

# The Randomized Competitive Ratio of Weighted $k$ -Server Is at Least Exponential

Nikhil Ayyadevara ✉

Indian Institute of Technology, New Delhi, India

Ashish Chiplunkar ✉

Indian Institute of Technology, New Delhi, India

---

## Abstract

The weighted  $k$ -server problem is a natural generalization of the  $k$ -server problem in which the cost incurred in moving a server is the distance traveled times the weight of the server. Even after almost three decades since the seminal work of Fiat and Ricklin (1994), the competitive ratio of this problem remains poorly understood even on the simplest class of metric spaces – the uniform metric spaces. In particular, in the case of randomized algorithms against the oblivious adversary, neither a better upper bound than the doubly exponential deterministic upper bound, nor a better lower bound than the logarithmic lower bound of unweighted  $k$ -server, is known. In this paper, we make significant progress towards understanding the randomized competitive ratio of weighted  $k$ -server on uniform metrics. We cut down the triply exponential gap between the upper and lower bound to a singly exponential gap by proving that the competitive ratio is at least exponential in  $k$ , substantially improving on the previously known lower bound of about  $\ln k$ .

**2012 ACM Subject Classification** Theory of computation → Online algorithms

**Keywords and phrases** weighted  $k$ -server, competitive analysis

**Digital Object Identifier** 10.4230/LIPIcs.ESA.2021.9

**Related Version** *Full Version*: <https://arxiv.org/abs/2102.11119>

**Funding** *Ashish Chiplunkar*: partially supported by the Pankaj Gupta New Faculty Fellowship.

## 1 Introduction

The  $k$ -server problem of Manasse, McGeoch, and Sleator [12] is one of the cleanest, simple-looking, and yet profound problems in online computation, and has been actively studied for over three decades. The  $k$ -server problem concerns deciding movements of  $k$  mobile servers on an underlying metric space to serve a sequence of online requests. Each request is issued at some point of the metric space, and to serve such a request, a server must move to the requested point (unless a server is already present there). The cost incurred in the movement of a server is equal to the distance through which the server moves, and the goal is to minimize the total cost.

Since an online algorithm is required to take its decisions only based on the past inputs, it cannot output the optimal solution, in general. An online algorithm for a minimization problem is said to be  $c$ -competitive if, on any instance, it produces a solution whose (expected) cost is at most  $c$  times the cost of the optimum solution. The competitive ratio of an algorithm is the minimum (technically, the infimum of all)  $c$  such that the algorithm is  $c$ -competitive. The deterministic (resp. randomized) competitive ratio of an online minimization problem is the minimum (technically, the infimum of all)  $c$  for which a  $c$ -competitive deterministic (resp. randomized) algorithm exists. Note that, unless otherwise specified, we assume that in case of randomized algorithms, the adversarial input is *oblivious*, that is, constructed with the knowledge of the algorithm but not the random choices the algorithm makes.



© Nikhil Ayyadevara and Ashish Chiplunkar;  
licensed under Creative Commons License CC-BY 4.0  
29th Annual European Symposium on Algorithms (ESA 2021).

Editors: Petra Mutzel, Rasmus Pagh, and Grzegorz Herman; Article No. 9; pp. 9:1–9:11

Leibniz International Proceedings in Informatics



LIPIC Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In their seminal work, Manasse, McGeoch, and Sleator [12] proved that the deterministic competitive ratio of the  $k$ -server problem is at least  $k$  on every metric with more than  $k$  points. They conjectured that the deterministic competitive ratio is, in fact, equal to  $k$  on any metric. This conjecture is popularly called the deterministic  $k$ -server conjecture and it remains unresolved to date. The deterministic algorithm with the best known competitive ratio of  $2k - 1$  is due to Koutsoupias and Papadimitriou [10]. Surprisingly, no better algorithm is known even using randomization. The randomized  $k$ -server conjecture states that a randomized algorithm with competitive ratio  $O(\log k)$  exists on all metrics, and this remains unresolved after some recent progress [5, 11]. The  $k$ -server problem on uniform metric spaces is particularly interesting because it is equivalent to the paging problem. In this case, several deterministic algorithms including Least-Recently-Used (LRU) and First-In-First-Out (FIFO) are known to be  $k$ -competitive. The randomized competitive ratio is known to be exactly  $H(k) = \sum_{i=1}^k 1/i \approx \ln k$ , where the lower bound is due to Fiat et al. [8] and the upper bound is due to [13, 1].

The weighted  $k$ -server problem is a natural generalization of the  $k$ -server problem where the objective is to minimize the weighted sum of the movements of servers. Specifically, the  $k$  servers have weights  $\beta_1 \leq \dots \leq \beta_k$ , and the cost of moving the  $i$ 'th server is  $\beta_i$  times the distance through which it moves. It is easy to see that a  $c$ -competitive  $k$ -server algorithm has competitive ratio at most  $c\beta_k/\beta_1$  for the weighted  $k$ -server problem, and therefore, the challenge is to design an algorithm with competitive ratio independent of the servers' weights. Surprisingly, this innocuous-looking introduction of weights into the  $k$ -server problem makes it incredibly difficult, and a competitive algorithm is known only for  $k \leq 2$  [14] (of which, the  $k = 1$  case is trivial).

## 1.1 Weighted $k$ -Server on Uniform Metrics

Owing to the difficulty of the weighted  $k$ -server problem on general metrics, the problem becomes particularly interesting on uniform metrics. The weighted  $k$ -server problem on uniform metric spaces models the paging problem where the cost of page replacement is determined by the location where the replacement takes place. Note that this problem is different from weighted caching [15], where the cost of page replacement is determined by the pages that get swapped in and out.

