

# An Algorithmic Approach to Search Games: Finding Solutions Using Best Response Oracles

Lisa Hellerstein\*

Thomas Lidbetter†

February 22, 2018

## Abstract

We present efficient algorithms for computing optimal or approximately optimal strategies in a zero-sum game for which Player I has  $n$  pure strategies and Player II has an arbitrary number of pure strategies. We assume that for any given mixed strategy of Player I, a best response or “approximate” best response of Player II can be found by an oracle in time polynomial in  $n$ . We then show how our algorithms may be applied to several search games with applications to security and counter-terrorism.

## 1 Introduction

Consider a zero-sum game with positive integer payoffs in which Player I (the maximizer) has  $n$  pure strategies and Player II (the minimizer) has an arbitrary number of pure strategies. The payoff function is denoted by  $C$ , so that if Players I and II play mixed strategies given by probability vectors  $\mathbf{x}$  and  $\mathbf{y}$ , the payoff is given by  $C(\mathbf{x}, \mathbf{y})$ . Whenever we refer to games in this paper we assume they are of this form. Optimal mixed strategies and the value of the game can be found solving a linear program, but if Player II’s strategy set is very large compared to  $n$  then the computational time required to do this may not be polynomial in  $n$ .

Examples of this nature frequently arise in the study of finite search games for an immobile Hider. In such games, which are played between a Hider and a Searcher, the Hider typically has a strategy set corresponding to  $n$  locations in which to hide and the Searcher’s strategy set corresponds to a permutation of those locations (or a subset thereof), corresponding to the order in which she searches the locations. The order of search may be restricted due to a network structure of the search space or other restrictions on the mode of search. Notwithstanding this, the size of the Searcher’s strategy set is usually exponential in  $n$ . The payoff of the game is generally either some cost incurred by the Searcher in searching for the Hider, which she wishes to minimize, or the probability of detection, which she wishes to maximize.

Search games are motivated by wide-ranging problems in national security, counter-terrorism, search-and-rescue, biology and others. Summaries of the literature on search games for an immobile Hider can be found in Alpern and Gal (2003), Gal (2011) and Hohzaki (2016).

---

\*Department of Computer Science and Engineering, NYU Tandon School of Engineering, 6 Metrotech Center, Brooklyn, NY 11201, USA

†Department of Management Science and Information Systems, Rutgers Business School, Newark, NJ, USA

In many such games, it is possible to exploit the structure of the game in order to solve the game (that is, find optimal strategies and the value) or find bounds on the value. An example of this is the classic search game for an immobile Hider on a network studied in Gal (1979) and Gal (2000). More recent examples can be found in Alpern (2016), Alpern and Lidbetter (2015, 2013a,b), Angelopoulos *et al.* (2016), Baston and Kikuta (2013, 2015), Fokkink *et al.* (2016a) and Lin and Singham (2016). Many of these game are infinite, in the sense that one or both of the players has a strategy set of infinite cardinality, but in this paper we restrict ourselves to finite games, which lend themselves better to an algorithmic approach.

Exploiting the structure of the game may not be possible in some games, or it may limit what can be achieved. However, in some games it may be possible to efficiently solve the one-sided problem of finding an optimal best response of Player II for any given mixed strategy of Player I: that is, finding a pure strategy for Player II which has minimal expected payoff against the given mixed strategy of Player I. We will refer to this problem as the *best response problem*. In this paper we study games for which there exists an oracle that can efficiently solve the best response problem. More generally, we consider games such that there is an oracle that, for any mixed strategy of Player I, can find a pure strategy for Player II that ensures the payoff is no more than  $\alpha$  times the best response payoff. We call such an oracle an  $\alpha$ -*approximate best response oracle*. So if  $\alpha = 1$ , then an  $\alpha$ -approximate best response oracle finds best responses; in this case we simply call it a *best response oracle*.

We show that for games with an  $\alpha$ -approximate best response oracle, there are efficient algorithms for finding  $\alpha$ -optimal strategies for both players. By this we mean a mixed strategy for Player I that guarantees a payoff of at least  $V^*/\alpha$  and a mixed strategy for Player II that guarantees a payoff of at most  $\alpha V^*$ , where  $V^*$  is the value of the game.

Our algorithms follow easily from previous work, and our main contribution is the wide range of applications to the field of search games, detailed in Section 5. The examples we give include known search games with known solutions, known search games with unknown solutions and new search games, all of which we are able to solve efficiently or find approximate solutions for with our algorithms. We hope that by giving these techniques more prominence, they might be used by others for the purpose of solving search games, or indeed games in other fields.

## 1.1 Example: Searching in Boxes

Before describing our algorithms in detail, we illustrate the type of game we will consider with an example of a search game that is easy to state, and that has the characteristics of the type of games our results can be applied to. Although its solution is already known, it is a fundamental game and provides a good introduction to the more complex games we will discuss in Section 5.

An item (for example an explosive device) is hidden in one of  $n$  boxes with costs  $c_1, \dots, c_n$ , so that the Hider's strategy set is the set of boxes,  $[n] = \{1, \dots, n\}$ . A pure strategy for the Searcher is simply an ordering of the boxes, so the Searcher's strategy set is the set of all permutations  $\pi : [n] \rightarrow [n]$ . In other words,  $\pi(j)$  is the  $j$ th location to be inspected. Clearly the Hider's strategy set has size  $n$  and the Searcher's strategy set has size  $n!$ . For given strategies  $j$  of the Hider and  $\pi$  of the Searcher, the payoff  $C(j, \pi)$  of the game, which we call the *search cost* is the sum of the costs of all the boxes in the ordering up to and including the box that contains the hidden object. That is,

$$C(j, \pi) = \sum_{i \leq \pi^{-1}(j)} c_{\pi(i)}.$$

We call this game BOX. The solution of BOX is already known in closed form. The Hider’s optimal strategy is to choose box  $j$  with probability proportional to  $c_j$  as shown in Alpern and Lidbetter (2013b), where it was also proved that this is the unique optimal strategy. The Searcher does not have a unique optimal strategy, but three different optimal strategies can be found in Alpern and Lidbetter (2013b), Lidbetter (2013) and (implicitly) in Condon *et al.* (2009). The value of the game is also given in closed form (in two different ways) in Alpern and Lidbetter (2013b) and Lidbetter (2013). It is somewhat surprising, given the simplicity with which the game can be described, that prior to these publications the game had not been solved, to the best of our knowledge.

We now show that one could alternatively find optimal strategies in this game by applying our algorithms. To do this, we must show that there is a best response oracle. This is the problem of finding a permutation  $\pi$  that minimizes  $C(\mathbf{x}, \pi)$  for a given Hider mixed strategy  $\mathbf{x}$ . It is a classic search problem solved by Blackwell, reported in Matula (1964). In fact, Blackwell solved a more general problem, which we describe later in Subsection 5.7. For this version of the problem, the solution is to search the boxes in non-increasing order of the index  $x_j/c_j$ .

We can arrive at the solution in another way, taken from the scheduling literature. Consider the problem of scheduling  $n$  jobs with processing times  $t_j$  and weights  $w_j$ , that correspond to the relative importance of the jobs. The object is to determine the schedule that orders the jobs in such a way as to minimize the weighted sum of the completion times. This problem is usually denoted  $1||\sum w_j C_j$  in the scheduling literature, and has the well known solution, given by Smith (1956), of scheduling the jobs in non-decreasing order of the ratio  $t_j/w_j$ . Of course, viewing the costs  $c_j$  in the best response problem for BOX as processing times, and viewing the probabilities  $x_j$  as weights, these two problems and their solutions are equivalent.

