Competitive Algorithms for Generalized k-Server in Uniform Metrics[*]

Nikhil Bansal[†]   Marek Eliáš[†]   Grigorios Koumoutsos[†]

Jesper Nederlof[†]

TU Eindhoven, Netherlands

July 18, 2017

Abstract

The generalized k-server problem is a far-reaching extension of the k-server problem with several applications. Here, each server s_i lies in its own metric space M_i. A request is a k-tuple $r = (r_1, r_2, \ldots, r_k)$ and to serve it, we need to move some server s_i to the point $r_i \in M_i$, and the goal is to minimize the total distance traveled by the servers. Despite much work, no $f(k)$-competitive algorithm is known for the problem for $k > 2$ servers, even for special cases such as uniform metrics and lines.

Here, we consider the problem in uniform metrics and give the first $f(k)$-competitive algorithms for general k. In particular, we obtain deterministic and randomized algorithms with competitive ratio $O(k2^k)$ and $O(k^3 \log k)$ respectively. Our deterministic bound is based on a novel application of the polynomial method to online algorithms, and essentially matches the long-known lower bound of $2^k - 1$. We also give a $2^{2^{O(k)}}$-competitive deterministic algorithm for weighted uniform metrics, which also essentially matches the recent doubly exponential lower bound for the problem.

1 Introduction

The k-server problem was proposed by Manasse et al. [24] as a far-reaching generalization of many online problems, and its study has led to various remarkable developments [6, 20, 21, 3]. In this problem, we are given k-servers $s_1, \ldots, s_k$ located at points of a metric space M. At each time step a request arrives at some point of M and must be served by moving some server there. The goal is to minimize the total distance traveled by the servers.

Koutsoupias and Taylor [23] introduced a substantial generalization of the k-server problem, called the *generalized k-server problem*. Here, each server s_i lies in its own metric space M_i, with its own distance function d_i. A request is a k-tuple $r = (r_1, r_2, \ldots, r_k)$ and must be served by moving some server s_i to the point $r_i \in M_i$. Note that the standard k-server problem corresponds to the special case when all the metrics are identical, $M_1 = \ldots = M_k = M$, and the requests are of the form $(r, r, \ldots, r)$, i.e., the k-tuple is identical in each coordinate.

The generalized k-server problem can model a rich class of online problems, for which the techniques developed for the standard k-server problem do not apply, see e.g. [23]. For that reason, it is widely believed that a deeper understanding of this problem should lead to powerful new techniques for designing online algorithms [23, 27]. According to Koutsoupias and Taylor [23], this problem "may act as a stepping stone towards building a robust (and less ad hoc) theory of online computation".

1.1 Previous Work

The k-server problem. The k-server problem has been extensively studied (an excellent reference is [6]). The initial work focused on special metrics such as uniform metrics and lines, and optimum competitive ratios were obtained in many cases [10, 11, 22]. A particularly interesting case is that of uniform metrics, which corresponds to the very well-studied *paging* problem, where tight k-competitive deterministic [30] and $O(\log k)$-competitive randomized algorithms [14, 26, 1] are known.

For general metrics, Koutsoupias and Papadimitriou [21] showed in a breakthrough result that the Work Function Algorithm (WFA) is $(2k - 1)$-competitive in any metric space. This essentially matches the lower bound of k for any deterministic algorithm [24]. More recently, a polylog(k, n) randomized competitive algorithm was obtained [3] where n in the number of points in M.

The generalized k-server problem. This problem is much less understood. In their seminal paper, Koutsoupias and Taylor [23] studied the special case where $k = 2$ and both the metrics M_1 and M_2 are lines. This is called CNN problem and it has attracted a lot of attention [2, 9, 18, 17]. They showed that, even for this special case, many successful k-server algorithms or their natural generalizations are not competitive.

Lower Bounds: For uniform metrics, Koutsoupias and Taylor [23] showed that even when each M_i contains $n = 2$ points, the competitive ratio is at least $2^k - 1$. For general metrics, the best known lower bound is $2^{2^{\Omega(k)}}$ [4], and comes from the weighted k-server problem (the weighted variant of the standard k-server problem). This problem corresponds to generalized-k-server where the metric spaces are scaled copies of each other, i.e. $M_i = w_i M$ for some fixed M, and the requests have the form $(r, \ldots, r)$.

Upper Bounds: Despite considerable efforts, competitive algorithms[1] are known only for the case of $k = 2$ servers [29, 27, 28]. In a breakthrough result, Sitters and Stougie [29] obtained a $O(1)$-competitive algorithm for $k = 2$ in any metric space. Recently, Sitters [27] showed that the generalized WFA is also $O(1)$-competitive for $k = 2$ by a careful and subtle analysis of the structure of work functions. Despite this progress, no $f(k)$-competitive algorithms are known for $k > 2$, even for special cases such as uniform metrics and lines.

1.2 Our Results

We consider the generalized k-server problem on uniform metrics and obtain the first $f(k)$-competitive algorithms for general k, whose competitive ratios almost match the known lower bounds.

Perhaps surprisingly, there turn out to be two very different settings for uniform metrics:

1. When all the metric spaces $M_1, \ldots, M_k$ are uniform (possibly with different number of points) with identical pairwise distance, say 1. We call this the *uniform metric* case.

2. When the metric spaces M_i are all uniform, but have different scales, i.e. all pairwise distances in M_i are w_i. We call this the *weighted uniform metric* case.

Our first result is the following.

Theorem 1.1. *There is a $O(k2^k)$ competitive deterministic algorithm for the generalized k-server problem in the uniform metric case.*

This almost matches the $2^k - 1$ lower bound due to [23] (we describe this instructive and simple lower bound instance in the Appendix for completeness).

The proof of Theorem 1.1 is based on a general combinatorial argument about how the set of feasible states evolves as requests arrive. Specifically, we divide the execution of the algorithm in phases, and consider the beginning of a phase when all the MSS states are feasible (e.g. the cost is 0 and not ∞). As requests arrive, the set of states that remain valid for all requests during this phase can only reduce. In particular, for this problem we show that any sequence of requests that causes the feasible state space to strictly reduce at each step, can have length at most 2^k until all states becomes infeasible.

Interestingly, this argument is based on a novel application of the polynomial or the rank method from linear algebra [19, 25, 16]. While the rank method has led to some spectacular recent successes in combinatorics and computer science [12, 13], we are not aware of any previous applications to online algorithms. We feel our approach could be useful for other online problems that can be modeled as Metrical Service Systems by analyzing the combinatorial structure in a similar way.

