THEORY OF COMPUTING, Volume 18 (23), 2022, pp. 1–24
www.theoryofcomputing.org

Pure Entropic Regularization for Metrical Task Systems

Christian Coester* James R. Lee†

Received July 23, 2019; Revised October 8, 2020; Published December 29, 2022

Abstract. We show that on every n-point HST metric, there is a randomized online algorithm for metrical task systems (MTS) that is 1-competitive for service costs and $O(\log n)$-competitive for movement costs. In general, these refined guarantees are optimal up to the implicit constant. While an $O(\log n)$-competitive algorithm for MTS on HST metrics was developed by Bubeck et al. (SODA'19), that approach could only establish an $O((\log n)^2)$-competitive ratio when the service costs are required to be $O(1)$-competitive. Our algorithm can be viewed as an instantiation of online mirror descent with the regularizer derived from a multiscale conditional entropy.

In fact, our algorithm satisfies a set of even more refined guarantees; we are able to exploit this property to combine it with known random embedding theorems and obtain, for *any* n-point metric space, a randomized algorithm that is 1-competitive for service costs and $O((\log n)^2)$-competitive for movement costs.

An extended abstract of this paper appeared in the Proceedings of the 32nd Ann. Conference on Learning Theory (COLT 2019) [17].

*Supported by EPSRC Award 1652110. Part of this work was carried out while C. Coester was visiting the University of Washington, hosted by J. R. Lee.

†Supported by NSF grants CCF-1616297 and CCF-1407779 and a Simons Investigator Award

ACM Classification: F.2.0

AMS Classification: 68W27

Key words and phrases: online algorithms, competitive analysis, mirror descent, metrical task systems, decision making under uncertainty

 DOI: 10.4086/toc.2022.v018a023

1 Introduction

Let (X, d_X) be a finite metric space with $|X| = n > 1$. The Metrical Task Systems (MTS) problem, introduced in [11] is described as follows. The input is a sequence $\langle c_t : X \to \mathbb{R}_+ : t \geq 1\rangle$ of nonnegative cost functions on the state space X. At every time t, an online algorithm maintains a state $\rho_t \in X$.

The corresponding cost is the sum of a *service cost* $c_t(\rho_t)$ and a *movement cost* $d_X(\rho_{t-1}, \rho_t)$. Formally, an *online algorithm* is a sequence of mappings $\boldsymbol{\rho} = \langle \rho_1, \rho_2, \ldots, \rangle$ where, for every $t \geq 1$, $\rho_t : (\mathbb{R}_+^X)^t \to X$ maps a sequence of cost functions $\langle c_1, \ldots, c_t\rangle$ to a state. The initial state $\rho_0 \in X$ is fixed. The *total cost of the algorithm* $\boldsymbol{\rho}$ *in servicing* $\boldsymbol{c} = \langle c_t : t \geq 1\rangle$ is defined as:

$$\mathrm{cost}_{\boldsymbol{\rho}}(\boldsymbol{c}) := \sum_{t\geq 1} \left[c_t(\rho_t(c_1, \ldots, c_t)) + d_X(\rho_{t-1}(c_1, \ldots, c_{t-1}), \rho_t(c_1, \ldots, c_t))\right].$$

The cost of the *offline optimum*, denoted $\mathrm{cost}^*(\boldsymbol{c})$, is the infimum of $\sum_{t\geq 1}[c_t(\rho_t) + d_X(\rho_{t-1}, \rho_t)]$ over *any* sequence $\langle \rho_t : t \geq 1\rangle$ of states. A *randomized online algorithm* $\boldsymbol{\rho}$ is said to be α*-competitive* if for every $\rho_0 \in X$, there is a constant $\beta > 0$ such that for all cost sequences $\boldsymbol{c}$:

$$\mathbb{E}\left[\mathrm{cost}_{\boldsymbol{\rho}}(\boldsymbol{c})\right] \leq \alpha \cdot \mathrm{cost}^*(\boldsymbol{c}) + \beta.$$

For the n-point uniform metric, a simple coupon-collector argument shows that the competitive ratio is $\Omega(\log n)$, and this is tight [11]. A long-standing conjecture is that this $\Theta(\log n)$ competitive ratio holds for an arbitrary n-point metric space. The lower bound has almost been established [8, 9]; for any n-point metric space, the competitive ratio is $\Omega(\log n/\log\log n)$. Following a long sequence of works (see, e. g., [20, 10, 7, 6, 19, 18]), an upper bound of $O((\log n)^2)$ was shown in [13].

Relation to adversarial multi-arm bandits MTS is naturally related to the adversarial setting of the classical multi-arm bandits model in sequential decision making, and provides a very general framework for "bandits with switching costs." Unlike in the setting of regret minimization, where one competes against the best static strategy in hindsight (see, e. g., [12]), competitive analysis compares the performance of an online algorithm to the best *dynamical* offline algorithm.

Thus this model emphasizes the importance of an adaptivity in the face of changing environments. For MTS, the online algorithm has *full information*: access to the complete cost function c_t is available when deciding on a point $\rho_t(c_1, \ldots, c_t) \in X$ at which to play. And yet one of the fascinating relationships between MTS and adversarial bandits is the parallel between adaptivity—being willing to "try out" new strategies—and the classical exploration/exploitation tradeoff that occurs in models where one only has access to partial information about the loss functions.

HST metrics The methods of [5] show that the competitive ratio for MTS is $O(\log n)$ on weighted star metrics. Recently, the authors of [13] generalized this result by designing an algorithm with competitive ratio $O(\mathfrak{D}_T \log n)$ on any weighted n-point tree metric with combinatorial depth $\mathfrak{D}_T$. We now discuss a special class of metrics.

Let $T = (V, E)$ be a finite tree with root $\mathbb{r}$ and vertex weights $\{w_u > 0 : u \in V\}$, let $\mathcal{L} \subseteq V$ denote the leaves of T, and suppose that the vertex weights on T are non-increasing along root-leaf paths. Consider the metric space $(\mathcal{L}, d_T)$, where $d_T(\ell, \ell')$ is the weighted length of the path connecting ℓ and ℓ' when the edge from a node u to its parent is w_u. We will use $\mathfrak{D}_T$ for the combinatorial (i. e., unweighted) depth of T.

$(\mathcal{L}, d_T)$ is called an *HST metric* (or, equivalently for finite metric spaces, an *ultrametric*). If, for some $\tau > 1$, the weights on T satisfy the stronger inequality $w_v \leq w_u/\tau$ whenever v is a child of u, the space $(\mathcal{L}, d_T)$ is said to be a *τ-HST metric.* Such metric spaces play a special role in MTS since every n-point metric space can be probabilistically approximated by a distribution over such spaces [6, 18]. Indeed, the $O((\log n)^2)$-competitive ratio for general metric spaces established in [13] is a consequence of their $O(\log n)$-competitive algorithm for HSTs.

1.1 Refined guarantees

The authors of [4] observe that there is a more refined way to analyze competive algorithms for MTS. For a randomized online algorithm $\boldsymbol{\rho}$ and a cost sequence $\boldsymbol{c}$, we denote, respectively, $\mathsf{S}_{\boldsymbol{\rho}}(\boldsymbol{c})$ and $\mathsf{M}_{\boldsymbol{\rho}}(\boldsymbol{c})$ for the (expected) service cost and movement cost, that is

$$\mathsf{S}_{\boldsymbol{\rho}}(\boldsymbol{c}) := \mathbb{E} \sum_{t \geq 1} c_t(\rho_t) \quad \text{and} \quad \mathsf{M}_{\boldsymbol{\rho}}(\boldsymbol{c}) := \mathbb{E} \sum_{t \geq 1} d_X(\rho_{t-1}, \rho_t)\,.$$

If there are numbers $\alpha, \alpha', \beta, \beta' > 0$ such that for every cost $\boldsymbol{c}$, it holds that

$$\begin{aligned} \mathsf{S}_{\boldsymbol{\rho}}(\boldsymbol{c}) &\leq \alpha \cdot \mathrm{cost}^*(\boldsymbol{c}) + \beta \\ \mathsf{M}_{\boldsymbol{\rho}}(\boldsymbol{c}) &\leq \alpha' \cdot \mathrm{cost}^*(\boldsymbol{c}) + \beta', \end{aligned}$$

one says that $\boldsymbol{\rho}$ is *α-competitive for service costs* and *α'-competitive for movement costs.*

In [4], it is shown that on every n-point HST metric, and for every $\varepsilon > 0$, there is an online algorithm that is simultaneously $(1 + \varepsilon)$-competitive for service costs and $O((\log(n/\varepsilon))^2)$-competitive for movement costs. The authors of [13] improve this slightly to show that actually there is an online algorithm that is simultaneously 1-competitive for service costs and $O((\log n)^2)$-competitive for movement costs. We obtain the optimal refined guarantees.

Theorem 1.1. *On any n-point HST metric X, there is a randomized online algorithm that is* 1*-competitive for service costs and $O(\log n)$-competitive for movement costs.*

Remark 1.2 (Optimality of the refined guarantees)**.** Any finitely competitive algorithm for MTS on an n-point uniform metric cannot be better than $\Omega(\log n)$-competitive for movement costs, regardless of its competitive ratio for service costs. This is because this lower bound holds even if the cost functions only take values 0 and ∞. Moreover, it cannot be better than 1-competitive for service costs, regardless of its competitive ratio for movement costs. To see this, consider the case where each cost function is the constant function 1.

Finely competitive guarantees Suppose that for some numbers $\alpha_0, \alpha_1, \gamma, \beta, \beta' > 0$, a randomized online algorithm $\boldsymbol{\rho}$ satisfies, for every cost $\boldsymbol{c}$ and *every* offline algorithm $\boldsymbol{\rho}^*$:

$$\mathsf{S}_\rho(\boldsymbol{c}) \le \alpha_0 \mathsf{S}_{\rho^*}(\boldsymbol{c}) + \alpha_1 \mathsf{M}_{\rho^*}(\boldsymbol{c}) + \beta \tag{1.1}$$

$$\mathsf{M}_\rho(\boldsymbol{c}) \le \gamma \mathsf{S}_\rho(\boldsymbol{c}) + \beta' . \tag{1.2}$$

In this case, we say that $\boldsymbol{\rho}$ *is* $(\alpha_0, \alpha_1, \gamma)$*-finely competitive.* We establish the following.

