the best possible oblivious algorithm w.r.t. the lightness parameter.

THEOREM 1.8. *For a sequence of points in a metric space (X,d), we can maintain a perfect matching of weight $O(\log|S_i|)\cdot\mathrm{Cost}(\mathrm{MST}(S_i))$ using recourse $O(\log^2|S_i|)$ where S_i is the set of the first $2i$ points.*

1.2 Technical Ideas

Online Padded Decompositions and Online Embedding into HST. A (Δ,β)*-padded decomposition* of a metric space (X,d_X) is a random partition of X into clusters of diameter at most Δ such that for a ball $B=B_X(v,r)$ of radius r centered at v, the probability that the points of B are split between different clusters is at most $\beta\cdot\frac{r}{\Delta}$. A metric space is β*-decomposable* if it admits a (Δ,β)-padded decomposition for every $\Delta>0$. Building on previous work [Bar96, Rao99, KLMN04], the main ingredient in all our metric embeddings (in particular the deterministic Theorem 1.1) are padded decompositions. Every n-point metric space is $O(\log n)$-decomposable [Bar96], while every metric space with doubling dimension ddim is $O(\mathrm{ddim})$-decomposable [GKL03, Fil19]. Roughly speaking, Bartal's padded decomposition [Bar96] works using a ball growing technique: Take an arbitrary order over the metric points $x_1,\dots,x_n$, sample radii $R_1,\dots,R_n$ from exponential distribution with parameter $\Theta(\log n)$, and successively construct clusters $C_i=B_X(x_i,R_i\cdot\Delta)\setminus\bigcup_{j<i}C_j$. In words, there are n clusters centered in the points of X. Each point y joins the first cluster C_i such that $d_X(y,x_i)\le R_i\cdot\Delta$. With high probability, it holds that $\max R_i\le\frac{1}{2}$ and thus the diameter of all the resulting clusters is bounded by Δ. Further, consider a ball $B=B_X(v,r)$, and suppose that $u\in B$ is the first point to join a cluster C_i. Then $R_i\ge\frac{d_X(x_i,u)}{\Delta}$. By the triangle inequality, $R_i\ge\frac{d_X(x_i,u)}{\Delta}+2r$ will imply that all the points in B will join C_i. By the memoryless property of the exponential distribution, the probability that $R_i<\frac{d_X(x_i,u)}{\Delta}+2r$ is at most $O(\log n)\cdot\frac{2r}{\Delta}$.

An HST (hierarchically separated tree, see Definition 2) is in essence just a hierarchical partition. Bartal's embedding into HST then works by creating separating decompositions for all possible distance scales $\Delta_i=2^i$, $i\in\mathbb{Z}$. That is, each i-level cluster is partitioned into $(i-1)$-level clusters using the padded decomposition described above. Let i be the first scale such that u and v were separated, the distance between them in the HST will be 2^{i+1}. Thus the expected distance between u,v is bounded by[10]

$$\sum_{i\ge\log d_X(u,v)}2^{i+1}\cdot\Pr[u,v\text{ are separated at level }i]\le\sum_{i\ge\log d_X(u,v)}2^{i+1}\cdot O(\log n)\cdot\frac{d_X(u,v)}{2^i}\le O(\log\Phi\cdot\log n)\,. \tag{1.1}$$

Indyk et al. [IMSZ10] observed that, given metric points $x_1,\dots,x_n$ in an online fashion, we can still preform Bartal's padded decomposition, since the order of the points were arbitrary. However, [IMSZ10] required prior knowledge of n and the aspect ratio Φ (in order to determine the parameter of the exponential distribution, and the relevant scales). Later, Bartal et al. [BFU20] observed that sampling R_j using exponential distribution with parameter $O(\log j)$ will still return a padded decomposition, and thus removed the requirement to know n in

[9]The weight of an MST is at least as large as the weight of the minimum perfect matching. Indeed, following the approach of Christofides algorithm, double each edge of the MST to obtain an Euler cycles C of weight $2\,\mathrm{Cost}(\mathrm{MST})$. Let $x_1,\dots,x_{2n}$ be the order of the points in the order of their first occurrence along the tour. The two matchings $M_1=(x_1,x_2),\dots,(x_{2n-1},x_{2n})$ and $M_2=(x_2,x_3),\dots,(x_{2n-2},x_{2n-1}),(x_{2n},x_1)$ combined have weight at most $\mathrm{Cost}(C)=2\,\mathrm{Cost}(\mathrm{MST})$. In particular, the weight of the minimum perfect matching is at most $\mathrm{Cost}(\mathrm{MST})$.

[10]Bartal [Bar96] obtains expected distortion $O(\log^2 n)$ by contracting all pairs at distance at most $\frac{\Delta_i}{\mathrm{poly}(n)}$ before preforming the decomposition at scale i. The effect of this contraction on pairs at distance $\tilde{\Theta}(\Delta_i)$ is negligible, while the contraction ensures that u,v have nonzero probability of being separated in only $O(\log n)$ different scales.

Copyright © 2024
Copyright for this paper is retained by authors

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

advance. To remove the requirement to know Φ in advance, [BFU20] simply forced R_1 to be $\Omega(1)$, and thus ensuring that all the partitions above a certain threshold are trivial.

Gupta, Krauthgamer, and Lee [GKL03] showed that every metric space with doubling dimension ddim is $O(\mathrm{ddim})$-decomposable. Their decomposition follows the approach from [CKR04], which samples a global permutation to decide where to cluster each point. Later, Filtser [Fil19] used the random shifts clustering algorithm of Miller, Peng, and Xu [MPX13] to obtain a similar decomposition with strong diameter guarantee.[11] However, both these decompositions are crucially global and centralized, and it is impossible to execute them in an online fashion. Indyk et al. [IMSZ10] studied online embeddings of doubling metrics into ultrametrics. However, lacking good padded decompositions for doubling spaces, they ended up using a similar partition based approach, which lead to an expected distortion $2^{O(\mathrm{ddim})} \cdot \log \Phi$.

We show that one can construct a padded decomposition with padding parameter $O(\mathrm{ddim})$ using the ball growing approach of Bartal [Bar96]: Sample the radii using an exponential distribution with parameter $O(\mathrm{ddim})$. This is crucial, as such a decomposition can be executed in an online fashion. Furthermore, one does not need to know the doubling dimension in advance. It is enough to use the doubling dimension of the metric space induced on the points seen so far. Interestingly, even if the doubling dimension eventually will turn out to be $O(\log n)$, the decomposition will have the optimal $O(\log n)$ parameter. Using these decomopsitions, we construct an HST in an online fashion and obtain Theorem 1.3 by replacing the $O(\log n)$ factor in inequality (1.1) by $O(\mathrm{ddim})$.