Next, we consider randomized algorithms against oblivious adversaries.

Theorem 1.2. *There is a randomized algorithm for the generalized k-server problem on uniform metrics with competitive ratio $O(k^3 \log k)$.*

The rank method above does not seem to be useful in the randomized setting as it only bounds the number of requests until the set of feasible states becomes empty, and does not give any structural information about how the set of states evolves over time. As we observe in Section 3, a $o(2^k)$ guarantee cannot be obtained without using such structural information. So

---

[1]Those with competitive ratio $f(k)$ that only depends on k. Note that an $n^k - 1$ competitive algorithm follows trivially, as the problem can be viewed as Metrical Service System (MSS) on n^k states, where $n = \max_{i=1}^{k} |M_i|$.

we explore the properties of this evolution more carefully and use it to design the randomized algorithm in Theorem 1.2.

In the Appendix, we also give a related lower bound. In particular, we note that an $\Omega(k/\ln^2 k)$ lower bound on the competitive ratio of any randomized algorithm follows directly by combining the lower bound instance of [23] with the results of [5].

Finally, we consider the weighted uniform metric case.

Theorem 1.3. *There is a $2^{2^{k+3}}$ competitive algorithm for generalized k-server on weighted uniform metrics.*

Theorem 1.3 follows by observing that a natural modification of an algorithm due to Fiat and Ricklin [15] for weighted k-server on uniform metrics also works for the more general generalized k-server setting. Our proof is essentially the same as that of [15], with some arguments streamlined and an improved competitive ratio[2]. Finally, note that the $2^{2^{\Omega(k)}}$ lower bound [4] for weighted k-server on uniform metrics implies that Theorem 1.3 is essentially optimal.

2 Deterministic algorithm for uniform metrics

In this section we prove Theorem 1.1. Recall that each M_i is the uniform metric with unit distance. We assume that all metrics have $n = \max_{i=1}^k |M_i|$ points (if for some metric $|M_i| < n$, we can add some extra points that are never requested). We use $[n]$ to denote $\{1, \ldots, n\}$. As the requests are arbitrary k-tuples and each metric M_i is uniform, we can relabel the points arbitrarily and hence assume that the set of points in each M_i is $[n]$. At any time t, the state of an algorithm can be described by the k-tuple $q^t = (q_1^t, \ldots, q_k^t)$ where for each $i \in [k]$, $q_i^t \in [n]$ denotes the location of server i. Let $r^t = (r_1^t, \ldots, r_k^t)$ denote the request vector at time t. We need to find a state with the following property:

Definition 2.1. *A state q^t satisfies (or is feasible for) the request r^t if $q_i^t = r_i^t$ for some $i \in [k]$.*

Moreover, if the state changes from q^t to q^{t+1}, the algorithm pays the Hamming distance

$$d(q^{t+1}, q^t) = |\{i : q_i^{t+1} \neq q_i^t\}|,$$

between q^t and q^{t+1}.

We describe a generic algorithm below that works in phases in Algorithm 1. We will show that during each phase the offline moves at least once and hence pays at least 1, while the online algorithm changes its state at most 2^k times and hence pays at most $k2^k$ as the Hamming distance between any two states is at most k. This will be sufficient as the offline optimum will need to change its state at least once as no state satisfies all requests, and it follows that our algorithm pays at most $(c^* + 1)k2^k$, where c^* denotes the optimal cost. Here the $+1$ accounts for the last (possible unfinished) phase.

We call this algorithm *generic* as it can pick any arbitrary point q as long as it is feasible for $r^1, \ldots, r^t$. Note that this algorithm captures a wide variety of natural algorithms including (variants) of the Work Function Algorithm.

Fix some phase that we wish to analyze, and let ℓ denote its length. Without loss of generality, we can assume that r^t always causes q^t to move (removing such requests does not reduce the online cost, and can only help the offline adversary). So the online algorithm moves

---

[2]It was first pointed out to us by Chiplunkar [8] that the competitive ratio $2^{2^{4k}}$ claimed in [15] can be improved to $2^{2^{k+O(1)}}$.

Algorithm 1: A deterministic $O(k2^k)$ competitive algorithm.

If a phase begins, the algorithm starts in some arbitrary q^1.

At each time t when a request r^t arrives do the following.

if the current state q^t does not satisfy the current request r^t **then**

    if there exists a state q that satisfies all requests $r^1, \ldots, r^t$ **then**

        Set $q^{t+1} = q$.

    else

        Set q^{t+1} to be an arbitrary location satisfying (only) r^t.

        End the current phase.

else

    Set $q^{t+1} = q^t$.

---

exactly ℓ times. Moreover, the adversary must move at least once during the phase as no location exists that satisfies all the requests $r^1, \ldots, r^\ell$ that arrive during the phase.

It suffices to show the following.

Theorem 2.2. *For any phase as defined above, its length satisfies $\ell \leq 2^k$.*

Proof. We use the rank method. Let $x = (x_1, \ldots, x_k), y = (y_1, \ldots, y_k)$ be points in $\mathbb{R}^k$, and consider the $2k$-variate degree k polynomial $p : \mathbb{R}^{2k} \to \mathbb{R}$,

$$p(x, y) := \prod_{i \in [k]} (x_i - y_i).$$

The key property of p is that a state $q \in [n]^k$ satisfies a request $r \in [n]^k$ iff $p(q, r) = 0$.

We now construct a matrix M that captures the dynamics of the online algorithm during a phase. Let $M \in \mathbb{R}^{\ell \times \ell}$ be an $\ell \times \ell$ matrix, where columns correspond to the states and rows to the requests, with entries $M[t, t'] = p(q^t, r^{t'})$, i.e., the $[t, t']$ entry of M corresponds to the evaluation of p on q^t and $r^{t'}$

Claim 2.3. *M is an upper triangular matrix with non-zero diagonal.*

Proof. At any time $t = 1, \ldots, \ell$, as the current state q^t does not satisfy the request r^t, it must be that $p(q^t, r^t) \neq 0$.

