TUDelft

Delft University of Technology

The search value of a set

Fokkink, Robbert; Kikuta, Ken; Ramsey, David

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

CrossMark

The search value of a set

Robbert Fokkink[1] · Ken Kikuta[2] · David Ramsey[3]

Abstract We study search games in which the hider may hide in a finite number of locations. We assume that the cost of searching these locations does not depend on the order in which the locations are searched. From these assumptions we derive that the cost function is submodular, thus placing search games with an immobile hider in the context of coalitional games.

Keywords Search game · Submodular function · Zero-sum game · Coalitional game

1 Introduction

Alpern and Asic (1985) defined the search value of a network by means of a search game that takes place on the network. In the present paper we define the search value of a set $\mathcal{V}(X)$ by means of a search game on a set X. The payoff of the search game is given by a submodular cost function $f : 2^X \to \mathbb{R}_{\geq 0}$. We prove that $\frac{f(X)}{2} < \mathcal{V}(X) \leq f(X)$ and that an optimal search corresponds to a directed random walk on a Hasse diagram. We conjecture that an optimal hider strategy is contained in the core of the game, which is defined as the polyhedron of probability vectors p_i satisfying the constraints $\sum_{i \in A} p_i \leq f(A)$ for all $A \subset X$.

✉ Robbert Fokkink
r.j.fokkink@tudelft.nl

Ken Kikuta
kikuta@biz.u-hyogo.ac.jp

David Ramsey
david.ramsey@pwr.edu.pl

[1] Delft Institute of Applied Mathematics, P.O. Box 5031, 2600 GA Delft, Netherlands

[2] School of Business Administration, University of Hyogo, Gakuennishimachi 8-2-1, Nishi-ku, Kobe 651-2197, Japan

[3] Department of Operations Research, Wroclaw University of Science and Technology, Wroclaw, Poland

We consider a zero-sum two-player game between Searcher and Hider. Hider chooses a place to hide from a finite number of locations. Searcher then goes through these locations one by one, and the game ends as soon as Searcher selects Hider's location. This game is known as a *search game on discrete locations with an immobile hider*. There exists an extensive literature on the topic, see Alpern and Gal (2003, Ch. 3) and Alpern et al. (2013, Ch. 1), but this concerns games on networks or metric spaces. In this paper we want to study search games that are not necessarily placed on a network, and therefore we follow an axiomatic approach, imposing general conditions only. Our results are motivated by a recent study of Alpern and Lidbetter on expanding search Alpern and Lidbetter (2013).

We number the hiding locations 1 to n. So Hider chooses an element from $X = \{1, \ldots, n\}$ and Searcher chooses a permutation π on n elements. The cost function $f : 2^X \to \mathbb{R}$ represents the cost of the search operation. We assume that the cost depends only on the locations that have been searched, and not on the order in which they have been searched. If Hider's location is $\pi(j)$, then the Hider receives the payoff $f(\{\pi(1), \ldots, \pi(j)\})$. In this game, Hider wants to maximize and Searcher wants to minimize the total cost of the search operation.

To simplify our notation, we will often omit brackets for singletons and write $f(x)$ instead of the more accurate $f(\{x\})$.

2 Conditions on the payoff function

We assume that the search is carried out by a team of agents who can coordinate their search operation, which is translated into the following informal conditions on f:

  (i) Searching nothing costs nothing
 (ii) Searching more costs more
(iii) Searching costs less if more has been searched

Since the game ends as soon as Searcher finds Hider, the value $f(\emptyset)$ never occurs as a payoff in the game, and condition (i) could be omitted. However, if $f(\emptyset) \neq 0$ then we can simply redefine the cost function by subtracting $f(\emptyset)$ from every other value. Therefore, we we may impose condition (i) without loss of generality. Condition (ii) says that searching is never to Searcher's benefit. This is a natural condition, but one can imagine situations in which this is not true, when there are benefits other than finding the Hider. It may be worthwhile to study games in which (ii) is omitted. The third condition says that if $A' \subset A$, then the marginal cost of searching B after searching A is at most equal to the marginal cost of searching B after searching A'.

