Theoretical Computer Science 569 (2015) 24–42

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Searching on a line: A complete characterization of the optimal solution ☆

Prosenjit Bose [a,1], Jean-Lou De Carufel [a,*,2], Stephane Durocher [b,1]

[a] *School of Computer Science, Carleton University, Ottawa, Canada*
[b] *Department of Computer Science, University of Manitoba, Winnipeg, Canada*

ARTICLE INFO

Article history:
Received 23 January 2014
Received in revised form 28 October 2014
Accepted 9 December 2014
Available online 16 December 2014
Communicated by G. Ausiello

Keywords:
Computational geometry
Online algorithms
Optimization
Fibonacci polynomials

ABSTRACT

We revisit the problem of searching for a target at an unknown location on a line when given upper and lower bounds on the distance D that separates the initial position of the searcher from the target. Prior to this work, only asymptotic bounds were known for the optimal competitive ratio achievable by any search strategy in the worst case. We present the first tight bounds on the exact optimal competitive ratio achievable, parameterized in terms of the given bounds on D, along with an optimal search strategy that achieves this competitive ratio. We prove that this optimal strategy is unique. We characterize the conditions under which an optimal strategy can be computed exactly and, when it cannot, we explain how numerical methods can be used efficiently. In addition, we answer several related open questions, including the *maximal reach problem*, and we discuss how to generalize these results to m rays, for any $m \geq 2$.

1. Introduction

Search problems are broadly studied within computer science. A fundamental search problem, which is the focus of this paper, is to specify how a searcher should move to find an immobile target at an unknown location on a line such that the total relative distance travelled by the searcher is minimized in the worst case [4,12,15]. The searcher is required to move continuously on the line, i.e., discontinuous jumps, such as random access in an array, are not possible. Thus, a search corresponds to a sequence of alternating left and right displacements by the searcher. This class of geometric search problems was introduced by Bellman [5] who first formulated the problem of searching for the boundary of a region from an unknown random point within its interior. Since then, many variants of the line search problem have been studied, including multiple rays sharing a common endpoint (as opposed to a line, which corresponds to two rays), multiple targets, multiple searchers, moving targets, and randomized search strategies (e.g., [1–4,6,8–11,14–16]).

For any given search strategy f and any given target location, we consider the ratio A/D, where A denotes the total length of the search path travelled by a searcher before reaching the target by applying strategy f, and D corresponds to the minimum travel distance necessary to reach the target. That is, the searcher and target initially lie a distance D from each other on a line, but the searcher knows neither the value D nor whether the target lies to its left or right. The *competitive*

☆ A preliminary version appeared in [7].

* Corresponding author.
E-mail address: jdecaruf@cg.scs.carleton.ca (J.-L. De Carufel).

[1] This research has been funded by the Natural Sciences and Engineering Research Council (NSERC).
[2] This research has been funded by the Fonds québécois de la recherche sur la nature et les technologies (FQRNT).

http://dx.doi.org/10.1016/j.tcs.2014.12.007
0304-3975/

ratio of a search strategy f, denoted $CR(f)$, is measured by the supremum of the ratios achieved over all possible target locations. Observe that $CR(f)$ is unbounded if D can be assigned any arbitrary real value; specifically, the searcher must know a lower bound $\lambda \le D$. Thus, it is natural to consider scenarios where the searcher has additional information about the distance to the target. In particular, in many instances the searcher can estimate good lower and upper bounds on D. Given a lower bound $D \ge \lambda$, Baeza-Yates et al. [4] show that any optimal strategy achieves a competitive ratio of 9. They describe such a strategy, which we call the *Power of Two* strategy. Furthermore, they observe that when D is known to the searcher, it suffices to travel a distance of $3D$ in the worst case, achieving a competitive ratio of 3.