We could also consider the more general box searching game introduced and solved by Lidbetter (2013) of searching for  $k \geq 1$  objects hidden in  $n$  boxes. However, in this case our approach cannot be used to solve the game, since the solution of the best response problem is not known. In fact, the computational complexity of this problem is unknown, even for  $k = 2$ , and we view this as an interesting open problem. More precisely, suppose two items are hidden in  $n$  boxes with search costs according to a known distribution, so that the probability they are in boxes  $i$  and  $j$  is some  $p_{ij}$ , for  $i \leq j$ . What ordering of the boxes minimizes the total expected cost of finding both objects?

## 1.2 Overview of Paper

We take two approaches to using an  $\alpha$ -approximate best response oracle to find approximate solutions to games. In Section 3, we present an algorithm using an existing *ellipsoid approach* from the approximation algorithms literature. This approach uses the ellipsoid algorithm with an approximate separation oracle to achieve approximately optimal solutions to an LP and to its dual. The approach was previously used in Jansen (2003), Carr and Vempala (2002), Friggstad and Swamy (2012) and Feldman *et al.* (2012).

In Section 4 we use a *multiplicative weights update* approach: a good general survey of this method can be found in Arora *et al.* (2012). Approximately 20 years ago, Freund and Schapire (1999, 1996) gave a multiplicative update algorithm that computes near-optimal strategies for two-player zero-sum games. Their algorithm yields strategies with payoffs that are within an *additive*  $\varepsilon$  of the game value  $V$ . Specifically, the strategy for Player I guarantees an expected payoff of at least  $V - \varepsilon$ , and the strategy for Player II guarantees an expected payoff of at most  $V + \varepsilon$ . The algorithm assumes that the payoffs are in the interval  $[0, 1]$ , so  $\varepsilon$  represents a fraction of the range  $[0, 1]$  of payoffs.

Surprisingly, the approach of Freund and Schapire does not seem to have been extended to yield multiplicative approximations: that is, strategies that guarantee an expected payoff within a multiplicative factor  $(1 + \varepsilon)$  of the value. Multiplicative approximations are more common in the literature on approximation algorithms. In Section 4, we address this gap in the literature.

The work of Section 4 is an application of known results and techniques. We use a multiplicative update rule of Garg and Könemann (2007), and we rely somewhat on their analysis, though their work does not automatically imply the existence of approximately optimal strategies for Player I. We therefore combine this with a variation of the analysis of Freund and Schapire (1999) as presented by Arora *et al.* (2012). We make some changes, due to the form of our multiplicative approximation factor, and the fact that our algorithm is designed to accommodate an approximate best response oracle. However, our primary contribution is what we believe is the first self-contained exposition of an algorithm and its analysis for computing multiplicative approximations of optimal strategies in a zero-sum game.

We then apply our algorithms to a variety of games in Section 5.

## 2 Preliminaries

In what follows, we consider a zero-sum game between a maximizing Player I with  $n$  pure strategies and a minimizing Player II with an number of strategies that that could be as large as  $2^{p(n)}$ , for some polynomial  $p$ . The payoff function is  $C$  and all the payoffs are positive integers. We assume there is an  $\alpha$ -approximate best response oracle, for some constant  $\alpha \geq 1$ . That is, there is a polynomial time algorithm (polynomial in  $n$ ) which, for any given mixed strategy  $\mathbf{x}$  of Player I, computes a pure strategy  $j_\alpha(\mathbf{x})$  such that  $C(\mathbf{x}, j_\alpha(\mathbf{x})) \leq \alpha \min_i C(\mathbf{x}, i)$ . In other words, the oracle computes a pure strategy for Player II whose payoff against  $\mathbf{x}$  is within a factor  $\alpha$  of the payoff given by Player II's best response.

Let  $V^*$  be the value of the game. We say that a strategy  $\hat{\mathbf{x}}$  for Player I is  $\alpha$ -optimal if  $C(\hat{\mathbf{x}}, \mathbf{y}) \geq V^*/\alpha$  for any strategy  $\mathbf{y}$  of Player II. We say a strategy for Player II is  $\alpha$ -optimal if  $C(\mathbf{x}, \hat{\mathbf{y}}) \leq \alpha V^*$  for any strategy  $\mathbf{x}$  of Player I.

We assume that mixed strategies  $\mathbf{x}$  and  $\mathbf{y}$  for the two players are represented in sparse form, as a list containing the index of each non-zero entry, together with its value. This is especially important for  $\mathbf{y}$ , as the number of pure strategies for Player II may not be polynomial in  $n$ , and we want the running time to depend polynomially on  $n$ .

Let  $A$  denote the game matrix. We assume an oracle for computing  $A[i, j]$ . That is, we assume there is a polynomial-time algorithm (polynomial in  $n$ ) that takes as input (representations of) pure strategies  $i$  and  $j$  of Players I and II, and outputs  $A[i, j]$ .

## 3 Ellipsoid Approach

Let  $\mu = \max_{i,j} A[i, j]$  be the largest entry in the game matrix. We assume this quantity is given as part of the input, or that it can be computed in time polynomial in  $n$ . In fact, it is sufficient to compute an upper bound  $\mu$  that is polynomial in  $\max_{i,j} A[i, j]$  and  $n$ .

Consider the standard LP representing the problem of finding an optimal strategy for Player I.

---

**LPI:** Maximize  $V$

such that

- (1)  $\sum_i x_i A[i, j] \geq V$  for all  $j$
  - (2)  $\sum_i x_i = 1$
  - (3)  $x_i \geq 0$  for all  $i$
- 

A separation oracle for LPI takes as input an assignment to the variables of the LP, and either reports that it is a feasible solution, or returns a hyperplane separating the assignment from the feasible region. If we had a separation oracle for LPI, we could use the ellipsoid algorithm to obtain an optimal strategy for Player I.

Although we do not have such an oracle, we can use the  $\alpha$ -approximate best response as an approximate separation oracle for LPI, as follows. A query to a separation oracle for LPI corresponds to a pair  $\langle \mathbf{x}, V \rangle$ , where the first element is the assignment to the  $x_i$  and the second is the assignment to  $V$ . Given such a query, if a constraint of type (2) or (3) of LPI is violated, return it as a separating hyperplane. Otherwise,  $\mathbf{x}$  corresponds to a mixed strategy of Player I. Query the  $\alpha$ -approximate best response oracle on  $\mathbf{x}$ . Let  $j^* = j_\alpha(\mathbf{x})$  be the returned strategy for Player II. Then check whether  $\sum_i x_i A(i, j^*) \geq V$  is satisfied by the queried assignment  $\langle \mathbf{x}, V \rangle$ . If not, return  $\sum_i x_i A(i, j^*) \geq V$  as a violated constraint. Otherwise, report (possibly incorrectly) that  $\langle \mathbf{x}, V \rangle$  is a feasible solution. Although  $\langle \mathbf{x}, V \rangle$  may not be feasible,  $\langle \mathbf{x}/\alpha, V \rangle$  is feasible, because  $j^*$  was returned by the  $\alpha$ -approximate best response oracle.

