Towards the k-Server Conjecture: A Unifying Potential, Pushing the Frontier to the Circle

Christian Coester ✉
CWI, Amsterdam, The Netherlands

Elias Koutsoupias ✉
University of Oxford, UK

Abstract

The k-server conjecture, first posed by Manasse, McGeoch and Sleator in 1988, states that a k-competitive deterministic algorithm for the k-server problem exists. It is conjectured that the work function algorithm (WFA) achieves this guarantee, a multi-purpose algorithm with applications to various online problems. This has been shown for several special cases: $k = 2$, $(k + 1)$-point metrics, $(k + 2)$-point metrics, the line metric, weighted star metrics, and $k = 3$ in the Manhattan plane.

The known proofs of these results are based on potential functions tied to each particular special case, thus requiring six different potential functions for the six cases. We present a *single* potential function proving k-competitiveness of WFA for all these cases. We also use this potential to show k-competitiveness of WFA on multiray spaces and for $k = 3$ on trees. While the DoubleCoverage algorithm was known to be k-competitive for these latter cases, it has been open for WFA. Our potential captures a type of lazy adversary and thus shows that in all settled cases, the worst-case adversary is lazy. Chrobak and Larmore conjectured in 1992 that a potential capturing the lazy adversary would resolve the k-server conjecture.

To our major surprise, this is not the case, as we show (using connections to the k-taxi problem) that our potential fails for three servers on the circle. Thus, our potential highlights laziness of the adversary as a fundamental property that is shared by all settled cases but violated in general. On the one hand, this weakens our confidence in the validity of the k-server conjecture. On the other hand, if the k-server conjecture holds, then we believe it can be proved by a variant of our potential.

2012 ACM Subject Classification Theory of computation → K-server algorithms

Keywords and phrases Online algorithms, competitive analysis, k-server, work function algorithm

Digital Object Identifier 10.4230/LIPIcs.ICALP.2021.57

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: `https://arxiv.org/abs/2102.10474` [14]

Funding *Christian Coester*: Supported by NWO VICI grant 639.023.812.

1 Introduction

The k-server problem, introduced by Manasse, McGoech and Sleator [23], is one of the most fundamental problems in online optimization and contains other problems like paging or weighted paging as important special cases. It is defined as follows: k servers are located in a metric space. One by one, points of the metric space are requested, and each request must be served upon arrival by moving one of the servers to the requested point. The problem is typically considered *online*, where the choice of this server has to be made without knowledge of future requests. The goal is to minimize the total distance traveled by all servers.

When Manasse, McGeoch and Sleator [23] introduced the k-server problem, they showed that on any metric space with $n \geq k+1$ points[1], every deterministic online algorithm has competitive ratio at least k. They showed that this lower bound is tight when $k = 2$ or $k = n-1$ by giving a k-competitive algorithm for these cases and boldly conjectured that a k-competitive online algorithm exists for the general case. This conjecture became known as the famous *k-server conjecture* and has been a driving force in online optimization, making the k-server problem perhaps the most studied problem in the field. It has often been referred to as "the holy grail of competitive analysis", and many techniques developed for the k-server problem have later found applications to other problems.

Chrobak, Karloff, Payne, and Vishwanathan [10] designed the elegant *Double Coverage* algorithm to achieve the optimal competitive ratio of k on the line metric. Shortly after, Chrobak and Larmore [11] extended this algorithm to tree metrics, again matching the lower bound of k. The first algorithm for general metrics with a competitive ratio depending only on k was found by Fiat, Rabani and Ravid [16], achieving a competitive ratio exponential in k. Significant progress was made by Koutsoupias and Papadimitriou [21], showing that a competitive ratio of $2k-1$ is achievable on general metric spaces.

While this reduces the gap between the upper and lower bound to a factor of 2, it remains open to determine the exact competitive ratio. The lack of a proof of the k-server conjecture is even more puzzling given that the algorithm conjectured to achieve the competitive ratio of k has been known for 30 years: The *work function algorithm (WFA)*. It is this algorithm that achieves the aforementioned upper bound of $2k-1$ [21]. Its definition is generic[2], with applications reaching far beyond the k-server problem. For instance, WFA achieves the optimal competitive ratio for metrical task systems [7, 6], the closely related *generalized WFA* has been applied successfully to the weighted k-server problem [3], the generalized 2-server problem [26] and layered graph traversal [9], and work functions have also played a crucial role in recent breakthroughs for convex body chasing [1, 25]. Given these connections, an exact understanding of the WFA for the k-server problem is likely to have a wider impact on online optimization in general.

WFA is known to achieve the tight competitive ratio of k for the following special cases, which impose restrictions on the number of servers and/or the type of metric space:

- $k = 2$ [12]
- $k = n-1$ (folklore; see e.g. [19])
- $k = n-2$ [20, 4]
- line metric [4]
- weighted star metrics [4]
- $k = 3$ in the Manhattan plane [5]

While there has been a lack of progress on the k-server conjecture for about two decades, tremendous progress has been achieved for the *randomized* k-server problem in recent years [2, 8, 22], leading to algorithms with polylogarithmic competitive ratios.

1.1 Our contribution

Our contribution consists of three parts.

(a) The known proofs of the aforementioned six special cases where WFA is k-competitive all use a different potential function, and thus do not seem to point towards a potential function that can solve the k-server conjecture in the general case. We present a single potential function that proves the k-server conjecture for all these cases.

[1] On metric spaces with $n \leq k$ points, the k-server problem is trivial.
[2] WFA always chooses the action that would be best if the future were a mirror image of the past.

(b) Tree metrics are the only special case of the k-server problem where WFA is not known to be k-competitive but another algorithm is (namely, the Double Coverage algorithm [11]). In [4], the question whether WFA is k-competitive on trees was raised as an intermediate step towards solving the k-server conjecture. In this direction, we use our potential function to show that WFA is k-competitive on multiray spaces (a type of tree metrics that generalizes the line and weighted star metrics) and for $k = 3$ on general trees. Our proofs employ the quasi-convexity property of work functions in several new ways.

(c) Chrobak and Larmore [12] formulated three conjectures which say, essentially, that the "adversary is lazy" in the sense that at any time, the worst-case continuation of the request sequence begins with many requests to the k offline server locations (forcing any sensible algorithm to converge to this configuration) before other points are requested. They verified their conjectures on tens of thousands of small metric spaces. In [5], a stronger statement was considered (ignoring the question what kind of work functions are "reachable"), which fails in general but which they conjectured to be true on the circle metric. We reject all these conjectures by showing that for $k = 3$, our potential captures exactly this lazy adversary (and a more restricted adversary for general k), but that it fails on the circle by giving an explicit request sequence as a counterexample. This highlights an important conceptual separation between all cases where k-competitiveness of WFA has been shown and the general case. We believe this property constitutes the main difficulty in resolving the k-server conjecture, and it suggests the circle as the main testing ground for further progress. Our method of constructing the counterexample is based on a connection with the k-taxi problem [13], which we use to generate phenomena of large metric spaces on a much smaller metric space.

1.2 Overview

We provide various definitions and lemmas in Section 2. In Section 3 we formally define our potential in two equivalent ways and show the basic way to use it to prove k-competitiveness. In Section 4, we relate our potential to the lazy adversary potential that was defined implicitly by Chrobak and Larmore. We prove k-competitiveness on multiray spaces in Section 5. This is our most involved proof, and implies the previously known k-competitiveness on the line and weighted stars as special cases. The proof for $k = 3$ on trees and proofs for previously known special cases using our potential can be found in the full version of our paper [14]. In Section 6, we describe ideas of a counter-example to our potential for $k = 3$ on the circle, implying that the adversary is not lazy in this case, unlike the cases where WFA is known to be k-competitive. Details of this construction are given in the full version [14].

2 Preliminaries

Basic notation and abuse of notation. We use (M, d) to denote the metric space, where d is the distance function. We denote by $n = |M|$ its size and by $\Delta = \max_{x,y} d(x, y)$ its diameter. For $x, y \in M$, we will often use the shorthand notation $xy := d(x, y)$. A multiset $C \subseteq M$ of k points is called a *configuration*, representing the location of k servers. We denote by $\mathcal{C}_M^k$ the set of all configurations. For two configurations, X and Y, we denote by $d(X, Y)$ the value of their minimum matching. For notational convenience, we often use the empty space as a union operator on elements of M. For example, we often write $x_1x_2 \dots x_i$ instead of $\{x_1, x_2, \dots, x_i\}$ when it is clear from the context that the set is meant. Similarly, given also a multiset C, we may write $Cx_1 \dots x_i$ instead of $C \cup \{x_1, \dots, x_i\}$. For $x \in M$ and $i \in \mathbb{N}_0$, we write x^i for the multiset containing i copies of x.

For a set $S \subseteq M$, let $\mathrm{clique}(S)$ be the sum of pairwise distances of the points in S.

The k-server problem. An instance of the k-server problem is defined by a metric space (M, d), an initial configuration $C_0 \in \mathcal{C}_M^k$ and a sequence $r_1, r_2, \ldots, r_T \in M$ of requests. A feasible solution is a sequence $C_1, C_2, \ldots C_T$ of configurations such that $r_t \in C_t$ for all $t = 1, \ldots, T$. The cost of this solution is the sum $\sum_{t=1}^{T} d(C_{t-1}, C_t)$.

