

Fusible HSTs and the randomized k -server conjecture

James R. Lee
University of Washington

Warning: There is a gap in Section 5.3.2 of this draft. Please [look here](#) for up to date information.

Abstract

We exhibit an $O((\log k)^6)$ -competitive randomized algorithm for the k -server problem on *any* metric space. It is shown that a potential-based algorithm for the fractional k -server problem on hierarchically separated trees (HSTs) with competitive ratio $f(k)$ can be used to obtain a randomized algorithm for any metric space with competitive ratio $f(k)^2 O((\log k)^2)$. Employing the $O((\log k)^2)$ -competitive algorithm for HSTs from our joint work with Bubeck, Cohen, Lee, and Mądry (2017) yields the claimed bound.

1	Introduction	2
1.1	HST embeddings	5
1.2	Cluster fusion	7
1.3	Embeddings, isoperimetry, and scales	8
1.4	Preliminaries	11
2	Fusible HSTs	12
2.1	Universal HSTs	12
2.2	Stochastic HST embeddings	14
2.3	The potential axioms	15
3	Construction of the embedding	16
3.1	Embedding components	16
3.2	The online algorithm	20
4	Distortion analysis	21
4.1	Active scales	22
4.2	The expected stretch	23
5	The competitive ratio	24
5.1	The HST potential	24
5.2	The accuracy potential	26
5.3	The fission potential	30
5.4	Competitive analysis	34
5.5	Rounding under fusion	35
6	Reductions	37
6.1	Mass at internal nodes	38
6.2	Extra server mass	39
6.3	Verification of the potential axioms for [BCL ⁺ 17]	40

1 Introduction

An *online algorithm* is one that receives a sequence of inputs $\langle x_1, x_2, \dots \rangle$ at discrete times $t \in \{1, 2, \dots\}$. At every time step t , the algorithm takes some feasible action based only on the inputs $\langle x_1, x_2, \dots, x_t \rangle$ 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].

The k -server problem. 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 $k \geq 1$ 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 $\langle \sigma_t \in X : t \geq 0 \rangle$ 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 $\sigma_t \in X$. In other words, a feasible state ρ_t is one that *services* σ_t :

$$\sigma_t \in \{(\rho_t)_1, \dots, (\rho_t)_k\}.$$

Formally, an *online algorithm* is a sequence of mappings $\rho = \langle \rho_1, \rho_2, \dots \rangle$ where, for every $t \geq 1$, $\rho_t : X^t \rightarrow X^k$ maps a request sequence $\langle \sigma_1, \dots, \sigma_t \rangle$ 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 ρ in servicing $\sigma = \langle \sigma_t : t \geq 1 \rangle$* is defined as the sum of the movements of all the servers:

$$\text{cost}_\rho(\sigma; k, \rho_0) := \sum_{t \geq 1} d_{X^k}(\rho_t(\sigma_1, \dots, \sigma_t), \rho_{t-1}(\sigma_1, \dots, \sigma_{t-1})), \quad (1.1)$$

where

$$d_{X^k}((x_1, \dots, x_k), (y_1, \dots, y_k)) := \sum_{i=1}^k d_X(x_i, y_i) \quad \forall x_1, \dots, x_k, y_1, \dots, y_k \in X.$$

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

$$\text{cost}^*(\sigma; k, \rho_0) := \inf_{\langle \rho_1, \rho_2, \dots \rangle} \sum_{t \geq 1} d_{X^k}(\rho_t, \rho_{t-1}),$$

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

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

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

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 :

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

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 $D_k(X, d_X)$ denote the infimum of competitive ratios achievable by deterministic online algorithms, and let $R_k(X, d_X)$ 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 (1.1), we sum over all times $t \geq 1$. This is simply to avoid the notational clutter caused by an upper time horizon. One can replace a finite sequence $\langle \sigma_1, \sigma_2, \dots, \sigma_t \rangle$ of requests by the infinite sequence $\langle \sigma_1, \sigma_2, \dots, \sigma_t, \sigma_t, \sigma_t, \dots \rangle$, where the final request is repeated.

The authors of [MMS90] showed that if (X, d_X) is an arbitrary metric space and $|X| > k$, then $D_k(X) \geq k$. They conjectured that this is tight.

Conjecture 1.1 (*k-server conjecture, [MMS90]*). *For every metric space X with $|X| > k \geq 1$, it holds that*

$$D_k(X) = k.$$

Fiat, Rabani, and Ravid [FRR94] were the first to show that $D_k(X) < \infty$ for every metric space; they gave the explicit bound $D_k(X) \leq k^{O(k)}$. While Conjecture 1.1 is still open, it is now known to be true within a factor of 2.

Theorem 1.2 (Koutsoupias-Papadimitriou, [KP95]). *For every metric space X and $k \geq 1$, it holds that*

$$D_k(X) \leq 2k - 1.$$

Paging and randomization. Let \mathcal{U}_n denote the metric space on $\{1, 2, \dots, n\}$ 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 A , define the request sequence that at time $t \geq 1$ makes a request at the unique site in \mathcal{U}_{k+1} at which A does *not* have a server.

Clearly 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 $\{1, 2, \dots, k+1\}$ 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 shown that allowing an online algorithm to make random choices helps dramatically in general.

Theorem 1.3 ([FKL⁺91, MS91]). *For every $n > k \geq 1$:*

$$R_k(\mathcal{U}_n) = 1 + \frac{1}{2} + \dots + \frac{1}{k}.$$

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 1.4. *For any metric space X and $k \geq 2$ such that $|X| > k$, it holds that*

$$R_k(X) \geq \Omega\left(\frac{\log k}{\log \log k}\right).$$

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

Conjecture 1.5 (Randomized k -server conjecture). *For every metric space X and $k \geq 2$:*

$$R_k(X) \leq O(\log k).$$

The possibility that $R_k(X) \leq (\log k)^{O(1)}$ 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 1.6 (Main theorem). *For every metric space X and $k \geq 2$:*

$$R_k(X) \leq O\left((\log k)^6\right).$$

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

Theorem 1.6 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 $w : X \rightarrow \mathbb{R}_+$. Define the distance $d_w(x, y) := \max\{w(x), w(y)\}$. We refer to this as a *weighted star metric*.

Theorem 1.7 (Bansal-Buchbinder-Naor, [BBN12]). *If X is a weighted star metric and $k \geq 2$, then*

$$R_k(X) \leq O(\log k).$$

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

Theorem 1.8 (Bansal-Buchbinder-Mądry-Naor, [BBMN15]). *For every $k \geq 2$ and finite metric space X , it holds that*

$$R_k(X) \leq (\log |X|)^{O(1)}.$$

Finally, in joint work with Bubeck, Cohen, Lee, and Mądry [BCL⁺17], we obtain a cardinality-independent bound when X is an ultrametric. This last result will form an essential component of our arguments.

Theorem 1.9 ([BCL⁺17]). *For every $k \geq 2$ and every ultrametric space X , it holds that*

$$R_k(X) \leq O\left((\log k)^2\right).$$

1.1 HST embeddings

The significance of ultrametrics in [Theorem 1.9](#) stems from their pivotal role in online algorithms for k -server. Consider a rooted tree $\mathcal{T} = (V, E)$ equipped with non-negative vertex weights $\{w_u \geq 0 : u \in V\}$ such that the weights are non-increasing along every root-leaf path. Let $\mathcal{L} \subseteq V$ denote the set of leaves of \mathcal{T} , and define an ultrametric on \mathcal{L} by

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

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

If it holds for some $\tau \geq 1$ that $w_v \leq w_u/\tau$ whenever v is a child of u , then (\mathcal{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 1.10. *Suppose that (X, d) is a finite metric space. Then for every $k \geq 2$:*

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

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

Clearly in conjunction with [Theorem 1.9](#), this yields $\mathbb{R}_k(X) \leq O((\log k)^2 \log |X|)$ for any finite metric space X . The reduction from general finite metric spaces to ultrametrics implicit in [Theorem 1.10](#) 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⁺17](#)], we showed how a *dynamic* embedding of a metric space into ultrametrics could overcome the distortion barrier.

Theorem 1.11 ([[BCL⁺17](#)]). *For every $k \geq 2$ and every finite metric space (X, d) :*

$$\mathbb{R}_k(X, d) \leq O\left((\log k)^3 \log(1 + \mathcal{A}_X)\right), \tag{1.2}$$

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. Indeed, suppose that $\{G_n\}$ is a family of expander graphs with uniformly bounded degrees and such that G_n has n vertices. Let (V_n, d_n) denote the induced shortest-path metric on the vertices of G_n . It is well-known that a probabilistic embedding into ultrametrics incurs distortion $\Omega(\log n)$, while (1.2) yields $\mathbb{R}_k(V_n, d_n) \leq O((\log k)^3 \log \log n)$.

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 $\alpha : X \rightarrow \mathcal{L}$, a k -server configuration ρ_t^α corresponding to an online algorithm for the request sequence mapped into \mathcal{L} via α .

Define the *annealed server measure* \bar{v}_t to be the measure on X that results from averaging the configurations $\alpha^{-1}(\rho_t^\alpha)$ over \mathcal{D}_t . Now one would like to update $\mathcal{D}_t \mapsto \mathcal{D}_{t+1}$ based on the measure \bar{v}_t . Ideally, the measure \bar{v}_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, as we will now indicate, even if we are allowed to see the entire request sequence in advance, there is no embedding $\alpha : X \rightarrow \mathcal{L}$ that can avoid distorting distances 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, α) , where ρ is a k -server configuration and $\alpha : X \rightarrow \mathcal{L}$ is an embedding. Now let \bar{v}_t denote $\alpha^{-1}(\rho)$ averaged over $\hat{\mathcal{D}}_t$.

The distribution $\hat{\mathcal{D}}_{t+1}$ is then sampled by a two-step process: $(\rho, \alpha) \mapsto (\hat{\rho}, \alpha) \mapsto (\hat{\rho}, \hat{\alpha})$. 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 α so that it more accurately approximates X according to the annealed server measure \bar{v}_t . A key property of the transformation $\alpha \mapsto \hat{\alpha}$ is that it should not induce any movement when the configuration is pulled back to X , i.e., $\alpha^{-1}(\hat{\rho}) = \hat{\alpha}^{-1}(\hat{\rho})$.

The limitations of dynamic HST embeddings. Consider a bounded metric space (X, d_X) (i.e., one with finite diameter) and a fixed (possibly infinite) τ -HST metric space $(\mathcal{L}, d_{\text{hst}})$ with $\tau \geq 2$. We may consider a fixed HST because one can choose a universal target space without loss of generality; see [Section 2.1](#). We will assume that every leaf has a unique preimage, i.e., there is a surjection $\beta : \mathcal{L} \rightarrow X$, and that the map β is 1-Lipschitz:

$$d_X(\beta(x), \beta(y)) \leq d_{\text{hst}}(x, y) \quad \forall x, y \in \mathcal{L}. \quad (1.3)$$

Given a request sequence $\sigma = \langle \sigma_1, \sigma_2, \dots \rangle$ in X , one can consider a random sequence $\alpha = \langle \alpha_1, \alpha_2, \dots \rangle$ of points in \mathcal{L} with the property that $\beta(\alpha_t) = \sigma_t$ for each $t \geq 1$. Say that α is *oblivious* if there is a single random map $F : X \rightarrow \mathcal{L}$ chosen independently of σ and $\alpha_t := F(\sigma_t)$. Say that α is *adapted to the request sequence* if α_t depends only on $\langle \sigma_1, \sigma_2, \dots, \sigma_t \rangle$ (and possibly some additional independent randomness).

Finally, say that α has *k -server distortion at most D* if there is a constant $c > 0$ such that for every request sequence σ :

$$\mathbb{E} \left[\text{cost}_{d_{\text{hst}}}^*(\alpha; k) \right] \leq D \cdot \text{cost}_X^*(\sigma; k) + c. \quad (1.4)$$

If α is adapted to the request sequence and has k -server distortion at most D , then a C -competitive k -server algorithm on $(\mathcal{L}, d_{\text{hst}})$ yields a CD -competitive algorithm for the k -server problem on (X, d_X) since (1.3) allows us to pull the server trajectories back to X at no additional cost.

In [\[BCL⁺17\]](#), it is shown that such adapted sequences α with k -server distortion $D \leq O(\log(k) \log(1 + \mathcal{A}_X))$. Unfortunately, this model is too weak to obtain [Theorem 1.6](#) even when X is the unit circle (or the real line), even for the case of $k = 1$ server. This is for a simple reason: Even if we don't require the sequence α to be adapted (i.e., we are given the entire request sequence in advance), there are request sequences σ so that if (1.4) holds, then $D \geq \Omega(\log \mathcal{A}_X)$.

Lemma 1.12. *For every $\mathcal{A} \geq 2$, there is a set of points X on the unit circle with $\mathcal{A}_X \leq \mathcal{A}$ and so that for any α satisfying (1.4) for every σ with $k = 1$, it holds that*

$$D \geq \Omega(\log_\tau \mathcal{A}) \geq \Omega(\log_\tau |X|).$$




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

We sketch the straightforward proof, as it will motivate our modification of the dynamic embedding model and its subsequent analysis. Fix some $n \geq 2$ and consider a request sequence $\sigma = \langle \sigma_1, \sigma_2, \dots, \sigma_n \rangle$, where $\sigma_t = e^{-2\pi i \frac{t}{n}} \in S^1$, and S^1 denotes the unit circle in the complex plane equipped with its radial metric d_{S^1} . (In other words, the requests come consecutively at n equally placed points on a unit circle.)

Clearly $\text{cost}_{S^1}^*(\sigma; 1) \leq O(1)$. We claim that for any sequence of leaves $\langle \alpha_1, \alpha_2, \dots, \alpha_n \rangle$ satisfying

$$d_{\text{hst}}(\alpha_i, \alpha_j) \geq d_{S^1}(\sigma_i, \sigma_j) \quad \forall i, j, \quad (1.5)$$

it holds that

$$\sum_{t=1}^{n-1} d_{\text{hst}}(\alpha_t, \alpha_{t+1}) \geq \Omega(\log n).$$

Indeed, this is immediate: For every $1 \leq j \leq \lfloor \log_\tau n \rfloor$, by (1.5), the sequence $\langle \alpha_1, \alpha_2, \dots, \alpha_n \rangle$ of leaves must exit a subtree of diameter at least τ^{-j} at least $\Omega(\tau^j)$ times, implying that

$$\sum_{t=1}^{n-1} d_{\text{hst}}(\alpha_t, \alpha_{t+1}) \geq \sum_{j=1}^{\lfloor \log_\tau n \rfloor} \tau^{-j} \tau^j \geq \Omega(\log_\tau n).$$

1.2 Cluster fusion

Consider again the example of the preceding section, but now it will be helpful to think about a continuous path: Suppose that $\sigma : [0, \infty) \rightarrow S^1$ is a point that moves clockwise at unit speed. Recall that $(\mathcal{L}, d_{\text{hst}})$ is a τ -HST metric.

A non-contractive embedding $\alpha : S^1 \rightarrow \mathcal{L}$ induces a sequence of partitions $\{P_j : j \geq 0\}$ of S^1 , where $P_0 = S^1$, for every $j \geq 0$, P_{j+1} is a refinement of P_j , and where every set $S \in P_j$ has diameter at most $2\pi\tau^{-j}$. When $\sigma(t)$ approaches the boundary of P_j , the image $\alpha(\sigma(t))$ stands to incur d_{hst} movement $\approx \tau^{-j}$ as $\alpha(\sigma(t))$ switches sets of the partition P_j . In order to prevent this, we will *fuse* together the two sets of P_j whose boundary $\sigma(t)$ is about to cross. See Figure 1.

When $\sigma(t)$ is safely past the boundary, we need to *unfuse* these sets so that we are prepared to fuse across the next P_j boundary. Failing to do this, we might start fusing a long chain of sets; having sets of unbounded diameter in P_j would prevent us from maintaining a non-contractive embedding into \mathcal{L} . We will soon describe a model that supports fusion and fission of sets in the target HST.

Potential-based algorithms for HSTs. Once we allow ourselves such operations, it no longer seems possible to use a competitive HST algorithm as a black box. Indeed, such an algorithm maintains internal state, and there is no reason it should continue to operate meaningfully under a sudden unexpected change to this state (resulting from the fusion of clusters).

Thus we will assume the existence of an HST algorithm that maintains a configuration χ and whose operation can be described as a function $(\chi, \sigma) \mapsto \chi'$ that maps a pair (χ, σ) to a new configuration χ' , where $\sigma \in \mathcal{L}$ is the request to be serviced, and χ' induces a fractional k -server measure $\mu^{\chi'}$ that services σ . (See [Section 2](#) for a discussion of fractional k -server measures; for the present discussion, one can think of $\mu^{\chi'}$ as simply a k -server configuration.) Moreover, we will assume that the HST algorithm's competitiveness is witnessed by a potential function $\Phi(\theta^*; \chi)$ that tracks the “discrepancy” between the server state induced by χ and the server state θ^* of the optimal offline algorithm.

Crucially, we will assume that Φ decreases monotonically under fusion operations applied simultaneously to both θ^* and (the measure underlying) χ . If Φ is thought of as a measure of discrepancy with respect to the underlying HST, then this makes sense: When two clusters are fused, the corresponding notion of discrepancy becomes more coarse (meaning that it is less able to distinguish θ^* from μ^χ).

We also need to assume that Φ is relatively stable under operations that correspond to fission of clusters. We state the required properties formally in [Section 2.3](#). In [Section 6.3](#), we confirm that the algorithm establishing [Theorem 1.9](#) satisfies these properties.

1.3 Embeddings, isoperimetry, and scales

Suppose now that (X, d_X) has diameter at most one. Let $\mathcal{P} = \langle P_j : j \geq 0 \rangle$ denote a sequence of partitions of X so that for each $j \geq 0$, if $S \in P_j$ then $\text{diam}_X(S) \leq \tau^{-j}$. For a partition P of X and $x \in X$, let $P(x)$ denote the unique set of P containing x .

One can define a τ -HST metric on X by

$$d_{\text{hst}}^{\mathcal{P}}(x, y) := \tau^{-\min\{j \geq 0 : P_j(x) \neq P_j(y)\}}.$$

If X is finite, then by choosing the partitions P_j appropriately at random, one can additionally obtain the property that

$$\mathbb{P}[P_j(x) \neq P_j(y)] \leq \frac{d_X(x, y)}{\tau^j} O(\log |X|) \quad \forall j \geq 0. \quad (1.6)$$

Such random partitions are now ubiquitous in many areas; see, for instance, [\[FRT04, KLMN05, LN05\]](#) for applications in algorithms and metric embedding theory.

In particular, summing over the values of $j \geq 0$ such that $\tau^{-j} \geq \min_{x \neq y \in X} d_X(x, y)$, [\(1.6\)](#) implies that for any $x, y \in X$,

$$\mathbb{E}[d_{\text{hst}}^{\mathcal{P}}(x, y)] \leq O(\log |X|) \cdot O(\log \mathcal{A}_X), \quad (1.7)$$

where we recall the *aspect ratio* of X from [\(1.2\)](#).

Both distortion factors in [\(1.7\)](#) are troublesome, but there is now a well-understood theory of how they arise. See, for instance, the elegant argument of [\[FRT04\]](#) which indicates that they cannot arise simultaneously: If the random partitions $\{P_j : j \geq 0\}$ are chosen carefully, then one can achieve the bound $\mathbb{E}[d_{\text{hst}}^{\mathcal{P}}(x, y)] \leq O(\log |X|) d_X(x, y)$.

The $O(\log |X|)$ factor inherited from [\(1.6\)](#) might be called the “isoperimetric” obstruction. For instance, it can be replaced by a universal constant if $X = \mathbb{R}$, but it is necessary if (X, d_X) is the

shortest-path metric on an expander graph or the ℓ_1 metric on $\{0, 1\}^d$ for some $d \geq 1$. The $O(\log \mathcal{A}_X)$ factor could be called the “multiscale” obstruction, and it arises whenever the underlying metric space contains paths, i.e., sequences $x_1, x_2, \dots, x_n \in X$ along which the triangle inequality is approximately tight:

$$d_X(x_1, x_n) \approx d_X(x_1, x_2) + d_X(x_2, x_3) + \dots + d_X(x_{n-1}, x_n).$$

1.3.1 The isoperimetric obstruction

If, instead of choosing a static embedding, we imagine maintaining an embedding for the purposes of solving the k -server problem, then it is not unreasonable to expect that the HST embedding only needs to “track” $O(k)$ regions at every scale.

Indeed, consider a sequence $\sigma = \langle \sigma_1, \sigma_2, \dots, \sigma_N \rangle$ of requests in X . Let $S \subseteq \{\sigma_1, \sigma_2, \dots, \sigma_N\}$ denote a τ^{-j} -separated subset of the requests (so that $d_X(\sigma, \sigma') \geq \tau^{-j}$ for $\sigma \neq \sigma' \in S$). If $|S| \geq k + h$, then *any* sequence $\rho_1, \rho_2, \dots, \rho_N \in X^k$ of k -server configurations that services σ must incur total movement at least $h\tau^{-j}$.

Thus when the request sequence is sufficiently spread out at scale τ^{-j} , the optimum offline algorithm must be incurring proportional cost. This allows us to track only $O(k)$ regions, and pay some cost whenever we have to alter the embedding to incorporate new regions (and discard old ones); that “edit cost” can be charged to the movement cost of the optimum. This suggests one might replace $O(\log |X|)$ by $O(\log k)$ in (1.7) and, indeed, this is the content of [Theorem 1.11](#).