The seminal paper of Fiat and Ricklin [9] gave a deterministic algorithm for weighted  $k$ -server on uniform metrics whose competitive ratio is doubly exponential in  $k$ : about  $3^{4^k}/3$  specifically, but can be improved to  $2^{2^{k+2}} = 16^{2^k}$  due to the result of Bansal et al. [3] for a more general problem. The fact that the deterministic competitive ratio is indeed doubly exponential in  $k$  was established only recently by Bansal et al. [2], who proved a lower bound of  $2^{2^{k-4}}$ , improving the previously known lower bound of  $(k+1)!/2$  due to Fiat and Ricklin [9].

The only known algorithm for the weighted  $k$ -server problem on uniform metrics which makes non-trivial use of randomness is by Chiplunkar and Vishwanathan [6]. This algorithm also achieves a doubly exponential competitive ratio of about  $c^{2^k}$  for  $c \approx 1.59792$ . It is, in fact, a *randomized memoryless algorithm* generalizing the algorithm by Chrobak and Sgall [7] for  $k = 2$ , and it achieves the competitive ratio against a stronger form of adversary called *adaptive online adversary*<sup>1</sup>. Chiplunkar and Vishwanathan also proved

---

<sup>1</sup> An adaptive online adversary can see the movements of the algorithm's servers even though the algorithm is randomized. However, the adversary must also serve its requests in an online manner. The algorithm's

that no randomized memoryless algorithm can achieve a better competitive ratio against adaptive online adversaries. However, even when an algorithm is allowed to use both memory and randomness, and the adversary is oblivious, no better upper bound is known. More embarrassingly, for randomized algorithms, no better lower bound than the logarithmic lower bound of (unweighted)  $k$ -server on uniform metrics is known, thus, leaving a triply exponential gap between the upper and lower bounds.

In this paper, we cut down the triply exponential gap between the best known bounds on the randomized competitive ratio of weighted  $k$ -server on uniform metrics by a doubly exponential improvement in the lower bound. We prove,

► **Theorem 1.** *The competitive ratio of any randomized algorithm for weighted  $k$ -server on uniform metrics is at least exponential in  $k$ , even when the algorithm is allowed to use memory and the adversary is oblivious.*

Due to our result, we now have only a singly exponential gap between the best known upper and lower bounds on the randomized competitive ratio of weighted  $k$ -server on uniform metrics.

## 1.2 Comparison with the Deterministic Lower Bound

Our proof of the randomized lower bound for weighted  $k$ -server is inspired by the proof of the deterministic lower bound by Bansal et al. [2]. Both proofs give adversaries which run recursively defined strategies relying crucially on a certain set-system  $\mathcal{Q}$ . However, our proof differs in the following aspects.

1. The adversary in the deterministic lower bound proof is able to carefully pick from  $\mathcal{Q}$  a set of points that does not contain points covered by the algorithm's heavier servers, and run its strategy on that set. In contrast, our adversary is oblivious and is unable to see the positions of the algorithm's servers. Therefore, it merely picks a random set from  $\mathcal{Q}$  and hopes that none of the points in that set is covered by the algorithm's heavier servers.
2. The strategy of Bansal et al. to defeat deterministic algorithms ensures that whenever an adversary's server other than the heaviest moves, it is accompanied by an eventual movement of a heavier server of the algorithm. Therefore, assuming that the weights of the servers are well-separated, their task reduces to proving that the heaviest server of the algorithm moves a large number of times as compared to the heaviest server of the adversary. On the other hand, we are unable to charge the movement of an adversary's server to the movement of an algorithm's heavier server. Consequently, we need to carefully track the contributions of all  $k$  servers towards the algorithm's and the adversary's costs.

## 2 Preliminaries

Let the weights of the  $k$  servers be  $1, \beta, \beta^2, \dots, \beta^{k-1}$  for some large integer  $\beta$  which we will fix later. Define the sequence  $n_0, n_1, \dots$  inductively as follows.  $n_0 = 1$ , and for  $\ell > 0$ ,

$$n_\ell = \left( \left\lceil \frac{n_{\ell-1}}{2} \right\rceil + 1 \right) \cdot \left( \left\lfloor \frac{n_{\ell-1}}{2} \right\rfloor + 1 \right).$$

---

cost is compared with the cost of the adversary's online solution to determine the competitive ratio.

## 9:4 The Randomized Competitive Ratio of Weighted $k$ -Server Is at Least Exponential

Observe that  $n_k$  grows doubly exponentially with respect to  $k$ . Since  $n_\ell \geq n_{\ell-1}^2/4$ , it is easy to prove using induction that  $n_\ell \geq 4 \cdot (64^{1/32})^{2^\ell}$  for all  $\ell \geq 5$ . Let  $H$  denote the harmonic function, that is,  $H(n) = \sum_{i=1}^n 1/i$ . It is known that  $H(n) \geq \ln n$ . We will establish Theorem 1 by proving the following bound.

► **Theorem 2.** *The randomized competitive ratio of weighted  $k$ -server on uniform metric spaces is at least  $H(n_{k-1}) = \Omega(2^k)$ .*

We use the following version of Yao's principle to prove the above bound.