The work function algorithm (WFA). Given an instance of the k-server problem, the *work function* w_t at time t is the function that maps any configuration C to the minimal cost of serving the first t requests and subsequently ending in configuration C. Formally,

$$w_t(C) := \min_{\substack{C_1,\ldots,C_t \\ \forall \tau\colon r_\tau \in C_\tau}} \sum_{\tau=1}^{t} d(C_{\tau-1}, C_\tau) + d(C_t, C).$$

The work function algorithm (WFA) selects $C_t \ni r_t$ so as to minimize $d(C_{t-1}, C_t) + w_t(C_t)$, with ties broken arbitrarily.

Quasiconvexity. A function $w\colon \mathcal{C}_M^k \to \mathbb{R}$ is called *quasiconvex* if for any configurations X and Y there exists a bijection $\mu\colon X \to Y$ such that for any $A \subseteq X$,

$$w(X) + w(Y) \geq w(A \cup \mu(X \setminus A)) + w(\mu(A) \cup (X \setminus A)).$$

It was shown in [21] that if w is quasiconvex, then μ can be chosen such that $\mu(x) = x$ for all $x \in X \cap Y$. More importantly, it was shown in [21] that any work function is quasiconvex.

Fundamentals about work functions. A function $w\colon \mathcal{C}_M^k \to \mathbb{R}$ is 1-*Lipschitz* if

$$w(X) - w(Y) \leq d(X, Y) \tag{1}$$

for all configurations X and Y. By triangle inequality, every work function is 1-Lipschitz.

Let $\mathcal{Q}_M^k$ be the set of functions $w\colon \mathcal{C}_M^k \to \mathbb{R}$ that are quasiconvex. Let $\mathcal{W}_M^k \subseteq \mathcal{Q}_M^k$ be the subset of functions that are additionally 1-Lipschitz. We may drop k and/or M from the notation when they are clear from the context or immaterial. For $w \in \mathcal{W}$ and configurations X and Y, we say that Y *supports* X if (1) holds with equality. Note that if Y supports X in w_t, then the cheapest way of serving the first t requests and ending in configuration X is equal to the cheapest way of serving the first t requests and then first going to Y and then to X. Thus, if Y supports X, then there is no reason for an offline algorithm to be in configuration X because it is at least as good to be in configuration Y and delay the move from Y to X until later.

The *support* of w, denoted $\mathrm{supp}(w)$, is the set of all configurations that are *not* supported by any other configuration. Intuitively, $\mathrm{supp}(w_t)$ are the possible configurations where an optimal offline algorithm might be at time t. Clearly,

$$w(X) = \min_{Y \in \mathrm{supp}(w)} w(Y) + d(X, Y)$$

for any configuration X. In particular, any work function is fully specified by its support and the values it takes on support configurations.

For $r \in M$, let $\mathcal{W}_M^k(r) \subseteq \mathcal{W}_M^k$ be the subset of 1-Lipschitz, quasiconvex functions with the property that every support configuration contains r. Again, we may drop k and/or M from the notation. Note that the work function w_t at time t is in $\mathcal{W}(r_t)$.

There exists a simple update rule to compute the new work function when a new request is issued. For $w \in \mathcal{W}_M$ and $r \in M$, the updated work function $w \wedge r \in \mathcal{W}(r)$ is defined by

$$w \wedge r(C) = \min_{X \ni r} w(X) + d(X, C).$$

It is easy to see that $w_t = w_{t-1} \wedge r_t$. A basic observation is that if $r_t \in C$, then $w_{t-1}(C) = w_t(C)$. Another basic property is that $w_t(C) \geq w_{t-1}(C)$.

2.1 Extended cost, minimizers and duality

The following lemma was proved by Chrobak and Larmore [12] (see also [19]):

▶ **Lemma 1** (Extended cost lemma). *If for every k-server instance on a metric space M it holds that*

$$\sum_{t=1}^{T} \max_X [w_t(X) - w_{t-1}(X)] \leq (\rho + 1) \cdot \min_X w_T(X) + c_M$$

for some constant c_M independent of the request sequence, then WFA is ρ-competitive on M.

The power of this lemma is that it reduces the task of proving competitiveness of WFA to a property of work functions. In particular, we do not need to keep track of the actual configurations of the online and offline algorithm. The quantity $\max_X [w_t(X) - w_{t-1}(X)]$ is also called the *extended cost* of the tth request, and the proof of the lemma is based on the fact that the total extended cost over all requests is an upper bound on the sum of WFA's cost and the optimal offline cost.

For a work function $w \in \mathcal{W}_M^k$ and a point $y \in M$, we call a configuration $X \in \arg\min w(X) - d(y^k, X)$ a *minimizer of w with respect to y*. There is a direct connection between minimizers and the configurations X maximizing the extended cost. This is captured by the duality lemma, which was first proved in [21]. We give a slightly stronger version of the duality lemma by stating it as an equivalence rather than an implication.

▶ **Lemma 2** (Duality lemma). *Let $w \in \mathcal{W}_M$ and $r \in M$. Define $w' = w \wedge r$. Then $A \in \arg\min_X w(X) - d(r^k, X)$ if and only if the following two conditions hold:*

$$A \in \arg\max_X w'(X) - w(X) \tag{2}$$

$$A \in \arg\min_X w'(X) - d(r^k, X) \tag{3}$$

Proof. The "only if" direction is the duality lemma of [21], where it was shown that if $A \in \arg\min_X w(X) - d(r^k, X)$ then for every configuration B

$$w'(A) + w(B) \geq w(A) + w'(B), \qquad \text{and} \tag{4}$$

$$w'(B) - d(r^k, B) \geq w'(A) - d(r^k, A). \tag{5}$$

By summing these two constraints we get $w(B) - d(r^k, B) \geq w(A) - d(r^k, A)$, which shows the other direction. ◀

It is interesting that the proof of the duality lemma does not use the fact that d is a distance, i.e., it satisfies the triangle inequality.

2.2 Additional properties of work functions

In this section, we provide additional properties of work functions that follow from the quasiconvexity property. We will use these properties to prove k-competitiveness on multiray spaces and for $k = 3$ on trees.

The notion of quasiconvex or quasiconcave functions appears in many different areas and was discovered independently a few times. As a result, they appear with different terminology in literature. For example, in the early 1980s Celso and Crawford [17] defined a related notion as a sufficient condition to the existence of Walrasian Equilibria and called a similar notion gross substitute functions[3]; in 1990, Dress and Wenzel [15] related them to a variant of the greedy algorithm and called them valuated matroids; Koutsoupias and Papadimitriou [21] defined them in the context of online algorithms for the k-server problem and called them quasiconvex. They have also played a central role in discrete optimization [24].

▶ **Lemma 3.** *Let $w \in \mathcal{Q}$. Let $X \in \arg\min w(X)$, and let $x \in X$. Then there exists $Y \in \arg\min_{Y \not\ni x} w(Y)$ such that $X - x \subset Y$.*

Proof. Let Y be chosen such that $X \cap Y$ is maximal under inclusion. Suppose towards a contradiction that there exists $x' \in (X - x) \setminus Y$. By quasiconvexity, there exists $y' \in Y \setminus X$ such that $w(X) + w(Y) \geq w(X - x' + y') + w(Y - y' + x')$. By choice of X, we have $w(X - x' + y') \geq w(X)$. Combining these last two inequalities, we get $w(Y) \geq w(Y - y' + x')$. But $Y - y' + x' \not\ni x$ and $X \cap Y \subsetneq X \cap (Y - y' + x')$, so this contradicts the choice of Y. ◀

▶ **Lemma 4.** *Let $w \in \mathcal{Q}^k_M$. Let $X \in \arg\min w(X)$, and let $A \subset M$ be a (multi)set of cardinality $|A| < k$. Then there exists $Y \in \arg\min_{Y \supset A} w(Y)$ such that $Y - A \subseteq X - A$.*

Proof. Let Y be chosen such that $(Y - A) \setminus (X - A)$ is minimal under inclusion and suppose towards a contradiction that there exists $y \in (Y - A) \setminus (X - A)$. By quasiconvexity, there exists $x \in X \setminus Y$ such that $w(X) + w(Y) \geq w(X - x + y) + w(Y - y + x)$. By choice of X, we have $w(X - x + y) \geq w(X)$. Combining these inequalities, we get $w(Y) \geq w(Y - y + x)$. But this contradicts the choice of Y since we would rather have chosen $Y - y + x$. ◀

▶ **Lemma 5.** *Let $w \in \mathcal{W}^k_M(r)$, let $X \subseteq M$ be a k-point multiset and $x, y \in X$. If X resolves[4] from x in w, then also $X - y + x$ resolves from x in w.*

Proof. Suppose that instead, $X - y + x$ resolves from some $z \in X - y - x$. Consider the $(k-3)$-point multiset $C := X - y - x - z$. Then

$$\begin{aligned} w(X) + w(X - y + x) &= w(Cxyz) + w(Cx^2z) \\ &= w(Cyzr) + w(Cx^2r) + rx + rz \\ &\geq w(Cxyr) + w(Cxzr) + rx + rz \\ &\geq w(Cxyz) + w(Cx^2z), \end{aligned}$$

where the first inequality is by quasiconvexity and the last by 1-Lipschitzness of w. Since the second and the last expression are the same, we have equality in all steps. But then the last step shows that Cx^2z resolves from x. Since $Cx^2z = X - y + x$, the lemma follows. ◀

---

[3] Gross substitute functions are real functions defined for all subsets of a ground set V, whose restriction to subsets of each size k are quasiconvex.

[4] When $w(X) = w(X - x + y) + xy$, we say that X "resolves from x to y". If $y = r$ is the last request, we simply say that X resolves from x.