Unfortunately, this argument overcharges the optimum cost by a factor proportional to $\log \mathcal{A}_X$, since it is not possible to naively perform the same charging argument for all scales simultaneously. If one considers $X = \mathbb{R}$ equipped with its usual metric, then even in the case of $k = 1$, the request sequence $\{0, 1, 2, \dots, N\}$ would incur $\asymp N$ charge at each of the $\asymp \log N$ scales even though the optimum only moves distance N .

To address this, we employ a sophisticated dynamic embedding and a charging scheme that tracks the relationship between the movement at various scales. This is encapsulated in the “accuracy potential” of [Section 5.2](#).

1.3.2 The multiscale obstruction

It is the more daunting multiscale obstruction that motivates a model in which we can fuse together sibling clusters in the HST embedding.

Our underlying idea is simple: Suppose that \bar{v}_t is the annealed server measure described earlier. Consider a ball B in X and the ball λB (with the same center, and with a λ times larger radius), where λ is a large constant. If it holds that for some small $\delta > 0$,

$$\bar{v}_t(B) \geq (1 - \delta)\bar{v}_t(\lambda B), \tag{1.8}$$

let us say that the ball B is *heavy* (with respect to \bar{v}_t). If B is heavy, it indicates that we would prefer a random partition to “cut around B ” in the light annulus $\lambda B \setminus B$.

We will enforce this by “fusing” together all the sets in P_j that intersect B into one supercluster; see [Figure 2\(a\)](#). The condition (1.8) directly implies that disjoint heavy balls must be far apart, and thus for λ chosen large enough, we avoid the problem of having chains of fusions that produce sets of unbounded diameter.

This also addresses the multiscale obstruction: At every scale τ^{-j} where the ball $B_X(x, \tau^{-j})$ is heavy, we fuse the clusters near x , and therefore do not pay the separation penalty in (1.6). At how

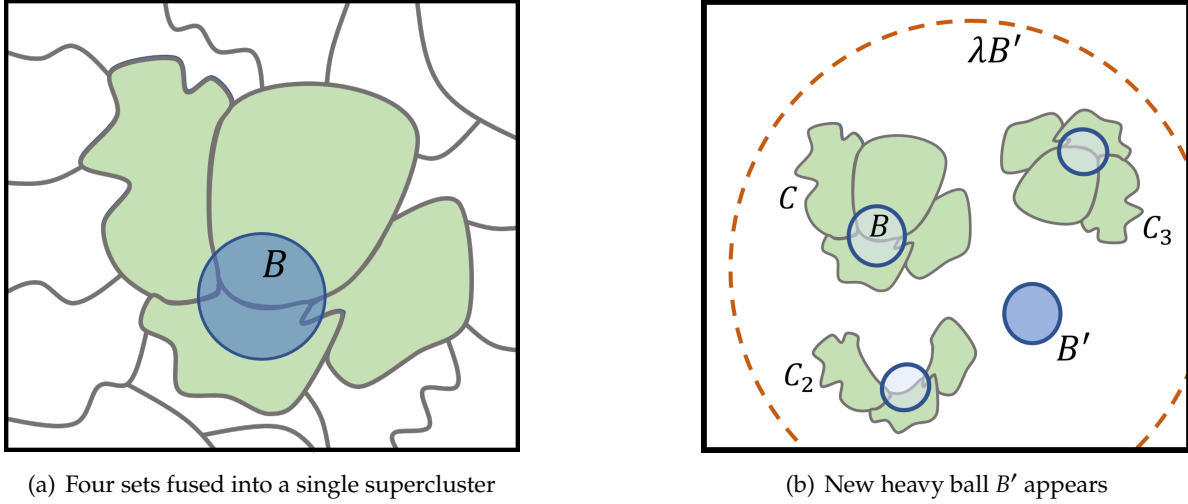


Figure 2: Fusion along heavy balls

many scales $j \in \{0, 1, 2, \dots\}$ can the ball $B_X(x, \tau^{-j})$ be light? It is easy to see that the answer is $O\left(\frac{1}{\delta} \log \frac{\bar{v}_t(X)}{\bar{v}_t(x)}\right)$ since, at every light scale, a δ -fraction of the mass is lost when zooming into x from radius $\lambda\tau^{-j}$ to radius τ^{-j} . We have $\bar{v}_t(X) = k$, and when x has been the site of a recent request, it will hold roughly that $\bar{v}_t(x) \geq 1/2$. Thus the number of non-trivial scales at which x is not fused with its neighbor clusters is only $O\left(\frac{1}{\delta} \log k\right)$. When combined with our solution to the isoperimetric obstruction, this leads to a bound of $O\left(\frac{1}{\delta} (\log k)^2\right)$ in (1.7).

Paying for cluster fission. As mentioned previously, the difficulty comes when a ball that was once heavy becomes light, and then we must “unfuse” the underlying clusters. This fission cost will be charged against the transportation cost of the sequence of measures $\langle \bar{v}_t : t \geq 0 \rangle$. We only unfuse the clusters corresponding to a heavy ball B when eventually some ball B' with $\text{diam}_X(B) \asymp \text{diam}_X(B')$ becomes heavy and satisfies $\sqrt{\lambda}B \subseteq \lambda B' \setminus B'$. See Figure 2(b). It is intuitively clear that this requires significant movement of the measure \bar{v}_t on which heaviness is based. We will employ two properties in order to charge the cost of fission against this movement:

1. If B is heavy with respect to \bar{v}_{t_0} and B' is heavy with respect to \bar{v}_{t_1} for $t_1 > t_0$, we can charge this against

$$\text{diam}_X(B') \bar{v}_t(B') \tag{1.9}$$

transportation cost incurred by $\langle \bar{v}_t : t \in [t_0, t_1] \rangle$ in the creation of the heavy ball B' .

2. When B' becomes heavy with respect to \bar{v}_t , we will need to unfuse any previous heavy balls B_1, B_2, \dots, B_m satisfying $\sqrt{\lambda}B_i \subseteq \lambda B' \setminus B'$. If C_i is the supercluster that was formed when B_i became heavy at some earlier time, then for λ chosen large enough,

$$C_1 \cup \dots \cup C_m \subseteq \sqrt{\lambda}B_1 \cup \dots \cup \sqrt{\lambda}B_m \subseteq \lambda B' \setminus B',$$

and therefore it holds that

$$\bar{v}_t(C_1) + \dots + \bar{v}_t(C_m) \leq \delta \bar{v}_t(B'). \tag{1.10}$$

We will assume that fission of a supercluster C only “costs” us $f_4(k) \text{diam}_X(C) \bar{v}_t(C)$ for some function $f_4 : \mathbb{N} \rightarrow [1, \infty)$ satisfying $f_4(k) \leq (\log k)^{O(1)}$. (See Axiom (A4) in [Section 2.3](#).)

In this case, (1.10) implies that the total cost of fission is at most

$$\delta \bar{v}_t(B') \max_i \{\text{diam}_X(C_i)\} \leq O(\delta) \bar{v}_t(B') \text{diam}_X(B'). \quad (1.11)$$

Thus by choosing $\delta > 0$ small enough, we can ensure that the cost of the fission in (2) is paid for by the transportation cost incurred in (1). The formal charging argument occurs using the “fission potential” introduced in [Section 5](#).

One should note that, unlike in [Section 1.3.1](#), where the geometry of the request sequence allows us to charge against the transportation cost of the optimal offline algorithm, here we only charge against the transportation cost of $\langle \bar{v}_t : t \geq 0 \rangle$ (which is essentially the movement cost our online algorithm has incurred). Thus it is essential that we are allowed to take $\delta > 0$ in (1.11) to be small (in fact, we will take $\delta \asymp 1/f_4(k)$).

The main theorem and algorithmic considerations. The methods outlined in the preceding sections allow us to obtain the following.

Theorem 1.13. *There is a constant $C > 1$ such that for every $k \geq 2$, the following holds. On every metric space (X, d_X) , there is a $C(\log k)^6$ -competitive randomized algorithm for the k -server problem on X .*

We remark that if we are allowed to solve a *fractional relaxation* of the k -server problem (see [Section 2](#) for a discussion of fractional k -server measures), then the algorithm described here can be implemented so that it responds to a request in time *polynomial in k* . (Here, we treat specification of a request $\sigma \in X$ and computation of a distance $d_X(x, y)$ as unit cost operations.)

The reason is simple: At any point in time, our algorithm only maintains a distribution over HST embeddings of $k^{O(1)}$ points in X . It is not hard to see that the distribution need only be supported on $k^{O(1)}$ different HSTs (as in, e.g., [CCG⁺98]). Moreover, the HST algorithm of [BCL⁺17] (to which we eventually appeal) performs a fractional update in $k^{O(1)}$ time on a $k^{O(1)}$ -vertex HSTs.

On the other hand, rounding a sequence of fractional server measures online (cf. [Theorem 2.5](#)) to a random integral measure currently requires time $k^{O(k)}$ per request.

1.4 Preliminaries

Let us write $\mathbb{R}_+ := [0, \infty)$ and $\mathbb{Z}_+ := \mathbb{Z} \cap \mathbb{R}_+$. Consider a set X . We use $\mathbb{M}(X)$ to denote the space of measures on X whose support is at most countable. Denote by $\mathbb{M}_k(X) \subseteq \mathbb{M}(X)$ the subset of countably-additive measures $\mu \in \mathbb{M}(X)$ that satisfy $\mu(X) = k$. Since our measures have at most countable support, when $x \in X$, we will often write $\mu(x)$ for $\mu(\{x\})$. For $\mu \in \mathbb{M}(X)$, define $\text{supp}(\mu) := \{x \in X : \mu(x) > 0\}$. 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 X, Y are two spaces, $F : X \rightarrow Y$, and $\mu \in \mathbb{M}(X)$, then we use $F\#\mu$ to denote the *pushforward* measure:

$$F\#\mu(S) := \mu(F^{-1}(S)) \quad \forall S \subseteq Y.$$

Note that if μ is integral, then so is $F\#\mu$. If $\boldsymbol{\mu} = \langle \mu_1, \mu_2, \dots \rangle$ is a sequence of measures, we define $F\#\boldsymbol{\mu} := \langle F\#\mu_1, F\#\mu_2, \dots \rangle$.

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

$$\text{cost}_X(\boldsymbol{\mu}) := \sum_{t \geq 1} W_X^1(\mu_t, \mu_{t+1}),$$

where $W_X^1(\mu, \nu)$ is the L^1 -transportation distance between μ and ν in X . This is sometimes referred to as the Wasserstein 1-distance or the Earthmover distance. One can consult the book [Vi103] for an introduction to the geometry of optimal transportation. Note that we only deal here with countably-supported measures, so our considerations are elementary. (The reader should also note that with slightly more notational overhead, one could assume that all encountered measures have finite support.)

The following claim is straightforward. If (X, d_X) and (Y, d_Y) are two metric spaces and $F : X \rightarrow Y$, one defines

$$\|F\|_{\text{Lip}} := \sup_{x \neq y \in X} \frac{d_Y(F(x), F(y))}{d_X(x, y)}.$$

Claim 1.14. For any sequence μ , it holds that

$$\text{cost}_Y(F\#\mu) \leq \|F\|_{\text{Lip}} \cdot \text{cost}_X(\mu).$$

For $x \in X$ and $r \geq 0$, we denote the ball $B_X(x, r) := \{y \in X : d_X(x, y) \leq r\}$ 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 $d_X(S, T) := \inf\{d_X(x, y) : x \in S, y \in T\}$.

For two non-negative expressions $E, E' \geq 0$, we write $E \leq O(E')$ to denote that there is a universal constant $C > 0$ such that $E \leq CE'$. We also write $E \asymp E'$ to denote the conjunction of $E \leq O(E')$ and $E' \leq O(E)$.

2 Fusible HSTs

Fix a metric space (X, d_X) with diameter at most one. Consider a global filtration $\mathcal{F} = \langle \mathcal{F}_1, \mathcal{F}_2, \dots \rangle$ where $\mathcal{F}_1 \subseteq \mathcal{F}_2 \subseteq \dots$, and \mathcal{F}_t represents information about the request sequence up to time t . Denote the request sequence $\sigma = \langle \sigma_1, \sigma_2, \dots \rangle$ with $\sigma_t \in X$ for all $t \geq 1$. We use $\sigma_{[s, t]}$ to denote the subsequence $\langle \sigma_s, \sigma_{s+1}, \dots, \sigma_t \rangle$. Say that a sequence $\rho = \langle \rho_0, \rho_1, \rho_2, \dots \rangle$ is \mathcal{F} -adapted if each object ρ_t is possibly a function of $\sigma_{[1, t]}$ (but *not* the future $\sigma_{t+1}, \sigma_{t+2}, \dots$).

A notable observation is that in many cases it suffices to maintain a *fractional* k -server state, as opposed to a (random) integral state; one then rounds, in an \mathcal{F} -adapted manner, the fractional solution to a random integral solution without blowing up the expected cost. This idea appears in [BBK99] and is made explicit in [BBN12] for weighted star metrics. In [BBMN15], it is extended to HST metrics. See [Theorem 2.5](#) below for a variant tailored to our setting.

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

2.1 Universal HSTs

It will be convenient for us to have a fixed HST into which our embeddings map requests. To accommodate request sequences of arbitrary length, the HST will be infinite, but the measure maintained by our algorithm will always be supported on a finite set of leaves (which are themselves a subset of the request sequence seen so far).

Fix some number $\tau \geq 2$. A sequence of subsets $\xi = \langle \xi_0, \xi_1, \xi_2, \dots \rangle$ of X is a τ -chain if

$$X = \xi_0 \supseteq \xi_1 \supseteq \xi_2 \supseteq \dots,$$

and $\text{diam}_X(\xi_j) \leq \tau^{-j}$ for all $j \geq 0$. If ξ is a finite sequence, we refer to ξ as *finite chain* and let $\text{len}(\xi)$ denote its length (otherwise set $\text{len}(\xi) := +\infty$). Define the *bottom* of ξ by $\mathbb{b}(\xi) := \bigcap_{i \geq 1} \xi_i$. Observe that for a finite τ -chain ξ ,

$$\text{diam}_X(\mathbb{b}(\xi)) \leq \tau^{-\text{len}(\xi)}. \quad (2.1)$$

A *decorated τ -chain* is a sequence $\hat{\xi} = \langle (\xi_0; 0), \hat{\xi}_1, \hat{\xi}_2, \dots \rangle$ where $\hat{\xi}_i = (\xi_i; \eta_i)$ for $i \geq 1$, $\langle \xi_0, \xi_1, \dots \rangle$ is a τ -chain, and $\langle \eta_i \in \mathbb{Z}_+ : i \geq 1 \rangle$ are arbitrary labels. We use $\text{len}(\hat{\xi})$ and $\mathbb{b}(\hat{\xi})$ to denote the corresponding quantities for the underlying undecorated chain. We denote $\eta((\xi_i, \eta_i)) := \eta_i$.

Remark 2.1 (The decorations). We note that the decorations $\{\eta_i\}$ will play a minor role in our arguments. One could take $\eta_i \in \{0, 1\}$ for all $i \geq 1$. We emphasize, in [Section 3.1.4](#) and [Section 6](#), the two places where they are used. One could do without them entirely, but they make some arguments substantially shorter.

Let $V_{\mathbb{T}}$ denote the set of finite decorated τ -chains in X . Define a rooted tree structure on \mathbb{T} as follows. The root of \mathbb{T} is the length-one chain $(X, 0)$ (with label 0). For two chains $\xi, \xi' \in V_{\mathbb{T}}$: ξ' is a child of ξ if ξ is a prefix of ξ' and $\text{len}(\xi') = \text{len}(\xi) + 1$. Let \mathbb{T} denote the rooted tree structure with vertex set $V_{\mathbb{T}}$. Let $V_{\mathbb{T}}^j \subseteq V_{\mathbb{T}}$ denote the set of τ -chains of length j . A decorated τ -chain ξ is a *leaf chain* if $\text{len}(\xi) = \infty$ and $|\mathbb{b}(\xi)| = 1$. Let $\mathcal{L}_{\mathbb{T}}$ denote the set of leaf chains. We denote the *extended vertex set* $\mathcal{V}_{\mathbb{T}} := \mathcal{L}_{\mathbb{T}} \cup V_{\mathbb{T}}$.

For two *distinct* chains $\xi, \xi' \in \mathcal{V}_{\mathbb{T}}$, define their *least common ancestor* $\text{lca}(\xi, \xi') \in V_{\mathbb{T}}$ as the maximal finite chain $(\xi_0, \xi_1, \dots, \xi_L)$ that is a prefix of both ξ and ξ' . This allows us to define a τ -HST metric on $\mathcal{V}_{\mathbb{T}}$ by

$$\text{dist}_{\mathbb{T}}(\xi, \xi') := \tau^{-\text{len}(\text{lca}(\xi, \xi'))}.$$

We call the pair $(\mathbb{T}, \text{dist}_{\mathbb{T}})$ the *universal τ -HST on (X, d_X)* . For succinctness, we will employ the notations $\text{cost}_{\mathbb{T}} := \text{cost}_{(\mathcal{V}_{\mathbb{T}}, \text{dist}_{\mathbb{T}})}$ and $W_{\mathbb{T}}^1 := W_{(\mathcal{V}_{\mathbb{T}}, \text{dist}_{\mathbb{T}})}^1$. We use $\mathcal{V}_{\mathbb{T}}^0 \subseteq \mathcal{V}_{\mathbb{T}}$ to denote the subset of chains whose decorations are identically 0 and $\mathcal{L}_{\mathbb{T}}^0 := \mathcal{L}_{\mathbb{T}} \cap \mathcal{V}_{\mathbb{T}}^0$.

Pushing measures to X . Define the map $\beta : \mathcal{L}_{\mathbb{T}} \rightarrow X$ as follows: $\beta(\xi)$ is the unique element in $\mathbb{b}(\xi)$.

Claim 2.2. β is 1-Lipschitz as a map from $(\mathcal{L}_{\mathbb{T}}, \text{dist}_{\mathbb{T}})$ to (X, d_X) .

Proof. Consider $\xi, \xi' \in \mathcal{L}_{\mathbb{T}}$. Let $\hat{\xi} := \text{lca}(\xi, \xi')$. Then by definition, $\beta(\xi), \beta(\xi') \in \mathbb{b}(\hat{\xi})$, hence

$$d_X(\beta(\xi), \beta(\xi')) \leq \text{diam}_X(\mathbb{b}(\hat{\xi})) \stackrel{(2.1)}{\leq} \tau^{-\text{len}(\hat{\xi})} = \text{dist}_{\mathbb{T}}(\xi, \xi'). \quad \square$$

If one considers a measure $\mu \in \mathbb{M}(\mathcal{L}_{\mathbb{T}})$, then the pushforward $\beta\#\mu$ gives a canonical way of transporting that measure to X .

Fusion maps and canonical injections. For $\xi \in V_{\mathbb{T}}$, define

$$\begin{aligned} \mathcal{V}_{\mathbb{T}}(\xi) &:= \{\xi' \in \mathcal{V}_{\mathbb{T}} : \xi \text{ is a prefix of } \xi'\} \\ \mathcal{L}_{\mathbb{T}}(\xi) &:= \mathcal{V}_{\mathbb{T}}(\xi) \cap \mathcal{L}_{\mathbb{T}}. \end{aligned}$$

Consider $j \geq 1$ and siblings $\xi, \xi' \in V_{\mathbb{T}}^j$ with $\mathbb{b}(\xi) \subseteq \mathbb{b}(\xi')$. Then there is a canonical mapping $\varphi_{\xi \leftrightarrow \xi'} : \mathcal{V}_{\mathbb{T}} \rightarrow \mathcal{V}_{\mathbb{T}}$ defined as follows: $\varphi_{\xi \leftrightarrow \xi'}|_{\mathcal{V}_{\mathbb{T}} \setminus \mathcal{V}_{\mathbb{T}}(\xi)}$ is the identity, and

$$\langle \xi_0, \xi_1, \dots, \xi_{j-1}, (\mathbb{b}(\xi); \eta(\xi)), \xi_{j+1}, \xi_{j+2}, \dots \rangle \in \mathcal{V}_{\mathbb{T}}(\xi)$$

is mapped to

$$\langle \xi_0, \xi_1, \dots, \xi_{j-1}, (\text{lb}(\xi'); \eta(\xi')), \xi_{j+1}, \xi_{j+2}, \dots \rangle \in \mathcal{V}_{\mathbb{T}}(\xi').$$

We refer to $\varphi_{\xi \hookrightarrow \xi'}$ as the *canonical injection of ξ into ξ'* . A map $\varphi : \mathcal{V}_{\mathbb{T}} \rightarrow \mathcal{V}_{\mathbb{T}}$ is called a *fusion map* if it is the composition of finitely many canonical injections (in particular, the identity map is a fusion map).

The importance of fusion maps is encapsulated in the following lemma. It asserts that transporting a leaf measure under a fusion map does not induce movement when the measure is pushed from $\mathcal{L}_{\mathbb{T}}$ to X . Its truth is immediate from the fact that if φ is a fusion map, then for every $\xi \in \mathcal{L}_{\mathbb{T}}$, $\beta(\varphi(\xi)) = \beta(\xi)$.

Lemma 2.3. *If $\mu \in \mathbb{M}(\mathcal{L}_{\mathbb{T}})$ and φ is a fusion map, then $\beta\#\varphi\#\mu = \beta\#\mu$.*

Remark 2.4 (Tree terminology). Despite the orientation of trees found in nature, we will sometimes informally refer to the root as at the “top” of the tree and the leaves at the “bottom.”