► **Proposition 3 (Yao's principle).** *Suppose there exists a probability distribution  $\mathcal{D}$  on the instances of an online minimization problem such that for every deterministic online algorithm  $A$ , we have,*

$$\mathbb{E}_{I \sim \mathcal{D}}[A(I)] > \alpha \cdot \mathbb{E}_{I \sim \mathcal{D}}[OPT(I)],$$

where  $A(I)$  is the cost of the algorithm's solution and  $OPT(I)$  is the cost of an optimal solution to instance  $I$ . Then the problem does not have an  $\alpha$ -competitive randomized online algorithm.

Thus, in order to prove Theorem 2, our task is exhibit a distribution on instances of weighted  $k$ -server on a uniform metric space such that the expected cost of any deterministic online algorithm is greater than  $H(n_{k-1})$  times the expectation of the optimum cost. To construct our distribution on instances, we use a combinatorial result with a constructive proof given by Bansal et al. [2]. We reproduce its proof in Appendix A for completeness. The result is as follows.

► **Lemma 4.** *Let  $\ell \in \mathbb{N}$  and let  $P$  be a set of  $n_\ell$  points. There exists a set-system  $\mathcal{Q}_\ell \subseteq 2^P$  satisfying the following properties.*

1.  $\mathcal{Q}_\ell$  contains  $\lceil n_{\ell-1}/2 \rceil + 1$  sets, each of size  $n_{\ell-1}$ .
2. For every  $p \in P$ , there exists a set in  $\mathcal{Q}_\ell$  not containing  $p$ .
3. For every  $p \in P$ , there exists a  $q \in P$  such that every set in  $\mathcal{Q}_\ell$  contains at least one of  $p$  and  $q$ .

### 3 Adversarial Strategy and Analysis

Consider the uniform metric space on a set  $S$  of  $n_{k-1} + 1$  points. Our adversarial input distribution is generated by the procedure `adversary` which uses a recursive procedure `strategy`, an oblivious version of its counterpart in Bansal et al. [2]. These procedures are defined as follows.

■ **Procedure 1** `adversary`.

---

**repeat** *infinitely many times*

- ┌ Pick a point  $p$  uniformly at random from  $S$  (with replacement);
  - └ Call `strategy`( $k - 1, S \setminus \{p\}$ );
-

■ **Procedure 2**  $\text{strategy}(\ell, P)$  (Promise:  $|P| = n_\ell$ ).

---

```

if  $\ell = 0$  (and therefore,  $|P| = n_0 = 1$ ) then
  | Request the unique point in  $P$ ;
else
  | Construct the set-system  $\mathcal{Q}_\ell \subseteq 2^P$  using Lemma 4;
  | repeat  $(\beta - 1) \cdot (\lceil n_{\ell-1}/2 \rceil + 1)$  times
  | | Pick a set  $P'$  uniformly at random from  $\mathcal{Q}_\ell$  (with replacement);
  | | Call  $\text{strategy}(\ell - 1, P')$ ;

```

---

Procedure  $\text{strategy}$  gets as input a non-negative number  $\ell$  and a set  $P$  of  $n_\ell$  points. In the base case where  $\ell = 0$ , the procedure issues a request to the unique point in  $P$ . In the inductive case where  $\ell > 0$ , the procedure constructs the set-system  $\mathcal{Q}_\ell$  with properties stated in Lemma 4 on the set  $P$ . Then it repeatedly gives recursive calls, passing  $\ell - 1$  in place of  $\ell$ , on sets chosen uniformly at random from  $\mathcal{Q}_\ell$ . Recall that these sets have size  $n_{\ell-1}$ , as required. Procedure  $\text{adversary}$  takes a uniform metric space on  $n_{k-1} + 1$  points. It repeatedly picks a point  $p$  uniformly at random and calls the procedure  $\text{strategy}$  on the set of points other than  $p$ .

For analysis, fix an arbitrary deterministic online algorithm and the initial positions of its servers. We first consider requests given by one execution of procedure  $\text{strategy}(\ell, P)$ , and bound the number of movements of the algorithm's servers to serve those requests.

► **Lemma 5.** *For every  $\ell \in \{0, \dots, k-1\}$  the following holds. Let  $\rho_0$  be an arbitrary sequence of requests and  $L$  be the set of positions of the algorithm's heaviest  $k - \ell$  servers after serving  $\rho_0$ . Let  $P$  be an arbitrary set of  $n_\ell$  points disjoint from  $L$ . Suppose  $\rho_0$  is followed by a random sequence  $\rho$  of requests given by a  $\text{strategy}(\ell, P)$  call. For  $i = 1, \dots, k$ , let the random variable  $X_i$  denote the number of movements of the algorithm's  $i$ 'th lightest server while the algorithm serves  $\rho$ . Then we have,*

$$\sum_{i=1}^k \beta^{\min(i-1, \ell)} \cdot \mathbb{E}[X_i] \geq (\beta - 1)^\ell.$$

We defer the proof of this lemma to Appendix B. On a high level, the proof goes as follows. If the algorithm moves one of its heaviest  $k - \ell$  servers while it serves  $\rho$ , then it pays a lot already. If not, it must serve  $\rho$  using its lightest  $\ell$  servers only. In this case, each recursive call given by the  $\text{strategy}(\ell, P)$  call is, with sufficient probability, on a set  $P'$  not containing the location of the algorithm's  $\ell$ 'th lightest server. This enables us to use induction hypothesis to bound the algorithm's cost in each recursive call.