3 The potential

We provide two different, but equivalent definitions of our potential function. The first formulation views the potential through the lens of the *m-evader problem*, which is equivalent to the k-server problem when $m = n - k$. Thereafter, we will give a more compact and equivalent formulation of the same potential in the k-server view based on extending the metric space by adding antipodal points.

3.1 The evader potential

The m-evader problem is defined similarly to the k-server problem, but instead of k servers there are m evaders in the metric space, which must occupy m different points at all times. When a point r is requested, rather than moving a server towards r, an evader that might be located at r has to move to a different point. The equivalence between the k-server problem and the $(n-k)$-evader problem follows by identifying a server configuration C with the evader configuration $M \setminus C$.[5] Given a k-server work function w, we denote by $\hat{w}$ the corresponding evader work function, defined by $\hat{w}(C) := w(M \setminus C)$.

In the evader view, the potential $\hat{\Phi}$ is defined as follows. Let $y = (y_1, \dots, y_n)$ be a permutation of the points of the metric space M. Let

$$\hat{\Phi}_y(\hat{w}) := \mathrm{clique}(y_1 \dots y_{n-k-1}) + \sum_{i=n-k}^{n} \min_{\substack{C \subseteq \{y_1,\dots,y_i\} \\ |C|=n-k}} \left(\hat{w}(C) + d(C, y_i^{n-k})\right)$$

$$\hat{\Phi}(\hat{w}) := \min_y \hat{\Phi}_y(\hat{w}). \tag{6}$$

▶ **Theorem 6.** *Let (M,d) be an n-point metric space. If for every $r \in M$ and every work function $w \in \mathcal{W}_M^k(r)$ it holds that $\hat{\Phi}(\hat{w}) = \hat{\Phi}_y(\hat{w})$ for a permutation y of M with $y_n = r$, then WFA is k-competitive on M.*

Proof. Consider a k-server instance on M with a request sequence $r_1, \dots, r_T$ and associated sequence of work functions $w_0, \dots, w_T$. We first show that at each time t, the change in potential is an upper bound on the extended cost.

By the premise of the lemma, $\hat{\Phi}(\hat{w}_t) = \hat{\Phi}_y(\hat{w}_t)$ for some y with $y_n = r_t$. Thus,

$$\begin{aligned}
\hat{\Phi}(\hat{w}_t) - \hat{\Phi}(\hat{w}_{t-1}) &\geq \hat{\Phi}_y(\hat{w}_t) - \hat{\Phi}_y(\hat{w}_{t-1}) \\
&\geq \min_{\substack{C \subseteq M \\ |C|=n-k}} \left(\hat{w}_t(C) + d(C, r_t^{n-k})\right) - \min_{\substack{C \subseteq M \\ |C|=n-k}} \left(\hat{w}_{t-1}(C) + d(C, r_t^{n-k})\right) \\
&= \min_{\substack{X \subseteq M \\ |X|=k}} \left(w_t(X) - d(X, r_t^k)\right) - \min_{\substack{X \subseteq M \\ |X|=k}} \left(w_{t-1}(X) - d(X, r_t^k)\right) \\
&= \max_X w_t(X) - w_{t-1}(X),
\end{aligned}$$

where the first inequality uses $\hat{\Phi}(\hat{w}_{t-1}) \leq \hat{\Phi}_y(\hat{w}_{t-1})$, the second inequality uses $y_n = r_t$ and the fact that $\hat{w}_{t-1}(C) \leq \hat{w}_t(C)$ for each C, the first equation translates evader work functions to server work functions and uses $d(C, r_t^{n-k}) = d(M, r_t^n) - d(M \setminus C, r_t^k)$, and the second equation is due to the duality lemma, which says that the same X can be chosen in both minima and the maximum. So indeed, the potential change upper bounds the extended cost.

[5] This identification requires the server configuration to be a set rather than a multiset. This is no restriction on the power of k-server algorithms (online or offline).

Now, we can bound the total extended cost by

$$\sum_{t=1}^{T} \max_X \left[w_t(X) - w_{t-1}(X)\right] \le \hat{\Phi}(\hat{w}_T)$$
$$\le (k+1)\cdot \min_X w_T(X) + c_M,$$

where the last inequality is due to the fact that $\hat{\Phi}(\hat{w}_T)$ is a sum of distances (which are absorbed by the constant c_M) and $k+1$ work function values, each of which differs from $\min_X w_T(X)$ by at most k times the diameter of M due to 1-Lipschitzness of w_T (and the diameters are also absorbed by c_M). The theorem follows from the extended cost lemma. ◀

3.2 The k-server potential

We now derive an equivalent but simpler expression for the aforementioned potential. To formulate it, we need the notion of antipodal points.

Let Δ be the diameter of M. A point $\bar{p} \in M$ is called the *antipode* of another point $p \in M$ if for each $x \in M$, $px + x\bar{p} = p\bar{p} = \Delta$. On some metric spaces such as the circle, each point has an antipode. As mentioned in [18], every metric space can be extended so that each point has an antipode: To achieve this, add to M another copy of the same points, $\bar{M} = \{\bar{p} \colon p \in M\}$, and define distances by $\bar{p}\bar{q} = pq$ and $\bar{p}q = 2\Delta - pq$ for $p, q \in M$. It is easy to check that $M \cup \bar{M}$ is a metric space where $\bar{p}$ and p are antipodes of each other.

Consider a metric space M where every point has an antipode. Let $x_1, \dots, x_k \in M$. We define the k-server potential Φ via

$$\Phi_{x_1,\dots,x_k}(w) := \sum_{i=0}^{k} w(\bar{x}_i^i x_{i+1} \dots x_k)$$
$$\Phi(w) := \min_{x_1,\dots,x_k} \Phi_{x_1,\dots,x_k}(w). \tag{7}$$

The following lemma states that the two potential functions differ by a fixed constant depending on M and are therefore equivalent.

▶ **Lemma 7.** *Let M be a pseudo-metric space of diameter Δ where every point has an antipode and there are k copies of each point.*[6] *For any work function $w \in \mathcal{W}_M^k$ and any permutation $y = (y_1, \dots, y_n)$ of M,*

$$\Phi_{y_{n-k+1}\cdots y_n}(w) = \hat{\Phi}_y(\hat{w}) - \mathrm{clique}(M) + \frac{k(k+1)}{2}\Delta.$$

Proof. Subtracting clique(M) from the evader potential and using server work functions instead of evader work functions, we have

$$\hat{\Phi}_y(\hat{w}) - \mathrm{clique}(M) = \sum_{i=n-k}^{n} \min_{\substack{C \supseteq \{y_{i+1},\dots,y_n\} \\ |C|=k}} \left(w(C) - \sum_{p \in C \cap \{y_1,\dots,y_i\}} py_i \right).$$

[6] It is only a pseudo-metric because the distance between two copies of the same point is 0. We use the assumption of several copies of the same point because the definition of $\Phi_{x_1,\dots,x_k}$ allows points to repeat, whereas Φ_y requires y to be a permutation.

Notice that the minimum in the summand for i is achieved when $C \setminus \{y_{i+1}, \dots, y_n\}$ consists of $k-n+i$ copies of the antipodal point $\bar{y}_i$. Thus, the expression is equal to

$$\sum_{i=n-k}^{n} \left(w(\bar{y}_i{}^{k-n+i} y_{i+1} \dots y_n) - (k-n+i)\Delta\right) = \sum_{i=n-k}^{n} w(\bar{y}_i{}^{k-n+i} y_{i+1} \dots y_n) - \frac{k(k+1)}{2}\Delta$$

$$= \Phi_{y_{n-k+1} \dots y_n}(w) - \frac{k(k+1)}{2}\Delta. \quad ◀$$

▶ **Corollary 8.** *Let (M,d) be a metric space where every point has an antipode. If for every $r \in M$ and every work function $w \in \mathcal{W}_M^k(r)$ it holds that $\Phi(w) = \Phi_{x_1 \dots x_k}(w)$ for some $x_1, \dots x_k \in M$ with $x_k = r$, then WFA is k-competitive on M.*

4 Interpretation as a lazy adversary potential

4.1 The implicitly defined potential by Chrobak and Larmore

Chrobak and Larmore [12] gave an *implicit* definition of a potential that they conjectured to prove the k-server conjecture. This potential captures exactly a type of lazy adversary. To give a precise definition, we first need some additional notation.