The formal conditions on the cost function are

  (i) $f(\emptyset) = 0$
 (ii) If $A \subseteq B$ then $f(A) \leq f(B)$
(iii) $f(A \cup B) + f(A \cap B) \leq f(A) + f(B)$

Our translation into the mathematical conditions (i) and (ii) is obvious, but (iii) demands an explanation. Let f_A be the *marginal cost function* defined as $f_A(B) = f(A \cup B) - f(A)$ for $A, B \subset X$. Then from above $f_A(B) \leq f_{A'}(B)$ if $A' \subset A$. However, this is equivalent to the condition that we impose here, which says that f is *submodular*. Such functions arise naturally from many different optimization problems Fujishige (2005) and in cooperative games Shapley (1971). There exists an extensive literature on such functions. That is why we state mathematical condition (iii) in this form. Of course, we still need to prove that the submodularity of f is equivalent to the condition that $f_A(B) \leq f_{A'}(B)$ if $A' \subset A$:

Lemma 1 *Condition (iii) is equivalent to $f_A(B) \leq f_{A'}(B)$ if $A' \subseteq A$.*

Proof It is known that condition (iii) is equivalent to $f_A(\{x\}) \leq f_{A'}(\{x\})$ for $A' \subseteq A$ and for each $x \in X$, see Schrijver (2003, Thm 44.1). We prove that this implies that $f_A(B) \leq f_{A'}(B)$ for any $B \subset X$, by induction on the cardinality of B. Pick $x \in B$ and set $B' = B \setminus \{x\}$. Then $f_A(B) = f_{A \cup \{x\}}(B') + f_A(\{x\})$. By our induction hypothesis, $f_{A \cup \{x\}}(B')$ is monotonically non-increasing with respect to A. So $f_A(B)$ is a sum of two monotonically non-increasing functions, and the result follows. □

Definition 1 Let $f : 2^X \to \mathbb{R}_{\geq 0}$ be a cost-function that satisfies our axioms. The *search value* $\mathcal{V}$ is equal to the value of the search game with cost function f.

3 The game's strategies

We describe the strategies of the players. Hider selects an element $x \in X$. Searcher selects an increasing chain $\emptyset = A_0 \subset A_1 \subset \cdots \subset A_k = X$. If A_i is the first element of the chain that contains x, then $f(A_i)$ is the payoff. It is Hider's reward and Searcher's cost.

We will allow the cardinality of elements of the chain to increase by more than one: if the search is carried out by a team, different locations may be searched at the same time. To describe such chains, we use lattices. A family of sets $\mathcal{L}$ is a lattice if it is closed under union and intersection. $\mathcal{L}$ can be illustrated by a *Hasse diagram*, which is a directed graph with vertex set $\mathcal{L}$ and edges between $A, B \in \mathcal{L}$ if $A \subset B$ and there is no $C \in \mathcal{L}$ such that $A \subset C \subset B$. In other words, B covers A. The Hasse diagram is directed upward from the root $\emptyset$ to the top X. A chain $\emptyset = A_0 \subset A_1 \subset \cdots \subset A_n = X$ corresponds to a path in the Hasse diagram that starts at the root and ends at the top. So a pure search strategy corresponds to a walk on the Hasse diagram from the bottom $\emptyset$ to the top X. If Hider chooses $x \in X$, then Searcher's payoff is $f(A_i)$ for the minimal i such that $x \in A_i$. We add weights to the edges, to compute this payoff. If B covers A then the edge AB has weight $f_A(B)$, i.e., the marginal search cost of B if A has been searched. The payoff is equal to the sum of all the weights in the path up until the first vertex that contains Hider's location.

Each lattice of sets $\mathcal{L}$ can be represented by 2^Y, where Y is the set of atoms of $\mathcal{L}$. So we may restrict our attention to chains that increase one by one. However, allowing chains to increase by more than one element makes it easier to define the search game, as we will illustrate in the third example in the next section. If $\mathcal{L}$ is equal to 2^X, then a pure search strategy corresponds to a permutation π and $\pi(1), \pi(2), \ldots, \pi(n)$ give the order in which the locations are searched. We will consider search games on the lattice 2^X only, unless explicitly stated otherwise. If we want to emphasize that the chain increases one by one, then we say that the search is sequential.