2.2 Stochastic HST embeddings

Let \mathbb{T} denote the universal τ -HST for (X, d_X) and some $\tau \geq 6$. A *stochastic HST embedding from X into \mathbb{T}* is a random \mathcal{F} -adapted sequence $\alpha = \langle \alpha_t : X \rightarrow \mathcal{L}_{\mathbb{T}} \mid t \geq 0 \rangle$ such that with probability one:

$$\text{lb}(\alpha_t(x)) = \{x\} \quad \forall x \in X, t \geq 0 \quad (2.2)$$

$$\alpha_t(\sigma_t) \in \mathcal{L}_{\mathbb{T}}^0 \quad \forall t \geq 1. \quad (2.3)$$

This yields a (random) request sequence $\alpha(\sigma) := \langle \alpha_1(\sigma_1), \alpha_2(\sigma_2), \dots \rangle$ in $\mathcal{L}_{\mathbb{T}}^0$. We remark that requests are restricted to map to 0-decorated leaves simply because we will use the decorations for “bookkeeping.”

The cost modulo fusion. We will consider fractional k -server algorithms μ with $\mu_t \in \mathbb{M}_k(\mathcal{L}_{\mathbb{T}})$. [Lemma 2.3](#) motivates the following notion of cost in which fusions are “free.” To that end, let us define the *reduced transportation distance*

$$W_{\mathbb{T}}^F(\mu \rightarrow \mu') := \inf \{W_{\mathbb{T}}^1(\varphi\#\mu, \mu') : \varphi \text{ a fusion map}\},$$

where the notation is meant to indicate that the “distance” is not symmetric in μ and μ' . One can think of this definition as follows: When moving from μ to μ' , without incurring movement cost, we are allowed to first apply a fusion map.

For an \mathcal{F} -adapted sequence of measures:

$$\mu = \langle \mu_t \in \mathbb{M}_k(\mathcal{L}_{\mathbb{T}}) : t \geq 0 \rangle, \quad (2.4)$$

define the *reduced cost*:

$$\text{cost}_{\mathbb{T}}^F(\mu) := \sum_{t \geq 0} W_{\mathbb{T}}^F(\mu_t \rightarrow \mu_{t+1}).$$

The next result is proved in [Section 5.5](#).

Theorem 2.5 (Online rounding under fusions). *For every sequence μ as in (2.4), there exists a random integral \mathcal{F} -adapted sequence $\hat{\mu} = \langle \hat{\mu}_t \in \widehat{\mathbb{M}}_k(X) : t \geq 0 \rangle$ such that*

$$\mu_t(\ell) \geq 1 \implies \hat{\mu}_t(\beta(\ell)) \geq 1 \quad \forall \ell \in \mathcal{L}_{\mathbb{T}}, t \geq 1,$$

and

$$\mathbb{E}[\text{cost}_X(\nu)] \leq O(1) \cdot \text{cost}_{\mathbb{T}}^F(\mu).$$

Our goal is now to construct a pair (μ, α) so that α is a stochastic HST embedding α from X into \mathbb{T} and μ is a random fractional k -server algorithm (as in (2.4)) satisfying: For every request sequence σ ,

1. μ services $\alpha(\sigma)$ with probability one, and
2. $\mathbb{E} [\text{cost}_{\mathbb{T}}^F(\mu)] \leq O((\log k)^6) \text{cost}_X^*(\sigma; k) + c$,

where $c = c(\rho_0)$ is a constant depending on the initial configuration of servers.

Combined with [Theorem 2.5](#), this yields an $O((\log k)^6)$ -competitive randomized algorithm for the k -server problem on X . In order to reach such a conclusion, we now assume the existence of an HST algorithm that satisfies a certain set of assumptions.

2.3 The potential axioms

We will assume we have a fractional k -server algorithm that operates in the following way.

- There is a configuration space Γ and a transition function $\gamma : \Gamma \times \mathcal{L}_{\mathbb{T}}^0 \rightarrow \Gamma$.
- Every configuration $\chi \in \Gamma$ has an associated fractional server measure $\mu^\chi \in \mathbb{M}_k(\mathcal{L}_{\mathbb{T}}^0)$. Upon receiving a request $\sigma \in \mathcal{L}_{\mathbb{T}}^0$, the algorithm updates its configuration to $\chi' := \gamma(\chi, \sigma)$ such that $\mu^{\chi'}(\sigma) \geq 1$.
- For some leaf $\ell_0 \in \mathcal{L}_{\mathbb{T}}$, there exists a configuration $\chi_0 \in \Gamma$ such that $\mu^{\chi_0}(\ell_0) = k$.

Moreover, there is a potential function $\Phi : \widehat{\mathbb{M}}_k(\mathcal{L}_{\mathbb{T}}) \times \Gamma \rightarrow \mathbb{R}_+$ that satisfies the following assumptions for some functions $f_1, f_4 : \mathbb{N} \rightarrow [1, \infty)$.

- (A1) **Movement of the “optimum” cannot increase the potential too much.** For any states $\theta, \theta' \in \widehat{\mathbb{M}}_k(\mathcal{L}_{\mathbb{T}})$ and configuration $\chi \in \Gamma$:

$$|\Phi(\theta; \chi) - \Phi(\theta'; \chi)| \leq f_1(k) W_{\mathbb{T}}^1(\theta, \theta').$$

In other words, Φ is $f_1(k)$ -Lipschitz in its first coordinate.

- (A2) **Movement of the algorithm decreases the potential.** For every $\sigma \in \mathcal{L}_{\mathbb{T}}$ and $\theta \in \widehat{\mathbb{M}}_k(\mathcal{L}_{\mathbb{T}})$ satisfying $\theta(\sigma) \geq 1$, the following holds. Denoting $\chi' := \gamma(\chi, \sigma)$, we have

$$\Phi(\theta; \chi') - \Phi(\theta; \chi) \leq -W_{\mathbb{T}}^1(\mu^\chi, \mu^{\chi'}).$$

- (A3) **Fusion is free.** For any fusion map φ and configuration $\chi \in \Gamma$, there is a configuration $\chi(\varphi) \in \Gamma$ such that $\mu^{\chi(\varphi)} = \varphi \# \mu^\chi$, and moreover

$$\Phi(\varphi \# \theta; \chi(\varphi)) \leq \Phi(\theta; \chi) \quad \forall \theta \in \widehat{\mathbb{M}}_k(\mathcal{L}_{\mathbb{T}}). \quad (2.5)$$

If one thinks of $\Phi(\theta; \chi)$ as the “discrepancy” between θ and μ^χ , then fusion corresponds to coarsening the discrepancy measure, which should make them appear more similar (hence the inequality in (2.5)).

(A4) **Stability under local edits.** Consider $\xi^0 \in V_{\mathbb{T}}^j$ and a child $\xi^1 \in V_{\mathbb{T}}^{j+1}$. Let $F : \mathcal{L}_{\mathbb{T}} \rightarrow \mathcal{L}_{\mathbb{T}}$ be any mapping that satisfies $F(\xi) = \xi$ for $\xi \notin \mathcal{L}_{\mathbb{T}}(\xi^1)$ and $F(\mathcal{L}_{\mathbb{T}}(\xi^1)) \subseteq \mathcal{L}_{\mathbb{T}}(\xi^0)$. Then for any $\chi \in \Gamma$ and $\theta \in \widehat{\mathbb{M}}_k(\mathcal{L}_{\mathbb{T}})$, it holds that

$$\Phi(F\#\theta; \chi) - \Phi(\theta; \chi) \leq f_4(k)\tau^{-j}\mu^\chi(\mathcal{L}_{\mathbb{T}}(\xi^1)).$$

This says that moving the θ -mass on ξ^1 arbitrarily underneath ξ^0 affects the potential by a controlled amount. Note that (1) would give $f_1(k)\tau^{-j}\theta(\mathcal{L}_{\mathbb{T}}(\xi^1))$ on the RHS since $\text{diam}_{\mathbb{T}}(\mathcal{L}_{\mathbb{T}}(\xi)) \leq \tau^{-j}$, but this control is in terms of $\mu^\chi(\mathcal{V}_{\mathbb{T}}(\xi^1))$.

The algorithm of [BCL⁺17] achieves these with $f_1(k), f_4(k) \leq O((\log k)^2)$. See [Section 6.3](#).

Theorem 2.6. *For any bounded metric space (X, d_X) , the following holds. If there is a transition function $\gamma : \Gamma \times \mathcal{L}_{\mathbb{T}}^0 \rightarrow \Gamma$ and a potential Φ satisfying Axioms (A0)–(A4) for some functions $f_1(k), f_4(k) \leq (\log k)^{O(1)}$, then there is an $O(f_1(k)f_4(k)(\log k)^2)$ -competitive randomized algorithm for the k -server problem on X .*

Corollary 2.7. *There is an $O((\log k)^6)$ -competitive randomized algorithm for the k -server problem on any bounded metric space.*

The extension to unbounded metric spaces is addressed in [Section 6.3.1](#).

3 Construction of the embedding

We will construct, inductively, a stochastic HST embedding $\alpha = \langle \alpha_t : X \rightarrow \mathcal{L}_{\mathbb{T}} \mid t \geq 0 \rangle$ and a random fractional k -server algorithm $\mu = \langle \mu_t \in \mathbb{M}_k(\mathcal{L}_{\mathbb{T}}) : t \geq 0 \rangle$ that services $\alpha(\sigma)$. Let ν^* denote an optimal offline integral k -server algorithm for σ in X . Without loss of generality, we may assume that ν^* is *lazy*: It responds to requests by moving at most one server per time step.

Theorem 3.1. *There is a constant $B \geq 1$ such that under the assumptions of [Theorem 2.6](#), there is a pair (μ, α) so that for every initial configuration ν_0^* and request sequence σ :*

$$\mathbb{E} [\text{cost}_{\mathbb{T}}^F(\mu)] \leq B f_1(k) f_4(k) (\log k)^2 \text{cost}_X(\nu^*) + 4\Phi(\nu_0^*; \chi_0),$$

where χ_0 is the initial configuration guaranteed in [Section 2.3](#) and $\mu_0 = \mu^{\chi_0}$.

For later use, define $\mu^* = \langle \mu_t : t \geq 0 \rangle$ as the pushforward of the offline optimum under the embedding: $\mu_t^* := \alpha_t \# \nu_t^*$, and denote $\nu_t := \beta \# \mu_t$ for $t \geq 0$.

3.1 Embedding components

We first describe some primitives that will be used in the construction of the stochastic HST embedding α .

3.1.1 Carving out semi-partitions

A *semi-partition* P of X is a collection of pairwise disjoint subsets of X . For such a semi-partition, denote

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

Define $[P] \subseteq X$ by $[P] := \bigcup_{S \in P} S$. We will sometimes think of P as a function that takes $x \in [P]$ to the unique set $P(x) \in P$ containing x . If $x \notin [P]$, we take $P(x) := \emptyset$. If P, P' are two semi-partitions, say that P is a *refinement* of P' if for every $S \in P$, there is an $\hat{S} \in P'$ such that $S \subseteq \hat{S}$. Say that \hat{P} is Δ -*bounded* if $S \in \hat{P} \implies \text{diam}_X(S) \leq \Delta$.

Consider a triple (C, R, π) where $C \subseteq X$ is a finite set, $R : C \rightarrow \mathbb{R}_+$, and $\pi : [|C|] \rightarrow C$ is a bijection. This defines a semi-partition into at most $|C|$ sets by iteratively carving out balls:

$$\hat{P}(C, R, \pi) := \left\{ B_X(\pi(i), R(\pi(i))) \setminus \bigcup_{h < i} B_X(\pi(h), R(\pi(h))) : i = 1, 2, \dots, |C| \right\}$$

By construction, $\hat{P}(C, R, \pi)$ is $(2 \max_{x \in C} R(x))$ -bounded.

3.1.2 Heavy nets and the annealed measures

Recall that we will define the pair (μ, α) inductively. For $t \geq 0$, denote by $\bar{\nu}_t \in \mathbb{M}_k(X)$ the measure

$$\bar{\nu}_t := \mathbb{E}[\nu_t].$$

Let $\lambda := \max(9, \tau)^2$ and consider $0 < \delta < 1/2$. We will choose δ later so that $\delta \asymp 1/f_4(k)$.

Say that a subset $S \subseteq X$ is r -*separated* if $x, y \in S \implies d_X(x, y) > r$. Say that a pair (x, r) with $x \in X$ and $r > 0$ is t -*heavy* if

$$\bar{\nu}_t(B_X(x, r)) \geq (1 - \delta) \bar{\nu}_t(B_X(x, \lambda r)). \quad (3.1)$$

We will also refer to a ball $B = B_X(x, r)$ as t -*heavy* if (x, r) is t -heavy, but in such cases the center and radius will be specified (as a set, B does not necessarily have a unique center or radius).

A set $\Lambda \subseteq X$ is called a t -*heavy r -net* if it is $3r$ -separated and satisfies

$$\left(x, \frac{r}{2\sqrt{\lambda}} \right) \text{ is } t\text{-heavy} \implies d_X(x, \Lambda) \leq \frac{r}{\sqrt{\lambda}}. \quad (3.2)$$

3.1.3 Cluster fusion

Given a semi-partition \hat{P} , a finite set of representatives $\Lambda \subseteq X$, and a radius $r > 0$, we now define the r -*fusion* of \hat{P} along Λ as follows. For $x \in \Lambda$, define

$$U_x := B_X(x, r) \cup \bigcup_{\substack{S \in \hat{P}: \\ B_X(x, r) \cap S \neq \emptyset}} S. \quad (3.3)$$

See [Figure 3](#).

Define the collection of fused clusters:

$$\mathcal{H}(\hat{P}, \Lambda, r) := \{U_x : x \in \Lambda\}, \quad (3.4)$$

and the semi-partition (cf. [Lemma 3.2](#)) of fused and unfused clusters:

$$\hat{Q}(\hat{P}, \Lambda, r) := \mathcal{H}(\hat{P}, \Lambda, r) \cup \left\{ S \in \hat{P} : S \cap \left[\mathcal{H}(\hat{P}, \Lambda, r) \right] = \emptyset \right\}.$$

The idea here is that in passing from \hat{P} to \hat{Q} , all the sets $S \in \hat{P}$ that intersect some ball $B_X(x, r)$ for $x \in \Lambda$ are “fused” into a single set U_x . (For technical reasons—see [Lemma 4.7](#) below—the ball itself is also fused in.)




Figure 3: Clusters U_{x_1} and U_{x_2} created by fusing the clusters intersecting $B_X(x_1, r)$ and $B_X(x_2, r)$.

Lemma 3.2. *If \hat{P} is Δ -bounded and Λ is $(r + \Delta)$ -separated, then $\hat{Q}(\hat{P}, \Lambda, r)$ is a $2(r + \Delta)$ -bounded semi-partition.*

Proof. Observe that $\text{diam}_X(U_x) \leq 2(r + \Delta)$. Moreover, every $y \in U_x$ satisfies $d_X(x, y) \leq r + \Delta$, hence if Λ is $(r + \Delta)$ -separated, then the sets $\{U_x : x \in \Lambda\}$ are pairwise disjoint. \square

3.1.4 Refinement and HST embeddings

Consider now a sequence $\hat{Q} = \langle \hat{Q}^j : j \in \mathbb{Z}_+ \rangle$ of semi-partitions of X such that $\hat{Q}^0 = \{X\}$ and

$$\hat{Q}^j \text{ is } \tau^{-j}\text{-bounded for all } j \geq 1. \quad (3.5)$$

We use these to define a sequence $Q = \langle Q^j : j \in \mathbb{Z}_+ \rangle$ of successively refined full partitions of X as follows.

First, we complete each semi-partition to a full partition \bar{Q}_t^j by adding singleton clusters:

$$\bar{Q}^j := \hat{Q}^j \cup \left\{ \{x\} : x \in X \setminus [\hat{Q}^j] \right\} \quad \forall j \in \mathbb{Z}_+.$$

Now we inductively define $Q^0 := \bar{Q}^0$ and for $j \geq 1$:

$$Q^j := \{S \cap S' : S \in \bar{Q}^j, S' \in Q^{j-1}\}.$$

This ensures that for each $j \in \mathbb{Z}_+$, Q^{j+1} is a refinement of Q^j .

For $x \in X$, define

$$\text{rank}^{\hat{Q}}(x) := \max \left\{ j \in \mathbb{Z}_+ : x \in \bigcap_{i \leq j} [\hat{Q}^i] \right\}. \quad (3.6)$$

We can now define an embedding $\alpha^{\hat{Q}} : X \rightarrow \mathcal{L}_{\mathbb{T}}$ by

$$\alpha^{\hat{Q}}(x) := \left\langle \left(Q^0(x); 0 \right), \left(Q^1(x); 0 \right), \dots, \left(Q^r(x); 0 \right), \left(Q^{r+1}(x); 1 \right), \left(Q^{r+2}(x); 1 \right), \dots \right\rangle, \quad (3.7)$$

where $r = \text{rank}^{\hat{Q}}(x)$.

One should verify that the latter sequence is indeed a decorated leaf chain by construction and (3.5). This is the only place that we make use of decorated chains in the proof of [Theorem 2.6](#). The particular form of (3.7) will be employed to prove (5.11) which asserts that the Φ_t potential does not increase under insertions (essentially because we have assumed our algorithm is sensible, thus it does not place mass in subtrees with a non-zero decoration).

Later, we will use the following basic fact.

Lemma 3.3. *For every sequence \hat{Q} of semi-partitions:*

$$\text{dist}_{\mathbb{T}} \left(\alpha^{\hat{Q}}(x), \alpha^{\hat{Q}}(y) \right) \leq 2\tau \sum_{j \geq 1} \tau^{-j} \Delta_{\hat{Q}_j}(x, y) \quad \forall x, y \in X.$$

Proof. Consider $x, y \in X$ and suppose that $\text{dist}_{\mathbb{T}} \left(\alpha^{\hat{Q}}(x), \alpha^{\hat{Q}}(y) \right) = \tau^{-\ell}$ for some $\ell > 0$. It is straightforward to check that $\ell + 1 = \min \{ j : \Delta_{\hat{Q}_j}(x, y) > 0 \}$. \square

3.1.5 Truncated exponential radii

For every $j \in \mathbb{Z}$, consider the probability distribution γ_j with density:

$$d\gamma_j(r) := \frac{K\tau^j \log K}{K-1} \exp \left(-r\tau^j \log K \right) \mathbb{1}_{[0, \tau^{-j}]}(r).$$

This is simply an exponential distribution truncated at τ^{-j} . Bartal [[Bar96](#)] showed that such distributions are extremely useful in the construction of random HST embeddings.

Lemma 3.4. *Consider a finite set $C \subseteq X$ and a permutation $\pi : [|C|] \rightarrow C$. Choose $\hat{R} : C \rightarrow \mathbb{R}_+$ so that $\{\hat{R}(x) : x \in C\}$ are independent random variables with law γ_j , and define $R(x) := \hat{R}(x) + \tau^{-j}$. Then $\hat{P} := \hat{P}(C, R, \pi)$ is a $4\tau^{-j}$ -bounded semi-partition with probability one, and moreover for every $x, y \in X$:*

$$\mathbb{P}[\Delta_{\hat{P}}(x, y) > 0] \leq O(\log(|C| + 1)) d_X(x, y) \tau^j. \quad (3.8)$$

If $\Lambda \subseteq X$ is any $6\tau^{-j}$ -separated set, then the $2\tau^{-j}$ -fusion of \hat{P} along Λ :

$$\hat{Q} := \hat{Q}(\hat{P}, \Lambda, 2\tau^{-j})$$

is a $16\tau^{-j}$ -bounded semi-partition of X . If $d_X(x, C \cup \Lambda) \leq \tau^{-j}$ and $y \in X$, then:

$$\mathbb{P}[\Delta_{\hat{Q}}(x, y) > 0] \leq O(\log(|C| + 1)) d_X(x, y) \tau^j. \quad (3.9)$$

Proof. The fact that \hat{P} is a $4\tau^{-j}$ -bounded semi-partition follows immediately from the fact that γ_j is supported on $[0, \tau^{-j}]$. Moreover, (3.8) is a standard bound (see, e.g., [[BCL⁺17](#), Lem 4.8]).

That \hat{Q} is a $16\tau^{-j}$ -bounded semi-partition follows from [Lemma 3.2](#). Let us now verify (3.9). We may assume that $d_X(x, y) \leq \tau^{-j}$, else the claim is vacuous. If $d_X(x, \Lambda) \leq \tau^{-j}$, then $x, y \in B_X(z, 2\tau^{-j})$ for some $z \in \Lambda$, hence $x, y \in [\hat{Q}]$ and $\hat{Q}(x) = \hat{Q}(y)$ because $x, y \in U_z$ (recall (3.3)).

Now assume that $d_X(x, C) \leq \tau^{-j}$. Observe that in this case, $x \in [\hat{P}]$ with probability one and by construction of the fusion, $\Delta_{\hat{Q}}(x, y) \leq \Delta_{\hat{P}}(x, y)$, meaning that (3.9) follows from (3.8). \square

3.2 The online algorithm

For $j, t \in \mathbb{Z}_+$, we will maintain several random \mathcal{F} -adapted sequences: Centers $C_t^j \subseteq X$, along with radii $R_t^j : C_t^j \rightarrow \mathbb{R}_+$, permutations $\pi_t^j : [|C_t^j|] \rightarrow C_t^j$, and t -heavy τ^{-j} -nets Λ_t^j . These give rise to semi-partitions $\hat{P}_t^j := \hat{P}(C_t^j, R_t^j, \pi_t^j)$ and fusions $\hat{Q}_t^j := \hat{Q}(\hat{P}_t^j, \Lambda_{t-1}^j, 2\tau^{-j-1})$, along with embeddings

$$\alpha_t := \alpha^{\hat{Q}_t},$$

where $\hat{Q}_t := \langle \hat{Q}_t^j : j \in \mathbb{Z}_+ \rangle$.