For $r \in M$ and a work function $w \in \mathcal{W}$, denote by $\nabla(w,r) := \max_A (w \wedge r)(A) - w(A)$ the extended cost of request r on w. For a request sequence $\rho = (r_1, \dots, r_T) \in M^*$, let

$$\nabla(w,\rho) := \sum_{t=1}^{T} \nabla(w_{t-1}, r_t)$$

be the total extended cost, where $w_t = w \wedge r_1 \wedge r_2 \wedge \dots \wedge r_t$ is the updated work function after the first t requests. The potential conjectured by Chrobak and Larmore is given by

$$\tilde{\Phi}(w) := \min_X \tilde{\Phi}_X(w)$$

where the maximum is taken over configurations X and

$$\tilde{\Phi}_X(w) := -\operatorname{clique}(X) + (k+1)w(X) - \sup_{\rho \in X^*} \nabla(w,\rho).$$

Because of the term $\sup_{\rho \in X^*} \nabla(w,\rho)$, this potential captures exactly the worst-case extended cost when the future request sequence consists only of points in X, until the work function is a cone[7] with support $\{X\}$. An adversary constructing such a request sequence can be thought of as "lazy" because it wants to force the online algorithm to the offline configuration X before it requests different points. The additional term $\operatorname{clique}(X)$ is needed because of extended cost being incurred when passing from one cone to a different cone. The definition of $\tilde{\Phi}$ is only implicit because of the supremum over request sequences $\rho \in X^*$. It was conjectured in [12] that $\tilde{\Phi}(w \wedge r) - \tilde{\Phi}(w) \geq \nabla(w,r)$ for any (reachable) work function w and request r. This would imply the k-server conjecture similarly to the proof of Theorem 6. They also conjectured that $\tilde{\Phi}_X(w \wedge r)$ is minimized for a configuration X containing r, and more specifically that it is minimized by a configuration $X \in \operatorname{supp}(w \wedge r)$. This would imply the previous conjectures. We show that for $k = 3$, the potential $\tilde{\Phi}$ matches our potential Φ. For $k \geq 4$, our potential captures a more restricted type of lazy adversary. As we will

[7] A work function is a *cone* if its support contains only a single configuration.

show in Section 6 that our potential fails to bound the extended cost for $k = 3$ on the circle, this disproves the conjectures from [12] and yields the surprising insight that the worst-case adversary on the circle is not lazy – unlike the adversary for all cases where WFA is known to be k-competitive.

4.2 Relationship to our potential

Our next lemma shows that our potential Φ captures a more restricted adversarial strategy, where the configuration X is ordered as $x_1, \ldots, x_k$ and the next request in ρ is always to the point x_i with i maximal that leads to a change of the work function. We will show later that for $k = 3$, this imposes no additional restriction.

For fixed $x_1, \ldots, x_k \in M$ and a work function $w \in \mathcal{W}_M^k$, define a request sequence $r_1, r_2, \ldots, r_T$ as follows. Let $w_t = w \wedge r_1 \wedge r_2 \wedge \cdots \wedge r_t$ be the updated work function after the first t requests. We define $r_t = x_i$ for i maximal such that $w_{t-1} \wedge x_i \neq w_{t-1}$; if no such i exists, the request sequence ends, $T = t - 1$, and w_T is a cone with support $\{\{x_1, \ldots, x_k\}\}$.

▶ **Lemma 9.**

$$\Phi_{x_1,\ldots,x_k}(w) = \frac{k(k+1)}{2}\Delta - \text{clique}(x_1, \ldots, x_k) + (k+1)w(x_1 \ldots x_k) - \sum_{t=1}^{T} \nabla(w_{t-1}, r_t).$$

Proof. It suffices to show

$$\Phi_{x_1,\ldots,x_k}(w_t) = \Phi_{x_1,\ldots,x_k}(w_{t-1}) + \nabla(w_{t-1}, r_t) \tag{8}$$

$$\Phi_{x_1,\ldots,x_k}(w_T) = (k+1)w(x_1 \ldots x_k) + \frac{k(k+1)}{2}\Delta - \sum_{1 \le i < j \le k} x_i x_j. \tag{9}$$

For equation (9), we have

$$\begin{aligned}
\Phi_{x_1,\ldots,x_k}(w_T) &= \sum_{j=0}^{k} w_T(\bar{x}_j^j x_{j+1} \ldots x_k) \\
&= (k+1)w_T(x_1, \ldots, x_k) + \sum_{1 \le i \le j \le k} x_i \bar{x}_j \\
&= (k+1)w(x_1, \ldots, x_k) + \sum_{1 \le i \le j \le k} (\Delta - x_i x_j) \\
&= (k+1)w(x_1, \ldots, x_k) + \frac{k(k+1)}{2}\Delta - \sum_{1 \le i < j \le k} x_i x_j.
\end{aligned}$$

We now show equation (8). Let i be such that $r_t = x_i$. Then,

$$\begin{aligned}
\Phi_{x_1,\ldots,x_k}(w_t) &= \sum_{j=0}^{k} (w_{t-1} \wedge x_i)(\bar{x}_j^j x_{j+1} \ldots x_k) \\
&= \sum_{j=0}^{i-1} w_{t-1}(\bar{x}_j^j x_{j+1} \ldots x_k) + \sum_{j=i}^{k} (w_{t-1} \wedge x_i)(\bar{x}_j^j x_{j+1} \ldots x_k).
\end{aligned} \tag{10}$$

By maximality of i, $x_{j+1} \ldots x_k$ is contained in every support configuration of w_{t-1}. Thus, $\bar{x}_i^i x_{i+1} \ldots x_k$ is a minimizer of w_{t-1} with respect to x_i and hence

$$(w_{t-1} \wedge x_i)(\bar{x}_i^i x_{i+1} \ldots x_k) = w_{t-1}(\bar{x}_i^i x_{i+1} \ldots x_k) + \nabla(w_{t-1}, r_t) \tag{11}$$

by the duality lemma.

We claim that

$$(w_{t-1} \wedge x_i)(\bar{x}_j^j x_{j+1} \dots x_k) = w_{t-1}(\bar{x}_j^j x_{j+1} \dots x_k) \qquad \forall j = i+1, \dots, k. \tag{12}$$

Assuming this is true, we obtain (8) by substituting (11) and (12) into (10).

It remains to show (12). Since $w_{t-1} \wedge x_i \geq w_{t-1}$, the direction "$\geq$" is obvious. For the other direction, since $x_j x_{j+1} \dots x_k$ is contained in every support configuration of w_{t-1},

$$\begin{aligned} w_{t-1}(\bar{x}_j^j x_{j+1} \dots x_k) &= w_{t-1}(\bar{x}_j^{j-1} x_j x_{j+1} \dots x_k) + \bar{x}_j x_j \\ &\geq w_{t-1}(\bar{x}_j^{j-1} x_i x_{j+1} \dots x_k) - x_i x_j + \bar{x}_j x_j \\ &= (w_{t-1} \wedge x_i)(\bar{x}_j^{j-1} x_i x_{j+1} \dots x_k) + x_i \bar{x}_j \\ &\geq (w_{t-1} \wedge x_i)(\bar{x}_j^j x_{j+1} \dots x_k). \end{aligned}$$ ◀

▶ **Lemma 10.** *Let $X \subset M$ with $|X| = 3$ and $r \in X$ be fixed and let $w \in \mathcal{W}_M^3(r)$. For a bijection $\pi\colon \{1, \dots, 3\} \to X$, write $\Phi_\pi := \Phi_{\pi(1)\pi(2)\pi(3)}$. Then*

$$\min_{\pi\colon \pi(3)=r} \Phi_\pi(w) = \min_\pi \Phi_\pi(w).$$

Proof. Let π be a minimizer of the right hand side. If $\pi(k) = r$, we are done. The case $\pi(k-1) = r$ is also easy, using the fact that r is contained in every support configuration. The remaining case $\pi(k-2) = r$ is non-trivial. Let $y := \pi(k-1)$ and $z := \pi(k)$. We will construct a permutation π' with $\pi'(3) = r$ and $\Phi_\pi(w) \geq \Phi_{\pi'}(w)$. This will only affect the last three terms in the sum of the definition of Φ,

$$w(\bar{r}^{k-2}yz) + w(\bar{y}^{k-1}z) + w(\bar{z}^k).$$

If $w(\bar{y}^{k-1}z) = w(\bar{y}^{k-2}rz) + \bar{y}r$, then

$$\begin{aligned} w(\bar{r}^{k-2}yz) + w(\bar{y}^{k-1}z) + w(\bar{z}^k) &= w(\bar{r}^{k-2}yz) + w(\bar{y}^{k-2}rz) + w(\bar{z}^k) + \bar{y}r \\ &\geq w(\bar{y}^{k-2}rz) + w(\bar{r}^{k-1}z) + w(\bar{z}^k) \end{aligned}$$

where the inequality uses $\bar{y}r = y\bar{r}$. This corresponds to a permutation with r in the next-to-last position, and it is easy to push it from there to the last position.

So we can assume $w(\bar{y}^{k-1}z) = w(\bar{y}^{k-1}r) + zr$. Thus

$$\begin{aligned} w(\bar{r}^{k-2}yz) + w(\bar{y}^{k-1}z) + w(\bar{z}^k) &= w(\bar{r}^{k-2}yz) + w(\bar{y}^{k-1}r) + w(\bar{z}^{k-1}r) + zr + \bar{z}r \\ &= w(\bar{r}^{k-2}yz) + w(\bar{y}^{k-1}r) + w(\bar{z}^{k-1}r) + \Delta. \end{aligned} \tag{13}$$

In the last expression, y and z are symmetric, so we can assume

$$w(\bar{r}^{k-2}yz) = w(\bar{r}^{k-2}rz) + yr. \tag{14}$$

By quasi-convexity and Lipschitzness of the work function (and $\bar{y}\bar{r} = yr$, $\bar{r}r = \Delta$),

$$\begin{aligned} w(\bar{y}^{k-1}r) + w(\bar{r}^{k-2}rz) &\geq w(\bar{y}^{k-2}zr) + w(\bar{r}^{k-2}\bar{y}r) \\ &\geq w(\bar{y}^{k-2}zr) + w(\bar{r}^k) - yr - \Delta \end{aligned} \tag{15}$$

Combining (13), (14) and (15), we get

$$w(\bar{r}^{k-2}yz) + w(\bar{y}^{k-1}z) + w(\bar{z}^k) \geq w(\bar{y}^{k-2}zr) + w(\bar{z}^{k-1}r) + w(\bar{r}^k),$$

corresponding to the permutation $(\pi(1), \pi(2), \pi(3)) = (y, z, r)$. ◀

We remark (without proof) that the above lemma fails for $k = 4$.