Intuitively, Lemma 5 gives a lower bound of  $(\beta - 1)^\ell$  on the expected cost incurred by the algorithm in serving requests given by a  $\text{strategy}(\ell, P)$  call, but with the following caveat: movements of the heaviest  $k - \ell - 1$  servers are charged at a discounted rate of  $\beta^\ell$ . However, when  $\ell$  is instantiated to  $k - 1$  in particular, no discount remains applicable. Thus,  $(\beta - 1)^{k-1}$  becomes a lower bound on the expected actual cost of the algorithm in serving requests given by a  $\text{strategy}(k - 1, P)$  call. With this observation, we immediately get the following bound on the expected cost of the algorithm in serving requests given by each  $\text{strategy}$  call made by the procedure  $\text{adversary}$ .

► **Corollary 6** (to Lemma 5). *The expected cost of the algorithm in serving requests given by each  $\text{strategy}$  call made by adversary is at least  $(\beta - 1)^{k-1} / (n_{k-1} + 1)$ .*

**Proof.** Consider any  $\text{strategy}(k-1, S \setminus \{p\})$  call, where  $p$  is a uniformly random point in  $S$ . Let  $r$  be the location of the algorithm's heaviest server at the time the call is made. Then  $\Pr[r \notin S \setminus \{p\}] = \Pr[p = r] = 1/|S| = 1/(n_{k-1} + 1)$ . Lemma 5 implies that conditioned on  $r \notin S \setminus \{p\}$ , the expected cost of the algorithm in serving requests given by the  $\text{strategy}(k-1, S \setminus \{p\})$  call is at least  $(\beta - 1)^{k-1}$ . Thus, the claim follows.  $\blacktriangleleft$

Let us now turn our attention towards the adversary's cost. We will show how the adversary, having the ability to see the future requests, can ensure that whenever  $\text{strategy}(\ell, P)$  is called, it has at least one server other than its  $\ell$  lightest servers occupying a point in  $P$  already. On the contrary, recall that in Lemma 5, we relied on the algorithm not having any of its servers except the  $\ell$  lightest ones occupying points in  $P$  at the time  $\text{strategy}(\ell, P)$  is called. Intuitively, the adversary is able to obtain advantage over the algorithm by having one server other than the  $\ell$  lightest ones in  $P$  whereas the algorithm has none.

► **Lemma 7.** *Define the sequence  $c_0, c_1, \dots$  inductively as follows:  $c_0 = 0$ , and for  $\ell > 0$ ,*

$$c_\ell = \beta^{\ell-1} + \beta \cdot (\lceil n_{\ell-1}/2 \rceil + 1) \cdot c_{\ell-1}.$$

*Suppose that the adversary has at least one server other than its  $\ell$  lightest servers occupying some point in  $P$  at the time  $\text{strategy}(\ell, P)$  is called. Then the adversary is able to serve all requests given in this call with cost at most  $c_\ell$  by moving only its  $\ell$  lightest servers.*

**Proof.** We prove the claim by induction on  $\ell$ . For the base case, suppose  $\ell = 0$ . Then  $|P| = 1$  and by assumption, the adversary has at least one server at the unique point in  $P$ . Therefore, the adversary can serve the unique request given by  $\text{strategy}(0, P)$  with cost  $c_0 = 0$ , without moving any server.

For the inductive case, suppose  $\ell > 0$ . We have assumed that the adversary has at least one server other than its lightest  $\ell$  servers occupying some point  $p$  in  $P$ . By the third property of the set-system  $\mathcal{Q}_\ell$  from Lemma 4, there exists a point  $q \in P$  such that each set in  $\mathcal{Q}_\ell$  contains at least one of  $p$  and  $q$ . The adversary moves its  $\ell$ 'th lightest server to such a point  $q$  and keeps it there until the end of the  $\text{strategy}(\ell, P)$  call. Due to this movement, the adversary incurs cost  $\beta^{\ell-1}$ , the first term in the definition of  $c_\ell$ . As a result, both  $p$  and  $q$  become occupied by the adversary's servers other than the  $\ell - 1$  lightest ones. We now show how the requests in all recursive calls made by  $\text{strategy}(\ell, P)$  can be served by moving the  $\ell - 1$  lightest servers only.

Consider any of the recursive calls made by  $\text{strategy}(\ell, P)$ . The set  $P' \in \mathcal{Q}_\ell$  on which this call is made contains at least one of  $p$  and  $q$ . Both  $p$  and  $q$  were occupied by the adversary's servers other than the  $\ell - 1$  lightest ones before  $\text{strategy}(\ell, P)$  made its first recursive call. All the previous recursive calls were served by moving only the  $\ell - 1$  lightest servers. Thus, at the time the current recursive call  $\text{strategy}(\ell - 1, P')$  is made, points  $p$  and  $q$  are still occupied by the adversary's servers other than the  $\ell - 1$  lightest ones. Therefore, at least one of these servers occupies a point in  $P'$ . By induction hypothesis, the adversary can serve all requests in the current recursive call  $\text{strategy}(\ell - 1, P')$  with cost at most  $c_{\ell-1}$  by moving only the  $\ell - 1$  lightest servers. Since the number of such recursive calls is  $(\beta - 1) \cdot (\lceil n_{\ell-1}/2 \rceil + 1) \leq \beta \cdot (\lceil n_{\ell-1}/2 \rceil + 1)$ , the adversary serves all requests made in these calls with cost at most  $\beta \cdot (\lceil n_{\ell-1}/2 \rceil + 1) \cdot c_{\ell-1}$ , the second term in the expression for  $c_\ell$ .  $\blacktriangleleft$

We now use Corollary 6 and Lemma 7 to prove Theorem 2.

