

Stochastic k -Server: How Should Uber Work?^{*†}

Sina Dehghani¹, Soheil Ehsani², Mohammad Hajiaghayi³,
Vahid Liaghat⁴, and Saeed Seddighin⁵

- 1 University of Maryland, College Park, MD, USA
dehghani@cs.umd.edu
- 2 University of Maryland, College Park, MD, USA
ehsani@cs.umd.edu
- 3 University of Maryland, College Park, MD, USA
hajiagha@cs.umd.edu
- 4 Facebook, Menlo Park, CA, USA
vliaghat@gmail.com
- 5 University of Maryland, College Park, MD, USA
saeedrez@cs.umd.edu

Abstract

In this paper we study a stochastic variant of the celebrated k -server problem. In the k -server problem, we are required to minimize the total movement of k servers that are serving an on-line sequence of t requests in a metric. In the stochastic setting we are given t independent distributions $\langle P_1, P_2, \dots, P_t \rangle$ in advance, and at every time step i a request is drawn from P_i .

Designing the optimal online algorithm in such setting is NP-hard, therefore the emphasis of our work is on designing an approximately optimal online algorithm. We first show a structural characterization for a certain class of *non-adaptive* online algorithms. We prove that in general metrics, the best of such algorithms has a cost of no worse than three times that of the optimal online algorithm. Next, we present an integer program that finds the optimal algorithm of this class for any arbitrary metric. Finally by rounding the solution of the linear relaxation of this program, we present an online algorithm for the stochastic k -server problem with an approximation factor of 3 in the line and circle metrics and factor of $O(\log n)$ in a general metric of size n . In this way, we achieve an approximation factor that is independent of k , the number of servers.

Moreover, we define the *Uber* problem, motivated by extraordinary growth of online network transportation services. In the Uber problem, each demand consists of two points -a source and a destination- in the metric. Serving a demand is to move a server to its source and then to its destination. The objective is again minimizing the total movement of the k given servers. We show that given an α -approximation algorithm for the k -server problem, we can obtain an $(\alpha + 2)$ -approximation algorithm for the Uber problem. Motivated by the fact that demands are usually highly correlated with the time (e.g. what day of the week or what time of the day the demand has arrived), we study the *stochastic Uber* problem. Using our results for stochastic k -server we can obtain a 5-approximation algorithm for the stochastic Uber problem in line and circle metrics, and a $O(\log n)$ -approximation algorithm for general metrics.

Furthermore, we extend our results to the correlated setting where the probability of a request arriving at a certain point depends not only on the time step but also on the previously arrived requests.

1998 ACM Subject Classification Computer science education

* A full version of the paper is available at <https://arxiv.org/abs/1705.05755>.

† Supported in part by NSF CAREER award CCF-1053605, NSF BIGDATA grant IIS-1546108, NSF AF:Medium grant CCF-1161365, DARPA GRAPHS/AFOSR grant FA9550-12-1-0423, and another DARPA SIMPLEX grant.



© Sina Dehghani, Soheil Ehsani, Mohammadtaghi Hajiaghayi, Vahid Liaghat, and Saeed Seddighin;
licensed under Creative Commons License CC-BY

44th International Colloquium on Automata, Languages, and Programming (ICALP 2017).

Editors: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl;

Article No. 126; pp. 126:1–126:14



Leibniz International Proceedings in Informatics
LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Keywords and phrases k -server, stochastic, competitive ratio, online algorithm, Uber

Digital Object Identifier 10.4230/LIPIcs.ICALP.2017.126

1 Introduction

The k -server problem is one of the most fundamental problems in online computation that has been extensively studied in the past decades. In the k -server problem we have k mobile servers on a metric space \mathcal{M} . We receive an online sequence of t requests where the i^{th} request is a point $r_i \in \mathcal{M}$. Upon the arrival of r_i , we need to move a server to r_i , at a cost equal to the distance from the current position of the server to r_i . The goal is to minimize the total cost of serving all requests.

Manasse, McGeoch, and Sleator [31] introduced the k -server problem as a natural generalization of several online problems, and a building block for other problems such as the metrical task systems. They considered the adversarial model, in which the online algorithm has no knowledge of the future requests. Following the proposition of Sleator and Tarjan [34], they evaluate the performance of an online algorithm using competitive analysis. In this model, an online algorithm ALG is compared to an *offline* optimum algorithm OPT which is aware of the entire input in advance. For a sequence of requests ρ , let $|\text{ALG}(\rho)|$ and $|\text{OPT}(\rho)|$ denote the total cost of ALG and OPT for serving ρ . An algorithm is *c-competitive* if for every ρ , $|\text{ALG}(\rho)| \leq c|\text{OPT}(\rho)| + c_0$ where c_0 is independent of ρ .

Manasse *et al.* [31] showed a lower bound of k for the competitive ratio of any deterministic algorithm in any metric space with at least $k + 1$ points. The celebrated *k-server conjecture* states that this bound is tight for general metrics. For several years the known upper bounds were all exponential in k , until a major breakthrough was achieved by Koutsoupias and Papadimitriou [29], who showed that the so-called *work function algorithm* is $(2k - 1)$ -competitive. Proving the tight competitive ratio has been the “holy grail” of the field in the past two decades. This challenge has led to the study of the problem in special spaces such as the uniform metric (also known as the paging problem), line, circle, and trees metrics (see [15, 16] and references therein). We also refer the reader to Section 1.3 for a short survey of randomized algorithms, particularly the recent result of Bansal, Buchbinder, Madry, and Naor [7] which achieves the competitive ratio of $O(\log^3 n \log^2 k)$ for discrete metrics that comprise n points.