On the other hand, for $t = 2, \ldots, \ell$, the state q^t was chosen such that it satisfied all the previous requests t' for $t' < t$. This gives that $M[t, t'] = 0$ for $t' < t$ and hence all the entries below the diagonal are 0. $\square$

As the determinant of any upper-triangular matrix is the product of its diagonal entries, this implies that M has non-zero determinant and has full rank, $\mathrm{rk}(M) = \ell$.

However, we can use the structure of p to show that the rank of M is at most 2^k in a fairly straight manner[3]. In particular, we give an explicit factorization of M as $M = AB$, where A is $\ell \times 2^k$ matrix and M is a $2^k \times \ell$ matrix. Clearly, as any $m \times n$ matrix has rank at most $\min(m, n)$, both A and B have rank at most 2^k. Moreover, as $\mathrm{rk}(AB) \leq \min(\mathrm{rk}(A), \mathrm{rk}(B))$, this implies $\mathrm{rk}(M) \leq 2^k$. It remains to show the factorization.

---

[3]Curiously, this particular rank upper bound was used in a previous work for answering a question the a completely different setting about the parameterized complexity of graph coloring parameterized by cutwidth [31].

Indeed, if we express $p(x, y)$ in terms of its 2^k monomials, we can write

$$p(x, y) = \sum_{S \subseteq [k]} (-1)^{k - |S|} X_S Y_{[k] \setminus S},$$

where $X_S = \prod_{i \in S} x_i$ with $X_\emptyset = 1$, and Y_S is defined analogously.

Now, let A be the $\ell \times 2^k$ matrix with rows indexed by time t and columns by subsets $S \in 2^{[k]}$, with the entries

$$A[t, S] = q_S^t := \prod_{i \in S} q_i^t.$$

Similarly, let B be the $2^k \times \ell$ matrix with rows indexed by subsets $S \in 2^{[k]}$ and columns indexed by time t'. We define

$$B[S, t'] = (-1)^{k - |S|} r_{[k] \setminus S}^{t'} := (-1)^{k - |S|} \prod_{i \in [k] \setminus S} r_i^{t'}.$$

Then, for any $t, t' \in [\ell]$,

$$M[t, t'] = p(q^t, r^{t'}) = \sum_{S \subseteq [k]} (-1)^{k - |S|} q_S^t r_{[k] \setminus S}^{t'} = \sum_{S \subseteq [k]} A[t, S] B[S, t'] = (AB)[t, t'].$$

and hence $M = AB$ as claimed. $\qquad\square$

We remark that an alternate way to view this result is that the length of any request sequence that causes the set of feasible states to strictly decrease at each step can be at most 2^k.

3 Randomized algorithm for uniform metrics

A natural way to randomize the algorithm above would be to pick a state uniformly at random among all the states that are feasible for all the requests thus far in the current phase. The standard randomized uniform MTS analysis [7] implies that this online algorithm would move $O(\log(n^k)) = O(k \log n)$ times. However, this guarantee is not useful if $n \gg \exp(\exp(k))$.

Perhaps surprisingly, even if we use the fact from Section 2 that the set of feasible states can shrink at most 2^k times, this does not suffice to give a randomized $o(2^k)$ guarantee. Indeed, consider the algorithm that picks a random state among the feasible ones in the current phase. If, at each step $t = 1, \ldots, \ell$, half of the feasible states become infeasible (expect the last step when all states become infeasible), then the algorithm must move with probability at least $1/2$ at each step, and hence incur an expected $\Omega(\ell) = \Omega(2^k)$ cost during the phase.

So proving a better guarantee would require showing that the scenario above cannot happen. In particular, we need a more precise understanding of how the set of feasible states evolves over time, rather than simply a bound on the number of requests in a phase.

To this end, in Lemmas 3.1 and 3.3 below, we impose some stronger subspace-like structure over the set of feasible states. Then, we use this structure to design a variant of the natural randomized algorithm above, that directly works with these subspaces.

Spaces of configurations. Let U_i denote the set of points in M_i. We can think of $U_i = [n]$, but U_i makes the notation clear. We call state in $\prod_{i=1}^{k} U_i = [n]^k$ a configuration. Here we slightly abuse notation by letting $\prod$ denote the generalized Cartesian product. It will be useful

to consider sets of configurations where some server locations are fixed at some particular location. For a vector $v \in \prod_{i=1}^{k}(U_i \cup \{*\})$, we define the space

$$S(v) := \left\{ c \in \prod_{i=1}^{k} U_i \;\middle|\; c_i = v_i \; \forall i \text{ s.t. } v_i \neq * \right\}.$$

A coordinate i with $v_i = *$ is called *free* and the corresponding server can be located at an arbitrary point of U_i. The number of free coordinates in the space $S(v)$ we call dimension and denote it with $\dim(S(v))$.

Let us consider a d-dimensional space S and a request r such that some configuration $c \in S$ is not feasible for r. Then, we claim that a vast majority of configurations from S are infeasible for r, as stated in the following lemma. We denote $F(r)$ the set of configuration satisfying r.

Lemma 3.1. *Let S be a d-dimensional space and let r be a request which makes some configuration $c \in S$ infeasible. Then, there exist d subspaces $S_1, \ldots, S_d$, each of dimension $d-1$, such that we have $S \cap F(r) = S_1 \cup \cdots \cup S_d$.*

Note that if all the metric spaces U_i contain n points, then $|S_i| = \frac{1}{n}|S|$ for each $i = 1, \ldots, d$.

Proof. By reordering the coordinates, we can assume that the first d coordinates of S are free and S corresponds to the vector $(*, \ldots, *, s_{d+1}, \ldots, s_k)$, for some $s_{d+1}, \ldots, s_k$. Let $r = (r_1, \ldots, r_k)$.

Consider the subspaces $S(v_1), \ldots, S(v_d)$, where

$$v_1 = (r_1, *, \ldots, *, s_{d+1}, \ldots, s_k), \ldots, v_d = (*, \ldots, *, r_d, s_{d+1}, \ldots, s_k).$$

Clearly, any configuration contained in $S(v_1) \cup \ldots \cup S(v_d)$, is feasible for r. Conversely, as there exists $c \in S$ infeasible for r, we have $s_i = c_i \neq r_i$ for each $i = d+1, \ldots, k$. This already implies that each configuration from S feasible for r must belong to $S(v_1) \cup \ldots \cup S(v_d)$: whenever $c' \in S$ is feasible for r, it needs to have $c'_i = r_i$ for some $i \in \{1, \ldots, d\}$ and therefore $c' \in S(v_i)$. $\qquad\square$

Spaces of feasible configurations. During each phase, we maintain a set $\mathcal{F}^t$ of spaces containing configurations which were feasible with respect to the requests $r^1, \ldots, r^t$. In the beginning of the phase, we set $\mathcal{F}^1 = \{(r_1^1, *, \ldots, *), \ldots, (*, \ldots, *, r_k^1)\}$, and, at time t, we update it in the following way. We remove all spaces of dimension 0 whose single configuration is infeasible w.r.t. r^t. In addition, we replace each $S \in \mathcal{F}^{t-1}$ of dimension $s > 0$ which contains some infeasible configuration by $S_1, \ldots, S_d$ according to the Lemma 3.1. The following observation follows easily from Lemma 3.1.