Theorem 1.3. *On any n-point HST metric X, for every $\kappa \ge 1$, there is an online randomized algorithm $\boldsymbol{\rho}$ that is $\big(1, 1/\kappa, O(\kappa \log n)\big)$-finely competitive. In fact, one can take $\beta = 0$ and $\beta' \le O(\kappa \operatorname{diam}(X))$.*

Combined with the random embedding from [18], this yields the following consequence for general n-point metric spaces.

Corollary 1.4. *On any n-point metric space, there is an online randomized algorithm that is 1-competitive for service costs and $O((\log n)^2)$-competitive for movement costs.*

Proof. Consider an n-point metric space (X, d_X). It is known [18] that there exists a random HST metric (T, d_T) so that $\mathcal{L}(T) = X$ and for all $x, y \in X$:

1. $\Pr[d_T(x, y) \ge d_X(x, y)] = 1$,
2. $\mathbb{E}[d_T(x, y)] \le D \cdot d_X(x, y)$,

and $D \le O(\log n)$.

Let $\boldsymbol{\rho}_T$ be the randomized algorithm for (T, d_T) guaranteed by Theorem 1.3 with $\kappa = D$. Let $\boldsymbol{\rho}$ denote the algorithm that results from sampling (T, d_T) and then using $\boldsymbol{\rho}_T$. We use M^T to denote movement cost measured in d_T and M^X for movement cost measured in d_X.

Then for any cost $\boldsymbol{c}$ and any offline algorithm $\boldsymbol{\rho}^*$, we have

$$\begin{aligned}\mathsf{S}_\rho(\boldsymbol{c}) = \mathbb{E}[\mathsf{S}_{\rho_T}(\boldsymbol{c})] &\le \mathsf{S}_{\rho^*}(\boldsymbol{c}) + \kappa^{-1}\, \mathbb{E}[\mathsf{M}^T_{\rho^*}(\boldsymbol{c})] + O(1) \\ &\le \mathsf{S}_{\rho^*}(\boldsymbol{c}) + \kappa^{-1} D \mathsf{M}^X_{\rho^*}(\boldsymbol{c}) + O(1) \\ &= \mathsf{S}_{\rho^*}(\boldsymbol{c}) + \mathsf{M}^X_{\rho^*}(\boldsymbol{c}) + O(1)\,,\end{aligned}$$

and

$$\mathsf{M}^X_\rho(\boldsymbol{c}) = \mathbb{E}[\mathsf{M}^X_{\rho_T}(\boldsymbol{c})] \le \mathbb{E}[\mathsf{M}^T_{\rho_T}(\boldsymbol{c})] \le O(\kappa \log n)\, \mathbb{E}[\mathsf{S}_{\rho_T}(\boldsymbol{c})] + O(1),$$

completing the proof. □

1.2 The fractional model on trees

We will work in the following deterministic fractional setting, which is equivalent to the randomized integral setting described earlier (see [13, §2]). The state of a fractional algorithm is given by a point in the polytope

$$\mathsf{K}_T := \left\{ x \in \mathbb{R}_+^V : x_{\mathbb{r}} = 1,\ x_u = \sum_{v \in \chi(u)} x_v \quad \forall u \in V \setminus \mathcal{L} \right\}, \tag{1.3}$$

where we use $\chi(u)$ for the set of children of u in T. For $u \neq \mathbb{r}$, we will also write $\mathsf{p}(u)$ for the parent of u in T.

A state $x \in \mathsf{K}_T$ corresponds to the situation that the state of a randomized integral algorithm is a leaf descendant of u with probability x_u. Note that K_T is simply an affine encoding of the probability simplex on $\mathcal{L}$. In the fractional setting, changing from state x to x' incurs movement cost $\|x - x'\|_{\ell_1(w)}$, where

$$\|z\|_{\ell_1(w)} := \sum_{u \in V} w_u |z_u|$$

denotes the weighted ℓ_1-norm on $\mathbb{R}^V$.

1.3 Mirror descent, metric filtrations, and regularization

Following [13], our algorithm is based on the mirror descent framework as established in [14]. This is a method for regularized online convex optimization, an approach that was previously explored for competitive analysis in [1, 15].

A central component of mirror descent is choosing the appropriate mirror map (which we will often refer to as the "regularizer"). This is a strictly convex function $\Phi : \mathsf{K}_T \to \mathbb{R}$ that endows K_T with a geometric (Riemannian) structure, specifying how to perform constrained vector flow. In other words, it specifies how one can move in a preferred direction while remaining inside K_T.

The paper [13] employs the following regularizer:

$$\Phi_0(x) := \frac{1}{\eta} \sum_{u \in V \setminus \{\mathbb{r}\}} w_u \left(x_u + \delta_u\right) \log\left(x_u + \delta_u\right) , \tag{1.4}$$

with $\eta = \Theta(\log |\mathcal{L}|)$ and $\delta_u = |\mathcal{L}_u|/|\mathcal{L}|$, where $\mathcal{L}_u$ is the set of leaves in the subtree rooted at u.

1.3.1 Metric filtrations

It is straightforward that one can think of Φ_0 as a type of multiscale entropy (this is the *negative* of the associated Shannon entropy, since we use the analyst's convention that the entropy is convex). To understand this notion, let us forget momentarily the weights on T. Then the structure of T gives a natural filtration over probability measures on the leaves $\mathcal{L}$. Suppose that $\boldsymbol{X}$ is a random variable taking values in $\mathcal{L}$ and, for $u \in V$, denote by $\mathcal{E}_u$ the event $\{\boldsymbol{X} \in \mathcal{L}_u\}$. Then the chain rule for Shannon entropy yields

$$\sum_{\ell \in \mathcal{L}} \Pr[\mathcal{E}_\ell] \log \frac{1}{\Pr[\mathcal{E}_\ell]} = \sum_{u \in V \setminus \{\mathbb{r}\}} \Pr[\mathcal{E}_u] \log \frac{\Pr[\mathcal{E}_{\mathsf{p}(u)}]}{\Pr[\mathcal{E}_u]}.$$

If we now imagine that uncertainty at higher scales is more costly than uncertainty at lower scales, then we might define an analogous *weighted* entropy by

$$\sum_{u \in V \setminus \{\mathbb{r}\}} w_u \Pr[\mathcal{E}_u] \log \frac{\Pr[\mathcal{E}_{\mathsf{p}(u)}]}{\Pr[\mathcal{E}_u]}. \tag{1.5}$$

Such a notion is natural in the context of "metric learning" problems.

Ignoring the $\{\delta_u\}$ values for a moment, consider that (1.4) is not analogous to (1.5). Indeed, it corresponds to the quantity

$$\sum_{u \in V \setminus \{\mathbb{r}\}} w_u \Pr[\mathcal{E}_u] \log \frac{1}{\Pr[\mathcal{E}_u]}, \tag{1.6}$$

and now one can see a fundamental reason why the algorithm associated to (1.4) only achieves an $O(\mathfrak{D}_T \log n)$ competitive ratio, where $\mathfrak{D}_T$ is the combinatorial depth of T: The quantity (1.6) *overmeasures* the metric uncertainty.

Suppose that $\boldsymbol{X}$ is a uniformly random leaf. Then $\sum_{\ell \in \mathcal{L}} \Pr[\mathcal{E}_\ell] \log \frac{1}{\Pr[\mathcal{E}_\ell]} = \log n$, where $n = |\mathcal{L}|$. But, in general, one could have $\sum_{u \in V} \Pr[\mathcal{E}_u] \log \frac{1}{\Pr[\mathcal{E}_u]} \geq \Omega(\mathfrak{D}_T \log n)$. This fact was not lost on the authors of [13], but they bypass the problem by combining mirror descent on stars with a recursive composition method called "unfair gluing."

1.3.2 Multiscale conditional entropy

We employ a regularizer that is a more faithful analog of (1.5):

$$\Phi(x) := \sum_{u \in V \setminus \{\mathbb{r}\}} \frac{w_u}{\eta_u} \left(x_u + \delta_u x_{\mathsf{p}(u)}\right) \log\left(\frac{x_u}{x_{\mathsf{p}(u)}} + \delta_u\right), \tag{1.7}$$

where $\mathsf{p}(u)$ denotes the parent of u.

If one ignores the additional parameters $\{\eta_u \geq 1, \delta_u > 0\}$, this is precisely the negative weighted Shannon entropy written according to the chain rule. Here, we set

$$\theta_u := \frac{|\mathcal{L}_u|}{|\mathcal{L}_{\mathsf{p}(u)}|} \tag{1.8}$$

$$\eta_u := 1 + \log(1/\theta_u) \tag{1.9}$$

$$\delta_u := \theta_u / \eta_u \,. \tag{1.10}$$

The numbers $\{\theta_u\}$ are the conditional probabilites of the uniform distribution on leaves. The $\{\delta_u\}$ values are employed as "noise" added to the entropy calculation. Such noise is a fundamental aspect for competitive analysis, and distinguishes it from the application of mirror descent to regret minimization problems (see, e. g., [12]).[1] The effect of these noise parameters appears ubiquitously in applications of the primal-dual method to competitive analysis (see [16]), and manifests itself as an additive term in the update rules (see equation (1.11) below). Intuitively, it ensures that the conditional probability $\frac{x_u}{x_{\mathsf{p}(u)}}$ is updated fast enough even when it is close to 0.

Finally, the numbers $\{\eta_u : u \in V\}$ are commonly referred to as "learning rates" in the study of online learning. They represent the rate at which information is discounted in the resulting

[1]One finds aspects of this "mixing with the uniform distribution" in the bandits setting as well, but used for variance reduction, a seemingly very different purpose.

algorithm; for MTS, this corresponds to the relative importance of costs arriving now vs. costs that arrived in the past.

1.3.3 The dynamics

We will derive in Section 3 the following *continuous time* evolution of the resulting mirror descent algorithm $(x(t) \in \mathsf{K}_T : t \in [0, \infty))$ for a cost path $c\colon [0, \infty) \to \mathbb{R}_+^{\mathcal{L}}$:

$$\partial_t \left(\frac{x_u(t)}{x_{\mathsf{p}(u)}(t)} \right) = \frac{\eta_u}{w_u} \left(\frac{x_u(t)}{x_{\mathsf{p}(u)}(t)} + \delta_u \right) \left(\beta_{\mathsf{p}(u)}(t) - \sum_{\ell \in \mathcal{L}_u} \frac{x_\ell(t)}{x_u(t)} c_\ell(t) \right) \tag{1.11}$$

Here, $\beta_{\mathsf{p}(u)}(t)$ is a Lagrangian multiplier that ensures conservation of conditional probability:

$$\sum_{v \in \chi(\mathsf{p}(u))} \partial_t \left(\frac{x_v(t)}{x_{\mathsf{p}(u)}(t)} \right) = 0\,.$$

One can see that the evolution is being driven by the expected instantaneous cost incurred conditioned on the current state being in the subtree rooted at u.