With this approximate separation oracle, we can apply the ellipsoid approach from the approximation algorithms literature, referenced above in Section 1.2. (In fact, Feldman *et al.* (2012) uses a minor variant of the approach that yields a somewhat smaller restricted dual LP. For simplicity of presentation, we do not discuss this variant here.)

It yields an algorithm that computes  $\alpha$ -optimal strategies for Players I and II. The algorithm makes use of the dual of LPI, which computes an optimal strategy for Player II. We present the dual LP as LPII.

---

**LPII:** Minimize  $V$

such that

- (1)  $\sum_j y_j A[i, j] \leq V$  for all  $i$
  - (2)  $\sum_j y_j = 1$
  - (3)  $y_j \geq 0$  for all  $j$
- 

The algorithm is as follows. To compute an  $\alpha$ -approximate strategy for Player I, run the ellipsoid algorithm using the approximate separation oracle for LPI described above. Let  $S$  be the set of violated constraints returned by the oracle. Let  $\hat{\mathbf{x}}$  denote the strategy that is returned. Return it as the  $\alpha$ -optimal strategy for Player I.

Let  $S'$  be the set of variables  $y_j$  in LPII, corresponding to the constraints in  $S$ . Let LPIIRestrict denote the restriction of LPII that is produced by setting all  $y_j \notin S'$  to 0. Generate LPIIRestrict

explicitly, computing  $A[i, j]$  for all  $j \in S'$  and for all  $i$ , to obtain the constraints. Solve this LP using a standard polynomial-time LP algorithm, such as the ellipsoid algorithm, to obtain an optimal solution  $\hat{\mathbf{y}}$  for LPIRestrict. Return  $\hat{\mathbf{y}}$  as the strategy for Player II.

This algorithm is the basis for the following theorem. In the proof, we provide the analysis of the algorithm for the benefit of the reader; the analysis is essentially the same as that used by e.g., Jansen (2003).

**Theorem 1** *Let  $V^*$  be the value of the game. There exists an algorithm that computes strategies  $\hat{\mathbf{x}}$  and  $\hat{\mathbf{y}}$  for Players I and II such that*

$$\begin{aligned} C(\mathbf{x}, \hat{\mathbf{y}}) &\leq \alpha V^* \text{ for any strategy } \mathbf{x} \text{ of Player I and} \\ C(\hat{\mathbf{x}}, \mathbf{y}) &\geq \frac{V^*}{\alpha} \text{ for any strategy } \mathbf{y} \text{ of Player II.} \end{aligned}$$

*This algorithm, which uses the ellipsoid algorithm, runs in time polynomial in  $n$  and  $\log \mu$ , where  $\mu$  is the largest entry in the game matrix.*

**Proof:** Proof. The algorithm presented above computes the strategies  $\hat{\mathbf{x}}$  and  $\hat{\mathbf{y}}$ . We now prove that  $\hat{\mathbf{x}}$  and  $\hat{\mathbf{y}}$  are  $\alpha$ -optimal strategies for Players I and II respectively. To show that  $\hat{\mathbf{x}}$  is an  $\alpha$ -optimal strategy for Player I, let  $\text{LPI}(v)$  denote the feasibility problem for the set of constraints of LPI, with  $V = v$ . Let  $P(v) = \{\mathbf{x} \mid \mathbf{x} \text{ is a feasible solution to } \text{LPI}(v)\}$ . First consider what happens if you run the ellipsoid algorithm using an exact separation oracle for LPI. It performs a binary search on the values in an interval  $I := \{0, \gamma, 2\gamma, \dots, \mu\}$ , where  $\gamma$  and  $1/\gamma$  is polynomial in  $n$  and  $\log \mu$ . The optimal value of LPI is guaranteed to be a value in  $I$ . During the binary search, for each tested value  $v$  in  $I$ , the ellipsoid algorithm makes a sequence  $Q_v$  of queries  $\langle \mathbf{x}, V \rangle$  to the separation oracle, where in each query,  $V = v$ . The querying continues until either (a) the separation oracle returns the queried assignment  $\langle \mathbf{x}, v \rangle$  (indicating that  $\mathbf{x} \in P(v)$ ) or (b) the set of violated constraints returned by the separation oracle, in answer to the queries in  $Q_v$ , implies that  $P(v) = \emptyset$  (because there is no assignment satisfying all constraints in that set). The binary search finds the largest value of  $v \in I$  for which  $P(v) \neq \emptyset$ . The ellipsoid algorithm then returns this largest value  $v$ , along with the associated  $\mathbf{x} \in \text{LPI}(v)$ , as the optimal solution to LPI.

Now consider what happens when the ellipsoid algorithm is instead run with the approximate separation oracle described above. Recall that  $V^*$  is the optimal value of LPI. The binary search is run, and ends by returning a final value of  $v$ , with an associated assignment  $\mathbf{x}$ . Let  $v^f$  denote this final value, and let  $\hat{\mathbf{x}}$  be the associated assignment. By the properties of the approximate separation oracle, and the behavior of binary search,  $\hat{\mathbf{x}} \in P(v^f/\alpha)$  and  $P(v^f + \gamma) = \emptyset$ . Because  $P(v^f + \gamma) = \emptyset$ ,

$$v^f \geq V^*. \tag{1}$$

Further, because  $\hat{\mathbf{x}} \in P(v^f/\alpha)$ , it follows that for all strategies  $\mathbf{y}$  of Player II,

$$C(\hat{\mathbf{x}}, \mathbf{y}) \geq v^f/\alpha, \tag{2}$$

and hence from (1),

$$C(\hat{\mathbf{x}}, \mathbf{y}) \geq V^*/\alpha \tag{3}$$

for all strategies  $\mathbf{y}$  of Player II. Thus  $\hat{\mathbf{x}}$  is an  $\alpha$ -optimal strategy for Player I.

By the well-known results in Grotschel *et al.* (1988) on the runtime of the ellipsoid algorithm, the above computation of  $\hat{\mathbf{x}}$  takes time polynomial in  $n$  and  $\log \mu$ .

The proof that  $\hat{\mathbf{y}}$  is an  $\alpha$ -approximate strategy for Player II is as follows. Let  $V^R$  denote the optimal value of LPIIRestrict. Since  $\hat{\mathbf{y}}$  is an optimal solution to LPIIRestrict,

$$C(\mathbf{x}, \hat{\mathbf{y}}) \leq V^R \tag{4}$$

for all strategies  $\mathbf{x}$  of Player I.

Because  $\hat{\mathbf{x}} \in P(v^f/\alpha)$ ,

$$v^f/\alpha \leq V^*. \tag{5}$$

Since LPII is dual to LPI, the optimal value of LPII is also  $V^*$ .

Let LPIRelax denote the dual of LPIIRestrict obtained from LPI by removing the constraints  $\sum_i x_i A[i, j] \geq V$  for  $j \notin S'$ . By duality,  $V^R$  is also the optimal value for LPIRelax. The violated constraints returned when running the ellipsoid algorithm with the approximate separation oracle for LPI are also violated constraints of LPIRelax. Therefore, the answers given by the approximate separation oracle for LPI could also have been given for LPIRelax. Since the execution of the ellipsoid algorithm depends entirely on the answers to the separation oracle queries, and not on other properties of the LP, it follows by the same argument used to prove (1) that

$$v^f \geq V^R. \tag{6}$$

Combining that with (5), we get  $V^R \leq \alpha V^*$ . From (4), it immediately follows that  $\hat{\mathbf{y}}$  is an  $\alpha$ -optimal strategy for Player II.

