

Shortest Paths without a Map, but with an Entropic Regularizer

Sébastien Bubeck
ML Foundations Group
Microsoft Research
Redmond, WA, United States
sebubeck@microsoft.com

Christian Coester
Department of Computer Science
University of Oxford
Oxford, United Kingdom
christian.coester@cs.ox.ac.uk

Yuval Rabani
School of Computer Science & Engr.
The Hebrew University of Jerusalem
Jerusalem, Israel
yrabani@cs.huji.ac.il

Abstract—In a 1989 paper titled “shortest paths without a map”, Papadimitriou and Yannakakis introduced an online model of searching in a weighted layered graph for a target node, while attempting to minimize the total length of the path traversed by the searcher. This problem, later called layered graph traversal, is parametrized by the maximum cardinality k of a layer of the input graph. It is an online setting for dynamic programming, and it is known to be a rather general and fundamental model of online computing, which includes as special cases other acclaimed models. The deterministic competitive ratio for this problem was soon discovered to be exponential in k , and it is now nearly resolved: it lies between $\Omega(2^k)$ and $O(k2^k)$. Regarding the randomized competitive ratio, in 1993 Ramesh proved, surprisingly, that this ratio has to be at least $\Omega(k^2/\log^{1+\varepsilon} k)$ (for any constant $\varepsilon > 0$). In the same paper, Ramesh also gave an $O(k^{13})$ -competitive randomized online algorithm. Since 1993, no progress has been reported on the randomized competitive ratio of layered graph traversal. In this work we show how to apply the mirror descent framework on a carefully selected evolving metric space, and obtain an $O(k^2)$ -competitive randomized online algorithm, nearly matching the known lower bound on the randomized competitive ratio.

Index Terms—online computing, competitive analysis, shortest path

I. INTRODUCTION

Our results: In this paper we present a randomized $O(k^2)$ -competitive online algorithm for width k layered graph traversal. The problem, whose history is discussed in detail below, is an online version of the conventional shortest path problem (see [28]), and of Bellman’s dynamic programming paradigm. It is therefore a very general framework for decision-making under uncertainty of the future, and it includes as special cases other celebrated models of online computing, such as metrical task systems. The following 1989 words of Papadimitriou and Yannakakis [26] still resonate today: “the techniques developed [for layered graph traversal] will add to the scarce rigorous methodological arsenal of Artificial Intelligence.” Our new upper bound on the competitive ratio nearly matches the lower bound of $\Omega(k^2/\log^{1+\varepsilon} k)$, for

The second author’s research was done while he was at Tel Aviv University, supported by the Israel Academy of Sciences and Humanities & Council for Higher Education Excellence Fellowship Program for International Postdoctoral Researchers. The third author’s research was supported in part by ISF grants 3565-21 and 389-22, and by BSF grant 2018687.

all $\varepsilon > 0$, given by Ramesh [27]. It improves substantially over the previously known $O(k^{13})$ upper bound from the same paper.

Problem definition: In layered graph traversal, a searcher attempts to find a short path from a starting node a to a target node b in an undirected graph G with non-negative integer edge weights $w : E(G) \rightarrow \mathbb{N}$. It is assumed that the vertices of G are partitioned into layers such that edges exist only between vertices of consecutive layers. The searcher knows initially only an upper bound k on the number of nodes in a layer (this is called the *width* of G). The search begins at a , which is the only node in the first layer of the graph (indexed layer 0). We may assume that all the nodes of G are reachable from a . When the searcher first reaches a node in layer i , the next layer $i+1$, the edges between layer i and layer $i+1$, and the weights of these edges, are revealed to the searcher.¹ At this point, the searcher has to move forward from the current position (a node in layer i) to a new position (a node in layer $i+1$). This can be done along a shortest path through the layers revealed so far.² The target node b occupies the last layer of the graph; the number of layers is unknown to the searcher. The target gets revealed when the preceding layer is reached. The goal of the searcher is to traverse a path of total weight as close as possible to the shortest path connecting a and b .³ The searcher is said to be C -competitive for some $C = C(k)$ iff for every width k input (G, w) , the searcher’s path weight is at most $C \cdot w_G(a, b)$, where $w_G(a, b)$ denotes the distance under the weight function w between a and b . The competitive ratio of the problem is the best C for which there is a C -competitive searcher.

Evidently, rational weights can be converted to integer weights if we know (or even update online) a common denominator, and irrational weights can be approximated to within any desired accuracy using rational weights. Moreover, any bipartite graph is layered, and any graph can be con-

¹So, in particular, when the search starts at a , the searcher knows all the nodes of layer 1, all the edges connecting a to a node in layer 1, and all the weights of such edges.

²Notice that this is not necessarily the shortest path in G between the nodes, because such a path may have to go through future layers that have not been revealed so far. In fact, some nodes in layer $i+1$ might not even be reachable at this point.

³Note that the traversed path is not required to be simple or level-monotone.

verted into a bipartite graph by subdividing edges. Thus, the layered structure of the input graph is intended primarily to parametrize the way in which the graph is revealed over time to the online algorithm. Also notice that if the width k is not known, it can be estimated, for instance by doubling the guess each time it is refuted. Therefore, the assumptions made in the definition of the problem are not overly restrictive.

Motivation and past work: The problem has its origins in a paper of Baeza-Yates et al. [3], and possibly earlier in game theory [21]. In [3], motivated by applications in robotics, the special case of a graph consisting of k disjoint paths is proposed, under the name *the lost cow problem*. The paper gives a deterministic 9-competitive algorithm for $k = 2$, and more generally a deterministic $2\frac{k^k}{(k-1)^{k-1}} + 1 \approx 2ek + 1$ competitive algorithm for arbitrary k . A year later, Papadimitriou and Yannakakis introduced the problem of layered graph traversal [26]. Their paper shows that for $k = 2$ the upper bound of 9 still holds, and that the results of [3] are optimal for disjoint paths. Unfortunately, the upper bound of 9 in [26] is the trivial consequence of the observation that for $k = 2$, the general case reduces to the disjoint paths case. This is not true for general k . Indeed, Fiat et al. [18] give a 2^{k-2} lower bound, and an $O(9^k)$ upper bound on the deterministic competitive ratio in the general case. The upper bound was improved by Ramesh [27] to $O(k^3 2^k)$, and further by Burley [12] to $O(k 2^k)$. Thus, currently the deterministic case is nearly resolved, asymptotically: the competitive ratio lies in $\Omega(2^k) \cap O(k 2^k)$.

Investigation of the randomized competitive ratio was initiated in the aforementioned [18] that gives a $\frac{k}{2}$ lower bound for the general case, and asymptotically tight $\Theta(\log k)$ upper and lower bounds for the disjoint paths case. In the randomized case, the searcher’s strategy is a distribution over moves, and it is C -competitive iff for every width k input (G, w) , the expected weight of the searcher’s path is at most $C \cdot w_G(a, b)$. It is a ubiquitous phenomenon of online computing that randomization improves the competitive ratio immensely, often guaranteeing exponential asymptotic improvement (as happens in the disjoint paths case of layered graph traversal). To understand why this might happen, one can view the problem as a game between the designer of G and the searcher in G . The game alternates between the designer introducing a new layer and the searcher moving to a node in the new layer. The designer is oblivious to the searcher’s moves. Randomization obscures the predictability of the searcher’s moves, and thus weakens the power of the designer.⁴ Following the results in [18] and the recurring phenomenon of exponential improvement, a natural conjecture would have been that the randomized competitive ratio of layered graph traversal is $\Theta(k)$. However, this natural conjecture was rather quickly refuted in the aforementioned [27] that surprisingly improves the lower bound to $\Omega(k^2 / \log^{1+\varepsilon} k)$, which holds for all $\varepsilon > 0$. Moreover, the same paper also gives an

upper bound of $O(k^{13})$. Thus, even though for the general case of layered graph traversal the randomized competitive ratio cannot be logarithmic in the deterministic competitive ratio, it is polylogarithmic in that ratio. The results of [27] on randomized layered graph traversal have not since been improved prior to our current paper.

Computing the optimal offline solution, a shortest path from source to target in a weighted layered graph, is a simple example and also a generic framework for dynamic programming [8]. The online version has applications to the design and analysis of hybrid algorithms. In particular, the disjoint paths case has applications in derandomizing online algorithms [20], in the design of divide-and-conquer online algorithms [2], [19], and in the design of advice and learning augmented online algorithms [1], [5], [25]. In this context, Kao et al. [23] resolve exactly the randomized competitive ratio of width 2 layered graph traversal: it is roughly 4.59112, precisely the solution for x in the equation $\ln(x - 1) = \frac{x}{x-1}$; see also [13]. For more in this vein, see also Kao et al. [22].

