2018 IEEE 59th Annual Symposium on Foundations of Computer Science

Fusible HSTs and the randomized k-server conjecture

[extended abstract]

James R. Lee
Computer Science
University of Washington
Seattle, USA
Email: jrl@cs.washington.edu

***Abstract*—We exhibit a poly($\log k$)-competitive randomized algorithm for the k-server problem on *any* metric space. The best previous result independent of the geometry of the underlying metric space is the $2k-1$ competitive ratio established for the *deterministic* work function algorithm by Koutsoupias and Papadimitriou (1995). Even for the special case when the underlying metric space is the real line, the best known competitive ratio was k. Since deterministic algorithms can do no better than k on any metric space with at least $k+1$ points, this establishes that for every metric space on which the problem is non-trivial, randomized algorithms give an exponential improvement over deterministic algorithms.**

Our algorithm maintains an approximation of the underlying metric space by a distribution over HSTs. The granularity and accuracy of the approximation is adjusted dynamically according to the aggregate behavior of the HST algorithms. In short: We try to obtain more accurate approximations at the locations and scales where the "action" is happening. Thus a crucial component of our approach is the $O((\log k)^2)$-competitive randomized algorithm for HSTs obtained in our previous work with Bubeck, Cohen, Lee, and Mądry, and its "multiscale information theory" perspective.

***Keywords*-online algorithms; competitive analysis; the k-server problem**

I. Introduction

An *online algorithm* is one that receives a sequence of inputs $\langle x_1, x_2, \ldots\rangle$ at discrete times $t \in \{1, 2, \ldots\}$. At every time step t, the algorithm takes some feasible action based only on the inputs $\langle x_1, x_2, \ldots, x_t\rangle$ it has seen so far. There is a cost associated with every feasible action, and the objective of an algorithm is to minimize the average cost per time step. This performance can be compared to the optimal *offline* algorithm which is allowed to decide on a sequence of feasible actions *given the entire input sequence* in advance.

Roughly speaking, an online algorithm is C-*competitive* if, on any valid input sequence, its average cost per time step is at most a factor C more than that of the optimal offline algorithm for the same sequence. The best achievable factor C is referred to as the *competitive ratio* of the underlying problem. It bounds the detrimental effects of uncertainty on optimization. Algorithms designed in the online model tend to trade off the benefits of acting locally to minimize cost while hedging against uncertainty in the future. We refer to the book [BE98].

Perhaps the most well-studied problem in this area is the *k-server problem* proposed by Manasse, McGeoch, and Sleator [MMS90] as a significant generalization of various other online problems. The authors of [BBN10] refer to it as the "holy grail" of online algorithms.

Fix an integer $k \geqslant 1$ and let (X, d_X) denote an arbitrary metric space. We will assume that all metric spaces occurring in the paper have at least two points. The input is a sequence $\langle \sigma_t \in X : t \geqslant 0\rangle$ of *requests.* At every time t, an online algorithm maintains a state $\rho_t \in X^k$ which can be thought of as the location of k *servers* in the space X. At time t, the algorithm is required to have a server at the requested site $\sigma_t \in X$. I.e., a feasible state ρ_t is one that *services* σ_t*:*

$$\sigma_t \in \left\{(\rho_t)_1, \ldots, (\rho_t)_k\right\}.$$

Formally, an *online algorithm* is a sequence of mappings $\boldsymbol{\rho} = \langle \rho_1, \rho_2, \ldots, \rangle$ where, for every $t \geqslant 1$, $\rho_t : X^t \to X^k$ maps a request sequence $\langle \sigma_1, \ldots, \sigma_t\rangle$ to a k-server state that services σ_t. In general, $\rho_0 \in X^k$ will denote some initial state of the algorithm.

The *cost of the algorithm* $\boldsymbol{\rho}$ *in servicing* $\boldsymbol{\sigma} = \langle \sigma_t : t \geqslant 1\rangle$ is the quantity $\mathrm{cost}_{\boldsymbol{\rho}}(\boldsymbol{\sigma}; k, \rho_0)$ defined as

$$\sum_{t \geqslant 1} d_{X^k}(\rho_t(\sigma_1, \ldots, \sigma_t), \rho_{t-1}(\sigma_1, \ldots, \sigma_{t-1})), \tag{I.1}$$

where $d_{X^k}\left((x_1, \ldots, x_k), (y_1, \ldots, y_k)\right) := \sum_{i=1}^k d_X(x_i, y_i)$.

For a given request sequence $\boldsymbol{\sigma} = \langle \sigma_t : t \geqslant 1\rangle$ and initial configuration ρ_0, denote the cost of the *offline optimum* by

$$\mathrm{cost}^*(\boldsymbol{\sigma}; k, \rho_0) := \inf_{\langle \rho_1, \rho_2, \ldots\rangle} \sum_{t \geqslant 1} d_{X^k}\left(\rho_t, \rho_{t-1}\right),$$

where the infimum is over all sequences $\langle \rho_1, \rho_2, \ldots\rangle$ such that ρ_t services σ_t for each $t \geqslant 1$.

2575-8454/18/$31.00 ©2018 IEEE
DOI 10.1109/FOCS.2018.00049

An online algorithm ρ is said to be *C-competitive* if, for every initial configuration $\rho_0 \in X^k$, there is a number $c_0 = c_0(\rho_0) \geqslant 0$ such that

$$\text{cost}_\rho(\sigma; k, \rho_0) \leqslant C \cdot \text{cost}^*(\sigma; k, \rho_0) + c_0$$

for all request sequences σ. A *randomized online algorithm* ρ is a random online algorithm that is feasible with probability one. Such an algorithm is said to be *C-competitive* if for every $\rho_0 \in X^k$, there is a number $c_0 = c_0(\rho_0) > 0$ such that for all σ:

$$\mathbb{E}\left[\text{cost}_\rho(\sigma; k, \rho_0)\right] \leqslant C \cdot \text{cost}^*(\sigma; k, \rho_0) + c_0 .$$

The initial configuration ρ_0 will play a minor role in our arguments, and we will usually leave it implicit, using instead the notations $\text{cost}_\rho(\sigma; k)$ and $\text{cost}^*(\sigma; k)$. Let $\mathsf{D}_k(X, d_X)$ denote the infimum of competitive ratios achievable by deterministic online algorithms, and let $\mathsf{R}_k(X, d_X)$ denote the infimum over randomized online algorithms. When the metric d_X on X is clear from context, we will often omit it from our notation.

One should note that in defining (I.1), we sum over all times $t \geqslant 1$. This is simply to avoid the notational clutter caused by an upper time horizon. One can replace a finite sequence $\langle \sigma_1, \sigma_2, \ldots, \sigma_t \rangle$ of requests by the infinite sequence $\langle \sigma_1, \sigma_2, \ldots, \sigma_t, \sigma_t, \sigma_t, \ldots \rangle$, where the final request is repeated indefinitely.

The authors of [MMS90] showed that if (X, d_X) is an arbitrary metric space and $|X| > k$, then $\mathsf{D}_k(X) \geqslant k$. They conjectured that this it tight.

Conjecture I.1 (k-server conjecture, [MMS90])**.** *For every metric space X with $|X| > k \geqslant 1$, it holds that*

$$\mathsf{D}_k(X) = k .$$

Fiat, Rabani, and Ravid [FRR94] were the first to show that $\mathsf{D}_k(X) < \infty$ for every metric space; they gave the explicit bound $\mathsf{D}_k(X) \leqslant k^{O(k)}$. While Conjecture I.1 is still open, it is now known to be true within a factor of 2.

Theorem I.2 (Koutsoupias-Papadimitriou, [KP95])**.** *For every metric space X and $k \geqslant 1$, it holds that*

$$\mathsf{D}_k(X) \leqslant 2k - 1 .$$

Paging and randomization. Let $\mathcal{U}_n$ denote the metric space on $\{1, 2, \ldots, n\}$ equipped with the uniform metric $d(i, j) = \mathbb{1}_{\{i \neq j\}}$. The special case of the k-server problem when $X = \mathcal{U}_n$ is called *k-paging*. Note that an adversarial request sequence for a deterministic online algorithm can be constructed by basing future requests on the current state of the algorithm.

Consider, for instance, the following lower bound for $\mathcal{U}_{k+1} \subseteq \mathcal{U}_n$ (for $n > k$). For any deterministic algorithm $\mathcal{A}$, define the request sequence that at time $t \geqslant 1$ makes a request at the unique site in $\mathcal{U}_{k+1}$ at which $\mathcal{A}$ does *not* have a server.

Clearly $\mathcal{A}$ incurs movement cost exactly t up to time t. On the other hand, the algorithm that starts with its servers at k uniformly random points in $\{1, 2, \ldots, k+1\}$ and moves a uniformly random server to service the request (whenever there is not already a server there) has expected movement cost t/k. Thus there is some (deterministic) offline algorithm with cost t/k up to time t. Moreover, manifestly there is also a randomized *online* algorithm that achieves cost $1/k$ per time step in expectation.

And indeed, in the setting of k-paging, it was show that allowing an online algorithm to make random choices helps dramatically in general.

Theorem I.3 ([FKL+91], [MS91])**.** *For every $n > k \geqslant 1$:*

$$\mathsf{R}_k(\mathcal{U}_n) = 1 + \frac{1}{2} + \cdots + \frac{1}{k} .$$

Work of Karloff, Rabani, and Ravid [KRR94] exploited a "metric Ramsey dichotomy" to give a lower bound on the randomized competitive ratio for *any sufficiently large metric space.* The works [BBM06], [BLMN05] made substantial advances along this front, obtaining the following.