Because the computation of  $\hat{\mathbf{x}}$  takes time polynomial in  $n$  and  $\log \mu$ , the number of constraints in LPIIRestrict is polynomial in those parameters. It follows that computing  $\hat{\mathbf{y}}$  takes time polynomial in those quantities.  $\square$

## 4 Multiplicative Weights Approach

We now present an alternative approach to solving (or approximately solving) games using an  $\alpha$ -approximate best response oracle. This approach is more “combinatorial” in nature, in the sense that it does not depend on the values of the payoffs in the same way as the approach of Section 3.

Our algorithm takes the same high level approach as the algorithm of Freund and Schapire (1999), which finds strategies that approximate the optimal strategies within an *additive*  $\varepsilon$ . Player I begins with a uniformly random strategy. Player I then repeatedly plays mixed strategies, to which Player II replies with best responses. In each round, Player I updates her strategy, using some multiplicative update rule based on the best response of Player II. The more successful the pure strategy is (that is, the higher payoff it achieves), the more weight Player I gives it in the next round.

The algorithm is based on a simple observation: finding optimal strategies for Player II is easily reducible to the problem of solving a *Packing LP*. An approximately optimal solution to the latter problem is given by a multiplicative weights update algorithm, due to Garg and Könemann (2007). Our algorithm uses essentially the same update rule, and we also rely on part of their analysis to show that this gives an approximately optimal solution for Player II. We show how the multiplicative weights update algorithm also yields an approximately optimal strategy for Player

I. We believe the strategy for Player I to be an original contribution, which does not immediately follow from the analysis of Garg and Könemann (2007).

The relation between our games and packing LPs is as follows. A packing LP is a linear program of the form  $\max\{\mathbf{c}^T \mathbf{y} \mid A\mathbf{y} \leq \mathbf{b}, \mathbf{y} \geq 0\}$  where  $A$ ,  $\mathbf{b}$ , and  $\mathbf{c}$  all have positive entries. If  $A$  is the payoff matrix of a two-player zero-sum game,  $\mathbf{b}^T = (1, \dots, 1)$ , and  $\mathbf{c}^T = (1, \dots, 1)$ , and  $\mathbf{y}^*$  is the optimal solution to the LP, it is easy to show that the strategy that chooses each column  $j$  with probability proportional to  $y_j^*$  is an optimal strategy for the Player II. Thus one can reduce the problem of computing an optimal strategy for Player II to the problem of solving the packing LP. Because the reduction is so simple (just a scaling of the variables), we integrate it into our algorithm, rather than performing it explicitly as a separate step. We note that a special case of this reduction was exploited previously by Condon *et al.* (2009) in the solution of a variant of the game BOX. However, the resulting specialized packing LP was then solved by different methods.

We first describe our algorithm in detail, then give a bound on the run time, and finally prove that it works.

Fix  $\eta > 0$  and  $\delta > 0$  (we will specify their values later) and let  $\mathbf{x}^{(1)}$  be the mixed strategy for Player I given by  $x_i^{(1)} = 1/n$  for all  $i$ .

1. Set  $t = 1$ .
2. Use the  $\alpha$ -approximate best response oracle to compute  $j^{(t)} = j_\alpha(\mathbf{x}^{(t)})$ .
3. Set  $M^{(t)} = \max_i C(i, j^{(t)})$ .
4. Define  $x_i^{(t+1)} = \left( \frac{1 + \eta C(i, j^{(t)}) / M^{(t)}}{1 + \eta C(\mathbf{x}^{(t)}, j^{(t)}) / M^{(t)}} \right) x_i^{(t)}$ . (The expression in the denominator is a normalizing factor to ensure that  $\mathbf{x}^{(t+1)}$  is a mixed strategy.)
5. If  $f(t) := \prod_{t' \leq t} \left( 1 + \eta C(\mathbf{x}^{(t')}, j^{(t')}) / M^{(t')} \right) > 1/(\delta n)$ , then set  $T = t$  and stop. Else increase  $t$  by 1 and return to Step 2.

Let  $\hat{\mathbf{y}}$  be the mixed strategy that chooses pure strategy  $j^{(t)}$  with probability proportional to  $1/M^{(t)}$ . More precisely,  $\hat{y}_j$  is proportional to  $\sum_{\{t: j^{(t)}=j\}} 1/M^{(t)}$ . Let  $\hat{t}$  be the value of  $t$  that maximizes  $C(\mathbf{x}^{(t)}, j^{(t)})$  and let  $\hat{\mathbf{x}} = \mathbf{x}^{(\hat{t})}$ . We will later show that these strategies  $\hat{\mathbf{x}}$  and  $\hat{\mathbf{y}}$  are approximately optimal.

The proof of the following lemma is from Garg and Könemann (2007) (following the presentation in Liu *et al.* (2008)).

**Lemma 1** *The number of iterations  $T$  of the algorithm satisfies*

$$T \leq n \left( 1 - \frac{\ln \delta}{\ln(1 + \eta)} \right). \quad (7)$$

**Proof:** Proof. For  $i = 1, \dots, n$ , let  $z_i^{(1)} = \delta$  and for  $t \geq 2$  let  $\mathbf{z}^{(t)} = \delta n f(t-1) \mathbf{x}^{(t)}$ . It follows that

$$z_i^{(t+1)} = (1 + \eta C(i, j^{(t)}) / M^{(t)}) z_i^{(t)}. \quad (8)$$



Note that the stopping condition  $f(t) > 1/(\delta n)$  is equivalent to the condition  $\sum_i z_i^{(t+1)} > 1$ , since

$$\sum_i z_i^{(t+1)} = \sum_i \delta n f(t) x_i^{(t+1)} = \delta n f(t).$$

Therefore, since the algorithm terminates on the  $T$ th iteration, we must have  $\sum_i z_i^{(T)} \leq 1$  and  $\sum_i z_i^{(T+1)} > 1$ .

In each iteration of the algorithm, there is some pure strategy  $i$  for Player I such that  $C(i, j^{(t)}) = M^{(t)}$ . Call such a strategy a *bottleneck strategy*. If  $i$  is a bottleneck strategy in iteration  $t$  then  $z_i^{(t+1)} = (1 + \eta)z_i^{(t)}$ , so if  $i$  is a bottleneck strategy a total of  $m$  times then  $z_i^{(T+1)} \geq \delta(1 + \eta)^m$ . But also  $z_i^{(T+1)} \leq 1 + \eta$ , since otherwise  $z_i^{(T)} > 1$ , contradicting  $\sum_i z_i^{(T)} \leq 1$ . Putting together these inequalities,

$$\begin{aligned} \delta(1 + \eta)^m &\leq z_i^{(T+1)} \leq 1 + \eta, \text{ so} \\ m &\leq 1 - \frac{\ln \delta}{\ln(1 + \eta)}. \end{aligned}$$

Inequality (7) follows from the fact that Player I has  $n$  pure strategies and in each iteration at least one of them is a bottleneck strategy.  $\square$

We now prove that for an appropriate choice of the parameters  $\delta$  and  $\eta$ , the strategies produced by the algorithm are approximately optimal and that the total runtime is not too large.