Online Deterministic Embedding into Euclidean Space. Bourgain's [Bou85] optimal embedding into Euclidean space with distortion $O(\log n)$ is a Fréchet type embedding. Specifically, it samples subsets uniformly with different densities, and sets each coordinate to be equal to the distance to a certain sampled subset. This is a global, centralized approach that cannot be executed in an online fashion. In contrast, Rao's [Rao99] classic embedding of the shortest path metric of planar graphs into ℓ_2 with distortion $O(\sqrt{\log n})$ is based on padded decompositions. Indyk et al. [IMSZ10] followed the padded decomposition based approach of Rao [Rao99]. Roughly speaking, Rao's approach is to create padded decompositions $\mathcal{P}$ for all possible distance scales, where we have a distinct coordinate for each partition $\mathcal{P}$. For every cluster $C \in \mathcal{P}$, assign a random coefficient $\alpha_C \in \{\pm 1\}$. Finally, for every $x \in C$, assign value $\alpha_C \cdot h_C(x)$, where $h_C(x)$ is the distance from x to the "boundary" of C. One can show that the distance in every coordinate is never expanding, while for two points x, y, in the scale $d_X(x,y)$, with constant probability, x, y will be separated, and x will be at least $\Omega(\frac{d_X(x,y)}{\log n})$ away from the boundary of its cluster—thus we will get some contribution to the distance. By sampling many such decompositions, one can get concentration, and thus an embedding with distortion $O(\log n \cdot \log \Phi)$ against an oblivious adversary. Indyk et al. did not suggest any way to cope with an adaptive adversary. Newman and Rabinovich [NR20] provided an $\Omega(\sqrt{n})$ lower bound for such an embedding, and conjectured that $\mathrm{poly}(n)$ distortion should always suffice. However, they did not suggest any way to achieve it.

Next Indyk et al. [IMSZ10] moved to embedding doubling spaces. Lacking good online padded decompositions, they observed that an isometric embedding of ultrametric into ℓ_2 can be maintained in an online fashion, and thus getting expected distortion $2^{O(\mathrm{ddim})} \cdot \log \Phi$ against an oblivious adversary. Plugging in our new padded decomposition into Rao's approach, one can get (worst case w.h.p.) distortion $O(\mathrm{ddim} \cdot \sqrt{\log \Phi})$ against an oblivious adversary. But how can one construct deterministic embedding to cope with an adaptive adversary?

Our solution is to create a "layer of abstraction". The previous ideas provide an embedding with expected distortion $O(\mathrm{ddim} \cdot \sqrt{\log \Phi})$. Specifically, we get expansion $O(\sqrt{\log \Phi})$ in the worst case, and contraction $O(\mathrm{ddim})$ in expectation. The only randomness is over the choice of radii in the ball growing that creates the padded decompositions (and some additional boolean parameters). Given a new point x_i, the expected squared distance $\mathbb{E}_f[\|f(x_i) - f(x_j)\|_2^2]$ can be computed exactly, as it only depends on the points that have arrived so far, with no randomness involved. In a sense, instead of mapping a metric point x_i into a vector in ℓ_2, we map it into a well-defined function $f_i : (r_1, r_2, \dots, r_i) \to \ell_2$. These functions are in the function space L_2, and the distance

$$\|f_i - f_j\|_2 = \left(\int_{r_1,\dots,r_i} \|f_i(r_1,\dots,r_i) - f_j(r_1,\dots,r_j)\|_2^2 \right)^{\frac{1}{2}}$$

[11] Given an edge-weighted graph $G = (V, E, w)$ and a cluster $C \subseteq V$, the *(weak) diameter* of C is $\max_{u,v \in C} d_G(u,v)$ the maximum pairwise distance in C w.r.t. the shortest path metric of G. The *strong diameter* of C is $\max_{u,v \in C} d_{G[C]}(u,v)$ the maximum pairwise distance in C w.r.t. the shortest path metric of the induced subgraph $G[C]$.

Copyright © 2024
Copyright for this paper is retained by authors

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

equals the expected distance by the random metric embedding. However, we want to return vectors and not complicated functions. In fact, the only required information is the L_2 distance between these functions, which define an Euclidean distance matrix. Given such a matrix, one can find a set of vectors implementing it (which is unique up to rotation and translation). Furthermore, these vectors can be efficiently and deterministically computed in an online fashion!

Online Minimum-Weight Perfect Matchings in Metric Spaces. Given a sequence of metric points $s_1, \ldots, s_{2i}$ in an online fashion from an unknown metric space, we use the online embedding algorithm to embed them into an ultrametric or the real line (with some distortion), and maintain a matching with recourse on the embedded points. If we can maintain a good approximation for the online MWPM in an ultametric (or in $\mathbb{R}$), then we can maintain the same approximation ratio, with the distortion of the embedding as an overhead, for the online MWPM problem.

Optimal Matchings on Trees: Inward Matchings. An ultrametric is represented by an hierarchically well-separated tree T, a rooted tree with exponentially decaying edge weights. As points arrive in an online fashion, the online embedding algorithm may successively add new leaves to T; and the points are embedded in the nodes of T. We show that a simple greedy matching is optimal (i.e., has minimum-weight) in an ultrametric, and can easily be updated with recourse proportional to the height of T. Specifically, an *inward matching*, introduced here, maintains the invariant that the points in each subtree induce a near-perfect matching. When a pair of new points arrive, we can restore this property by traversing the shortest paths between the corresponding nodes in T. Consequently, an inward matching can be maintained with recourse $O(h)$, where h is the height of T. We have $h \le O(\log \Phi_U)$, where Φ_U is the aspect ratio of the ultrametric U, which in turn is bounded by $O(\log \Phi_X)$, the aspect ratio of the metric space X induced by the input points seen so far.

Heavy-Path Decomposition on HSTs. Bartal et al. [BFU20] showed that the distortion of any online metric embedding algorithm into trees depends on the aspect ratio Φ, in particular, the factor $\log \Phi$ in our distortion bounds in Theorems 1.3–1.4 is unavoidable. It is unclear whether any dependence on Φ is necessary for the bounds for the online MWPM problem. We can eliminate the dependence on Φ for the recourse, while maintaining the same approximation guarantee (which, however, still depends on the distortion, hence on Φ).