We will also maintain a sequence $\langle \chi_t \in \Gamma : t \geq 0 \rangle$ of configurations. These yield our sequence $\mu = \langle \mu_t \in \mathbb{M}_k(\mathcal{V}_\mathbb{T}) : t \geq 0 \rangle$ of induced fractional k -server measures: $\mu_t := \mu^{\chi_t}$.

Initialization. Let $\chi_0 \in \Gamma$ and $\ell_0 \in \mathcal{L}_\mathbb{T}$ be the configuration and leaf promised in [Section 2.3](#).

- For all $j \geq 1$, define $C_0^j := \{\ell_0\}$ and $\Lambda_0^j := \{\ell_0\}$.
- For all $t \geq 0$, $C_t^0 := \emptyset$, $\Lambda_t^0 := \emptyset$, and $\hat{P}_t^0 := X$
- Define $\chi_{-1} := \chi_0$ and $\mu_{-1} := \mu_0$. For all $j \geq 0$, define $C_{-1}^j := C_0^j$, $\hat{P}_{-1}^j = \hat{P}_0^j$, $\Lambda_{-1}^j := \Lambda_0^j$.

Request. Suppose we receive a request $\sigma_t \in X$ for some $t \geq 1$. For $j \geq 1$, denote

$$\mathbf{I}_t^j := \begin{cases} 1 & d_X(\sigma_t, C_{t-1}^j \cup \Lambda_{t-1}^j) > \tau^{-j-1} \\ 0 & \text{otherwise.} \end{cases}$$

Deletions. In the next definition, $K \geq 1$ is a parameter that will be chosen later (our choice will satisfy $K \leq k^{O(1)}$). For every $j \geq 1$:

$$C_{t,\text{del}}^j := \begin{cases} C_{t-1}^j & \mathbf{I}_t^j = 0 \text{ or } |C_{t-1}^j| < K \\ C_{t-1}^j \setminus \{z_t^j\} & \text{otherwise,} \end{cases}$$

where $z_t^j \in C_{t-1}^j$ is chosen uniformly at random. Denote $\hat{P}_{t,\text{del}}^j := \hat{P}(C_{t,\text{del}}^j, R_{t-1}^j, \pi_{t-1}^j)$.

Fission. Denote

$$\hat{Q}_{t,\text{fis}}^j := \hat{Q}(\hat{P}_{t,\text{del}}^j, \Lambda_{t-1}^j \cap \Lambda_{t-2}^j, 2\tau^{-j-1}).$$

This is the semi-partition $\hat{P}_{t,\text{del}}^j$ fused only along the centers that survive from time $t-2$ to $t-1$.

Insertions. For every $j \geq 1$, if $\mathbf{I}_t^j = 1$, we define:

$$\begin{aligned} C_t^j &:= C_{t,\text{del}}^j \cup \{\sigma_t\} \\ \pi_t^j(|C_t^j|) &:= \sigma_t \\ R_t^j(\sigma_t) &:= \tau^{-j-1} + Z_t^j, \end{aligned}$$

where Z_t^j is sampled independently with law γ_{j+1} . If $\mathbf{I}_t^j = 0$, then $C_t^j := C_{t-1}^j$.

In either case, we define π_t^j so that it induces the same ordering on $C_t^j \setminus \{\sigma_t\}$ as π_{t-1}^j , and R_t^j so that $R_t^j|_{C_t^j \setminus \{\sigma_t\}} = R_{t-1}^j|_{C_{t-1}^j \setminus \{\sigma_t\}}$.

Fusion. Consider the semi-partition $\hat{Q}_t^j = \hat{Q}(\hat{P}_t^j, \Lambda_{t-1}^j, 2\tau^{-j-1})$ and its prefused version:

$$\hat{Q}_{t,\text{pre}}^j := \hat{Q}_{t,\text{fis}}^j \cup \left\{ \hat{P}_t^j(\sigma_t) \setminus \left[\hat{Q}_{t,\text{fis}}^j \right] \right\} \cup \left\{ B_X(x, 2\tau^{-j-1}) \setminus \left[\hat{Q}_{t,\text{fis}}^j \right] : x \in \Lambda_{t-1}^j \setminus \Lambda_{t-2}^j \right\}. \quad (3.10)$$

Define $\hat{Q}_{t,\text{pre}} := \langle \hat{Q}_{t,\text{pre}}^j : j \geq 1 \rangle$. We have $[\hat{Q}_t^j] = [\hat{Q}_{t,\text{pre}}^j]$ and $\hat{Q}_{t,\text{pre}}^j$ is a refinement of \hat{Q}_t^j by construction.

Thus we can realize \hat{Q}_t from $\hat{Q}_{t,\text{pre}}$ via an iterative merging of pairs of siblings. Note that this can be expressed as a composition of canonical injections; to merge siblings $\xi, \xi' \in V_{\mathbb{T}}^j$ with $\text{diam}_X(\mathfrak{b}(\xi) \cup \mathfrak{b}(\xi')) \leq \tau^{-j}$, we fuse ξ and ξ' into their common 0-decorated sibling $(\mathfrak{b}(\xi) \cup \mathfrak{b}(\xi'), 0) \in V_{\mathbb{T}}^j$. Let φ_t denote the corresponding fusion map (recall that a fusion map is a composition of canonical injections). Using Axiom (A3), this yields a configuration $\chi_{t-1}(\varphi_t) \in \Gamma$ such that $\mu^{\chi_{t-1}(\varphi_t)} = \varphi_t \# \mu_{t-1}$ and (2.5) is satisfied.

HST evolution. We update the configuration:

$$\chi_t := \gamma(\chi_{t-1}(\varphi_t), \alpha_t(\sigma_t)).$$

Heavy net maintenance. Now we specify how to update Λ_{t-1}^j to Λ_t^j .

For $j = 1, 2, \dots$, do the following:

1. Set $\tilde{\Lambda}_t^j := \Lambda_{t-1}^j$.
2. While there is some $x \in X$ such that $(x, \frac{\tau^{-j}}{2\sqrt{\lambda}})$ is t -heavy and $d_X(x, \tilde{\Lambda}_t^j) > \frac{\tau^{-j}}{\sqrt{\lambda}}$:
 - (a) Remove from $\tilde{\Lambda}_t^j$ all $y \in X$ such that $d_X(x, y) < \frac{\sqrt{\lambda}}{3}\tau^{-j}$.
 - (b) $\tilde{\Lambda}_t^j := \tilde{\Lambda}_t^j \cup \{x\}$.
3. Set $\Lambda_t^j := \tilde{\Lambda}_t^j$.

4 Distortion analysis

Our first goal is to establish a bound on how much α distorts distance in expectation. Let us first verify a few basic properties of the embedding algorithm from Section 3.2.

Lemma 4.1. *Assume that $\tau \geq 12$ and $\lambda \geq 81$. Then for each $j \geq 1$ and $t \geq 1$, it holds that*

1. Λ_t^j is a t -heavy τ^{-j} -net.
2. \hat{Q}_t^j is a τ^{-j} -bounded semi-partition.

Proof. Λ_t^j is explicitly constructed to satisfy (3.2) and to be $3r$ -separated with $r = \tau^{-j}$, as long as $\lambda \geq 81$. We need to verify that the construction is well-defined, i.e., that the loop defining Λ_t^j always terminates.

To prove this, it suffices to show that if $y \in X$ is removed in step 2(a), then $(y, \frac{\tau^{-j}}{2\sqrt{\lambda}})$ is not t -heavy. To that end, it suffices to show that there cannot be two points $x, y \in X$ and a radius $r > 0$ satisfying

$$\frac{\lambda}{3}r \geq d_X(x, y) > 2r \quad \text{and} \quad (x, r) \text{ and } (y, r) \text{ are } t\text{-heavy.}$$

Note that under these assumptions, $B_X(x, r) \cap B_X(y, r) = \emptyset$, but $B_X(x, \lambda r) \supseteq B_X(y, r)$ and $B_X(y, \lambda r) \supseteq B_X(x, r)$. Therefore it cannot be that both (x, r) and (y, r) are t -heavy as long as $\delta < 1/2$ (recall (3.1)).

Now the fact that \hat{Q}_t^j is a τ^{-j} -bounded semi-partition follows from Lemma 3.4. \square

We want to distinguish two types of randomness used in the algorithm. There is the probability space underlying the choice of elements z_t^j in the deletion step which we denote by Ω^{del} . All other randomness is denoted by Ω^{hst} .

Fact 4.2. *The random variables C_t^j and Λ_t^j are independent of Ω^{hst} . Note that Λ_t^j is defined using \bar{v}_t , but this measure is constructed by averaging over Ω^{hst} .*

4.1 Active scales

Define the functions $\rho, \hat{\rho} : X \times \mathbb{M}_k(X) \rightarrow \mathbb{R}_+$ by

$$\begin{aligned}\rho(x, v) &:= \sup \{r : v(B_X(x, r)) < 1/2\} , \\ \hat{\rho}(x, v) &:= \inf \{W_X^1(v, v') : v'(x) \geq 1\} .\end{aligned}$$

The next lemma is straightforward from the definitions.

Lemma 4.3. *The following hold true.*

1. For any $v \in \mathbb{M}_k(X)$, the maps $x \mapsto \rho(x, v)$ and $x \mapsto \hat{\rho}(x, v)$ are 1-Lipschitz on (X, d_X) .
2. For any $x \in X$, the map $v \mapsto \hat{\rho}(x, v)$ is 1-Lipschitz on $(\mathbb{M}_k(X), W_X^1)$.
3. For any $v \in \mathbb{M}_k(X)$ and $x \in X$:

$$\hat{\rho}(x, v) \geq \frac{\rho(x, v)}{2} . \tag{4.1}$$

Make the further definitions: For $t \geq 1$,

$$\begin{aligned}\rho_t(\xi) &:= \rho(\xi, \bar{v}_t) \\ \hat{\rho}_t(\xi) &:= \hat{\rho}(\xi, \bar{v}_t) .\end{aligned}$$

Lemma 4.4. *For every $t \geq 1$, it holds that*

$$\frac{\rho_{t-1}(\sigma_t)}{2} \leq \hat{\rho}_{t-1}(\sigma_t) \leq \mathbb{E} \left[W_{\mathbb{T}}^1(\mu^{\chi_{t-1}(\varphi_t)}, \mu^{\chi_t}) \right] .$$

Proof. The first inequality follows from (4.1). To prove the second, write

$$W_{\mathbb{T}}^1(\mu^{\chi_{t-1}(\varphi_t)}, \mu^{\chi_t}) \geq W_X^1(\beta\#\mu^{\chi_{t-1}(\varphi_t)}, \beta\#\mu^{\chi_t}) = W_X^1(\beta\#\mu^{\chi_{t-1}}, \beta\#\mu^{\chi_t}) \geq \hat{\rho}(\sigma_t, \beta\#\mu^{\chi_{t-1}}) .$$

where the last inequality follows from $\mu^{\chi_t}(\sigma_t) \geq 1$.

Now convexity of the Wasserstein distance yields

$$\mathbb{E} [\hat{\rho}(\sigma_t, \beta\#\mu^{\chi_{t-1}})] \geq \hat{\rho}(\sigma_t, \mathbb{E} \beta\#\mu^{\chi_{t-1}}) = \hat{\rho}(\sigma_t, \beta\#\mathbb{E} \mu^{\chi_{t-1}}) = \hat{\rho}(\sigma_t, \bar{v}_{t-1}) . \quad \square$$

Corollary 4.5. *It holds that*

$$\sum_{t \geq 1} \rho_{t-1}(\sigma_t) \leq 2 \sum_{t \geq 1} \hat{\rho}_{t-1}(\sigma_t) \leq 2 \mathbb{E} [\text{cost}_{\mathbb{T}}^{\mathbb{F}}(\boldsymbol{\mu})] .$$

Definition 4.6. Say that a point $x \in X$ is (j, t) -heavy if $d_X(x, \Lambda_t^{j-1}) \leq \tau^{-j-1}$. If $d_X(x, \Lambda_t^{j-1}) \geq \frac{1}{2}\tau^{-j-1}$, say that x is (j, t) -light. (A point can be both heavy and light.)

We record a fact that follows from our construction of \hat{Q}_t (cf. (3.3)).

Lemma 4.7. *If $x \in X$ is (j, t) -heavy, then $B_X(x, \tau^{-j-1}) \subseteq \hat{Q}_t^j(x)$.*

Denote $\eta := (32kf_1(k))^{-1}$; we may assume that $\eta \geq k^{-O(1)}$. We now define a subset $\mathbb{J}_t(x) \subseteq \mathbb{Z}_+$ of “active” scales for a given $x \in X$:

$$\begin{aligned} \mathbb{L}_t(x) &:= \{j \in \mathbb{Z}_+ : x \text{ is } (j, t)\text{-light}\} , \\ \mathbb{J}_t(x) &:= \{j \in \mathbb{Z}_+ : \tau^{-j} > \eta \rho_t(x)\} \cap \mathbb{L}_t(x) , \end{aligned}$$

The next lemma is an essential component of all our arguments: For every $x \in X$, there are only $O(\frac{1}{\delta} \log k)$ active scales.

Lemma 4.8. *For every $x \in X$ and $t \geq 0$,*

$$|\mathbb{J}_t(x)| \leq O\left(\frac{\log k}{\delta} + \log \frac{1}{\eta}\right) \leq O\left(\frac{\log k}{\delta}\right) .$$

Proof. If x is (j, t) -light, it means that $B_X\left(x, \frac{\tau^{-j}}{2\sqrt{\lambda}}\right)$ is not t -heavy, which means that

$$\bar{v}_t\left(B_X\left(x, \frac{\tau^{-j}}{2\sqrt{\lambda}}\right)\right) < (1 - \delta)\bar{v}_t\left(B_X\left(x, \frac{\sqrt{\lambda}}{2}\tau^{-j}\right)\right) \leq (1 - \delta)\bar{v}_t\left(B_X\left(x, \tau^{-j}\right)\right) .$$

Since $\bar{v}_t(X) = k$ and $\bar{v}_t(B_X(x, \rho_t(x))) \geq 1/2$, the result follows using $\lambda, \tau \leq O(1)$ and the fact that there are only $O(\log \frac{1}{\eta})$ additional scales between $\eta \rho_t(x)$ and $\rho_t(x)$. \square

4.2 The expected stretch

Let us now establish the central claim of this section.

Lemma 4.9. *For every $t \geq 1$ and every $x \in X$, it holds that*

$$\mathbb{E}_{\Omega^{\text{hst}}} \left[\text{dist}_{\mathbb{T}}(\alpha_t(x), \alpha_t(\sigma_t)) \right] \leq O\left(\frac{1}{\delta} \log(k) \log(K)\right) d_X(x, \sigma_t) + 2\eta \rho_{t-1}(\sigma_t) .$$

Before proving the lemma, we state a consequence. It uses the definition of η and the fact that ν^* is lazy.

Corollary 4.10. *For every $t \geq 1$, it holds that*

$$\mathbb{E}_{\Omega^{\text{hst}}} \left[W_{\mathbb{T}}^1(\alpha_t \# \nu_t^*, \alpha_t \# \nu_{t-1}^*) \right] \leq O\left(\frac{1}{\delta} \log(k) \log(K)\right) W_X^1(\nu_t^*, \nu_{t-1}^*) + \frac{\rho_{t-1}(\sigma_t)}{16f_1(k)} .$$

Proof of Lemma 4.9. Let $M := \max(\eta\rho_{t-1}(\sigma_t), 2\tau d_X(x, \sigma_t))$ and $j_0 := \max\{j \in \mathbb{Z}_+ : \tau^{-j} \geq M\}$. From Lemma 3.3, it holds that

$$\begin{aligned} \frac{1}{2\tau} \text{dist}_\mathbb{T}(\alpha_t(x), \alpha_t(\sigma_t)) &\leq \sum_{j \geq 1} \tau^{-j} \Delta_{\hat{Q}_t^j}(x, \sigma_t) \\ &\leq \eta\rho_{t-1}(\sigma_t) + 2\tau d_X(x, \sigma_t) + \sum_{j=1}^{j_0} \tau^{-j} \Delta_{\hat{Q}_t^j}(x, \sigma_t). \end{aligned}$$

Note that $j \leq j_0$ implies $x \in B_X(\sigma_t, \frac{1}{2}\tau^{-j-1})$. Thus if additionally $j \notin \mathbb{J}_{t-1}(\sigma_t)$, then Lemma 4.7 asserts that $\Delta_{\hat{Q}_t^j}(x, \sigma_t) = 0$. Therefore:

$$\sum_{j=1}^{j_0} \tau^{-j} \Delta_{\hat{Q}_t^j}(x, \sigma_t) \leq \sum_{j \in \mathbb{J}_{t-1}(\sigma_t)} \tau^{-j} \Delta_{\hat{Q}_t^j}(x, \sigma_t).$$

Now Lemma 3.4 (specifically (3.9)) gives, for every $j \geq 1$:

$$\mathbb{E}_{\Omega^{\text{hst}}} \left[\tau^{-j} \Delta_{\hat{Q}_t^j}(x, \sigma_t) \right] \leq O(\log K) d_X(x, \sigma_t).$$

Therefore:

$$\mathbb{E}_{\Omega^{\text{hst}}} \left[\sum_{j \in \mathbb{J}_{t-1}(\sigma_t)} \tau^{-j} \Delta_{\hat{Q}_t^j}(x, y) \right] \leq O(\log K) |\mathbb{J}_{t-1}(\sigma_t)| d_X(x, \sigma_t) \leq O\left(\frac{1}{\delta} \log(k) \log(K)\right) d_X(x, \sigma_t),$$

where the final inequality uses Lemma 4.8. □

5 The competitive ratio

In order to prove our main result (Theorem 3.1), we will relate $\text{cost}_X(\nu^*)$ and $\mathbb{E}[\text{cost}_\mathbb{T}^F(\mu)]$ using three potential functions.

5.1 The HST potential

Let us define the primary potential function:

$$\Phi_t := \Phi(\alpha_t \# \nu_t^*; \chi_t).$$

In the course of analyzing $\Delta_t \Phi_t := \Phi_t - \Phi_{t-1}$, we will define a number of auxiliary objects that will be used in the remainder of Section 5. Let us define $\alpha_t^{\text{pre}} := \alpha_{\hat{Q}_t, \text{pre}}$, where we recall the prefusion semi-partitions $\hat{Q}_{t, \text{pre}}$ from (3.10).

We can then express

$$\Delta_t \Phi_t = \left[\Phi(\alpha_t \# \nu_t^*; \chi_t) - \Phi(\alpha_t \# \nu_t^*; \chi_{t-1}(\varphi_t)) \right] \quad [\mu \text{ movement}] \quad (5.1)$$

$$+ \left[\Phi(\alpha_t \# \nu_t^*; \chi_{t-1}(\varphi_t)) - \Phi(\alpha_t \# \nu_{t-1}^*; \chi_{t-1}(\varphi_t)) \right] \quad [\nu^* \text{ movement}] \quad (5.2)$$

$$+ \left[\Phi(\alpha_t \# \nu_{t-1}^*; \chi_{t-1}(\varphi_t)) - \Phi(\alpha_t^{\text{pre}} \# \nu_{t-1}^*; \chi_{t-1}) \right] \quad [\text{fusion}] \quad (5.3)$$

$$+ [\Phi(\alpha_t^{\text{pre}} \# v_{t-1}^*; \chi_{t-1}) - \Phi(\alpha_{t-1} \# v_{t-1}^*; \chi_{t-1})]. \quad \text{[dynamic update]} \quad (5.4)$$

Addressing (5.1), Axiom (A2) gives

$$\Phi(\alpha_t \# v_t^*; \chi_t) - \Phi(\alpha_t \# v_t^*; \chi_{t-1}(\varphi_t)) \leq -W_{\mathbb{T}}^1(\mu_t, \mu^{\chi_{t-1}(\varphi_t)}). \quad (5.5)$$

Axiom (A1) yields

$$\Phi(\alpha_t \# v_t^*; \chi_{t-1}(\varphi_t)) - \Phi(\alpha_t \# v_{t-1}^*; \chi_{t-1}(\varphi_t)) \leq f_1(k) W_{\mathbb{T}}^1(\alpha_t \# v_t^*, \alpha_t \# v_{t-1}^*),$$

and this allows us to control (5.2) using [Corollary 4.10](#):

$$\begin{aligned} \mathbb{E}_{\Omega^{\text{hst}}} [\Phi(\alpha_t \# v_t^*; \chi_{t-1}(\varphi_t)) - \Phi(\alpha_t \# v_{t-1}^*; \chi_{t-1}(\varphi_t))] \\ \leq O(f_1(k)^{\frac{1}{\delta}} \log(k) \log(K)) W_X^1(v_t^*, v_{t-1}^*) + \frac{\rho_{t-1}(\sigma_t)}{16}. \end{aligned} \quad (5.6)$$

Axiom (A3) asserts that the term in (5.3) is non-positive; this uses the fact that $\alpha_t \# v_{t-1}^* = \varphi_t \# \alpha_t^{\text{pre}} \# v_{t-1}^*$.

Thus we have satisfactorily controlled $\Delta_t \Phi_t$ except for the term (5.4) corresponding to the dynamic modifications we make to the embedding (insertions, deletions, and fission operations).