One should interpret equation (1.11) only when $x(t)$ lies in the relative interior of K_T. Otherwise, the conditional probabilities are ill-defined. One way to rectify this is to prevent $x(t)$ from hitting the relative boundary of K_T at all. It is possible to adaptively modify the cost functions by a suitably small perturbation so as to guarantee this property and, at the same time, ensure that the total discrepancy between the modified and true service cost is a small additive constant.

Instead, we will follow a different approach, by extending the dynamics to an analogous system of conditional probabilities $\{q_u(t) : u \in V \setminus \{\mathbb{r}\}\}$:

$$\partial_t q_u(t) = \frac{\eta_u}{w_u} \left(q_u(t) + \delta_u \right) \left(\beta_{\mathsf{p}(u)}(t) - \hat{c}_u(t) + \alpha_u(t) \right), \tag{1.12}$$

where $q_u(t) = \frac{x_u(t)}{x_{\mathsf{p}(u)}(t)}$ whenever $x_{\mathsf{p}(u)}(t) > 0$, $\alpha_u(t)$ is a Lagrangian multiplier for the constraint $q_u(t) \geq 0$, and $\hat{c}_u(t)$ is the "derived" cost in the subtree rooted at u:

$$\hat{c}_u(t) := \sum_{\ell \in \mathcal{L}_u} q_{\ell|u}(t) c_\ell(t)$$
$$q_{\ell|u}(t) := \prod_{v \in \gamma_{u,\ell} \setminus \{u\}} q_v(t)\,,$$

where $\gamma_{u,\ell}$ is the unique simple u-ℓ path in T.

Stated this way, the mirror descent algorithm can be envisioned as running a "weighted star" algorithm on the conditional probabilities at every internal node of T, with the derived costs at an internal node u given by the average cost of the current strategy for playing one unit of mass in the subtree rooted at u.

In the next section, we will implement and analyze a discretization of (1.12) using Bregman projections. Since our regularizer Φ and convex body K_T do not satisfy the assumptions underlying the existence and uniqueness theorem of [14], we need to construct a solution to (1.12) and, indeed, taking the discretization parameter in our algorithm to zero, one establishes a solution of bounded variation; see Section 3.3.

The major benefit of the formulations (1.11) and (1.12) is in motivating such an algorithm and prescribing the derived costs. In Section 3, we describe how these dynamics can be predicted from the definition (1.7).

2 The MTS algorithm

We will first establish some generic machinery which, at this point, is not specific to MTS yet. Consider a convex polytope $\mathsf{K}_0 \subseteq \mathbb{R}^n$, define $\mathsf{K} := \mathsf{K}_0 \cap \mathbb{R}^n_+$, and assume that K is compact. Suppose additionally that $\Phi : \mathcal{D} \to \mathbb{R}$ is differentiable and strictly convex in an open neighborhood $\mathcal{D} \supseteq \mathsf{K}$.

Let us write D_Φ for the corresponding Bregman divergence

$$\mathsf{D}_\Phi(y \,\|\, x) := \Phi(y) - \Phi(x) - \langle \nabla\Phi(x), y - x \rangle \,,$$

which is non-negative due to convexity of Φ. Then for $x, y, z \in \mathsf{K}$, we have:

$$\mathsf{D}_\Phi(z \,\|\, y) - \mathsf{D}_\Phi(z \,\|\, x) = -\Phi(y) + \Phi(x) - \langle \nabla\Phi(y), z - y \rangle + \langle \nabla\Phi(x), z - x \rangle. \tag{2.1}$$

For a vector $c \in \mathbb{R}^n$ and $x \in \mathsf{K}$, define the projection

$$\Pi^c_{\mathsf{K}}(x) := \operatorname{argmin} \left\{ \mathsf{D}_\Phi(y \,\|\, x) + \langle c, y \rangle : y \in \mathsf{K} \right\}.$$

Since K is compact and Φ is strictly convex, there is a unique minimizer $y^* \in \mathsf{K}$.

For $x \in \mathsf{K}$, recall the definition of the normal cone at x:

$$\mathsf{N}_{\mathsf{K}}(x) = \{ p \in \mathbb{R}^n : \langle p, y - x \rangle \le 0 \text{ for all } y \in \mathsf{K} \}.$$

Given a representation of K by inequality constraints, $\mathsf{K} = \{x \in \mathbb{R}^n : Ax \le b\}$ for $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^n$, it holds

$$\mathsf{N}_{\mathsf{K}}(x) = \{A^T y : y \ge 0 \text{ and } y^T(Ax - b) = 0\}.$$

The KKT conditions yield

$$\nabla\Phi(y^*) = \nabla\Phi(x) - c - \lambda^* \,, \tag{2.2}$$

where $\lambda^* \in \mathsf{N}_{\mathsf{K}}(y^*)$. Since $\mathsf{N}_{\mathsf{K}}(y^*) = \mathsf{N}_{\mathsf{K}_0}(y^*) + \mathsf{N}_{\mathbb{R}^n_+}(y^*)$, we can can decompose $\lambda^* = \beta - \alpha$ with $\beta \in \mathsf{N}_{\mathsf{K}_0}(y^*)$ and $-\alpha \in \mathsf{N}_{\mathbb{R}^n_+}(y^*)$. In particular, we have $\alpha \ge 0$ and $\alpha_i > 0 \implies y^*_i = 0$ for every $i = 1, \ldots, n$.

Substituting this into equation (2.1) gives

$$\begin{aligned}\mathsf{D}_\Phi(z \,\|\, y^*) - \mathsf{D}_\Phi(z \,\|\, x) &= -\Phi(y^*) + \Phi(x) + \langle \nabla\Phi(x), y^* - x\rangle + \langle c - \alpha + \beta, z - y^*\rangle \\ &\le -\mathsf{D}_\Phi(y^* \,\|\, x) + \langle c - \alpha, z - y^*\rangle,\end{aligned}$$

where the inequality comes from $\langle \beta, z - y^*\rangle \le 0$ since $z \in \mathsf{K}$ and $\beta \in \mathsf{N}_\mathsf{K}(y^*)$. We have proved the following.

Lemma 2.1. *For any $x, z \in \mathsf{K}$, and $c \in \mathbb{R}^n$, let $y^* = \Pi^c_\mathsf{K}(x)$ and λ^* be as in (2.2). Then for any $\alpha \in -\mathsf{N}_{\mathbb{R}^+_n}(y^*)$ such that $\lambda^* + \alpha \in \mathsf{N}_{\mathsf{K}_0}(y^*)$, it holds that*

$$\mathsf{D}_\Phi(z \,\|\, y^*) - \mathsf{D}_\Phi(z \,\|\, x) \le \langle c - \alpha, z - y^*\rangle.$$

2.1 Iterative Bregman projections

We describe now a discretization of the algorithm from the introduction. This discretization will mimic the continuous dynamics if the entries of each individual cost vector are small. We can achieve this by splitting each cost vector into several copies of scaled down versions of itself, as discussed in Section 2.3. In Section 3.3, we will give a formal argument that this indeed yields a discretization of the continuous dynamics from the introduction.

Fix a tree T and recall the definition of K_T from (1.3). Let Q_T denote the collection of vectors $q \in \mathbb{R}_+^{V\setminus\{\mathbb{r}\}}$ such that for all $u \in V \setminus \mathcal{L}$,

$$\sum_{v\in\chi(u)} q_v = 1.$$

For $q \in Q_T$ and $u \in V \setminus \mathcal{L}$, we use $q^{(u)} \in \mathbb{R}_+^{\chi(u)}$ to denote the vector defined by $q_v^{(u)} := q_v$ for $v \in \chi(u)$, and define the corresponding probability simplex $Q_T^{(u)} := \{q^{(u)} : q \in Q_T\}$. We will use $\Delta : Q_T \to \mathsf{K}_T$ for the map which sends $q \in Q_T$ to the (unique) $x = \Delta(q) \in \mathsf{K}_T$ such that

$$x_v = x_u q_v \qquad \forall u \in V \setminus \mathcal{L}, v \in \chi(u).$$

Note that q contains more information than x; the map Δ fails to be invertible whenever there is some $u \in V \setminus \mathcal{L}$ with $x_u = 0$.

Fix $\kappa \ge 1$. On the open domain $\mathcal{D}^{(u)} = (-\min_{v\in\chi(u)} \delta_v, \infty)^{\chi(u)}$, for δ_v as given in equation (1.10), define the strictly convex function $\Phi^{(u)} : \mathcal{D}^{(u)} \to \mathbb{R}$ by

$$\Phi^{(u)}(p) := \frac{1}{\kappa} \sum_{v\in\chi(u)} \frac{w_v}{\eta_v} (p_v + \delta_v) \log (p_v + \delta_v) .$$

Denote the corresponding Bregman divergence on $Q_T^{(u)}$ by

$$\mathsf{D}^{(u)}(p \,\|\, p') = \frac{1}{\kappa} \sum_{v\in\chi(u)} \frac{w_v}{\eta_v} \left[(p_v + \delta_v) \log \frac{p_v + \delta_v}{p'_v + \delta_v} + p'_v - p_v \right].$$

We now define an algorithm that takes a point $q \in Q_T$ and a cost vector $c \in \mathbb{R}_+^{\mathcal{L}}$ and outputs a point $p = \mathcal{A}(q, c) \in Q_T$. Fix $\langle u_1, u_2, \ldots, u_N \rangle$ a topological ordering of $V \setminus \mathcal{L}$ such that every child in T occurs before its parent. We define p inductively as follows. Let $\hat{c}_\ell := c_\ell$ for $\ell \in \mathcal{L}$. For every $j = 1, 2, \ldots, N$:

$$\hat{c}_v^{(u_j)} := \hat{c}_v \qquad \forall v \in \chi(u_j) \tag{2.3}$$

$$p^{(u_j)} := \operatorname{argmin} \left\{ \mathsf{D}^{(u_j)}\left(p \,\|\, q^{(u_j)}\right) + \left\langle p, \hat{c}^{(u_j)} \right\rangle \;\middle|\; p \in Q_T^{(u_j)} \right\} \tag{2.4}$$

$$\hat{c}_{u_j} := \sum_{v \in \chi(u_j)} p_v^{(u_j)} \hat{c}_v \tag{2.5}$$

Let $\alpha^{(u_j)}$ be the vector of Lagrange multipliers corresponding to the nonnegativity constraints in equation (2.4) (recall Lemma 2.1). One should note that in this setting (a probability simplex), the nonnegativity multipliers are unique and thus well-defined.

