Doctopic: Algorithms, automata, complexity and games ARTICLE IN PRESS TCS:11045

Theoretical Computer Science ••• (••••) •••–•••

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Infinite linear programming and online searching with turn cost

Spyros Angelopoulos [a,*,1], Diogo Arsénio [b], Christoph Dürr [a,1]

[a] Sorbonne Universités, UPMC Univ Paris 06, CNRS, LIP6, Paris, France
[b] CNRS and Institut de Mathématiques, Université Paris Diderot, France

ARTICLE INFO

Article history:
Received 13 May 2014
Received in revised form 29 February 2016
Accepted 15 January 2017
Available online xxxx
Communicated by S. Sen

Keywords:
Search and exploration problems
Infinite linear programming
Competitive analysis of online algorithms

ABSTRACT

We consider the problem of searching for a hidden target in an environment that consists of a set of concurrent rays. Every time the searcher turns direction, it incurs a fixed cost. The objective is to derive a search strategy for locating the target as efficiently as possible, and the performance of the strategy is evaluated by means of the well-established competitive ratio. In this paper we revisit an approach due to Demaine et al. [8] based on infinite linear-programming formulations of this problem. We first demonstrate that their definition of duality in infinite LPs can lead to erroneous results. We then provide a non-trivial correction which establishes the optimality of a certain round-robin search strategy.

1. Introduction

Searching for a hidden object is a task that is often encountered in everyday situations. It is thus not surprising that computational aspects of search problems have attracted significant attention. The broad setting can be described by three components: a *domain* (i.e., environment) which may be known or unknown; a hidden, immobile target that lies in some unknown position of the environment; and a mobile searcher (e.g. a robot), initially placed at some predefined position of the domain. The objective is to develop a search strategy for locating the target as efficiently as possible. As in [8] we are interested in the case of unbounded domains.

One of the earliest examples of search problems is the *linear search problem*, proposed in [6] and independently in [4]. Here, the target is hidden at some unknown position of the infinite line, and at distance h from a given point designated as the origin, whereas the searcher is initially located at the origin. The objective is to design a search strategy (namely, an algorithm that describes the movement of the searcher on the infinite line) that minimizes the *competitive ratio* of the strategy: the latter is defined as the worst-case ratio of the overall travel cost of the searcher divided by the distance h. A natural generalization of the linear-search problem is the *star search* or *ray search* problem. In this setting, we are given a set of m infinite rays with a common origin O, and a searcher which is initially placed at the origin. The target is located at distance h from O, however the searcher has no knowledge of the ray on which the target lies. A search strategy is an algorithm that specifies how the searcher traverses the rays, and the competitive ratio is defined as the worst-case ratio of the first time a searcher locates the target, over the optimal distance h.

* Corresponding author.
E-mail addresses: spyros.angelopoulos@lip6.fr (S. Angelopoulos), arsenio@math.univ-paris-diderot.fr (D. Arsénio), christoph.durr@lip6.fr (C. Dürr).

[1] Research supported by project ANR-11-BS02-0015 "New Techniques in Online Computation" (NeTOC).

http://dx.doi.org/10.1016/j.tcs.2017.01.013
0304-3975/© 2017 Elsevier B.V. All rights reserved.

Please cite this article in press as: S. Angelopoulos et al., Infinite linear programming and online searching with turn cost, Theoret. Comput. Sci. (2017), http://dx.doi.org/10.1016/j.tcs.2017.01.013

In this paper we study the setting in which the searcher incurs a fixed *turn cost* upon changing direction, and the overall travel cost for a searcher is the sum of the individual search and turn costs. This formulation models the often-encountered setting in which changing a searcher's direction is a time-consuming operation which cannot be ignored; for instance, a robot cannot turn instantaneously. Following [8], we assume that there are costs d_1 and d_2 for turning at a ray and at the origin, respectively; hence the turn cost incurred by a searcher on a single ray exploration is $d = d_1 + d_2$.