5.1.1 Dynamic updates: Analyzing (5.4)

We need to define a few more intermediate embeddings, so let us denote:

$$\begin{aligned} \hat{Q}_{t,\text{del}}^j &:= \hat{Q}(\hat{P}_{t,\text{del}}^j, \Lambda_{t-1}^j, 2\tau^{-j-1}) \\ \hat{Q}_{t,\text{del}} &:= \langle \hat{Q}_{t,\text{del}}^j : j \geq 1 \rangle, & \alpha_t^{\text{del}} &:= \alpha^{\hat{Q}_{t,\text{del}}} \\ \hat{Q}_{t,\text{fis}} &:= \langle \hat{Q}_{t,\text{fis}}^j : j \geq 1 \rangle, & \alpha_t^{\text{fis}} &:= \alpha^{\hat{Q}_{t,\text{fis}}}. \end{aligned}$$

We will further decompose (5.4):

$$\begin{aligned} \Phi(\alpha_t^{\text{pre}} \# v_{t-1}^*; \chi_{t-1}) - \Phi(\alpha_{t-1} \# v_{t-1}^*; \chi_{t-1}) \\ = [\Phi(\alpha_t^{\text{pre}} \# v_{t-1}^*; \chi_{t-1}) - \Phi(\alpha_t^{\text{fis}} \# v_{t-1}^*; \chi_{t-1})] \quad \text{[insertion]} \end{aligned} \quad (5.7)$$

$$+ [\Phi(\alpha_t^{\text{fis}} \# v_{t-1}^*; \chi_{t-1}) - \Phi(\alpha_t^{\text{del}} \# v_{t-1}^*; \chi_{t-1})] \quad \text{[fission]} \quad (5.8)$$

$$+ [\Phi(\alpha_t^{\text{del}} \# v_{t-1}^*; \chi_{t-1}) - \Phi(\alpha_{t-1} \# v_{t-1}^*; \chi_{t-1})]. \quad \text{[deletion]} \quad (5.9)$$

Insertions (5.7). We now claim that

$$\Phi(\alpha_t^{\text{pre}} \# v_{t-1}^*; \chi_{t-1}) - \Phi(\alpha_t^{\text{fis}} \# v_{t-1}^*; \chi_{t-1}) \leq \sum_{j \geq 0} \left(\Phi(\alpha_t^j \# v_{t-1}^*; \chi_{t-1}) - \Phi(\alpha_t^{j-1} \# v_{t-1}^*; \chi_{t-1}) \right)_+ \quad (5.10)$$

$$\leq 0, \quad (5.11)$$

where $\alpha_t^0 = \alpha_t^{\text{fis}}$ and α_t^j results from α_t^{j-1} by incorporating the possible insertion of a set $\{S_j\} = \hat{P}_t^j \setminus \hat{P}_{t,\text{del}}^j$. Thus α_t^j and α_t^{j-1} agree outside $S_j \setminus [\hat{Q}_{t,\text{fis}}^j]$.

By definition of the embedding (recall (3.7)), $\alpha_t^j(x) \neq \alpha_t^{j-1}(x)$ only at points $x \in \text{supp}(v_{t-1}^*)$ with rank less than j (cf. (3.6)), but the images of all such points lie outside $\mathcal{L}_{\mathbb{T}}^0$ by construction. Since $\mu^{\chi_{t-1}}$ is supported on $\mathcal{L}_{\mathbb{T}}^0$, Axiom (A4) implies that each term in (5.10) is zero, establishing (5.11).

Deletions (5.9). Denote by

$$j_t^* = \min\{j \geq 1 : \mathbf{I}_t^j = 1\} \quad (5.12)$$

be the highest scale at which an insertion occurs (we take $j_t^* := 0$ if no such j exists). If $\mathbf{I}_t^j = 1$, then some cluster $S_j \in \hat{\mathcal{P}}_t^j$ with center z_t^j is possibly deleted. Therefore:

$$W_{\mathbb{T}}^1(\alpha_t^{\text{del}} \# v_{t-1}^*, \alpha_{t-1} \# v_{t-1}^*) \leq \sum_{j \geq 1} \mathbb{1}_{\{\mathbf{I}_t^j = 1\}} v_{t-1}^*(S_j) \tau^{-j+1}.$$

Recall that $v_{t-1}^*(X) = k$. Since we remove a uniformly random level- j cluster and there are at least K of them (if a deletion takes place), it holds that

$$\mathbb{E}_{\Omega^{\text{del}}} [W_{\mathbb{T}}^1(\alpha_t^{\text{del}} \# v_{t-1}^*, \alpha_{t-1} \# v_{t-1}^*) \mid j_t^*] \leq \frac{k}{K} \sum_{j \geq 1} \tau^{-j+1} \mathbf{I}_t^j \leq \frac{2\tau k}{K} \tau^{-j_t^*} \mathbb{1}_{\{j_t^* > 0\}}.$$

where we take expectation only over the random choice of which cluster to delete. In particular, using Axiom (A1), this implies that

$$\mathbb{E}_{\Omega^{\text{del}}} [\Phi(\alpha_t^{\text{del}} \# v_{t-1}^*; \chi_{t-1}) - \Phi(\alpha_{t-1} \# v_{t-1}^*; \chi_{t-1}) \mid j_t^*] \leq \frac{2\tau f_1(k)k}{K} \tau^{-j_t^*} \mathbb{1}_{\{j_t^* > 0\}}. \quad (5.13)$$

It is at this point that we are no longer able to continue analyzing $\Delta_t \Phi_t$ locally in time. Deletions can increase Φ_t , and we need a way of obtaining credit for this increase from prior moment of μ . A similar fact is true for the analysis of (5.8).

We encapsulate the contents of this section as follows. Let $\Delta_t^{\text{fis}} \Phi_t$ denote the expression in (5.8), and define $\Delta \Phi := \sum_{t \geq 1} \Delta_t \Phi_t$.

Lemma 5.1. *There is a constant $C_\Phi \geq 1$ such that*

$$\begin{aligned} \mathbb{E}[\Delta \Phi] &\leq -\frac{3}{4} \mathbb{E}[\text{cost}_{\mathbb{T}}^{\text{F}}(\mu)] + C_\Phi f_1(k) \frac{1}{8} \log(k) \log(K) \text{cost}_X(v^*) \\ &\quad + \sum_{t \geq 1} \mathbb{E} \left[\frac{2\tau f_1(k)k}{K} \tau^{-j_t^*} \mathbb{1}_{\{j_t^* > 0\}} + \Delta_t^{\text{fis}} \Phi_t \right]. \end{aligned}$$

Proof. We sum the inequalities (5.5), (5.6), (5.11), and (5.13), and use the fact (5.3) is non-positive. Summing the right-hand side of (5.5) and taking expectations gives precisely $-\mathbb{E}[\text{cost}_{\mathbb{T}}^{\text{F}}(\mu)]$. We have also employed Corollary 4.5 to bound the $\rho_{t-1}(\sigma_t)$ term from (5.6). \square

5.2 The accuracy potential

The accuracy potential Ψ_t^A will help us track the cost of insertions and deletions. It measures how accurately the tree structure induced by the semi-partitions $\hat{\mathcal{Q}}_t$ represents the fractional server measure. One could effectively ignore Ψ_t^A upon a first reading; using a cruder bound, one loses an $O(\log \log \mathcal{A}_X)$ factor in the competitive ratio.

Finally, for $x, y \in X$, denote the truncated distance function:

$$d_X^j(x, y) := \min(\tau^{-j}, d_X(x, y)).$$

For $\mu \in \mathbb{M}(\mathcal{L}_\mathbb{T})$, a function $\hat{\rho} : X \rightarrow \mathbb{R}_+$, and sequences of finite subsets $C = \langle C^j \subseteq X : j \geq 1 \rangle$ and $\Lambda = \langle \Lambda^j \subseteq X : j \geq 1 \rangle$, we define:

$$\begin{aligned} \psi^A(\mu; \hat{\rho}, C, \Lambda) &:= \sum_{x \in X} \beta \# \mu(x) \sum_{j \geq 1} d_X^j(x, C^j) \cdot \tau^j \left(d_X^j(x, \Lambda^j) - 2\eta \hat{\rho}(x) - \frac{1}{2} \tau^{-j-1} \right)_+ \\ \Psi_t^A &:= \psi^A(\mu_t; \hat{\rho}_{t-1}, C_t, \Lambda_{t-1}), \end{aligned} \quad (5.14)$$

where $C_t := \langle C_t^j : j \geq 1 \rangle$ and $\Lambda_t := \langle \Lambda_t^j : j \geq 1 \rangle$.

The next lemma follows from $\beta \# \varphi \# \mu = \beta \# \mu$ when φ is any fusion map and $\mu \in \mathbb{M}(\mathcal{L}_\mathbb{T})$.

Lemma 5.2. *If φ is a fusion map, then*

$$\psi^A(\varphi \# \mu; \hat{\rho}, C, \Lambda) = \psi^A(\mu; \hat{\rho}, C, \Lambda).$$

Remark 5.3 (Accuracy potential). Recall that an insertion occurs at level j when $d_X(\sigma_t, C_{t-1}^j \cup \Lambda_{t-1}^j) > \tau^{-j-1}$. Such an insertion does not increase the potential Φ (recall (5.11)), but it triggers a level- j deletion which might adversely increase Φ (recall (5.13)). The potential Ψ_t^A measures how accurately the sets $C_t^j \cup \Lambda_{t-1}^j$ approximate the measure $\beta \# \mu_t$.

We know that the underlying k -server algorithm satisfies $\mu_t(\alpha_t(\sigma_t)) \geq 1$, and therefore it should be that either the HST algorithm moves substantially in response to a level- j insertion or the accuracy improves (because $\sigma_t \in C_t^j$), yielding a lower Ψ_t^A value. This gain is used to charge the adverse effects of deletion against the movement of the HST algorithm.

Lemma 5.4. *For every $t \geq 0$ and every sequence C , the map $\mu \mapsto \psi^A(\mu; \hat{\rho}_{t-1}, C, \Lambda_{t-1})$ is $O(\frac{1}{8} \log k)$ -Lipschitz on $(\mathbb{M}(\mathcal{L}_\mathbb{T}), W_\mathbb{T}^1)$.*

Proof. Define $\psi_j : X \rightarrow \mathbb{R}_+$ by

$$\psi_j(x) := d_X^j(x, C^j) \cdot \tau^j \left(d_X^j(x, \Lambda_{t-1}^j) - 2\eta \hat{\rho}_{t-1}(x) - \frac{1}{2} \tau^{-j-1} \right)_+.$$

Consider any $\ell, \ell' \in \mathcal{L}_\mathbb{T}$ and $x = \beta(\ell), y = \beta(\ell')$. Let $\mu' = s(\mathbb{1}_\ell - \mathbb{1}_{\ell'})$ for some $s \neq 0$ and write:

$$\begin{aligned} \frac{1}{|s|} \left| \psi^A(v + v'; \hat{\rho}_{t-1}, C, \Lambda_{t-1}) - \psi^A(v'; \hat{\rho}_{t-1}, C, \Lambda_{t-1}) \right| &\leq \sum_{j \geq 0} \left| \psi_j(\xi) - \psi_j(\xi') \right| \\ &\leq \sum_{j \in \mathbb{J}_{t-1}(x) \cup \mathbb{J}_{t-1}(y)} \left| \psi_j(\xi) - \psi_j(\xi') \right| \\ &\leq O\left(\frac{1}{8} \log k\right) \sup_{j \geq 1} \left| \psi_j(\xi) - \psi_j(\xi') \right|, \end{aligned}$$

where in the second inequality, have used the definition of \mathbb{J}_{t-1} and Lemma 4.3(3), and in the last inequality, Lemma 4.8.

We are left to show that $\|\psi_j\|_{\text{Lip}} \leq 4$ for every $j \geq 1$. Consider that $\hat{\rho}_{t-1}$ is 1-Lipschitz (cf. Lemma 4.3(1)), as are the maps $x \mapsto d_X^j(x, C^j)$ and $x \mapsto d_X^j(x, \Lambda_{t-1}^j)$. Factor $\psi_j = f_j g_j$ with $f_j = d_X^j(x, C^j)$. Then:

$$\|\psi_j\|_{\text{Lip}} \leq \|f_j\|_\infty \|g_j\|_{\text{Lip}} + \|g_j\|_\infty \|f_j\|_{\text{Lip}} \leq 4,$$

completing the proof. \square

Lemma 5.5. For every $\mu \in \mathbb{M}_k(\mathcal{L}_{\mathbb{T}})$ and $t \geq 1$, it holds that

$$|\psi^A(\mu; \hat{\rho}_{t-1}, \mathbf{C}_{t,\text{del}}, \mathbf{\Lambda}_{t-2}) - \psi^A(\mu; \hat{\rho}_{t-2}, \mathbf{C}_{t,\text{del}}, \mathbf{\Lambda}_{t-2})| \leq O(\eta \frac{k}{\delta} \log k) \mathbb{E} \left[W_{\mathbb{T}}^1(\mu^{\chi_{t-2}(\varphi_{t-1})}, \mu^{\chi_{t-1}}) \right].$$

Proof. Since $\hat{\rho}_{-1} = \hat{\rho}_0$, we may assume that $t \geq 2$. Write

$$\begin{aligned} & |\psi^A(\mu; \hat{\rho}_{t-1}, \mathbf{C}_{t,\text{del}}, \mathbf{\Lambda}_{t-2}) - \psi^A(\mu; \hat{\rho}_{t-2}, \mathbf{C}_{t,\text{del}}, \mathbf{\Lambda}_{t-2})| \\ & \leq \eta \beta \# \mu(X) \sup_{x \in X} (|\mathbb{J}_{t-1}(x)| + |\mathbb{J}_{t-2}(x)|) \sup_{x \in X} |\hat{\rho}_{t-1}(x) - \hat{\rho}_{t-2}(x)| \\ & \leq O(\eta k \frac{1}{\delta} \log k) W_X^1(\beta \# \bar{\mu}_{t-1}, \beta \# \bar{\mu}_{t-2}), \end{aligned}$$

where the second inequality is [Lemma 4.8](#), and the last inequality is a consequence of [Lemma 4.3\(2\)](#). By convexity of the Wasserstein distance, we have

$$\begin{aligned} W_X^1(\beta \# \bar{\mu}_{t-1}, \beta \# \bar{\mu}_{t-2}) & \leq \mathbb{E} \left[W_X^1(\beta \# \mu_{t-1}, \beta \# \mu_{t-2}) \right] \\ & = \mathbb{E} \left[W_X^1(\beta \# \mu^{\chi_{t-1}}, \beta \# \mu^{\chi_{t-2}}) \right] \\ & = \mathbb{E} \left[W_X^1(\beta \# \mu^{\chi_{t-1}}, \beta \# \mu^{\chi_{t-2}(\varphi_{t-1})}) \right], \end{aligned}$$

where the last inequality uses $\beta \# \mu^{\chi_{t-2}(\varphi_{t-1})} = \beta \# \varphi_{t-1} \# \mu^{\chi_{t-2}} = \beta \# \mu^{\chi_{t-2}}$. Now the desired result follows from the fact that β is 1-Lipschitz. \square

5.2.1 Analysis

For $t \geq 1$, define $\Delta_t \Psi_t^A := \Psi_t^A - \Psi_{t-1}^A$. We decompose:

$$\begin{aligned} \Delta_t \Psi_t^A & = \psi^A(\mu_t; \hat{\rho}_{t-1}, \mathbf{C}_t, \mathbf{\Lambda}_{t-1}) - \psi^A(\mu_{t-1}; \hat{\rho}_{t-2}, \mathbf{C}_{t-1}, \mathbf{\Lambda}_{t-2}) \\ & = \left[\psi^A(\mu_t; \hat{\rho}_{t-1}, \mathbf{C}_t, \mathbf{\Lambda}_{t-1}) - \psi^A(\mu^{\chi_{t-1}(\varphi_t)}; \hat{\rho}_{t-1}, \mathbf{C}_t, \mathbf{\Lambda}_{t-1}) \right] \quad [\mu \text{ movement}] \quad (5.15) \end{aligned}$$

$$+ \left[\psi^A(\mu^{\chi_{t-1}(\varphi_t)}; \hat{\rho}_{t-1}, \mathbf{C}_t, \mathbf{\Lambda}_{t-1}) - \psi^A(\mu^{\chi_{t-1}(\varphi_t)}; \hat{\rho}_{t-2}, \mathbf{C}_t, \mathbf{\Lambda}_{t-1}) \right] \quad [\text{isolation}] \quad (5.16)$$

$$+ \left[\psi^A(\mu^{\chi_{t-1}(\varphi_t)}; \hat{\rho}_{t-2}, \mathbf{C}_t, \mathbf{\Lambda}_{t-1}) - \psi^A(\mu_{t-1}; \hat{\rho}_{t-2}, \mathbf{C}_t, \mathbf{\Lambda}_{t-1}) \right] \quad [\text{fusion}] \quad (5.17)$$

$$+ \left[\psi^A(\mu_{t-1}; \hat{\rho}_{t-2}, \mathbf{C}_t, \mathbf{\Lambda}_{t-1}) - \psi^A(\mu_{t-1}; \hat{\rho}_{t-2}, \mathbf{C}_{t,\text{del}}, \mathbf{\Lambda}_{t-1}) \right] \quad [\text{insertion}] \quad (5.18)$$

$$+ \left[\psi^A(\mu_{t-1}; \hat{\rho}_{t-2}, \mathbf{C}_{t,\text{del}}, \mathbf{\Lambda}_{t-1}) - \psi^A(\mu_{t-1}; \hat{\rho}_{t-2}, \mathbf{C}_{t,\text{del}}, \mathbf{\Lambda}_{t-2}) \right] \quad [\text{fission}] \quad (5.19)$$

$$+ \left[\psi^A(\mu_{t-1}; \hat{\rho}_{t-2}, \mathbf{C}_{t,\text{del}}, \mathbf{\Lambda}_{t-2}) - \psi^A(\mu_{t-1}; \hat{\rho}_{t-2}, \mathbf{C}_{t-1}, \mathbf{\Lambda}_{t-2}) \right] \quad [\text{deletion}] \quad (5.20)$$

Observe that:

1. From [Lemma 5.4](#), we have (5.15) $\leq O(\frac{1}{\delta} \log k) W_{\mathbb{T}}^1(\mu_t, \mu^{\chi_{t-1}(\varphi_t)})$.
2. And [Lemma 5.2](#) gives (5.17) ≤ 0 .
3. From [Lemma 5.5](#), we conclude that (5.16) $\leq O(\eta(k/\delta) \log k) W_{\mathbb{T}}^1(\mu_{t-1}, \mu^{\chi_{t-2}(\varphi_{t-1})})$.

Deletion. If $\mathbf{I}_t^j = 1$, then some center $z_t^j \in C_{t-1}^j$ is possibly deleted in passing from C_{t-1}^j to $C_{t,\text{del}}^j$.

For each $x \in X$, let $z_{x,j} \in C_{t-1}^j$ denote some center for which $d_X(z_{x,j}, x) = d_X(C_{t-1}^j, x)$, and for $z \in C_{t-1}^j$, let $X_j(z) := \{x \in X : z_{x,j} = z\}$. Then for each $j \geq 1$, $\{X_j(z) : z \in C_{t-1}^j\}$ forms a partition of X , and we have:

$$\psi^A(\mu_{t-1}; \hat{\rho}_{t-2}, C_{t,\text{del}}, \Lambda_{t-2}) - \psi^A(\mu_{t-1}; \hat{\rho}_{t-2}, C_{t-1}, \Lambda_{t-2}) \leq \sum_{j \geq 1} \mathbb{1}_{\{\mathbf{I}_t^j=1\}} \nu_{t-1}(X_j(z_t^j)) \tau^{-j}.$$

Recall that $\nu_{t-1}(X) = k$. Since z_t^j is chosen uniformly at random from C_{t-1}^j and $|C_{t-1}^j| = K$ if a deletion takes place, it holds that

$$\mathbb{E}_{\Omega^{\text{del}}} [\psi^A(\mu_{t-1}; \hat{\rho}_{t-2}, C_{t,\text{del}}, \Lambda_{t-2}) - \psi^A(\mu_{t-1}; \hat{\rho}_{t-2}, C_{t-1}, \Lambda_{t-2}) \mid j_t^*] \leq \frac{2\tau k}{K} \tau^{-j_t^*} \mathbb{1}_{\{j_t^* > 0\}}. \quad (5.21)$$

where we take expectation only over the random choice of which cluster to delete, and we recall the definition of j_t^* from (5.12).

Insertion. We now analyze the effect of inserting σ_t . This is the most delicate part of the analysis of $\Delta_t \Psi_t^A$.

Lemma 5.6. *For every $t \geq 1$, it holds that*

$$(5.18) \leq -\frac{\tau^{-j_t^*-1}}{8} \mathbb{1}_{\{j_t^* > 0\}} + W_{\mathbb{T}}^1(\mu_t, \mu^{\chi_{t-1}(\varphi_t)}) + 2\eta \hat{\rho}_{t-2}(\sigma_t). \quad (5.22)$$

Proof. Fix $j \geq 1$. Suppose that $\mathbf{I}_t^j = 1$ and denote

$$\psi(x; C) := d_X^j(x, C^j) \cdot \tau^j \left(d_X^j(x, \Lambda_{t-1}^j) - 2\eta \hat{\rho}_{t-2}(x) - \frac{1}{2} \tau^{-j-1} \right)_+.$$

Consider some $x \in \mathcal{V}_{\mathbb{T}}$ and let $\hat{x} \in C_{t-1}^j$ be such that $d_X(x, C_{t-1}^j) = d_X(x, \hat{x})$. Since σ_t is inserted into C_t^j , it must hold that $d_X(\sigma_t, \hat{x}) > \tau^{-j-1}$ and $d_X(\sigma_t, \Lambda_{t-1}^j) > \tau^{-j-1}$. Therefore either:

1. $d_X(x, \sigma_t) \geq \frac{1}{4} \tau^{-j-1}$, or
2. $d_X(x, \sigma_t) \leq d_X(x, \hat{x}) - \frac{1}{2} \tau^{-j-1}$, and

$$d_X(x, \Lambda_{t-1}^j) \geq d_X(\sigma_t, \Lambda_{t-1}^j) - d_X(x, \sigma_t) \geq \frac{3}{4} \tau^{-j-1}.$$

In either case, we can conclude that for any $x \in X$,

$$\psi(x; C_t^j) - \psi(x; C_{t,\text{del}}^j) \leq -\frac{1}{8} \tau^{-j-1} + d_X(x, \sigma_t) + 2\eta \hat{\rho}_{t-2}(x). \quad (5.23)$$

The value $\psi(x; C_t^j) - \psi(x; C_{t,\text{del}}^j)$ is never positive (since $C_{t,\text{del}}^j \subseteq C_t^j$), so we will only use the contribution (5.23) for certain $x \in X$.