We represent a search strategy by a function $f : \mathbb{N} \to \mathbb{R}^+$. Given such a function, a searcher travels a distance of $f(0)$ in one direction from the origin (say, to the right), returns to the origin, travels a distance of $f(1)$ in the opposite direction (to the left), returns to the origin, and so on, until reaching the target. We refer to $f(i)$ as the distance the searcher travels from the origin during the i-th iteration. The corresponding function for the Power of Two strategy of Baeza-Yates et al. is $f(i) = 2^i\lambda$. Showing that every optimal strategy achieves a competitive ratio of exactly 9 relies on the fact that no upper bound on D is specified [4]. Therefore, it is natural to ask whether a search strategy can achieve a better competitive ratio when provided both lower and upper bounds $\lambda \le D \le \Lambda$. In what follows, we use the scale invariant $\rho = \Lambda/\lambda$ to indicate how good the upper bound is. When $\rho = 1$, then $\Lambda = \lambda$. This corresponds to the case where the searcher knows D. Moreover, if we let $\rho \to \infty$, this represents the case where the searcher is not given any upper bound.

Given R, the *maximal reach problem*, examined by Hipke et al. [12], is to identify the largest bound Λ such that there exists a search strategy that finds any target within distance $D \le \Lambda$ with competitive ratio at most R. López-Ortiz and Schuierer [15] study the maximal reach problem on m rays, from which they deduce that the competitive ratio $CR(f_{opt})$ of any optimal strategy f_{opt} is at least

$$1 + 2\frac{m^m}{(m-1)^{m-1}} - O\left(\frac{1}{\log^2 \rho}\right).$$

When $m = 2$, the corresponding lower bound becomes

$$9 - O\left(\frac{1}{\log^2 \rho}\right).$$

They also provide a general strategy that achieves this asymptotic behaviour on m concurrent rays, given by

$$f(i) = \sqrt{1 + \frac{i}{m}}\left(\frac{m}{m-1}\right)^i \lambda.$$

Again, for $m = 2$ this is

$$f(i) = \sqrt{1 + \frac{i}{2}}\, 2^i \lambda.$$

Surprisingly, this general strategy is independent of ρ. In essence, it ignores any upper bound on D, regardless of how tight it is. Thus, we examine whether there exists a better search strategy that depends on ρ, thereby using both the upper and lower bounds on D. Furthermore, previous lower bounds on $CR(f_{opt})$ have an asymptotic dependence on ρ applying only to large values of ρ, corresponding to having only coarse bounds on D. Can we express tight bounds on $CR(f_{opt})$ in terms of ρ?

Let $f_{opt}(i) = a_i\lambda$ denote an optimal strategy for given values λ and Λ. Since f_{opt} is optimal and $D \ge \lambda$, we must have $f_{opt}(i) \ge \lambda$ for all $i \ge 0$. Therefore, $a_i \ge 1$ for all $i \ge 0$. Moreover, for any possible position of the target, the strategy f_{opt} must eventually reach it. Hence, there must be two integers i and i' of different parities such that $f_{opt}(i) \ge \Lambda$ and $f_{opt}(i') \ge \Lambda$, which implies that $a_i \ge \rho$ and $a_{i'} \ge \rho$. Since f_{opt} is optimal, we have $f_{opt}(i) = f_{opt}(i') = \Lambda$. Moreover, let n be the smallest integer such that $f_{opt}(n) = \Lambda$ (equivalently, $a_n = \rho$). Since f_{opt} is optimal, we have $f_{opt}(n+1) = \Lambda$ (equivalently, $a_{n+1} = \rho$). Consequently, $n + 2$ is the number of iterations necessary to reach the target with strategy f_{opt} in the worst case (recall that the sequence starts at $i = 0$). The question is now to determine the sequence $\{a_i\}_{i=0}^n$ that defines f_{opt}. López-Ortiz and Schuierer [15] provide an algorithm to compute the maximal reach for a given competitive ratio together with a strategy corresponding to this maximal reach. They state that the value n and the sequence $\{a_i\}_{i=0}^n$ can be computed using binary search, which increases the running time proportionally to $\log\rho$. Can we find a faster algorithm for computing f_{opt}? Since in general, a_0 is the root of a polynomial equation of unbounded degree (see Theorem 1), a binary search is equivalent to the bisection method for solving polynomial equations. However, the bisection method is a slowly converging numerical method. Can the computational efficiency be improved? Moreover, given ε, can we bound the number of steps necessary for a root-finding algorithm to identify a solution within tolerance ε of the exact value?

1.1. Overview of results

We address all of the questions raised above in Section 2. We characterize f_{opt} by computing the sequence $\{a_i\}_{i=0}^n$ for the optimal strategy. We do this by computing the number of iterations $n + 2$ needed to find the target in the worst case.