We denote $\alpha = \alpha^{q,c} \in \mathbb{R}_+^V$ as the vector given by $\alpha_v := \alpha_v^{(\mathsf{p}(v))}$ for $v \neq \mathbb{r}$ and $\alpha_{\mathbb{r}} := 0$. Recall the complementary slackness conditions:

$$\alpha_v > 0 \implies p_v = 0. \tag{2.6}$$

For $v \in \chi(u)$, calculate

$$\left(\nabla \Phi^{(u)}(p)\right)_v = \frac{1}{\kappa} \frac{w_v}{\eta_v} \left(1 + \log(p_v + \delta_v)\right).$$

Then using equation (2.2), we can write the algorithm as follows:

For $j = 1, 2, \ldots, N$:
 For $v \in \chi(u_j)$:
  $p_v^{(u_j)} := (q_v^{(u_j)} + \delta_v) \exp\left(\kappa \frac{\eta_v}{w_v}\left(\beta_{u_j} - (\hat{c}_v - \alpha_v)\right)\right) - \delta_v,$
 $\hat{c}_{u_j} := \sum_{v \in \chi(u_j)} p_v^{(u_j)} \hat{c}_v.$

where $\beta_{u_j} \geq 0$ is the multiplier for the constraint $\sum_{v \in \chi(u_j)} q_v^{(u_j)} \geq 1$. There is no multiplier for the constraint $\sum_{v \in \chi(u_j)} q_v^{(u_j)} \leq 1$ because this constraint will be satisfied automatically and is therefore not needed in (2.4): If it were violated, decreasing some p_v with $p_v > q_v^{(u_j)}$ would yield a strictly better solution to the minimization problem (2.4).

2.2 The global divergence

For $z \in \mathsf{K}_T$ and $q \in Q_T$, define the global divergence function

$$\tilde{\mathsf{D}}(z \,\|\, q) := \frac{1}{\kappa} \sum_{u \notin \mathcal{L}} \sum_{v \in \chi(u)} \frac{w_v}{\eta_v} \left[(z_v + \delta_v z_u) \log\left(\frac{\frac{z_v}{z_u} + \delta_v}{q_v + \delta_v} \right) + z_u q_v - z_v \right],$$

where we use the convention that $0 \log\left(\frac{0}{0}+\delta_v\right) = \lim_{\varepsilon\to 0} \varepsilon \log\left(\frac{0}{\varepsilon}+\delta_v\right) = 0$.

Note that $\tilde{\mathsf{D}}$ is the Bregman divergence associated to the regularizer (1.7) (divided by κ) when $\frac{x_v}{x_u}$ is replaced by q_v. One can write

$$\tilde{\mathsf{D}}(z \,\|\, q) = \sum_{u\notin\mathcal{L}} z_u \mathsf{D}^{(u)}(p \,\|\, q)\,,$$

where $p \in \Delta^{-1}(z)$. In other words, $p \in Q_T$ is any point satisfying $z_v = p_v z_u$ for all $u \notin \mathcal{L}$ and $v \in \chi(u)$,

We will use $\tilde{\mathsf{D}}$ as a potential function to prove inequality (1.1), and z will denote the configuration of some offline algorithm. Note that the state of the online algorithm is encoded by $q \in Q_T$, which contains more information than its configuration $\Delta(q) \in K_T$.

The next lemma shows that when an offline algorithm moves, the change in potential is bounded by $O(1/\kappa)$ times the offline movement cost.

Lemma 2.2. *It holds that for any $q \in Q_T$ and $z, z' \in \mathsf{K}_T$,*

$$\left|\tilde{\mathsf{D}}(z \,\|\, q) - \tilde{\mathsf{D}}(z' \,\|\, q)\right| \le \frac{1}{\kappa}\left(2+\frac{4}{\tau}\right)\|z-z'\|_{\ell_1(w)}\,.$$

Proof. Consider a differentiable map $z\colon [0,1] \to \mathbb{R}^V_{++}$ such that $\sum_{v\in\chi(u)} z_v(t) \le z_u(t)$ for each t and $u \notin \mathcal{L}$. It suffices to show that for each t and every fixed $q \in Q_T$,

$$\kappa\left|\partial_t \tilde{\mathsf{D}}(z(t) \,\|\, q)\right| \le \left(2+\frac{4}{\tau}\right)\|z'(t)\|_{\ell_1(w)}\,.$$

Moreover, it suffices to address the case when there is at most one $u \in V$ with $z'_u(t) \neq 0$.

A direct calculation gives

$$\begin{aligned}\kappa\partial_t \tilde{\mathsf{D}}(z(t) \,\|\, q) &= \frac{w_u}{\eta_u} z'_u(t) \log\left(\frac{z_u(t)/z_{\mathsf{p}(u)}(t)+\delta_u}{q_u+\delta_u}\right) \\ &\quad + \sum_{v\in\chi(u)} \frac{w_v}{\eta_v}\left[\delta_v z'_u(t)\log\left(\frac{z_v(t)/z_u(t)+\delta_v}{q_v+\delta_v}\right) + z'_u(t)\left(q_v - \frac{z_v(t)}{z_u(t)}\right)\right]. \end{aligned} \tag{2.7}$$

Let us now use definitions (1.9) and (1.10) to observe that

$$\frac{1}{\eta_v}\left|\log\frac{p_v+\delta_v}{q_v+\delta_v}\right| \le \frac{1}{\eta_v}\log\frac{1+\delta_v}{\delta_v} \le 2.$$

Using this in equation (2.7) yields

$$\kappa\left|\partial_t \tilde{\mathsf{D}}(z(t) \,\|\, q)\right| \le w_u|z'_u(t)|\left(2+\frac{1}{\tau}\sum_{v\in\chi(u)}\left(2\delta_v + \left|q_v - \frac{z_v(t)}{z_u(t)}\right|\right)\right) \le w_u|z'_u(t)|\left(2+\frac{4}{\tau}\right),$$

where the last inequality uses $\sum_{v\in\chi(u)} \delta_v \le \sum_{v\in\chi(u)} \theta_v \le 1$ and $\sum_{v\in\chi(u)} z_v(t) \le z_u(t)$. □

We will sometimes implicitly restrict vectors $x \in \mathbb{R}^V$ to the subspace spanned by $\{e_\ell : \ell \in \mathcal{L}\}$. In this case, we employ the notation

$$\langle x, y \rangle_{\mathcal{L}} := \sum_{\ell \in \mathcal{L}} x_\ell y_\ell,$$

when either vector lies in $\mathbb{R}^V$ or $\mathbb{R}^{\mathcal{L}}$.

According to the following lemma, the change in potential due to movement of the online algorithm is bounded by the difference in service cost between the offline and online algorithm.

Lemma 2.3. *For any cost vector $c \in \mathbb{R}_+^{\mathcal{L}}$, $z \in \mathsf{K}_T$, and $q \in Q_T$, it holds that if $p = \mathcal{A}(q, c)$, then*

$$\tilde{\mathsf{D}}(z \,\|\, p) - \tilde{\mathsf{D}}(z \,\|\, q) \leq \langle c, z - \Delta(p) \rangle_{\mathcal{L}}\,.$$

Proof. Fix $q \in Q_T$ and $c \in \mathbb{R}_+^{\mathcal{L}}$. Let $\alpha = \alpha^{q,c}$ denote the vector of multipliers defined in Section 2.1. For $u \in V \setminus \mathcal{L}$ with $z_u > 0$, define $z^{(u)} \in Q_T^{(u)}$ by

$$z_v^{(u)} := \frac{z_v}{z_u}.$$

Then Lemma 2.1 gives

$$\mathsf{D}^{(u)}\left(z^{(u)} \,\|\, p^{(u)}\right) - \mathsf{D}^{(u)}\left(z^{(u)} \,\|\, q^{(u)}\right) \leq \left\langle \hat{c}^{(u)} - \alpha^{(u)}, z^{(u)} - p^{(u)} \right\rangle_{\chi(u)},$$

where we use $\langle \cdot, \cdot \rangle_{\chi(u)}$ for the standard inner product on $\mathbb{R}^{\chi(u)}$. Multiplying by z_u and summing yields

$$\begin{aligned}\tilde{\mathsf{D}}(z \,\|\, p) - \tilde{\mathsf{D}}(z \,\|\, q) &\leq \sum_{u \notin \mathcal{L}} z_u \left\langle \hat{c}^{(u)} - \alpha^{(u)}, z^{(u)} - p^{(u)} \right\rangle_{\chi(u)} \\ &= \sum_{u \notin \mathcal{L}} \sum_{v \in \chi(u)} (\hat{c}_v^{(u)} - \alpha_v^{(u)}) z_v - \sum_{u \notin \mathcal{L}} z_u \sum_{v \in \chi(u)} (\hat{c}_v^{(u)} - \alpha_v^{(u)}) p_v\,.\end{aligned}$$

Note that from implication (2.6), the latter expression is

$$\sum_{u \notin \mathcal{L}} z_u \sum_{v \in \chi(u)} \hat{c}_v^{(u)} p_v \overset{(2.5)}{=} \sum_{u \notin \mathcal{L}} z_u \hat{c}_u\,.$$

Noting that $\hat{c}_{\mathbb{r}} = \sum_{\ell \in \mathcal{L}} \Delta(p)_\ell c_\ell$, this gives

$$\tilde{\mathsf{D}}(z \,\|\, p) - \tilde{\mathsf{D}}(z \,\|\, q) \leq \sum_{u \neq \mathbb{r}} (\hat{c}_u - \alpha_u) z_u - \sum_{u \notin \mathcal{L}} z_u \hat{c}_u \leq \langle c, z - \Delta(p) \rangle_{\mathcal{L}}\,. \qquad \square$$

2.3 Algorithm and competitive analysis

Let us now outline the proof of inequality (1.2). First, we perform a standard reduction that allows us to bound only the "positive" movement costs when the algorithm moves from x to y. Its proof is straightforward.

Lemma 2.4. *For $x, y \in \mathsf{K}_T$ it holds that*

$$\|x - y\|_{\ell_1(w)} = 2\,\|(x - y)_+\|_{\ell_1(w)} + [\psi(y) - \psi(x)],$$

where $\psi(x) := \sum_{u \neq \mathbb{r}} w_u x_u$ for $x \in \mathsf{K}_T$.