By using (5.23) for the mass moving to σ_t in the passage from $\nu_{t-1} \rightarrow \nu_t$, we can write for every $j \geq 1$:

$$\psi^A(\mu_{t-1}; \rho_{t-1}, C_t, \Lambda_{t-1}) - \psi^A(\mu_{t-1}; \rho_{t-1}, C_{t,\text{del}}, \Lambda_{t-1})$$

$$\begin{aligned}
&\leq \left(-\frac{\tau^{-j-1}}{8} + W_X^1(v_t, v_{t-1}) + 2\eta\hat{\rho}_{t-2}(\sigma_t) \right) \mathbb{1}_{\{l_t^j=1\}} \\
&\leq \left(-\frac{\tau^{-j-1}}{8} + W_{\mathbb{T}}^1(\mu_t, \mu^{\chi_{t-1}(\varphi_t)}) + 2\eta\hat{\rho}_{t-2}(\sigma_t) \right) \mathbb{1}_{\{l_t^j=1\}}.
\end{aligned}$$

Summing over $j \geq 1$ yields

$$(5.18) \leq -\frac{\tau^{-j_t^*-1}}{8} \mathbb{1}_{\{j_t^*>0\}} + W_{\mathbb{T}}^1(\mu_t, \mu^{\chi_{t-1}(\varphi_t)}) + 2\eta\hat{\rho}_{t-2}(\sigma_t),$$

completing the proof. \square

We encapsulate the contents of this section in the following lemma. Define $\Delta\Psi^A := \sum_{t \geq 0} \Delta_t \Psi_t^A$. Write $\Delta_t^{\text{fis}} \Psi_t^A$ for (5.19).

Lemma 5.7. *There is a constant $C_A \geq 1$ such that if $K \geq 32\tau k$, then*

$$\mathbb{E} [\Delta\Psi^A] \leq -\frac{1}{16} \sum_{t \geq 1} \mathbb{E} \left[\tau^{-j_t^*} \mathbb{1}_{\{j_t^*>0\}} \right] + C_A \left(\frac{1}{\delta} \log k \right) \mathbb{E} [\text{cost}_{\mathbb{T}}^F(\mu)] + \mathbb{E} \left[\sum_{t \geq 1} \Delta_t^{\text{fis}} \Psi_t^A \right]$$

Proof. Sum (5.21) and (5.22), along with the bounds (1)–(3) derived at the beginning of the section, and apply Corollary 4.5 to bound the sum over $\hat{\rho}_{t-2}(\sigma_t)$. Finally, use the fact that $\eta \leq \frac{1}{32k}$. \square

5.3 The fission potential

The fission potential is central to our approach; it allows us to charge the change in $\Delta_t^{\text{fis}} \Phi_t$ due to breaking previously fused clusters against the movement of μ . Recall (3.4) and denote

$$\mathcal{H}_t^j := \mathcal{H}(\hat{p}_t^j, \Lambda_{t-1}^j, 2\tau^{-j-1}).$$

Observe that $[\mathcal{H}_t^j] \subseteq [\hat{Q}_t^j]$ is the subset of points that participate in a fused cluster in \hat{Q}_t^j .

Given $\mu \in \mathbb{M}(\mathcal{L}_{\mathbb{T}})$, a sequence $\hat{\mathcal{P}} = \langle \hat{p}^j : j \geq 1 \rangle$ of semi-partitions of X , and $\Lambda = \langle \Lambda^j \subseteq X : j \geq 1 \rangle$ a sequence of finite subsets, define:

$$\begin{aligned}
\psi^F(\mu; \hat{\mathcal{P}}, \Lambda) &:= - \sum_{j \geq 1} \tau^{-j} \beta \# \mu \left(\mathcal{H}(\hat{p}^j, \Lambda^j, 2\tau^{-j-1}) \right) \\
\Psi^F &:= \psi^F(\mu_t; \hat{\mathcal{P}}_t, \Lambda_{t-1}) \\
&= - \sum_{j \geq 1} \tau^{-j} \beta \# \mu \left(\mathcal{H}_t^j \right).
\end{aligned}$$

Remark 5.8 (Fission potential). The Ψ_t^F potential rewards us for fusing a cluster that contains significant ν_t mass. This will pay for the adverse effects of fission on the Φ potential as long as when we unfuse clusters, we are always doing it in order to fuse new clusters with much greater mass. This is why we fuse near the centers of heavy balls (which triggers a fission in the “light” annuli around the heavy ball).

As in Lemma 5.2, the proof of the next lemma follows from $\beta \# \mu = \beta \# \varphi \# \mu$ for any fusion map φ .

Lemma 5.9. *If φ is a fusion map, then*

$$\psi^F(\varphi\#\mu; \hat{\mathcal{P}}, \Lambda) = \psi^F(\mu; \hat{\mathcal{P}}, \Lambda).$$

Lemma 5.10. *The map $\mu \mapsto \psi^F(\mu; \hat{\mathcal{P}}, \Lambda)$ is 2-Lipschitz on $(\mathbb{M}(\mathcal{L}_\top), W_\top^1)$.*

Proof. Consider $\mu' = \mu + s(\mathbb{1}_\ell - \mathbb{1}_{\ell'})$ for some $s \in \mathbb{R}$ and $\ell, \ell' \in \mathcal{L}_\top$ with $\text{dist}_\top(\ell, \ell') = \tau^{-j}$. Manifestly:
 $\left| \psi^F(\mu; \hat{\mathcal{P}}, \Lambda) - \psi^F(\mu'; \hat{\mathcal{P}}, \Lambda) \right| \leq |s| \sum_{i \geq j} \tau^{-i} \leq 2|s| \tau^{-j}$. \square

5.3.1 Analysis

For $t \geq 1$, define $\Delta_t \Psi_t^F := \Psi_t^F - \Psi_{t-1}^F$. We decompose:

$$\begin{aligned} \Delta_t \Psi_t^F &:= \psi^F(\mu_t; \hat{\mathcal{P}}_t, \Lambda_{t-1}) - \psi^F(\mu_{t-1}, \hat{\mathcal{P}}_{t-1}, \Lambda_{t-2}) \\ &= \left[\psi^F(\mu_t; \hat{\mathcal{P}}_t, \Lambda_{t-1}) - \psi^F(\mu^{\chi_{t-1}(\varphi_t)}; \hat{\mathcal{P}}_t, \Lambda_{t-1}) \right] && [\mu \text{ movement}] \end{aligned} \quad (5.24)$$

$$+ \left[\psi^F(\mu^{\chi_{t-1}(\varphi_t)}; \hat{\mathcal{P}}_t, \Lambda_{t-1}) - \psi^F(\mu_{t-1}; \hat{\mathcal{P}}_t, \Lambda_{t-1}) \right] \quad [\text{fusion}] \quad (5.25)$$

$$+ \left[\psi^F(\mu_{t-1}; \hat{\mathcal{P}}_t, \Lambda_{t-1}) - \psi^F(\mu_{t-1}; \hat{\mathcal{P}}_t, \Lambda_{t-2}) \right] \quad [\text{heavy net update}] \quad (5.26)$$

$$+ \left[\psi^F(\mu_{t-1}; \hat{\mathcal{P}}_t, \Lambda_{t-2}) - \psi^F(\mu_{t-1}; \hat{\mathcal{P}}_{\text{del}, t}, \Lambda_{t-2}) \right] \quad [\text{insertion}] \quad (5.27)$$

$$+ \left[\psi^F(\mu_{t-1}; \hat{\mathcal{P}}_{\text{del}, t}, \Lambda_{t-2}) - \psi^F(\mu_{t-1}, \hat{\mathcal{P}}_{t-1}, \Lambda_{t-2}) \right]. \quad [\text{deletion}] \quad (5.28)$$

Observe that:

1. [Lemma 5.10](#) yields (5.24) $\leq 2W_\top^1(\mu_t, \mu^{\chi_{t-1}(\varphi_t)})$.
2. From [Lemma 5.9](#), we conclude that (5.25) ≤ 0 .
3. Moreover, (5.27) ≤ 0 because insertion can only enlarge the set of points that participate in a fused cluster.

Define $\Delta \Psi^F := \sum_{t \geq 1} \Delta_t \Psi_t^F$ and $\Delta_t^{\text{fis}} \Psi_t^F$ to be the expression in (5.26).

Lemma 5.11. *It holds that*

$$\mathbb{E} [\Delta \Psi^F] \leq 2 \mathbb{E} [\text{cost}_\top^F(\mu)] + \frac{2\tau k}{K} \sum_{t \geq 1} \mathbb{E} \left[\tau^{-j_t^*} \mathbb{1}_{\{j_t^* > 0\}} \right] + \sum_{t \geq 1} \mathbb{E} [\Delta_t^{\text{fis}} \Psi_t^F]$$

Proof. We are left to analyze (5.28). This argument is very similar to the deletion analysis in [Section 5.1.1](#) and [Section 5.2.1](#). If we delete a level- j cluster in moving from \hat{P}_{t-1}^j to $\hat{P}_{t, \text{del}}^j$, then the expected potential change (over the random choice of which cluster to delete) is $\tau^{-j} \frac{k}{K}$. Summing yields the desired bound. \square

5.3.2 Fusion and fission

Our final task is to analyze the quantities $\Delta_t^{\text{fis}} \Phi_t$, $\Delta_t^{\text{fis}} \Psi_t^A$, and $\Delta_t^{\text{fis}} \Psi_t^F$.




Figure 4: New heavy balls $B_1, B_2 \in \mathcal{B}_{\mathcal{V}}^j$ are responsible for ejecting some of the previously heavy balls in $\mathcal{B}_{\mathcal{U}}^j$.

Lemma 5.12. For every $t \geq 1$ and numbers $0 < c_A, c_F < 1$, if

$$\delta \leq \frac{c_F}{4(f_4(k) + c_A)},$$

then

$$\mathbb{E} \left[\Delta_t^{\text{fis}} \left(\Phi_t + c_F \Psi_t^F + c_A \Psi_t^A \right) \right] \leq 0.$$

Toward that end, fix $t \geq 2$. Let $\mathcal{U}^j := \Lambda_{t-2}^j \setminus \Lambda_{t-1}^j$ denote the set of heavy net points that are ejected in the “heavy net maintenance” phase of time step $t - 1$. Let $\mathcal{V}^j := \Lambda_{t-1}^j \setminus \Lambda_{t-2}^j$. Every $u \in \mathcal{U}^j$ is ejected because of some newly added point $\hat{u} \in \mathcal{V}^j$ with $d_X(u, \hat{u}) \leq \frac{\sqrt{\lambda}}{3} \tau^{-j}$. Denote

$$\begin{aligned} \mathcal{B}_{\mathcal{U}}^j &:= \{B_X(u, \tau^{-j}) : u \in \mathcal{U}^j\}, \\ \mathcal{B}_{\mathcal{V}}^j &:= \left\{ B_X \left(v, \frac{\tau^{-j}}{2\sqrt{\lambda}} \right) : v \in \mathcal{V}^j \right\}. \end{aligned}$$

Note that since Λ_{t-1}^j and Λ_{t-2}^j are heavy τ^{-j} -nets (cf. Lemma 4.1(1)), they are $3\tau^{-j}$ -separated, and thus the balls in $\mathcal{B}_{\mathcal{U}}^j$ are pairwise disjoint, and the same holds for the balls in $\mathcal{B}_{\mathcal{V}}^j$. See Figure 4.

Since $\lambda \geq 36$, we have $B_X(u, \tau^{-j}) \subseteq B_X(\hat{u}, \lambda \frac{\tau^{-j}}{2\sqrt{\lambda}})$ for each $u \in \mathcal{U}^j$. Therefore:

$$\bigcup_{B \in \mathcal{B}_{\mathcal{U}}^j} B \subseteq \bigcup_{B \in \mathcal{B}_{\mathcal{V}}^j} (\lambda B \setminus B). \quad (5.29)$$

For $j \in \mathbb{Z}_+$, define

$$S_{\text{out}}^j := \bigcup_{B \in \mathcal{B}_{\mathcal{V}}^j} (\lambda B \setminus B),$$

$$S_{\text{in}}^j := \bigcup_{B \in \mathcal{B}_{\mathcal{V}}^j} B.$$

The next three lemmas will yield the proof of [Lemma 5.12](#).

Lemma 5.13. *For every $t \geq 2$:*

$$\Delta_t^{\text{fis}} \Psi_t^A \leq \sum_{j \geq 1} \tau^{-j} \nu_{t-1} \left(S_{\text{out}}^j \right).$$

Proof. By definition:

$$\begin{aligned} \Delta_t^{\text{fis}} \Psi_t^A &= \psi^A(\mu_{t-1}; \hat{\rho}_{t-1}, \mathbf{C}_{t,\text{del}}, \Lambda_{t-1}) - \psi^A(\mu_{t-1}; \hat{\rho}_{t-1}, \mathbf{C}_{t,\text{del}}, \Lambda_{t-2}) \\ &= \psi^A(\mu_{t-1}; \hat{\rho}_{t-1}, \mathbf{C}_{t,\text{del}}, \Lambda_{t-1}) - \psi^A(\mu_{t-1}; \hat{\rho}_{t-1}, \mathbf{C}_{t,\text{del}}, \Lambda_{t-2} \setminus \mathcal{U}^j) \\ &\quad + \psi^A(\mu_{t-1}; \hat{\rho}_{t-1}, \mathbf{C}_{t,\text{del}}, \Lambda_{t-2} \setminus \mathcal{U}^j) - \psi^A(\mu_{t-1}; \hat{\rho}_{t-1}, \mathbf{C}_{t,\text{del}}, \Lambda_{t-2}). \end{aligned}$$

Since the first term involves the addition of points in \mathcal{V}^j , it is non-positive. Thus we focus on the second term.

In order for the x term in Ψ_t^A to be affected by the change from Λ_{t-2} to $\Lambda_{t-2} \setminus \mathcal{U}^j$, it must be that $d_X(x, \mathcal{U}^j) \leq \frac{1}{2} \tau^{-j}$, therefore

$$\Delta_t^{\text{fis}} \Psi_t^A \leq \sum_{j \geq 1} \tau^{-j} \nu_{t-1} \left(B_X(\mathcal{U}^j, \tau^{-j}) \right).$$

In conjunction with [\(5.29\)](#), this completes the proof. \square

The next lemma is the primary way that Axiom (A4) is employed.

Lemma 5.14. *For every $t \geq 1$:*

$$\Delta_t^{\text{fis}} \Phi_t \leq f_4(k) \sum_{j \geq 1} \tau^{-j} \nu_{t-1} \left(S_{\text{out}}^j \right).$$

Proof. By definition

$$\Delta_t^{\text{fis}} \Phi_t := \Phi(\alpha_t^{\text{fis}} \# \nu_{t-1}^*; \chi_{t-1}) - \Phi(\alpha_t^{\text{del}} \# \nu_{t-1}^*; \chi_{t-1}).$$

Observe that for each $j \geq 1$, the change from $\hat{\mathcal{Q}}_{\text{del},t}$ to $\hat{\mathcal{Q}}_{\text{fis},t}$ (which induces the change from α_t^{del} to α_t^{fis}) results from “unfusing” along the points of \mathcal{U}^j . Since $\hat{\mathcal{Q}}_{t,\text{del}}^j$ and $\hat{\mathcal{Q}}_{t,\text{fis}}^j$ induce the same semi-partition on $X \setminus \left[\mathcal{B}_{\mathcal{U}}^j \right]$, Axiom (A4) in conjunction with [\(5.29\)](#) yields the desired result. \square

The final lemma is key: The introduction of a new heavy ball yields a large decrease in potential.

Lemma 5.15. *For every $t \geq 2$:*

$$\Delta_t^{\text{fis}} \Psi_t^F \leq \sum_{j \geq 1} \tau^{-j} \left(\nu_{t-1}(S_{\text{out}}^j) - \nu_{t-1}(S_{\text{in}}^j) \right).$$

Proof. Recall that

$$\Delta_t^{\text{fis}} \Psi_t^F = \psi^F(\mu_{t-1}; \hat{\mathcal{P}}_t, \Lambda_{t-1}) - \psi^F(\mu_{t-1}; \hat{\mathcal{P}}_t, \Lambda_{t-2}),$$

where

$$\psi_t^F(\mu_{t-1}; \hat{\mathcal{P}}_t, \Lambda) = - \sum_{j \geq 1} \tau^{-j} \beta_{\#} \mu_{t-1} \left(\mathcal{H}(\hat{\mathcal{P}}_t^j, \Lambda^j, 2\tau^{-j-1}) \right).$$

Each $B \in \mathcal{B}_{\mathcal{V}}^j$ contributes at most $-\tau^{-j} v_{t-1}(B)$ to the potential, while we gain at most $\tau^{-j} v_{t-1}(S_{\text{out}}^j)$. \square

Proof of Lemma 5.12. Combining the preceding three lemmas gives, for any $0 < c_A, c_F < 1$:

$$\Delta_t^{\text{fis}} \left(\Phi_t + c_F \Psi_t^F + c_A \Psi_t^A \right) \leq \sum_{j \geq 1} \sum_{B \in \mathcal{B}_{\mathcal{V}}^j} \tau^{-j} \left((c_F + f_4(k) + c_A) v_{t-1}(\lambda B \setminus B) - c_F v_{t-1}(B) \right). \quad (5.30)$$

Now observe that since $\mathcal{B}_{\mathcal{V}}^j$ consists of $(t-1)$ -heavy balls, it holds that for every $B \in \mathcal{B}_{\mathcal{V}}^j$,

$$\bar{v}_{t-1}(\lambda B \setminus B) \leq \delta(1 - \delta) \bar{v}_{t-1}(B) \leq 2\delta \bar{v}_{t-1}(B).$$

Therefore taking expectations in (5.30) yields

$$\mathbb{E} \left[\Delta_t^{\text{fis}} \left(\Phi_t + c_F \Psi_t^F + c_A \Psi_t^A \right) \right] \leq \sum_{j \geq 1} \tau^{-j} \sum_{B \in \mathcal{B}_{\mathcal{V}}^j} \bar{v}_{t-1}(B) [\delta(f_4(k) + c_A) - (1 - 2\delta)c]. \quad (5.31)$$

If we now choose

$$\delta \leq \frac{c_F}{4(f_4(k) + c_A)},$$

then (5.31) becomes at most zero, completing the proof of Lemma 5.12. \square

5.4 Competitive analysis

Let us now prove Theorem 2.6.