Moreover, layered graph traversal is a very general model of online computing. For example, many online problems can be represented as chasing finite subsets of points in a metric space. This problem, introduced by Chrobak and Larmore [13], [14] under the name *metrical service systems*, is **equivalent** to layered graph traversal [18]. The width k of the layered graph instance corresponds to the maximum cardinality of any request of the metrical service systems instance. (See Section IV.) Metrical service systems are a special case of metrical task systems, introduced by Borodin et al. [9]. Width k layered graph traversal includes as a special case metrical task systems in k -point metric spaces.⁵ There is a tight bound of $2k - 1$ on the deterministic competitive ratio of metrical task systems in any k -point metric [9], and the randomized competitive ratio lies between an $\Omega(\log k / \log \log k)$ lower bound (Bartal et al. [6], [7]) and an $O(\log^2 k)$ upper bound (Bubeck et al. [10]). Thus, width k layered graph traversal is strictly a more general problem than k -point metrical task systems. Another closely related problem is the k -taxi problem, whose best known lower bound for deterministic algorithms is obtained via a reduction from layered graph traversal [15].

Our methods: Our techniques are based on the method of online mirror descent with entropic regularization that was pioneered by Bubeck et al. [10], [11] in the context of the k -server problem and metrical task systems, and further explored in this context in a number of recent papers [4], [16], [17]. It is known that layered graph traversal is equivalent to its special case where the input graph is a tree [18]. Based on this reduction, we reduce width k layered graph traversal to a problem that we name the (depth k) *evolving tree game*. In this game, one player, representing the algorithm, occupies a (non-root) leaf in an evolving rooted, edge weighted, bounded

⁴This can be formalized through an appropriate definition of the designer’s information sets.


⁵Indeed, this implies that while metrical service systems on k -point metrics are a special case of metrical task systems on k -point metrics, also metrical task systems on k -point metrics are a special case of metrical service systems using k -point requests.

degree tree of depth $\leq k$. Its opponent, the adversary, is allowed to change the metric and topology of the tree using the following repertoire of operations: (i) increase the weight of an edge incident to a leaf (see Fig. 1a); (ii) delete a leaf and the incident edge, and smooth the tree at the parent if its degree is now 2 (see Fig. 1b);⁶ (iii) create two (or more) new leaves and connect them with weight 0 edges to an existing leaf whose combinatorial depth is strictly smaller than k (see Fig. 1c). The algorithm may move from leaf to leaf at any time, incurring *movement cost* equal to the weight of the path between the leaves. If the algorithm occupies a leaf at the endpoint of a growing weight edge, it pays the increase in weight. We call this the *service cost* of the algorithm. If the algorithm occupies a leaf that is being deleted, it must move to a different leaf prior to the execution of the topology change. If the algorithm occupies a leaf that is being converted into an internal node (because new leaves are appended to it), the algorithm must move to a leaf after the execution of the topology change. At the end of the game, the total (movement + service) cost of the algorithm is compared against the adversary's cost, which is the weight of the lightest root-to-leaf path.⁷


Mirror descent is used to generate a fractional online solution to the evolving tree game. The algorithm maintains a probability distribution on the leaves. A fractional solution can be converted easily on-the-fly into a randomized algorithm. As in [10], [11] the analysis of our fractional algorithm for the dynamic tree game is based on a potential function that combines (in our case, a modification of) the Bregman divergence associated with the entropic regularizer with a weighted depth potential. The Bregman divergence is used to bound the algorithm's *service cost* against the adversary's cost. The weighted depth potential is used to bound the algorithm's *movement cost* against its own service cost.

However, in our setting, in contrast to [10], [11], the metric on the set of leaves, and even the topology of the underlying tree, change dynamically. This poses a few new challenges to the approach. In particular, the potential function that works for metrical task systems is not invariant under the topology changes that are needed here. We resolve this problem by working with revised edge weights that slightly over-estimate the true edge weights. When a topology change would lead to an increase of the potential function (by reducing the combinatorial depth of some vertices), we prevent such an increase by downscaling the affected revised edge weights appropriately.


Even so, the extra cost incurred by the perturbation of entropy, which is required to handle distributions close to the boundary, cannot be handled in the same manner as in [10], [11]. This issue is fixed by modifying both the Bregman divergence and the control function of the mirror



(a) Continuous step.



(b) Delete step.



(c) Fork step.

Fig. 1: Three types of metric and topology changes of the tree.

descent dynamic. The latter damps down the movement of the algorithm when it incurs service cost at a rate close to 0.

In the competitive ratio of $O(k^2)$, one factor k comes from the maximal depth k of the tree. The other factor k is due to the fact that the perturbation that we require is exponentially small in k , on account of having to set it consistently against unknown future topology changes. (One way to think about it is that we are handling in the background a complete tree of depth k , only part of which is currently revealed. In fact, our bounds do not require that the number of leaves be restricted to k , as is the case with layered graph traversal.) We note that implementing the mirror descent approach in evolving trees is a major challenge to the design and analysis of online algorithms for online problems in metric spaces (e.g., the k -server problem, see [24], where an approach based on mirror descent in an evolving tree is also studied). Our ideas may prove applicable to other problems.

Organization: The rest of this paper is organized as follows. In Section II we define and analyze the evolving tree game. In Section III we motivate the evolving tree algorithm and analysis. In Section IV we discuss the application to

⁶Smoothing is the reverse operation of subdividing. In other words, smoothing is merging the two edges incident to a degree 2 node. We maintain w.l.o.g. the invariant that the tree has no degree 2 node.

⁷In fact, our algorithm can handle also the operation of reducing the weight of an edge, under the assumption that this operation incurs on both players a cost equal to the reduction in weight, if performed at their location.

layered graph traversal/small set chasing.

II. THE EVOLVING TREE GAME

For a rooted edge-weighted tree $T = (V, E)$, we denote by r its root, by $V^0 := V \setminus \{r\}$ the set of non-root vertices and by $\mathcal{L} \subseteq V^0$ the set of leaves. For $u \in V^0$, we denote by p_u the parent of u and by w_u the length of the edge connecting u to p_u .

The evolving tree game is a two person continuous time game between an adversary and an algorithm. The adversary grows a rooted edge-weighted tree $T = (V, E)$ of bounded degree. Without loss of generality, we enforce that the root r always has degree 1, and we denote its single child by c_r . Initially $V = \{r, c_r\}$, and the two nodes are connected by a zero-weight edge. The root r will be fixed throughout the game, but the identity of its child c_r may change as the game progresses. The game has continuous steps and discrete steps (see Fig. 1):

- **Continuous step:** The adversary picks a leaf ℓ and increases the weight w_ℓ of the edge incident on ℓ at a fixed rate of $w'_\ell = 1$ for a finite time interval.
- **Discrete step:** There are two types of discrete steps:
 - **Delete step:** The adversary chooses a leaf $\ell \in \mathcal{L}$, $\ell \neq c_r$, and deletes ℓ and its incident edge from T . If the parent p_ℓ of ℓ remains with a single child c , the adversary smooths T at p_ℓ as follows: it merges the two edges $\{c, p_\ell\}$ and $\{p_\ell, p_{p_\ell}\}$ into a single edge $\{c, p_{p_\ell}\}$, removing the vertex p_ℓ , and assigns $w_c \leftarrow w_c + w_{p_\ell}$.
 - **Fork step:** The adversary generates two or more new nodes and connects all of them to an existing leaf $\ell \in \mathcal{L}$ with edges of weight 0. Notice that this removes ℓ from \mathcal{L} and adds the new nodes to \mathcal{L} .

The continuous and discrete steps may be interleaved arbitrarily by the adversary.

A pure strategy of the algorithm maps the timeline to a leaf of T that exists at that time. Thus, the start of the game (time 0 of step 1) is mapped to c_r , and at all times the algorithm occupies a leaf and may move from leaf to leaf. If the algorithm occupies a leaf ℓ continuously while w_ℓ grows, it pays the increase in weight (we call this the *service cost*). If the algorithm moves from a leaf ℓ_1 to a leaf ℓ_2 , it pays the total weight of the path in T between ℓ_1 and ℓ_2 at the time of the move (we call this the *movement cost*). A mixed strategy/randomized algorithm is, as usual, a probability distribution over pure strategies. A fractional strategy maps the timeline to probability distributions over the existing leaves. Writing x_u for the total probability of the leaves in the subtree of u , this means that a fractional strategy maintains at all times a point in the changing polytope

$$K(T) := \left\{ x \in \mathbb{R}_+^{V^0} : \sum_{v: p(v)=u} x_v = x_u \forall u \in V \setminus \mathcal{L} \right\}, \quad (1)$$

where we view $x_r := 1$ as a constant.

Notice that the tree T , the weight function w , the point $x \in K(T)$, and the derived parameters are all functions of the adversary's step and the time t within a continuous step. Thus, we shall use henceforth the following notation: $T(j, 0)$, $w(j, 0)$, $x(j, 0)$, etc. to denote the values of these parameters at the start of step j of the adversary. If step j is a continuous step of duration τ , then for $t \in [0, \tau]$, we use $T(j, t)$, $w(j, t)$, $x(j, t)$, etc. to denote the values of these parameters at time t since the start of step j . If it is not required to mention the parameters (j, t) for clarity, we omit them from our notation.

We require that for a fixed continuous step j , the function $x(j, \cdot): [0, \tau] \rightarrow K(T)$ is absolutely continuous, and hence differentiable almost everywhere. Notice that the polytope $K(T)$ is fixed during a continuous step, so this requirement is well-defined. We denote by x' the derivative of x with respect to t , and similarly we denote by w' the derivative of w with respect to t . The cost of the algorithm during the continuous step j is

$$\int_0^\tau \sum_{v \in V^0} (w'_v(j, t) x_v(j, t) + w_v(j, t) |x'_v(j, t)|) dt.$$

Notice that the first summand (the *service cost*) is non-zero only at the single leaf ℓ for which $w_\ell(j, t)$ is growing.