Instead of an optimal matching on the HST T, we maintain an 2-approximate minimum-weight perfect matching. We use the classical *heavy-path decomposition* of the tree T, due to Sleator and Tarjan [ST83], which is a partition of the vertices into subsets that each induce a path (*heavy path*); the key property is that every path in T intersects only $O(\log n)$ heavy paths, regardless of the height of h. The heavy path decomposition can be maintained dynamically with $O(\log n)$ split-merge operations over the paths.

We relax the definition of inward matchings such that at most one edge of the marching can pass between any two adjacent heavy paths, but we impose only mild conditions within each heavy path. A charging scheme shows that the relaxed inward matching is a 2-approximation of the MWPM.

LEMMA 1.1. *In k-HST with n nodes, where new nodes are added in an online fashion, one can maintain an 2-approximate minimum-weight perfect matching with recourse $O(\log^3 n)$.*

On each heavy path, we maintain a matching designed for points on a real line, with $O(\log^2 n)$ depth[12], which supports split-merge operations in $O(\log^2 n)$ changes in the matching (i.e., recourse). Overall, we can maintain a 2-approximate minimum-weight matching on T with worst-case recourse $O(\log^3 n)$.

Minimum-Weight Matching on a Real Line: Reduction to Depth. We reduce the online MWPM problem to a purely combinatorial setting. For a set of edges E on a finite set $S \subset \mathbb{R}$, we say that E is *laminar* if there are no two edges a_1b_1 and a_2b_2 such that $a_1 < a_2 < b_1 < b_2$ (i.e., no two *interleaving* or *crossing* edges). Containment defines a partial order: We say that $a_1b_1 \le a_2b_2$ (resp., $a_1b_1 \prec a_2b_2$) if the interval a_2b_2 contains (resp., properly contains) the interval a_1b_1. The Hasse diagram of a laminar set E of edges is a forest of rooted trees $F(E)$ on E, where a directed edge (a_1b_1, a_2b_2) in $F(E)$ means that a_2b_2 is the shortest interval that strictly contains a_1b_1. The *depth* of E is the depth of the forest $F(E)$; equivalently, the depth of E is the maximum number of pairwise overlapping edges in E. Based on this, we show that for a dynamic point set S on $\mathbb{R}$, one can maintain a laminar near-perfect matching of depth $O(\log n)$ such that it modifies (adds or deletes) at most $O(\log^2 n)$ edges in each step.

[12] The *depth* of a matching on n points in $\mathbb{R}$ (or a path) is the maximum number of pairwise overlapping edges.

Copyright © 2024
Copyright for this paper is retained by authors

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Importantly, the laminar property and the depth of the matching depend only on the order of the points in S, and the real coordinates do not matter. While it is not difficult to maintain a laminar near-perfect matching. However, controlling the depth is challenging. We introduce the notion of *virtual edges*, which is the key technical tool for maintaining logarithmic depth. We maintain a set of invariants that ensure that, for a nested sequence of edges yields a nested sequence of virtual edges with the additional property that they have exponentially increasing lengths. We argue that if a near-perfect matching with virtual edges satisfies the invariants then the depth of such matching is logarithmic.

We reduce the case of general metric space to a line metric using a result by Gu et al. [GGK16]: Given a sequence of metric points, one can maintain a spanning tree of weight $O(\mathrm{Cost}(\mathrm{MST}))$ in an online fashion (insertion only), with constant recourse per point. We maintain an Euler tour[13] $\mathcal{E}(T)$ for the tree T produced by their algorithm. The Euler tour $\mathcal{E}(T)$ induces a Hamilton path $\mathcal{P}(T)$ (according to the order of first appearance). Intuitively, we treat the metric points as if they were points on a line ordered according to $\mathcal{P}(T)$. We show that each edge deletion and each edge insertion in the forest F incurs $O(1)$ edge insertions or deletions in the tour $\mathcal{E}(T)$ and path $\mathcal{P}(T)$, a very limited change! Thus we can use our data structure for the line to maintain a near-perfect matching for the points w.r.t. $\mathcal{P}(T)$. As the total weight of this path is $O(\mathrm{Cost}(\mathrm{MST}))$, we obtain a perfect matching of weight $O(\log n) \cdot \mathrm{Cost}(\mathrm{MST})$ using recourse $O(\log^2 n)$.

Lower Bounds for Competitive Ratio and Recourse on Minimum-Weight Matchings. For integers $r \geq 2$ and $n \geq 10r$, we show that an adaptive adversary can construct a sequence S_n of integer points on the real line such that any deterministic online perfect matching algorithm with recourse r per arrival maintains a perfect matching of weight $\Omega\left(\frac{\mathrm{diam}(S_n)\log n}{r \log r}\right)$. The adversary presents points in $k = \Theta\left(\frac{\log n}{r \log r}\right)$ rounds. In round 0, the points are consecutive integers. In subsequent rounds, the number of points decreases exponentially, but the spacing between them increases. If the algorithm matches new points among themselves in every round, the weight of the resulting matching would be $\Omega(k \cdot \mathrm{OPT})$. The weight could be improved with recourse, however, the number of new points rapidly decreases, and they do not generate enough recourse to make amends: We show that the weight increases by $\Omega(\mathrm{OPT})$ in every round. There is one twist in the adversarial strategy, which makes it adaptive: If the matching M_{i-1} at the beginning of round i could possibly "absorb" the points of round i (in the sense that the weight would not increase by $\Omega(\mathrm{OPT})$), then we show that M_{i-1} already contains many long edges and $\mathrm{Cost}(M_{i-1}) \geq \Omega(i \cdot \mathrm{OPT})$: In this case, the adversary can simply skip the next round.

In fact, this lower bound construction extends to oblivious adversaries by skipping some of the rounds randomly. Moreover, since, for a set of points in the real line, the minimum weight of a perfect matching is trivially bounded by the diameter of the point set, we conclude that the lightness and competitive ratio of any online algorithm with constant recourse is $\Omega(\log n)$; and an $O(1)$-competitive algorithm would require recourse at least $r = \Omega(\log n / \log\log n)$.