The line metric (or Euclidean 1-dimensional metric space) is of particular interest for developing new ideas. Chrobak, Karloof, Payne, and Vishwnathan [15] were the first to settle the conjecture in the line by designing an elegant k -competitive algorithm. Chrobak and Larmore [16] generalized this approach to tree metrics. Later, Bartal and Koutsoupias [10] proved that the work function algorithm is also k -competitive in line. Focusing on the special case of $k = 2$ in line, Bartal *et al.* [9] show that, using randomized algorithms, one can break the barrier of lower bound k by giving a 1.98-competitive algorithm for the case where we only have two servers.

Despite the strong lower bounds for the k -server problem, there are heuristics algorithms that are *constant* competitive in practice. For example, for the paging problem- the special case of uniform metric- the least recently used (LRU) strategy is shown to be experimentally constant competitive (see Section 1.3). In this paper we present an algorithm and run it on real world data to measure its empirical performance. In particular we use the distribution of car accidents obtained from road safety data. Our experiments illustrate our algorithm is performing even better in practice. See the full version of the paper for more details.

The idea of comparing the performance of an online algorithm (with zero-knowledge of the future) to the request-aware offline optimum has led to crisp and clean solutions. However, that is not without its downsides. The results in the online model are often very pessimistic leading to theoretical guarantees that are hardly comparable to experimental results. Indeed, one way to tighten this gap is to use stochastic information about the input data as we describe in this paper.

We should also point out that the competitive analysis is not the only possible or necessarily the most suitable approach for this problem. Since the distributions from which the input is generated are known, one can use dynamic programming (or enumeration of future events) to derive the optimal movement of servers. Unfortunately, finding such an optimal online solution using the distributions is an NP-hard problem¹, thus the dynamic programming or any other approach takes exponential time. This raises the question that how well one can perform in comparison to the best online solution. In the rest of the paper we formally define the model and address this question.

A natural and well-motivated generalization of k -server is to assume the demands are two points instead of just one, consisting of a source and a destination. To serve a demand we need to move a server to the source and then move it to the destination. We call this problem the *Uber* problem. One can see, the *Uber* problem is the same as k -server when the sources and the destinations are the same. We also show that, given an α -approximation algorithm for the k -server problem, we can obtain a $(\alpha + 2)$ -approximation algorithm for the *Uber* problem. Thus our results for k -server also apply to the *Uber* problem.

1.1 The Stochastic Model

In this paper, we study the *stochastic k -server problem* where the input is not chosen adversarially, but consists of draws from given probability distributions. This problem has lots of applications such as network transportations and equipment replacement in data centers. The current mega data centers contain hundreds of thousands of servers and switches with limited life-span. For example servers usually retire after at most three years. The only efficient way to scale up the maintenance in data centers is by automation, and robots are designed to handle maintenance tasks such as repairs or manual operations on servers. The replacement process can be modeled as requests that should be satisfied by robots, and robots can be modeled as servers. This problem also has applications in physical networks. As an example, suppose we model a shopping service (e.g. Google Express) as a k -server problem in which we receive an online sequence of shopping requests for different stores. We have k shopping cars (i.e., servers) that can serve the requests by traveling to the stores. It is quiet natural to assume that on a certain time of the week/day, the requests arrive from a distribution that can be discovered by analyzing the history. For example, an *Uber* request is more likely to be from suburb to midtown in the morning, and from midtown to suburb at night. We formalize this stochastic information as follows.

For every $i \in [1 \cdots t]$, a discrete probability distribution P_i is given in advance from which request r_i will be drawn at time step i . The distributions are chosen by the adversary and are assumed to be independent but not necessarily identical. This model is inspired by the

¹ Reduction from k -median to Stochastic k -server: to find the k median of set S of vertices, one can construct an instance of stochastic k -server with $t = 1$ and $P_1(v) = 1/|S|$ for every $v \in S$. The best initialization of the servers gives the optimum solution to k -median of S .

well-studied model of *prophet inequalities*² [30, 25]. As mentioned before, the case of line metric has proven to be a very interesting restricted case for studying the k -server problem. In this paper, we focus mainly on the class of line metric though our results carry over to circle metric and general metrics as well.

In the adversarial model, the competitive ratio seems to be the only well-defined notion for analyzing the performance of online algorithms. However, in the presence of stochastic information, one can derive a much better benchmark that allows us to make fine-grained distinctions between the online algorithms. We recall that in the offline setting, for a class of algorithms \mathcal{C} , the natural notion to measure the performance of an algorithm $\text{ALG} \in \mathcal{C}$ is the *approximation ratio* defined as the worse case ratio of $|\text{ALG}|$ to $|\text{OPT}(\mathcal{C})|$ where $\text{OPT}(\mathcal{C})$ is the optimal algorithm in the class. In this paper, we also measure the performance of an online algorithm by its approximation ratio—compared to the *optimal online solution*. We note that given distributions P_1, \dots, P_t , one can iteratively compute the optimal online solution by solving the following exponential-size dynamic program: for every $i \in [0 \dots t]$ and every possible placement A of k servers (called a *configuration*) on the metric, let $\tau(i, A)$ denote the minimum expected cost of an online algorithm for serving the first i requests and then moving the servers to configuration A . Note that $\tau(i, A)$ can inductively be computed via the following recursive formula

$$\tau(i, A) = \min_B \tau(i-1, B) + \mathbb{E}_{r_i \sim P_i} [\text{min. distance from } B \text{ to } A \text{ subject to serving } r_i] ,$$

where $\tau(0, A)$ is initially zero for every A .

1.2 Our Results³

Our first main result is designing a constant approximation algorithm in the line metric when the distributions for different time steps are not necessarily identical.

► **Theorem 1.** *There exists a 3-approximation online algorithm for the stochastic k -server problem in the line metric. The running time is polynomial in k and the sum of the sizes of the supports of input distributions. The same guarantee holds for the circle metric.*

For the general metric, we present an algorithm with a logarithmic approximation guarantee.