In a discrete step j , the topology of the tree and thus the polytope $K(T)$ changes. The old tree is $T(j, 0)$ and the new tree is $T(j+1, 0)$. In a delete step, when a leaf ℓ is deleted, the algorithm first has to move from its old position $x(j, 0) \in K(T(j, 0))$ to some position $x \in K(T(j, 0))$ with $x_\ell = 0$. The cost of moving from $x(j, 0)$ to x is given by

$$\sum_{v \in V^0} w_v(j, 0) |x_v(j, 0) - x_v|.$$

The new state $x(j+1, 0)$ is the projection of x onto the new polytope $K(T(j+1, 0))$, where the ℓ -coordinate and possibly (if smoothing happens) the p_ℓ -coordinate are removed.

In a fork step, the algorithm chooses as its new position any point in $K(T(j+1, 0))$ whose projection onto $K(T(j, 0))$ is the old position of the algorithm. No cost is incurred here (since the new leaves are appended at distance 0).

The following lemma is analogous to Lemma 9. Its proof is very similar. It is omitted here; we actually do not need this claim to prove the main result of the paper.

Lemma 1. *For every fractional strategy of the algorithm there is a mixed strategy incurring the same cost in expectation.*

Our main result in this section is the following theorem.

Theorem 1. *For every $k \in \mathbb{N}$ and for every $\varepsilon > 0$ there exists a fractional strategy of the algorithm with the following performance guarantee. For every pure strategy of the adversary that grows trees of depth at most k , the cost C of the algorithm satisfies*

$$C \leq O(k^2 \log d_{\max}) \cdot (\text{opt} + \varepsilon),$$

where opt is the minimum distance in the final tree from the root to a leaf, and d_{\max} is the maximum degree of a node at any point during the game.

Notice that for every strategy of the adversary, there exists a pure strategy that pays exactly opt service cost and zero movement cost. We will refer to this strategy as the *optimal play*. The algorithm cannot in general choose the optimal play because the evolution of the tree only gets revealed step-by-step.

A. Additional notation

Let j be any step, and let t be a time in that step (so, if j is a discrete step then $t = 0$, and if j is a continuous step of duration τ then $t \in [0, \tau]$). For a vertex $u \in V(j, t)$ we denote by $h_u(j, t)$ its combinatorial depth, i.e., the number of edges on the path from r to u in $T(j, t)$. Instead of the actual edge weights $w_u(j, t)$, our algorithm will be based on revised edge weights defined as

$$\tilde{w}_u(j, t) := \frac{2k - 1}{2k - h_u(j, t)} (w_u(j, t) + \varepsilon 2^{-j_u}),$$

where $j_u \in \mathbb{N}$ is the step number when u was created (or 0 if u existed in the initial tree). The purpose of the term $\varepsilon 2^{-j_u}$ is to ensure that $\tilde{w}_u(j, t)$ is strictly positive.

For $u \in V^0(j, t)$, we also define a shift parameter by induction on $h_u(j, t)$, as follows. For $u = c_r(j, t)$, $\delta_u(j, t) = 1$. For other u , $\delta_u(j, t) = \delta_{p_u}(j, t)/(d_{p_u}(j, t) - 1)$, where $p_u = p_u(j, t)$, and $d_{p_u}(j, t)$ is the degree of p_u in $T(j, t)$ (i.e., $d_{p_u}(j, t) - 1$ is the number of children of p_u in $T(j, t)$; note that every non-leaf node in $V^0(j, t)$ has degree at least 3). Observe that by definition $\delta(j, t) \in K(T(j, t))$. As mentioned earlier, we often omit the parameters (j, t) from our notation, unless they are required for clarity.

B. The algorithm

We consider four distinct types of steps: continuous steps, fork steps, deadend steps, and merge steps. A delete step is implemented by executing a deadend step, and if needed followed by a merge step. It is convenient to examine the two operations required to implement a delete step separately.

Continuous step: In a continuous step, the weight w_ℓ of some leaf $\ell \in \mathcal{L}$ grows continuously (and thus $\tilde{w}'_\ell > 0$). In this case, for $u \in V^0$ we update the fractional mass in the subtree below u at rate

$$x'_u = -\frac{2x_u}{\tilde{w}_u} \tilde{w}'_u + \frac{x_u + \delta_u}{\tilde{w}_u} (\lambda_{p_u} - \lambda_u), \quad (2)$$

where $\lambda_u = 0$ for $u \in \mathcal{L}$ and $\lambda_u \geq 0$ for $u \in V \setminus \mathcal{L}$ are chosen such that x remains in the polytope $K(T)$. We will show in Section III-D that such λ exists (as a function of time).⁸

Fork step: In a fork step, new leaves $\ell_1, \ell_2, \dots, \ell_q$ (for some $q \geq 2$) are spawned as children of a leaf u (so that u is no longer a leaf). They are “born” with $w_{\ell_1} = w_{\ell_2} = \dots = w_{\ell_q} = 0$ and $x_{\ell_1} = x_{\ell_2} = \dots = x_{\ell_q} = x_u/q$.

⁸The dynamic of x corresponds to running mirror descent with regularizer $\Phi_t(z) = \sum_{u \in V} \tilde{w}_u(z_u + \delta_u) \log(z_u + \delta_u)$, using the growth rate \tilde{w}' of the approximate weights as cost function, and scaling the rate of movement by a factor $\frac{2x_\ell}{x_\ell + \delta_\ell}$ when ℓ is the leaf whose edge grows. See Section III.

Deadend step: In a deadend step, we delete a leaf $\ell \neq c_r$. To achieve this, we first compute the limit of a continuous step where the weight \tilde{w}_ℓ grows to infinity, ensuring that the mass x_ℓ tends to 0. This, of course, changes the mass at other nodes, and we update x to be the limit of this process. Then, we remove the leaf ℓ along with the edge $\{\ell, p_\ell\}$ from the tree. Notice that this changes the degree d_{p_ℓ} . Therefore, it triggers a discontinuous change in the shift parameter δ_u for every vertex u that is a descendant of p_ℓ .

Merge step: A merge step immediately follows a deadend step if, after the removal of the edge $\{\ell, p_\ell\}$, the vertex $v = p_\ell$ has only one child c left. Notice that $v \neq r$. We merge the two edges $\{c, v\}$ of weight w_c and $\{v, p_v\}$ of weight w_v into a single edge $\{c, p_v\}$ of weight $w_c + w_v$. The two edges that were merged and the vertex v are removed from T . This decrements by 1 the combinatorial depth h_u of every vertex u in the subtree rooted at c . Thus, it triggers a discontinuous change in the revised weight \tilde{w}_u , for every vertex u in this subtree.

C. Competitive analysis

The analysis of the algorithm is based on a potential function argument. Let $y \in K(T)$ denote the state of the optimal play. Note that as the optimal play is a pure strategy, the vector y is simply the indicator function for the nodes on some root-to-leaf path. We define a potential function $P = P_{k,T,w}(x, y)$, where $x \in K(T)$ and we prove that the algorithm’s cost plus the change in P is at most $O(k^2 \log d_{\max})$ times the optimal cost, where d_{\max} is the maximum degree of a node of T . This, along with the fact that P is bounded, implies $O(k^2 \log d_{\max})$ competitiveness. In Section III we motivate the construction of the potential function P . Here, we simply define it as follows:

$$P := 2 \sum_{u \in V^0} \tilde{w}_u \left(4k y_u \log \frac{1 + \delta_u}{x_u + \delta_u} + (2k - h_u) x_u \right). \quad (3)$$

We now consider the cost and potential change for each of the different steps separately.

1) Continuous step:

Bounding the algorithm’s cost: Let ℓ be the leaf whose weight w_ℓ is growing, and recall that c_r is the current neighbor of the root r . By definition of the game, the algorithm pays two types of cost. Firstly, it pays for the mass x_ℓ at the leaf ℓ moving away from the root at rate w'_ℓ . Secondly, it pays for moving mass from ℓ to other leaves. Notice that $x_{c_r} = 1$ stays fixed. Let $C = C(j, t)$ denote the total cost that the algorithm accumulates in the current step j , up to the current time t .