► **Theorem 2.** *The randomized competitive ratio of weighted  $k$ -server on uniform metric spaces is at least  $H(n_{k-1}) = \Omega(2^k)$ .*

**Proof.** We track the costs incurred by the algorithm and the adversary per  $\text{strategy}(k-1, P)$  call made by the procedure `adversary`, and show that the former is at least  $H(n_{k-1})$  times the latter.

Here is how the adversary serves the requests. Let  $q$  denote the position of the adversary's heaviest server at the time a  $\text{strategy}(k-1, P)$  call is made. If  $P = S \setminus \{q\}$ , that is, the random point sampled from  $S$  turns out to be  $q$ , then the adversary finds the point  $q'$  which is sampled farthest in future by the procedure `adversary`, and moves its heaviest server there. These are the only movements of the adversary's heaviest server. By the standard coupon-collector argument, the expected number of samples from the current sample of  $q$  to  $q'$  is  $(n_{k-1} + 1)H(n_{k-1})$ , because  $|S| = n_{k-1} + 1$  and we have already sampled  $q$ . Thus, in the long run, the cost of the adversary resulting from moving its heaviest server, per  $\text{strategy}$  call made by the procedure `adversary`, is  $\beta^{k-1}/((n_{k-1} + 1)H(n_{k-1}))$ .

By moving its heaviest server as described above, the adversary ensures the following. Before the adversary starts serving requests given by a  $\text{strategy}(k-1, S \setminus \{p\})$  call, its heaviest server is located at some point different from  $p$ , and therefore, in  $S \setminus \{p\}$ . By Lemma 7, the adversary is able to serve requests given by each  $\text{strategy}(k-1, S \setminus \{p\})$  with cost at most  $c_{k-1}$  without moving its heaviest server. In other words, the contribution of the adversary's servers other than the heaviest towards its cost per  $\text{strategy}$  call is at most  $c_{k-1}$ .

Thus, the adversary's cost per  $\text{strategy}$  call made by the procedure `adversary` is at most  $\beta^{k-1}/((n_{k-1} + 1) \cdot H(n_{k-1})) + c_{k-1}$ , which, by unrolling the recurrence in the statement of Lemma 7, is given by

$$\frac{\beta^{k-1}}{(n_{k-1} + 1) \cdot H(n_{k-1})} + c_{k-1} = \frac{\beta^{k-1}}{(n_{k-1} + 1) \cdot H(n_{k-1})} + \beta^{k-2} \cdot \sum_{i=1}^{k-1} \prod_{j=i}^{k-2} \left( \left\lceil \frac{n_j}{2} \right\rceil + 1 \right).$$

Let  $\varepsilon$  be an arbitrarily small positive number. By choosing

$$\beta = \lceil \varepsilon^{-1} \rceil \cdot (n_{k-1} + 1) \cdot H(n_{k-1}) \cdot \sum_{i=1}^{k-1} \prod_{j=i}^{k-2} \left( \left\lceil \frac{n_j}{2} \right\rceil + 1 \right),$$

the adversary's cost per  $\text{strategy}$  call is bounded from above by

$$\frac{\beta^{k-1} \cdot (1 + \varepsilon)}{(n_{k-1} + 1) \cdot H(n_{k-1})}.$$

On the other hand, recall from Corollary 6 that the expected cost of the algorithm per  $\text{strategy}$  call made by the procedure `adversary` is at least

$$\frac{(\beta - 1)^{k-1}}{n_{k-1} + 1} = \frac{\beta^{k-1}}{n_{k-1} + 1} \cdot \left( 1 - \frac{1}{\beta} \right)^{k-1} \geq \frac{\beta^{k-1}}{n_{k-1} + 1} \cdot \left( 1 - \frac{k-1}{\beta} \right) \geq \frac{\beta^{k-1} \cdot (1 - \varepsilon)}{n_{k-1} + 1},$$

because  $\beta \gg k/\varepsilon$ . Thus, modulo the  $(1 \pm \varepsilon)$  factors, the algorithm's cost per  $\text{strategy}$  call is at least  $H(n_{k-1})$  times the adversary's cost per  $\text{strategy}$  call. Since  $\varepsilon$  is arbitrarily small, we use Proposition 3 to conclude that the competitive ratio of any randomized online algorithm for weighted  $k$ -server on uniform metrics is at least  $H(n_{k-1})$ . ◀

## 4 Concluding Remarks

Given our lower bound on the randomized competitive ratio of weighted  $k$ -server on uniform metric spaces, the gap between the known upper and lower bounds has reduced from three orders of exponentiation to one. The natural question that needs to be investigated is to determine the randomized competitive ratio, or at least, prove upper and lower bounds that match in the order of exponentiation.



Our result also sheds light on the randomized competitive ratio of a generalization of the weighted  $k$ -server problem on uniform metrics called the generalized  $k$ -server problem on weighted uniform metrics. In this problem  $k$  servers are restricted to move in  $k$  different uniform metric spaces that are scaled copies of one another. A request contains one point from each copy and to serve it, one of the points must be covered by the server moving in its copy. Our lower bound directly applies to the generalized  $k$ -server problem on weighted uniform metrics and improves the previously known lower bound<sup>2</sup> of  $\Omega(k/\log^2 k)$  by Bansal et al. [3] to exponential in  $k$ . This also proves that the generalized  $k$ -server problem on weighted uniform metrics is qualitatively harder than its unweighted counterpart, the generalized  $k$ -server problem on uniform metrics, which has randomized competitive ratio  $O(k^2 \log k)$  due to Beinowski, Jeż, and Schmidt [4].