► **Theorem 2.** *There exists a $O(\log n)$ -approximation online algorithm for the stochastic k -server problem in a general metric of size n .*

We prove the theorems using two important structural results. The first key ingredient is a general reduction from class of online algorithms to a restricted class of *non-adaptive algorithms* while losing only a constant factor in the approximation ratio. Recall that a configuration is a placement of k -servers on the metric. We say an algorithm ALG is *non-adaptive* if it follows the following procedure: ALG pre-computes a sequence of configurations A_0, A_1, \dots, A_t . We start by placing the k -servers on A_0 . Upon the arrival of r_i , (i) we move the servers to configuration A_i ; next (ii) we move the closest server s to r_i ; and finally (iii) we return s to its original position in A_i . We first prove the following structural result.

² In the prophet inequality setting, given (not necessarily identical) distributions P_1, \dots, P_t , an online sequence of values x_1, \dots, x_n where x_i is drawn from P_i , an onlooker has to choose one item from the succession of the values, where x_i is revealed at step i . The onlooker can choose a value only at the time of arrival. The goal is to maximize the chosen value.

³ In the interest of space, we have omitted some of the proofs. We refer the reader to the full version of the paper in order to see all of the proofs.

► **Theorem 3.** *For the stochastic k -server problem in the general metric, the optimal non-adaptive online algorithm is within 3-approximation of the optimal online algorithm.*

Using the aforementioned reduction, we focus on designing the optimal non-adaptive algorithm. We begin by formulating the problem as an integer program. The second ingredient is to use the relaxation of this program to formalize a natural *fractional variant* of the problem. In this variant, a configuration is a fractional assignment of *server mass* to the points of the metric such that the total mass is k . To serve a request at point r_i , we need to move some of the mass to have at least one amount of server mass on r_i . The cost of moving the server mass is naturally defined as the integral of the movement of infinitesimal pieces of the server mass. By solving the linear relaxation of the integer program, we achieve the optimal fractional non-adaptive algorithm. We finally prove Theorems 1 and 2 by leveraging the following rounding techniques. The rounding method in line has been also observed by Türkoglu [35]. We provide the proof for the case of line in Section 5 for the sake of completeness. The rounding method for general metrics is via the well-known embedding of a metric into a distribution of well-separated trees while losing a logarithmic factor in the distortion. Bansal *et al.* [7] use a natural rounding method similar to that of Blum, Burch, and Kalai [12] to show that any fractional k -server movement on well-separated trees can be rounded to an integral counterpart by losing only a constant factor.

► **Theorem 4** (first proven in [35]). *Let ALG_f denote a fractional k -server algorithm in the line, or circle. One can use ALG_f to derive a randomized integral algorithm ALG such that for every request sequence σ , $\mathbb{E}[|\text{ALG}(\sigma)|] = |\text{ALG}_f(\sigma)|$. The expectation is over the internal randomness of ALG . Furthermore, in the stochastic model ALG can be derandomized.*

► **Theorem 5** (proven in [7]). *Let ALG_f denote a fractional k -server algorithm in any metric. One can use ALG_f to derive a randomized integral algorithm ALG such that for every request sequence σ , $\mathbb{E}[|\text{ALG}(\sigma)|] \leq O(\log n) |\text{ALG}_f(\sigma)|$.*

We also show that having an α -approximation algorithm for k -server, we can obtain a $(\alpha + 2)$ -approximation for the Uber problem, using a simple reduction.

► **Theorem 6.** *Let ALG denote an α -approximation algorithm for k -server. One can use ALG to derive a $(\alpha + 2)$ -approximation algorithm for the Uber problem.*

Proof. Consider an instance of the Uber problem I_U . Let s_i and t_i denote the i -th source and destination, respectively. We generate an instance of the k -server problem I_k by removing every t_i from I_U . In other words the demands are s_i 's. We use ALG to provide a solution for I_U as follows. For satisfying the i -th demand, we use ALG to move a server to s_i . Then using the shortest path from s_i to t_i , we move that server to t_i and then return it back to s_i . Let OPT_U and OPT_k denote the cost of the optimal solutions for I_U and I_k , respectively. Let $d(s_i, t_i)$ denote the distance of t_i from s_i in the metric. Let C denote the total movement of the servers. We have,

$$\begin{aligned} \text{OPT}_U &\geq \text{OPT}_k. \\ \text{OPT}_U &\geq \sum_i d(s_i, t_i). \\ C &\leq \alpha \text{OPT}_k + 2 \sum_i d(s_i, t_i) \leq (\alpha + 2) \text{OPT}_U. \end{aligned} \quad \blacktriangleleft$$

1.3 Further Related Work

The randomized algorithms often perform much better in the online paradigm. For the k -server problem, a lower bound of $\Omega(\log k)$ is shown by [28] for the competitive ratio of randomized algorithms in most common metrics. Despite the exponential gap, compared to the lower bound of deterministic algorithms, very little is known about the competitiveness of randomized algorithms. In fact, the only known algorithms with competitive ratios below k , work either in the uniform metric (also known as the paging problem [21, 32, 2, 8]), a metric comprising $k + 1$ points [23], and two servers on the line [9]. Two decades after the introduction of the k -server problem, a major breakthrough was achieved by Bansal *et al.* [7] in discrete metrics with sub-exponential size. If \mathcal{M} comprise n points, their randomized algorithm achieves a competitive ratio of $O(\log^3 n \log^2 k)$.

The case of uniform metric has been extensively studied under various stochastic models motivated by the applications in computer caching. Koutsoupias and Papadimitriou [29] consider two refinements of the competitive analysis for server problems. First, they consider the *diffuse adversary* model. In this model, at every step i the adversary chooses a distribution D_i over the uniform metric of the paging problem. Then the i^{th} request is drawn from D_i which needs to be served. The distribution D_i is not known to the online algorithm and it may depend on the previous requests. However, in their paper, they consider the case wherein it is guaranteed that for every point p , $D_i(p) \leq \epsilon$ for a small enough ϵ ; i.e., the next request is not predictable with absolute certainty for the adversary. The results of Koutsoupias and Papadimitriou and later Young [36] shows that the optimum competitive ratio in this setting is close to $1 + \Theta(k\epsilon)$.