1.3 Related Work

Online Minimum-Weight Perfect Matching with Delays. Similarly to our model, Emek et al. [EKW16] considered online minimum-weight perfect matchings in a metric space. However, they allow *delays* instead of recourse: The decisions of the online matching algorithm are irrevocable, but may be delayed, incurring a time penalty of t_i if a point s_i remains unmatched for t_i units of time. The objective is to minimize the *sum* of the weight and all time penalties. Emek et al. [EKW16] show that a randomized algorithm (against an oblivious adversary) can achieve a competitive ratio $O(\log^2 n + \log \Phi)$ in this model, where Φ is the aspect ratio of the metric space (which can be unbounded as a function of n). Later, Azar et al. [ACK17], improved the competitive ratio to $O(\log n)$. Ashlagi et al. [AAC+17] studied the bipartite version of this problem; and Mari et al. [MPRS23] considered the stochastic version of this problem where the input requests follow Poisson arrival process. Recently, Deryckere and Umboh [DU23] initiated the study of online problems with set delay, where the delay cost at any given time is an arbitrary function of the set of pending requests. However, time penalties cannot be directly compared to the recourse model. An advantage of the recourse model is that it allows us to maintain perfect matching explicitly at all times, as opposed to the delay model where some points might remain unmatched at every step.

Online algorithms with recourse have been studied extensively over the years; see [IW91, MSVW16, GGK16,

[13] A DFS traversal of a tree T (starting from an arbitrary root) defines an *Euler tour* $\mathcal{E}(T)$ that traverses every edge of T precisely twice (once in each direction).

Copyright © 2024
Copyright for this paper is retained by authors

GGK+22, BGW21, BHR19]. The question is, given the power of hindsight, how much one can improve the solution of an online algorithm [GGK16].

Online MST with Recourse. In the online minimum spanning tree (MST) problem, points in a metric space arrive one by one, and we need to connect each new point to a previous point to maintain a spanning tree. Without recourse, Imase and Waxman [IW91] showed that a natural greedy algorithm is $O(\log n)$-competitive, and this bound is the best possible (see also [AA93, DT07]). They also showed how to maintain a 2-competitive tree with recourse $O(n^{3/2})$ over the first n arrivals for every n. Therefore, the amortized budget, i.e., the average number of swaps per arrival, is $O(\sqrt{n})$. Later, Megow et al. [MSVW16] substantially improved this result. They gave an algorithm with a constant amortized budget bound. In a breakthrough, Gu et al. [GGK16] showed that one can maintain a spanning tree of weight $O(\mathrm{Cost}(\mathrm{MST}))$ with recourse $O(1)$ per point.

When new points arrive in the online model, the weight of the MST may decrease. However, it cannot decrease by a factor more than $\frac{1}{2}$. Indeed, the decrease is bounded by the *Steiner ratio* [GP68], which is the infimum of the ratio between the weight of a Steiner tree and the MST for a finite point set, and is at least $\frac{1}{2}$ in any metric space. Similarly, in the online traveling salesman problem (TSP), where the length of the optimal TSP tour increases monotonically as new points arrive, one can maintain an $O(1)$-competitive solution with constant recourse (see Section 1.3).

Online TSP. In the online traveling salesman problem (TSP), points of a metric space arrive one by one, and we need to maintain a traveling salesman tour (or path) including the new point. Rosenkrantz et al. [RSI77] showed that a natural greedy algorithm with *one* recourse per point insertion (replacing one edge by two new edges) is $O(\log n)$ competitive, and there is a lower bound of $\Omega(\log n/\log\log n)$ even in Euclidean plane [Aza94, BKP94]. However, as an Euler tour around an MST 2-approximates the weight of a TSP tour, the online MST algorithm by Gu et al. [GGK16] immediately yields a $O(1)$-competitive algorithm with recourse $O(1)$.

Kalyanasundaram and Pruhs [KP94] studied a variant of online TSP, where new cities are revealed locally during the traversal of a tour (i.e., an arrival at a city reveals any adjacent cities that must also be visited). Jaillet and Lu [JL11] studied the online TSP with service flexibility, where they introduced a sound theoretical model to incorporate "yes-no" decisions on which requests to serve, together with an online strategy to visit the accepted requests.

Greedy Matchings. Online algorithms with or without recourse often make greedy choices. For the online MWPM, the following greedy approach with constant recourse seems intuitive: Suppose points p_1 and p_2 arrive when our current matching on S_i is M_i. Then we find a closest neighbor for p_1 and p_1, resp., say a_1 and a_2 in S_{i+1}, delete any current edges $a_1b_1, a_2b_2 \in M_i$, and add all edges of a minimum-weight matching on $\{a_1, a_2, b_1, b_2, p_1, p_2\}$ to the matching.

An online greedy approach would, at best, "approximate" an offline greedy solution. The *offline greedy* algorithm successively adds an edge ab between the closest pair of vertices and removes both a and b from further consideration. Reingold and Tarjan [RT81] showed, however, that the greedy algorithm on $2n$ points in a metric space achieves an $O(n^{\log\frac{3}{2}})$-approximation, where $\log\frac{3}{2} \approx 0.58496$, and this bound is the best possible already on the real line. Frieze, McDiarmid, and Reed [FMR90] later showed that for integers $S_n = \{1, 2, \ldots, 2n\} \subset \mathbb{R}$, the offline greedy algorithm returns an $O(\log n)$-approximation, and this bound is tight (if ties are broken arbitrarily when multiple point pairs attain the minimum distance).

Metric Embeddings. There is a vast literature on metric embeddings that we will not attempt to cover here. We refer to the extended book chapter [Mat13], and some of the recent papers for an overview [AFGN22, FGK20, Fil21, FL21]. See also the recent FOCS22 workshop. In the context of online embeddings, embeddings into low dimensional ℓ_1, ℓ_2, and ℓ_∞ normed spaces were studied [IMSZ10, NR20]. In particular, every tree metric admits an isometric (with distortion 1) online embedding into ℓ_1 [NR20]. A significant part of the metric embeddings literature is concerned with the embedding of topologically restricted metric spaces, such as planar graphs, minor free graphs, and graphs with bounded treewidth/pathwidth [Rao99, KLMN04, Fil20, FKS19, CFKL20, FL22]. However, at present, these embeddings do not have online counterparts. The reason is perhaps the lack of a good online version of a padded decompositions for such spaces [KPR93, FT03, AGG+19, Fil19]. Designing online padded decompositions for such spaces is a fascinating open problem.

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Copyright © 2024
Copyright for this paper is retained by authors

2 Preliminaries

Ultrametrics. An ultrametric (X, d) is a metric space satisfying a strong form of the triangle inequality, that is, for all $x, y, z \in X$, $d(x, z) \leq \max\{d(x, y), d(y, z)\}$. A related notion is a k-hierarchically well-separated tree (k-HST).