4 Some sample games

4.1 Sequentially searching three locations for the price of two

The cost function is defined on subsets of $\{1, 2, 3\}$. We write $f(1) = x$, $f(2) = y$, $f(3) = z$ and we assume that $x \leq y \leq z$. We define the function to be additive on doubletons, but $f\{1, 2, 3\})$ is the sum of the two largest costs. Hence the name 'searching three for the price of two'. We leave it to the reader to verify that f satisfies conditions $(i), (ii), (iii)$.

Fig. 1 The weighted Hasse diagram for the 3 for 2 search game. The edge-weights correspond to the marginal search costs
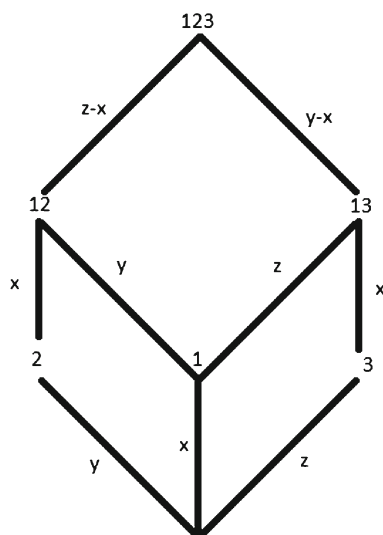
Searcher has six pure strategies. If Searcher chooses a permutation with $\pi(3) = 1$, then he pays the maximal cost $y + z$ unless he finds the Hider immediately. So switching $\pi(2)$ and $\pi(3)$ does not harm Searcher if $\pi(3) = 1$. Therefore, the pure strategies in which $\pi(3) = 1$ are dominated. We may assume that location 1 is never the last to be searched, reducing the strategy space to four permutations. The game can thus be represented by the Hasse diagram in Fig. 1.

In the degenerate case that $x = y = z$, the unique optimal mixed strategy for Hider is to hide equiprobably in one of the three locations. It is optimal for Searcher to select one of the three vertices 1, 2, 3 equiprobably, but the choice of the second vertex is unimportant, since the cost of searching two is the same as the cost of searching three in this degenerate case. It follows that the optimal search strategy is not unique.

The more general game of searching n for the price of k seems to be very difficult to solve. A further analysis of a related search game is contained in Fokkink et al. (2015).

4.2 Expanding search on a tree

The expanding search game on edge-weighted networks was introduced by Alpern and Lidbetter in Alpern and Lidbetter (2013). The game on a weighted tree was completely solved in that paper. In this game, the hiding locations X are the leaves of the tree. For $A \subset X$ the search cost $f(A)$ is defined as the total weight of the edges in the minimal subtree containing A and the root of the tree. We denote this minimal subtree by $T(A)$. The tree in Fig. 2 illustrates this game. For instance, searching the subset $\{1, 3\}$ costs $a + b + d + e$, and searching $\{1, 2\}$ costs $a + b + c$.

To show that this game fits into our framework, we need to prove that f is submodular. Since $T(A) \cup T(B)$ is a tree that contains $A \cup B$, we conclude that $T(A \cup B) \subseteq T(A) \cup T(B)$. Similarly, since $T(A) \cap T(B)$ is a tree, it follows that $T(A \cap B) \subseteq T(A) \cap T(B)$. Each edge weight in $T(A \cap B)$ occurs twice in the sum $f(A) + f(B)$, and therefore $f(A) + f(B) - f(A \cap B)$ sums all the edge weights in $T(A) \cup T(B)$ at least once. Therefore $f(A) + f(B) - f(A \cap B) \geq f(A \cup B)$, i.e., the weight function is submodular.

Fig. 2 Expanding search on a tree
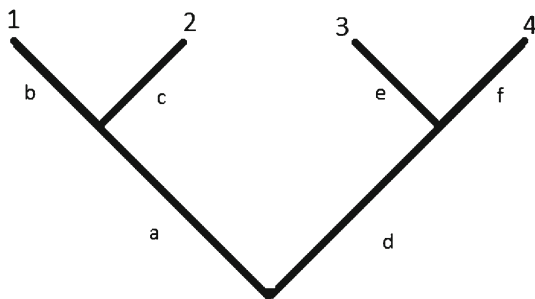
Fig. 3 Expanding search on a tree with leaves $\{1, x, 4\}$
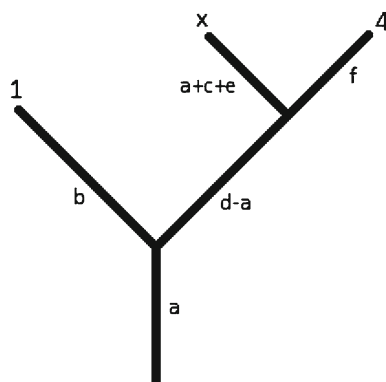
4.3 Restricted expanding search