We now state the key technical lemma which controls the positive movement cost by the service cost. To this end, we employ an auxiliary potential function $\Psi : Q_T \to \mathbb{R}$ defined by

$$\Psi_u(q) := -\Delta(q)_u\, \mathsf{D}^{(u)}\left(\theta^{(u)} \,\|\, q^{(u)}\right)$$
$$\Psi(q) := \sum_{u \notin \mathcal{L}} \Psi_u(q).$$

Intuitively, $\Psi(q)$ is a measure of difference between the online configuration q and the uniform distribution over leaves (whose conditional probabilities are given by θ).

Let us give a brief explanation of the need for Ψ. Our addition of "noise" to the multiscale conditional entropy is to achieve the smoothness property established in Lemma 2.2. But this has the adverse effect of increasing the movement cost of the algorithm, as one can see from the δ_u term in (1.11). This additional movement cannot be easily charged against the service cost in the regime where the noise term is dominant: $\frac{x_u(t)}{x_{\mathsf{p}(u)}(t)} \ll \delta_u$. On the other hand, this additional movement has the effect of further decreasing $\frac{x_u(t)}{x_{\mathsf{p}(u)}(t)}$, which drives the conditional probabilities at $\mathsf{p}(u)$ away from the uniform distribution, decreasing Ψ. A formal statement appears later in Lemma 2.11.

For the next two results, take any $q \in Q_T$ and cost $c \in \mathbb{R}_+^{\mathcal{L}}$, and denote $p = \mathcal{A}(q, c)$, $x = \Delta(q)$, $y = \Delta(p)$.

Lemma 2.5 (Movement analysis)**.** *It holds that*

$$\frac{\tau - 3}{\kappa\tau}\left\|(x - y)_+\right\|_{\ell_1(w)} \leq (2\mathfrak{D}_T + \log n)\langle c, x\rangle_{\mathcal{L}} + [\Psi(q) - \Psi(p)]\,.$$

This lemma will be proved in Section 2.4. Let us first see that it can be used to establish bounds on the competitive ratio. Define $w_{\min} := \min\{w_\ell : \ell \in \mathcal{L}\}$ and

$$\varepsilon_T := \frac{w_{\min}}{2(2\mathfrak{D}_T + \log n)}\frac{\tau - 3}{\tau\kappa}.$$

Theorem 2.6. *For any $z \in \mathsf{K}_T$:*

$$\langle c, y\rangle_{\mathcal{L}} \le \langle c, z\rangle_{\mathcal{L}} + \left[\tilde{\mathsf{D}}(z \,\|\, q) - \tilde{\mathsf{D}}(z \,\|\, p)\right] \tag{2.8}$$

$$\kappa^{-1} \|x - y\|_{\ell_1(w)} \le [\psi(y) - \psi(x)] + \frac{2\tau}{\tau - 3}\left([\Psi(q) - \Psi(p)] + (2\mathfrak{D}_T + \log n)\langle c, x\rangle_{\mathcal{L}}\right) \tag{2.9}$$

Moreover, if $\|c\|_\infty \le \varepsilon_T$, then

$$\kappa^{-1} \|x - y\|_{\ell_1(w)} \le [\psi(y) - \psi(x)] + \frac{4\tau}{\tau - 3}\left([\Psi(q) - \Psi(p)] + (2\mathfrak{D}_T + \log n)\langle c, y\rangle_{\mathcal{L}}\right). \tag{2.10}$$

Proof. Inequality (2.8) follows from Lemma 2.3, and inequality (2.9) follows from Lemma 2.5 and Lemma 2.4. To see that inequality (2.10) follows from inequality (2.9) and Lemma 2.5, use the fact that

$$\langle c, x\rangle_{\mathcal{L}} \le \langle c, y\rangle_{\mathcal{L}} + \frac{\|c\|_\infty}{w_{\min}} \left\|(x - y)_+\right\|_{\ell_1(w)}. \qquad \square$$

In light of Theorem 2.6, we can respond to a cost function $c \in \mathbb{R}_+^{\mathcal{L}}$ by splitting it into M pieces $c_1, c_2, \ldots, c_M$ where $M = \lceil \|c\|_\infty / \varepsilon_T \rceil$. Now define $q_i := \mathcal{A}(q_{i-1}, c/M)$, $q_0 := q$ and $\bar{\mathcal{A}}(q, c) := q_M$.

Theorem 2.7. *Fix $\tau \ge 4$. Consider the algorithm that begins in some configuration $q_0 \in Q_T$. If $c_t \in \mathbb{R}_+^{\mathcal{L}}$ is the cost function that arrives at time t, denote $q_t := \bar{\mathcal{A}}(q_{t-1}, c_t)$. Then the sequence $\langle \Delta(q_0), \Delta(q_1), \ldots\rangle$ is an online algorithm that is $(1, O(1/\kappa), O(\kappa(\mathfrak{D}_T + \log n)))$-finely competitive.*

We prove this momentarily. The following fact is well-known and, in conjunction with the preceding theorem, yields the validity of Theorems 1.1 and 1.3.

Lemma 2.8 (See, e.g., [3, Thm. 2.4]). *If $(\mathcal{L}, d_T)$ is an HST metric, then there is another weighted tree T' with leaf set $\mathcal{L}$ such that*

1. *$(\mathcal{L}, d_{T'})$ is a 7-HST metric.*
2. *$\mathfrak{D}_{T'} \le \log_2 |\mathcal{L}|$*
3. *All the leaves of T' have depth $\mathfrak{D}_{T'}$.*
4. *$d_T(\ell, \ell') \le d_{T'}(\ell, \ell') \le O(d_T(\ell, \ell'))$ for all $\ell, \ell' \in \mathcal{L}$.*

Proof sketch. Replace every weight w_v in T with $\hat{w}_v := 7^{\lceil \log_7 w_v \rceil}$ and iteratively contract every edge $(p(u), u)$ with $\hat{w}_{p(u)} = \hat{w}_u$ and $u \notin \mathcal{L}$. The resulting weighted tree T_1 is a 7-HST by construction.

Now iteratively contract every edge $(p(u), u)$ in T_1 for which $|\mathcal{L}_u^{T_1}| > \frac{1}{2}|\mathcal{L}_{p(u)}^{T_1}|$. The resulting tree T' has depth $\mathfrak{D}_{T'} \le \log_2 |\mathcal{L}|$. Finally, one can achieve property (3) by increasing the depth of every root-leaf path to $\mathfrak{D}_{T'}$ using vertex weights that decrease by a factor of 7 along the path. $\square$

Proof of Theorem 2.7. Consider a sequence $\langle c_t : t \geq 1\rangle$ of cost functions. By splitting the costs into smaller pieces, we may assume that $\|c_t\|_\infty \leq \varepsilon_T$ for all $t \geq 1$.

Let $\{z_t^*\}$ denote some offline algorithm with $z_0^* = \Delta(q_0)$, and let $\{x_t = \Delta(q_t)\}$ denote our online algorithm. Then using $\tilde{\mathsf{D}}(z_0^* \,\|\, x_0) = 0$ along with inequality (2.8) and Lemma 2.2 yields, for any time $t_1 \geq 1$,

$$\begin{aligned}\sum_{t=1}^{t_1}\langle c_t, x_t\rangle_{\mathcal{L}} &\leq \sum_{t=1}^{t_1}\langle c_t, z_t^*\rangle_{\mathcal{L}} - \tilde{\mathsf{D}}(z_{t_1}^* \,\|\, q_{t_1}) + O(1/\kappa)\sum_{t=1}^{t_1}\|z_t^* - z_{t-1}^*\|_{\ell_1(w)} \\ &\leq \sum_{t=1}^{t_1}\langle c_t, z_t^*\rangle_{\mathcal{L}} + O(1/\kappa)\sum_{t=1}^{t_1}\|z_t^* - z_{t-1}^*\|_{\ell_1(w)},\end{aligned}$$

where we have used $\tilde{\mathsf{D}}(z \,\|\, q) \geq 0$ for all $z \in \mathsf{K}_T$ and $q \in Q_T$. This verifies inequality (1.1) with $\alpha_0 = 1$, $\alpha_1 = O(1/\kappa)$, and $\beta = 0$. Moreover, inequality (2.10) gives

$$\frac{1}{\kappa}\sum_{t=1}^{t_1}\|x_t - x_{t-1}\|_{\ell_1(w)} \leq [\psi(x_{t_1}) - \psi(x_0)] + \frac{4\tau}{\tau-3}[\Psi(q_0) - \Psi(q_{t_1})] + \left(2\mathfrak{D}_T + \log n\right)\sum_{t=1}^{t_1}\langle c_t, x_t\rangle_{\mathcal{L}},$$

verifying inequality (1.2) with $\alpha_1 \leq O(\kappa(\mathfrak{D}_T + \log n))$ and $\beta' \leq O(\kappa \max_{v \neq \mathbb{r}} w_v)$ (see Lemma 2.10 below). □

2.4 Movement analysis

It remains to prove Lemma 2.5. Recall that $q \in Q_T, c \in \mathbb{R}_+^{\mathcal{L}}$ and $p = \mathcal{A}(q, c), x = \Delta(q), y = \Delta(p)$.

The KKT conditions (see equation (2.2)) give: For every $v \in \chi(u)$,

$$\frac{1}{\kappa}\frac{w_v}{\eta_v}\log\left(\frac{p_v + \delta_v}{q_v + \delta_v}\right) = \beta_u - \hat{c}_v + \alpha_v\,, \tag{2.11}$$

where $\beta_u \geq 0$ is the multiplier corresponding to the constraint $\sum_{v \in \chi(u)} q_v \geq 1$.