The second refinement introduced in [29] restricts the optimal solution to having lookahead at most ℓ . Hence, one can define a *comparative ratio* which indicates the worst-case ratio of the cost of the best online solution to the best solution with lookahead ℓ . They show that for the k -server problem, and more generally the metrical task system problem, there are online algorithms that admit a comparative ratio of $2\ell + 1$; for some instances this ratio is tight.

Various other models of restricting the adversary (access graph model [14, 26, 22], fault rate model [27, 6, 19], etc) have also been considered for the paging problem (see [33, 11] and references therein for a further survey of these results). Unfortunately, many of the stochastic settings considered for the paging problem do not seem to have a natural generalization beyond the uniform metric setting. For example, in the diffuse adversary model, most of the studied distributions do not weaken the adversary in the general metric. In this paper, we look for polynomial-time *approximation* algorithms in the class of online algorithms that have access to the distributions.

We would like to mention that various online problems have been previously considered under prophet inequality model or i.i.d. model (where all distributions are identical). The maximum matching problem, scheduling, and online network design has been extensively studied in these models (see e.g. [4, 3, 5, 17, 1, 18]). In the graph connectivity problems, Garg, Gupta, Leonardi, and Sankowski [24] consider the online variants of Steiner tree and several related problems under the i.i.d. stochastic model. In the adversarial model, there exists an $\Omega(\log n)$ lower bound on the competitive ratio of any online algorithm, where n is the number of demands. However, Garg *et al.* show that under the i.i.d. assumption, these problems admit online algorithms with constant or $O(\log \log n)$ competitive ratios. We refer the reader to the excellent book by Borodin and El-Yaniv [13] for further study of online problems.

2 Preliminaries

In this section we formally define the stochastic k -server problem. The classical k -server problem is defined on a metric \mathcal{M} which consists of points that could be infinitely many. For every two points x and y in metric \mathcal{M} , let $d(x, y)$ denote the distance of x from y which is a symmetric function and satisfies the triangle inequality. More precisely for every three points x , y , and z we have

$$d(x, x) = 0 \tag{1}$$

$$d(x, y) = d(y, x) \tag{2}$$

$$d(x, y) + d(y, z) \geq d(x, z). \tag{3}$$

In the k -server problem the goal is to place k servers on k points of the metric, and move these servers to satisfy the requests. We refer to every placement of the servers on the metric points by a *configuration*. Let $\rho = \langle r_1, r_2, \dots, r_t \rangle$ be a sequence of requests, the goal of the k -server problem is to find configurations $\langle A_0, A_1, A_2, \dots, A_t \rangle$ such that for every i there exists a server on point r_i in configuration A_i . We say such a list of configurations is *valid* for the given list of requests. A valid sequence of configurations is optimal if $\sum d(A_{i-1}, A_i)$ is minimized where $d(X, Y)$ stands for the minimum cost of moving servers from configuration X to configuration Y . An optimal sequence $\langle A_0, A_1, \dots, A_t \rangle$ of configurations is called an optimal offline solution of OFKS(\mathcal{M}, ρ) when ρ is known in advance. We refer to the optimal cost of such movements with $|\text{OFKS}(\mathcal{M}, \rho)| = \sum d(A_{i-1}, A_i)$.

We also define the notion of *fractional configuration* as an assignment of the metric points to non-negative real numbers. More precisely, each number specifies a mass of fractional server on a point. Every fractional solution adheres to the following condition: The total sum of the values assigned to all points is exactly equal to k . Analogously, a fractional configuration serves a request r_i if there is a mass of size at least 1 of server assigned to point r_i . An offline fractional solution of the k -server problem for a given sequence of requests ρ is defined as a sequence of fractional configurations $\langle A_0, A_1, \dots, A_t \rangle$ such that A_i serves r_i .

In the online k -server problem, however, we are not given the whole sequence of requests in the beginning, but we will be informed of every request once its realization is drawn. An algorithm \mathcal{A} is an online algorithm for the k -server problem if it reports a configuration A_0 as an initial configuration and upon realization of every request r_i it returns a configuration A_i such that $\langle A_0, A_1, \dots, A_i \rangle$ is valid for $\langle r_1, r_2, \dots, r_i \rangle$. If \mathcal{A} is deterministic, it generates a unique sequence of configurations for every sequence of requests. Let $\mathcal{A}(\mathcal{M}, \rho)$ be the sequence that \mathcal{A} generates for requests in ρ and $|\mathcal{A}(\mathcal{M}, \rho)|$ denote its cost.

In the online stochastic k -server problem, in addition to metric \mathcal{M} , we are also given t independent probability distributions $\langle P_1, P_2, \dots, P_t \rangle$ which show the probability that every request r_i is realized on a point of the metric at each time. An algorithm \mathcal{A} is an online algorithm for such a setting, if it generates a configuration for every request r_i not solely based on $\langle r_1, r_2, \dots, r_i \rangle$ and $\langle A_0, A_1, \dots, A_{i-1} \rangle$ but also with respect to the probability distributions. Similarly, we define the cost of an online algorithm \mathcal{A} for a given sequence of requests ρ with $|\mathcal{A}(\mathcal{M}, \rho, \langle P_1, P_2, \dots, P_t \rangle)|$. We define the expected cost of an algorithm \mathcal{A} on metric \mathcal{M} and with probability distributions $\langle P_1, P_2, \dots, P_t \rangle$ by

$$|\mathcal{A}(\mathcal{M}, \langle P_1, P_2, \dots, P_t \rangle)| = \mathbb{E}_{\forall i, r_i \sim P_i} |\mathcal{A}(\mathcal{M}, \rho, \langle P_1, P_2, \dots, P_t \rangle)|.$$

For every metric \mathcal{M} and probability distributions $\langle P_1, P_2, \dots, P_t \rangle$ we refer to the online algorithm with the minimum expected cost by $\text{OPT}_{\mathcal{M}, \langle P_1, P_2, \dots, P_t \rangle}$.