Observation 3.2. *Let us consider a phase with requests $r^1, \ldots, r^\ell$. A configuration c is feasible with respect to the requests $r^1, \ldots, r^t$ if and only if c belongs to some space in $\mathcal{F}^t$.*

An alternative deterministic algorithm. Based on $\mathcal{F}^t$, we can design an alternative deterministic algorithm that has a competitive ratio of $3k!$. This is worse than Algorithm 1 but will be very useful to obtain our randomized algorithm. To serve a request at time t, it chooses some space $Q^t \in \mathcal{F}^t$ and moves to an arbitrary $q^t \in Q^t$. Whenever Q^{t-1} no more belongs to $\mathcal{F}^t$, it moves to another space Q^t regardless whether q^{t-1} stayed feasible or not, see Algorithm 2 for details. While, this is not an optimal behaviour, a primitive exploitation of the structure of $\mathcal{F}^t$ already gives a reasonably good algorithm.

The following lemma bounds the maximum number of distinct spaces which can appear in $\mathcal{F}^t$ during one phase. In fact, it already implies that the competitive ratio of Algorithm 2 is at most $k! \cdot \sum_{d=0}^{k-1} \frac{1}{d!} \leq 3k!$.

6

Algorithm 2: Alternative deterministic algorithm.

---

at time t:

foreach $S \in \mathcal{F}^{t-1}$ containing some infeasible configuration **do**   `// update` $\mathcal{F}^t$ `for` r^t
  └ replace S by $S_1, \ldots, S_d$ according to Lemma 3.1

if $\mathcal{F}^t = \emptyset$ **then**   `// start a new phase, if needed`
  └ Set $\mathcal{F}^t = \{S((r_1^t, *, \ldots, *)), \ldots, S((*, \ldots, *, r_k^t))\}$

if $Q^{t-1} \in \mathcal{F}^t$ **then**   `// serve the request`
  └ set $Q^t := Q^{t-1}$ and $q^t := q^{t-1}$
else
  └ choose arbitrary $Q^t \in \mathcal{F}^t$ and move to an arbitrary $q^t \in Q^t$

Lemma 3.3. *Let us consider a phase with requests* $r^1, \ldots, r^\ell$. *Then* $\bigcup_{t=1}^{\ell} \mathcal{F}^t$ *contains at most* $k!/d!$ *spaces of dimension* d.

Proof. We proceed by induction on d. In the beginning, we have $k = k!/(k-1)!$ spaces of dimension $k-1$ in $\mathcal{F}^1$ and, by Lemma 3.1, all spaces added later have strictly lower dimension.

By the way $\mathcal{F}^t$ is updated, each $(d-1)$-dimensional space is created from some d-dimensional space already present in $\bigcup_{t=1}^{\ell} \mathcal{F}^t$. By the inductive hypothesis, there could be at most $k!/d!$ distinct d-dimensional spaces and Lemma 3.1 implies that each of them creates at most d distinct $(d-1)$-dimensional spaces. Therefore, there can be at most $\frac{k!}{d!}d = \frac{k!}{(d-1)!}$ spaces of dimension $d-1$ in $\bigcup_{t=1}^{\ell} \mathcal{F}^t$. □

Randomized algorithm. Now we transform Algorithm 2 into a randomized one. Let m_t denote the largest dimension among all the spaces in $\mathcal{F}^t$ and let $\mathcal{M}^t$ denote the set of spaces of dimension m_t in $\mathcal{F}^t$.

The algorithm works as follows: Whenever moving, it picks a space Q^t from $\mathcal{M}^t$ uniformly at random, and moves to some arbitrary $q^t \in Q^t$. As the choice of q^t is arbitrary, whenever some configuration from Q^t becomes infeasible, the algorithm assumes that q^t is infeasible as well[4].

Algorithm 3: Randomized Algorithm for Uniform metrics.

---

at time t:

foreach $S \in \mathcal{F}^{t-1}$ containing some infeasible configuration **do**   `// update` $\mathcal{F}^t$ `for` r^t
  └ replace S by $S_1, \ldots, S_d$ according to Lemma 3.1

if $\mathcal{F}^t = \emptyset$ **then**   `// start a new phase, if needed`
  └ Set $\mathcal{F}^t = \{S((r_1^t, *, \ldots, *)), \ldots, S((*, \ldots, *, r_k^t))\}$

if $Q^{t-1} \in \mathcal{M}^t$ **then**   `// serve the request`
  └ set $Q^t := Q^{t-1}$ and $q^t := q^{t-1}$
else
  └ Choose a space Q^t from $\mathcal{M}^t$ uniformly at random
  └ Move to an arbitrary $q^t \in Q^t$

At each time t, ALG is located at some configuration q^t contained in some space in $\mathcal{F}^t$ which

---

[4]This is done to keep the calculations simple, as the chance of Q^t being removed from $\mathcal{F}$ and q^t staying feasible is negligible when $k \ll n$.

implies that its position is feasible with respect to the current request r^t, see Lemma 3.2. Here is the key property about the state of ALG.

Lemma 3.4. *At each time t, the probability of Q^t being equal to some fixed $S \in \mathcal{M}^t$ is $1/|\mathcal{M}^t|$.*

Proof. If ALG moved at time t, the statement follows trivially, since Q^t was chosen from $\mathcal{M}^t$ uniformly at random. So, let us condition on the event that $Q^t = Q^{t-1}$.