Lemma 2.9. *It holds that $\alpha_v \leq \hat{c}_v$ for all $v \in V \setminus \{\mathbb{r}\}$.*

Proof. Note that $\hat{c}_v \geq 0$ by construction. Thus if $\alpha_v = 0$, we are done. Otherwise, by complementary slackness, it must be that $p_v = 0$, and therefore $\log(\frac{p_v+\delta_v}{q_v+\delta_v}) \leq 0$. Since $\beta_{\mathsf{p}(v)} \geq 0$, equation (2.11) implies that $\alpha_v \leq \hat{c}_v$. □

Define $\sigma_v := \log\left(\frac{p_v+\delta_v}{q_v+\delta_v}\right)$ so that

$$q_v - p_v = (q_v + \delta_v)(1 - \mathrm{e}^{\sigma_v}). \tag{2.12}$$

Recall that for $v \in \chi(u)$, we have $x_v = q_v x_u$ and $y_v = p_v y_u$, thus

$$x_v - y_v = x_u(q_v - p_v) + p_v(x_u - y_u) = (x_v + \delta_v x_u)(1 - e^{\sigma_v}) + p_v(x_u - y_u).$$

In particular,

$$\begin{aligned} w_v\,(x_v - y_v)_+ &\le w_v(x_v + \delta_v x_u)(1 - \mathrm{e}^{\sigma_v})_+ + w_v p_v\,(x_u - y_u)_+ \\ &\le w_v(x_v + \delta_v x_u)(1 - \mathrm{e}^{\sigma_v})_+ + \frac{w_u}{\tau} p_v\,(x_u - y_u)_+ \,. \end{aligned}$$

Using $\sum_{v \in \chi(u)} p_v = 1$ and summing over all vertices yields

$$\sum_{v \neq \mathbb{r}} w_v\,(x_v - y_v)_+ \le \sum_{v \neq \mathbb{r}} w_v(x_v + \delta_v x_{\mathsf{p}(v)})(1 - \mathrm{e}^{\sigma_v})_+ + \frac{1}{\tau} \sum_{v \neq \mathbb{r}} w_v\,(x_v - y_v)_+ \,,$$

hence

$$\begin{aligned} \sum_{v \neq \mathbb{r}} w_v\,(x_v - y_v)_+ &\le \frac{\tau}{\tau - 1} \sum_{v \neq \mathbb{r}} w_v(x_v + \delta_v x_{\mathsf{p}(v)})(1 - \mathrm{e}^{\sigma_v})_+ \\ &\le \frac{\tau}{\tau - 1} \sum_{v \neq \mathbb{r}} w_v(x_v + \delta_v x_{\mathsf{p}(v)})\,(\sigma_v)_- \\ &\le \frac{\kappa\tau}{\tau - 1} \left(\sum_{v \neq \mathbb{r}} \eta_v x_v \hat{c}_v + \sum_{u \notin \mathcal{L}} x_u \sum_{v \in \chi(u)} \theta_v(\hat{c}_v - \alpha_v) \right), \end{aligned} \tag{2.13}$$

where the last line uses Lemma 2.9 and equation (2.11), to bound $w_v(\sigma_v)_- \le \kappa\eta_v\,(\hat{c}_v - \alpha_v)$.

Note that

$$\sum_{v \neq \mathbb{r}} \eta_v x_v \hat{c}_v \le \sum_{\ell \in \mathcal{L}} c_\ell x_\ell \sum_{v \in \gamma_{\mathbb{r},\ell} \setminus \{\mathbb{r}\}} \eta_v \le (\mathfrak{D}_T + \log n)\,\langle c, x \rangle \,, \tag{2.14}$$

since for any $\ell \in \mathcal{L}$, it holds that

$$\sum_{v \in \gamma_{\mathbb{r},\ell} \setminus \{\mathbb{r}\}} \eta_v = \mathfrak{D}_T(\ell) + \sum_{v \in \gamma_{\mathbb{r},\ell} \setminus \{\mathbb{r}\}} \log \frac{|\mathcal{L}_{\mathsf{p}(v)}|}{|\mathcal{L}_v|} = \mathfrak{D}_T(\ell) + \log n,$$

where $\mathfrak{D}_T(\ell)$ is the combinatorial depth of ℓ.

The second sum in (2.13) can be interpreted as the service cost of hybrid configurations of q and θ: While $\sum_{v \in \chi(u)} x_v \hat{c}_v$ is the service cost of x in $\mathcal{L}_u$, the term $x_u \sum_{v \in \chi(u)} \theta_v \hat{c}_v$ is the service cost in $\mathcal{L}_u$ of the modification of x whose conditional probabilities at the children of u are given by $\theta^{(u)}$ rather than $q^{(u)}$. To bound this hybrid service cost, we will employ the auxiliary potential Ψ.

2.4.1 The hybrid cost

We require the following elementary estimate.

Lemma 2.10. *For $u \notin \mathcal{L}$ it holds that*

$$\max\left\{ \mathsf{D}^{(u)}(r \,\|\, p) : r, p \in Q_T^{(u)} \right\} \le \frac{2}{\kappa} \frac{w_u}{\tau}.$$

Proof. Define $\phi_v : (-\delta_v, \infty) \to \mathbb{R}$ by

$$\phi_v(p) := \frac{1}{\eta_v}(p_v + \delta_v)\log(p_v + \delta_v),$$

and let

$$\mathsf{D}_{\phi_v}(q_v \,\|\, p_v) = \frac{1}{\eta_v}\left[(q_v + \delta_v)\log\frac{q_v + \delta_v}{p_v + \delta_v} + (p_v - q_v)\right]$$

denote the corresponding Bregman divergence. Then for $q_v, p_v \geq 0$, it holds that $\mathsf{D}_{\phi_v}(q_v \,\|\, p_v) \geq 0$ since ϕ_v is convex on $\mathbb{R}_+$. Employing the τ-HST property of T, this implies that

$$\mathsf{D}^{(u)}(r \,\|\, p) = \frac{1}{\kappa}\sum_{v \in \chi(u)} w_v \mathsf{D}_{\phi_v}(r_v \,\|\, p_v) \leq \frac{w_u}{\kappa\tau}\sum_{v \in \chi(u)} \mathsf{D}_{\phi_v}(r_v \,\|\, p_v).$$

Define $F : Q_T^{(u)} \times Q_T^{(u)} \to \mathbb{R}_+$ by $F(r, p) := \sum_{v \in \chi(u)} \mathsf{D}_{\phi_v}(r_v \,\|\, p_v)$. The map $r \mapsto F(r, p)$ is convex in general (for any Bregman divergence). The map $p \mapsto F(r, p)$ is convex as well, as this holds for each map $p_v \mapsto \mathsf{D}_{\phi_v}(q_v \,\|\, p_v)$ since $-\log(x)$ is convex on $\mathbb{R}_{++}$. Since the maximum of a convex function on the a polytope is achieved at an extreme point, we have

$$\begin{aligned}\max\left\{F(r, p) : r, p \in Q_T^{(u)}\right\} &\leq \max_{\substack{v, v' \in \chi(u) \\ v \neq v'}}\left[\frac{1}{\eta_v}\left((1 + \delta_v)\log\frac{1 + \delta_v}{\delta_v} - 1\right) + \frac{1}{\eta_{v'}}\left(\delta_{v'}\log\frac{\delta_{v'}}{1 + \delta_{v'}} + 1\right)\right] \\ &\leq 2.\end{aligned}$$

□

The next lemma is crucial: It relates the service cost (with respect to the reduced cost $\hat{c} - \alpha$) of the hybrid configurations to the service cost of the actual configuration and the movement cost.

Lemma 2.11. *For any $u \notin \mathcal{L}$, it holds that*

$$\Psi_u(p) - \Psi_u(q) \leq \frac{2}{\kappa}\frac{w_u}{\tau}(x_u - y_u)_+ + \sum_{v \in \chi(u)} (\hat{c}_v - \alpha_v)\left[x_v - \theta_v x_u\right]. \tag{2.15}$$

Proof. Write

$$\begin{aligned}\Psi_u(p) - \Psi_u(q) &= x_u \mathsf{D}^{(u)}\left(\theta^{(u)} \,\|\, q^{(u)}\right) - y_u \mathsf{D}^{(u)}\left(\theta^{(u)} \,\|\, p^{(u)}\right) \\ &= (x_u - y_u)\mathsf{D}^{(u)}(\theta^{(u)} \,\|\, p^{(u)}) + x_u\left[\mathsf{D}^{(u)}(\theta^{(u)} \,\|\, q^{(u)}) - \mathsf{D}^{(u)}(\theta^{(u)} \,\|\, p^{(u)})\right].\end{aligned}$$

Using Lemma 2.10, the first term is bounded by $\frac{2}{\kappa}\frac{w_u}{\tau}(x_u - y_u)_+$.

Let us now bound the second term. Using $1 + t \le e^t$, we have

$$\begin{aligned}\kappa x_u \left[\mathsf{D}^{(u)}(\theta^{(u)} \,\|\, q^{(u)}) - \mathsf{D}^{(u)}(\theta^{(u)} \,\|\, p^{(u)})\right] &= x_u \sum_{v\in\chi(u)} \frac{w_v}{\eta_v}\left[(\theta_v+\delta_v)\log\frac{p_v+\delta_v}{q_v+\delta_v} + q_v - p_v\right] \\ &\overset{(2.12)}{=} x_u \sum_{v\in\chi(u)} \frac{w_v}{\eta_v}\left[(\theta_v+\delta_v)\sigma_v + (q_v+\delta_v)(1-e^{\sigma_v})\right] \\ &\le x_u \sum_{v\in\chi(u)} \frac{w_v}{\eta_v}\sigma_v(\theta_v - q_v) \\ &= \sum_{v\in\chi(u)} \frac{w_v}{\eta_v}\sigma_v\left[\theta_v x_u - x_v\right].\end{aligned}$$

To finish the proof, observe that from equation (2.11),

$$\sum_{v\in\chi(u)} \frac{w_v}{\eta_v}\sigma_v\left[\theta_v x_u - x_v\right] = \kappa \sum_{v\in\chi(u)} (\beta_u - \hat{c}_v + \alpha_v)\left[\theta_v x_u - x_v\right] = \kappa \sum_{v\in\chi(u)} (\alpha_v - \hat{c}_v)\left[\theta_v x_u - x_v\right],$$

where the last equality uses $\sum_{v\in\chi(u)} x_v = x_u$ and $\sum_{v\in\chi(u)} \theta_v = 1$ (from (1.8)). ◻

Using the lemma gives

$$\begin{aligned}\sum_{u\notin\mathcal{L}} x_u \sum_{v\in\chi(u)} \theta_v(\hat{c}_v - \alpha_v) &\overset{(2.15)}{\le} [\Psi(q) - \Psi(p)] + \frac{2}{\kappa\tau}\left\|(\Delta(q) - \Delta(p))_+\right\|_{\ell_1(w)} + \sum_{v\neq\mathbb{r}} \hat{c}_v x_v \\ &\le [\Psi(q) - \Psi(p)] + \frac{2}{\kappa\tau}\left\|(\Delta(q) - \Delta(p))_+\right\|_{\ell_1(w)} + \mathfrak{D}_T\langle c, x\rangle_{\mathcal{L}}.\end{aligned}$$

Combining this inequality with inequality (2.13) and inequality (2.14) gives

$$\kappa^{-1}\left\|(x-y)_+\right\|_{\ell_1(w)} \le \frac{\tau}{\tau-1}\left[\left(2\mathfrak{D}_T + \log n\right)\langle c, x\rangle_{\mathcal{L}} + (\Psi(q) - \Psi(p)) + \frac{2}{\kappa\tau}\left\|(x-y)_+\right\|_{\ell_1(w)}\right], \tag{2.16}$$

completing the verification of Lemma 2.5.