1.1. Related work

It has long been known that *geometric* strategies are optimal for linear search [5], a result that was extended initially in [9] as well as in [3] and [11] to the m-ray setting. In this class of strategies, the searcher performs a round-robin exploration of rays with distances forming a geometric sequence (i.e., of the form $a^0, a^1, a^2, \ldots$ for some $a > 1$). In particular, Gal [9] showed an optimal geometric strategy of competitive ratio

$$1 + 2M, \text{ with } M = \frac{a^m}{a-1} \text{ and } a = \frac{m}{m-1}. \tag{1}$$

Other related work includes the study of randomization [20,14], multi-searcher strategies [18], the variant in which an upper bound is known on the distance of the target from the origin [17,7], the variant in which some probabilistic information on target placement is known [11,12], the related problem of designing *hybrid algorithms* [13], and more recently, the study of new performance measures [16,19]. For an overview of results on ray-searching we refer the reader to Chapter 9 of the textbook by Gal and Alpern [1].

The above results assume no turn cost. For given turn cost d, [8] studies ray searching using an approach based on infinite linear-program (LP) formulations. More specifically, in order to lower-bound the cost of any search strategy, they define an infinite series of linear programs, with each linear program describing a (progressively better) set of adversarial target placements. At the limit, the optimal value of the infinite LP gives the strongest lower bound. The approach of [8] consists of solving experimentally this series of finite LPs, then guessing a solution to the infinite LP, and finally providing a proof of optimality based on appropriate duality properties of infinite LP formulations. More precisely, they claim that for every search strategy there is a placement of the target at a certain distance h from the origin such that the strategy incurs a (tight) cost of $(1 + 2M)h + (M - m)d$. Furthermore, they show that this bound is tight, by providing a matching round-robin (and near-cyclical) strategy.

1.2. Contribution of this paper

We begin by revisiting the technique of infinite-LP formulations by Demaine et al. in the context of m-ray searching with turn cost, and by pointing out some subtle pitfalls. Specifically, we give a dual solution that is feasible for the infinite LP of [8] and whose objective value is larger than the upper bound on the search cost that is shown in [8]. This contradiction clearly demonstrates that we cannot rely on dual solutions to the infinite LP; instead one must insist on solutions that are feasible for any finite formulation, and evaluate the objective value that is attained at the limit.

In more technical terms, in order to establish the feasibility of a crucial dual constraint, one needs to study the infinite sequence that is generated by a specific linear recurrence relation. In particular, we seek appropriate initial data for the recurrence relation (at a trade-off relation with the objective value of the LP) such that the generated sequence observes certain limit properties. We prove that the choice of initial data must be such that the sequence in question eventually becomes negative, in stark contrast to [8] which stipulates that the sequence must be strictly positive. This gives rise to a problem related to linear recurrences, which we address using tools from linear algebra and complex analysis.

The remainder of the paper is structured as follows. Section 2 demonstrates the caveats of infinite LPs and shows that duality in infinite LP formulations of the problem is not upheld. Section 3 shows how to remedy this problem; more precisely, we show how to obtain a feasible dual solution for every finite LP formulation, which suffices for obtaining the desired result. Section 4 addresses the technical details behind the construction of the dual solution, and provides a self-contained study of the underlying recurrence relation.

2. LP formulations and the caveats of infinite LPs

2.1. Preliminaries and definitions

We first provide some preliminary facts and definitions concerning m-ray searching. We say that a search algorithm is (α, β)-*competitive* if for any placement of the target at distance h, the search cost is at most $\alpha h + \beta$. In particular, under the assumptions that the target is never placed within distances smaller than a specified fixed constant and for zero turn cost, previous work has established optimal $(1 + 2M, 0)$-competitive algorithms. The question we address is then the following: What is the smallest B such that an algorithm is $(1 + 2M, B)$ competitive, with no assumptions on the target placement? Note that this question becomes non-trivial only in the presence of turn costs, since otherwise B is zero, as argued in [8]. Note also that M is upper-bounded by a linear function in m, since $M = \frac{m^m}{(m-1)^{m-1}}$, therefore $m < M \leq \mathrm{e} \cdot m$, where e is the Euler constant.