Lemma 2. *The rate at which C increases with t is*

$$C' \leq 3\tilde{w}'_\ell x_\ell + 2 \sum_{u \in V^0} (x_u + \delta_u) \lambda_u.$$

Proof. We have

$$\begin{aligned} C' &= w'_\ell x_\ell + \sum_{u \in V^0 \setminus \{c_r\}} w_u |x'_u| \\ &\leq \tilde{w}'_\ell x_\ell + \sum_{u \in V^0 \setminus \{c_r\}} \tilde{w}_u |x'_u| \end{aligned} \quad (4)$$

$$= \tilde{w}'_\ell x_\ell + \sum_{u \in V^0 \setminus \{c_r\}} |-2x_u \tilde{w}'_u + (x_u + \delta_u)(\lambda_{p_u} - \lambda_u)| \quad (5)$$

$$\leq 3\tilde{w}'_\ell x_\ell + \sum_{u \in V^0 \setminus \{c_r\}} (x_u + \delta_u)(\lambda_{p_u} + \lambda_u) \quad (6)$$

$$\leq 3\tilde{w}'_\ell x_\ell + 2 \sum_{u \in V^0} (x_u + \delta_u) \lambda_u. \quad (7)$$

Inequality (4) uses the fact that $w_u \leq \tilde{w}_u$ and $w'_\ell \leq \tilde{w}'_\ell$. Equation (5) uses the definition of the dynamic in Equation (2). Inequality (6) uses the triangle inequality. Finally, Inequality (7) uses the fact that $x, \delta \in K(T)$, so $\sum_{v: p_v=u} (x_v + \delta_v) = x_u + \delta_u$ for all $u \in V \setminus \mathcal{L}$. \square

Change of potential: We decompose the potential function as $P = 4kD - 2\Psi$, where

$$D := \sum_{u \in V^0} \tilde{w}_u \left(2y_u \log \frac{1 + \delta_u}{x_u + \delta_u} + x_u \right)$$

and

$$\Psi := \sum_{u \in V^0} h_u \tilde{w}_u x_u.$$

We first analyze the rate of change of Ψ .

Lemma 3. *The rate at which Ψ changes satisfies:*

$$\Psi' \geq -k\tilde{w}'_\ell x_\ell + \sum_{u \in V^0} \lambda_u (x_u + \delta_u).$$

Proof. We have

$$\begin{aligned} \Psi' &= h_\ell \tilde{w}'_\ell x_\ell + \sum_{u \in V^0 \setminus \{c_r\}} h_u \tilde{w}_u x'_u \\ &= h_\ell \tilde{w}'_\ell x_\ell + \sum_{u \in V^0 \setminus \{c_r\}} h_u (-2x_u \tilde{w}'_u + (x_u + \delta_u)(\lambda_{p_u} - \lambda_u)) \\ &= -h_\ell \tilde{w}'_\ell x_\ell + \sum_{u \in V^0 \setminus \{c_r\}} h_u (x_u + \delta_u)(\lambda_{p_u} - \lambda_u) \\ &\geq -k\tilde{w}'_\ell x_\ell + \sum_{u \in V^0} \lambda_u \left((h_u + 1) \sum_{v: p_v=u} (x_v + \delta_v) - h_u (x_u + \delta_u) \right) \quad (8) \\ &= -k\tilde{w}'_\ell x_\ell + \sum_{u \in V^0} \lambda_u (x_u + \delta_u). \end{aligned}$$

Here, Inequality (8) uses the fact that $h_\ell \leq k$ and, for $u = p_v$, $h_v = h_u + 1$. Equation (9) uses the previously noted fact that, as $x, \delta \in K(T)$, then for all $u \notin \mathcal{L}$, $\sum_{v: p_v=u} (x_v + \delta_v) = x_u + \delta_u$ (and if $u \in \mathcal{L}$, then $\lambda_u = 0$). \square

Next, we analyze the rate of change of D .

Lemma 4. *The rate at which D changes satisfies:*

$$D' \leq -\tilde{w}'_\ell x_\ell + 2(2 + k \log d_{\max}) y_\ell \tilde{w}'_\ell.$$

Proof. We have

$$\begin{aligned} D' &= \tilde{w}'_\ell \left(2y_\ell \log \frac{1 + \delta_\ell}{x_\ell + \delta_\ell} + x_\ell \right) + \sum_{u \in V^0 \setminus \{c_r\}} \tilde{w}_u x'_u \left(\frac{-2y_u}{x_u + \delta_u} + 1 \right) \\ &= \tilde{w}'_\ell \left(2y_\ell \log \frac{1 + \delta_\ell}{x_\ell + \delta_\ell} + x_\ell \right) + \\ &\quad + \sum_{u \in V^0 \setminus \{c_r\}} (-2x_u \tilde{w}'_u + (x_u + \delta_u)(\lambda_{p_u} - \lambda_u)) \left(\frac{-2y_u}{x_u + \delta_u} + 1 \right) \\ &= -\tilde{w}'_\ell x_\ell + 2y_\ell \tilde{w}'_\ell \left(\log \frac{1 + \delta_\ell}{x_\ell + \delta_\ell} + \frac{2x_\ell}{x_\ell + \delta_\ell} \right) + \\ &\quad + \sum_{u \in V^0 \setminus \{c_r\}} (\lambda_{p_u} - \lambda_u) (-2y_u + x_u + \delta_u) \\ &\leq -\tilde{w}'_\ell x_\ell + 2y_\ell \tilde{w}'_\ell (2 + k \log d_{\max}) + \\ &\quad + \sum_{u \in V^0} \lambda_u \left(2y_u - x_u - \delta_u - \sum_{v: p_v=u} (2y_v - x_v - \delta_v) \right) \quad (10) \\ &= -\tilde{w}'_\ell x_\ell + 2y_\ell \tilde{w}'_\ell (2 + k \log d_{\max}). \quad (11) \end{aligned}$$

Inequality (10) uses the fact that $\delta_\ell \geq (d_{\max})^{1-h_\ell} \geq (d_{\max})^{1-k}$ and $2y_{c_r} - x_{c_r} - \delta_{c_r} = 0$. Equation (11) uses the fact that $x, y, \delta \in K(T)$, hence for $u \in V \setminus \mathcal{L}$,

$$2y_u - x_u - \delta_u = \sum_{v: p_v=u} (2y_v - x_v - \delta_v)$$

(and for $u \in \mathcal{L}$, $\lambda_u = 0$). \square

We obtain the following lemma, which bounds the algorithm's cost and change in potential against the service cost of the optimal play.

Lemma 5. *For every $k \geq 2$, it holds that*

$$C' + P' \leq O(k^2 \log d_{\max}) w'_\ell y_\ell.$$

Proof. Combine Equations (7), (9), and (11), and recall that $P = 4kD - 2\Psi$. We get

$$\begin{aligned} C' + P' &\leq (2k + 3 - 4k) \tilde{w}'_\ell x_\ell + 8k(2 + k \log d_{\max}) \tilde{w}'_\ell y_\ell \\ &\leq O(k^2 \log d_{\max}) w'_\ell y_\ell, \end{aligned} \quad (12)$$

where in the last inequality we use $\tilde{w}'_\ell < 2w'_\ell$. \square

2) *Fork step:*

Fork steps may increase the value of the potential function P , because the new edges have revised weight > 0 . The following lemma bounds this increase.

(9) **Lemma 6.** *The total increase in P due to all fork steps is at most $\varepsilon \cdot O(k^2 \log d_{\max})$.*

Proof. Consider a fork step that attaches new leaves ℓ_1, \dots, ℓ_q to a leaf u . The new leaves are born with revised edge weights $\frac{2k-1}{2k-h_u-1} \varepsilon 2^{-j} \leq \varepsilon 2^{-j+1}$, where j is the current step number. Since $\sum_{i=1}^q y_{\ell_i} = y_u \leq 1$ and $\sum_{i=1}^q x_{\ell_i} = x_u \leq 1$, the change ΔP in P satisfies

$$\begin{aligned} \Delta P &\leq \varepsilon 2^{-j+2} \cdot \left(4k \log \frac{1 + \delta_u/q}{\delta_u/q} + 2k - h_u - 1 \right) \\ &\leq \varepsilon 2^{-j+2} \cdot (2k + 4k^2 \log d_{\max}), \end{aligned}$$

where the last inequality follows from $\delta_u/q \geq (d_{\max})^{1-k}$. As the step number j is different in all fork steps, the total cost of all fork steps is at most

$$\varepsilon \cdot (2k + 4k^2 \log d_{\max}) \sum_{j=1}^{\infty} 2^{-j+2} = \varepsilon \cdot O(k^2 \log d_{\max}),$$

as claimed. \square

3) Deadend step:

Recall that when a leaf ℓ is deleted, we first compute the limit of a continuous step as the weight \tilde{w}_ℓ grows to infinity. Let \bar{x} be the mass distribution that the algorithm converges to when \tilde{w}_ℓ approaches infinity.