Theorem I.4. *For any metric space X and $k \geqslant 2$ such that $|X| > k$, it holds that*

$$\mathsf{R}_k(X) \geqslant \Omega\left(\frac{\log k}{\log \log k}\right) .$$

In light of a lack of further examples, a folklore conjecture arose (see, for instance, [Kou09, Conj. 2]).

Conjecture I.5 (Randomized k-server conjecture)**.** *For every metric space X and $k \geqslant 2$:*

$$\mathsf{R}_k(X) \leqslant O(\log k) .$$

The possibility that $\mathsf{R}_k(X) \leqslant (\log k)^{O(1)}$ is stated explicitly many times in the literature; see, e.g., [BBK99] and [BE98, Ques. 11.1]. Our main theorem asserts that, indeed, randomization helps dramatically for every metric space.

Theorem I.6 (Main theorem)**.** *For every metric space X and $k \geqslant 2$:*

$$\mathsf{R}_k(X) \leqslant (\log k)^{O(1)}$$

Even when $X = \mathbb{R}$, the best previous upper bound was inherited from the deterministic setting [CKPV91]: $\mathsf{R}_k(\mathbb{R}) \leqslant \mathsf{D}_k(\mathbb{R}) = k$.

Theorem I.6 owes much to three recent works that each dramatically improve our understanding of the k-server problem. The first is the successful resolution of the randomized k-server conjecture for an important special case called *weighted paging*. Consider a set X and a non-negative weight $w : X \to \mathbb{R}_+$. Define the distance $d_w(x, y) := \max\{w(x), w(y)\}$. We refer to this as a *weighted star metric.*

Theorem I.7 (Bansal-Buchbinder-Naor, [BBN12]). *If X is a weighted star metric and $k \geqslant 2$, then*

$$\mathsf{R}_k(X) \leqslant O(\log k).$$

The second recent breakthrough [BBMN15] shows that when X is finite, the competitive ratio can be bounded by polylogarithmic factors in $|X|$.

Theorem I.8 (Bansal-Buchbinder-Mądry-Naor). *For every $k \geqslant 2$ and finite metric space X, it holds that*

$$\mathsf{R}_k(X) \leqslant \left(\log |X|\right)^{O(1)}$$

Finally, in joint work with Bubeck, Cohen, Lee, and Mądry [BCL+18], we obtain a cardinality-independent bound when X is an ultrametric.

Theorem I.9 ([BCL+18]). *For every $k \geqslant 2$ and every ultrametric space X, it holds that*

$$\mathsf{R}_k(X) \leqslant O\left((\log k)^2\right).$$

A. Approximation by HSTs

The significance of ultrametrics in Theorem I.9 stems from their pivotal role in online algorithms for k-server. Consider a rooted tree $T = (V, E)$ equipped with positive vertex weights $\{w_u > 0 : u \in V\}$ such that the weights are non-increasing along every root-leaf path. Let $\mathcal{L} \subseteq V$ denote the set of leaves of T, and define an ultrametric on $\mathcal{L}$ by

$$d_w(\ell, \ell') := w_{\mathrm{lca}(\ell,\ell')},$$

where $\mathrm{lca}(u, v)$ denotes the least common ancestor of $u, v \in V$ in T.

If it holds for some $\tau \geqslant 1$ that $w_v \leqslant w_u/\tau$ whenever v is a child of u, then (T, w) is called a *τ-hierarchically separated tree* (τ-HST) and $(\mathcal{L}, d_w)$ is referred to as a *τ-HST metric space.* (For finite metric spaces, the notion of an ultrametric and a 1-HST are equivalent.)

This notion was introduced in a seminal work of Bartal [Bar96], [Bar98] along with the powerful tool of probabilistic embeddings into random HSTs. Moreover, he showed that every n-point metric space embeds into a distribution over random HSTs with $O(\log n \log\log n)$ distortion. Using the optimal $O(\log n)$ distortion bound from [FRT04] yields the following consequence.

Theorem I.10. *For every $k \geqslant 2$:*

$$\mathsf{R}_k(X, d) \leqslant O(\log |X|) \cdot \sup_{(L,d')} \mathsf{R}_k(L, d'),$$

where the supremum is over all ultrametrics (L, d') with $|L| = |X|$.

Clearly in conjunction with Theorem I.9, this yields $\mathsf{R}_k(X) \leqslant O\left((\log k)^2 \log |X|\right)$ for any finite metric space X. The reduction from general finite metric spaces to ultrametrics implicit in Theorem I.10 is oblivious to the request sequence; one chooses a single random embedding from X into an HST metric $(\mathcal{L}, d_w)$, and then simulates an online algorithm for the request sequence mapped into $(\mathcal{L}, d_w)$. This is both useful and problematic, as no such approach can yield a bound that does not depend on the cardinality of X; there are many families of metric spaces for which the $O(\log |X|)$ distortion bound is tight.

In [BCL+18], we showed how a *dynamic* embedding of a metric space into ultametrics could overcome the distortion barrier.

Theorem I.11 ([BCL+18]). *For every $k \geqslant 2$ and every finite metric space (X, d):*

$$\mathsf{R}_k(X, d) \leqslant O\left((\log k)^3 \log(1 + \mathcal{A}_X)\right), \tag{I.2}$$

where

$$\mathcal{A}_X := \frac{\max_{x,y \in X} d(x, y)}{\min_{x \neq y \in X} d(x, y)}.$$

The dependence of the competitive ratio on $\mathcal{A}_X$ is still problematic, but one should note that the resulting bound could not be achieved with an oblivious embedding (this can be observed for a family of bounded-degree expander graphs).

B. Experts over HSTs

A natural approach is to construct an online algorithm that maintains, at every time step, a distribution $\mathcal{D}_t$ over embeddings into an HST metric $(\mathcal{L}, d_w)$ and for each embedding $F : X \to \mathcal{L}$, a k-server configuration ρ_t^F corresponding to an online algorithm for the request sequence mapped into $\mathcal{L}$ via F.

Define the *annealed server measure* $\bar{\mu}_t$ to be the measure on X that results from averaging the configurations $F^{-1}(\rho_t^F)$ over $\mathcal{D}_t$. Now one would like to update $\mathcal{D}_t \mapsto \mathcal{D}_{t+1}$ based on the measure $\bar{\mu}_t$. Ideally, the measure $\bar{\mu}_t$ would indicate which pieces of the space X are important to approximate well, allowing an embedding sampled randomly from $\mathcal{D}_t$ to bypass the distortion lower bounds.

Problematically, even if we are allowed to see the entire request sequence in advance, there is no random embedding $F : X \to \mathcal{L}$ that can avoid distorting distances in expectation by less than $\Omega(\log \mathcal{A}_X)$, even when $X \subseteq \mathbb{R}$. In the language of online learning, there is no good "expert."

At a high level, our solution to this problem is to enlarge the class of experts: We maintain instead a distribution $\hat{\mathcal{D}}_t$ on pairs (ρ, F), where ρ is a (fractional)

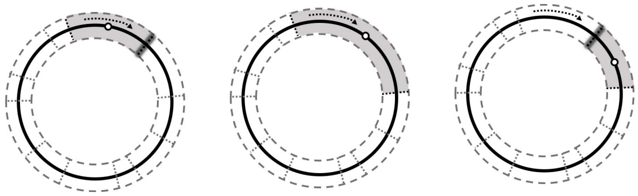

Figure 1. Fusion and fission of two clusters as a server approaches the boundary and then departs.

k-server configuration and $F : X \to \mathcal{L}$ is an embedding. Now let $\bar{\mu}_t$ denote $F^{-1}(\rho)$ averaged over $\hat{\mathcal{D}}_t$.

The distribution $\hat{\mathcal{D}}_{t+1}$ is then sampled by a two-step process: $(\rho, F) \mapsto (\hat{\rho}, F) \mapsto (\hat{\rho}, \hat{F})$. The first step corresponds to updating the k-server configuration to service the request σ_t that arrives at time t. The second step is new: We alter the embedding F so that it more accurately approximates X according to the annealed server measure $\bar{\mu}_t$. An essential feature of this approach is that $\bar{\mu}_t$ is a function only of the distribution $\hat{\mathcal{D}}_t$ which is a function of the request sequence: We are running many HST algorithms in parallel and using the *aggregate behavior* of the algorithms to reshape the underlying HSTs.

Our main primitive for altering the geometry of the HSTs is by *fusion* and *fission* of the metric partitions that underly the HST structure. As a very elementary example, see Figure 1 where a single set is fused to allow the server to continue moving without encountering a partition boundary. We will perform such operations at every location and scale of the metric space according to the annealed server measure.

Executing this strategy is difficult because of a fundamental tension: While, on the one hand, the geometry of the HSTs must be altered to maintain a strong approximation of the metric space, these alterations should not (i) induce too much additional movement cost, and (ii) should not destroy too much of the information that the HST algorithm has learned about the optimal solution.

For the sake of this extended abstract, we will only prove Theorem I.6 under the following assumptions:

1) The metric space (X, d_X) is finite (though our competitive ratio will not depend on $|X|$)
2) The metric space (X, d_X) is *decomposable* in the sense described in Section III-B. This holds, for instance, when $X \subseteq \mathbb{R}^d$ for any $d \geqslant 1$. This allows us to avoid handling insertion and deletion operations; conceptually, at least, such operations are handled by the argument underlying Theorem I.11
3) We will only provide a *fractional* k-server algorithm; see Section I-C. Such an algorithm needs to be rounded online to a random integral algorithm. In [BBMN15], it is shown how to accomplish this while blowing up the expected movement cost by at most an $O(1)$ factor.
4) Our fractional algorithm will only $(1-\frac{1}{2k})$-*serve* the requests, i.e., we will only place $1-\frac{1}{2k}$ measure of fractional servers at the requested site. In [BCL+18], we show how such an algorithm can be converted online to a fractional k-server algorithm that fully serves the requests, increasing the movement cost by a factor of at most 2.