An alternative way to represent a solution of the k -server problem is as a vector of configurations $\langle B_0, B_1, \dots, B_t \rangle$ such that B_i does not necessarily serve request r_i . The cost of such solution is equal to $\sum d(B_{i-1}, B_i) + \sum 2d(B_i, r_i)$ where $d(B_i, r_i)$ is the minimum distance of a server in configuration B_i to request r_i . The additional cost of $2d(B_i, r_i)$ can be thought of as moving a server from B_i to serve r_i and returning it back to its original position. Thus, every such representation of a solution can be transformed to the other representation. Similarly, $d(B_i, r_i)$ for a fractional configuration B_i is the minimum cost which is incurred by placing a mass 1 of server at point r_i . We use letter B for the configurations of such solutions throughout the paper.

In this paper the emphasis is on the stochastic k -server problem on the line metric. We define the line metric \mathcal{L} as a metric of points from $-\infty$ to $+\infty$ such that the distance of two points x and y is always equal to $|x - y|$. Moreover, we show that deterministic algorithms are as powerful as randomized algorithms in this setting, therefore we only focus on deterministic algorithms in this paper. Thus, from here on, we omit the term deterministic and every time we use the word algorithm we mean a deterministic algorithm unless otherwise is explicitly mentioned.

3 Structural Characterization

Recall that an online algorithm \mathcal{A} has to fulfill the task of reporting a configuration A_i upon arrival of request r_i based on $\langle A_0, A_1, \dots, A_{i-1} \rangle$, $\langle r_1, r_2, \dots, r_i \rangle$, and $\langle P_1, P_2, \dots, P_t \rangle$. We say an algorithm \mathcal{B} is *request oblivious*, if it reports configuration B_i regardless of request r_i . As such, \mathcal{B} generates configurations $\langle B_0, B_1, \dots, B_t \rangle$ for a sequence of requests $\langle r_1, r_2, \dots, r_t \rangle$ and the cost of such configuration is $\sum d(B_{i-1}, B_i) + \sum 2d(B_i, r_i)$. More precisely, no matter what request r_i is, \mathcal{B} will generate the same configuration for a given list of past configurations $\langle B_0, B_1, \dots, B_{i-1} \rangle$, a given sequence of past requests $\langle r_1, r_2, \dots, r_{i-1} \rangle$, and the sequence of probability distributions $\langle P_1, P_2, \dots, P_t \rangle$. In the following we show that every online algorithm \mathcal{A} can turn into a request oblivious algorithm $\mathcal{B}_{\mathcal{A}}$ that has a cost of at most $3|\mathcal{A}(\mathcal{M}, \rho, \langle P_1, P_2, \dots, P_t \rangle)|$ for a given sequence of requests ρ .

► **Lemma 7.** *Let \mathcal{A} be an online algorithm for the stochastic k -server problem. For any metric \mathcal{M} , there exists a request oblivious algorithm $\mathcal{B}_{\mathcal{A}}$ such that*

$$|\mathcal{B}_{\mathcal{A}}(\mathcal{M}, \langle P_1, P_2, \dots, P_t \rangle)| \leq 3|\mathcal{A}(\mathcal{M}, \langle P_1, P_2, \dots, P_t \rangle)|.$$

An immediate corollary of Lemma 7 is that the optimal request oblivious algorithm has a cost of at most $3|\text{OPT}_{\mathcal{M}, \langle P_1, P_2, \dots, P_t \rangle}(\mathcal{M}, \langle P_1, P_2, \dots, P_t \rangle)|$. Therefore, if we only focus on the request oblivious algorithms, we only lose a factor of 3 in comparison to the optimal online algorithm. The following lemma states a key structural lemma for an optimal request oblivious algorithm.

► **Lemma 8.** *For every request oblivious algorithm \mathcal{B} , there exists a randomized request oblivious algorithm \mathcal{B}' with the same expected cost which is not only oblivious to the last request, but also oblivious to all requests that have come prior to this.*

Proof. For any given request oblivious online algorithm \mathcal{B} , we construct an online algorithm \mathcal{B}' which is oblivious to all of the requests as follows: For an input $\langle B_1, B_2, \dots, B_{i-1} \rangle$ of configurations and probability distributions $\langle P_1, P_2, \dots, P_t \rangle$, draw a sequence of requests $\langle r_1, r_2, \dots, r_i \rangle$ from $\langle P_1, P_2, \dots, P_t \rangle$ conditioned on the constraint that \mathcal{B} would generate configurations $\langle B_1, B_2, \dots, B_{i-1} \rangle$ for requests $\langle r_1, r_2, \dots, r_{i-1} \rangle$. Now, report the output of \mathcal{B} for inputs $\langle B_1, B_2, \dots, B_{i-1} \rangle$, $\langle r_1, r_2, \dots, r_i \rangle$, and $\langle P_1, P_2, \dots, P_t \rangle$.

We define the cost of step i of algorithm B' as $d(B_{i-1}, B_i) + 2d(B_i, r_i)$. Due to the construction of algorithm B' , the expected cost of this algorithm at every step i for a random sequence of requests is equal to the expected cost of algorithm B for a random sequence of requests drawn from $\langle P_1, P_2, \dots, P_t \rangle$. Therefore, the expected cost of both algorithms for a random sequence of requests are equal and thus $|\mathcal{B}(\mathcal{M}, \langle P_1, P_2, \dots, P_t \rangle)| = |\mathcal{B}'(\mathcal{M}, \langle P_1, P_2, \dots, P_t \rangle)|$. ◀

Lemma 8 states that there always exists an optimal randomized request oblivious online algorithm that returns the configurations regardless of the requests. We call such an algorithm *non-adaptive*. Since a non-adaptive algorithm is indifferent to the sequence of the requests, we can assume it always generates a sequence of configurations just based on the distributions. For an optimal of such algorithms, all such sequence of configurations should be optimal as well. Therefore, there always exists an optimal non-adaptive online algorithm which is deterministic. By Lemma 7 not only do we know the optimal request oblivious algorithm is at most 3-approximation, but also the same holds for the optimal non-adaptive algorithm.