DEFINITION 2. (k-HST) *A metric* (X, d_X) *is a* k-hierarchically well-separated tree *(*k-HST*) if there exists a bijection* φ *from* X *to leaves of a rooted tree* T *in which:*

1. *each node* $v \in T$ *is associated with a label* Γ_v *such that* $\Gamma_v = 0$ *if* v *is a leaf, and* $\Gamma(v) \geq k\Gamma(u)$ *if* v *is an internal node and* u *is any child of* v*;*

2. $d_X(x, y) = \Gamma(\mathrm{lca}(\varphi(x), \varphi(y)))$ *where* $\mathrm{lca}(u, v)$ *is the least common ancestor of any two given nodes* u, v *in* T*.*

It is well known that any ultrametric is a 1-HST, and any k-HST is an ultrametric (see [BLMN05]). Note that a 1-HST induces a laminar partition of the metric space. For an internal node $v \in T$, denote by X_v the set of leaves in the subtree rooted at v. We will use the terms ultrametric and 1-HST interchangeably throughout the paper.

Doubling Dimension. The doubling dimension of a metric space is a measure of its local "growth rate". A metric space (X, d) has doubling constant λ if for every $x \in X$ and radius $r > 0$, the ball $B(x, 2r)$ can be covered by λ balls of radius r. The doubling dimension is defined as $\mathrm{ddim} = \log_2 \lambda$. A d-dimensional ℓ_p space has $\mathrm{ddim} = \Theta(d)$, and every n point metric has $\mathrm{ddim} = O(\log n)$. Even though it is NP-hard to compute the doubling dimension of a metric space [GK13], in polynomial time, one can compute an $O(1)$-approximation [HM06, Theorem 9.1]. The following lemma gives the standard packing property of doubling metrics (see, e.g., [GKL03, Proposition 1.1].).

LEMMA 2.1. (PACKING PROPERTY) *Let* (X, d) *be a metric space with doubling dimension* ddim*. If* $S \subseteq X$ *is a subset of points with minimum interpoint distance* r *that is contained in a ball of radius* R*, then* $|S| \leq 2^{\mathrm{ddim} \cdot \lceil \log \frac{2R}{r} \rceil}$.

3 Conclusions

We introduced the problem of online minimum-weight perfect matchings for a sequence of $2n$ points in a metric space. In contrast to the online MST and TSP, where $O(1)$-competitive algorithms with recourse $O(1)$ are known, we showed that the competitive ratio *or* the recourse must be at least polylogarithmic. We also devised polylogarithmic upper bounds for the competitive ratio and lightness, resp., against oblivious and adaptive adversaries, using polylogarithmic recourse. Closing the gaps between the upper and lower bounds are obvious open problems, both in general metrics and in special cases such as Euclidean spaces or in ultrametrics. We highlight a few specific open problems.

1. We have shown (Propositiom 1.1) that recourse $r = 1$ per point pair is not enough for a competitive algorithm. Is there a competitive algorithm with recourse $r = O(1)$?

2. What is the minimum recourse for an $O(1)$-competitive algorithm against an adaptive (resp. oblivious) adversary? Our Theorem 1.7 gives a lower bound of $\Omega(\log n / \log\log n)$, but we are unaware of any nontrivial upper bound.

3. Are the optimal trade-offs different for competitive ratio and for lightness? Does it take more recourse to maintain a perfect matching of weight $O(\varrho \cdot \mathrm{OPT})$ than one of weight $O(\varrho \cdot \mathrm{MST})$ for any ratio $\varrho \geq 1$? Our lower bound (Theorem 1.7) does not distinguish between competitive ratio and lightness, but in general the ratio $\frac{\mathrm{MST}}{\mathrm{OPT}}$ is unbounded.

4. For maintaining a minimum-weight near-perfect matching on a fully dynamic point set (with insertions *and* deletions), what are the best possible trade-offs between the approximation ratio (or lightness) and the number of changes in the matching? Our 1D data structure can handle both point insertions and deletions, and maintains a matching of lightness $O(\log n)$, but the problem remains open in other metric spaces.

5. Consider the adversarial model, sometimes called *prefix-model*, which is weaker than the oblivious model. Here, the metric space (X, δ) and the entire sequence of arriving points $x_1, x_2, \ldots, x_{2n}$ are known in

Copyright © 2024
Copyright for this paper is retained by authors

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

advance. The goal is to construct a sequence of n matchings $M_1, M_2, \ldots, M_n$, where M_i is a perfect matching for $S_i = \{x_1, \ldots, x_{2i}\}$, while minimizing the competitive ratio $\max_i \frac{\mathrm{Cost}(M_i)}{\mathrm{OPT}(S_i)}$, and the maximum recourse $\max_i |M_i \setminus M_{i+1}|$. Note that our lower bound from Proposition 1.1 holds in this model, as well, and thus a single recourse is not enough for a competitive algorithm. Clearly, we can use the same algorithm we used for oblivious routing (Theorem 1.5), while using the embedding of [FRT04] instead of [BFU20], and thus improving the competitive ratio to $O(\log n)$. Is it possible to further improve on either the competitive ratio or the recourse?

6. Online Decomposition of Minor-Free Graphs. Consider the shortest path metric of a fixed minor-free graph (e.g., planar graphs). Such metrics enjoy a good padded decomposition scheme [KPR93, AGG+19, Fil19], and therefore also better embeddings into Euclidean space, compared to general metric spaces [Rao99]. However, no online version of such a decomposition is known, even for subfamilies such as planar graphs, and bounded treewidth/pathwidth graphs. Note that, given such a decomposition, the framework in this paper will imply good online metric embeddings into both HST and Euclidean spaces. Thus we find the construction of such decompositions to be a fascinating open problem. One might be tempted to think that the KPR [KPR93] decomposition can be easily implemented in an online fashion. This would make sense as the choice of a center in every ring in the decomposition is arbitrary. However, unfortunately only the first rings/chops are constructed w.r.t. the original metric. All the other steps are preformed w.r.t. induced subgraphs. It is unclear how to obtain distances, or even an approximation in an induced subgraph defined by a chop in an online fashion.

7. Krauthgamer et al. [KLMN04] showed that every metric space with doubling dimension ddim can be embedded into Euclidean space with distortion $O(\sqrt{\mathrm{ddim} \cdot \log n})$; and this bound is known to be tight [JLM11]. In contrast, our online embedding has distortion $O(\mathrm{ddim} \cdot \sqrt{\log \Phi})$. In Theorme 1.2 we showed that the dependence on the aspect ratio Φ is tight. It would be interesting to see whether the dependence on the doubling dimension can be improved to $\sqrt{\mathrm{ddim}}$.