The reader should note that the last two items above have to be adapted to our setting (which allows the algorithm to dynamically alter the geometry of the underlying HSTs).

In this setting, we exhibit here an algorithm with competitive ratio $O((\log k)^8 \log\log k)$. For the general case (which requires dispensing with assumption (2) above), this yields a competitive ratio of $O((\log k)^9 \log\log k)$.

C. *Preliminaries*

Let us write $\mathbb{R}_+ := [0, \infty)$ and $\mathbb{Z}_+ := \mathbb{Z} \cap \mathbb{R}_+$. Consider a finite metric space (X, d_X) with diameter at most one. We use $\mathbb{M}(X)$ to denote the space of measures on X. Denote by $\mathbb{M}_k(X) \subseteq \mathbb{M}(X)$ the subset of measures $\mu \in \mathbb{M}(X)$ that satisfy $\mu(X) = k$. When $x \in X$, we will often write $\mu(x)$ for $\mu(\{x\})$. Denote by $\widehat{\mathbb{M}}(X)$ the set of *integral* measures on X, i.e., those $\mu \in \mathbb{M}(X)$ which take values in $\mathbb{Z}_+$, and similarly $\widehat{\mathbb{M}}_k(X) := \mathbb{M}_k(X) \cap \widehat{\mathbb{M}}(X)$.

If $\boldsymbol{\mu}$ is a sequence of measures in $\mathbb{M}_k(X)$, we write

$$\mathrm{cost}_X(\boldsymbol{\mu}) := \sum_{t \geqslant 1} W^1_X(\mu_t, \mu_{t+1}),$$

where $W^1_X(\mu, \nu)$ is the L^1-transportation distance between μ and ν in X (aka, the Earthmover distance).

For $x \in X$ and $r \geqslant 0$, we denote the ball $B_X(x, r) := \{y \in X : d_X(x, y) \leqslant r\}$ and for $S \subseteq X$, the neighborhood $B_X(S, r) := \bigcup_{x \in S} B_X(x, r)$. For two subsets $S, T \subseteq X$, we write $d_X(S, T) := \inf\{d_X(x, y) : x \in S, y \in T\}$.

Fractional k-server algorithms. Consider a global filtration $\mathbb{F} = \langle \mathbb{F}_1, \mathbb{F}_2, \ldots \rangle$ where $\mathbb{F}_1 \subseteq \mathbb{F}_2 \subseteq \cdots$, and $\mathbb{F}_t$ represents information about the request sequence up to time t. Denote the request sequence $\boldsymbol{\sigma} = \langle \sigma_1, \sigma_2, \ldots \rangle$ with $\sigma_t \in X$ for all $t \geqslant 1$. We use $\boldsymbol{\sigma}_{[s,t]}$ to denote the subsequence $\langle \sigma_s, \sigma_{s+1}, \ldots, \sigma_t \rangle$. Say that a sequence $\boldsymbol{\rho} = \langle \rho_0, \rho_1, \rho_2, \ldots \rangle$ is $\mathbb{F}$-*adapted* if each object ρ_t is possibly a function of $\boldsymbol{\sigma}_{[1,t]}$ (but *not* the future $\sigma_{t+1}, \sigma_{t+2}, \ldots$).

An *offline fractional k-server algorithm* (for $\boldsymbol{\sigma}$) is a sequence of measures $\boldsymbol{\mu} = \langle \mu_0, \mu_1, \mu_2, \ldots \rangle$ such that $\mu_t \in \mathbb{M}_k(X)$ for all $t \geqslant 1$, and such that $\mu_t(\sigma_t) \geqslant 1$ holds for every $t \geqslant 1$. We say that $\boldsymbol{\mu}$ is *integral* if each measure μ_t takes values in $\mathbb{Z}_+$. An *online fractional k-server algorithm* is such a sequence $\boldsymbol{\mu}$ that is additionally $\mathbb{F}$-adapted. We will use the term *fractional k-server algorithm* to mean an online algorithm and explicitly use "offline" for the former notion.

II. Fractional k-server and the allocation polytope

Let $T = (V, E)$ be a finite, rooted tree with positive vertex weights $\{w_v : v \in V\}$ and root $r \in V$. We will further assume throughout that $w_v = w_{v'}$ whenever $v, v' \in V$ have the same (combinatorial) depth in T. We equip V with the metric $\mathrm{dist}_{T,w}$ defined as the weighted-path distance where an edge (u, v) (directed away from the root) has weight w_v.

Let $\mathcal{L} \subseteq V$ denote the set of leaves, and for $u \in V$, let $\mathcal{L}_u$ denote the set of leaves in the subtree rooted at u. We write $\mathsf{ch}(u)$ for the set of children of u in T, $\mathsf{p}(u)$ for the parent of u whenever $u \neq r$, and $\mathsf{des}(u)$ for all the vertices in the subtree rooted at u.

Consider an element

$$z = \langle z_{v,i} \in [0,1] : v \in V, i = 1, 2, \ldots \rangle \in \ell^\infty(V \times \mathbb{Z}_+)$$

For $v \in V$, we denote $z_v := \sum_{i \geqslant 1} z_{v,i}$, and write

$$\chi(v) := \mathsf{ch}(v) \times \mathbb{Z}_+ \,.$$

Let A denote the closed convex set of all z such that z has only finitely many non-zero coordinates and satisfies the family of constraints:

$$\begin{aligned}
z_{r,i} &= \begin{cases} 1 & i \in \{1, 2, \ldots, k\} \\ 0 & i > k \end{cases} \\
z_u &= \sum_{v \in \mathsf{ch}(u)} z_v && \forall u \in V \\
\sum_{i \leqslant |S|} z_{u,i} &\geqslant \sum_{(v,j) \in S} z_{v,j} && \forall S \subseteq \chi(u) \\
z_{u,i} &\geqslant z_{u,i+1} && \forall u \in V, i \geqslant 1
\end{aligned}$$

To every such $z \in \mathsf{A}$, we can associate a unique measure $\nu^z \in \mathbb{M}_k(\mathcal{L})$ such that $\nu^z(S) = \sum_{\ell \in S} z_\ell$ for every subset $S \subseteq \mathcal{L}$.

If $\mu \in \widehat{\mathbb{M}}_k(\mathcal{L})$ is an integral k-server measure, we associate to it the point $\hat{z}^\mu \in \mathsf{A}$ defined as follows:

$$\hat{z}^\mu_{v,i} = \begin{cases} 1 & \mu(\mathcal{L}_v) \geqslant i \\ 0 & \text{otherwise.} \end{cases}$$

We denote $\hat{\mathsf{A}} := \left\{ \hat{z}^\mu : \mu \in \widehat{\mathbb{M}}_k(\mathcal{L}) \right\} \subseteq \mathsf{A}$.

A. The BCLLM $(1-\eta)$-serving algorithm

Let $\eta := \frac{1}{2k}$. Define the mapping $\mathsf{D}_\eta : \hat{\mathsf{A}} \times \mathsf{A} \to \mathbb{R}_+$ so that

$$\mathsf{D}_\eta(\hat{z}; z) := \sum_{u \in V} w_u \sum_{i \geqslant 1} (1 - \hat{z}_{u,i} + \eta) \log \frac{1 - \hat{z}_{u,i} + \eta}{1 - z_{u,i} + \eta}$$

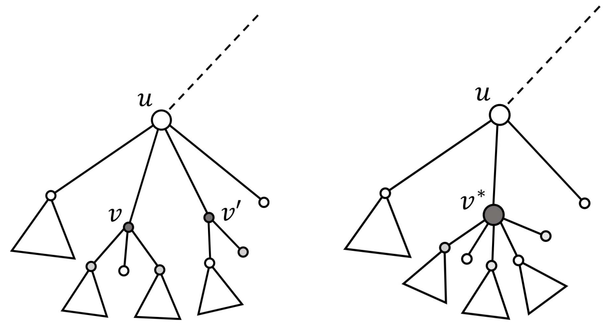

Figure 2. Fusion of T along $\{v, v'\} \implies T_{vv'}$

We also define $\mathsf{D} : \hat{\mathsf{A}} \times \mathsf{A} \to \mathbb{R}_+$ and $\Psi : \mathsf{A} \to \mathbb{R}_+$ by

$$\begin{aligned}
\Psi(z) &:= \sum_{u \in V \setminus \{r\}} \left(w_u + w_{\mathsf{p}(u)}\right)(z_u + 1/4) \log \frac{k + 1/4}{z_u + 1/4} \\
\mathsf{D}(\hat{z}; z) &:= \alpha_0 \mathsf{D}_\eta(\hat{z}; z) + \alpha_1 \Psi(z) \,,
\end{aligned}$$

where $\alpha_1 = \Theta(1)$ and $\alpha_0 = \Theta(\log k)$ is a number depending only on k.

The [BCL+18] algorithm can be interpreted as a mapping $\Gamma : \mathsf{A} \times \mathcal{L} \to \mathsf{A}$ such that for all $\ell \in \mathcal{L}$, if $z' = \Gamma(z, \ell)$, then the following holds.

Lemma II.1. *If (T, w) is a 4-HST, then:*

(T1) $z'_{\ell,1} = 1 - \eta$
(T2) For any $\hat{z}$ with $\hat{z}_{\ell,1} = 1$, it holds that

$$\mathsf{D}(\hat{z}; z') - \mathsf{D}(\hat{z}; z) \leqslant -W^1_T(z, z') \,.$$

The next lemma is straightforward from the definition.