► **Theorem 9.** *There exists a sequence of configurations $\langle B_0, B_1, \dots, B_t \rangle$ such that an online algorithm which starts with B_0 and always returns configuration B_i upon arrival of request r_i has an approximation factor of at most 3.*

4 Fractional Solutions

In this section we provide a fractional online algorithm for the k -server problem that can be implemented in polynomial time. Note that by Theorem 9 we know that there exist configurations $\langle B_1, B_2, \dots, B_t \rangle$ such that the expected cost of a non-adaptive algorithm that always returns these configurations is at most 3 times the cost of an optimal online algorithm. Therefore, we write an integer program to find such configurations with the least expected cost. Next, we provide a relaxed LP of the integer program and show that every feasible solution of such LP corresponds to a fractional online algorithm for the stochastic k -server problem. Hence, solving such a linear program, that can be done in polynomial time, gives us a fractional online algorithm for the problem.

4.1 Linear Program

Recall that given t independent distributions $\langle P_1, \dots, P_t \rangle$ for online stochastic k -server, an adaptive algorithm can be represented by $t + 1$ configurations $\langle B_0, \dots, B_t \rangle$. Upon the arrival of each request r_i , we move the servers from configuration B_{i-1} to B_i and then one server serves r_i and goes back to its position in B_i . The objective is to find the configurations such that the cost of moving to new configurations in addition to the expected cost of serving the requests is minimized. Therefore the problem can be formulated in an offline manner. First we provide an integer program in order to find a vector of configurations with the least cost.

The decision variables of the program represent the configurations, the movement of servers from one configuration to another, and the way that each possible request is served. In particular, at each time step τ :

- For each node v there is a variable $b_{\tau,v} \in N$ denoting the number of servers on node v .
- For each pair of nodes u and v , there is a movement variable $f_{\tau,u,v} \in N$ denoting the number of servers going from u to v for the next round.
- For each node v and possible request node r , there is a variable $x_{\tau,v,r} \in \{0, 1\}$ denoting whether r is served by v or not.

In the following integer program, the first set of constraints ensures the number of servers on nodes at each time is updated correctly according to the movement variables. The second set of constraints ensures that each possible request is served by at least one server. The third set of constraints ensures that no possible request is served by an empty node. By the definition, the cost of a sequence of configurations $\langle B_0, \dots, B_t \rangle$ is $\sum_{i=1}^t d(B_{i-1}, B_i) + 2 \sum_{i=1}^t d(B_i, r_i)$. Thus the objective is to minimize the expression

$$\sum_{\tau} \sum_{u,v} f_{\tau,u,v} d(u,v) + 2 \sum_{\tau} \sum_v \sum_r x_{\tau,v,r} \Pr(z \sim P_{\tau} = r) d(v,r),$$

where $\Pr(z \sim P_{\tau} = r)$ denotes the probability that r is requested at time τ .

$$\begin{aligned} \min. \quad & \sum_{\tau} \sum_{u,v} f_{\tau,u,v} d(u,v) + 2 \sum_{\tau} \sum_v \sum_r x_{\tau,v,r} \Pr(z \sim P_{\tau} = r) d(v,r) \\ \forall \tau, v \quad & b_{\tau+1,v} = b_{\tau,v} + \sum_u f_{\tau,u,v} - \sum_u f_{\tau,v,u} \\ \forall \tau, u, v \quad & \sum_v x_{\tau,v,r} \geq 1. \\ \forall \tau, v, r \quad & x_{\tau,v,r} \leq b_{\tau,v}. \\ \forall \tau \quad & \sum_v b_{\tau,v} \leq k. \\ \forall \tau, v, r \quad & x_{\tau,v,r} \in \{0, 1\}. \\ \forall \tau, u, v \quad & f_{\tau,u,v} \in N. \\ \forall \tau, v \quad & b_{\tau,v} \in N. \end{aligned}$$

Now we consider the following relaxation of the above integer program.

$$\begin{aligned} \min. \quad & \sum_{\tau} \sum_{u,v} f_{\tau,u,v} d(u,v) + 2 \sum_{\tau} \sum_v \sum_r x_{\tau,v,r} \Pr(z \sim P_{\tau} = r) d(v,r) \\ \forall \tau, v \quad & b_{\tau+1,v} = b_{\tau,v} + \sum_u f_{\tau,u,v} - \sum_u f_{\tau,v,u} \\ \forall \tau, u, v \quad & \sum_v x_{\tau,v,r} \geq 1. \\ \forall \tau, v, r \quad & x_{\tau,v,r} \leq b_{\tau,v}. \\ \forall \tau \quad & \sum_v b_{\tau,v} \leq k. \end{aligned}$$

5 Reduction from Integral k -server to Fractional k -server

In this section we show how we can obtain an integral algorithm for the stochastic k -server problem from a fractional algorithm. We first show that every fractional algorithm for the line metric can be modified to an integral algorithm with the same cost. Next, we study the problem on HST metrics; we give a rounding method that produces an integral algorithm from a fractional algorithm while losing a constant factor. Finally, we leverage the previously known embedding techniques to show every metric can be embedded into HST's with a distortion of at most $O(\log n)$. This will lead to a rounding method for obtaining an integral algorithm from every fractional algorithm on general metrics while losing a factor of at most $O(\log n)$. Combining this with the 3 approximation fractional algorithm that we provide

in Section 4, we achieve an $O(\log n)$ approximation algorithm for the stochastic k -server problem on general graphs.