Now, the algorithm does not change state if and only if $Q^{t-1} \in \mathcal{M}^t$. Moreover, in this case m_t does not change, and $\mathcal{M}^t \subset \mathcal{M}^{t-1}$. By induction, Q^{t-1} is distributed uniformly within $\mathcal{M}^{t-1}$, and hence conditioned on $Q^{t-1} \in \mathcal{M}^t$, Q^t is uniformly distributed within $\mathcal{M}^t$. $\qquad \square$

Proof of Theorem 1.2. At the end of each phase (except possibly for the last unfinished phase), the set of feasible states $\mathcal{F}^t = \emptyset$, and hence OPT must pay at least 1 during each of those phases. Denoting N the number of phases needed to serve the entire request sequence, we have $\mathrm{cost}(\mathrm{OPT}) \geq (N-1)$. On the other hand, the expected online cost is at most,

$$E[\mathrm{cost}(\mathrm{ALG})] \leq c(N-1) + c \leq c\,\mathrm{cost}(\mathrm{OPT}) + c,$$

where c denotes the expected cost of ALG in one phase. This implies that ALG is c-competitive, and strictly $2c$-competitive (as the offline must move at least once, if the online algorithm pays a non-zero cost).

Now we prove that c is at most $O(k^3 \log k)$. To show this, we use a potential function

$$\Phi(t) = H(|\mathcal{M}^t|) + \sum_{d=0}^{m_t-1} H(k!/d!),$$

where $H(n)$ denotes the nth harmonic number. As the beginning of the phase, $\Phi(1) \leq kH(k!) \leq k(\log k! + 1) = O(k^2 \log k)$ as $|\mathcal{M}^1| \leq k!$ and $m_1 \leq k-1$. Moreover the phase ends whenever $\Phi(t)$ decreases to 0. Therefore, it is enough to show that, at each time t, the expected cost incurred by the algorithm is at most k times the decrease of the potential. We distinguish two cases.

If $m_t = m_{t-1}$, let us denote $b = |\mathcal{M}^{t-1}| - |\mathcal{M}^t|$. If $b > 0$, the potential decreases, and its change can be bounded as

$$\Delta\Phi \leq H(|\mathcal{M}^t|) - H(|\mathcal{M}^{t-1}|) = -\frac{1}{|\mathcal{M}^t|+1} - \frac{1}{|\mathcal{M}^t|+2} - \cdots - \frac{1}{|\mathcal{M}^t|+b} \leq -b \cdot \frac{1}{|\mathcal{M}^{t-1}|}.$$

On the other hand, the expected cost of ALG is at most k times the probability that it has to move, which is exactly $P[A_t \in \mathcal{M}^{t-1} \setminus \mathcal{M}^t] = b/|\mathcal{M}^{t-1}|$ using Lemma 3.4. Thus the expected cost of the algorithm is at most $k \cdot b/|\mathcal{M}^{t-1}|$, which is at most $k \cdot (-\Delta\Phi)$.

In the second case, we have $m_t < m_{t-1}$. By Lemma 3.3, we know that $|\mathcal{M}^t| \leq k!/m_t!$ and hence

$$\Delta\Phi = \Phi(t) - \Phi(t-1) = H(|\mathcal{M}^t|) - H(|\mathcal{M}^{t-1}|) - H(k!/m_t!) \leq -H(|\mathcal{M}^{t-1}|) \leq -1,$$

since $|\mathcal{M}^{t-1}| \geq 1$ and therefore $H(|\mathcal{M}^{t-1}|) \geq 1$. As the expected cost incurred by the algorithm is at most k, this is at most $k \cdot (-\Delta\Phi)$. $\qquad \square$

4 Algorithm for weighted uniform metrics

In this section we prove Theorem 1.3. Our algorithm is a natural extension of the algorithm of Fiat and Ricklin [15] for the weighted k-server problem on uniform metrics.

High-level idea. The algorithm is defined by a recursive construction based on the following idea. First, we can assume that the weights of the metric spaces are highly separated, i.e., $w_1 \ll w_2 \ll \ldots \ll w_k$ (if they are not we can make them separated while losing some additional factors). So in any reasonable solution, the server s_k lying in metric M_k should move much less often than the other servers. For that reason, the algorithm moves s_k only when the accumulated cost of the other $k-1$ servers reaches w_k. Choosing where to move s_k turns out to be a crucial decision. For that reason, (in each "level k-phase") during the first part of the request sequence when the algorithm only uses $k-1$ servers, it counts how many times each point of M_k is requested. We call this "learning subphase". Intuitively, points of M_k which are requested a lot are "good candidates" to place s_k. Now, during the next $c(k)$ (to be defined later) subphases, s_k visits the $c(k)$ most requested points. This way, it visits all "important locations" of M_k. A similar strategy is repeated recursively using $k-1$ servers within each subphase.

Notation and Preliminaries. We denote by s_i^{ALG} and s_i^{ADV} the server of the algorithm (resp. adversary) that lies in metric space M_i. Sometimes we drop the superscript and simply use s_i when the context is clear. We set $R_k := 2^{2^{k+2}}$ and $c(k) := 2^{2^{k+1}-3}$. Note that $c(1) = 2$ and that for all i,

$$4(c(i)+1) \cdot c(i) \le 8c(i)^2 = c(i+1). \tag{1}$$

Moreover, for all $i \ge 2$, we have

$$R_i = 8 \cdot c(i) \cdot R_{i-1}. \tag{2}$$

We assume (by rounding the weights if necessary) that $w_1 = 1$ and that for $2 \le i \le k$, w_i is an integral multiple of $2(1 + c(i-1)) \cdot w_{i-1}$. Let m_i denote the ratio $w_i/(2(1 + c(i-1)) \cdot w_{i-1})$.

The rounding can increase the weight of each server at most by a factor of $4^{k-1}c(k-1) \cdot \ldots \cdot c(1) \le R_{k-1}$. So, proving a competitive ratio R_k for an instance with rounded weights will imply a competitive ratio $R_k \cdot R_{k-1} < (R_k)^2$ for arbitrary weights.

Finally, we assume that in every request ALG needs to move a server. This is without loss of generality: requests served by the algorithm without moving a server do not affect its cost and can only increase the optimal cost. This assumption will play an important role in the algorithm below.

4.1 Algorithm Description

The algorithm is defined recursively, where ALG_i denotes the algorithm using servers $s_1, \ldots, s_i$. An execution of ALG_i is divided into phases. The phases are independent of each other and the overall algorithm is completely determined by describing how each phase works. We now describe the phases.