---

### References

- 1 Dimitris Achlioptas, Marek Chrobak, and John Noga. Competitive analysis of randomized paging algorithms. *Theor. Comput. Sci.*, 234(1-2):203–218, 2000. doi:10.1016/S0304-3975(98)00116-9.
- 2 Nikhil Bansal, Marek Eliás, and Grigorios Koumoutsos. Weighted  $k$ -server bounds via combinatorial dichotomies. In *FOCS*, pages 493–504, 2017. doi:10.1109/FOCS.2017.52.
- 3 Nikhil Bansal, Marek Eliás, Grigorios Koumoutsos, and Jesper Nederlof. Competitive algorithms for generalized  $k$ -server in uniform metrics. In *SODA*, pages 992–1001, 2018. doi:10.1137/1.9781611975031.64.
- 4 Marcin Bienkowski, Łukasz Jeż, and Paweł Schmidt. Slaying hydrae: Improved bounds for generalized  $k$ -server in uniform metrics. In *ISAAC*, volume 149 of *LIPICs*, pages 14:1–14:14, 2019. doi:10.4230/LIPICs.ISAAC.2019.14.
- 5 Sébastien Bubeck, Michael B. Cohen, Yin Tat Lee, James R. Lee, and Aleksander Madry.  $k$ -server via multiscale entropic regularization. In *STOC*, pages 3–16, 2018. doi:10.1145/3188745.3188798.
- 6 Ashish Chiplunkar and Sundar Vishwanathan. Randomized memoryless algorithms for the weighted and the generalized  $k$ -server problems. *ACM Trans. Algorithms*, 16(1):14:1–14:28, 2020. doi:10.1145/3365002.
- 7 Marek Chrobak and Jiří Sgall. The weighted 2-server problem. *Theoretical Computer Science*, 324(2-3):289–312, 2004. doi:10.1016/j.tcs.2004.05.020.
- 8 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. doi:10.1016/0196-6774(91)90041-V.
- 9 Amos Fiat and Moty Ricklin. Competitive algorithms for the weighted server problem. *Theoretical Computer Science*, 130(1):85–99, 1994. doi:10.1016/0304-3975(94)90154-6.
- 10 Elias Koutsoupias and Christos H. Papadimitriou. On the  $k$ -server conjecture. *J. ACM*, 42(5):971–983, 1995. doi:10.1145/210118.210128.
- 11 James R. Lee. Fusible HSTs and the randomized  $k$ -server conjecture. In *FOCS*, pages 438–449, 2018. doi:10.1109/FOCS.2018.00049.
- 12 Mark S. Manasse, Lyle A. McGeoch, and Daniel Dominic Sleator. Competitive algorithms for on-line problems. In *STOC*, pages 322–333, 1988. doi:10.1145/62212.62243.
- 13 Lyle A. McGeoch and Daniel Dominic Sleator. A strongly competitive randomized paging algorithm. *Algorithmica*, 6(6):816–825, 1991. doi:10.1007/BF01759073.
- 14 René Sitters. The generalized work function algorithm is competitive for the generalized 2-server problem. *SIAM J. Comput.*, 43(1):96–125, 2014. doi:10.1137/120885309.
- 15 Neal E. Young. On-line file caching. *Algorithmica*, 33(3):371–383, 2002. doi:10.1007/s00453-001-0124-5.

---

<sup>2</sup> This bound, in fact, holds for the unweighted counterpart, and to the best of the authors’ knowledge, no better bound for the weighted problem was known.



## A Set-system Construction

► **Lemma 4.** *Let  $\ell \in \mathbb{N}$  and let  $P$  be a set of  $n_\ell$  points. There exists a set-system  $\mathcal{Q}_\ell \subseteq 2^P$  satisfying the following properties.*

1.  $\mathcal{Q}_\ell$  contains  $\lceil n_{\ell-1}/2 \rceil + 1$  sets, each of size  $n_{\ell-1}$ .
2. For every  $p \in P$ , there exists a set in  $\mathcal{Q}_\ell$  not containing  $p$ .
3. For every  $p \in P$ , there exists a  $q \in P$  such that every set in  $\mathcal{Q}_\ell$  contains at least one of  $p$  and  $q$ .

**Proof (Bansal et al. [2]).** Construct the set-system  $\mathcal{Q}_\ell$  as follows. Recall that

$$|P| = n_\ell = \left( \left\lceil \frac{n_{\ell-1}}{2} \right\rceil + 1 \right) \cdot \left( \left\lfloor \frac{n_{\ell-1}}{2} \right\rfloor + 1 \right).$$

Let  $M$  be an arbitrary subset of  $P$  having size  $\lceil n_{\ell-1}/2 \rceil + 1$ , so that

$$|P \setminus M| = \left( \left\lceil \frac{n_{\ell-1}}{2} \right\rceil + 1 \right) \cdot \left\lfloor \frac{n_{\ell-1}}{2} \right\rfloor.$$

Partition  $P \setminus M$  into  $\lfloor n_{\ell-1}/2 \rfloor + 1$  sets of size  $\lfloor n_{\ell-1}/2 \rfloor$  each, and for each  $r \in M$ , name a distinct set in the partition  $P'_r$ . Next, for each  $r \in M$ , define  $P_r = (M \setminus \{r\}) \cup P'_r$ , and let  $\mathcal{Q}_\ell = \{P_r \mid r \in M\}$ .