References

[AA93] Noga Alon and Yossi Azar. On-line Steiner trees in the Euclidean plane. *Discrete & Computational Geometry*, 10:113–121, 1993. doi:10.1007/BF02573969. 11

[AAC+17] Itai Ashlagi, Yossi Azar, Moses Charikar, Ashish Chiplunkar, Ofir Geri, Haim Kaplan, Rahul Makhijani, Yuyi Wang, and Roger Wattenhofer. Min-cost bipartite perfect matching with delays. In *Proc. Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, (APPROX/RANDOM)*, volume 81 of *LIPIcs*, pages 1:1–1:20. Schloss Dagstuhl, 2017. doi:10.4230/LIPIcs.APPROX-RANDOM.2017.1. 10

[AAF+97] James Aspnes, Yossi Azar, Amos Fiat, Serge Plotkin, and Orli Waarts. On-line routing of virtual circuits with applications to load balancing and machine scheduling. *Journal of the ACM (JACM)*, 44(3):486–504, 1997. doi:10.1145/258128.258201. 2

[ACK17] Yossi Azar, Ashish Chiplunkar, and Haim Kaplan. Polylogarithmic bounds on the competitiveness of min-cost perfect matching with delays. In *Proc. 28th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1051–1061, 2017. doi:10.1137/1.9781611974782.67. 2, 10

[AFGN22] Ittai Abraham, Arnold Filtser, Anupam Gupta, and Ofer Neiman. Metric embedding via shortest path decompositions. *SIAM J. Comput.*, 51(2):290–314, 2022. doi:10.1137/19m1296021. 11

[AGG+19] Ittai Abraham, Cyril Gavoille, Anupam Gupta, Ofer Neiman, and Kunal Talwar. Cops, robbers, and threatening skeletons: Padded decomposition for minor-free graphs. *SIAM J. Comput.*, 48(3):1120–1145, 2019. doi:10.1137/17M1112406. 11, 13

[Alb03] Susanne Albers. Online algorithms: A survey. *Mathematical Programming, Ser. B*, 97:3–26, 2003. doi:10.1007/s10107-003-0436-0. 2

[AS22] Susanne Albers and Sebastian Schubert. Tight bounds for online matching in bounded-degree graphs with vertex capacities. In *Proc. 30th European Symposium on Algorithms (ESA)*, volume 244 of *LIPIcs*, pages 4:1–4:16. Schloss Dagstuhl, 2022. doi:10.4230/LIPIcs.ESA.2022.4. 2

[Aza94] Yossi Azar. Lower bounds for insertion methods for TSP. *Comb. Probab. Comput.*, 3:285–292, 1994. doi:10.1017/S096354830000119X. 11

Copyright © 2024
Copyright for this paper is retained by authors

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

[Aza05] Yossi Azar. On-line load balancing. In Amos Fiat and Gerhard J. Woeginger, editors, *Online algorithms: The state of the art*, volume 1442 of *LNCS*, pages 178–195. Springer, Heidelberg, 2005. doi:10.1007/BFb0029569. 2

[Bar96] Yair Bartal. Probabilistic approximations of metric spaces and its algorithmic applications. In *Proc. 37th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 184–193, 1996. doi:10.1109/SFCS.1996.548477. 3, 7, 8

[Bar98] Yair Bartal. On approximating arbitrary metrices by tree metrics. In Jeffrey Scott Vitter, editor, *Proc. 30th ACM Symposium on the Theory of Computing (STOC)*, pages 161–168, 1998. doi:10.1145/276698.276725. 3

[Bar04] Yair Bartal. Graph decomposition lemmas and their role in metric embedding methods. In *Proc. 12th European Symposium on Algorithms (ESA)*, volume 3221 of *LNCS*, pages 89–97. Springer, 2004. doi:10.1007/978-3-540-30140-0_10. 3

[BBMN15] Nikhil Bansal, Niv Buchbinder, Aleksander Madry, and Joseph Naor. A polylogarithmic-competitive algorithm for the k-server problem. *J. ACM*, 62(5):40:1–40:49, 2015. doi:10.1145/2783434. 3

[BCR23] Sébastien Bubeck, Christian Coester, and Yuval Rabani. The randomized k-server conjecture is false! In *Proc. 55th ACM Symposium on Theory of Computing (STOC)*, pages 581–594, 2023. doi:10.1145/3564246.3585132. 2

[BE98] Allan Borodin and Ran El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998. 1

[BFK+17] Joan Boyar, Lene M Favrholdt, Christian Kudahl, Kim S Larsen, and Jesper W Mikkelsen. Online algorithms with advice: A survey. *ACM Computing Surveys (CSUR)*, 50(2):1–34, 2017. doi:10.1145/3056461. 2

[BFN19] Yair Bartal, Arnold Filtser, and Ofer Neiman. On notions of distortion and an almost minimum spanning tree with constant average distortion. *J. Comput. Syst. Sci.*, 105:116–129, 2019. doi:10.1016/j.jcss.2019.04.006. 5

[BFU20] Yair Bartal, Nova Fandina, and Seeun William Umboh. Online probabilistic metric embedding: A general framework for bypassing inherent bounds. In *Proc. 31st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1538–1557, 2020. doi:10.1137/1.9781611975994.95. 3, 5, 7, 8, 9, 13

[BGW21] Sayan Bhattacharya, Fabrizio Grandoni, and David Wajc. Online edge coloring algorithms via the nibble method. In *Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2830–2842, 2021. doi:10.1137/1.9781611976465.168. 11

[BHR19] Aaron Bernstein, Jacob Holm, and Eva Rotenberg. Online bipartite matching with amortized $o(\log^2 n)$ replacements. *Journal of the ACM (JACM)*, 66(5):1–23, 2019. doi:10.1145/3344999. 11

[BKP94] Vineet Bafna, Bala Kalyanasundaram, and Kirk Pruhs. Not all insertion methods yield constant approximate tours in the euclidean plane. *Theor. Comput. Sci.*, 125(2):345–360, 1994. doi:10.1016/0304-3975(94)90257-7. 11

[BLMN05] Yair Bartal, Nathan Linial, Manor Mendel, and Assaf Naor. Some low distortion metric Ramsey problems. *Discret. Comput. Geom.*, 33(1):27–41, 2005. doi:10.1007/s00454-004-1100-z. 12

[BN09] Niv Buchbinder and Joseph Seffi Naor. The design of competitive online algorithms via a primal–dual approach. *Foundations and Trends in Theoretical Computer Science*, 3(2–3):93–263, 2009. doi:10.1561/0400000024. 2