ALG_1 is very simple; given any request, ALG_1 moves the server to the requested point. For purposes of analysis, we divide the execution of ALG_1 into phases, where each phase consists of $2(c(1)+1) = 6$ requests.

Phase of ALG_1:
for $j = 1$ **to** $2(c(1)+1)$ **do**
Request arrives to point p: Move s_1 to p.
Terminate Phase

We now define a phase of ALG_i for $i \ge 2$. Each phase of ALG_i consists of exactly $c(i)+1$ subphases. The first subphase within a phase is special and we call it the *learning subphase*. During each subphase we execute ALG_{i-1} until the cost incurred is exactly w_i.

During the learning subphase, for each point $p \in M_i$, ALG_i maintains a count $m(p)$ of the number of requests r where p is requested in M_i, i.e. $r(i) = p$. Let us order the points of M_i as $p_1, \ldots, p_n$ such that $m(p_1) \geq \ldots \geq m(p_n)$ (ties are broken arbitrarily). We assume that $|M_i| \geq c(i)$ (if M_i has fewer points, we add some dummy points that are never requested). Let P be the set of $c(i)$ most requested points during the learning subphase, i.e. $P = \{p_1, \ldots, p_{c(i)}\}$.

For the rest of the phase ALG_i repeats the following $c(i)$ times: it moves s_i to a point $p \in P$ that it has not visited during this phase, and starts the next subphase (i.e. it calls ALG_{i-1} until its cost reaches w_i). The figure below shows the algorithm.

Phase of ALG_i, $i \geq 2$:

Move s_i to an arbitrary point of M_i;
Run ALG_{i-1} until cost incurred equals w_i ;   `// Learning subphase`
For $p \in M_i$, $m(p) \leftarrow$ # of requests such that $r(i) = p$;   `// Assume` $m(p_1) \geq \ldots \geq m(p_n)$
$P \leftarrow \{p_1, \ldots, p_{c(i)}\}$;
for $j = 1$ **to** $c(i)$ **do**
    Move s_i to an arbitrary point $p \in P$;
    $P \leftarrow P - p$;
    Run ALG_{i-1} until cost incurred equals w_i ;   `// (j + 1)th subphase`
Terminate Phase

4.2 Analysis

We first note some basic properties that follow directly by the construction of the algorithm. Call a phase of ALG_i, $i \geq 2$ *complete*, if all its subphases are finished. Similarly, a phase of ALG_1 is complete if it served exactly 6 requests.

Observation 4.1. *For $i \geq 2$, a complete phase of ALG_i consists of $(c(i) + 1)$ subphases.*

Observation 4.2. *For $i \geq 2$, the cost incurred to serve all the requests of a subphase of ALG_i is w_i.*

These observations give the following corollary.

Corollary 4.3. *For $i \geq 1$, the cost incurred by ALG_i to serve requests of a phase is $2(c(i)+1)w_i$.*

Proof. For $i = 1$ this holds by definition of the phase. For $i \geq 2$, a phase consists of $(c(i) + 1)$ subphases. Before each subphase ALG_i moves server s_i, which costs w_i, and moreover ALG_{i-1} also incurs cost w_i. $\qquad\square$

Using this, we get the following two simple properties.

Lemma 4.4. *By definition of ALG, the following properties hold:*

1. *A subphase of ALG_i, $i \geq 2$, consists of m_i complete phases of ALG_{i-1}.*

2. *All complete phases of ALG_i, $i \geq 1$, consist of the same number of requests.*

Proof. The first property uses the rounding of the weights. By Corollary 4.3, each phase of ALG_{i-1} costs $2(c(i-1)+1)w_{i-1}$ and, in each subphase of ALG_i, the cost incurred by ALG_{i-1} is w_i. So there are exactly $w_i/(2(c(i-1)+1)w_{i-1}) = m_i$ phases of ALG_{i-1}.

The property above, combined with Observation 4.1 implies that a complete phase of ALG_i contains $m_i \cdot (c(i) + 1)$ complete phases ALG_{i-1}. Now, the second property follows directly by induction: each phase of ALG_1 consists of $2(c(1) + 1) = 6$ requests, and each phase of ALG_i consists of $m_i(c(i) + 1)$ phases of ALG_{i-1}. $\qquad\square$

10

Consider a phase of ALG_i. The next lemma shows that, for any point $p \in M_i$, there exists a subphase where it is not requested too many times. This crucially uses the assumption that ALG_i has to move a server in every request.

Lemma 4.5. *Consider a complete phase of* ALG_i, $i \geq 2$. *For any point* $p \in M_i$, *there exists a subphase such that at most* $1/c(i)$ *fraction of the requests have* $r(i) = p$.

Proof. Let P be the set of $c(i)$ most requested points of M_i during the learning subphase. We consider two cases: if $p \in P$, there exists a subphase where s_i^{ALG} is located at p. During this subphase there are no requests such that $r(i) = p$, by our assumption that the algorithm moves some server at every request. Otherwise, if $p \notin P$, then during the learning subphase, the fraction of requests such that $r(i) = p$ is no more than $1/c(i)$. $\qquad\square$

To prove the competitiveness of ALG_k with respect to the optimal offline solution ADV_k, the proof uses a subtle induction on k. Clearly, one cannot compare ALG_i, for $i < k$ against ADV_k, since the latter has more servers and its cost could be arbitrarily lower. So the idea is to compare ALG_i against ADV_i, an adversary with servers $s_1, \ldots, s_i$, while ensuring that ADV_i is an accurate estimate of ADV_k during time intervals when ALG_i is called by ALG_k. To achieve this, the inductive hypothesis is required to satisfy certain properties described below. For a fixed phase, let $\mathrm{cost}(\mathrm{ALG}_i)$ and $\mathrm{cost}(\mathrm{ADV}_i)$ denote the cost of ALG_i and ADV_i respectively.

(i) **Initial Configuration of** ADV_i**.** Algorithm ALG_i (for $i < k$), is called several times during a phase of ALG_k. As we don't know the current configuration of ADV_i each time ALG_i is called, we require that for every complete phase, $\mathrm{cost}(\mathrm{ALG}_i) \leq R_i \cdot \mathrm{cost}(\mathrm{ADV}_i)$, for any initial configuration of ADV_i.