By the following corollary, for $k = 3$ it holds that Φ is an explicit expression for the implicit potential of [12].

▶ **Corollary 11.** *For $k = 3$,*

$$\Phi(w) = 6\Delta + \min_{\substack{X:\,|X|=3\\ r_1,\dots,r_T\in X}} \left[4w(X) - \text{clique}(X) - \sum_{t=1}^{T} \nabla(w \wedge r_1 \dots r_{t-1}, r_t)\right]. \tag{16}$$

Proof. The direction "$\geq$" follows from Lemma 9. For direction "$\leq$", select X and $r_1, \dots, r_T$ to minimize the right hand side. Let $w_t = w \wedge r_1 \dots r_t$. By minimality of the right hand side, w_T is a cone at X. Let $\Phi_X = \min \Phi_{x_1x_2x_3}$, with the minimum taken over permutations x_1, x_2, x_3 of X. By Lemma 10, we have $\Phi_X(w_t) = \Phi_{xyr_t}(w_t)$ for some $x, y \in X$. Thus,

$$\begin{aligned}\Phi_X(w_t) - \Phi_X(w_{t-1}) &\geq \Phi_{xyr_t}(w_t) - \Phi_{xyr_t}(w_{t-1})\\ &= w_t(\bar{r}_t^3) - w_{t-1}(\bar{r}_t^3)\\ &= \nabla(w_{t-1}, r_t).\end{aligned}$$

Hence,

$$\begin{aligned}\Phi(w) \leq \Phi_X(w) = \Phi_X(w_T) &- \sum_{t=1}^{T} [\Phi_X(w_t) - \Phi_X(w_{t-1})]\\ \leq \Phi_X(w_T) &- \sum_{t=1}^{T} \nabla(w_{t-1}, r_t),\end{aligned}$$

which is equal to the right hand side of (16) by Lemma 9 and since w_T is a cone at X. ◀

5 Multi-ray spaces

A multi-ray space is a tree of depth 1 whose edges have infinite length and where requests can appear at arbitrary locations along the edges. We call these edges rays.

We will show in this section that WFA is k-competitive on multiray spaces. Note that a multiray space with only 2 rays is equal to the line metric. A subset of a multi-ray space containing only one point from each ray is a weighted star. Our proof therefore recovers the known proofs that WFA is k-competitive on the line and on weighted stars as special cases.

We denote by c the center/root of the multi-ray space, i.e., the origin of the rays. We can assume that every ray has finite length by considering only a sufficiently long part that all requests fall into. We call the endpoint of a ray that is not the center a *leaf*. Denote by $\mathcal{L}$ the set of leaves. For $w \in \mathcal{W}^k$, define $m_w(X) := w(X) - d(c^k, X)$. Note that m_w is also quasiconvex. As we use the server definition of the potential, we augment the multi-ray space by adding antipodes as discussed earlier. In the definition (6), we require the points $x_1, \dots, x_k$ to be chosen from the original metric space M. This corresponds to requiring the permutation in the evader potential to end with k points from the original metric space, which does not affect the proof of Theorem 6.

The proof that WFA is k-competitive on multi-ray spaces proceeds along the following three main steps:

1. First we establish properties of $\Phi_{x_1\dots x_k}$ when $x_i = \ell_i$ are leaves. In particular, we express $\Phi_{\ell_1\dots\ell_k}$ in terms of m_w, and show that $\ell_1, \dots, \ell_k$ can be permuted under certain conditions.
2. We then show by induction on k that $\Phi_{x_1\dots x_k}(w)$ is indeed minimized when $x_1, \dots, x_k$ are leaves and $\min_X m_w(X) = m_w(x_1 \dots x_k)$.
3. Finally, we show that $\Phi_{x_1\dots x_k}(w)$ is also minimized for some $x_1, \dots, x_k$ where only $x_1, \dots, x_{k-1}$ are leaves whereas $x_k = r$ is the last request.

Step 1: Properties of $\Phi_{x_1 \dots x_k}$ when x_i are leaves

▶ **Lemma 12.** *Let $w \in \mathcal{W}^k$. There exist leaves $\ell_1, \dots, \ell_k$ such that $\min_X m_w(X) = m_w(\ell_1 \dots \ell_k)$.*

Proof. Follows from the fact that since w is 1-Lipschitz, $m_w(X)$ cannot increase when a point in X moves away from c towards a leaf. ◀

▶ **Lemma 13.** *For any $w \in \mathcal{W}^k$, a leaf $\ell \in \mathcal{L}$ and $x_{i+1}, \dots, x_k \in M$,*

$$w(\bar{\ell}^i x_{i+1} \dots x_k) = \min_{\substack{X \supseteq x_{i+1} \dots x_k: \\ X - x_{i+1} \dots x_k \subseteq \mathcal{L} - \ell}} m_w(X) + i(\Delta - c\ell) + \sum_{j=i+1}^{k} cx_j.$$

Proof.

$$\begin{aligned} w(\bar{\ell}^i x_{i+1} \dots x_k) &= \min_{X \supseteq x_{i+1} \dots x_k} w(X) + d(X - x_{i+1} \dots x_k, \bar{\ell}^i) \\ &= \min_{X \supseteq x_{i+1} \dots x_k} w(X) + i\Delta - d(X - x_{i+1} \dots x_k, \ell^i) \end{aligned}$$

We claim that the minimum is achieved by some X with $X - x_{i+1} \dots x_k \subseteq \mathcal{L} - \ell$: Indeed, if there is some $x \in X - x_{i+1} \dots x_k$ that is not in $\mathcal{L} - \ell$, sliding x away from ℓ along a path to some other leaf increases $d(X - x_{i+1} \dots x_k, \ell^i)$ by the distance moved, and it increases $w(X)$ by at most this distance, so the whole term cannot increase.

This also means that every distance in $d(X - x_{i+1} \dots x_k, \ell^i)$ goes across the center, i.e.,

$$d(X - x_{i+1} \dots x_k, \ell^i) = i \cdot c\ell + d(X - x_{i+1} \dots x_k, c^i).$$

The lemma now follows by definition of m_w. ◀

As a consequence of Lemma 13, we obtain the following expression for $\Phi_{\ell_1 \dots \ell_k}$ whenever $\ell_1, \dots, \ell_k$ are leaves:

$$\Phi_{\ell_1 \dots \ell_k}(w) = \frac{k(k+1)}{2}\Delta + \sum_{i=0}^{k} \min_{\substack{X_i \supseteq \ell_{i+1} \dots \ell_k: \\ X_i - \ell_{i+1} \dots \ell_k \subseteq \mathcal{L} - \ell_i}} m_w(X_i). \tag{17}$$

The following symmetry and monotonicity properties allow us to reorder $\ell_1, \dots, \ell_k$ under certain circumstances.

▶ **Lemma 14** (Symmetry and Monotonicity Lemma)**.** *Let $w \in \mathcal{W}^k$ and let $\ell_1, \dots, \ell_k$ be leaves such that $\min_X m_w(X) = m_w(\ell_1 \dots \ell_k)$. The following properties hold:*
Symmetry: *$\Phi_{\ell_1 \dots \ell_k}(w)$ is constant under permutation of $\ell_1, \dots, \ell_k$.*
Monotonicity: *For any leaf ℓ, $\Phi_{\ell_1 \dots \ell_{k-1} \ell}(w) \geq \Phi_{\ell \ell_1 \dots \ell_{k-1}}(w) \geq \Phi_{\ell_1 \dots \ell_k}(w)$.*

Proof. For the symmetry property and the first inequality of the monotonicity property, we proceed by induction on k. The base case $k = 1$ is trivial. For the induction step, it suffices to show that $\Phi_{\ell_1 \dots \ell_{k-1} \ell}(w) \geq \Phi_{\ell_1 \dots \ell_{k-2} \ell \ell_{k-1}}(w)$, with equality if $\ell = \ell_k$. The lemma then follows by invoking the induction hypothesis on $\tilde{w} = w(\,\cdot\, \ell_{k-1}) \in \mathcal{W}^{k-1}$, observing that $\min_X m_{\tilde{w}}(X) = m_{\tilde{w}}(\ell_1 \dots \ell_{k-2} \ell_k)$, and that $\Phi_{x_1 \dots x_{k-1} \ell_{k-1}}(w) - \Phi_{x_1 \dots x_{k-1}}(\tilde{w}) = w(\bar{\ell}_{k-1}^k)$ is a constant function of $x_1, \dots, x_{k-1}$ (and thus $\Phi_{x_1 \dots x_{k-1} \ell_{k-1}}(w)$ and $\Phi_{x_1 \dots x_{k-1}}(\tilde{w})$ are affected in the same way when $x_1, \dots, x_{k-1}$ are permuted).