5.1 Integrals Are as Strong as Fractionals On the Line

In this section we show every fractional algorithm on the line metric can be derandomized to an integral solution with the same expected cost. The rounding method is as follows: For every fractional configuration A , we provide an integral configuration $I(A)$ such that (i) the distance of two configurations A_1 and A_2 is equal to the expected distance of two configurations $I(A_1)$ and $I(A_2)$. (ii) for every point x in the metric that A has a server mass of size at least 1 on x , there exists a server on point x in $I(A)$.

Let for every point x in the metric, $A(v)$ denote the amount of server mass on node v of the line. For every fractional configuration B , we define a mass function $f_A : (0, k] \rightarrow V$ as follows. $f_A(x) = v_j$ if and only if j is the minimum integer such that $\sum_{i=1}^{j-1} A(i) < x$ and $\sum_{i=1}^j A(i) \geq x$. Intuitively, if one gathers the server mass by sweeping the line from left to right, $f_A(x)$ is the first position on which we have gathered x amount of server mass. The rounding algorithm is as follows:

- Pick a random real number r in the interval $[0, 1)$.
- $I(A)$ contains k servers on positions $f_A(r), f_A(r + 1), f_A(r + 2), \dots, f_A(r + k - 1)$.

Note that the rounding method uses the same r for all of the configurations. More precisely, we draw r from $[0, 1)$ at first and use this number to construct the integral configurations from fractional configurations. The following two lemmas show that both of the properties hold for the rounding algorithm we proposed.

► **Lemma 10.** *Let A be a fractional configuration and x be a point such that $A(x) \geq 1$. Then $I(A)$ has a server on x .*

Proof. Due to the construction of our rounding method, for every two consecutive servers a and b in $I(A)$, the total mass of servers after a and before b in the fractional solution is less than 1. Therefore, $I(A)$ should put a server on point x , otherwise the total mass of servers in the fractional solution between the first server before x and the first server after x would be at least 1. ◀

The next lemma shows that the rounding preserves the distances between the configurations in expectation.

► **Lemma 11.** *Let A_1 and A_2 be two fractional configurations and $|A_1 - A_2|$ be their distance. The following holds for the distances of the configurations*

$$\mathbb{E}|I(A_1) - I(A_2)| = |A_1 - A_2|.$$

Proof. The key point behind the proof of this lemma is that the distance of two fractional configurations A_1 and A_2 can be formulated as follows

$$|A_1 - A_2| = \int_0^1 |I_\omega(A_1) - I_\omega(A_2)| d\omega$$

where $I_\omega(A)$ stands for an integral configurations which places the servers on points $f_A(\omega), f_A(\omega + 1), f_A(\omega + 2), \dots, f_A(\omega + k - 1)$. Since at the beginning of the rounding method we draw r uniformly at random, the expected distance of the two rounded configurations is exactly equal to

$$\int_0^1 |I_\omega(A_1) - I_\omega(A_2)| d\omega$$

which is equal to the distance of A_1 from A_2 . ◀

► **Theorem 12.** *For any given fractional online algorithm \mathcal{A} for the k -server problem on the line metric, there exists an online integral solution for the same problem with the same expected cost.*

5.2 Reduction for General Graphs

An HST is a undirected rooted tree in which every leaf represents a point in the metric and the distance of a pair of points in the metric is equal to the distance of the corresponding leaves in the tree. In an HST, weights of the edges are uniquely determined by the depth of the vertices they connect. More precisely, in a σ -HST the weight of an edges between a vertex v and its children is equal to σ^{h-d_v} where h stands for the height of the tree and d_v denotes the depth of vertex v .

Since HSTs are very well structured, designing algorithms on HSTs is relatively easier in comparison to a more complex metric. Therefore, a classic method for alleviating the complexity of the problems is to first embed the metrics into HSTs with a low distortion and then solve the problems on these trees.

Perhaps the most important property of the HSTs is the following:

► **Observation 1.** For every pair of leaves $u, v \in T$ of an HST, the distance of u and v is uniquely determined by the depth of their deepest common ancestor.

Note that, the higher the depth of the common ancestor is, the lower the distance of the leaves will be. Therefore, the closest leaves to a leaf v are the ones that share the most common ancestors with v . Bansal *et al.* propose a method for rounding every fractional solution of the k -server problem to an integral solution losing at most a constant factor [7].

► **Theorem 13** ([7]). *Let T be a σ -HST with n leaves, $\sigma > 5$, and let $A = \langle A_0, A_1, A_2, \dots, A_t \rangle$ be a sequence of fractional configurations. There is an online procedure that maintains a sequence of randomized k -server configurations $S = \langle S_0, S_1, S_2, \dots, S_t \rangle$ satisfying the following two properties:*

- *At any time i , the state S_i is consistent with the fractional state A_i .*
- *If the fractional state changes from x_{i-1} to x_i at time i , incurring a movement cost of c_i , then the state S_{i-1} can be modified to a state S_i while incurring a cost of $O(c_i)$ in expectation.*

Embedding general metrics into trees and in particular HSTs has been the subject of many studies. The seminal work of Fakcharoenphol *et al.* [20] has shown that any metric can be randomly embedded to σ -HSTs with distortion $O(\frac{\sigma \log n}{\log \sigma})$.

► **Theorem 14** ([20]). *There exists a probabilistic method to embed an arbitrary metric \mathcal{M} into σ -HSTs with distortion $\frac{\sigma \log n}{\log \sigma}$.*

Therefore, to round a fractional solution on a general metric, we first embed it into 6-HSTs with a distortion of at most $O(\log n)$ and then round the solution while losing only a constant factor. This will give us an integral algorithm that has an expected cost of at most $O(\log n)$ times the optimal.

► **Theorem 15.** *For any given fractional online algorithm \mathcal{A} for the k -server problem on an arbitrary metric, there exists an online integral solution for the same problem having a cost of no worse than $O(\log n)$ times the cost of \mathcal{A} in expectation.*

Acknowledgment. We would like to thank Shi Li for having helpful discussions.