Lemma 7. *The limit \bar{x} satisfies $\bar{x}_\ell = 0$. Hence, \bar{x} with the ℓ -coordinate removed is a valid mass distribution in the new polytope $K(T)$. Also, a deadend step decreases P by at least the cost the algorithm incurs to move to \bar{x} .*

Proof. Note that $y_\ell = 0$ for the “dying” leaf ℓ . Thus, by Lemma 5, the cost of the algorithm during the growth of \tilde{w}_ℓ is bounded by the decrease of P during that time. Clearly, P can only decrease by a finite amount (as it remains non-negative) and thus the algorithm’s cost is finitely bounded. But this means that the mass at ℓ must tend to 0, since otherwise the service cost would be infinite. Moreover, notice that the growth of \tilde{w}_ℓ is just a simulation and the algorithm doesn’t pay the service cost, only the cost of moving from its state x at the start of the simulation to the limit state \bar{x} . However, this movement cost is at most the total cost to the algorithm during the simulation, and P decreases by at least the total cost. Finally, at \bar{x} , the term in P for ℓ equals 0, so removing it does not increase P . Also, for every vertex u in a subtree rooted at a sibling of ℓ the term δ_u increases (as the degree d_{p_ℓ} decreases by 1). However, this too cannot increase P (as $x_u \leq 1$). \square

4) Merge step:

Lemma 8. *A merge step does not increase P .*

Proof. Let j be the step number in which the merge happens. Substituting the expression for the revised weights, the potential P can be written as

$$P = 2 \sum_{u \in V^0} (w_u + 2^{-ju}) \left(\frac{2k-1}{2k-h_u} 4ky_u \log \frac{1+\delta_u}{x_u+\delta_u} + (2k-1)x_u \right).$$

Consider the two edges $\{c, v\}$ and $\{v, p_v\}$ that are to be merged, where $v = p_c(j, 0)$. Firstly, for each vertex u in the subtree of c (including c itself), its depth h_u decreases by 1. This cannot increase P . Notice also that as $d_v = 2$, we have $\delta_c = \delta_v$ and the merge does not change any δ_u . The new value $h_c(j+1, 0)$ equals the old value $h_v(j, 0)$. Note also that $y_c = y_v$ and $x_c = x_v$ because c is the only child of v . Thus, merging the two edges of lengths w_c and w_v into a single edge of length $w_c + w_v$, and removing vertex v , only leads to a further decrease in P resulting from the disappearance of the 2^{-jv} term. \square

5) Putting it together:

Proof of Theorem 1. By Lemmas 5, 6, 7 and 8,

$$C \leq O(k^2 \log d_{\max}) \text{opt} + P_0 - P_f + \varepsilon \cdot O(k^2 \log d_{\max}),$$

where P_0 and P_f are the initial and final value of P , respectively. Now, observe that $P_0 = \varepsilon \cdot O(k)$ and $P_f \geq 0$. \square

III. DERIVATION OF ALGORITHM AND POTENTIAL FUNCTION

We now describe how we derived the algorithm and potential function from the last section, and justify the existence of λ .

A. Online mirror descent

Our algorithm is based on the online mirror descent framework of [10], [11]. In general, an algorithm in this framework is specified by a convex body $K \subset \mathbb{R}^n$, a suitable strongly convex function $\Phi: K \rightarrow \mathbb{R}$ (called *regularizer*) and a map $f: [0, \infty) \times K \rightarrow \mathbb{R}^n$ (called *control function*). The algorithm corresponding to K , Φ and f is the (usually unique) solution $x: [0, \infty) \rightarrow K$ to the following differential inclusion:

$$\nabla^2 \Phi(x(t)) \cdot x'(t) \in f(t, x(t)) - N_K(x(t)), \quad (13)$$

where $\nabla^2 \Phi(x)$ denotes the Hessian of Φ at x and

$$N_K(x) := \{\mu \in \mathbb{R}^n : \langle \mu, y - x \rangle \leq 0, \forall y \in K\}$$

is the normal cone of K at $x(t)$. Intuitively, (13) means that x tries to move in direction $f(t, x(t))$, with the normal cone term $N_K(x(t))$ ensuring that $x(t) \in K$ can be maintained, and multiplication by the positive definite matrix $\nabla^2 \Phi(x(t))$ corresponding to a distortion of the direction in which x is moving.

A benefit of the online mirror descent framework is that there exists a default potential function for its analysis, namely the Bregman divergence associated to Φ , defined as

$$D_\Phi(y||x) := \Phi(y) - \Phi(x) + \langle \nabla \Phi(x), x - y \rangle$$

for $x, y \in K$. Plugging in $x = x(t)$, the change of the Bregman divergence as a function of time is

$$\frac{d}{dt} D_\Phi(y||x(t)) = \langle \nabla^2 \Phi(x(t)) \cdot x'(t), x(t) - y \rangle \quad (14)$$

$$= \langle f(t, x(t)) - \mu(t), x(t) - y \rangle$$

for some $\mu(t) \in N_K(x(t))$ (15)

$$\leq \langle f(t, x(t)), x(t) - y \rangle, \quad (16)$$

where (14) follows from the definition of D_Φ and the chain rule, (15) follows from (13), and (16) follows from the definition of $N_K(x(t))$.

B. Charging service cost for evolving trees

In the evolving tree game, we have $K = K(T)$. For a continuous step, it would seem natural to choose $f(t, x(t)) = -w'(t)$, so that (16) implies that the online service cost $\langle w'(t), x(t) \rangle$ plus change in the potential $D_\Phi(y||x(t))$ is at most the offline service cost $\langle w'(t), y \rangle$. For the regularizer Φ (which should be chosen in a way that also allows to bound the movement cost later), the choice analogous to [10] would be

$$\Phi(x) := \sum_{u \in V^0} w_u(x_u + \delta_u) \log(x_u + \delta_u).$$

However, since Φ (and thus D_Φ) depends on w , the evolution of w leads to an additional change of D_Φ , which the bound (16) does not account for as it assumes the regularizer Φ to be fixed. To determine this additional change, first observe that by a simple calculation

$$D_\Phi(y||x) = \sum_{u \in V^0} w_u \left[(y_u + \delta_u) \log \frac{y_u + \delta_u}{x_u + \delta_u} + x_u - y_u \right].$$

When w_ℓ increases at rate 1, this potential increases at rate

$$(y_\ell + \delta_\ell) \log \frac{y_\ell + \delta_\ell}{x_\ell + \delta_\ell} + x_\ell - y_\ell.$$

The good news is that the part

$$y_\ell \log \frac{y_\ell + \delta_\ell}{x_\ell + \delta_\ell} \leq y_\ell \cdot O(\log \frac{1}{\delta_\ell}) \leq O(k)y_\ell$$

can be charged to the offline service cost, which increases at rate y_ℓ . The term $-y_\ell$ also does no harm as it is non-positive. The term x_ℓ might seem to be a slight worry, because it is equal to the online service cost, which is precisely the quantity that the change in potential is supposed to cancel. It means that effectively we have to cancel two times the online service cost, which can be achieved by accelerating the movement of the algorithm by a factor 2 (by multiplying the control function by a factor 2). The main worry is the remaining term $\delta_\ell \log \frac{y_\ell + \delta_\ell}{x_\ell + \delta_\ell}$, which does not seem controllable. We would therefore prefer to have a potential that has the δ_u terms only inside but not outside the log.

Removing this term (and, for simplicity, omitting the y_u at the end of the potential, which does not play any important role), our desired potential would then be a sum of the following two terms $L(t)$ and $M(t)$:

$$\begin{aligned} L(t) &:= \sum_{u \in V^0} w_u(t) y_u \log \frac{y_u + \delta_u}{x_u(t) + \delta_u} \\ M(t) &:= \sum_{u \in V^0} w_u(t) x_u(t) \end{aligned}$$

Let us study again why these terms are useful as part of the classical Bregman divergence potential by calculating their change. Dropping t from the notation, and using that $\nabla^2 \Phi(x)$ is the diagonal matrix with entries $\frac{w_u}{x_u + \delta_u}$, we have

$$\begin{aligned} L' &= \langle w', y \rangle O(k) - \langle y, \nabla^2 \Phi(x) \cdot x' \rangle \\ &= \langle w', y \rangle O(k) - \langle y, f - \mu \rangle \end{aligned}$$

and

$$\begin{aligned} M' &= \langle w', x \rangle + \langle w, x' \rangle \\ &= \langle w', x \rangle + \langle x + \delta, \nabla^2 \Phi(x) \cdot x' \rangle \\ &= \langle w', x \rangle + \langle x + \delta, f - \mu \rangle \end{aligned}$$

for some $\mu \in N_K(x)$.

For a convex body K of the form $K = \{x \in \mathbb{R}^n : Ax \leq b\}$ where $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, the normal cone is given by

$$N_K(x) = \{A^T \lambda \mid \lambda \in \mathbb{R}_+^m, \langle \lambda, Ax - b \rangle = 0\}.$$

The entries of λ are called *Lagrange multipliers*. In our case, we will have $x_u > 0$ for all $u \in V^0$, so the Lagrange multipliers corresponding to the constraints $x_u \geq 0$ will be zero. So the only tight constraints are the equality constraints, and the normal cone corresponding to any such x is given by

$$N_K(x) = \{(\lambda_u - \lambda_{p(u)})_{u \in V^0} \mid \lambda \in \mathbb{R}^V, \lambda_u = 0 \forall u \in \mathcal{L}\}. \quad (17)$$

Since $\delta \in K$ and $\delta_u > 0$ for all u , we thus have $N_K(x) = N_K(\delta)$. Hence, we can cancel the μ terms in L' and M' by taking the potential $D = 2L + M$, so that

$$\begin{aligned} D' &= 2L' + M' = \langle w', y \rangle O(k) + \langle w', x \rangle + \langle x + \delta - 2y, f - \mu \rangle \\ &\leq \langle w', y \rangle O(k) + \langle w', x \rangle + \langle x + \delta - 2y, f \rangle, \end{aligned}$$

where the inequality uses that $\mu \in N_K(x)$ and $\mu \in N_K(\delta)$. Recalling that $w'_u = \mathbb{1}_{u=\ell}$, and choosing $f = -2w'$ as the control function, we get

$$D' \leq y_\ell O(k) - x_\ell, \quad (18)$$

i.e., the potential charges the online service cost to $O(k)$ times the offline service cost. Indeed, the potential function D used in Section II-C1 is given by $D = 2L + M$, up to the replacement of w by \tilde{w} . Moreover we note that (18) remains true with the control function $f = -\frac{2x_\ell}{x_\ell + \delta_\ell} w'$, which will be helpful for the movement as we discuss next.