Note that only the two terms involving X_{k-1} and X_k in (17) are affected when the last two leaves are swapped. For two leaves y and z, let

$$f(y,z) := \min_{\substack{Y\ni z:\\ Y-z\subseteq\mathcal{L}-y}} m_w(Y) + \min_{Z\subseteq\mathcal{L}-z} m_w(Z)$$

We only need to show that $f(\ell_{k-1},\ell)\geq f(\ell,\ell_{k-1})$, and that this holds with equality if $\ell=\ell_k$. Assume $\ell_{k-1}\neq\ell$ as otherwise there is nothing to show. Then

$$\begin{aligned}
&f(\ell_{k-1},\ell)-f(\ell,\ell_{k-1})\\
&\quad=\min_{\substack{Y_1\ni\ell:\\ Y_1\subseteq\mathcal{L}-\ell_{k-1}}} m_w(Y_1)+\min_{Z_1\subseteq\mathcal{L}-\ell} m_w(Z_1)-\min_{\substack{Y_2\ni\ell_{k-1}:\\ Y_2\subseteq\mathcal{L}-\ell}} m_w(Y_2)-\min_{Z_2\subseteq\mathcal{L}-\ell_{k-1}} m_w(Z_2)\\
&\quad\geq\min_{Z_1\subseteq\mathcal{L}-\ell} m_w(Z_1)-\min_{\substack{Y_2\ni\ell_{k-1}:\\ Y_2\subseteq\mathcal{L}-\ell}} m_w(Y_2)\\
&\quad=0,
\end{aligned}$$

where the last equation follows by applying Lemma 3 to the quasi-convex function m_w. If $\ell=\ell_k$, then the same argument shows that the inequality can be replaced by equality.

It remains to show the second inequality of the monotonicity property. Due to the symmetry property, and by a renaming of leaves, it suffices to show that $\Phi_{\ell\ell_2\ldots\ell_k}(w)\geq\Phi_{\ell_1\ell_2\ldots\ell_k}(w)$. Assume $\ell\neq\ell_1$, otherwise we are done. In (17), the only terms affected when the first leaf is replaced are the ones involving X_0 and X_1. Let X_0 and X_1 be these sets in $\Phi_{\ell_1\ell_2\ldots\ell_k}(w)$ and X_0' and X_1' those in $\Phi_{\ell\ell_2\ldots\ell_k}(w)$. Then $X_0=\ell_1\ldots\ell_k$, $X_0'=\ell\ell_2\ldots\ell_k$, and since $\min_X m_w(X)=m_w(\ell_1\ldots\ell_k)$, we can choose $X_1'=\ell_1\ldots\ell_k$. Moreover, X_0' satisfies the requirements of X_1 (apart from minimality, possibly), hence $m_w(X_1)\leq m_w(X_0')$. Thus,

$$\Phi_{\ell\ell_2\ldots\ell_k}(w)-\Phi_{\ell_1\ell_2\ldots\ell_k}(w)=m_w(X_0')+m_w(X_1')-m_w(X_0)-m_w(X_1)\geq 0. \qquad ◀$$

Step 2: $x_1,\ldots,x_k$ are indeed leaves

▶ **Lemma 15.** *Let $w\in\mathcal{W}^k$. Let $\ell_1,\ldots,\ell_k$ be leaves such that $\min_X m_w(X)=m_w(\ell_1\ldots\ell_k)$. Then $\Phi(w)=\Phi_{\ell_1\ldots\ell_k}(w)$.*

Proof. By induction on k. The base case $k=0$ is trivial. For the induction step, fix x such that $\Phi(w)=\Phi_{x_1\ldots x_{k-1}x}(w)$ for some $x_1,\ldots,x_{k-1}$. Consider the function $\tilde{w}=w(\cdot x)\in\mathcal{W}^{k-1}$. By Lemma 4, $\min m_{\tilde{w}}(X)=m_{\tilde{w}}(\ell_1\ldots\ell_k-\ell')$ for some $\ell'\in\ell_1\ldots\ell_k$. By Lemma 14, we can assume without loss of generality that $\ell'=\ell_k$, i.e., $\min m_{\tilde{w}}(X)=m_{\tilde{w}}(\ell_1\ldots\ell_{k-1})$. By the induction hypothesis, $\Phi(\tilde{w})=\Phi_{\ell_1\ldots\ell_{k-1}}(\tilde{w})$. Hence,

$$\begin{aligned}
\Phi(w)&=\min_{x_1\ldots x_{k-1}}\Phi_{x_1\ldots x_{k-1}x}(w)\\
&=\min_{x_1\ldots x_{k-1}}\Phi_{x_1\ldots x_{k-1}}(\tilde{w})+w(\bar{x}^k)\\
&=\Phi_{\ell_1\ldots\ell_{k-1}}(\tilde{w})+w(\bar{x}^k)\\
&=\Phi_{\ell_1\ldots\ell_{k-1}x}(w).
\end{aligned}$$

We will now transform the last expression in several steps with the goal of eventually replacing x by ℓ_k.

Denote by ℓ the leaf below x. The goal of the following transformations is to replace x by ℓ. For some $a\in\{0,1,\ldots,k\}$ we have

$$w(\bar{x}^k)=w(\bar{\ell}^a\ell^{k-a})+a\cdot x\ell+(k-a)\cdot x\bar{\ell}.$$

The symmetry property of Lemma 14 allows us to assume that $\ell_1,\dots,\ell_{s-1}$ are all different from ℓ and $\ell_s=\ell_{s+1}=\dots=\ell_{k-1}=\ell$ for some $s\in\{1,\dots,k\}$.

As an intermediate step, we show by (backwards) induction on $j=k,k-1,\dots,\max\{s,a\}$ that

$$\Phi(w)\ge\sum_{i=0}^{j-1}w(\bar\ell_i^i\ell_{i+1}\dots\ell_{k-1}x)+\sum_{i=j}^{k-1}w(\bar\ell_i^{i+1}\ell_{i+1}\dots\ell_{k-1})$$
$$+w(\bar\ell^a\ell^{j-a}\ell_j\dots\ell_{k-1})+a\cdot x\ell+(j-a)\cdot x\bar\ell \tag{18}$$

The base case $k=j$ follows from the previous equation. Suppose now that (18) holds for some $j>\max\{s,a\}$. From $\ell_{j-1}=\ell$ we get

$$w(\bar\ell_{j-1}^{j-1}\ell_j\dots\ell_{k-1}x)\ge w(\bar\ell_{j-1}^{j}\ell_j\dots\ell_{k-1})-x\bar\ell$$

and

$$w(\bar\ell^a\ell^{j-a}\ell_j\dots\ell_{k-1})=w(\bar\ell^a\ell^{j-1-a}\ell_{j-1}\dots\ell_{k-1}).$$

The induction step of (18) follows by plugging these in to (18).

If $a\ge s$, then (18) for $j=a$ yields

$$\begin{aligned}\Phi(w)&\ge\sum_{i=0}^{a-1}w(\bar\ell_i^i\ell_{i+1}\dots\ell_{k-1}x)+\sum_{i=a}^{k-1}w(\bar\ell_i^{i+1}\ell_{i+1}\dots\ell_{k-1})+w(\bar\ell^a\ell_a\dots\ell_{k-1})+a\cdot x\ell\\&\ge\sum_{i=0}^{a-1}w(\bar\ell_i^i\ell_{i+1}\dots\ell_{k-1}\ell)+\sum_{i=a}^{k-1}w(\bar\ell_i^{i+1}\ell_{i+1}\dots\ell_{k-1})+w(\bar\ell^a\ell_a\dots\ell_{k-1})\\&=\Phi_{\ell_1\dots\ell_{k-1}\ell}(w),\end{aligned}$$

where we have used that $\ell_i=\ell$ for $i\ge a\ge s$. The lemma then follows from the monotonicity property of Lemma 14.

Otherwise, $a<s$, and from (18) for $j=s$ we get

$$\begin{aligned}\Phi(w)&\ge\sum_{i=0}^{s-1}w(\bar\ell_i^i\ell_{i+1}\dots\ell_{k-1}x)+\sum_{i=s}^{k-1}w(\bar\ell_i^{i+1}\ell_{i+1}\dots\ell_{k-1})\\&\qquad+w(\bar\ell^a\ell^{s-a}\ell_s\dots\ell_{k-1})+a\cdot x\ell+(s-a)\cdot x\bar\ell\\&\ge\sum_{i=0}^{s-a-1}w(\bar\ell_i^i\ell_{i+1}\dots\ell_{k-1}x)+\sum_{i=s-a}^{s-1}w(\bar\ell_i^i\ell_{i+1}\dots\ell_{k-1}\ell)+\sum_{i=s+1}^{k}w(\bar\ell^i\ell^{k-i})\\&\qquad+w(\bar\ell^a\ell^{s-a}\ell_s\dots\ell_{k-1})+(s-a)\cdot x\bar\ell.\end{aligned}$$

By Claim 16 below, replacing a by $a+1$ does not increase the latter quantity. Inductively we may therefore replace a by s to obtain

$$\begin{aligned}\Phi(w)&\ge\sum_{i=0}^{s-1}w(\bar\ell_i^i\ell_{i+1}\dots\ell_{k-1}\ell)+\sum_{i=s+1}^{k}w(\bar\ell^i\ell^{k-i})+w(\bar\ell^s\ell^{k-s})\\&=\Phi_{\ell_1\dots\ell_{k-1}\ell}(w).\end{aligned}$$

The monotonicity property of Lemma 14 completes the proof. ◀

▷ Claim 16. Let $0\le a<s\le k$ and $w\in\mathcal{W}^k$. Let $\ell_{s-a-1},\dots,\ell_{k-1}$ and ℓ be leaves such that $\ell_i\ne\ell$ for $i<s$, and let x be a point on the ray of ℓ. Then

$$\begin{aligned}&w(\bar\ell_{s-a-1}^{s-a-1}\ell_{s-a}\dots\ell_{k-1}x)+w(\bar\ell^a\ell^{s-a}\ell_s\dots\ell_{k-1})+x\bar\ell\\&\qquad\ge w(\bar\ell_{s-a-1}^{s-a-1}\ell_{s-a}\dots\ell_{k-1}\ell)+w(\bar\ell^{a+1}\ell^{s-a-1}\ell_s\dots\ell_{k-1}).\end{aligned}$$