3 Derivation of the dynamics and derived costs

For the sake of motivating the dynamics (1.11), we review the continuous-time mirror descent framework of [14]. Suppose that $\mathsf{K} \subseteq \mathbb{R}^N$ is a convex set. We recall again the definition of the *normal cone to* K *at* $x \in \mathsf{K}$, which is given by

$$N_{\mathsf{K}}(x) := (\mathsf{K} - x)^\circ = \left\{p \in \mathbb{R}^N : \langle p, y - x\rangle \le 0 \text{ for all } y \in \mathsf{K}\right\}.$$

Suppose additionally that $\Phi : \mathcal{D} \to \mathbb{R}$ is C^2 and strictly convex on an open neighborhood $\mathcal{D} \supseteq \mathsf{K}$ so that the Hessian $\nabla^2\Phi(x)$ is well-defined and positive definite on $\mathcal{D}$. Given a control function $F : [0,\infty) \times \mathsf{K} \to \mathbb{R}^N$ and an initial point $x_0 \in \mathsf{K}$, we will be concerned with absolutely continuous solutions $x : [0,\infty) \to \mathsf{K}$ to the differential inclusion

$$\begin{aligned} x(0) &= x_0, \\ \nabla^2\Phi(x(t))x'(t) &\in F(t, x(t)) - N_{\mathsf{K}}(x(t)). \end{aligned}$$

In other words, a trajectory that satisfies $x(0) = x_0$ and for almost every $t \geq 0$:

$$x'(t) = \nabla^2\Phi(x(t))^{-1} \left(F(t, x(t)) - \gamma(t)\right), \tag{3.1}$$

with $\gamma(t) \in N_{\mathsf{K}}(x(t))$.

Under suitably strong conditions on Φ and F, there is a unique absolutely continuous solution to equation (3.1) [14]. In our setup, these conditions are actually *not satisfied* unless we prevent the path x from hitting the relative boundary of K. Nevertheless, the formal calculation is elucidating and motivates the algorithm of Section 2. For simplicity, we assume $\kappa := 1$ in this section.

3.1 Hessian computation

Let us take Φ as in (1.7) and calculate $\nabla^2\Phi(x)$ for $x \in \mathbb{R}^V_{++}$. Fix $u \neq \mathbb{r}$. Then we have

$$\partial_u\Phi(x) = \frac{w_u}{\eta_u}\left(\log\left(\frac{x_u}{x_{\mathsf{p}(u)}} + \delta_u\right) + 1\right) + \sum_{v\in\chi(u)} \frac{w_v}{\eta_v}\left(\delta_v \log\left(\frac{x_v}{x_u} + \delta_v\right) - \frac{x_v}{x_u}\right). \tag{3.2}$$

Moreover, $\partial_{uv}\Phi(x) = 0$ unless $u = v$, $u \in \chi(v)$, or $v \in \chi(u)$, and in this case,

$$\begin{aligned} \partial_{uu}\Phi(x) &= \frac{w_u}{\eta_u(x_u + \delta_u x_{\mathsf{p}(u)})} + \sum_{v\in\chi(u)} \left(\frac{x_v}{x_u}\right)^2 \frac{w_v}{\eta_v(x_v + \delta_v x_u)} \\ \partial_{u,\mathsf{p}(u)}\Phi(x) = \ \partial_{\mathsf{p}(u),u}\Phi(x) &= -\frac{x_u}{x_{\mathsf{p}(u)}} \frac{w_u}{\eta_u(x_u + \delta_u x_{\mathsf{p}(u)})}. \end{aligned}$$

3.2 Explicit dynamics

We are now in a position to calculate the formal dynamics. Let us define the control by $F(\cdot, t) := -c(t)$. We claim that for $u \neq \mathbb{r}$,

$$\partial_t\left(\frac{x_u(t)}{x_{\mathsf{p}(u)}(t)}\right) = \frac{\eta_u}{w_u}\left(\frac{x_u(t)}{x_{\mathsf{p}(u)}(t)} + \delta_u\right)\left(\beta_{\mathsf{p}(u)}(t) - \sum_{\ell\in\mathcal{L}_u} \frac{x_\ell(t)}{x_u(t)} c_\ell\right), \tag{3.3}$$

where $\beta_u(t) \geq 0$ denotes the Lagrange multiplier corresponding to the constraint $x_u = \sum_{v\in\chi(u)} x_v$.

To verify equation (3.3), let us define, for $u \neq \mathbb{r}$,

$$\mathcal{E}(u) := \frac{w_u}{\eta_u} \frac{x_{\mathsf{p}(u)}(t)}{x_u(t) + \delta_u x_{\mathsf{p}(u)}(t)} \partial_t\left(\frac{x_u(t)}{x_{\mathsf{p}(u)}(t)}\right).$$

Then equation (3.3) is equivalent to the assertion that

$$\mathcal{E}(u) = \beta_{\mathsf{p}(u)}(t) - \sum_{\ell \in \mathcal{L}_u} \frac{x_\ell(t)}{x_u(t)} c_\ell(t). \tag{3.4}$$

Recalling equation (3.1), the equality $\left(\nabla^2 \Phi(x(t)) x'(t)\right)_u = (F(t, x(t)) - \gamma(t))_u$ is equivalent to

$$\mathcal{E}(\ell) = \beta_{\mathsf{p}(\ell)}(t) - c_\ell(t), \qquad \ell \in \mathcal{L}, \tag{3.5}$$

$$\mathcal{E}(u) - \sum_{v \in \chi(u)} \frac{x_v(t)}{x_u(t)} \mathcal{E}(v) = \beta_{\mathsf{p}(u)}(t) - \beta_u(t), \qquad u \in V \setminus (\mathcal{L} \cup \{\mathbb{r}\}). \tag{3.6}$$

Clearly equation (3.5) already confirms equation (3.4) for $\ell \in \mathcal{L}$.

Let us conclude by verifying equation (3.4) for all $u \notin \mathbb{r}$ by (reverse) induction on the depth. Employing equation (3.6) along with the validity of equation (3.4) for $\{\mathcal{E}(v) : v \in \chi(u)\}$ yields

$$\begin{aligned} \mathcal{E}(u) &= \beta_{\mathsf{p}(u)}(t) - \beta_u(t) + \sum_{v \in \chi(u)} \frac{x_v(t)}{x_u(t)} \left(\beta_u(t) - \sum_{\ell \in \mathcal{L}_u} \frac{x_\ell(t)}{x_u(t)} c_\ell(t) \right) \\ &= \beta_{\mathsf{p}(u)}(t) - \sum_{\ell \in \mathcal{L}_u} \frac{x_\ell(t)}{x_u(t)} c_\ell(t), \end{aligned}$$

where we used the fact that $x_u = \sum_{v \in \chi(u)} x_v$ for $x \in \mathsf{K}_T$.

3.3 Relationship between discrete and continuous dynamics

Recall the setup from Section 1.3.3. We consider a system of variables $\{q_u(t) : u \in V \setminus \{\mathbb{r}\}\}$ satisfying the differential equations

$$\partial_t q_u(t) = \frac{\eta_u}{w_u} (q_u(t) + \delta_u) \left(\beta_{\mathsf{p}(u)}(t) - \hat{c}_u(t) + \alpha_u(t) \right), \tag{3.7}$$

where $\alpha_u(t)$ is a Lagrangian multiplier for the constraint $q_u(t) \geq 0$, and $\hat{c}_u(t)$ is the "derived" cost in the subtree rooted at u:

$$\begin{aligned} \hat{c}_u(t) &:= \sum_{\ell \in \mathcal{L}_u} q_{\ell|u}(t) c_\ell(t) \\ q_{\ell|u}(t) &:= \prod_{v \in \gamma_{u,\ell} \setminus \{u\}} q_v(t), \end{aligned}$$

where $\gamma_{u,\ell}$ is the unique simple u-ℓ path in T. Now the values $q_{\ell|\mathbb{r}}$ give a probability distribution on the leaves.

Let us argue that when the discretization parameter of the algorithm presented in Section 2 goes to zero, one arrives at a solution to equation (3.7). Recall that in Section 2.3, we split each cost function $c \in \mathbb{R}_+^{\mathcal{L}}$ into M pieces $M^{-1}c$ and computed a sequence of configurations $q_0, \dots, q_M \in Q_T$. Define the piecewise-linear function $q_{(M)} : [0,1] \to Q_T$ by

$$q_{(M)}\left(\frac{j+\delta}{M}\right) := (1-\delta)q_j + \delta q_{j+1}, \qquad \delta \in [0,1], j \in \{0,\dots,M-1\}.$$

Recalling Section 2.1, we have

$$q_j^{(u)} := \operatorname{argmin}\left\{\mathsf{D}^{(u)}\left(p \,\|\, q_{j-1}^{(u)}\right) + \left\langle p, M^{-1}\hat{c}_j^{(u)}\right\rangle \;\middle|\; p \in Q_T^{(u)}\right\}, \tag{3.8}$$

where

$$\hat{c}_j^{(u)} = \sum_{\ell \in \mathcal{L}_v} (q_j)_{\ell|u} c_\ell.$$

Thus for $v \in \chi(u)$ and $j \geq 1$,

$$\left(q_j^{(u)}\right)_v = \left[\left(q_{j-1}^{(u)}\right)_v + \delta_v\right] \exp\left(\frac{\eta_v}{w_v}\left(\beta_u - (M^{-1}(\hat{c}_j^{(u)})_v - \alpha_v)\right)\right) - \delta_v.$$

One can now verify that there is a constant $L = L(c,T)$ such that

$$\left|(q_j)_v - (q_{j-1})_v\right| \leq \frac{L}{M}, \quad j \in \{1,\dots,M\}, v \in V \setminus \{\mathbb{r}\}.$$

In particular, we see that $q'_{(M)} \in L^\infty([0,1], \mathbb{R}^{V \setminus \{\mathbb{r}\}})$ for every $M \geq 1$ and, moreover,

$$\sup_{M \geq 1} \left\|q'_{(M)}\right\|_{L^\infty} < \infty. \tag{3.9}$$

Therefore by Arzelà–Ascoli (see, e.g., [2, Thm. 0.3.1]), there is a subsequence $\{M_k\}$ such that $q_{(M_k)}$ converges uniformly to a function $q : [0,1] \to Q_T$.