C. Bounding movement via damped control and revised weights

Besides the service cost, we also need to bound the movement cost. By (13) and (17) and since $\nabla^2 \Phi(x)$ is the diagonal matrix with entries $\frac{w_u}{x_u + \delta_u}$, the movement of the algorithm satisfies

$$\begin{aligned} w_u x'_u &= (x_u + \delta_u)(f_u + \lambda_{p(u)} - \lambda_u) \\ &= -2x_u w'_u + (x_u + \delta_u)(\lambda_{p(u)} - \lambda_u), \end{aligned} \quad (19)$$

where the last equation uses $f = -\frac{2x_\ell}{x_\ell + \delta_\ell} w'$. Up to the discrepancy between w and \tilde{w} , this is precisely Equation (2). Here, damping the control function f by the factor $\frac{x_\ell}{x_\ell + \delta_\ell}$ is crucial: Otherwise there would be additional movement of the form $\delta_\ell w'_\ell$. Although a similar δ -induced term in the movement exists also in [10], the argument in [10] to control such a term relies on w being fixed and would therefore fail in our case. Scaling by $\frac{x_\ell}{x_\ell + \delta_\ell}$ prevents such movement from occurring in the first place.

To bound the movement cost, [10] employs a *weighted depth potential* defined as

$$\Psi = \sum_{u \in V^0} h_u w_u x_u.$$

Our calculation in Lemma 3 suggests that we can use the same Ψ here, choosing the overall potential function as $P = 4kD - 2\Psi$. But now the problem is that the combinatorial depths h_u can change during merge steps, which would lead to an increase of the overall potential P . To counteract this, we use the revised weights \tilde{w}_u : The scaling by $\frac{2k-1}{2k-h_u}$ in their definition means that \tilde{w}_u slightly increases whenever h_u decreases, and overall this ensures that the potential P does not increase in a merge step. Since $\frac{2k-1}{2k-h_u} \in [1, 2]$, such scaling loses only a constant factor in the competitive ratio. The additional term 2^{-j_u} in the definition of the revised weights only serves to ensure that $\tilde{w}_u > 0$, so that Φ is strongly convex as required by the mirror descent framework.

D. Existence of the mirror descent path for time-varying Φ_t

To justify the existence of our algorithm, we need the following theorem, which generalizes [11, Theorem 2.2] to the setting where a fixed regularizer Φ is replaced by a time-varying regularizer Φ_t .

Theorem 2. *Let $K \subset \mathbb{R}^n$ be compact and convex, $f: [0, \infty) \times K \rightarrow \mathbb{R}^n$ continuous, $\Phi_t: K \rightarrow \mathbb{R}$ strongly convex for $t \geq 0$ and such that $(x, t) \mapsto \nabla^2 \Phi_t(x)^{-1}$ is continuous. Then for any $x_0 \in K$ there is an absolutely continuous solution $x: [0, \infty) \rightarrow K$ satisfying*

$$\begin{aligned} \nabla^2 \Phi_t(x(t)) \cdot x'(t) &\in f(t, x(t)) - N_K(x(t)) \\ x(0) &= x_0. \end{aligned}$$

If $(x, t) \mapsto \nabla^2 \Phi_t(x)^{-1}$ is Lipschitz and f locally Lipschitz, then the solution is unique.

Proof. It suffices to consider a finite time interval $[0, \tau]$. It was shown in [11, Theorem 5.7] that for $C \subset \mathbb{R}^n$ compact and convex, $H: C \rightarrow \{A \in \mathbb{R}^{n \times n} \mid A \succ 0\}$ continuous, $g: [0, \tau] \times C \rightarrow \mathbb{R}^n$ continuous and $y_0 \in C$, there is an absolutely continuous solution $y: [0, \tau] \rightarrow C$ satisfying

$$\begin{aligned} y'(t) &\in H(y(t)) \cdot (g(t, y(t)) - N_C(y(t))) \\ y(0) &= y_0. \end{aligned}$$

We choose $C = [-1, \tau + 1] \times K$,

$$H(t, x) = \begin{bmatrix} 1 & 0 \\ 0 & \nabla^2 \Phi_t(x)^{-1} \end{bmatrix},$$

$g(t, (s, x)) = (1, f(t, x))$ and $y_0 = (0, x_0)$. Decomposing the solution as $y(t) = (s(t), x(t))$ for $s(t) \in [-1, \tau + 1]$ and $x(t) \in K$, and noting that for $s(t) \in [0, \tau]$ we have $N_C(y(t)) = \{0\} \times N_K(x(t))$, we get

$$\begin{aligned} s(t) &= t \\ x'(t) &\in \nabla^2 \Phi_t(x)^{-1} \cdot (f(t, x(t)) - N_K(x(t))) \\ x(0) &= x_0. \end{aligned}$$

By [11, Theorem 5.9] the solution is unique provided H is Lipschitz and g locally Lipschitz, which is satisfied under the additional assumptions of the theorem. \square

For every continuous step,

$$\Phi_t = \sum_{u \in V^0} \tilde{w}_u(t) (x_u + \delta_u) \log(x_u + \delta_u)$$

and

$$f(t, x) = -\frac{2x_\ell(t)}{x_\ell(t) + \delta_\ell} w'(t)$$

satisfy the assumptions of the theorem. By the calculation in (19) (with w replaced by \tilde{w}), the corresponding well-defined algorithm is the one from equation (2). Note that Lagrange multipliers for the constraints $x_u \geq 0$ are indeed not needed (see below).

Sign of Lagrange multipliers: We stipulated in Section II-B that $\lambda_u \geq 0$ for $u \in V \setminus \mathcal{L}$, and we do not have any Lagrange multipliers for the constraints $x_u \geq 0$. To see this, it suffices to show that $\lambda_u \geq 0$ for $u \in V \setminus \mathcal{L}$ in the case that the constraints $x_u \geq 0$ are removed from K : If this is true, then (2) shows for any leaf $u \in \mathcal{L}$ that $x'_u < 0$ is possible only if $x_u > 0$ (since $\lambda_u = 0$ when u is a leaf, recalling (17)). Hence, $x_u \geq 0$ holds automatically for any leaf u , and thus also for internal vertices u due to the constraints of the polytope. Consequently, we do not need Lagrange multipliers for constraints $x_u \geq 0$.

The proof that $\lambda_u \geq 0$ for $u \in V \setminus \mathcal{L}$ is completely analogous to [4, Lemma 12]. As an alternative proof, one can also see this by replacing in K the constraints $\sum_{v: p(v)=u} x_v = x_u$ by $\sum_{v: p(v)=u} x_v \geq x_u$, which directly gives $\lambda_u \geq 0$ in (17); this still yields a feasible solution (with the constraints satisfied with equality) by arguments completely analogous to [10, Lemma 3.2].

IV. REDUCTIONS AND APPLICATIONS

In this section we show that layered graph traversal and small set chasing (a.k.a. metrical service systems) reduce to the evolving tree game. The reductions imply the following new bounds on the competitive ratio for these problems, the main result of this section.

Theorem 3. *There are randomized $O(k^2)$ -competitive online algorithms for traversing width k layered graphs, as well as for chasing sets of cardinality k in any metric space.*

A. Layered graph traversal

Recall the definition of the problem in the introduction. We will introduce useful notation. Let

$$V_0 = \{a\}, V_1, V_2, \dots, V_n = \{b\}$$

denote the layers of the input graph G , in consecutive order. Let E_1, E_2, \dots, E_n be the partition of the edge set of G , where for every $i = 1, 2, \dots, n$, every edge $e \in E_i$ has one endpoint in V_{i-1} and one endpoint in V_i . Also recall that $w: E \rightarrow \mathbb{N}$ is the weight function on the edges, and

$k = \max\{|V_i| : i = 0, 1, 2, \dots, n\}$ is the *width* of G . The input G is revealed gradually to the searcher. Let

$$G_i = (V_0 \cup V_1 \cup \dots \cup V_i, E_1 \cup E_2 \cup \dots \cup E_i)$$

denote the subgraph that is revealed up to and including step i . The searcher, currently at a vertex $v_{i-1} \in V_{i-1}$ chooses a path in G_i from v_{i-1} to a vertex $v_i \in V_i$. Let $w_{G_i}(v_{i-1}, v_i)$ denote the total weight of a shortest path from v_{i-1} to v_i in G_i . (Clearly, the searcher has no good reason to choose a longer path.) Formally, a pure strategy (a.k.a. deterministic algorithm) of the searcher is a function that maps, for all $i = 1, 2, \dots$, a layered graph G_i (given including its partition into a sequence of layers) to a vertex in V_i (i.e., the searcher's next move). A mixed strategy (a.k.a. randomized algorithm) of the searcher is a probability distribution over such functions.