Lemma II.2. *For any $\hat{z}, \hat{z}' \in \hat{\mathsf{A}}$ and $z \in \mathsf{A}$:*

$$|\mathsf{D}(\hat{z}; z) - \mathsf{D}(\hat{z}'; z)| \leqslant O((\log k)^2) W^1_T(\hat{z}, \hat{z}') \,.$$

B. Dynamic operations on T

We will now discuss changes to the structure of T; since there will be multiple trees in our discussion, we will use the notations $V^T, \mathcal{L}^T, \mathsf{ch}^T, \mathsf{A}^T$, etc., to distinguish the corresponding objects.

Fusion and zippering. Consider siblings $v, v' \in \mathsf{ch}^T(u)$ for some $u \in V^T$. We define the *fusion of T along* $\{v, v'\}$ as the tree $T_{vv'}$ with $V^{T_{vv'}} = (V^T \setminus \{v, v'\}) \cup \{v^*\}$, where $v^* \notin V^T$. The structure of $T_{vv'}$ is that of of the induced tree $T[V^T \setminus \{v, v'\}]$, except that v^* is a child of u and $\mathsf{ch}^{T_{vv'}}(v^*) = \mathsf{ch}^T(v) \cup \mathsf{ch}^T(v')$. See Figure 2. We take $w^{T_{vv'}}_{v^*} = w^T_v = w^T_{v'}$, and $w^{T_{vv'}}_y = w^T_y$ for $y \in V^T \setminus \{v, v'\}$.

For a sequence σ of numbers, let $\mathrm{sort}(\sigma)$ denote the sequence sorted in non-increasing order. For any $z \in \mathsf{A}^T$, we define a new point $z^{vv'} \in \mathsf{A}^{T_{vv'}}$ as follows: $z^{vv'}_{y,i} = z_{y,i}$ for $y \notin \{v, v'\}$, and

$$\langle z^{vv'}_{v^*,i} : i \in \mathbb{Z}_+\rangle = \mathrm{sort}(\langle z_{v,i}, z_{v',j} : i, j \in \mathbb{Z}_+\rangle)\,.$$

The next lemma is straightforward to verify.

Lemma II.3. *For any $\hat{z} \in \hat{\mathsf{A}}$ and $z \in \mathsf{A}$, it holds that*

$$\mathsf{D}^{T_{vv'}}\left(\hat{z}^{vv'}; z^{vv'}\right) \leqslant \mathsf{D}^{T}\left(\hat{z}; z\right).$$

Fission. We also define a fission operation on T. Suppose that $u \in V^T \setminus \{r\}$. Then a *fission of T at u* is any tree T' with $V^{T'} \supseteq V^T$ satisfying:

1) For all $v \in V^T \setminus \{\mathsf{p}(u), u\}$: $\mathsf{ch}^T(v) = \mathsf{ch}^{T'}(v)$.
2) It holds that

$$\bigcup_{v \in \mathsf{ch}^T(\mathsf{p}^T(u))} \mathsf{ch}^T(v) \subseteq \bigcup_{v \in \mathsf{ch}^{T'}(\mathsf{p}^{T'}(u))} \mathsf{ch}^{T'}(v)\,.$$

3) It holds that $v \in V^T \implies w^{T'}_v = w^T_v$ and also that $v \in \mathsf{ch}^{T'}(\mathsf{p}(u)) \implies w^{T'}_v w^T_u$.

For any such T', we say that *T' is a fission of T*. See Figure 3.

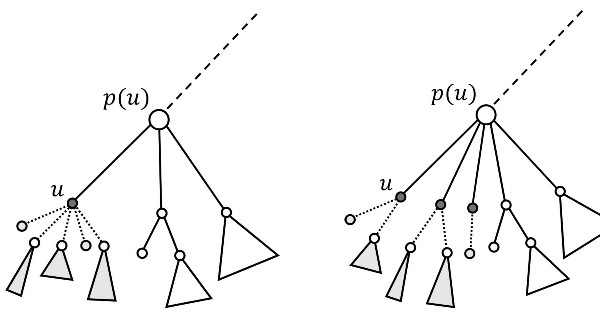

Figure 3. Fission of T at $u \implies T'$

For two sequences σ, σ', we write $\sigma \succeq \sigma'$ if $\sigma_i \geqslant \sigma'_i$ for every $i \geqslant 1$. If T' is a fission of T at u and $z \in \mathsf{A}^T$, then a point $z' \in \mathsf{A}^{T'}$ is *(T, T')-consistent with z* if $z'_{y,i} = z_{y,i}$ for every $i \geqslant 1$ and every $y \in V^T \setminus \mathsf{des}^T(u)$. We say that z' *(T, T')-dominantes* z if the following holds:

$$y \in \left(V^T \setminus \mathsf{des}^T(u)\right) \cup \left(V^T \cap \mathsf{des}^{T'}(u)\right) \implies \langle z'_{y,i} : i \geqslant 1\rangle \succeq \langle z_{y,i} : i \geqslant 1\rangle\,.$$

Note, in particular, that if z' -dominates z, then $z'_u = z_u$, i.e., all the fractional server mass below u in T sits below u in T'.

Example II.4. Consider the tree T with edges $\{(r, u), (u, x), (u, y)\}$ and the fission T' of T with edges $\{(r, u), (r, v), (u, x), (v, y)\}$. If $z_{r,1} = z_{r,2} = z_{u,1} = z_{u,2} = z_{x,1} = z_{y,1} = 1$ are the non-zero coordinates of $z \in \mathsf{A}^T$, and $z'_{r,1} = z'_{r,2} = z'_{u,1} = z'_{u,2} = z'_{x,1} = z'_{x,2} = 1$ are the non-zero coordinates of $z' \in \mathsf{A}^{T'}$, then z' is (T, T')-consistent with z and additionally (T, T')-dominates z.

On the other hand, if we define $z'_{r,1} = z'_{r,2} = z'_{u,1} = z'_{v,1} = z'_{x,1} = z'_{y,1} = 1$ as the non-zero coordinates of $z' \in \mathsf{A}^{T'}$, then z' is (T, T')-consistent with z but does not (T, T')-dominate z.

Lemma II.5. *Suppose that T is a 4-HST and that T' is a fission of T at $u \in V^T$. Consider points $\hat{z} \in \hat{\mathsf{A}}^T, \hat{z}' \in \hat{\mathsf{A}}^{T'}, z \in \mathsf{A}^T, z' \in \mathsf{A}^{T'}$ such that $\hat{z}'$ is (T, T')-consistent with $\hat{z}$, z' is (T, T')-consistent with z', and z' (T, T')-dominates z. Then:*

$$\mathsf{D}^{T'}(\hat{z}'; z') - \mathsf{D}^{T}(\hat{z}; z) \leqslant 4\alpha_0 w_u \min(z_u, \hat{z}_u)\left(1 - \frac{\hat{z}'_u}{\hat{z}_u}\right)\log\left(\frac{1+\eta}{\eta}\right).$$

III. Random partitions and HSTs

A *semipartition H of X* is a collection of pairwise disjoint subsets of X. Say that H is *Δ-bounded* if $S \in H \implies \mathrm{diam}_X(S) \leqslant \Delta$. Say that H is *r-separated* if $S, S' \in H$ and $S \neq S'$ implies $d_X(S, S') > r$. We use $[\![H]\!] := \bigcup_{S \in H} S$ to denote the set of points covered by H.

If P is a partition of X, we will sometimes think of P as a function that takes $x \in X$ to the unique set $P(x) \in P$ containing x. If $A \subseteq X$, we write $P(A)$ for the collection of sets $S \in P$ such that $S \cap A \neq \emptyset$. We define the *cut semimetric:*

$$\Delta_P(x, y) := \frac{1}{2}\sum_{S \in P} |\mathbb{1}_S(x) - \mathbb{1}_S(y)|\,.$$

Fix $\tau := 4$, and let $\mathcal{S} := \{0, -1, -2, -3, \ldots, j_{\min}\}$ denote the set of scales, where $j_{\min}$ is chosen so that $\tau^{j_{\min}} < \min_{x \neq y} d_X(x, y)$. Given a sequence $\mathcal{P} = \langle P^j : j \in \mathcal{S}\rangle$ of partitions of X, we define an associated distance on X:

$$\mathrm{dist}_{\mathcal{P}}(x, y) := \sum_{j \in \mathcal{S}} \tau^j \Delta_{P^j}(x, y)\,.$$

We write $\tilde{\mathcal{P}} = \langle \tilde{P}^j : j \in \mathcal{S}\rangle$ for the hierarchical refinement of $\mathcal{P}$. In other words, for each $j \in \mathcal{S}$, we define $\tilde{P}^j$ as the common refinement of the partitions $\langle P^i : i \geqslant j\rangle$.

Define a tree $T^{\mathcal{P}}$ whose nodes are the elements $\{(S, j) : S \in \tilde{P}^j, j \in \mathcal{S}\}$ and such that $(S', j-1)$ is a child of (S, j) whenever $S' \subseteq S$. There is a canonical mapping between the leaves of $T^{\mathcal{P}}$ (which are singleton sets) and the points of X. If we define $w^{\mathcal{P}}(S, j) := \tau^j$, then the resulting weighted tree $(T^{\mathcal{P}}, w^{\mathcal{P}})$ satisfies

$$\mathrm{dist}_{\mathcal{P}}(x, y) \leqslant \mathrm{dist}_{T^{\mathcal{P}}, w^{\mathcal{P}}}(x, y) \leqslant 2\,\mathrm{dist}_{\mathcal{P}}(x, y) \quad \forall x, y \in X\,.$$

Moreover, $(T^{\mathcal{P}}, w^{\mathcal{P}})$ is a τ-HST by construction.