[Bou85] J. Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel Journal of Mathematics*, 52(1):46–52, Mar 1985. doi:10.1007/BF02776078. 3, 8

[CFKL20] Vincent Cohen-Addad, Arnold Filtser, Philip N. Klein, and Hung Le. On light spanners, low-treewidth embeddings and efficient traversing in minor-free graphs. In *Proc. 61th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 589–600, 2020. See: conference version, arXiv version,. 11

[CKR04] Gruia Călinescu, Howard J. Karloff, and Yuval Rabani. Approximation algorithms for the 0-extension problem. *SIAM J. Comput.*, 34(2):358–372, 2004. doi:10.1137/S0097539701395978. 8

[DH09] Nikhil R. Devanur and Thomas P. Hayes. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *Proc. 10th ACM Conference on Electronic Commerce*, pages 71–78, 2009. doi:10.1145/1566374.1566384. 2

[DT07] Adrian Dumitrescu and Csaba D. Tóth. Analysis of two sweep-line algorithms for constructing spanning trees and Steiner trees. *J. Univers. Comput. Sci.*, 13(11):1615–1627, 2007. doi:10.3217/jucs-013-11-1615. 11

[DU23] Lindsey Deryckere and Seeun William Umboh. Online matching with set and concave delays. In *Proc. Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, volume 275 of *LIPIcs*, pages 17:1–17:17. Schloss Dagstuhl, 2023. doi:10.4230/LIPIcs.APPROX/RANDOM.2023.17. 2, 10

[EFN18] Michael Elkin, Arnold Filtser, and Ofer Neiman. Prioritized metric structures and embedding. *SIAM J. Comput.*, 47(3):829–858, 2018. doi:10.1137/17M1118749. 5

[EKW16] Yuval Emek, Shay Kutten, and Roger Wattenhofer. Online matching: haste makes waste! In *Proc. 48th ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 333–344, 2016. doi:10.1145/2897518.2897557. 2, 10

[EN22] Michael Elkin and Ofer Neiman. Lossless prioritized embeddings. *SIAM J. Discret. Math.*, 36(3):1529–1550, 2022. doi:10.1137/21m1436221. 5

[ERW10] Matthias Englert, Harald Räcke, and Matthias Westermann. Reordering buffers for general metric spaces. *Theory Comput.*, 6(1):27–46, 2010. doi:10.4086/toc.2010.v006a002. 3

[FGK20] Arnold Filtser, Lee-Ad Gottlieb, and Robert Krauthgamer. Labelings vs. embeddings: On distributed representations of distances. In *Proc. 31st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1063–1075,

Copyright © 2024
Copyright for this paper is retained by authors

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

2020. doi:10.1137/1.9781611975994.65. 5, 11
[Fil19] Arnold Filtser. On strong diameter padded decompositions. In *Proc. Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, volume 145 of *LIPIcs*, pages 6:1–6:21. Schloss Dagstuhl, 2019. doi:10.4230/LIPIcs.APPROX-RANDOM.2019.6. 7, 8, 11, 13
[Fil20] Arnold Filtser. A face cover perspective to ℓ_1 embeddings of planar graphs. In *Proc. 31st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1945–1954, 2020. doi:10.1137/1.9781611975994.120. 11
[Fil21] Arnold Filtser. Hop-constrained metric embeddings and their applications. In *Proc. 62nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 492–503, 2021. doi:10.1109/FOCS52979.2021.00056. 11
[FKS19] E. Fox-Epstein, P. N. Klein, and A. Schild. Embedding planar graphs into low-treewidth graphs with applications to efficient approximation schemes for metric problems. In *Proc. 30th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, page 1069–1088, 2019. doi:10.1137/1.9781611975482.66. 11
[FL21] Arnold Filtser and Hung Le. Clan embeddings into trees, and low treewidth graphs. In *Proc. 53rd ACM Symposium on Theory of Computing (STOC)*, pages 342–355, 2021. doi:10.1145/3406325.3451043. 11
[FL22] Arnold Filtser and Hung Le. Low treewidth embeddings of planar and minor-free metrics. In *Proc. 63rd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1081–1092, 2022. doi:10.1109/FOCS54457.2022.00105. 11
[FMR90] Alan M. Frieze, Colin McDiarmid, and Bruce A. Reed. Greedy matching on the line. *SIAM J. Comput.*, 19(4):666–672, 1990. doi:10.1137/0219045. 11
[FRT04] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.*, 69(3):485–497, 2004. doi:10.1016/j.jcss.2004.04.011. 3, 13
[FT03] J. Fakcharoenphol and K. Talwar. An improved decomposition theorem for graphs excluding a fixed minor. In *Proc. 6th International Workshop on Approximation Algorithms for Combinatorial Optimization (APPROX)*, volume 2764 of *LNCS*, pages 36–46. Springer, 2003. doi:10.1007/978-3-540-45198-3_4. 11
[GGK16] Albert Gu, Anupam Gupta, and Amit Kumar. The power of deferral: Maintaining a constant-competitive Steiner tree online. *SIAM J. Comput.*, 45(1):1–28, 2016. doi:10.1137/140955276. 3, 6, 10, 11
[GGK+22] Anupam Gupta, Vijaykrishna Gurunathan, Ravishankar Krishnaswamy, Amit Kumar, and Sahil Singla. Online discrepancy with recourse for vectors and graphs. In *Proc. 33rd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1356–1383, 2022. doi:10.1137/1.9781611977073.57. 11
[GK13] Lee-Ad Gottlieb and Robert Krauthgamer. Proximity algorithms for nearly doubling spaces. *SIAM J. Discret. Math.*, 27(4):1759–1769, 2013. doi:10.1137/120874242. 12
[GKL03] Anupam Gupta, Robert Krauthgamer, and James R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *Proc. 44th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 534–543, 2003. doi:10.1109/SFCS.2003.1238226. 7, 8, 12
[GKM+19] Buddhima Gamlath, Michael Kapralov, Andreas Maggiori, Ola Svensson, and David Wajc. Online matching with general arrivals. In *Proc. 60th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 26–37, 2019. doi:10.1109/FOCS.2019.00011. 2
[GM08] Gagan Goel and Aranyak Mehta. Online budgeted matching in random input models with applications to Adwords. In *Proc. 19th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 982–991, 2008. URL: https://dl.acm.org/doi/10.5555/1347082.1347189. 2
[GP68] E. N. Gilbert and H. O. Pollak. Steiner minimal trees. *SIAM Journal on Applied Mathematics*, 16(1):1–29, 1968. doi:10.1137/0116001. 11
[GW95] Michel X. Goemans and David P. Williamson. A general approximation technique for constrained forest problems. *SIAM J. Comput.*, 24(2):296–317, 1995. doi:10.1137/S0097539793242618. 5
[HM06] Sariel Har-Peled and Manor Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM J. Comput.*, 35(5):1148–1184, 2006. doi:10.1137/S0097539704446281. 12
[IMSZ10] Piotr Indyk, Avner Magen, Anastasios Sidiropoulos, and Anastasios Zouzias. Online embeddings. In *Proc. 13th APPROX*, volume 6302 of *LNCS*, pages 246–259. Springer, 2010. doi:10.1007/978-3-642-15369-3_19. 3, 4, 5, 7, 8, 11
[IW91] Makoto Imase and Bernard M. Waxman. Dynamic Steiner tree problem. *SIAM Journal on Discrete Mathematics*, 4(3):369–384, 1991. doi:10.1137/0404033. 11
[JL11] Patrick Jaillet and Xin Lu. Online traveling salesman problems with service flexibility. *Networks*, 58(2):137–146, 2011. doi:10.1002/net.20454. 11
[JLM11] Alexander Jaffe, James R. Lee, and Mohammad Moharrami. On the optimality of gluing over scales. *Discret. Comput. Geom.*, 46(2):270–282, 2011. doi:10.1007/s00454-011-9359-3. 13
[KKM+12] Maleq Khan, Fabian Kuhn, Dahlia Malkhi, Gopal Pandurangan, and Kunal Talwar. Efficient distributed approximation algorithms via probabilistic tree embeddings. *Distributed Computing*, 25(3):189–205, 2012. doi:10.1007/s00446-012-0157-9. 3
[KLMN04] Robert Krauthgamer, James R. Lee, Manor Mendel, and Assaf Naor. Measured descent: A new embedding