1) Fractional strategies:

Given a mixed strategy S of the searcher, we can define a sequence P_0, P_1, P_2, \dots , where P_i is a probability distribution over V_i . For every $v \in V_i$, $P_i(v)$ indicates the probability that the searcher's mixed strategy S chooses to move to v (i.e., $v_i = v$). A fractional strategy of the searcher is a function that maps, for all $i = 1, 2, \dots$, a layered graph G_i to a probability distribution P_i over V_i . For a fractional strategy choosing probability distributions $P_0, P_1, P_2, \dots, P_n$, we define its cost as follows. For $i = 1, 2, \dots, n$, let τ_i be a probability distribution over $V_{i-1} \times V_i$, with marginals P_{i-1} on V_{i-1} and P_i on V_i , that minimizes

$$w_{G_i, \tau_i}(P_{i-1}, P_i) = \sum_{u \in V_{i-1}} \sum_{v \in V_i} w_{G_i}(u, v) \tau_i(u, v).$$

The cost of the strategy is then defined as $\sum_{i=1}^n w_{G_i, \tau_i}(P_{i-1}, P_i)$.

The following lemma can be deduced through the reduction to small set chasing discussed later, the fact that small set chasing is a special case of metrical task systems, and a similar known result for metrical task systems. Here we give a straightforward direct proof.

Lemma 9. *For every fractional strategy of the searcher there is a mixed strategy incurring the same cost.*

Proof. Fix any fractional strategy of the searcher, and suppose that the searcher plays $P_0, P_1, P_2, \dots, P_n$ against a strategy G_n of the designer. I.e., the designer chooses the number of rounds n , and plays in round i the last layer of

$$G_i = (\{a\} \cup V_1 \cup V_2 \cup \dots \cup V_i, E_1 \cup E_2 \cup \dots \cup E_i).$$

The searcher responds with P_i , which is a function of G_i . Notice that when the designer reveals G_i , the searcher can compute τ_i , because that requires only the distance functions w_{G_i} and the marginal probability distributions P_{i-1} and P_i . We construct a mixed strategy of the searcher inductively as follows. It is sufficient to define, for every round i , the conditional probability distribution on the searcher's next move $v_i \in V_i$, given any possible play so far. Initially, at the

start of round 1, the searcher is deterministically at a . Suppose that the searcher reached a vertex $v_{i-1} \in V_{i-1}$. Then, we set

$$\Pr[v_i = v \in V_i \mid v_{i-1}] = \frac{\tau_i(v_{i-1}, v)}{P_{i-1}(v_{i-1})}.$$

Notice that the searcher can move from v_{i-1} to v_i along a path in G_i of length $w_{G_i}(v_{i-1}, v_i)$.

We now analyze the cost of the mixed strategy thus defined. We prove by induction over the number of rounds that in round i , for every pair of vertices $u \in V_{i-1}$ and $v \in V_i$, the probability that the searcher's chosen pure strategy (which is a random variable) reaches v is $P_i(v)$ and the probability that this strategy moves from u to v is $\tau_i(u, v)$ (the latter assertion is required to hold for $i > 0$). The base case is $i = 0$, which is trivial, as the searcher's initial position is a , $P_0(a) = 1$, and the statement about τ is vacuous. So, assume that the statement is true for $i - 1$. By the definition of the strategy, in round i , for every $v \in V_i$,

$$\begin{aligned} \Pr[v_i = v] &= \sum_{u \in V_{i-1}} \Pr[v_{i-1} = u] \cdot \Pr[v_i = v \mid v_{i-1} = u] \\ &= \sum_{u \in V_{i-1}} P_{i-1}(u) \cdot \frac{\tau_i(u, v)}{P_{i-1}(u)} \\ &= P_i(v), \end{aligned}$$

where the penultimate equality uses the induction hypothesis, and the final equality uses the condition on the marginals of τ_i at V_i . Similarly,

$$\begin{aligned} &\Pr[\text{the searcher moves from } u \text{ to } v] \\ &= \Pr[v_{i-1} = u] \cdot \Pr[\text{the searcher moves from } u \text{ to } v \mid v_{i-1} = u] \\ &= P_{i-1}(u) \cdot \frac{\tau_i(u, v)}{P_{i-1}(u)} \\ &= \tau_i(u, v). \end{aligned}$$

Thus, by linearity of expectation, the searcher's expected total cost is

$$\sum_{i=1}^n \sum_{u \in V_{i-1}} \sum_{v \in V_i} \tau_i(u, v) \cdot w_{G_i}(u, v),$$

and this is by definition equal to the cost of the searcher's fractional strategy. \square

2) Layered trees:

We now discuss special cases of layered graph traversal whose solution implies a solution to the general case. We begin with a definition.

Definition 1. *A rooted layered tree is an acyclic layered graph, where every vertex $v \neq a$ has exactly one neighbor in the preceding layer. We say that a is the root of such a tree.*

Theorem 4 (Fiat et al. [18, Section 2]). *Suppose that the designer is restricted to play a width k rooted layered tree with edge weights in $\{0, 1\}$, and suppose that there is a C -competitive (pure or mixed) strategy of the searcher for this restricted game. Then, there is a C -competitive (pure or mixed, respectively) strategy of the searcher for the general case,*

where the designer can play any width k layered graph with non-negative integer edge weights.

A width k rooted layered tree is binary iff every vertex has at most two neighbors in the following layer. (Thus, the degree of each node is at most 3.)

Corollary 1. *The conclusion of Theorem 4 holds if there is a C -competitive strategy of the searcher for the game restricted to the designer using width k rooted layered binary trees with edge weights in $\{0, 1\}$. Moreover, the conclusion holds if in addition we require that between two adjacent layers there is at most one edge of weight 1.*

Proof. Suppose that the designer plays an arbitrary width k layered tree. The searcher converts the tree on-the-fly into a width k layered binary tree, uses the strategy for binary trees, and maps the moves back to the input tree. The conversion is done as follows. Between every two layers that the designer generates, the searcher simulates $\lceil \log_2 k \rceil - 1$ additional layers. If a vertex $u \in V_{i-1}$ has $m \leq k$ neighbors $v_1, v_2, \dots, v_m \in V_i$, the searcher places in the simulated layers between V_{i-1} and V_i a layered binary tree rooted at u with v_1, v_2, \dots, v_m as its leaves. Notice that this can be done simultaneously for all such vertices in V_{i-1} without violating the width constraint in the simulated layers. The lengths of the new edges are all 0, except for the edges touching the leaves. For $j = 1, 2, \dots, m$, the edge touching v_j inherits the length $w(\{u, v_j\})$ of the original edge $\{u, v_j\}$. Clearly, any path traversed in the simulated tree corresponds to a path traversed in the input tree that has the same cost—simply delete from the path in the simulated tree the vertices in the simulated layers; the edges leading to them all have weight 0. The additional requirement is easily satisfied by now making the following change. Between every two layers $i - 1, i$ of the rooted layered binary tree insert $k - 1$ simulated layers. Replace the j -th edge between layer $i - 1, i$ (edges are indexed arbitrarily) by a length k path. If the original edge has weight 0, all the edges in the path have weight 0. If the original edge has weight 1, then all the edges in the path have weight 0, except for the j -th edge that has weight 1. \square

B. Small set chasing

This two-person game is defined with respect to an underlying metric space $\mathcal{M} = (X, \text{dist})$. The game alternates between the adversary and the algorithm. The latter starts at an arbitrary point $x_0 \in X$. The adversary decides on the number of rounds n that the game will be played (this choice is unknown to the algorithm until after round n). In round i of the game, the adversary chooses a finite set $X_i \subset X$. The algorithm must then move to a point $x_i \in X_i$. The game is parametrized by an upper bound k on $\max_{i=1}^n |X_i|$. The algorithm pays $\sum_{i=1}^n \text{dist}(x_{i-1}, x_i)$ and the adversary pays

$$\min \left\{ \sum_{i=1}^n \text{dist}(y_{i-1}, y_i) : y_0 = x_0 \wedge y_1 \in X_1 \wedge \dots \wedge y_n \in X_n \right\}.$$

Theorem 5 (Fiat et al. [18, Theorem 18]). *For every $k \in \mathbb{N}$ and for every $C = C(k)$, there exists a pure (mixed, re-*

spectively) C -competitive online algorithm for cardinality k set chasing in every metric space with integral distances iff there exists a pure (mixed, respectively) C -competitive online algorithm for width k layered graph traversal.

C. Reduction to evolving trees

The main result of this section, Theorem 3, is implied by the following reduction.