In the following sections, there will be parameters $\alpha, \gamma, \delta, \lambda, \varepsilon > 0$ of the orders: $\alpha \asymp 1, \gamma \asymp (\log k)^2, \delta \asymp (\log k)^{-2}, \lambda \asymp (\log k)^{O(1)}, \varepsilon = k^{-O(1)}$.

A. Cluster fusion

Suppose P is a semipartition of X. For a subset $S \subseteq X$, define

$$S^P := \bigcup_{\substack{C \in P: \\ C \cap S \neq \emptyset}} C\,.$$

This is the set S fused with all the sets $C \in P$ that intersect it. (See Figure 4.)

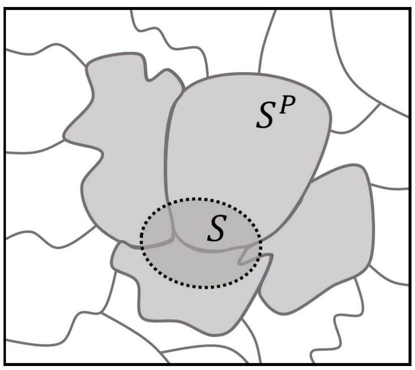

Figure 4. The set S fused with the sets of P it intersects.

Lemma III.1. *For any $\gamma > 1$ and $\Delta > 0$, the following holds. If P and H are semipartitions of X such that P is Δ-bounded and H is $\gamma\Delta$-bounded and $2(\gamma+1)\Delta$-separated, then the collection of subsets*

$$P[H] := \left\{S^P : S \in H\right\} \cup \left\{S' \in P : S' \cap S^P = \emptyset \quad \forall S \in H\right\}$$

is a $(\gamma+2)\Delta$-bounded semipartition of X.

B. Random partitions

We will suppose that we have a *random* τ-adic system $\mathcal{P} = \langle P^j : j \in \mathcal{S}\rangle$ of partitions, where P^j is τ^j-bounded and such such that for every $j \in \mathcal{S}$:

$$\mathbb{P}\left[P^j(x) \neq P^j(y)\right] \leqslant \beta \frac{d_X(x,y)}{\tau^j} \qquad \forall x, y \in X\,.$$

For instance, if $X \subseteq \mathbb{R}^d$, it is well-known that such a system exists with $\beta \leqslant O(\sqrt{d})$ [CCG+98].

We will also maintain an (*independent*) random collection $\mathcal{H} = \langle H^j : j \in \mathcal{S}\rangle$ such that H^j is $\gamma\tau^j$-bounded and $3\gamma\tau^j$-separated. We will also choose $\gamma \geqslant 4(\alpha+1)$ so that $3\gamma \geqslant 2(\gamma+2\alpha+2)$. Let $\hat{H}^j := \{B_X(S, (\alpha+1)\tau^j) : S \in H^j\}$. Note that $\hat{H}^j$ is $2(\gamma+1)\tau^j$-separated by construction.

We denote by $\mathcal{Q} := \langle Q^j : j \in \mathcal{S}\rangle$ the fused sequence of partitions formed by $Q^j := P^j[\hat{H}^j]$. The random fused sequence $\mathcal{Q}$ represents the random HST structure underlying our algorithm; we will dynamically modify the sequence $\mathcal{H}$ as requests arrive, thereby altering the structure of the corresponding HSTs.

C. Randomized heavy covers

For a point $x \in X$ and $\nu \in \mathbb{M}(X)$, define the *isolation radius:*

$$\rho(x,\nu) := \sup\{r : \nu(B_X(x,r)) < 1/4\}\,.$$

Recall that $\alpha, \lambda > 1$ and $\varepsilon > 0$ are parameters we will choose later. Suppose that $\nu \in \mathbb{M}(X)$. Say that a random subset $\Lambda \subseteq X$ is an *L-heavy cover for ν at scale r* if

$$\begin{aligned} & r > \rho(x,\nu) \implies \\ & \quad \mathbb{P}\left[d_X(x,\Lambda) > \alpha r\right] \leqslant L \log\left(\frac{\nu(B_X(x,\lambda r))}{\nu(B_X(x,r))}\right). \end{aligned}$$

The next lemma shows how fusing around a heavy cover is used to bound the expected stretch in the HST metric $\mathrm{dist}_{\mathcal{Q}}$.

Lemma III.2. *Consider some $\nu \in \mathbb{M}_k(X)$. Suppose that for every $j \in \mathcal{S}$, the set $[\![H^j]\!]$ is an L-heavy cover for ν at scale τ^j. Then for every $x, y \in X$ and $\varepsilon \in (0, 1/2]$, it holds that*

$$\begin{aligned} & \mathbb{E}[\mathrm{dist}_{\mathcal{Q}}(x,y)] \\ & \quad \leqslant O\big((\beta L \log(\lambda)\log k) + \log(\varepsilon^{-1})\big)\, d_X(x,y) + \varepsilon \rho(x,\nu)\,. \end{aligned}$$

Proof: Note that by construction, $d_X(x, [\![H^j]\!]) \leqslant \alpha\tau^j$ implies the containment $B_X(x,\tau^j) \subseteq Q^j(x)$. Define:

$$\begin{aligned} i_0 &:= \lceil \log_\tau \max((\varepsilon/2)\rho(x,\nu), d_X(x,y))\rceil \\ i_1 &:= \lceil \log_\tau \max(2\rho(x,\nu), d_X(x,y))\rceil. \end{aligned}$$

Denote $p_j(x) := \mathbb{P}\left[d_X(x, [\![H^j]\!]) > \alpha\tau^j\right]$ and write:

$$\begin{aligned} & \sum_{\substack{j \in \mathcal{S}: \\ j \geqslant i_1}} \tau^j\, \mathbb{E}\left[\Delta_{Q^j}(x,y)\right] \\ & \quad \leqslant \sum_{\substack{j \in \mathcal{S}: \\ j \geqslant i_1}} \tau^j p_j(x)\, \mathbb{E}\left[\Delta_{P^j}(x,y)\right] \\ & \quad \leqslant \beta d_X(x,y) \sum_{\substack{j \in \mathcal{S}: \\ j \geqslant i_1}} p_j(x) \\ & \quad \leqslant L\beta d_X(x,y) \sum_{\substack{j \in \mathcal{S}: \\ j \geqslant i_1}} \log\left(\frac{\nu(B(x,\lambda\tau^j))}{\nu(B(x,\tau^j))}\right) \\ & \quad \leqslant O(L\beta \log\lambda)\, d_X(x,y) \log\left(\frac{\nu(X)}{\nu(B(x, 2\rho(x,v)))}\right) \\ & \quad \leqslant O(L\beta \log(\lambda)\log(k))\, d_X(x,y)\,. \end{aligned}$$

In the third inequality, we have used the fact that $[\![H^j]\!]$ is an L-heavy cover for ν at scale τ^j and $j \geqslant i_1 \implies \tau^j > \rho(x,\nu)$.

We can also bound

$$\sum_{j=i_0}^{i_1-1} \tau^j \, \mathbb{E}\left[\Delta_{Q^j}(x,y)\right] \leqslant \sum_{j=i_0}^{i_1-1} \tau^j p_j(x) \leqslant O(\log \varepsilon^{-1}) d_X(x,y),$$

and, finally by a geometric summation:

$$\sum_{\substack{j\in\mathcal{S}:\\ j<i_0}} \tau^j \, \mathbb{E}\left[\Delta_{Q^j}(x,y)\right] \leqslant \sum_{\substack{j\in\mathcal{S}:\\ j<i_0}} \tau^j \leqslant 2d_X(x,y) + \varepsilon\rho(x,\nu)\,,$$

completing the proof.

IV. Heavy chains and the annealed measure

Recall that Q specifies a τ-HST (T^Q, w^Q) whose leaves are associated to points of X. Thus the request sequence σ naturally gives a request sequence on $\mathcal{L}^{T^Q}$. We will maintain a random fractional k-server measure $\mu^Q \in \mathbb{M}_k(X)$ and a corresponding random point $z^Q \in \mathsf{A}^{T^Q}$ such that $\nu^{z^Q} = \mu^Q$.

Define the *annealed measure* $\bar{\mu} \in \mathbb{M}_k(X)$ by

$$\bar{\mu} := \mathbb{E}\left[\mu^Q\right] .$$

We will not update Q directly, but instead the sets $\mathcal{H} = \langle H^j : j \in \mathcal{S}\rangle$ from which Q is formed.

Recall that, for every $j \in \mathcal{S}$, we need H^j to be $\gamma\tau^j$-bounded and $3\gamma\tau^j$-separated. We will require some additional properties. For convenience, we will assume that every $A \in H^j$ comes equipped with a representative $\mathcal{R}^j(A) \in A$. We will also assume that every H^j is partitioned into levels:

$$H^j = \bigcup_{\ell\geqslant 0} H^{j,\ell}\,,$$

and the following properties hold.

H1 For all $j \in \mathcal{S}$:

$$A, A' \in H^j \text{ and } A \cap A' = \emptyset \implies d_X(\mathcal{R}^j(A), \mathcal{R}^j(A')) > 5\gamma\tau^j.$$

H2 For all $j \in \mathcal{S}$ and $\ell \geqslant 0$: If $A \in H^{j,\ell}$ and $A' \in H^{j-1,\ell}$, then either $d_X(\mathcal{R}^j(A), \mathcal{R}^j(A')) > 5\gamma\tau^j$, or $A' \subseteq A$.