**Theorem 2** *Let  $V^*$  be the value of the game. For any  $\varepsilon > 0$ , there exist algorithms that compute strategies  $\hat{\mathbf{x}}$  and  $\hat{\mathbf{y}}$  for Player I and II such that*

$$\begin{aligned} C(\mathbf{x}, \hat{\mathbf{y}}) &\leq \alpha(1 + \varepsilon)V^* \text{ for any strategy } \mathbf{x} \text{ of Player I and} \\ C(\hat{\mathbf{x}}, \mathbf{y}) &\geq \frac{V^*}{\alpha(1 + \varepsilon)} \text{ for any strategy } \mathbf{y} \text{ of Player II.} \end{aligned}$$

*The algorithm that computes  $\hat{\mathbf{y}}$  runs in time polynomial in  $n$  and  $1/\varepsilon$ ; the algorithm that computes  $\hat{\mathbf{x}}$  runs in time polynomial in  $n$ ,  $1/\varepsilon$  and  $\alpha$ .*

**Proof:** Proof.

We have

$$x_i^{(T+1)} = \left( \prod_{t \leq T} \frac{1 + \eta C(i, j^{(t)})/M^{(t)}}{1 + \eta C(\mathbf{x}^{(t)}, j^{(t)})/M^{(t)}} \right) x_i^{(1)} = \left( \prod_{t \leq T} \frac{1 + \eta C(i, j^{(t)})/M^{(t)}}{1 + \eta C(\mathbf{x}^{(t)}, j^{(t)})/M^{(t)}} \right) \frac{1}{n}.$$

Since every coordinate of  $\mathbf{x}^{(T+1)}$  is at most 1, we have

$$n \geq n x_i^{(T+1)} = \frac{\prod_{t \leq T} 1 + \eta C(i, j^{(t)})/M^{(t)}}{\prod_{t \leq T} 1 + \eta C(\mathbf{x}^{(t)}, j^{(t)})/M^{(t)}} \geq \frac{(1 + \eta)^{\sum_{t \leq T} C(i, j^{(t)})/M^{(t)}}}{\exp(\eta \sum_{t \leq T} C(\mathbf{x}^{(t)}, j^{(t)})/M^{(t)})} \quad (9)$$

The second inequality above follows from the facts that

$$(1 + \eta z) \geq (1 + \eta)^z \text{ for } z \in [0, 1],$$

and

$$(1 + z) \leq \exp(z) \text{ for all } z.$$

Taking natural logarithms in (9), we get

$$\ln n \geq \ln(1 + \eta) \sum_{t \leq T} C(i, j^{(t)})/M^{(t)} - \eta \sum_{t \leq T} C(\mathbf{x}^{(t)}, j^{(t)})/M^{(t)}.$$

Now we use the fact that  $\ln(1 + \eta) \geq \eta - \eta^2$  and rearrange:

$$\sum_{t \leq T} C(\mathbf{x}^{(t)}, j^{(t)})/M^{(t)} \geq -\frac{\ln n}{\eta} + (1 - \eta) \sum_{t \leq T} C(i, j^{(t)})/M^{(t)}.$$

By the linearity of the cost function and the right-hand side, the inequality above is still true if we replace the terms  $C(i, j^{(t)})$  by  $C(\mathbf{x}, j^{(t)})$ , where  $\mathbf{x}$  is any arbitrary strategy for Player I. Combining this with the fact that  $C(\mathbf{x}^{(t)}, j^{(t)}) \leq \alpha V^*$ , we get

$$\alpha V^* \sum_{t \leq T} 1/M^{(t)} \geq \sum_{t \leq T} C(\mathbf{x}^{(t)}, j^{(t)})/M^{(t)} \geq -\frac{\ln n}{\eta} + (1 - \eta) \sum_{t \leq T} C(\mathbf{x}, j^{(t)})/M^{(t)}. \quad (10)$$

By definition of  $\hat{\mathbf{y}}$ ,

$$C(\mathbf{x}, \hat{\mathbf{y}}) = \frac{\sum_{t \leq T} C(\mathbf{x}, j^{(t)})/M^{(t)}}{\sum_{t \leq T} 1/M^{(t)}},$$

and substituting this into (10) and rearranging gives

$$\alpha V^* \geq -\frac{\ln n}{\eta \sum_{t \leq T} 1/M^{(t)}} + (1 - \eta) C(\mathbf{x}, \hat{\mathbf{y}}). \quad (11)$$

By the stopping condition,  $f(T) > 1/(\delta n)$ , so

$$\begin{aligned} \frac{1}{\delta n} &< \prod_{t \leq T} \left(1 + \eta C(\mathbf{x}^{(t)}, j^{(t)})/M^{(t)}\right) \\ &\leq \prod_{t \leq T} \left(1 + \eta \alpha V^*/M^{(t)}\right) \text{ (by definition of } j^{(t)}) \\ &\leq \exp\left(\sum_{t \leq T} \eta \alpha V^*/M^{(t)}\right), \text{ so taking logs and rearranging,} \\ \sum_{t \leq T} \frac{1}{M^{(t)}} &\geq -\frac{\ln(n\delta)}{\eta \alpha V^*}. \end{aligned} \quad (12)$$

We will see later than  $n\delta < 1$ , so that  $\ln(n\delta)$  is negative. Combining (12) with (11) gives

$$C(\mathbf{x}, \hat{\mathbf{y}}) \leq \frac{\alpha V^*}{(1 - \eta)(1 + \ln n / \ln \delta)}. \quad (13)$$

Let  $\eta = ((1 + \varepsilon)^{1/2} - 1)/2$  and let  $\delta = n^{-1/\eta}$ , and use  $(1 - \eta)^{-1} \leq (1 + 2\eta)$  to get

$$C(\mathbf{x}, \hat{\mathbf{y}}) \leq \frac{\alpha V^*}{(1 - \eta)^2} \leq (1 + 2\eta)^2 \alpha V^* = \alpha(1 + \varepsilon)V^*.$$

With these choices of  $\eta$  and  $\delta$ , by Lemma 1, the number of iterations  $T$  of the algorithm satisfies

$$\begin{aligned} T &\leq n \left( 1 + \frac{(1/\eta) \ln n}{\ln(1 + \eta)} \right) \\ &\leq n \left( 1 + (1/\eta + 1/\eta^2) \ln n \right) \quad (\text{using } \ln(1 + \eta) \geq \eta/(1 + \eta)) \\ &\leq n(1 + (8/\varepsilon + 8/\varepsilon^2) \ln n) \quad (\text{since } \eta \geq \varepsilon/8). \end{aligned}$$

This is clearly polynomial in  $n$  and  $1/\varepsilon$ .

We now turn to Player I's strategy  $\hat{\mathbf{x}}$ . We will choose different, smaller values for both  $\eta$  and  $\delta$  (to be specified later), giving a longer runtime for computing  $\hat{\mathbf{x}}$  that is polynomial in  $n$ ,  $1/\varepsilon$  and  $\alpha$ . If we wish to calculate both strategies we could either run the algorithm twice using the two different choices of  $\eta$  and  $\delta$ , or simply run the algorithm once using the second choice. This would give Player II's strategy as well as Player I's since the right-hand side of (13) is increasing in  $\delta$  and  $\varepsilon$ .