Proof. Consider the bijection from the definition of quasiconvexity between the two configurations on the left hand side.[8] By the pigeonhole principle, at least one of the $s-a$ copies of ℓ in the second configuration maps to some point $p \in \ell_{s-a} \ldots \ell_{s-1}x$ in the first configuration. Quasiconvexity gives

$$\begin{aligned} &w(\bar{\ell}_{s-a-1}^{s-a-1}\ell_{s-a} \ldots \ell_{k-1}x) + w(\bar{\ell}^a \ell^{s-a}\ell_s \ldots \ell_{k-1}) \\ &\qquad \geq w(\bar{\ell}_{s-a-1}^{s-a-1}\ell_{s-a} \ldots \ell_{k-1}\ell x - p) + w(\bar{\ell}^a \ell^{s-a-1}\ell_s \ldots \ell_{k-1}p). \end{aligned}$$

By 1-Lipschitzness of w, we get

$$w(\bar{\ell}_{s-a-1}^{s-a-1}\ell_{s-a} \ldots \ell_{k-1}\ell x - p) \geq w(\bar{\ell}_{s-a-1}^{s-a-1}\ell_{s-a} \ldots \ell_{k-1}\ell) - px$$

and

$$w(\bar{\ell}^a \ell^{s-a-1}\ell_s \ldots \ell_{k-1}p) \geq w(\bar{\ell}^{a+1}\ell^{s-a-1}\ell_s \ldots \ell_{k-1}) - p\bar{\ell}.$$

Since $\ell_i \neq \ell$ for $i < s$ and $p \in \ell_{s-a} \ldots \ell_{s-1}x$, the point x lies on the path from p to ℓ, i.e., $px + x\ell = p\ell$. Equivalently, $x\bar{\ell} = px + p\bar{\ell}$. The claim follows by combining these inequalities. ◁

Step 3: Alternatively, $x_k = r$

▶ **Lemma 17.** *For any $w \in \mathcal{W}^k(r)$, there exist leaves $\ell_1, \ldots, \ell_{k-1}$ such that $\Phi(w) = \Phi_{\ell_1 \ldots \ell_{k-1} r}(w)$.*

Proof. Since $w \in \mathcal{W}^k(r)$, we can choose $X \in \arg\min_X m_w(X)$ of the form $X = r\ell_2 \ldots \ell_k$ for $\ell_2, \ldots, \ell_k \in \mathcal{L}$. If $\ell := \ell_1$ is the leaf of the ray containing r, then clearly $\ell_1 \ldots \ell_k$ is also a minimizer of m_w and $\ell_1 \ldots \ell_k$ resolves from ℓ_1. Let $\ell_2, \ldots, \ell_k$ be ordered such that $\ell = \ell_1 = \cdots = \ell_s$ for some $s \geq 1$ and $\ell_i \neq \ell$ for $i > s$.

The main part of this proof is to show that there exists $a \in \{1, \ldots, s\}$ such that

$$w(\bar{\ell}_i^i \ell_{i+1} \ldots \ell_k) = \begin{cases} w(\bar{\ell}_i^i \ell_{i+2} \ldots \ell_k r) + r\ell & \text{if } i < a \\ w(\bar{\ell}_i^{i-1} \ell_{i+1} \ldots \ell_k r) + r\bar{\ell}_i & \text{if } i \geq a. \end{cases} \tag{19}$$

Before we prove this, let us see why it implies the lemma. By Lemma 15 and the fact that $\ell = \ell_1 = \cdots = \ell_a$, we have

$$\begin{aligned} \Phi(w) &= \Phi_{\ell_1 \ldots \ell_k}(w) \\ &= \sum_{i=0}^{k} w(\bar{\ell}_i^i \ell_{i+1} \ldots \ell_k) \\ &= \sum_{i=0}^{a-1} w(\bar{\ell}_i^i \ell_{i+2} \ldots \ell_k r) + \sum_{i=a}^{k} w(\bar{\ell}_i^{i-1} \ell_{i+1} \ldots \ell_k r) + a \cdot r\ell + \sum_{i=a}^{k} r\bar{\ell}_i \\ &= \sum_{i=0}^{k-1} w(\bar{\ell}_{i+1}^i \ell_{i+2} \ldots \ell_k r) + w(\bar{\ell}^{a-1} \ell_{a+1} \ldots \ell_k r) + (a-1) \cdot r\ell + \sum_{i=a+1}^{k} \bar{r}\ell_i + \Delta \\ &\geq \sum_{i=0}^{k-1} w(\bar{\ell}_{i+1}^i \ell_{i+2} \ldots \ell_k r) + w(\bar{r}^k) \\ &= \Phi_{\ell_2 \ldots \ell_k r}(w). \end{aligned}$$

[8] We remark that earlier proofs about competitiveness of the work function algorithm only used a weaker form of quasi-convexity and did not actually use the existence of such a bijection.

It remains to show (19). We choose a maximal such that $\bar{\ell}_{a-1}^{a-1}\ell_a \dots \ell_k$ resolves from ℓ. Recall from the start of this proof that $\ell_1 \dots \ell_k$ resolves from ℓ_1, so $a \geq 1$. Moreover, $a \leq s$ since $\ell_i \neq \ell$ for $i > s$. So $a \in \{1, \dots, s\}$ as required. The case "$i < a$" of (19) now follows by backwards induction on i, where the induction step is due to Lemma 5.

Consider now some $i > s$. Letting $\tilde{w} = w(\,\cdot\,\ell_{i+1} \dots \ell_k) \in \mathcal{W}^i_{\mathcal{L}+r}$, we have

$$w(\bar{\ell}_i^i \ell_{i+1} \dots \ell_k) = \tilde{w}(\bar{\ell}_i^i) = \min_{X \subseteq \mathcal{L}+r-\ell_i} m_{\tilde{w}}(X) + i(\Delta - c\ell_i),$$

where the last equation is proved similarly to Lemma 13, but we may allow X to contain the non-leaf r since it is on a different ray than ℓ_i (thanks to $i > s$). Since $\min_X m_w(X) = m_w(r\ell_2 \dots \ell_k)$, we have $\min_X m_{\tilde{w}}(X) = m_{\tilde{w}}(r\ell_2 \dots \ell_i)$, so by Lemma 3 the minimum under the restriction $X \subseteq \mathcal{L} + r - \ell_i$ is achieved for some X with $r \in X$. Thus,

$$\begin{aligned}
w(\bar{\ell}_i^i \ell_{i+1} \dots \ell_k) &= \min_{Y \subseteq \mathcal{L}-\ell_i} \tilde{w}(Yr) - d(Y, c^{i-1}) - r\ell_i + i(\Delta - c\ell_i) \\
&= \min_{Y \subseteq \mathcal{L}-\ell_i} \tilde{w}(Yr) + d(Y, \bar{\ell}_i^{i-1}) + \bar{\ell}_i r \\
&\geq \tilde{w}(\bar{\ell}_i^{i-1} r) + \bar{\ell}_i r \\
&= w(\bar{\ell}_i^{i-1} r \ell_{i+1} \dots \ell_k) + r\bar{\ell}_i.
\end{aligned}$$

Note that the inequality between the first and last expression cannot be strict due to 1-Lipschitzness of w, and their equality reveals that $\bar{\ell}_i^i \ell_{i+1} \dots \ell_k$ resolves from $\bar{\ell}_i$, as desired.

Finally, consider $i \in \{a, a+1, \dots, s\}$. Then $\ell_i = \ell$, and we need to show that $\bar{\ell}^i \ell_{i+1} \dots \ell_k$ resolves from $\bar{\ell}$. Suppose that it instead resolves from ℓ_h for some $h > i$. Since $a \leq s$ was chosen maximal, we know that $\ell_h \neq \ell$. By Lemma 13,

$$\begin{aligned}
w(\bar{\ell}^i \ell_{i+1} \dots \ell_k) &= w(\bar{\ell}^i \ell_{i+1} \dots \ell_{h-1} \ell_{h+1} \dots \ell_k r) + r\ell_h \\
&= \min_{\substack{X \supseteq \ell_{i+1} \dots \ell_{h-1} \ell_{h+1} \dots \ell_k r \\ X - \ell_{i+1} \dots \ell_{h-1} \ell_{h+1} \dots \ell_k r \subseteq \mathcal{L} - \ell}} m_w(X) + i(\Delta - c\ell) + \sum_{\substack{j=i+1 \\ j \neq h}}^{k} c\ell_j + cr + r\ell_h. \\
&= \min_{Y \not\ni \ell} m_{\tilde{w}}(Y) + i(\Delta - c\ell) + r\ell_h,
\end{aligned}$$

where $\tilde{w} = w(\,\cdot\,\ell_{i+1} \dots \ell_{h-1} \ell_{h+1} \dots \ell_k r) \in \mathcal{W}^i_{\mathcal{L}}$. Since $\min_X m_w(X) = m_w(r\ell_2 \dots \ell_k)$, we have that $\min_X m_{\tilde{w}}(X) = m_{\tilde{w}}(\ell_2 \dots \ell_i \ell_h)$, so by Lemma 3 the minimum under the restriction $Y \not\ni \ell$ is achieved for some Y with $\ell_h \in Y$. Letting $Y' = Y - \ell_h$, we get

$$\begin{aligned}
w(\bar{\ell}^i \ell_{i+1} \dots \ell_k) &= m_{\tilde{w}}(Y) + i(\Delta - c\ell) + r\ell_h \\
&= w(Y' \ell_{i+1} \dots \ell_k r) - d(Y, c^i) + i(\Delta - c\ell) + (rc + c\ell_h) \\
&= w(Y' \ell_{i+1} \dots \ell_k r) - \sum_{y \in Y'} y\ell + i\Delta - c\ell + rc \\
&\geq w(Y' \ell_{i+1} \dots \ell_k r) + \sum_{y \in Y'} y\bar{\ell} + \Delta - r\ell \\
&\geq w(\bar{\ell}^{i-1} \ell_{i+1} \dots \ell_k r) + r\bar{\ell},
\end{aligned}$$

where the second equation uses that $\ell \neq \ell_h$ and therefore c lies on the path from r to ℓ_h, and the third equation uses that y and ℓ are different leaves and therefore $yc + c\ell = y\ell$. Again, the inequality between the first and last expression cannot be strict due to 1-Lipschitzness of w, and their equality reveals that $\bar{\ell}^i \ell_{i+1} \dots \ell_k$ resolves from $\bar{\ell}$, completing the proof. ◀