Proof of Theorem 2.6. Use Lemma 5.1, Lemma 5.7, Lemma 5.11, to write, for any numbers $0 < c_A, c_F < 1$:

$$\begin{aligned} \mathbb{E} \left[\Delta \Phi + c_F \Delta \Psi^F + c_A \Delta \Psi^A \right] &\leq C_{\Phi} f_1(k) \frac{\log k}{\delta} \log(K) \text{cost}_X(\nu^*) \\ &\quad + \sum_{t \geq 1} \left(\frac{2\tau f_1(k)}{K} + c_F \frac{2\tau k}{K} - \frac{c_A}{8} \right) \mathbb{E} \left[\tau^{-j_t^*} \mathbb{1}_{\{j_t^* > 0\}} \right] \\ &\quad + \left(2c_F + c_A c_A \frac{\log k}{\delta} - \frac{3}{4} \right) \mathbb{E}[\text{cost}_{\mathbb{V}}^F(\mu)] \\ &\quad + \sum_{t \geq 1} \mathbb{E} \left[\Delta_t^{\text{fis}} \left(\Psi_t + c_F \Psi_t^F + c_A \Psi_t^A \right) \right]. \end{aligned}$$

Choosing $c_F := 1/8$ and $\delta := \frac{c_F}{4f_4(k) + c_A}$ and employing Lemma 5.12 yields

$$\mathbb{E} \left[\Delta \Phi + c_F \Delta \Psi^F + c_A \Delta \Psi^A \right] \leq C_{\Phi} f_1(k) \frac{\log k}{\delta} \log(K) \text{cost}_X(\nu^*)$$

$$\begin{aligned}
& + \sum_{t \geq 1} \left(\frac{2\tau f_1(k)}{K} + \frac{\tau k}{4K} - \frac{c_A}{8} \right) \mathbb{E} \left[\tau^{-j_t^*} \mathbb{1}_{\{j_t^* > 0\}} \right] \\
& + \left(c_A C_A \frac{\log k}{\delta} - \frac{1}{2} \right) \mathbb{E}[\text{cost}_{\mathbb{T}}^F(\boldsymbol{\mu})].
\end{aligned}$$

Setting $K := \frac{2\tau k f_1(k)}{32c_A}$ then gives

$$\mathbb{E} \left[\Delta\Phi + c_F \Delta\Psi^F + c_A \Delta\Psi^F \right] \leq C_{\Phi} f_1(k) \frac{\log k}{\delta} \log(K) \text{cost}_X(\boldsymbol{v}^*) + \left(c_A C_A \frac{\log k}{\delta} - \frac{1}{2} \right) \mathbb{E}[\text{cost}_{\mathbb{T}}^F(\boldsymbol{\mu})].$$

Finally, choose $c_A := (8C_A f_4(k) \log k)^{-1}$, yielding

$$\begin{aligned}
\mathbb{E} \left[\Delta\Phi + c_F \Delta\Psi^F + c_A \Delta\Psi^F \right] & \leq C_{\Phi} f_1(k) \frac{\log k}{\delta} \log(K) \text{cost}_X(\boldsymbol{v}^*) - \frac{1}{4} \mathbb{E}[\text{cost}_{\mathbb{T}}^F(\boldsymbol{\mu})] \\
& \leq O\left(f_1(k) f_4(k) (\log k)^2\right) \text{cost}_X(\boldsymbol{v}^*) - \frac{1}{4} \mathbb{E}[\text{cost}_{\mathbb{T}}^F(\boldsymbol{\mu})],
\end{aligned}$$

i.e.,

$$\mathbb{E}[\text{cost}_{\mathbb{T}}^F(\boldsymbol{\mu})] \leq O\left(f_1(k) f_4(k) (\log k)^2\right) \text{cost}_X(\boldsymbol{v}^*) - 4 \mathbb{E} \left[\Delta\Phi + c_F \Delta\Psi^F + c_A \Delta\Psi^F \right].$$

Now observe that due to starting in an initial configuration with μ_0 concentrated at a single leaf $\ell_0 \in \mathcal{L}_{\mathbb{T}}$ and $C_0^j = \Lambda_0^j = \{\ell_0\}$, it is the case that $\Delta\Psi^F, \Delta\Psi^A \geq 0$ because Ψ_0^F and Ψ_0^A both take their minimum value. Moreover, $\Phi \geq 0$ by assumption, and thus $-\Delta\Phi \leq \Phi(v_0^*; \chi_0)$, yielding the desired conclusion. \square

5.5 Rounding under fusion

Consider a pair of siblings $\xi^A, \xi^B \in V_{\mathbb{T}}^j$ with $\mathfrak{b}(\xi^A) \subseteq \mathfrak{b}(\xi^B)$ and the canonical injection $\varphi_{\xi^A \hookrightarrow \xi^B}$. Using auxiliary labels $\{1, 2\}$ (say), one can encode this injection by a multistep process:

$$\begin{aligned}
& \left\langle \hat{\xi}_0, \dots, \hat{\xi}_{j-1}, (\mathfrak{b}(\xi^A); 1), (\xi_{j+1}; 1), (\xi_{j+2}; 1), \dots \right\rangle \\
& \mapsto \left\langle \hat{\xi}_0, \dots, \hat{\xi}_{j-1}, (\mathfrak{b}(\xi^B); 2), (\xi_{j+1}; 1), (\xi_{j+2}; 1), \dots \right\rangle \\
& \mapsto \left\langle \hat{\xi}_0, \dots, \hat{\xi}_{j-1}, (\mathfrak{b}(\xi^B); 2), (\xi_{j+1}; 2), (\xi_{j+2}; 1) \dots \right\rangle \\
& \mapsto \left\langle \hat{\xi}_0, \dots, \hat{\xi}_{j-1}, (\mathfrak{b}(\xi^B); 2), (\xi_{j+1}; 2), (\xi_{j+2}; 2) \dots \right\rangle.
\end{aligned}$$

The idea is that only one label is changed from 1 to 2 at every step. (At the end, such atomic steps can be used to restore the original labeling.)

The advantage of this perspective is that if one is trying to prove monotonicity of some quantity under fusion maps, it suffices to establish monotonicity for canonical injections, and thus to establish it for one step of the above process. This corresponds to first “fusing” A into B but still distinguishing the children of B from those of A , then recursively fusing the children of B into the children of A , and so on. We will refer to such a step as a *primitive fusion of ξ^A into ξ^B* .

This will be a useful way of thinking in the next section, as well as in [Section 6.3](#).

5.5.1 Online rounding

The authors of [BBMN15, §5.2] present an online algorithm to round a fractional k -server algorithm on a τ -HST (for $\tau > 5$) to a random integral k -server algorithm in a way that the expected cost increases by at most an $O(1)$ factor. Unfortunately, this does not quite suffice for us, as our model allows cluster fusion.

Theorem 5.16 (HST rounding under fusions). *Consider an \mathcal{F} -adapted sequence $\mu = \langle \mu_t \in \mathbb{M}_k(\mathcal{L}_\top) : t \geq 0 \rangle$. There exists a random \mathcal{F} -adapted sequence $\hat{\mu} = \langle \hat{\mu}_t \in \widehat{\mathbb{M}}_k(\mathcal{L}_\top) : t \geq 0 \rangle$ such that for every $t \geq 0$: With probability one, for every $\xi \in \mathcal{V}_\top$,*

$$\hat{\mu}_t(\mathcal{L}_\top(\xi)) \in \left\{ \lfloor \mu_t(\mathcal{L}_\top(\xi)) \rfloor, \lceil \mu_t(\mathcal{L}_\top(\xi)) \rceil \right\}. \quad (5.32)$$

Moreover:

$$\mathbb{E} [\text{cost}_\top^F(\hat{\mu})] \leq O(1) \text{cost}_\top^F(\mu).$$

Proof. In [BBMN15, §5.2], the authors give a procedure for online rounding of a fractional k -server algorithm on HSTs to a distribution over integral algorithms that only loses an $O(1)$ factor in the expected cost. The key property maintained is that the integral algorithm is supported on balanced configurations with respect to the fractional algorithm, i.e., that (5.32) holds for every $\xi \in \mathcal{V}_\top$.

In order to extend this to our model, we need to give a method for the primitive fusion of two clusters while maintaining the balance property, Suppose that $\hat{\mu}$ is a random integral k -server measure that satisfies, for two siblings $\xi^A, \xi^B \in \mathcal{V}_\top$ with $\mathfrak{b}(\xi^A) \subseteq \mathfrak{b}(\xi^B)$,

$$\begin{aligned} \mathbb{E}[\hat{\mu}(\mathcal{L}_\top(\xi^A))] &= \mu(\mathcal{L}_\top(\xi^B)) \\ \mathbb{E}[\hat{\mu}(\mathcal{L}_\top(\xi^B))] &= \mu(\mathcal{L}_\top(\xi^B)), \end{aligned}$$

and with probability one, $\hat{\mu}$ satisfies the balance conditions:

$$\begin{aligned} \hat{\mu}(\mathcal{L}_\top(\xi^A)) &\in [\lfloor \mu(\mathcal{L}_\top(\xi^A)) \rfloor, \lceil \mu(\mathcal{L}_\top(\xi^A)) \rceil] \\ \hat{\mu}(\mathcal{L}_\top(\xi^B)) &\in [\lfloor \mu(\mathcal{L}_\top(\xi^B)) \rfloor, \lceil \mu(\mathcal{L}_\top(\xi^B)) \rceil]. \end{aligned}$$

For simplicity, let us denote

$$\begin{aligned} \hat{\mu}_A &:= \hat{\mu}(\mathcal{L}_\top(\xi^A)) \\ \hat{\mu}_B &:= \hat{\mu}(\mathcal{L}_\top(\xi^B)) \\ \mu_A &:= \mu(\mathcal{L}_\top(\xi^A)) \\ \mu_B &:= \mu(\mathcal{L}_\top(\xi^B)) \\ \varepsilon_A &:= \mu_A - \lfloor \mu_A \rfloor \\ \varepsilon_B &:= \mu_B - \lfloor \mu_B \rfloor. \end{aligned}$$

We need to produce a random variable (k_A, k_B) with the following properties:

1. $\text{supp}((k_A, k_B)) \subseteq \text{supp}((\hat{\mu}_A, \hat{\mu}_B))$
2. $\mathbb{P}(k_A = \lfloor \mu_A \rfloor) = \mathbb{P}(\hat{\mu}_A = \lfloor \mu_A \rfloor)$

3. $\mathbb{P}(k_B = \lfloor \mu_B \rfloor) = \mathbb{P}(\hat{\mu}_B = \lfloor \mu_B \rfloor)$
4. The balance condition is satisfied:

$$\mathbb{P}(k_A + k_B \in [\lfloor \mu_A + \mu_B \rfloor, \lceil \mu_A + \mu_B \rceil]) = 1.$$

We then define $\hat{\mu}$ of the fused cluster as $k_A + k_B$ and couple the distributions of the children accordingly using the conditional distributions $\hat{\mu} \mid \hat{\mu}_A = k_A$ and $\hat{\mu} \mid \hat{\mu}_B = k_B$. In this way, we preserve a balanced online rounding under a primitive fusion step. Note that we do not incur any reduced movement cost because we do not pay for the fusion (by definition of the reduced cost).

There are two cases. Note that the first case includes the situation in which one of μ_A or μ_B is an integer.

1. $\varepsilon_A + \varepsilon_B \leq 1$:

$$\begin{aligned} \mathbb{P}[(k_A, k_B) = (\lfloor \mu_A \rfloor, \lfloor \mu_B \rfloor)] &= \mathbb{P}(\hat{\mu}_A = \lfloor \mu_A \rfloor) + \mathbb{P}(\hat{\mu}_B = \lfloor \mu_B \rfloor) - 1 \\ \mathbb{P}[(k_A, k_B) = (\lfloor \mu_A \rfloor, \lceil \mu_B \rceil)] &= \mathbb{P}(\hat{\mu}_B = \lceil \mu_B \rceil) \mathbb{1}_{\{\varepsilon_B > 0\}} \\ \mathbb{P}[(k_A, k_B) = (\lceil \mu_A \rceil, \lfloor \mu_B \rfloor)] &= \mathbb{P}(\hat{\mu}_A = \lceil \mu_A \rceil) \mathbb{1}_{\{\varepsilon_A > 0\}} \\ \mathbb{P}[(k_A, k_B) = (\lceil \mu_A \rceil, \lceil \mu_B \rceil)] &= 0. \end{aligned}$$

2. $\varepsilon_A + \varepsilon_B > 1$:

$$\begin{aligned} \mathbb{P}[(k_A, k_B) = (\lfloor \mu_A \rfloor, \lfloor \mu_B \rfloor)] &= 0 \\ \mathbb{P}[(k_A, k_B) = (\lfloor \mu_A \rfloor, \lceil \mu_B \rceil)] &= \mathbb{P}(\hat{\mu}_A = \lfloor \mu_A \rfloor) \\ \mathbb{P}[(k_A, k_B) = (\lceil \mu_A \rceil, \lfloor \mu_B \rfloor)] &= \mathbb{P}(\hat{\mu}_B = \lfloor \mu_B \rfloor) \\ \mathbb{P}[(k_A, k_B) = (\lceil \mu_A \rceil, \lceil \mu_B \rceil)] &= \mathbb{P}(\hat{\mu}_A = \lceil \mu_A \rceil) + \mathbb{P}(\hat{\mu}_B = \lceil \mu_B \rceil) - 1. \quad \square \end{aligned}$$

The final lemma of this section completes the proof of [Theorem 2.5](#) in conjunction with [Theorem 5.16](#).

Lemma 5.17. *If $\hat{\mu} = \langle \hat{\mu}_t \in \widehat{\mathbb{M}}_k(\mathcal{L}_{\mathbb{T}}) : t \geq 0 \rangle$ is a sequence of integral measures, and $\nu = \langle \nu_t : t \geq 0 \rangle$ is defined by $\nu_t := \beta \# \mu_t$, then*

$$\text{cost}_X(\nu) \leq \text{cost}_{\mathbb{T}}^F(\mu').$$

Proof. This follows immediately from the fact that β is 1-Lipschitz and if φ is a fusion map and $\mu \in \mathbb{M}(\mathcal{L}_{\mathbb{T}})$, then $\varphi \# \mu \in \mathbb{M}(\mathcal{L}_{\mathbb{T}})$, and $\beta \# \varphi \# \mu = \beta \# \mu$. \square

6 Reductions

We now present some generic reductions that allow us to assume a weaker set of potential axioms. In [Section 6.3](#), these are used to apply our framework to the [\[BCL⁺17\]](#) algorithm.

6.1 Mass at internal nodes

We first discuss algorithms that maintain fractional server mass at internal nodes $V_{\mathbb{T}}$ of \mathbb{T} and only to pay for movement of server mass down the tree. Such algorithms can be incorporated into our framework as follows.

We now allow the fractional server measure μ^χ associated to a configuration $\chi \in \Gamma$ to place mass on both leaves and internal vertices of \mathbb{T} , i.e., $\mu^\chi \in \mathbb{M}_k(\mathcal{V}_{\mathbb{T}})$.

Definition 6.1. For two measures $\mu, \mu' \in \mathbb{M}(\mathcal{V}_{\mathbb{T}})$, say that μ *dominates* μ' if the following are satisfied:

1. $\mu(\mathcal{V}_{\mathbb{T}}) = \mu'(\mathcal{V}_{\mathbb{T}})$.
2. $\mu(\mathcal{V}_{\mathbb{T}}(\xi)) \leq \mu'(\mathcal{V}_{\mathbb{T}}(\xi))$ for all $\xi \in V_{\mathbb{T}}$.

Let $\llbracket \mu \rrbracket$ denote the collection of all measures $\mu' \in \mathbb{M}(\mathcal{V}_{\mathbb{T}})$ such that μ dominates μ' .

Note that if $\mu' \in \llbracket \mu \rrbracket$, then μ' can be obtained from μ by pushing mass “down the tree” (recall [Remark 2.4](#)). It is helpful to observe that if μ is supported on $\mathcal{L}_{\mathbb{T}}$, then $\llbracket \mu \rrbracket = \{\mu\}$ (in intuitive terms, mass supported on $\mathcal{L}_{\mathbb{T}}$ cannot be “pushed down” any further). We now introduce related modifications of Axiom (A3).

(A3^o) For any fusion map φ and configuration $\chi \in \Gamma$, there is a configuration $\chi(\varphi) \in \Gamma$ such that $\mu^{\chi(\varphi)} \in \llbracket \varphi \# \mu^\chi \rrbracket$, and moreover

$$\Phi(\varphi \# \theta; \chi(\varphi)) \leq \Phi(\theta; \chi) \quad \forall \theta \in \widehat{\mathbb{M}}_k(\mathcal{L}_{\mathbb{T}}). \quad (6.1)$$

Theorem 6.2. For any metric space (X, d_X) , the following holds. If there are functions

$$\Phi : \widehat{\mathbb{M}}_k(\mathcal{L}_{\mathbb{T}}) \times \Gamma \rightarrow \mathbb{R}, \quad \gamma : \Gamma \times \mathcal{L}_{\mathbb{T}}^0 \rightarrow \Gamma, \quad \chi \mapsto \mu^\chi \in \mathbb{M}_k(\mathcal{V}_{\mathbb{T}})$$

satisfying (A1), (A2), (A3^o), and (A4) for some functions $f_1(k), f_4(k)$, then there is an augmented state space $\hat{\Gamma}$, and functions

$$\hat{\Phi} : \widehat{\mathbb{M}}_k(\mathcal{L}_{\mathbb{T}}) \times \hat{\Gamma} \rightarrow \mathbb{R}, \quad \hat{\gamma} : \hat{\Gamma} \times \mathcal{L}_{\mathbb{T}}^0 \rightarrow \hat{\Gamma}, \quad \chi \mapsto \tilde{\mu}^\chi \in \mathbb{M}_k(\mathcal{L}_{\mathbb{T}})$$

satisfying axioms (A1)–(A4) with the functions $2f_1(k)$ and $2f_4(k)$.

Proof. Given the sequence of measures $\mu = \langle \mu_t : t \geq 0 \rangle$ corresponding to the online algorithm induced by the maps γ and $\chi \mapsto \mu^\chi$, with $\mu_0 \in \mathbb{M}_k(\mathcal{L}_{\mathbb{T}})$ and $\mu_t \in \mathbb{M}_k(\mathcal{V}_{\mathbb{T}})$ for $t \geq 1$, one can consider the corresponding \mathcal{F} -adapted lazy sequence $\tilde{\mu} = \langle \tilde{\mu}_t : t \geq 0 \rangle$ with $\{\tilde{\mu}_t : t \geq 0\} \subseteq \mathbb{M}_k(\mathcal{L}_{\mathbb{T}})$ that only moves mass between leaves.

Our augmented configuration space will be $\hat{\Gamma} := \Gamma \times \mathbb{M}_k(\mathcal{L}_{\mathbb{T}})$. For $\mu \in \mathbb{M}(\mathcal{V}_{\mathbb{T}})$, define the average “height” of μ :

$$\Psi^H(\mu) := \sum_{j \geq 0} \tau^{-j} \mu(V_{\mathbb{T}}^j),$$

and for $(\chi, \tilde{\mu}) \in \hat{\Gamma}$, define the modified potential

$$\hat{\Phi}(\theta; (\chi, \tilde{\mu})) := 2\Phi(\theta; \chi) + \Psi^H(\mu^\chi) + W_{\mathbb{T}}^1(\mu^\chi, \tilde{\mu}). \quad (6.2)$$

The validity of Axioms (A1) and (A4) is unchanged since $\hat{\Phi}$ does not introduce an additional dependence on its first argument.

Consider that when μ transports mass, $\hat{\Phi}$ does not increase because of (A2). On the other hand, when $\tilde{\mu}$ moves, $\hat{\Phi}$ decreases in proportion because of the last term in (6.2), and thus (A2) is satisfied for $\tilde{\mu}^\chi$ as well. Moreover, under a “push down” operation occurring as in (A3^o), the last term in (6.2) may increase, but then Ψ^H will decrease $\hat{\Phi}$ at least as much, meaning that (A3) is satisfied for $(\hat{\Phi}, \hat{\gamma}, \chi \mapsto \tilde{\mu}^\chi)$. \square

6.2 Extra server mass

We will now show that for any $0 < \varepsilon < 1$, it suffices to have a fractional $(k + \varepsilon)$ -server algorithm satisfying Axioms (A1)–(A4). The idea of converting a $(k + \varepsilon)$ -server algorithm to a k -server algorithm storing mass at internal nodes is taken from [BCL⁺17].

Define the rounding map $\rho : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ as follows: For $h \in \mathbb{Z}_+$, define $\rho|_{[h, h+\varepsilon]} = h$ and extend ρ affinely outside $\bigcup_{h \in \mathbb{Z}_+} [h, h + \varepsilon]$. Note that ρ is $\frac{1}{1-\varepsilon}$ -Lipshitz. For a measure $\mu \in \mathbb{M}(\mathcal{L}_\mathbb{T})$, define the measure $\Lambda_\varepsilon \mu \in \mathbb{M}(\mathcal{V}_\mathbb{T})$ by

$$\Lambda_\varepsilon \mu(\xi) := \begin{cases} \rho(\mu(\xi)) & \xi \in \mathcal{L}_\mathbb{T} \\ \rho(\mu(\mathcal{L}_\mathbb{T}(\xi))) - \sum_{\xi' \in \text{ch}(\xi)} \rho(\mu(\mathcal{L}_\mathbb{T}(\xi'))) & \text{otherwise.} \end{cases}$$

Note that ρ is superadditive, i.e., $\rho(y + y') \geq \rho(y) + \rho(y')$ for all $y, y' \in \mathbb{R}_+$, so $\Lambda_\varepsilon \mu$ does define a measure. Moreover, by construction we have:

$$\Lambda_\varepsilon \mu(\mathcal{V}_\mathbb{T}(\xi)) = \rho(\mu(\mathcal{L}_\mathbb{T}(\xi))) \quad \forall \xi \in \mathcal{V}_\mathbb{T},$$

and therefore

$$\Lambda_\varepsilon \mu(\mathcal{V}_\mathbb{T}) = \rho(\mu(\mathcal{L}_\mathbb{T})) = \rho(k + \varepsilon) = k,$$

thus $\Lambda_\varepsilon \mu \in \mathbb{M}_k(\mathcal{V}_\mathbb{T})$.

The next lemma follows from the fact that ρ is superadditive.

Lemma 6.3. *For any $\nu \in \mathbb{M}_{k+\varepsilon}(\mathcal{L}_\mathbb{T})$, and fusion map φ , it holds that*

$$\Lambda_\varepsilon(\varphi \# \nu) \in \llbracket \varphi \# \Lambda_\varepsilon \nu \rrbracket.$$

Theorem 6.4. *For any metric space (X, d_X) and $0 < \varepsilon < 1$, the following holds. If there is a transition function $\gamma : \Gamma \times \mathcal{L}_\mathbb{T}^0 \rightarrow \Gamma$, a potential Φ , and a map $\chi \mapsto \mu^\chi \in \mathbb{M}_{k+\varepsilon}(\mathcal{L}_\mathbb{T})$ satisfying Axioms (A0)–(A4) for some functions $f_1(k), f_4(k)$, then replacing the map $\chi \mapsto \mu^\chi$ by the map $\chi \mapsto \Lambda_\varepsilon(\mu^\chi)$ yields an algorithm satisfying Axioms (A1), (A2), (A3^o), and (A4) with the new potential $\hat{\Phi} = \frac{2}{1-\varepsilon}\Phi$ and the functions $\frac{2}{1-\varepsilon}f_1(k), \frac{2}{1-\varepsilon}f_4(k)$.*

Proof. Axiom (A3) is satisfied because of Lemma 6.3. Thus if we establish that

$$W_\mathbb{T}^1(\Lambda_\varepsilon \mu, \Lambda_\varepsilon \mu') \leq \frac{1}{1-\varepsilon} W_\mathbb{T}^1(\mu, \mu') \quad \forall \mu, \mu' \in \mathbb{M}(\mathcal{L}_\mathbb{T}), \quad (6.3)$$

it will show Axiom (A2) is satisfied with the potential $\hat{\Phi} = \frac{2}{1-\varepsilon}\Phi$, completing the proof.