(ii) **Adversary can ignore a fraction of requests.** During a phase of ALG_i, ADV_k may serve requests with servers $s_{i+1}, \ldots, s_k$, and hence the competitive ratio of ALG_i against ADV_i may not give any meaningful guarantee. To get around this, we will require that $\mathrm{cost}(\mathrm{ALG}_i) \leq R_i \cdot \mathrm{cost}(\mathrm{ADV}_i)$, even if the ADV_i ignores an $f(i) := 4/c(i+1)$ fraction of requests. This will allow us to use the inductive hypothesis for the phases of ALG_i where ADV_k uses servers $s_{i+1}, \ldots, s_k$ to serve at most $f(i)$ fraction of requests.

For a fixed phase, we say that ALG_i is strictly R_i-competitive against ADV_i, if $\mathrm{cost}(\mathrm{ALG}_i) \leq R_i \cdot \mathrm{cost}(\mathrm{ADV}_i)$. The key result is the following.

Theorem 4.6. *Consider a complete phase of* ALG_i. *Let* ADV_i *be an adversary with* i *servers that is allowed to choose any initial configuration and to ignore any* $4/c(i+1)$ *fraction of requests. Then,* ALG_i *is strictly* R_i-*competitive against* ADV_i.

Before proving this, let us note that this directly implies Theorem 1.3. Indeed, for any request sequence σ, all phases except possibly the last one, are complete, so $\mathrm{cost}(\mathrm{ALG}_k) \leq R_k \cdot \mathrm{cost}(\mathrm{ADV}_k)$. The cost of ALG_k for the last phase, is at most $2(c(k)+1)w_k$, which is a fixed additive term independent of the length of σ. So, $\mathrm{ALG}_k(\sigma) \leq R_k \cdot \mathrm{ADV}_k(\sigma) + 2(c(k)+1)w_k$, and ALG_k is R_k-competitive. Together with loss in rounding the weights, this gives a competitive ratio of at mot $(R_k)^2 \leq 2^{2^{k+3}}$ for arbitrary weights.

We now prove Theorem 4.6.

Proof of Theorem 4.6. We prove the theorem by induction on k.

Base case ($i = 1$): As $R_1 > 6$ and $4/c(2) = 1/8 \leq 1/3$, it suffices to show here that ALG_1 is strictly 6-competitive in a phase where ADV_1 can ignore at most $1/3$ fraction of requests,

11

for any starting point of $s_1^{\mathrm{ADV}_1}$. By Lemma 4.4, we have $\mathrm{cost}(\mathrm{ALG}_1) = 2(c(1) + 1) = 6$. We show that $\mathrm{cost}(\mathrm{ADV}_1) \geq 1$. Consider two consecutive requests r_{t-1}, r_t. By our assumption that ALG_1 has to move its server in every request, it must be that $r_{t-1} \neq r_t$. So, for any t if ADV_1 does not ignore both r_{t-1} and r_t, then it must pay 1 to serve r_t. Moreover, as the adverary can chose the initial server location, it may (only) serve the first request at zero cost. As a phase consists of 6 requests, ADV_i can ignore at most $6/3 = 2$ of them, so there are at most 4 requests that are either ignored or appear immediately after an ignored request. So among requests $r_2, \ldots, r_6$, there is at least one request r_t, such that both r_{t-1} and r_t are not ignored.

Inductive step: Assume inductively that ALG_{i-1} is strictly R_{i-1}-competitive against any adversary with $i-1$ servers that can ignore up to $4/c(i)$ fraction of requests.

Let us consider some phase at level i, and let I denote the set of requests that ADV_i chooses to ignore during the phase. We will show that $\mathrm{cost}(\mathrm{ADV}_i) \geq w_i/(2R_{i-1})$. This implies the theorem, as $\mathrm{cost}(\mathrm{ALG}_i) = 2(c(i) + 1)w_i$ by Lemma 4.4 and hence,

$$\frac{\mathrm{cost}(\mathrm{ALG}_i)}{\mathrm{cost}(\mathrm{ADV}_i)} \leq \frac{2(c(i) + 1)w_i}{w_i/(2R_{i-1})} = 4(c(i) + 1)R_{i-1} \leq 8 \cdot c(i) \cdot R_{i-1} = R_i.$$

First, if ADV_i moves server s_i during the phase, its cost is already at least w_i and hence more than $w_i/(2R_{i-1})$. So we can assume that s_i^{ADV} stays fixed at some point $p \in M_i$ during the entire phase. So, ADV_i is an adversary that uses $i-1$ servers and can ignore all requests with $r(i) = p$ and the requests of I. We will show that there is a subphase where $\mathrm{cost}(\mathrm{ADV}_i) \geq w_i/(2R_{i-1})$.

By Lemma 4.5, there exists a subphase, call it j, such that at most $1/c(i)$ fraction of the requests have $r(i) = p$. As all $c(i) + 1$ subphases have the same number of requests (by Lemma 4.4), even if all the requests of I belong to subphase j, they make up at most $(4 \cdot (c(i) + 1))/c(i + 1) \leq 1/c(i)$ fraction of its requests, where the inequality follows from equation (1). So overall during subphase j, ADV_i uses servers $s_1, \ldots, s_{i-1}$ and ignores at most $2/c(i)$ fraction of requests.

We now apply the inductive hypothesis together with an averaging argument. As subphase j consists of m_i phases of ALG_{i-1}, all of equal length, and ADV_i ignores at most $2/c(i)$ fraction of requests of the subphase, there are at most $m_i/2$ phases of ALG_{i-1} where it can ignore more than $4/c(i)$ fraction of requests. So, for at least $m_i/2$ phases of ALG_{i-1}, ADV_i uses $i-1$ servers and ignores no more than $4/c(i)$ fraction of requests. By the inductive hypothesis, ALG_{i-1} is strictly R_{i-1}-competitive against ADV_i in these phases. As the cost of ALG_{i-1} for each phase is the same (by Lemma 4.4), overall ALG_i is strictly $2R_{i-1}$ competitive during subphase j. As the cost of ALG_i during subphase j is w_i, we get that $\mathrm{cost}(\mathrm{ADV}_i) \geq w_i/2R_{i-1}$, as claimed. $\quad\square$

Acknowledgments

References

[1] Dimitris Achlioptas, Marek Chrobak, and John Noga. Competitive analysis of randomized paging algorithms. *Theor. Comput. Sci.*, 234(1-2):203–218, 2000.

[2] John Augustine and Nick Gravin. On the continuous CNN problem. In *ISAAC*, pages 254–265, 2010.