▶ **Theorem 18.** *WFA is k-competitive on multiray spaces.*

Proof. Follows from Lemma 17 and Corollary 8. ◀

6 Non-laziness of the worst-case adversary on the circle

▶ **Theorem 19.** *For $k = 3$ servers on the circle, there exists a reachable work function from where the worst-case adversarial continuation of the request sequence is not lazy. More precisely, there exists a request sequence such that the induced work functions w_t and w_{t+1} after time steps t and $t+1$ and the WFA configuration C_t after time step t satisfy $w_{t+1}(C_t) - w_t(C_t) > \Phi(w_{t+1}) - \Phi(w_t)$.*

In other words, the extended cost is strictly greater than the change in potential. Due to the interpretation of our potential (Section 4), this means that the worst-case continuation of the request sequence after time t is not lazy. If Theorem 19 could be strengthened such that the request sequence to reach w_t has extended cost equal to its induced change in potential, then this would disprove the premise of the extended cost lemma (because one could create a cyclic request sequence where extended cost is always at least the change in potential and exceeds it infinitely often; we remark that one can go from a cone work function to any other cone via a request sequence whose extended cost equals its potential change).

Note that Theorem 19 holds even if in the extended cost $\max_X w_{t+1}(X) - w_t(X)$ we replace X by the configuration C_t of WFA at time t. The significance of this is that the sum of the terms $w_{t+1}(C_t) - w_t(C_t)$ over all time steps is *equal* to the sum of WFA's cost and the optimal offline cost (up to a bounded additive error). Thus, proving violation of the premise of the extended cost lemma with X replaced by C_t would imply that WFA's competitive ratio is strictly greater than k.

The proof of Theorem 19 is given in the full version of our paper [14]. It is based on a tight connection between the k-server problem and the "easy" version of the k-taxi problem observed in [13]. The k-taxi problem is the generalization of the k-server problem where each request is not a single point, but a pair (s, t) of two points, representing the start s and destination t of a taxi request. To serve it, the algorithm has to select a server that first goes to s and then to t. In the "easy" version relevant for us, the cost is defined as the total distance traveled by servers.[9] As shown in [13], the easy k-taxi problem has exactly the same competitive ratio as the k-server problem. The idea of this reduction is that a k-taxi request (s, t) can be simulated by a sequence of many k-server requests along the shortest path from s to t. We extend this idea to show that we can use k-taxi requests to reach work functions that are arbitrarily close to work functions reachable via k-server requests. We then give an explicit counter-example using k-taxi requests.

We remark that up to symmetry and shift by an additive constant, for $k = 3$ there exist over 280,000 different work functions reachable by taxi requests with starts/destinations at the points on the circle considered in our construction (8 equally spaced points for the destinations and 16 equally spaced points for the starts – the aforementioned 8 points as well as the 8 intermediate points). Among these over 280,000 work functions, the pair of w_t and w_{t+1} from our construction is the *only* counterexample to laziness of the adversary. Using only k-server requests and no k-taxi requests, we were unable to find any counterexamples for n equally spaced points on the circle for the values of n that were computationally feasible for us to try. Of course, though, our approximability argument of k-taxi requests via k-server requests implies that such counterexamples do exist for n sufficiently large. Given the rarity of these counterexamples, it is not surprising that Chrobak and Larmore [12] who reported testing their conjecture on tens of thousands of small metric spaces in the early 90s did not find any counterexample.

[9] In contrast, the "hard" k-taxi problem defines the cost as only the overhead distance traveled while *not* carrying a passenger, i.e., the distance from s to t is excluded from the cost.

7 Conclusion

Our potential gives a unified perspective on all cases where WFA is known to be k-competitive. Unlike previous potentials, which were specific to their special case and had no clear intuition, our potential has a natural interpretation as capturing a lazy adversary. We remark that our potential also proves k-competitiveness on 6-point metric spaces. Since work functions, the WFA, and the generalized WFA are central to various online problems, similar potential functions may also prove useful to analyze different problems.

Since it was a major belief that a lazy adveresary would capture the worst case, our insights yield a qualitative explanation of the shortcomings of previous approaches and may point in a direction to overcome these shortcomings.

Our proof for $k = 3$ on trees relies on the fact that if (M, d) is a tree metric, then d is quasiconcave (i.e., $-d$ is quasiconvex). We are puzzled by the question whether this has any deeper connection to the quasiconvexity property of work functions and whether it is crucial for the existence of k-competitive algorithms. While the k-server problem is also k-competitive on some non-quasiconcave metrics (such as the cases $k = 2$ and $n = k - 2$), the reason for this might simply be due to the fact that the subspaces relevant in all proof steps are small (note that any 3-point metric is quasiconcave).

References

1. C. J. Argue, Anupam Gupta, Guru Guruganesh, and Ziye Tang. Chasing convex bodies with linear competitive ratio. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA '20*, pages 1519–1524, 2020.
2. Nikhil Bansal, Niv Buchbinder, Aleksander Madry, and Joseph Naor. A polylogarithmic-competitive algorithm for the k-server problem. *J. ACM*, 62(5):40:1–40:49, 2015.
3. Nikhil Bansal, Marek Eliás, and Grigorios Koumoutsos. Weighted k-server bounds via combinatorial dichotomies. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS '17*, 2017.
4. Yair Bartal and Elias Koutsoupias. On the competitive ratio of the work function algorithm for the k-server problem. *Theoretical Computer Science*, 324(2-3):337–345, 2004.
5. Wolfgang W. Bein, Marek Chrobak, and Lawrence L. Larmore. The 3-server problem in the plane. *Theor. Comput. Sci.*, 289(1):335–354, 2002.
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. Sébastien Bubeck, Michael B. Cohen, Yin Tat Lee, James R. Lee, and Aleksander Madry. k-server via multiscale entropic regularization. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018*, pages 3–16, 2018.
9. William R. Burley. Traversing layered graphs using the work function algorithm. *J. Algorithms*, 20(3):479–511, 1996.
10. Marek Chrobak, Howard Karloff, Tom Payne, and Sundar Vishwanathan. New results on server problems. *SIAM Journal on Discrete Mathematics*, 4(2):172–181, 1991.
11. 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.
12. Marek Chrobak and Lawrence L Larmore. The server problem and on-line games. In *On-line Algorithms, volume 7 of DIMACS Series in Discrete Mathematics and Theoretical Computer Science*. Citeseer, 1992.
13. Christian Coester and Elias Koutsoupias. The online k-taxi problem. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC '19*, 2019.

14 Christian Coester and Elias Koutsoupias. Towards the k-server conjecture: A unifying potential, pushing the frontier to the circle. *CoRR*, abs/2102.10474, 2021. `arXiv:2102.10474`.

15 Andreas W.M. Dress and Walter Wenzel. Valuated matroids: a new look at the greedy algorithm. *Applied Mathematics Letters*, 3(2):33–35, 1990.

16 Amos Fiat, Yuval Rabani, and Yiftach Ravid. Competitive k-server algorithms. *J. Comput. Syst. Sci.*, 48(3), 1994.

17 Alexander S. Kelso Jr and Vincent P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society*, pages 1483–1504, 1982.

18 Elias Koutsoupias. Weak adversaries for the k-server problem. In *40th Annual Symposium on Foundations of Computer Science, FOCS '99*, pages 444–449, 1999.

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

20 Elias Koutsoupias and Christos Papadimitriou. The 2-evader problem. *Information Processing Letters*, 57(5):249–252, 1996.

21 Elias Koutsoupias and Christos H. Papadimitriou. On the k-server conjecture. *J. ACM*, 42(5), 1995.

22 James R. Lee. Fusible HSTs and the randomized k-server conjecture. In *Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '18, pages 438–449, 2018.

23 Mark S. Manasse, Lyle A. McGeoch, and Daniel Dominic Sleator. Competitive algorithms for on-line problems. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, STOC '88*, 1988.

24 Kazuo Murota. *Discrete Convex Analysis*. Society for Industrial and Applied Mathematics, January 2003.

25 Mark Sellke. Chasing convex bodies optimally. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA '20*, pages 1509–1518, 2020.

26 René Sitters. The generalized work function algorithm is competitive for the generalized 2-server problem. *SIAM J. Comput.*, 43(1), 2014.