We now prove that  $\mathcal{Q}_\ell$  indeed satisfies the required properties. First, the number of sets in  $\mathcal{Q}_\ell$  is equal to  $|M| = \lceil n_{\ell-1}/2 \rceil + 1$ , and the size of each set  $P_r \in \mathcal{Q}_\ell$  is

$$|P_r| = |M| - 1 + |P'_r| = \left\lceil \frac{n_{\ell-1}}{2} \right\rceil + \left\lfloor \frac{n_{\ell-1}}{2} \right\rfloor = n_{\ell-1}.$$

For the second property, observe that a point  $p \in M$  is not contained in the corresponding set  $P_p \in \mathcal{Q}_\ell$ , whereas for a point  $p \in P'_r$ , the only set in  $\mathcal{Q}_\ell$  that contains  $p$  is  $P_r$ . For the third property, if  $p \in M$ , define  $q$  to be any other point in  $M$ , and if  $p \in P'_r$ , define  $q = r$ , and check that the property is indeed satisfied. ◀

## B Analysis of the Algorithm's Movements

We present the proof of Lemma 5 here, for which we need the following lemma.

► **Lemma 8.** *Let  $Z_1$  and  $Z_2$  be non-negative random variables and  $E$  be an event on a common sample space such that  $\mathbb{E}[Z_1 \mid E] \geq b$  and  $\mathbb{E}[Z_2 \mid \neg E] \geq b$  for some real number  $b$ . Then  $\mathbb{E}[Z_1 + Z_2] \geq b$ .*

**Proof.** We have,

$$\mathbb{E}[Z_1 + Z_2] = \mathbb{E}[Z_1 + Z_2 \mid E] \cdot \Pr[E] + \mathbb{E}[Z_1 + Z_2 \mid \neg E] \cdot \Pr[\neg E].$$

Since  $Z_1$  and  $Z_2$  are non-negative, we have,

$$\mathbb{E}[Z_1 + Z_2] \geq \mathbb{E}[Z_1 \mid E] \cdot \Pr[E] + \mathbb{E}[Z_2 \mid \neg E] \cdot \Pr[\neg E] \geq b \cdot \Pr[E] + b \cdot \Pr[\neg E] = b,$$

as required. ◀

► **Lemma 5.** *For every  $\ell \in \{0, \dots, k-1\}$  the following holds. Let  $\rho_0$  be an arbitrary sequence of requests and  $L$  be the set of positions of the algorithm's heaviest  $k - \ell$  servers after serving  $\rho_0$ . Let  $P$  be an arbitrary set of  $n_\ell$  points disjoint from  $L$ . Suppose  $\rho_0$  is followed by a random sequence  $\rho$  of requests given by a  $\text{strategy}(\ell, P)$  call. For  $i = 1, \dots, k$ , let the random*

## 9:10 The Randomized Competitive Ratio of Weighted $k$ -Server Is at Least Exponential

variable  $X_i$  denote the number of movements of the algorithm's  $i$ 'th lightest server while the algorithm serves  $\rho$ . Then we have,

$$\sum_{i=1}^k \beta^{\min(i-1, \ell)} \cdot \mathbb{E}[X_i] \geq (\beta - 1)^\ell.$$

**Proof.** We prove the claim by induction on  $\ell$ . For the base case, suppose  $\ell = 0$ . Then  $|P| = 1$ , and we are assured that  $L$ , the set of points occupied by all the algorithm's servers, is disjoint from  $P$ . In other words, none of the algorithm's servers occupies the unique point in  $P$ . Therefore, to serve the one request given by  $\text{strategy}(0, P)$ , the algorithm must move at least one of its servers, and thus,

$$\sum_{i=1}^k \beta^{\min(i-1, \ell)} \cdot \mathbb{E}[X_i] = \sum_{i=1}^k \mathbb{E}[X_i] \geq 1 = (\beta - 1)^\ell,$$

as required.

For the inductive case, suppose  $\ell > 0$ . We are assured that except for the lightest  $\ell$  servers, none of the servers of the algorithm occupies any point in  $P$  at the time the  $\text{strategy}(\ell, P)$  call is made. This call makes  $m = (\beta - 1) \cdot (\lceil n_{\ell-1}/2 \rceil + 1) = (\beta - 1) \cdot |\mathcal{Q}_\ell|$  recursive calls. For  $i = 1, \dots, k$  and  $j = 1, \dots, m$ , let the random variable  $Y_i^j$  denote the number of movements of the algorithm's  $i$ 'th lightest server to serve requests from the  $j$ 'th recursive call. Thus, for all  $i$ ,  $X_i = \sum_{j=1}^m Y_i^j$ .