Roughly speaking, the preceding conditions imply that the sets $\langle H^{j,\ell} : j \in \mathcal{S}\rangle$ form increasing chains (H2) and distinct chains are sufficiently separated (H1). Note that since H^j is a $\gamma\tau^j$-bounded semipartition, condition (H1) implies that H^j is $3\gamma\tau^j$-separated.

A. Randomized mass levels

Let $j_0 \geqslant 1$ be parameters we will choose soon. For a number $\theta \in [0,1]$ and $m \in \mathbb{R}_+$, denote

$$\begin{aligned}
\ell_\theta(m) &:= \lceil \theta + \log_8(m) \rceil \\
M^j(x) &:= \mathbb{E}\left[\mu\left(\llbracket \tilde{Q}^{j+j_0}(B_X(x, 5\gamma\tau^j))\rrbracket\right)\right] \\
L^j_\theta(x) &:= \ell_\theta(M^j(x)) \\
M^j_\theta(x) &:= 8^{L^j_\theta(x)}
\end{aligned}$$

For $A \in H^j$, we use the notations $M^j(A) := M^j(\mathcal{R}^j(A))$, $L^j_\theta(A) := L^j_\theta(\mathcal{R}^j(A))$, $M^j_\theta(A) := M^j_\theta(\mathcal{R}^j(A))$. We will make use of the following fact.

Fact IV.1. *If $S' \subseteq S \subseteq X$ and $i' \leqslant i$, then*

$$\llbracket \tilde{Q}^i(S) \rrbracket \supseteq \llbracket \tilde{Q}^{i'}(S') \rrbracket\,.$$

We choose $\theta \in [0,1]$ uniformly at random. If $A \in H^j$ and the value $L^j_\theta(A)$ changes under HST movement, we are going to remove A from H^j (see lines (2)–(3) in Algorithm 1 below). Note the following estimates:

$$\bar{\mu}\left(B_X(x, 5\gamma\tau^j)\right) \leqslant M^j(x) \leqslant \bar{\mu}\left(B_X(x, C\tau^j)\right) \tag{IV.1}$$

for some $C \leqslant O((\gamma+\alpha)\tau^{j_0})$. We will enforce the following additional requirement:

H3 For $j \in \mathcal{S}$ and $\ell \geqslant 0$: $A \in H^{j,\ell} \implies L^j_\theta(A) = \ell$.

The heavy cover property. For $x \in X$ and $j \in \mathcal{S}$, say that *x is τ^j-regular* if

C0 $\bar{\mu}(B(x,\tau^j)) > (1-\delta)\bar{\mu}(B(x,\lambda\tau^j))$.

C1 $\ell_\theta(\bar{\mu}(B_X(x,\tau^j))) = L^j_\theta(x)$, and

C2 $L^j_\theta(x) = L^{j+s}_\theta(y)$ $\forall y \in B(x, 2\gamma\tau^{j+s})$ and $s \in \{-1, 1\}$.

We will enforce the following property:

H4 For every $x \in X$ and $j \in \mathcal{S}$: If x is τ^j-regular, then $d_X(x, H^j) \leqslant \alpha\tau^j$.

This implies that H^j is an L-heavy cover for $\bar{\mu}$ at scale τ^j with $L \leqslant O(1/\delta)$, as follows. If (C0) doesn't hold, then $\log \frac{\bar{\mu}(B(x,\lambda\tau^j))}{\bar{\mu}(B(x,\tau^j))} > \Omega(\delta)$, and therefore the L-heavy cover condition with $L = O(1/\delta)$ is vacuous.

Assume now that (C0) does hold. Note from (IV.1) that

$$\max_{y\in B_X(x,2\gamma\tau^{j+s})} M^{j+s}_\theta(y) \leqslant \bar{\mu}\left(B\left(x, \lambda\tau^j\right)\right)$$

for some choice of $\lambda \leqslant O((\gamma+\alpha)\tau^{j_0})$. Moreover using the LHS of (IV.1):

$$\min_{y\in B_X(x,2\gamma\tau^{j+s})} M^{j+s}_\theta(y) \geqslant \bar{\mu}\left(B(x, 3\gamma\tau^{j+s}\right) \geqslant \bar{\mu}\left(B(x,\tau^j)\right)\,.$$

Thus the probability (over the choice of θ) that one of (C1) or (C2) is violated is at most

$$O\left(\log \frac{\bar{\mu}(B_X(x,\lambda\tau^j))}{\bar{\mu}(B_X(x,\tau^j))}\right),$$

as desired.

V. The algorithm and competitive analysis

The state of the HST algorithm is specified by z^Q and the corresponding fractional server measure μ^Q. We will use opt $\in \widehat{\mathbb{M}}_k(X)$ to denote some optimal offline solution.

A. Potential functions

First, we define a family of auxiliary potential functions. For $\mu \in \mathbb{M}(X)$, define the *fusion potential:*

$$\psi^F(\mu) := -\sum_j \tau^j \sum_{\ell \geqslant 0} \sum_{A \in H^{j,\ell}} \mu(Q^j(A)) .$$

Note that ψ^F is 1-Lipschitz in the $W^1_{T^Q}$ distance.

For $A \subseteq X$, denote the set of "children" of A:

$$C^{j-1}(A) := \left\{ S \in Q^{j-1} : d_X(S, A) \leqslant \tau^j \right\} ,$$

and define the *kinetic potential* by

$$\psi^K(\mu) := -\sum_j \tau^{j-1} \sum_{\ell \geqslant 0} \sum_{A \in H^j} \mu\left(\llbracket C^{j-1}(A) \rrbracket\right).$$

Note that ψ^K is 1-Lipschitz in the $W^1_{T^Q}$ distance.

Chain potentials. For $A \in H^{j,\ell}$, we use $H^{j-1}(A)$ to denote the unique (if it exists) set $A' \in H^{j-1,\ell}$ with $A' \subseteq A$. If no such A' exists, we take $H^{j-1}(A) = \emptyset$. Define:

$$\begin{aligned} \mathrm{opt}^j_\theta(A) &:= \min\left(\mathrm{opt}(Q^j(A)), M^j_\theta(A)\right) \\ E^j_\theta(A) &:= \mathrm{opt}^j_\theta(A) \log\left(\mathrm{opt}\left(Q^j(H^j(A))\right)\right) \end{aligned}$$

Moreover, denote:

$$\psi^C_* := \sum_{j \in \mathcal{S}} \sum_{A \in H^j} \left[\mathrm{diam}_X(H^{j-1}(A)) - \mathrm{diam}_X(A)\right] E^j_\theta(A) ,$$

$$\psi^C := \sum_{j \in \mathcal{S}} \sum_{A \in H^j} \left[\tau^j + \mathrm{diam}_X(A) - \mathrm{diam}_X(H^{j-1}(A))\right] M^j_\theta(A) .$$

(Note that we take $0 \log 0 = 0$ by convention.)

Fact V.1. *Since H^j is $\gamma\tau^j$-bounded, when* opt *moves distance D in T^Q, the value of ψ^C_* changes by at most $O(D\gamma \log k)$.*

Our overall potential function is given as follows.

Definition V.2. We define:

$$\Phi := \mathbb{E}\left[\mathsf{D}(\hat{z}^Q; z^Q) + \varepsilon_0\left(\psi^F(\mu) + \psi^K(\mu) + \varepsilon_1 \psi^C\right) + \frac{C_0 (\log k)^2}{\gamma} \psi^C_* \right],$$

where $\varepsilon_0, \varepsilon_1 > 0$ and $C_0 > 1$ are constants we will choose later, and $\hat{z}^Q$ is the element of $\hat{\mathsf{A}}^{T^Q}$ corresponding to opt.

B. Responding to a request

When a request $\sigma_t \in X$ arrives at time t, we use the algorithm from Section II-A to respond in each HST:

$$z^Q := \Gamma^{T^Q}(z^Q, \sigma_t) .$$

This induces a change in $\bar{\mu}$, so the sequence $\mathcal{H}$ perhaps no longer satisfies (H1)–(H4). Algorithm 1 addresses this.

We will use the following notation in our analysis: If H is a semipartition of X and $x \in X$, then $\mathcal{C}(H, x)$ is some set $A \in H$ such that $d_X(x, A) = d_X(x, \llbracket H \rrbracket)$. We will only use this notation in situations where the minimizing set $A \in H$ is unique.

C. Update analysis

Lines (19)–(22) ensure that H^j is $\gamma\tau^j$-bounded. The deletion in lines (6)–(10) ensures that (H1) holds as long as $\lambda \geqslant 5\gamma + \tau(\gamma + \alpha + 2)$. The first loop (1)–(3) enforces (H3), and the second loop (4)–(28) enforces (H4). (H2) is enforced downward by (9)–(11) and upward by the augmentation loop (18)–(28); moreover, DeleteBelow performs deletions in a downward-closed manner.

D. Mass level maintenance

Let us argue that the expected cost of the first loop (1)–(3) can be charged against a small portion of the expected movement cost of the HST algorithm in responding to σ_t.

Indeed, for any $m \geqslant m' > 1/4$:

$$\mathop{\mathbb{P}}_{\theta \in [0,1]}[\ell_\theta(m) \neq \ell_\theta(m')] \leqslant O(1) \left|\log \frac{m}{m'}\right| .$$

Therefore if p is the probability that $L^j_\theta(x)$ changes in updating μ^Q, then this entails average movement $p \cdot m\tau^{j+j_0}$ of server mass into or out of the sets $\tilde{Q}^{j+j_0}(B_X(x, 5\gamma\tau^j))$, where $m = \bar{\mu}(\tilde{Q}^{j+j_0}(B_X(x, 5\gamma\tau^j)))$. Moreover, applying this to $x = \mathcal{R}^j(A)$ for $A \in H^j$, we see that we can charge a $\mu^Q(\tilde{Q}^j(A))$ portion of this movement to A since the sets in $\{\tilde{Q}^j(A) : A \in H^j\}$ are pairwise disjoint.