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Copyright © 2024
Copyright for this paper is retained by authors

method for finite metrics. In *Proc. 45th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 434–443, 2004. doi:10.1109/FOCS.2004.41. 7, 11, 13

[KP94] Bala Kalyanasundaram and Kirk R. Pruhs. Constructing competitive tours from local information. *Theoretical Computer Science*, 130(1):125–138, 1994. doi:10.1016/0304-3975(94)90155-4. 11

[KP98] Bala Kalyanasundaram and Kirk Pruhs. On-line network optimization problems. In Amos Fiat and Gerhard J. Woeginger, editors, *Online Algorithms: The State of the Art*, pages 268–280. Springer, Berlin, 1998. doi:10.1007/BFb0029573. 2

[KPR93] Philip N. Klein, Serge A. Plotkin, and Satish Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proc. 25th ACM Symposium on Theory of Computing (STOC)*, pages 682–690, 1993. doi:10.1145/167088.167261. 11, 13

[KVV90] Richard M. Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proc. 22nd ACM Symposium on Theory of Computing (STOC)*, pages 352–358, 1990. doi:10.1145/100216.100262. 2

[LLMV20] Silvio Lattanzi, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii. Online scheduling via learned weights. In *Proc. 14th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1859–1877, 2020. doi:10.1137/1.9781611975994.11. 2

[LLR95] Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995. doi:10.1007/BF01200757. 3

[Mat13] Jiří Matoušek. Lecture notes on metric embeddings. Technical report, ETH Zürich, 2013. URL: https://kam.mff.cuni.cz/~matousek/ba-a4.pdf. 11

[MPRS23] Mathieu Mari, Michal Pawlowski, Runtian Ren, and Piotr Sankowski. Online matching with delays and stochastic arrival times. In *Proc. 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 976–984. ACM, 2023. doi:10.5555/3545946.3598737. 10

[MPX13] Gary L. Miller, Richard Peng, and Shen Chen Xu. Parallel graph decompositions using random shifts. In Guy E. Blelloch and Berthold Vöcking, editors, *Proc. 25th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 196–203, 2013. doi:10.1145/2486159.2486180. 8

[MSVV07] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *J. ACM*, 54(5):22–es, 2007. doi:10.1145/1284320.1284321. 2

[MSVW16] Nicole Megow, Martin Skutella, José Verschae, and Andreas Wiese. The power of recourse for online MST and TSP. *SIAM J. Comput.*, 45(3):859–880, 2016. doi:10.1137/130917703. 11

[NR20] Ilan Newman and Yuri Rabinovich. Online embedding of metrics. In *Proc. 17th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT)*, volume 162 of *LIPIcs*, pages 32:1–32:13. Schloss Dagstuhl, 2020. Full version available at arXiv:2303.15945. doi:10.4230/LIPIcs.SWAT.2020.32. 4, 5, 8, 11

[Rao99] Satish Rao. Small distortion and volume preserving embeddings for planar and Euclidean metrics. In *Proc. 15th Symposium on Computational Geometry (SoCG)*, pages 300–306. ACM, 1999. doi:10.1145/304893.304983. 7, 8, 11, 13

[RSI77] Daniel J. Rosenkrantz, Richard Edwin Stearns, and Philip M. Lewis II. An analysis of several heuristics for the traveling salesman problem. *SIAM J. Comput.*, 6(3):563–581, 1977. doi:10.1137/0206041. 11

[RT81] Edward M. Reingold and Robert Endre Tarjan. On a greedy heuristic for complete matching. *SIAM J. Comput.*, 10(4):676–681, 1981. doi:10.1137/0210050. 11

[SRP83] Kenneth J. Supowit, Edward M. Reingold, and David A. Plaisted. The travelling salesman problem and minimum matching in the unit square. *SIAM J. Comput.*, 12(1):144–156, 1983. doi:10.1137/0212009. 7

[SS89] J. Michael Steele and Timothy Law Snyder. Worst-case growth rates of some classical problems of combinatorial optimization. *SIAM J. Comput.*, 18(2):278–287, 1989. doi:10.1137/0218019. 7

[ST83] Daniel Dominic Sleator and Robert Endre Tarjan. A data structure for dynamic trees. *J. Comput. Syst. Sci.*, 26(3):362–391, 1983. doi:10.1016/0022-0000(83)90006-5. 9

[SWW95] David B. Shmoys, Joel Wein, and David P. Williamson. Scheduling parallel machines on-line. *SIAM Journal on Computing*, 24(6):1313–1331, 1995. doi:10.1137/S0097539793248317. 2

Downloaded 04/11/24 to 46.117.103.92 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Copyright © 2024
Copyright for this paper is retained by authors