First let  $\mathbf{x}^*$  be an optimal strategy of Player I. Then  $C(\mathbf{x}^*, j^{(t)}) \geq V^*$  because  $\mathbf{x}^*$  is optimal. Hence, substituting  $\mathbf{x} = \mathbf{x}^*$  into (10),

$$\sum_{t \leq T} C(\mathbf{x}^{(t)}, j^{(t)})/M^{(t)} \geq -\frac{\ln n}{\eta} + (1 - \eta) \sum_{t \leq T} C(\mathbf{x}^*, j^{(t)})/M^{(t)} \geq -\frac{\ln n}{\eta} + (1 - \eta)V^* \sum_{t \leq T} 1/M^{(t)}. \quad (14)$$

By definition of  $j^{(t)}$ , we have  $C(\hat{\mathbf{x}}, j^{(t)}) \leq \alpha C(\hat{\mathbf{x}}, \mathbf{y})$  for any strategy  $\mathbf{y}$  so

$$\begin{aligned} C(\hat{\mathbf{x}}, \mathbf{y}) &\geq \frac{1}{\alpha} C(\hat{\mathbf{x}}, j^{(t)}) \\ &\geq \frac{1}{\alpha \sum_{t \leq T} 1/M^{(t)}} \sum_{t \leq T} C(\mathbf{x}^{(t)}, j^{(t)})/M^{(t)} \quad (\text{by definition of } \hat{t}) \\ &\geq \frac{-\ln n/\eta}{\alpha \sum_{t \leq T} 1/M^{(t)}} + \frac{(1 - \eta)V^*}{\alpha} \quad (\text{by (14)}) \\ &\geq \frac{V^* \ln n}{\ln(n\delta)} + \frac{(1 - \eta)V^*}{\alpha} \quad (\text{by (12)}) \\ &= \frac{V^*}{\alpha} \left( 1 - \eta + \frac{\alpha \ln n}{\ln(n\delta)} \right). \end{aligned}$$

Again, take  $\eta = ((1 + \varepsilon)^{1/2} - 1)/2$  this time take  $\delta = n^{-1-\alpha/\eta}$ , so that

$$C(\hat{\mathbf{x}}, \mathbf{y}) \geq \frac{V^*}{\alpha} (1 - 2\eta) \geq \frac{V^*}{\alpha(1 + 2\eta)^2} = \frac{V^*}{\alpha(1 + \varepsilon)}.$$

For these values of  $\eta$  and  $\delta$ , the bound on the runtime given by Lemma 1 is

$$\begin{aligned} T &\leq n \left( 1 + \frac{(1 + \alpha/\eta) \ln n}{\ln(1 + \eta)} \right) \\ &\leq n \left( 1 + (1/\eta + 1/\eta^2)(\alpha + \eta) \ln n \right) \\ &\leq n(1 + (8/\varepsilon + 8/\varepsilon^2)(\alpha + \varepsilon) \ln n) \quad (\text{again using } \eta \geq \varepsilon/8 \text{ and also } \eta \leq \varepsilon). \end{aligned}$$

This time  $T$  is polynomial in  $n$ ,  $1/\varepsilon$  and  $\alpha$ . □

## 5 Applications to Search Games

In this section we describe how our algorithms can be applied to several known and new search games. Each game is played between a Hider, who corresponds to Player I in our theorems, and a Searcher, who corresponds to Player II. In this section we present several theorems concerning the existence of polynomial time algorithms for computing optimal or approximately optimal strategies for search games. These theorems refer to algorithms based on either the ellipsoid approach in Section 3 or the multiplicative weights approach in Section 4.

### 5.1 Searching in Boxes with Precedence Constraints

We begin with a generalization of the game BOX discussed in Subsection 1.1. In particular, we restrict the Searcher only to orderings that are consistent with some predefined precedence constraints on the boxes. That is, we suppose that there is a partial order  $\prec$  on the boxes so that box  $i$  can be searched before box  $j \neq i$  if and only if  $i \prec j$ . The Hider's strategy set and the payoff function remain unchanged. We call this game PREC, and while it has not been studied before in the form we define it in here, a more general version of it was studied in Fokkink *et al* (2016b), as we will discuss in the next subsection.

In order to apply our algorithms to PREC we must consider the best response problem of choosing the search that minimizes the expected search cost for a given Hider distribution  $\mathbf{x}$ . Similarly as for BOX, we can reframe this problem as the scheduling problem of finding a precedence-constrained ordering of a set of jobs with processing times and weights to minimize the sum of the weighted completion times. This problem is usually denoted  $1|prec|\sum w_j C_j$  in the scheduling literature. It is well known to be NP-hard, and many 2-approximation algorithms can be found, for instance Chekuri *et al.* (1999) and Schulz (1996). These algorithms can be used as a 2-approximate best response oracle, so we obtain our first new result for search games.

**Theorem 3** *There exist polynomial time algorithms that compute 2-approximate strategies for both players in PREC.*

### 5.2 The Submodular Search Game

This game was introduced by Fokkink *et al.* (2016a), and further studied by Fokkink *et al* (2016b). As in the game, BOX, the Hider's strategy set is  $[n]$  and the Searcher's strategy set is all permutations  $\pi$  of  $[n]$ . The difference is in the payoff function. Let  $f : S \rightarrow \mathbb{R}^+$  be a non-decreasing submodular set function on  $S$ , which corresponds to the cost of searching subsets of  $S$ . For a given hiding location  $j$  and a given permutation  $\pi$ , the payoff  $C(j, \pi)$  of the game is defined as the cost of the set of locations searched up to and including  $j$ . That is,

$$C(j, \pi) = f(\{i : i \leq \pi^{-1}(j)\}).$$

We call this game SUB. It is a further generalization of PREC. This can be seen by defining a cost function  $f$  on subsets  $S$  of  $[n]$  in PREC such that  $f(S)$  is the cost of all the elements of  $[n]$  in the precedence-closure of  $S$ . Then, under this cost function, optimal strategies for SUB will correspond to optimal strategies in PREC.

Fokkink *et al* (2016b) also consider the best response problem for SUB, and they prove that this problem can be solved approximately, within a factor of 2, generalizing the analogous result

in the scheduling literature. It follows that our algorithms can be used to calculate 2-approximate strategies for the players in SUB, and so we obtain our next new result for search games.

**Theorem 4** *There exist polynomial time algorithms that calculate 2-approximate strategies for both players in SUB.*

In Fokkink *et al* (2016b), the authors obtain  $1/(1 - k)$ -approximate strategies, where  $k \in [0, 1]$  is the *total curvature* of the function  $f$ . For  $k > 1/2$ , our theorem is an improvement on this.

### 5.3 Expanding Search

In PREC, the partial order  $\prec$  on  $[n]$  uniquely defines a directed acyclic graph, and if that graph is a tree then PREC is equivalent to the *expanding search game* on a tree, introduced in Alpern and Lidbetter (2013b). The game, which we will call EXP, can be played on any undirected graph with a distinguished vertex called the *root* and edges with costs corresponding to the time taken to traverse them. A pure strategy for the Hider is a vertex at which to hide, and a pure strategy for the Searcher is a sequence of edges, the first of which is incident to the root, and each other of which must be adjacent to a previously chosen edge. This type of search on a graph is called an *expanding search*. The payoff is the search cost, which is the sum of the costs of the edges chosen by the Searcher up to and including the first edge that is incident to the Hider’s chosen vertex.