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

[4] Nikhil Bansal, Marek Eliáš, and Grigorios Koumoutsos. Weighted k-server bounds via combinatorial dichotomies. *CoRR*, abs/1704.03318, To appear in FOCS'17.

[5] Yair Bartal, Béla Bollobás, and Manor Mendel. Ramsey-type theorems for metric spaces with applications to online problems. *J. Comput. Syst. Sci.*, 72(5):890–921, 2006.

[6] Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, 1998.

[7] Allan Borodin, Nathan Linial, and Michael E. Saks. An optimal on-line algorithm for metrical task system. *J. ACM*, 39(4):745–763, 1992.

[8] Ashish Chiplunkar. Personal Communication. Oct 2016.

[9] Marek Chrobak. SIGACT news online algorithms column 1. *SIGACT News*, 34(4):68–77, 2003.

[10] Marek Chrobak, Howard J. Karloff, Thomas H. Payne, and Sundar Vishwanathan. New results on server problems. *SIAM J. Discrete Math.*, 4(2):172–181, 1991.

[11] Marek Chrobak and Lawrence L. Larmore. An optimal on-line algorithm for k-servers on trees. *SIAM J. Comput.*, 20(1):144–148, 1991.

[12] Z. Dvir. On the size of Kakeya sets in finite fields. *J. Amer. Math. Soc.*, 22:1093–1097, 2009.

[13] J. S. Ellenberg and D. Gijswijt. On large subsets of F_q^n with no three-term arithmetic progression. *ArXiv e-prints*, arXiv:1605.09223, 2016.

[14] Amos Fiat, Richard M. Karp, Michael Luby, Lyle A. McGeoch, Daniel Dominic Sleator, and Neal E. Young. Competitive paging algorithms. *J. Algorithms*, 12(4):685–699, 1991.

[15] Amos Fiat and Moty Ricklin. Competitive algorithms for the weighted server problem. *Theor. Comput. Sci.*, 130(1):85–99, 1994.

[16] L. Guth. *Polynomial Methods in Combinatorics*. University Lecture Series. American Mathematical Society, 2016.

[17] Kazuo Iwama and Kouki Yonezawa. Axis-bound CNN problem. *IEICE TRANS*, pages 1–8, 2001.

[18] Kazuo Iwama and Kouki Yonezawa. The orthogonal CNN problem. *Inf. Process. Lett.*, 90(3):115–120, 2004.

[19] Stasys Jukna. *Extremal Combinatorics - With Applications in Computer Science*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2011.

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

[21] Elias Koutsoupias and Christos H. Papadimitriou. On the k-server conjecture. *J. ACM*, 42(5):971–983, 1995.

[22] Elias Koutsoupias and Christos H. Papadimitriou. The 2-evader problem. *Inf. Process. Lett.*, 57(5):249–252, 1996.

[23] Elias Koutsoupias and David Scot Taylor. The CNN problem and other k-server variants. *Theor. Comput. Sci.*, 324(2-3):347–359, 2004.

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

[25] Jiří Matoušek. *Thirty-three Miniatures: Mathematical and Algorithmic Applications of Linear Algebra*. American Mathematical Society, 2010.

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

[27] René Sitters. The generalized work function algorithm is competitive for the generalized 2-server problem. *SIAM J. Comput.*, 43(1):96–125, 2014.

[28] René Sitters, Leen Stougie, and Willem de Paepe. A competitive algorithm for the general 2-server problem. In *ICALP*, pages 624–636, 2003.

[29] René A. Sitters and Leen Stougie. The generalized two-server problem. *J. ACM*, 53(3):437–458, 2006.

[30] Daniel Dominic Sleator and Robert Endre Tarjan. Amortized efficiency of list update and paging rules. *Commun. ACM*, 28(2):202–208, 1985.

[31] Bas van Geffen, Bart Jansen, Noud de Kroon, Rolf Morel, and Jesper Nederlof. Optimal algorithms on graphs of bounded width (and degree): Cutwidth sometimes beats treewidth, but planarity does not help. Unpublished.

A   Lower Bounds

We present simple lower bounds on the competitive ratio of deterministic and randomized algorithms for the generalized k-server problem in uniform metrics.

Deterministic Algorithms.   We show a simple construction due to [23] that directly implies a $(2^k - 1)/k$ lower bound on the competitive ratio of deterministic algorithm. Using a more careful argument, [23] also improve this to $2^k - 1$.

Assume that each metric space M_i has $n = 2$ points, labeled by 0,1. A configuration of servers is a vector $c \in \{0,1\}^k$, so there are 2^k possible configurations. Now, a request $r = (r_1, \ldots, r_k)$ is unsatisfied if and only if the algorithm is in the antipodal configuration $\bar{r} = (1 - r_1, \ldots, 1 - r_k)$. Let ALG be any online algorithm and ADV be the adversary. Initially, ALG and ADV are in the same configuration. At each time step, if the current configuration of ALG is $a = (a_1, \ldots, a_k)$, the adversary requests $\bar{a}$ until ALG visits every configuration. If p is the configuration that ALG visits last, the adversary can simply move to p at the beginning, paying at most k, and satisfy all requests until ALG moves to p. On the other hand, ALG pays at least $2^k - 1$ until it reaches p. Once ALG and ADV are in the same configuration, the strategy repeats.

Randomized Algorithms.   Viewing generalized k-server as a metrical service system (MSS), we can get a non-trivial lower bound for randomized algorithms. In particular, we can apply the $\Omega(\frac{\log N}{\log^2 \log N})$ lower bound due to Bartal et al. [5] on the competitive ratio of any randomized online algorithm against oblivious adversaries, for any metrical task system on N states. Of course, the MSS corresponding to a generalized k-server instance is restricted as the cost vectors may not be completely arbitrary. However, we consider the case where all metrics M_i have $n = 2$ points. Let s be an arbitrary state among the $N = 2^k$ possible states. A request in the antipodal point $\bar{s}$ only penalizes s and has cost 0 for every other state. So the space of cost vectors here is rich enough to simulate any MSS on these N states[5].

This implies a $\Omega(\frac{k}{\log^2 k})$ lower bound for generalized k-server problem on uniform metrics.

---

[5]Note that if there is a general MSS request that has infinite cost on some subset S of states, then decomposing this into $|S|$ sequential requests where each of them penalizes exactly one state of S, can only make the competitive ratio worse.