Lemma 10. *Let $k \in \mathbb{N}$, let $C = C(k)$, and let $\varepsilon > 0$. Suppose that there exists a (pure, mixed, fractional, respectively) C -competitive strategy for the evolving tree game on binary trees of maximum depth k that always pays a cost of at most $C \cdot (\text{opt} + \varepsilon)$. Then, there exists a (pure, mixed, fractional, respectively) C -competitive strategy for traversing width k layered graphs with minimum non-zero edge weight at least ε .⁹*

Proof. Consider at first fractional strategies. By Lemma 9 and Corollary 1, we can restrict our attention to designing fractional strategies on width k rooted layered binary trees with edge weights in $\{0, 1\}$. Now, suppose that we have a fractional C -competitive strategy for the depth k evolving tree game. We use it to construct a fractional C -competitive strategy for the traversal of width k rooted layered binary trees as follows. To simplify the proof, add a virtual layer -1 containing a single node $p(a)$ connected to the source a with an edge of weight 0. We construct the layered graph strategy by induction over the current layer. Our induction hypothesis is that in the current state:

- 1) The evolving tree is homeomorphic to the layered subtree spanned by the paths from $p(a)$ to the nodes in the current layer.
- 2) In this homeomorphism, r is mapped to $p(a)$ and the leaves of the evolving tree are mapped to the leaves of the layered subtree, which are the nodes in the current layer.
- 3) In this homeomorphism, each edge of the evolving tree is mapped to a path of the same weight in the layered subtree.
- 4) The probability assigned to a leaf by the fractional strategy for the evolving tree is equal to the probability assigned to its homeomorphic image (a node in the current layer) by the fractional strategy for the layered tree.

Initially, the traversal algorithm occupies the source node a with probability 1. The evolving tree consists of the two initial nodes r and c_r , with a 0-weight edge connecting them. The homeomorphism maps r to $p(a)$ and c_r to a . The evolving tree algorithm occupies c_r with probability 1. Hence, the induction hypothesis is satisfied at the base of the induction. For the inductive step, consider the current layer $i - 1$, the new layer i and edges between them. If a node in layer $i - 1$ has no child in layer i , we delete the homeomorphic preimage (which must be a leaf and cannot be c_r) in the evolving tree. If a node v

⁹In particular, if the edge weights are integers, one can take $\varepsilon = 1$.

in layer $i - 1$ has two children in layer i , we execute a fork step where we generate two new leaves and connect them to the preimage of v (a leaf) in the evolving tree, and extend the homeomorphism to the new leaves in the obvious way. Otherwise, if a node v in layer $i - 1$ has a single child in layer i , we modify the homeomorphism to map the preimage of v to its child in layer i . After executing as many such discrete steps as needed, if there is a weight 1 edge connecting a node u in layer $i - 1$ to a node v in layer i , we execute a continuous step, increasing the weight of the edge incident on the homeomorphic preimage of v in the evolving tree (which must be a leaf) for a time interval of length 1. After executing all these steps, we simply copy the probability distribution on the leaves of the evolving tree to the homeomorphic images in layer i of the layered tree. This clearly satisfies all the induction hypotheses at layer i .

Notice that since the target b is assumed to be the only node in the last layer, when we reach it, the evolving tree is reduced to a single edge connecting r to the homeomorphic preimage of b . The weight of this edge equals the weight of the path in the layered tree from $p(a)$ to b , which is the same as the weight of the path from a to b (because the edge $\{p(a), a\}$ has weight 0). Moreover, the fractional strategy that is induced in the layered tree does not pay more than the fractional strategy in the evolving tree. Hence, it is C -competitive.

Finally, deterministic strategies are fractional strategies restricted to probabilities in $\{0, 1\}$, hence the claim for deterministic strategies is a corollary. This also implies the claim for mixed strategies, as they are probability distributions over pure strategies. \square

We note that the depth k evolving tree game is strictly more general than width k layered graph traversal. In particular, the evolving binary tree corresponding to the width k layered graph game has depth at most k and also at most k leaves. However, in general a depth k binary tree may have 2^{k-1} leaves. Our evolving tree algorithm and analysis applies without further restriction on the number of leaves.

REFERENCES

- [1] A. Antoniadis, C. Coester, M. Eliás, A. Polak, and B. Simon. Online metric algorithms with untrusted predictions. In *Proc. of the 37th Int'l Conf. on Machine Learning*, pages 345–355, 2020.
- [2] Y. Azar, A. Z. Broder, and M. S. Manasse. On-line choice of on-line algorithms. In *Proc. of the 4th Ann. ACM-SIAM Symp. on Discrete Algorithms*, pages 432–440, 1993.
- [3] R. A. Baeza-Yates, J. C. Culberson, and G. J. E. Rawlins. Searching with uncertainty. In *Proc. of Scandinavian Workshop on Algorithm Theory*, pages 176–189, 1988.
- [4] N. Bansal and C. Coester. Online metric allocation and time-varying regularization. In *30th Ann. European Symp. on Algorithms*, pages 13:1–13:13, 2022.
- [5] N. Bansal, C. Coester, R. Kumar, M. Purohit, and E. Vee. Learning-augmented weighted paging. In *Proc. of the 33rd Ann. ACM-SIAM Symp. on Discrete Algorithms*, 2022.
- [6] Y. Bartal, B. Bollobas, and M. Mendel. A Ramsey-type theorem for metric spaces and its applications for metrical task systems and related problems. In *Proc. of the 42nd Ann. IEEE Symp. on Foundations of Computer Science*, pages 396–405, 2001.
- [7] Y. Bartal, N. Linial, M. Mendel, and A. Naor. On metric Ramsey-type phenomena. In *Proc. of the 35th Ann. ACM Symp. on Theory of computing*, pages 463–472, June 2003.
- [8] R. Bellman. *Dynamic Programming*. Princeton University Press, 1957.
- [9] A. Borodin, N. Linial, and M. E. Saks. An optimal online algorithm for metrical task systems. In *Proc. of the 19th Ann. ACM Symp. on Theory of Computing*, pages 373–382, 1987.
- [10] S. Bubeck, M. B. Cohen, J. R. Lee, and Y.-T. Lee. Metrical task systems on trees via mirror descent and unfair gluing. In *Proc. of the 30th Ann. ACM-SIAM Symp. on Discrete Algorithms*, pages 89–97, 2019.
- [11] S. Bubeck, M. B. Cohen, Y.-T. Lee, J. R. Lee, and A. Madry. k -server via multiscale entropic regularization. In *Proc. of the 50th Ann. ACM Symp. on Theory of Computing*, pages 3–16, 2018.
- [12] W. R. Burley. Traversing layered graphs using the work function algorithm. *J. Algorithms*, 20(3):479–511, 1996.
- [13] M. Chrobak and L. Larmore. The server problem and on-line games. In *On-Line Algorithms: Proc. of a DIMACS Workshop*, pages 11–64. American Math. Society, 1991.
- [14] M. Chrobak and L. Larmore. Metrical service systems: deterministic strategies. Technical Report UCR-CS-93-1, UC Riverside, 1993.
- [15] C. Coester and E. Koutsoupias. The online k -taxi problem. In *Proc. of the 51st Ann. ACM Symp. on Theory of Computing*, pages 1136–1147, 2019.
- [16] C. Coester and J. R. Lee. Pure entropic regularization for metrical task systems. In *Proc. of the 32nd Conf. on Learning Theory*, pages 835–848, 2019.
- [17] F. Ebrahimnejad and J. R. Lee. Multiscale entropic regularization for MTS on general metric spaces. In *13th Innovations in Theoretical Computer Science Conference*, pages 60:1–60:21, 2022.
- [18] A. Fiat, D. P. Foster, H. J. Karloff, Y. Rabani, Y. Ravid, and S. Vishwanathan. Competitive algorithms for layered graph traversal. In *Proc. of the 32nd Ann. IEEE Symp. on Foundations of Computer Science*, pages 288–297, 1991.
- [19] A. Fiat, Y. Rabani, and Y. Ravid. Competitive k -server algorithms. In *Proc. of the 31st Ann. IEEE Symp. on Foundations of Computer Science*, pages 454–463, 1990.
- [20] A. Fiat, Y. Rabani, Y. Ravid, and B. Schieber. A deterministic $O(k^3)$ -competitive k -server algorithm for the circle. *Algorithmica*, 11:572–578, 1994.
- [21] S. Gal. *Search Games*. Academic Press, New York, NY, USA, 1980.
- [22] M.-Y. Kao, Y. Ma, M. Sipser, and Y. Yin. Optimal constructions of hybrid algorithms. In *Proc. of the 5th Ann. ACM-SIAM Symp. on Discrete Algorithms*, pages 372–381, 1994.
- [23] M.-Y. Kao, J. H. Reif, and S. R. Tate. Searching in an unknown environment: An optimal randomized algorithm for the cow-path problem. In *Proc. of the 4th Ann. ACM-SIAM Symp. on Discrete Algorithms*, pages 441–447, 1993.
- [24] J. R. Lee. Fusible HSTs and the randomized k -server conjecture. In *Proc. of the 59th Ann. IEEE Symp. on Foundations of Computer Science*, pages 438–449, 2018.
- [25] T. Lykouris and S. Vassilvitskii. Competitive caching with machine learned advice. *J. ACM*, 68(4):24:1–24:25, 2021.
- [26] C. H. Papadimitriou and M. Yannakakis. Shortest paths without a map. In *Proc. of the 16th Int'l Colloq. on Automata, Languages and Programming*, pages 610–620, 1989.
- [27] H. Ramesh. On traversing layered graphs on-line. In *Proc. of the 4th Ann. ACM-SIAM Symp. on Discrete Algorithms*, pages 412–421, 1993.
- [28] A. Schrijver. On the history of the shortest path problem. In *Documenta Mathematica*, Extra Volume: Optimization Stories, pages 155–167, 2012.