Expanding search differs from the usual model of search, where the Searcher chooses a walk on the graph, and is more appropriate for situations in which the area being searched expands at a constant rate, for example in coal mining, where the time taken to recommence digging from a point already reached is negligible compared to the time taken to dig. Expanding search is also a better model for other situations in which searching must be slow but movement can be quick, such as the search for landmines or search-and-rescue scenarios in which a region must be slowly combed. As explained in more detail in Alpern and Lidbetter (2013b), expanding search can also be interpreted as the search of a large team of agents, such as a SWAT team searching for a terrorist, who walk together on a graph and successively split up into smaller teams.

For EXP played on a tree, the best response problem is what Alpern and Lidbetter (2013b) call the *Bayesian problem* of finding a search that minimizes the expected payoff for a known hiding distribution on the vertices of the tree, and the authors give a solution to this problem. The solution to the equivalent scheduling problem  $1|prec|\sum w_j C_j$  for tree-like precedence constraints is also known, and can be solved using the so-called *Sidney decomposition* proposed by Sidney (1975). Therefore, our algorithms can be applied to find optimal strategies for the players in the game. However, optimal strategies have already been found in closed form for this game in Alpern and Lidbetter (2013b).

If we consider EXP on a more general graph, then the best response problem can be shown to be NP-hard (Dürr (2016)). Whether the problem can be efficiently approximated is an open question. If this were the case then our results would imply that approximately optimal strategies for EXP could also be found.

### 5.4 Expanding Search Ratio

This is a game that was introduced by Angelopoulos *et al.* (2016). The game is played on a graph with  $n$  vertices plus a root vertex and edges with costs. The Hider chooses one of the  $n$  non-root vertices of the tree and the Searcher chooses an expanding search, as in the expanding search game

considered in Alpern and Lidbetter (2013b). However, for a given vertex  $v$  and a given expanding search, the payoff is not the search cost, but the ratio of the search cost to the shortest path distance  $d(v)$  from the root to  $v$ . This payoff is called the *search ratio*. It is a measure of the “regret” incurred by the Searcher in finding the Hider, and so we call the game EXPr.

Angelopoulos *et al.* (2016) find a  $5/4$ -optimal strategy for the Searcher in EXPr for star graphs, which is generalized in Angelopoulos *et al.* (2017) to general unweighted graphs (where all the edges have equal cost) and tree graphs. Independently, Condon *et al.* (2009) have studied the same game for star graphs in the context of throughput maximization, and they give a full solution. But the solution of EXPr for any other class of graphs is unknown.

Consider the game played on a tree graph, and we will show that in this case, the best response problem can be solved in polynomial time. Indeed, suppose the Hider is located at vertex  $v$  with probability  $p(v)$ . Let  $x(v) = p(v)/d(v)$  and let  $\hat{\mathbf{x}}$  be the normalization of  $\mathbf{x}$ , so that  $\hat{\mathbf{x}}$  is a probability vector. Then the best response problem is the problem of finding the expanding search that minimizes the expected search ratio with respect to  $\mathbf{p}$ , which is clearly equivalent to the problem of finding the expanding search that minimizes the expected *search time* with respect to  $\hat{\mathbf{x}}$ . The latter problem is the best response problem for EXP, which, as discussed in Subsection 5.3, can be solved in polynomial time. Therefore we obtain the following theorem.

**Theorem 5** *There exist polynomial time algorithms that calculate optimal strategies for both players in EXPr played on a tree.*

For other types of graphs, using our results depends on finding approximations to the best response problem for EXP.

## 5.5 Search Games with Regret

The payoff matrix of the game EXPr discussed in the previous subsection, can be obtained from the payoff matrix of the game EXP by dividing each row of the matrix by a constant, which is equal to the shortest path distance to the vertex corresponding to that row. If we take any search game, we can define a “regret version” of that game by dividing the row corresponding to each Hider pure strategy  $j$  by some constant. There is usually a natural choice of constant, for example in SUB, we might divide the row corresponding to Hider strategy  $j$  by  $f(\{j\})$ , the cost that would be incurred by the Searcher if she knew where the Hider was (provided  $f(\{j\}) > 0$  for all  $j$ ).

If there is an  $\alpha$ -approximate best response oracle for some search game, then it follows that there is also an  $\alpha$ -approximate best response oracle for any regret version of that game, using similar reasoning as in Subsection 5.4. This implies, using our results, that we can find  $\alpha$ -optimal strategies in both games. So, for instance, we immediately find that for the regret version of SUB, there is an efficient algorithm that finds 2-approximate strategies.

We sum this up below.

**Theorem 6** *Suppose there is an  $\alpha$ -approximate best response oracle for some game we call GAME with payoff function  $C$ . Let GAMEr be the game defined by the payoff function  $C'(i, j) := k_i C(i, j)$ , where  $k_1, \dots, k_n$  are non-negative constants. Then there is an  $\alpha$ -approximate best response oracle for GAMEr, and hence there exist polynomial time algorithms for calculating  $\alpha$ -optimal strategies.*

## 5.6 Hide-Seek and Pursuit-Evasion

We now discuss a different type of game, introduced by Gal and Casas (2014) in the context of predator-prey interaction, though the game could equally apply to problems of national security such as the pursuit of a terrorist. As in the games discussed above, a Hider or hidden object is located in one of  $n$  locations (this time without search costs). The Searcher can search a subset of locations size at most  $k$ , for some given  $k$ , and if the Hider's location  $j$  lies in this set, a pursuit ensues and the Searcher captures him with probability  $p_j$ . The payoff of the game, which the Hider wishes to minimize and the Searcher to maximize, is the probability that the Hider is captured. Gal and Casas (2014) provide a closed form solution to this game.

Here we introduce a generalization of this game, in which the locations have search costs  $c_j$ . A pure strategy for the Searcher is a subset of locations, the sum of whose search costs does not exceed  $k$ . The payoff, as before, is the probability the Hider is captured. We call this game HSPE (hide-seek and pursuit-evasion). If all the locations' costs are equal to 1, then HSPE is equivalent to the original game of Gal and Casas (2014). Consider the best response problem, where the probability  $x_j$  that the Hider is in location  $j, j = 1, \dots, n$  is known and the problem is to choose a subset of locations that maximizes the probability of capture. Let  $w_j = p_j x_j$ . Then the problem is to choose a subset  $S$  of  $[n]$  of total cost  $c(S) := \sum_{j \in S} c_j$  at most  $k$ , that maximizes  $w(S) := \sum_{j \in S} w_j$ . This is the classic Knapsack problem, which has a fully polynomial time approximation scheme (see Vazirani (2013)). That is, given  $\varepsilon > 0$ , there exists an algorithm that is polynomial in  $n$  and  $1/\varepsilon$  that approximates the solution of the knapsack problem within a multiplicative factor of  $1 + \varepsilon$ . This implies the following.

**Theorem 7** *The problem of finding optimal strategies in the game HSPE has a fully polynomial time approximation scheme.*

## 5.7 An Infinite Game

We finish by discussing an infinite game studied by Lin and Singham (2015), for which our results do not directly apply because the Searcher has a strategy set of (countably infinite) cardinality. We are optimistic that our results could be extended to such games in future work, which is why we mention it here.