Since the unit ball of $L^\infty([0,1], \mathbb{R}^{V \setminus \{\mathbb{r}\}})$ is weakly compact (by the sequential Banach–Alaoglu Theorem; see, e.g., [2, Thm. 0.3.3]), we can pass to a further subsequence $\{M'_k\}$ along which $q'_{(M'_k)}$ converges weakly to some $h \in L^\infty([0,1], \mathbb{R}^{V \setminus \{\mathbb{r}\}})$. Moreover, since $q_{(M)}(b) - q_{(M)}(a) = \int_a^b q'_{(M)}(t)\,dt$ for all $0 \leq a < b \leq 1$, it follows that $q(b) - q(a) = \int_a^b h(t)\,dt$ as well, and therefore for almost all $t \in [0,1]$, we have $q'(t) = h(t)$.

If we similarly linearly interpolate the cost function to $\hat{c}_{(M)} : [0,1] \to \mathbb{R}_+^{V \setminus \{\mathbb{r}\}}$, then $\hat{c}_{(M_k)} \to \hat{c}$ along this sequence as well, and

$$\hat{c}^{(u)}(t) = \sum_{\ell \in \mathcal{L}_v} q_{\ell|u}(t) c_\ell.$$

Now the KKT conditions for optimality in (3.8) give

$$\nabla\Phi^{(u)}\left(q_j^{(u)}\right) - \nabla\Phi^{(u)}\left(q_{j-1}^{(u)}\right) + M^{-1}\hat{c}_j^{(u)} \in -\mathsf{N}_{Q_T^{(u)}}\left(q_j^{(u)}\right),$$

or equivalently,

$$\frac{\nabla\Phi^{(u)}\left(q_j^{(u)}\right)-\nabla\Phi^{(u)}\left(q_{j-1}^{(u)}\right)}{M^{-1}} \in -\hat{c}_j^{(u)} - \mathsf{N}_{Q_T^{(u)}}\left(q_j^{(u)}\right).$$

By standard results in differential inclusion theory (e. g., the Convergence Theorem [2, Thm. 1.4.1]), we conclude that $q : [0,1] \to Q_T$ solves the differential inclusion

$$\nabla^2\Phi^{(u)}\left(q^{(u)}(t)\right)\partial_t q^{(u)}(t) \in -\hat{c}^{(u)}(t) - \mathsf{N}_{Q_T^{(u)}}(q^{(u)}(t)).$$

Calculating the Hessian $\nabla^2\Phi^{(u)}$ reveals that $q(t)$ is a solution to equation (3.7).

References

[1] JACOB ABERNETHY, PETER BARTLETT, NIV BUCHBINDER, AND ISABELLE STANTON: A regularization approach to metrical task systems. In *Proc. 21st Internat. Conf. Algorithmic Learning Theory (ALT'10)*, pp. 270–284. Springer, 2010. [doi:10.1007/978-3-642-16108-7_23] 5

[2] JEAN-PIERRE AUBIN AND ARRIGO CELLINA: *Differential Inclusions: Set-Valued Maps and Viability Theory*. Springer, 1984. [doi:10.1007/978-3-642-69512-4] 21, 22

[3] NIKHIL BANSAL, NIV BUCHBINDER, ALEKSANDER MADRY, AND JOSEPH NAOR: A polylogarithmic-competitive algorithm for the k-server problem. *J. ACM*, 62(5):40:1–49, 2015. Preliminary version in FOCS'11. [doi:10.1145/2783434, arXiv:1110.1580] 14

[4] NIKHIL BANSAL, NIV BUCHBINDER, AND JOSEPH NAOR: Metrical task systems and the k-server problem on HSTs. In *Proc. 37th Internat. Colloq. on Automata, Languages, and Programming (ICALP'10)*, pp. 287–298. Springer, 2010. [doi:10.1007/978-3-642-14165-2_25] 3

[5] NIKHIL BANSAL, NIV BUCHBINDER, AND JOSEPH NAOR: A primal-dual randomized algorithm for weighted paging. *J. ACM*, 59(4):19:1–24, 2012. Preliminary version in FOCS'07. [doi:10.1145/2339123.2339126] 2

[6] YAIR BARTAL: Probabilistic approximations of metric spaces and its algorithmic applications. In *Proc. 37th FOCS*, pp. 184–193. IEEE Comp. Soc., 1996. [doi:10.1109/SFCS.1996.548477] 2, 3

[7] YAIR BARTAL, AVRIM BLUM, CARL BURCH, AND ANDREW TOMKINS: A polylog(n)-competitive algorithm for metrical task systems. In *Proc. 29th STOC*, pp. 711–719. ACM Press, 1997. [doi:10.1145/258533.258667] 2

[8] YAIR BARTAL, BÉLA BOLLOBÁS, AND MANOR MENDEL: Ramsey-type theorems for metric spaces with applications to online problems. *J. Comput. System Sci.*, 72(5):890–921, 2006. [doi:10.1016/j.jcss.2005.05.008, arXiv:cs/0406028] 2

[9] YAIR BARTAL, NATHAN LINIAL, MANOR MENDEL, AND ASSAF NAOR: On metric Ramsey-type phenomena. *Ann. Math.*, 162(2):643–709, 2005. Preliminary version in STOC'03. [doi:10.4007/annals.2005.162.643, arXiv:math/0406353] 2

[10] AVRIM BLUM, HOWARD KARLOFF, YUVAL RABANI, AND MICHAEL SAKS: A decomposition theorem for task systems and bounds for randomized server problems. *SIAM J. Comput.*, 30(5):1624–1661, 2000. [doi:10.1137/S0097539799351882] 2

[11] ALLAN BORODIN, NATHAN LINIAL, AND MICHAEL E. SAKS: An optimal on-line algorithm for metrical task system. *J. ACM*, 39(4):745–763, 1992. [doi:10.1145/146585.146588] 2

[12] SÉBASTIEN BUBECK AND NICOLÒ CESA-BIANCHI: Regret analysis of stochastic and non-stochastic multi-armed bandit problems. *Found. Trends Mach. Learning*, 5(1):1–122, 2012. [doi:10.1561/2200000024, arXiv:1204.5721] 2, 6

[13] SÉBASTIEN BUBECK, MICHAEL B. COHEN, JAMES R. LEE, AND YIN TAT LEE: Metrical task systems on trees via mirror descent and unfair gluing. *SIAM J. Comput.*, 50(3):909–923, 2021. Preliminary version in SODA'19. [doi:10.1137/19M1237879, arXiv:1807.04404] 2, 3, 4, 5, 6

[14] SÉBASTIEN BUBECK, MICHAEL B. COHEN, YIN TAT LEE, JAMES R. LEE, AND ALEKSANDER MADRY: k-server via multiscale entropic regularization. In *Proc. 50th STOC*, pp. 3–16. ACM Press, 2018. [doi:10.1145/3188745.3188798, arXiv:1711.01085] 5, 8, 18, 19

[15] NIV BUCHBINDER, SHAHAR CHEN, AND JOSEPH (SEFFI) NAOR: Competitive analysis via regularization. In *Proc. 25th Ann. ACM–SIAM Symp. on Discrete Algorithms (SODA'14)*, pp. 436–444. SIAM, 2014. [doi:10.1137/1.9781611973402.32] 5

[16] NIV BUCHBINDER AND JOSEPH NAOR: The design of competitive online algorithms via a primal-dual approach. *Found. Trends Theor. Comp. Sci.*, 3(2–3):93–263, 2009. [doi:10.1561/0400000024] 6

[17] CHRISTIAN COESTER AND JAMES R. LEE: Pure entropic regularization for metrical task systems. In *Proc. 32nd Ann. Conf. on Learning Theory (COLT'19)*, pp. 835–848. MLR Press, 2019. PMLR. 1

[18] JITTAT FAKCHAROENPHOL, SATISH RAO, AND KUNAL TALWAR: A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. System Sci.*, 69(3):485–497, 2004. Preliminary version in STOC'03. [doi:10.1016/j.jcss.2004.04.011] 2, 3, 4

[19] AMOS FIAT AND MANOR MENDEL: Better algorithms for unfair metrical task systems and applications. *SIAM J. Comput.*, 32(6):1403–1422, 2003. Preliminary version in STOC'00. [doi:10.1137/S0097539700376159, arXiv:cs/0406034] 2

[20] STEVE SEIDEN: Unfair problems and randomized algorithms for metrical task systems. *Inform. Comput.*, 148(2):219–240, 1999. [doi:10.1006/inco.1998.2744] 2

AUTHORS

Christian Coester
Associate professor
Department of Computer Science
University of Oxford
Oxford, United Kingdom
christian.coester@cs.ox.ac.uk
https://www.cs.ox.ac.uk/people/christian.coester/

James R. Lee
Professor
Paul G. Allen Center for Computer Science & Engineering
University of Washington
Seattle, Washington, USA
jrl@cs.washington.edu
https://homes.cs.washington.edu/~jrl/

ABOUT THE AUTHORS

CHRISTIAN COESTER is Associate Professor in the Department of Computer Science at the University of Oxford and Tutorial Fellow at St Anne's College. He received his Ph. D. from Oxford in 2020, under the supervision of Elias Koutsoupias. After stints as a postdoctoral researcher at CWI in Amsterdam and Tel Aviv University and as a lecturer at the University of Sheffield, he returned to Oxford in 2022. His research focuses on the design and analysis of algorithms, especially online algorithms and learning-augmented algorithms. Outside of his research, he enjoys sports, chess and playing the piano.

JAMES R. LEE is a Professor of Computer Science & Engineering at the University of Washington. He received his Ph. D. from the University of California, Berkeley in 2005, under the supervision of Christos Papadimitriou, followed by a postdoc at the Institute for Advanced Study in Princeton. His research interests are varied and eclectic, ranging from spectral graph algorithms to functional analysis, and from convex optimization to statistical physics. He challenged a class of undergrads to compete against the MTS algorithm in this paper. They fought (and coded) valiantly. They PyTorched and TensorFlowed. But in the end, the Theory won.