References

- 1 Melika Abolhasani, Soheil Ehsani, Hosein Esfandiari, MohammadTaghi Hajiaghayi, Robert Kleinberg, and Brendan Lucier. Beating $1 - 1/e$ for ordered prophets. *arXiv preprint arXiv:1704.05836*, 2017.
- 2 Dimitris Achlioptas, Marek Chrobak, and John Noga. Competitive analysis of randomized paging algorithms. *Theoretical Computer Science*, 234(1):203–218, 2000.
- 3 Saeed Alaei, Mohammad T Hajiaghayi, Vahid Liaghat, Dan Pei, and Barna Saha. Adcell: Ad allocation in cellular networks. In *Algorithms-ESA 2011*, pages 311–322. Springer, 2011.
- 4 Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. Online prophet-inequality matching with applications to ad allocation. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 18–35. ACM, 2012.
- 5 Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. The online stochastic generalized assignment problem. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 11–25. Springer, 2013.
- 6 Susanne Albers, Lene M Favrholdt, and Oliver Giel. On paging with locality of reference. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 258–267. ACM, 2002.
- 7 Nikhil Bansal, Niv Buchbinder, Aleksander Madry, and Joseph Naor. A polylogarithmic-competitive algorithm for the k-server problem. In Rafail Ostrovsky, editor, *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 267–276. IEEE Computer Society, 2011. doi:10.1109/FOCS.2011.63.
- 8 Nikhil Bansal, Niv Buchbinder, and Joseph Seffi Naor. A primal-dual randomized algorithm for weighted paging. *Journal of the ACM (JACM)*, 59(4):19, 2012.
- 9 Yair Bartal, Marek Chrobak, and Lawrence L Larmore. A randomized algorithm for two servers on the line. *Information and Computation*, 158(1):53–69, 2000.
- 10 Yair Bartal and Elias Koutsoupias. On the competitive ratio of the work function algorithm for the k-server problem. *Theoretical computer science*, 324(2):337–345, 2004.
- 11 Luca Becchetti. Modeling locality: A probabilistic analysis of lru and fwf. In *Algorithms-ESA 2004*, pages 98–109. Springer, 2004.
- 12 Avrim Blum, Carl Burch, and Adam Kalai. Finely-competitive paging. In *Foundations of Computer Science, 1999. 40th Annual Symposium on*, 1999.
- 13 Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. cambridge university press, 2005.
- 14 Allan Borodin, Sandy Irani, Prabhakar Raghavan, and Baruch Schieber. Competitive paging with locality of reference. *Journal of Computer and System Sciences*, 50(2):244–258, 1995.
- 15 Marek Chrobak, H Karloff, T Payne, and S Vishwnathan. New results on server problems. *SIAM Journal on Discrete Mathematics*, 4(2):172–181, 1991.
- 16 Marek Chrobak and Lawrence L Larmore. An optimal on-line algorithm for k servers on trees. *SIAM Journal on Computing*, 20(1):144–148, 1991.
- 17 Sina Dehghani, Soheil Ehsani, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Saeed Seddighin. Online survivable network design and prophets. 2015.
- 18 Sina Dehghani, Ian A Kash, and Peter Key. Online stochastic scheduling and pricing the clouds. 2017.
- 19 Peter J Denning. The working set model for program behavior. *Communications of the ACM*, 26(1):43–48, 1983.
- 20 Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 448–455. ACM, 2003.

- 21 Amos Fiat, Richard M Karp, Michael Luby, Lyle A McGeoch, Daniel D Sleator, and Neal E Young. Competitive paging algorithms. *Journal of Algorithms*, 12(4):685–699, 1991.
- 22 Amos Fiat and Manor Mendel. Truly online paging with locality of reference. In *Foundations of Computer Science, 1997. Proceedings., 38th Annual Symposium on*, pages 326–335. IEEE, 1997.
- 23 Amos Fiat and Manor Mendel. Better algorithms for unfair metrical task systems and applications. *SIAM Journal on Computing*, 32(6):1403–1422, 2003.
- 24 Naveen Garg, Anupam Gupta, Stefano Leonardi, and Piotr Sankowski. Stochastic analyses for online combinatorial optimization problems. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 942–951. Society for Industrial and Applied Mathematics, 2008.
- 25 Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI*, volume 7, pages 58–65, 2007.
- 26 Sandy Irani, Anna R Karlin, and Steven Phillips. Strongly competitive algorithms for paging with locality of reference. *SIAM Journal on Computing*, 25(3):477–497, 1996.
- 27 Anna R Karlin, Steven J Phillips, and Prabhakar Raghavan. Markov paging. *SIAM Journal on Computing*, 30(3):906–922, 2000.
- 28 Howard Karloff, Yuval Rabani, and Yiftach Ravid. Lower bounds for randomized k -server and motion-planning algorithms. *SIAM Journal on Computing*, 23(2):293–312, 1994.
- 29 Elias Koutsoupias and Christos H Papadimitriou. On the k -server conjecture. *Journal of the ACM (JACM)*, 42(5):971–983, 1995.
- 30 Ulrich Krengel, Louis Sucheston, et al. Semiamarts and finite values. *Bull. Amer. Math. Soc*, 83(4), 1977.
- 31 Mark S Manasse, Lyle A McGeoch, and Daniel D Sleator. Competitive algorithms for server problems. *Journal of Algorithms*, 11(2):208–230, 1990.
- 32 Lyle A McGeoch and Daniel D Sleator. A strongly competitive randomized paging algorithm. *Algorithmica*, 6(1-6):816–825, 1991.
- 33 Konstantinos Panagiotou and Alexander Souza. On adequate performance measures for paging. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 487–496. ACM, 2006.
- 34 Daniel D Sleator and Robert E Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, 1985.
- 35 Duru Türkoglu. The k -server problem and fractional analysis, 2005. Master’s Thesis, The University of Chicago. URL: <http://people.cs.uchicago.edu/~duru/papers/masters.pdf>.
- 36 Neal E Young. Bounding the diffuse adversary. In *SODA*, volume 98, pages 420–425, 1998.