In the game of Lin and Singham (2015), a Hider is located in one of  $n$  locations with search costs  $c_j$  and capture probabilities  $p_j$ , as in the game, HSPE of the previous subsection. The Searcher chooses an infinite sequence of locations, and each time the Searcher examines the Hider's location  $j$ , she finds him independently with probability  $p_j$ . The payoff of the game is the Searcher's expected cost of finding the Hider for the first time.

As Lin and Singham (2015) point out, the solution of the best response problem for this game is well known. Suppose the Hider is hidden in location  $j$  with probability  $x_j$ . Then, as showed by Blackwell (reported in Matula (1964)), the optimal policy for the Searcher is, at any time, to choose a location that maximizes the index  $x_j p_j / c_j$ , with the hiding probability  $x_j$  being updated according to Bayes' Law after each search. Of course, because the optimal policy is an infinite sequence, it cannot necessary be concisely expressed, but an approximately optimal solution may be found, and Lin and Singham (2015) exploit this fact to give an algorithm which they believe converges to an optimal search strategy in the game. We leave it for future work to develop a provably fully polynomial approximation scheme for the problems of determining optimal strategies in the game.

## 6 Conclusion

We have shown how we may use an algorithmic approach to find solutions or approximate solutions to search games, by exploiting oracles that find the best responses or approximate best responses. This has applications to the search games we have mentioned in Section 5, but we believe it may have further applications in other search games, and indeed in games studied in other fields of operations research such as security games.

## Acknowledgements

L. Hellerstein was partially supported by NSF Award IIS-1217968.

## References

- Alpern S (2016) Searching a Network Using Combinatorial Paths. *Oper. Res.* (in press).
- Alpern S, Gal S (2003). *The Theory of Search Games and Rendezvous*. Kluwer International Series in Operations Research and Management Science (Kluwer, Boston), 319.
- Alpern S, Lidbetter T (2013a) Searching a Variable Speed Network. *Math. Oper. Res.* 39(3):697–711.
- Alpern S, Lidbetter T (2013b) Mining Coal or Finding Terrorists: The Expanding Search Paradigm. *Oper. Res.* 61(2):265–279.
- Alpern S, Lidbetter T (2015) Optimal Trade-Off Between Speed and Acuity When Searching for a Small Object. *Oper. Res.* 63(1):122–133.
- Angelopoulos S, Dürr C, Lidbetter T (2016) The Expanding Search Ratio of a Graph. In *33rd International Symposium on Theoretical Aspects of Computer Science (STACS)*.
- Angelopoulos S, Dürr C, Lidbetter T (2017) The Expanding Search Ratio of a Graph. arXiv preprint arXiv:1602.06258.
- Arora S, Hazan E, Kale S (2012) The Multiplicative Weights Update Method: a Meta-Algorithm and Applications. *Theory of Computing*, 8(1):121–164.
- Baston V, Kikuta K (2013) Search Games on Networks with Travelling and Search Costs and with Arbitrary Searcher Starting Points. *Networks* 62(1):72–79.
- Baston V, Kikuta K (2015) Search Games on a Network with Travelling and Search Costs. *Int. J. Game Theory* 44(2):347–365.
- Carr R, Vempala S (2002) *Randomized metarounding*, *Random Struct. Algorithms*, 20(3):343–352.
- Chekuri C, Motwani R (1999) Precedence constrained scheduling to minimize sum of weighted completion times on a single machine. *Discr. Appl. Math.* 98(1):29–38.
- Condon A, Deshpande A, Hellerstein L, Wu N (2009) Algorithms for distributional and adversarial pipelined filter ordering problems. *ACM Transactions on Algorithms (TALG)* 5(2):24–34.
- Dürr C (2016) On the NP-hardness of the expanding search problem. Private communication.
- Feldman M, Kortsarz K, Nutov Z (2012) *Improved approximation algorithms for Directed Steiner Forest*, *J. Comput. Syst. Sci.* 78(1): 279–292.
- Freund Y, Schapire RE (1999) Adaptive Game Playing Using Multiplicative Weights. *Games and Economic Behavior* 29(1):79–103.
- Friggstad Z, Swamy C(2012) Approximation Algorithms for Regret-bounded Vehicle Routing and Applications to Distance-constrained Vehicle Routing. In *Proceedings of the Forty-sixth Annual ACM Symposium on Theory of Computing*.



- Fokkink R, Kikuta K, Ramsey D (2016a) The Search Value of a Set. *Ann. Oper. Res.* 1-11.
- Fokkink R, Lidbetter T, Végh L (2016b) Submodular Search is Scheduling, arXiv preprint arXiv:1607.07598.
- Freund Y and Schapire R (1996) Game Theory, On-Line Prediction and Boosting. In *Proceedings of the Ninth Annual Conference on Computational Learning Theory*.
- Hohzaki R (2016) Search Games: Literature and Survey. *J. Oper. Res. Soc. Jpn.* 59(1):1–34.
- Gal S (1979) Search Games with Mobile and Immobile Hider. *SIAM J. Control Optim.* 17(1):99–122.
- Gal S (2000) On the optimality of a simple strategy for searching graphs. *Internat. J. Game Theory* 29(4):533–542.
- Gal S (2011) *Search games*. Wiley Encyclopedia of Operations Research and Management Science, Cochran J, ed., Wiley.
- Gal S and Casas J (2014) Succession of hide–seek and pursuit–evasion at heterogenous locations. *J. R. Soc. Interf.* 11(94):20140062.
- Garg N and Könemann J (2007) Faster and Simpler Algorithms for Multicommodity Flow and Other Fractional Packing Problems. *SIAM J. Comput.* 37(2):630–652.
- Grötschel M, Lovász L, Schrijver A (1988) *Geometric Algorithms and Combinatorial Optimization*, Springer, Berlin.
- Jansen K (2003) *Approximate Strong Separation with Application in Fractional Graph Coloring and Preemptive Scheduling*, Theor. Comput. Sci.,302(1-3):239–256.
- Lidbetter T (2013) Search Games with Multiple Hidden Objects, *SIAM J. Control and Optim.* 51(4):3056–3074.
- Lin KY, Singham DI (2015) Robust Search Policies Against an Intelligent Evader. Technical report, Naval Postgraduate School, Monterey, California.
- Lin KY, Singham DI (2016) Finding a Hider by an Unknown Deadline. *Oper. Res. Lett.* 44(1):25–32.
- Liu Z, Parthasarathy S, Ranganathan A, Yang H (2008) A generic flow algorithm for shared filter ordering problems. In *Proceedings of the Twenty-Seventh ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*.
- Matula D (1964) A periodic optimal search. *Am. Math. Mon.* 71(1):15–21.
- Schulz AS (1996) Scheduling to minimize total weighted completion time: Performance guarantees of LP-based heuristics and lower bounds. In *Proceedings of the 5th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, 301–315.
- Sidney J (1975) Decomposition Algorithms for Single-Machine Sequencing with Precedence Relations and Deferral Costs. *Oper. Res.* 23(2):283–298.
- Smith WE (1956) Various optimizers for singlestage production. *Naval Res. Logist. Quarterly* 3(1-2):59–66.
- Vazirani VV (2013) *Approximation algorithms*, Springer Science & Business Media, Berlin.