On the other hand, the increase of Φ due to the deletion of A is bounded by $O(\tau^j (\log k)^2 \mu(\tilde{Q}^j(A)))$. Hence for $j_0 \asymp \log_\tau((\log k)^2) \asymp \log\log k$, this can be charged against an arbitrarily small fraction of the HST movement cost.

E. Gain from insertion

Our next goal is to show that $\mathbb{E}[\psi^F(\mu^Q) + \psi^K(\mu^Q)]$ decreases by $-\Omega(\tau^j)\bar{\mu}(B_X(x, \tau^j))$ whenever a $\bar{\mu}$-heavy pair (x, τ^j) is found with $\ell = L^j_\theta(x)$ and $d_X(x, \llbracket H^{j,\ell} \rrbracket) > \alpha\tau^j$.

If x is inserted via lines (9)–(11) or (13), then $d_X(x, \llbracket H^j \rrbracket) > 2\alpha\tau^j$, and the decrease comes in $\mathbb{E}[\psi^F(\mu^Q)]$ via the addition of $\{x\}$ to H^j.

Algorithm 1 Heavy chain maintenance

1: For $j \in \mathcal{S}, \ell \geqslant 0$:
2: While $\exists A \in H^{j,\ell}$ with $L_\theta^j(A) \neq \ell$:
3: Remove A from $H^{j,\ell}$
4: For $j \in \mathcal{S}$ in decreasing order:
5: While $\exists x \in X$ such that (x is τ^j-regular and $d_X(x, [\![H^j]\!]) > \alpha\tau^j$):
6: $\ell := L_\theta^j(x)$
7: If $d_X(x, [\![H^{j,\ell}]\!]) \leqslant 2\alpha\tau^j$, then:
8: $A^j := \mathcal{C}(H^{j,\ell}, x)$
9: Else if $d_X(x, [\![H^{j-1,\ell}]\!]) \leqslant 2\alpha\tau^j$, then:
10: $A^j := \mathcal{C}(H^{j-1,\ell}, x) \cup \{x\}$
11: $H^{j,\ell} := H^{j,\ell} \cup \{A^j\}$; $\mathcal{R}^j(A^j) := x$
12: Else:
13: $A^j := \{x\}$; $H^{j,\ell} := H^{j,\ell} \cup \{A^j\}$; $\mathcal{R}^j(A^j) := x$
14: If $d_X(x, [\![H^{j+1,\ell}]\!]) \leqslant 2\alpha\tau^{j+1}$, then:
15: $A^{j+1} := \mathcal{C}(H^{j+1,\ell}, x)$
16: Let $A^{j+1} \subseteq A^{j+2} \subseteq \cdots \subseteq A^h$ be a max. chain s.t. $A^i \in H^{i,\ell}$ for $i = j+1, j+2, \ldots, h$
17: Else: $h := j$
18: For $i = j, j+1, \ldots, h$:
19: If $\mathrm{diam}_X(A^i \cup \{x\}) > (\gamma - 2\alpha)\tau^i$, then:
20: If $L_\theta^i(x) = \ell$, then:
21: $H^{i,\ell} := (H^{i,\ell} \setminus \{A^i\}) \cup \{H^{i-1}(A^i) \cup \{x\}\}$
22: $\mathcal{R}^i(H^{i-1}(A^i) \cup \{x\}) := x$
23: DeleteBelow$(x, i+1, j)$
24: Else:
25: $H^{i,\ell} := H^{i,\ell} \setminus \{A^i\}$
26: Else:
27: $H^{i,\ell} := (H^{i,\ell} \setminus \{A^i\}) \cup \{A^i \cup \{x\}\}$
28: $\mathcal{R}^i(A^i \cup \{x\}) := \mathcal{R}^i(A^i)$

29: DeleteBelow(x, i, j):
30: For $s \leqslant h$:
31: While $\exists A \in H^s$ such that $d_X(x, A) \in [\tau^j + (\alpha+2)\tau^i, 6\gamma\tau^i]$:
32: For $s' \leqslant s$ and $A' \in H^{s'}$:
33: If $A' \subseteq A$, then $H^{s'} := H^{s'} \setminus \{A'\}$.

Let us now consider the other case (lines (7)–(8)): First, note that $B_X(x, \tau^j) \cap [\![C^{j-1}(A^j)]\!] = \emptyset$. Since x is τ^j-regular, any fused set $S \in Q^{j-1}$ with $d_X(x, S) \leqslant \tau^j$ must satisfy $L_\theta^j(S) = L_\theta^j(x)$ and hence $S \subseteq A$ by (H2). On the other hand, any unfused $S \in Q^{j-1}$ satisfies $\mathrm{diam}_X(S) \leqslant \tau^{j-1}$, hence if $d_X(x, S) \leqslant \tau^j$, then:

$$\begin{aligned} d_X(S, A^j) &> d_X(x, A^j) - d_X(x, S) - \mathrm{diam}_X(S) \\ &> \alpha\tau^j - \tau^j - \tau^{j-1} > \tau^j \end{aligned}$$

for $\alpha \geqslant 3$.

Next, observe that $B_X(x, \tau^j) \subseteq [\![C^{j-1}(A^j \cup \{x\})]\!]$. This holds because any set $S \in Q^{j-1}$ with $S \cap B_X(x, \tau^j) \neq \emptyset$ must satisfy $d_X(S, A^j \cup \{x\}) \leqslant d_X(S, x) \leqslant \tau^j$. Since $[\![C^{j-1}(A)]\!] = \emptyset$ and $B_X(x, \tau^j) \subseteq [\![C^{j-1}(A')]\!]$, we conclude that $\mathbb{E}[\psi^K(\mu^Q)]$ decreases by $-\Omega(\tau^j \bar{\mu}(x, \tau^j))$ when $A^j \cup \{x\}$ replaces A^j in H^j.

Note that ψ^C increases under insertion by $O(\tau^j M_\theta^j(x))$. By choosing ε_1 small enough, this can be charged against a portion of the $-\Omega(\tau^j \bar{\mu}(B(x, \tau^j))$ decrease in $\psi^F + \psi^K$.

F. Loss from deletion of fused clusters.

Note that when line (23) is executed, there is a corresponding increase in $\mathsf{D}(\hat{z}^Q; z^Q)$, as well as ψ^F and ψ^K. But the total expected increase is bounded by

$$\begin{aligned} O((\log k)^2)\tau^i \bar{\mu} & \left(B_X(x, 7\gamma\tau^i) \setminus B_X(x, \tau^j)\right) \\ & \leqslant O((\log k)^2 \tau^i \delta \bar{\mu}\left(B_X(x, \tau^j)\right), \end{aligned}$$

where the latter inequality follows from the fact that x is τ^j-regular and $L_\theta^i(x) = L_\theta^j(x)$.

On the other hand, for the removal of a fused set $A \in H^i$, there is a corresponding decrease in ψ^C of $-\Omega(\tau^i M_\theta^i(A))$. For any set $A \in H^i$ deleted by the invocation of (23), τ^j-regularity of x implies that $M_\theta^i(A) \geqslant \Omega(\bar{\mu}(B(x, \tau^j)))$.

Thus choosing $\delta \asymp \varepsilon_1 (\log k)^{-2}$ sufficiently small, the increase in $\mathsf{D} + \varepsilon_0(\psi^F + \psi^K)$ can be charged against the decrease in $\varepsilon_0\varepsilon_1\psi^C$.

G. Loss from level change

In line (25), a set A^i is removed from H^i because none of its points are τ^i-regular (since the check (20) failed). The cost of removing this set is $O(\tau^i M_\theta^i(A^i)(\log k)^2)$. On the other hand, its removal yields a commensurate decrease in ψ^C by an amount $-\Omega(\gamma\tau^i M_\theta^i(A^i))$. Thus for $\gamma \asymp (\log k)^2$ chosen sufficiently large, the decrease in ψ^C pays for the increase in the other potentials.

H. Augmentation

Note that the value $\psi^F(\mu^Q) + \psi^K(\mu^Q)$ does not decrease under the augmention step (25)–(26). By Lemma II.3, the same is true for D. It is straightforward to calculate that under augmentation, the change in ψ_*^C is bounded by $O(\tau^j \bar{\mu}(B(x, \tau^j)) \log k)$, and the change in ψ^C is bounded by $O(\tau^j \bar{\mu}(B(x, \tau^j)))$.

I. Shrinkage

Let us now analyze the potential change due to the shrinkage in lines (21)–(22). This step can be implemented on T^Q via a sequence of fission operations at nodes corresponding to sets in $\tilde{Q}^i(Q^i(A^i))$. For each fission T' of T along this sequence, we update so that $\boldsymbol{z}'^Q$ (T, T')-dominates $\boldsymbol{z}^Q$. In the setting of μ^Q, this means that the fractional server mass in the sets $\tilde{Q}^i(Q^i(A^i))$ is *moved into the sets* $\tilde{Q}^i(Q^i(A^{i-1}))$.