Consider an arbitrary  $j \in \{1, \dots, m\}$ . Let  $E_j$  denote the event that the random variables  $Y_i^{j'}$  are all 0 for all  $i > \ell$  and  $j' < j$ . In words,  $E_j$  is the event that none of the algorithm's heaviest  $k - \ell$  servers moves during the first  $j - 1$  recursive calls. Recall that originally these servers did not occupy any point in  $P$ . Therefore, if  $E_j$  happens, these servers are guaranteed to be out of the set  $P' \subseteq P$  on which the  $j$ 'th recursive call is made. Next, let  $E'_j$  denote the event that the  $j$ 'th recursive call is made on a set  $P'$  that does not contain the position of the algorithm's  $\ell$ 'th lightest server after the first  $j - 1$  recursive calls. Thus, if both  $E_j$  and  $E'_j$  happen, then  $P'$  is disjoint from the set of positions of the algorithm's  $k - \ell + 1$  heaviest servers. We can then apply the induction hypothesis to get,

$$\sum_{i=1}^k \beta^{\min(i-1, \ell-1)} \cdot \mathbb{E}[Y_i^j \mid E_j \wedge E'_j] \geq (\beta - 1)^{\ell-1}. \quad (1)$$

Next, let us understand the behavior of the random variables  $Y_i^j$  conditioned on  $E_j$  only. We have,

$$\mathbb{E}[Y_i^j \mid E_j] \geq \mathbb{E}[Y_i^j \mid E_j \wedge E'_j] \cdot \Pr[E'_j \mid E_j] \geq \frac{\mathbb{E}[Y_i^j \mid E_j \wedge E'_j]}{|\mathcal{Q}_\ell|}. \quad (2)$$

Here, the first inequality holds because  $Y_i^j$  is non-negative. The second inequality holds because, by the second property of the set-system  $\mathcal{Q}_\ell$  given by Lemma 4, for every possible history before the  $j$ 'th recursive call,  $\mathcal{Q}_\ell$  contains at least one set which does not contain the position of the algorithm's  $\ell$ 'th lightest server. This implies  $\Pr[E'_j \mid E_j] \geq 1/|\mathcal{Q}_\ell|$ . From Equation 1 and Equation 2, we get,

$$\sum_{i=1}^k \beta^{\min(i-1, \ell-1)} \cdot \mathbb{E}[Y_i^j \mid E_j] \geq \frac{(\beta - 1)^{\ell-1}}{|\mathcal{Q}_\ell|}. \quad (3)$$

By the non-negativity of the random variables  $Y_i^{j'}$  and the definition of  $E_j$ , we trivially have,

$$\sum_{i=\ell+1}^k \sum_{j'=1}^m \mathbb{E}[Y_i^{j'} \mid \neg E_j] \geq \sum_{i=\ell+1}^k \sum_{j'=1}^{j-1} \mathbb{E}[Y_i^{j'} \mid \neg E_j] \geq 1,$$

and hence,

$$\frac{(\beta - 1)^{\ell-1}}{|\mathcal{Q}_\ell|} \cdot \sum_{i=\ell+1}^k \sum_{j'=1}^m \mathbb{E}[Y_i^{j'} \mid \neg E_j] \geq \frac{(\beta - 1)^{\ell-1}}{|\mathcal{Q}_\ell|}. \quad (4)$$

From Equation 3 and Equation 4, using Lemma 8, we get,

$$\sum_{i=1}^k \beta^{\min(i-1, \ell-1)} \cdot \mathbb{E}[Y_i^j] + \frac{(\beta - 1)^{\ell-1}}{|\mathcal{Q}_\ell|} \cdot \sum_{i=\ell+1}^k \sum_{j'=1}^m \mathbb{E}[Y_i^{j'}] \geq \frac{(\beta - 1)^{\ell-1}}{|\mathcal{Q}_\ell|}.$$

The above inequality holds for all  $j \in \{1, \dots, m\}$ . Summing up over all  $j$  and recalling  $X_i = \sum_{j=1}^m Y_i^j$ , we get,

$$\sum_{i=1}^k \beta^{\min(i-1, \ell-1)} \cdot \mathbb{E}[X_i] + m \cdot \frac{(\beta - 1)^{\ell-1}}{|\mathcal{Q}_\ell|} \cdot \sum_{i=\ell+1}^k \mathbb{E}[X_i] \geq m \cdot \frac{(\beta - 1)^{\ell-1}}{|\mathcal{Q}_\ell|}.$$

Recall that  $m = (\beta - 1) \cdot |\mathcal{Q}_\ell|$ . Thus,

$$\sum_{i=1}^k \beta^{\min(i-1, \ell-1)} \cdot \mathbb{E}[X_i] + (\beta - 1)^\ell \cdot \sum_{i=\ell+1}^k \mathbb{E}[X_i] \geq (\beta - 1)^\ell. \quad (5)$$

Finally, note that for  $i > \ell$ ,  $\min(i - 1, \ell - 1) = \ell - 1$ , and since  $\ell \geq 1$ , we have  $(\beta - 1)^\ell \leq \beta^{\ell-1}(\beta - 1)$ . Therefore, the multiplier of the  $\mathbb{E}[X_i]$  term in Equation 5 is bounded as,

$$\beta^{\min(i-1, \ell-1)} + (\beta - 1)^\ell \leq \beta^{\ell-1} + \beta^{\ell-1}(\beta - 1) = \beta^\ell = \beta^{\min(i-1, \ell)}.$$

On the other hand, for  $i \leq \ell$ ,  $\min(i - 1, \ell - 1) = \min(i - 1, \ell)$ . Thus, for all  $i \in \{1, \dots, k\}$ , the multiplier of the  $\mathbb{E}[X_i]$  term in Equation 5 is at most  $\beta^{\min(i-1, \ell)}$ . Since the  $X_i$ 's are all non-negative, we get,

$$\sum_{i=1}^k \beta^{\min(i-1, \ell)} \cdot \mathbb{E}[X_i] \geq (\beta - 1)^\ell,$$

as required. ◀