Consider $\mu \in \mathbb{M}(\mathcal{V}_{\mathbb{T}})$ and denote

$$\|\mu\|_{\mathbb{T}} := \sum_{j \geq 1} \tau^{-j} \sum_{\xi \in V_{\mathbb{T}}} \mu(\mathcal{V}_{\mathbb{T}}(\xi)).$$

Then,

$$\frac{1}{2} \|\mu - \mu'\|_{\mathbb{T}} \leq W_{\mathbb{T}}^1(\mu, \mu') \leq \|\mu - \mu'\|_{\mathbb{T}} \quad \forall \mu, \mu' \in \mathbb{M}(\mathcal{V}_{\mathbb{T}}).$$

Now we can write:

$$W_{\mathbb{T}}^1(\Lambda_{\varepsilon}\mu, \Lambda_{\varepsilon}\mu') \leq \|\Lambda_{\varepsilon}\mu - \Lambda_{\varepsilon}\mu'\|_{\mathbb{T}} \leq \|\rho\|_{\text{Lip}} \|\mu - \mu'\|_{\mathbb{T}} \leq \frac{2}{1-\varepsilon} W_{\mathbb{T}}^1(\mu, \mu') \quad \forall \mu, \mu' \in \mathbb{M}(\mathcal{L}_{\mathbb{T}}),$$

verifying (6.3). □

By composing [Theorem 6.4](#) and [Theorem 6.2](#), we obtain the following.

Corollary 6.5. *Under the assumptions of [Theorem 6.4](#), there is a triple $(\Phi, \gamma, \chi \mapsto \mu^x \in \mathbb{M}_k(\mathcal{L}_{\mathbb{T}}))$ satisfying Axioms (A1)–(A4) with $\hat{f}_1(k) \leq \frac{4}{1-\varepsilon} f_1(k)$ and $\hat{f}_4(k) \leq \frac{4}{1-\varepsilon} f_4(k)$.*

6.3 Verification of the potential axioms for [\[BCL⁺17\]](#)

In light of [Corollary 6.5](#), it will suffice to demonstrate a fractional $(k + \varepsilon)$ -server algorithm satisfying (A1)–(A4) for some $0 < \varepsilon < 1$.

Consider an element

$$x = \langle x^{\xi,i} \in [0, 1] : \xi \in V_{\mathbb{T}}, i = 1, 2, \dots \rangle \subseteq \ell^{\infty}(V_{\mathbb{T}} \times \mathbb{Z}_+).$$

For $\xi \in V_{\mathbb{T}}$, write $\text{ch}(\xi)$ for the set of children of ξ in \mathbb{T} . Let K denote the closed convex set of such x that satisfy the following linear constraints for every $\xi \in V_{\mathbb{T}}$:

$$\begin{aligned} x^{X,i} &= \begin{cases} 0 & i \in \{1, 2, \dots, k\} \\ 1 & i > k, \end{cases} \\ \sum_{i \leq |S|} x^{\xi,i} &\leq \sum_{(\xi',j) \in S} x^{\xi',j} \quad \text{for all finite } S \subseteq \text{ch}(\xi) \times \mathbb{Z}_+. \end{aligned} \quad (6.4)$$

Let us furthermore define $z = z(x)$ by $z^{\xi,i} := \frac{1}{1-\delta} (1 - x^{\xi,i})$ and $z^{\xi} := \sum_{i \geq 1} z^{\xi,i}$. (Note that x and z are related by an invertible linear transformation, and thus we need only specify one of them in order to define the corresponding set of values.)

For a leaf $\ell = \langle \xi_0, \xi_1, \dots \rangle \in \mathcal{L}_{\mathbb{T}}$, we write

$$z^{\ell} := \lim_{j \rightarrow \infty} z_t^{\xi_j}.$$

Fix $\delta := \frac{1}{3k}$ and let $K_{\delta} \subseteq K$ denote the subset of $x \in K$ for which the set $\{\ell : z^{\ell} \neq 0\}$ is finite, as well as the sets $\{i : z^{\xi,i} \neq 0\}$ for each $\xi \in V_{\mathbb{T}}$, and furthermore:

$$z^{\ell} \leq 1 \quad \forall \ell \in \mathcal{L}_{\mathbb{T}}, \quad (6.5)$$

$$\sum_{\ell \in \mathcal{L}_{\mathbb{T}}} z^{\ell} = \frac{k}{1-\delta} = k + \varepsilon, \quad (6.6)$$

where we note that $\varepsilon := \frac{\delta k}{1-\delta} < 1$ for all $k \geq 1$.

Define the measure $\mu^x \in \mathbb{M}_{k+\varepsilon}(\mathcal{L}_{\mathbb{T}})$ by

$$\mu^x(S) := \sum_{\ell \in S} z^\ell \quad \forall S \subseteq \mathcal{L}_{\mathbb{T}}. \quad (6.7)$$

One should note that for $x \in K_\delta$, the inequalities (6.4) imply that for every $\xi \in V_{\mathbb{T}}$,

$$z^\xi \geq \sum_{\xi' \in \text{ch}(\xi)} z^{\xi'}, \quad (6.8)$$

and since $z^X = k + \varepsilon$, (6.6) implies that the inequality in (6.8) holds with equality. In other words, for every $\xi \in V_{\mathbb{T}}$, we have

$$z^\xi = \mu^x(\mathcal{L}_{\mathbb{T}}(\xi)). \quad (6.9)$$

The [BCL⁺17] algorithm. To each $\theta \in \widehat{\mathbb{M}}_k(\mathcal{L}_{\mathbb{T}})$, we associate a representation \hat{x}_θ as follows: For every $\xi \in V_{\mathbb{T}}$,

$$\hat{z}_\theta^\xi := \sum_{\ell \in \mathcal{L}_{\mathbb{T}}(\xi)} \hat{\theta}(\ell),$$

and for $\xi \in V_{\mathbb{T}}$ and $i \geq 1$:

$$\hat{x}_\theta^{\xi,i} = \begin{cases} 0 & \hat{z}_\theta^\xi \geq i \\ 1 & \text{otherwise.} \end{cases}$$

Let $\Gamma := K_\delta$, and define the potential:

$$\Phi(\theta; x) := C_0 D(\theta; x) - H(x),$$

where $C_0 \asymp \log k$, and

$$D(\theta; x) := \sum_{j \geq 1} \left(\tau^{-j} \sum_{\xi \in V_{\mathbb{T}}^j} \sum_{i \geq 1} (\hat{x}_\theta^{\xi,i} + \delta) \log \left(\frac{\hat{x}_\theta^{\xi,i} + \delta}{x^{\xi,i} + \delta} \right) \right),$$

$$H(x) := \sum_{j \geq 1} \tau^{-j} \sum_{\xi \in V_{\mathbb{T}}^j} \left[(z^\xi + (1 + \tau^{-1})\varepsilon) \log \frac{z^\xi + \varepsilon}{\varepsilon} + z^\xi \log(z^\xi + \varepsilon) \right].$$

and $\hat{\xi}$ denotes the parent of ξ in \mathbb{T} . One should note that this sum converges absolutely because the sets $\{\xi \in V_{\mathbb{T}}^j : z^\xi > 0\}$ are finite for every $j \geq 0$, and moreover z forms a measure of weight $k + \varepsilon$ at every level.

The [BCL⁺17] algorithm can be interpreted as a mapping $\gamma : \Gamma \times \mathcal{L}_{\mathbb{T}}^0 \rightarrow \Gamma$ that satisfies axioms (A1) and (A2).

Theorem 6.6 ([BCL⁺17]). *There is a mapping $\gamma : \Gamma \times \mathcal{L}_{\mathbb{T}}^0 \rightarrow \Gamma$ and such that the following hold for every $x \in \Gamma$.*

1. For any two states $\theta, \theta' \in \widehat{\mathbb{M}}_k(\mathcal{L}_{\mathbb{T}})$:

$$|\Phi(\theta; x) - \Phi(\theta'; x)| \leq O(\log k)^2 W_{\mathbb{T}}^1(\theta, \theta').$$

2. For every $\sigma \in \mathcal{L}_{\mathbb{T}}^0$, we have $\mu^{\gamma(x,\sigma)}(\sigma) \geq 1$.
3. For every $\sigma \in \mathcal{L}_{\mathbb{T}}^0$ and every integral measure $\theta \in \widehat{\mathbb{M}}_k(\mathcal{L}_{\mathbb{T}})$ satisfying $\theta(\sigma) \geq 1$:

$$\Phi(\theta; \gamma(x, \sigma)) - \Phi(\theta; x) \leq -W_{\mathbb{T}}^1(\mu^x, \mu^{\gamma(x,\sigma)}).$$

Moreover, the associated measures lie in the 0-decorated subtree: $\{\mu^x : x \in K_{\delta}\} \subseteq \mathbb{M}_k(\mathcal{L}_{\mathbb{T}}^0)$.

We are thus left to verify Axioms (A3) and (A4).

Axiom (A3). In order to demonstrate the validity of (A3), we need to give a way of updating the z -variables under a primitive fusion of ξ^A into ξ^B , where $\xi^A, \xi^B \in V_{\mathbb{T}}^j$ are siblings in \mathbb{T} with $\mathfrak{b}(A) \subseteq \mathfrak{b}(B)$. We will use \bar{z} to denote the variables after the fusion.

For any descendant ξ of ξ^A (including ξ^A itself), set $\bar{z}^{\xi,i} := 0$ for all $i \geq 1$. Let ξ' denote the application of a primitive fusion step to ξ (so that ξ' is a descendant of ξ^B). If $\xi' \neq \xi^B$, we set $\bar{z}^{\xi',i} := z^{\xi,i}$ for all $i \geq 1$. We now specify how to update the variables $\{z^{\xi^B,i} : i \geq 1\}$. All other variables remain unchanged.

Define the sequence $\langle \bar{z}^{\xi^B,1}, \bar{z}^{\xi^B,2}, \dots \rangle$ by sorting, in non-increasing order, the concatenation of the two sequences

$$\langle z^{\xi^B,i} : i \geq 1 \rangle, \quad \langle z^{\xi^A,i} : i \geq 1 \rangle. \quad (6.10)$$

(Recall that since $x \in K_{\delta}$, each such sequence has only finitely many non-zero values.)

Lemma 6.7. *It holds that $\bar{x} \in K_{\delta}$ and $\mu^{\bar{x}} = \varphi \# \mu^x$, where φ denotes the corresponding primitive fusion map. Furthermore for any $\theta \in \widehat{\mathbb{M}}_k(\mathcal{L}_{\mathbb{T}})$,*

$$\Phi(\varphi \# \theta; \bar{x}) \leq \Phi(\theta; x). \quad (6.11)$$

Proof. The fact that $\mu^{\bar{x}} = \varphi \# \mu^x$ is immediate from the construction. And from this, it follows that (6.6) holds. Thus we need only verify that $\bar{x} \in K$, and only for the third set of inequalities (6.4) is this slightly non-trivial.

By construction, it is straightforward that those inequalities hold for \bar{x} for any $\xi \in V_{\mathbb{T}}$ except $\xi = \xi^B$. For the parent $\hat{\xi}$ of ξ^A and ξ^B , the fact that we have merged the two child lists means that the inequalities (6.4) continue to hold for $\hat{\xi}$.

Thus we need only verify the inequalities for ξ^B . Note that one can rewrite the inequality in (6.4) as

$$\sum_{i=1}^{|S|} \bar{z}^{\xi^B,i} \geq \sum_{(\xi',j) \in S} \bar{z}^{\xi',j} \quad \forall \text{ finite } S \subseteq \text{ch}(\xi^B) \times \mathbb{Z}_+.$$

Since these inequalities hold for z , they also hold for \bar{z} because when ξ^A is fused into ξ^B , we sort the corresponding list of values in decreasing order.

Let us now prove (6.11). The potential $\Phi(\theta; x)$ is a sum of two expressions; the first depends on θ , whereas the second does not. To see that $H(\bar{x}) \geq H(x)$, apply the next lemma with $a = z^{\xi^A}, b = z^{\xi^B}, c = \tau^{-1}$.

Lemma 6.8. *For any numbers $a, b, c \geq 0$ and $0 \leq \varepsilon \leq 1$ such that $(1+c)\varepsilon \leq 1$, it holds that*

$$(a + b + (1+c)\varepsilon) \log(a + b + \varepsilon) \geq (a + (1+c)\varepsilon) \log(a + \varepsilon) + (b + (1+c)\varepsilon) \log(b + \varepsilon).$$

Proof. Without loss of generality, we may assume that $a \geq b$. Define

$$f(t) := (a + t + (1 + c)\varepsilon) \log(a + t + (1 + c)\varepsilon) + (b - t + (1 + c)\varepsilon) \log(b - t + (1 + c)\varepsilon),$$

and compute

$$\left. \frac{d}{dt} \right|_{t=0} f(t) = \log \frac{a + (1 + c)\varepsilon}{b + (1 + c)\varepsilon} \geq 0.$$

We conclude that

$$\begin{aligned} & (a + b + (1 + c)\varepsilon) \log(a + b + (1 + c)\varepsilon) + (1 + c)\varepsilon \log[(1 + c)\varepsilon] \\ & \geq (a + (1 + c)\varepsilon) \log(a + (1 + c)\varepsilon) + (b + (1 + c)\varepsilon) \log(b + (1 + c)\varepsilon). \end{aligned}$$

Since $(1 + c)\varepsilon \leq 1$ by assumption, this yields

$$\begin{aligned} & (a + b + (1 + c)\varepsilon) \log(a + b + (1 + c)\varepsilon) \\ & \geq (a + (1 + c)\varepsilon) \log(a + (1 + c)\varepsilon) + (b + (1 + c)\varepsilon) \log(b + (1 + c)\varepsilon). \end{aligned} \quad (6.12)$$

Observe also that

$$\begin{aligned} & (a + b + (1 + c)\varepsilon) \log\left(1 - \frac{c\varepsilon}{a + b + (1 + c)\varepsilon}\right) \\ & \geq (a + (1 + c)\varepsilon) \log\left(1 - \frac{c\varepsilon}{a + (1 + c)\varepsilon}\right) + (b + (1 + c)\varepsilon) \log\left(1 - \frac{c\varepsilon}{b + (1 + c)\varepsilon}\right). \end{aligned} \quad (6.13)$$

Adding (6.12) and (6.13) yields the desired result. \square

We now address $D(\theta; \mathbf{x})$. The only terms that change are the ones corresponding to ξ^A and ξ^B . Let $\theta' = \varphi\#\theta$. By construction:

$$\hat{x}_{\theta'}^{\xi,1} \leq \hat{x}_{\theta'}^{\xi,2} \leq \dots$$

Since we sort both $\hat{x}_{\theta'}$ and \bar{x} in increasing order, the value of the $D(\theta; \mathbf{x})$ decreases, verifying the claim. \square

Axiom (A4). Clearly the change $\theta \mapsto \theta'$ with $\theta, \theta' \in \widehat{\mathbb{M}}_k(\mathcal{L}_{\mathbb{T}})$ does not change H , so we need only analyze the first part of Φ . Consider $\xi^0 \in V_{\mathbb{T}}^{h-1}$ and a child $\xi^1 \in V_{\mathbb{T}}^h$. If $\theta' = F\#\theta$ where $F(\xi) = \xi$ for $\xi \notin \mathcal{V}_{\mathbb{T}}(\xi^1)$ and $F(\mathcal{V}_{\mathbb{T}}(\xi^1)) \subseteq \mathcal{V}_{\mathbb{T}}(\xi^0)$, then the value of Φ can change by at most

$$C_0 \left| \sum_{j \geq h} \tau^{-j} \sum_{i \geq 1} \sum_{\xi_2 \in V_{\mathbb{T}}^i: \xi_2 \leq \xi_1} \log(x^{\xi_2, i} + \delta) \right|, \quad (6.14)$$

where we have used the notation $\xi_2 \leq \xi_1$ to denote that ξ_2 is a descendant of ξ_1 (and we say that ξ_1 is a descendant of itself). The desired conclusion follows from the next fact.

Fact 6.9. For every $x \in [0, 1]$ and $\delta \in [0, \frac{1}{2}]$:

$$\log \frac{1 + \delta}{x + \delta} \leq \frac{1 - x}{1 - \delta} \log \frac{1}{\delta}$$

Using this and recalling that $\delta = \frac{1}{3k}$, (6.14) is bounded by

$$O(C_0 \log k) \sum_{j \geq h} \tau^{-j} \sum_{i \geq 1} \sum_{\xi_2 \in V_{\mathbb{T}}^j: \xi_2 \leq \xi_1} \frac{1 - x^{\xi_2, i}}{1 - \delta} \leq O((\log k)^2) \tau^{-h} \mu^x(\mathcal{V}_{\mathbb{T}}(\xi_1)),$$

where in the last inequality we used (6.9).

6.3.1 Extension to unbounded metric spaces

Note that the conclusion of [Theorem 2.6](#) has no dependence on the diameter of the space (X, d_X) , and our restriction to $\text{diam}(X, d_X) \leq 1$ was only a matter of scaling.

Throughout, we have used \mathbb{T} to denote the universal τ -HST over X , but let us now use the notation \mathbb{T}_X . To handle the case when X is unbounded, we consider a sequence of algorithms on the HSTs $\mathbb{T}_{B_0}, \mathbb{T}_{B_1}, \dots$ corresponding to bounded spaces $B_0 \subseteq B_1 \subseteq \dots \subseteq X$ defined by $B_t = \{\ell_0, \alpha_1(\sigma_1), \alpha_2(\sigma_2), \dots, \alpha_t(\sigma_t)\}$, where we recall that $\ell_0 \in \mathcal{L}_{\mathbb{T}}$ is the leaf guaranteed in [Section 2.3](#) and $\sigma = \langle \sigma_t : t \geq 1 \rangle$ is the request sequence.

The only necessity is that our assumed HST algorithm can be “isometrically transported” from \mathbb{T}_{B_t} into an isometric subtree of $\mathbb{T}_{B_{t+1}}$. One could easily formalize this property, but for simplicity we instead confirm that it holds for the algorithm described in the preceding section. Simply observe that we can extend the potentials $D(\theta; x)$ and $H(x)$ to sum over all $j \in \mathbb{Z}$. The first remains unchanged, while we define:

$$\tilde{H}(x) := \sum_{j \in \mathbb{Z}} \tau^{-j} \left(-c + \sum_{\xi \in V_{\mathbb{T}}^j} \left[\left(z^\xi + (1 + \tau^{-1})\varepsilon \right) \log \frac{z^\xi + \varepsilon}{\varepsilon} + z^\xi \log(z^\xi + \varepsilon) \right] \right),$$

where

$$c := (k + (1 + \tau^{-1})\varepsilon) \log \frac{k + \varepsilon}{\varepsilon} + k \log(k + \varepsilon).$$

This ensures that \tilde{H} is bounded whenever the leaf measure $\mu(\ell) = z^\ell$ has bounded support. (When $\tau^{-j} > \text{diam}_{\mathbb{T}}(\text{supp}(\mu))$, the corresponding term will be zero.)

Acknowledgements

I am grateful to my coauthors [[BCL⁺17](#)] for stimulating preliminary discussions, and Anna Karlin for useful comments on a preliminary draft. 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. [5](#), [19](#)
- [Bar98] Yair Bartal. On approximating arbitrary metrics by tree metrics. In *STOC '98 (Dallas, TX)*, pages 161–168. ACM, New York, 1998. [5](#)

- [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. [4](#), [12](#)
- [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. [4](#)
- [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. [4](#), [12](#), [36](#)
- [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. [2](#)
- [BBN12] Nikhil Bansal, Niv Buchbinder, and Joseph Naor. A primal-dual randomized algorithm for weighted paging. *J. ACM*, 59(4):Art. 19, 24, 2012. [4](#), [12](#)
- [BCL⁺17] Sebastien Bubeck, Michael B. Cohen, James R. Lee, Yin-Tat Lee, and Aleksander Mądry. k -server via multiscale entropic regularization. 2017. [1](#), [4](#), [5](#), [6](#), [11](#), [16](#), [19](#), [37](#), [39](#), [40](#), [41](#), [44](#)
- [BE98] Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, New York, 1998. [2](#), [4](#)
- [BLMN05] Yair Bartal, Nathan Linial, Manor Mendel, and Assaf Naor. On metric Ramsey-type phenomena. *Ann. of Math. (2)*, 162(2):643–709, 2005. [4](#)
- [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. [11](#)
- [CKPV91] M. Chrobak, H. Karloff, T. Payne, and S. Vishwanathan. New results on server problems. *SIAM J. Discrete Math.*, 4(2):172–181, 1991. [4](#)
- [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. [3](#)
- [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). [3](#)
- [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. [5](#), [8](#)
- [KLMN05] R. Krauthgamer, J. R. Lee, M. Mendel, and A. Naor. Measured descent: A new embedding method for finite metrics. *Geom. Funct. Anal.*, 15(4):839–858, 2005. [8](#)
- [Kou09] Elias Koutsoupias. The k -server problem. *Computer Science Review*, 3(2):105–118, 2009. [4](#)
- [KP95] Elias Koutsoupias and Christos H. Papadimitriou. On the k -server conjecture. *J. Assoc. Comput. Mach.*, 42(5):971–983, 1995. [3](#)

- [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. [4](#)
- [LN05] James R. Lee and Assaf Naor. Extending Lipschitz functions via random metric partitions. *Invent. Math.*, 160(1):59–95, 2005. [8](#)
- [MMS90] Mark S. Manasse, Lyle A. McGeoch, and Daniel D. Sleator. Competitive algorithms for server problems. *J. Algorithms*, 11(2):208–230, 1990. [2](#), [3](#)
- [MS91] Lyle A. McGeoch and Daniel D. Sleator. A strongly competitive randomized paging algorithm. *Algorithmica*, 6(6):816–825, 1991. [3](#)
- [Vil03] Cédric Villani. *Topics in optimal transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2003. [12](#)