From Lemma II.5, we know that the change D is bounded by

$$O(\tau^i(\log k)^2)\,\mathrm{opt}^i_\theta(A^i)\left(1-\frac{\mathrm{opt}(Q^{i-1}(A^{i-1}))}{\mathrm{opt}(Q^i(A^i))}\right) \qquad \text{(V.1)}$$
$$\leqslant O(\tau^i(\log k)^2)\,\mathrm{opt}^i_\theta(A^i)\log\left(\frac{\mathrm{opt}(Q^i(A^i))}{\mathrm{opt}(Q^{i-1}(A^{i-1}))}\right),$$

where we have used the fact that $1-\frac{a}{a+b}\leqslant\log(\frac{a+b}{a})$ for all $a,b\geqslant 0$.

Let us now calculate the corresponding change in ψ^C_*. Write $b_i := \mathrm{opt}(Q^i(A^i))$ and $a_i := \min(b_i, M^i_\theta(A^i))$. Then the change in ψ^C_* due to this shrinkage step is at most

$$(\mathrm{diam}_X(A^{i-1})-\mathrm{diam}_X(A^i))\left[a_i\log b_i - a_{i-1}\log b_{i-1}\right]. \qquad \text{(V.2)}$$

Since $\mathrm{diam}_X(A^{i-1})\leqslant\gamma\tau^{i-1}$ and

$$\mathrm{diam}_X(A^i)\geqslant\mathrm{diam}_X(A^i\cup\{x\})-2\alpha\tau^j\geqslant(\gamma-4\alpha)\tau^j,$$

we have

$$\mathrm{diam}_X(A^{i-1})-\mathrm{diam}_X(A^i)\leqslant-(\gamma-4\alpha)\tau^i/2\leqslant-\gamma\tau^i/4.$$

Moreover,

$$\begin{aligned}a_i\log b_i-a_{i-1}\log b_{i-1} &= (a_i-a_{i-1})\log b_i + a_{i-1}\log\tfrac{b_i}{b_{i-1}}\\ &\geqslant (a_i-a_{i-1})\log\tfrac{b_i}{b_{i-1}} + a_{i-1}\log\tfrac{b_i}{b_{i-1}}\\ &= \mathrm{opt}^i_\theta(A^i)\log\left(\frac{\mathrm{opt}(Q^i(A^i))}{\mathrm{opt}(Q^{i-1}(A^{i-1}))}\right).\end{aligned}$$

Thus for C_0 chosen large enough in Definition V.2 it holds that (V.2) compensates for (V.1).

Because we move the fractional server mass from $\tilde{Q}^i(Q^i(A^i))$ to $\tilde{Q}^i(Q^i(A^{i-1}))$, the sum $\psi^F(\mu^Q)+\psi^K(\mu^Q)$ decreases by $-\Omega(\tau^{i-1}d\mu)$, where $d\mu$ is the amount of displaced server mass. On the other hand, The corresponding movement cost in dist_Q is $O(\tau^i d\mu)$.

J. The competitive ratio

To analyze the competitive ratio, we consider the evolution of Φ as we respond to a request. When opt moves in (X,d_X) in response to the request σ_t, we can use Lemma II.2, Lemma III.2, and Fact V.1 to bound the change in Φ:

$$\begin{aligned}\Delta_{\mathrm{opt}}\,\mathbb{E}[\Phi]\leqslant O\Big((\log k)^3\Big[&\varepsilon\rho(\sigma_t,\bar{\mu}_{t-1})\\ &+D(\log(\varepsilon^{-1})+\beta(\log k)^3\log\log k)\Big]\Big).\end{aligned}$$

Let μ^Q_{t-1} and μ^Q_t be the fractional server measure before and after responding to a request and executing the Algorithm 1. Combining (T2) with the analysis of Sections V-D–V-I shows that for $\gamma\asymp(\log k)^2$ chosen sufficiently large and $\varepsilon_0>0$ chosen sufficiently small:

$$\Delta_{\mathrm{alg}}\,\mathbb{E}[\Phi]\leqslant-\frac{1}{2}\,\mathbb{E}\left[W^1_{TQ}(\mu^Q_{t-1},\mu^Q_t)\right]$$

Observing that

$$\mathbb{E}[W^1_{TQ}(\mu^Q_{t-1},\mu^Q_t)]\geqslant\frac{1}{2}\rho(\sigma_t,\bar{\mu}_{t-1}),$$

we can choose $\varepsilon=k^{-O(1)}$ sufficiently small so that

$$\begin{aligned}\mathbb{E}[\Phi_t-\Phi_{t-1}] &\leqslant O\big(\beta(\log k)^6\log\log k\big)\,W^1_X(\mathrm{opt}_{t-1},\mathrm{opt}_t)\\ &\quad-\frac{1}{4}\,\mathbb{E}\left[W^1_{TQ}(\mu^Q_{t-1},\mu^Q_t)\right]\\ &\leqslant O\big(\beta(\log k)^6\log\log k\big)\,W^1_X(\mathrm{opt}_{t-1},\mathrm{opt}_t)\\ &\quad-\Omega(\gamma^{-1})\,\mathbb{E}\left[W^1_X(\mu^Q_{t-1},\mu^Q_t)\right]\\ &\leqslant O\big(\beta(\log k)^6\log\log k\big)\,W^1_X(\mathrm{opt}_{t-1},\mathrm{opt}_t)\\ &\quad-\Omega(\gamma^{-1})W^1_X(\bar{\mu}_{t-1},\bar{\mu}_t).\end{aligned}$$

This yields the existence of an $O(\gamma\beta(\log k)^6\log\log k)=O(\beta(\log k)^8\log\log k)$-competitive (fractional) k-server algorithm on (X,d_X).

Acknowledgements

I am grateful to my coauthors [BCL+18] for stimulating preliminary discussions, and to Christian Coester for many invaluable comments and interactions regarding preliminary drafts. The author is supported by NSF grants CCF-1616297 and CCF-1407779, and a Simons Investigator Award.

References

[Bar96] Yair Bartal. Probabilistic approximations of metric spaces and its algorithmic applications. In *37th Annual Symposium on Foundations of Computer Science, FOCS '96, Burlington, Vermont, USA, 14-16 October, 1996*, pages 184–193, 1996. 3

[Bar98] Yair Bartal. On approximating arbitrary metrices by tree metrics. In *STOC '98 (Dallas, TX)*, pages 161–168. ACM, New York, 1998. 3

[BBK99] Avrim Blum, Carl Burch, and Adam Kalai. Finely-competitive paging. In *40th Annual Symposium on Foundations of Computer Science (New York, 1999)*, pages 450–457. IEEE Computer Soc., Los Alamitos, CA, 1999. 2

[BBM06] 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. 2

[BBMN15] Nikhil Bansal, Niv Buchbinder, Aleksander Madry, and Joseph Naor. A polylogarithmic-competitive algorithm for the k-server problem. *J. ACM*, 62(5):Art. 40, 49, 2015. 3, 4

[BBN10] Nikhil Bansal, Niv Buchbinder, and Joseph Naor. Metrical task systems and the k-server problem on HSTs. In *Automata, languages and programming. Part I*, volume 6198 of *Lecture Notes in Comput. Sci.*, pages 287–298. Springer, Berlin, 2010. 1

[BBN12] Nikhil Bansal, Niv Buchbinder, and Joseph Naor. A primal-dual randomized algorithm for weighted paging. *J. ACM*, 59(4):Art. 19, 24, 2012. 3

[BCL+18] Sébastien Bubeck, Michael B. Cohen, Yin Tat Lee, James R. Lee, and Aleksander Madry. k-server via multiscale entropic regularization. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 3–16, 2018. 3, 4, 5, 11

[BE98] Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, New York, 1998. 1, 2

[BLMN05] Yair Bartal, Nathan Linial, Manor Mendel, and Assaf Naor. On metric Ramsey-type phenomena. *Ann. of Math. (2)*, 162(2):643–709, 2005. 2

[CCG+98] Moses Charikar, Chandra Chekuri, Ashish Goel, Sudipto Guha, and Serge A. Plotkin. Approximating a finite metric by a small number of tree metrics. In *39th Annual Symposium on Foundations of Computer Science, FOCS '98, November 8-11, 1998, Palo Alto, California, USA*, pages 379–388, 1998. 7

[CKPV91] M. Chrobak, H. Karloff, T. Payne, and S. Vishwanathan. New results on server problems. *SIAM J. Discrete Math.*, 4(2):172–181, 1991. 2

[FKL+91] A. Fiat, R. M. Karp, M. Luby, L. A. McGeoch, D. D. Sleator, and Young N. E. Competitive paging algorithms. *J. Algorithms*, 12(4):685–699, 1991. 2

[FRR94] Amos Fiat, Yuval Rabani, and Yiftach Ravid. Competitive k-server algorithms. *J. Comput. System Sci.*, 48(3):410–428, 1994. 31st Annual Symposium on Foundations of Computer Science (FOCS) (St. Louis, MO, 1990). 2

[FRT04] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.*, 69(3):485–497, 2004. 3

[Kou09] Elias Koutsoupias. The k-server problem. *Computer Science Review*, 3(2):105–118, 2009. 2

[KP95] Elias Koutsoupias and Christos H. Papadimitriou. On the k-server conjecture. *J. Assoc. Comput. Mach.*, 42(5):971–983, 1995. 2

[KRR94] Howard Karloff, Yuval Rabani, and Yiftach Ravid. Lower bounds for randomized k-server and motion-planning algorithms. *SIAM J. Comput.*, 23(2):293–312, 1994. 2

[MMS90] Mark S. Manasse, Lyle A. McGeoch, and Daniel D. Sleator. Competitive algorithms for server problems. *J. Algorithms*, 11(2):208–230, 1990. 1, 2

[MS91] Lyle A. McGeoch and Daniel D. Sleator. A strongly competitive randomized paging algorithm. *Algorithmica*, 6(6):816–825, 1991. 2