We now give an example of a game with a lattice other that 2^X. If we modify the previous game by forcing Searcher to inspect locations 2 and 3 simultaneously, then the lattice $\mathcal{L}$ is generated by the atoms $\{1\}, \{2, 3\}, \{4\}$. If we denote the two locations 2, 3 by a single element x, then we get a sequential search game with locations $\{1, x, 4\}$. This game can still be described by expanding search on a tree. For instance, if $a \leq d$ then the cost function corresponds to the expanding search of the tree in Fig. 3.

4.4 Multiple objects

We conclude with another example of a search game in which the lattice is not the full lattice. Suppose Hider can select more than one location, so he can hide multiple objects as in Lidbetter (2013). For instance, suppose that there are two hidden objects. Then we have a search game on the product lattice $\{A \times A : A \subset X\} \subset 2^{X \times X}$. The game ends as soon as A contains both hidden objects.

5 Directed random walk and bounds on the value

In this section we derive some properties of the optimal strategies, from which we derive our bounds on the search value $\mathcal{V}$. A mixed Hider strategy is a probability distribution $\mathcal{P}$ on the set of locations X.

Lemma 2 *Let $\mathcal{P}$ be an optimal Hider strategy. Suppose that the cost function f is strictly increasing. Then all strategies are active, i.e., $\mathcal{P}(x) > 0$ for all $x \in X$.*

Note that by submodularity f is strictly increasing if and only if $f(X) > f(A)$ for all proper subsets $A \subset X$.

Proof We argue by contradiction and we suppose $\mathcal{P}(N) = 0$ for some $N \subset X$. We need to show that $\mathcal{P}$ is not optimal. Let $\emptyset = A_0 \subset A_1 \subset \cdots \subset A_n = X$ be a pure search strategy that is a best response against $\mathcal{P}$. The cost of this strategy is equal to $\sum_{x \in X} \mathcal{P}(x) f(A_x)$, where A_x denotes the first element of the chain which contains x. We modify the chain so that N is searched last. Define $A_i' = A_i \setminus N$ and adjust the indices such that $\emptyset = A_0' \subset A_1' \subset \cdots \subset A_k' = X \setminus N$ increases one by one. Extend the chain by adding the elements of N in an arbitrary order. Then $A_x' \subset A_x$ for all $x \notin N$ and we find the inequality

$$\sum_{x \in X} \mathcal{P}(x) f(A_x') = \sum_{x \in X \setminus N} \mathcal{P}(x) f(A_x') \leq \sum_{x \in X} \mathcal{P}(x) f(A_x).$$

Since the original search chain is a best response, the modified search chain cannot improve, and this must in fact be an equality. We conclude that $f(A_x') = f(A_x)$ for all $x \notin N$. But then $A_x' = A_x$ by our assumption that f is increasing. We conclude that in any best response, N is searched last.

An optimal Searcher strategy must mix between pure strategies that are best responses to $\mathcal{P}$. So if $\mathcal{P}$ is optimal, then N is searched last. Suppose $w \notin N$ and $z \in N$ and we select a Hider strategy $\mathcal{P}'$ that is equal to $\mathcal{P}$ for all elements, except $\mathcal{P}'(z) = \mathcal{P}(w)$ and $\mathcal{P}'(w) = 0$. A search that starts in $X \setminus N$ will have the strict inclusion $A_w \subset A_z$. Against $\mathcal{P}'$ the search cost changes by $\mathcal{P}'(z) f(A_z) - \mathcal{P}(w) f(A_w) > 0$, which contradicts our assumption that $\mathcal{P}$ is optimal. □

The condition that $f(X) > f(A)$ is sufficient for all Hider strategies to be active, but it is not necessary. In the 3-for-the-price-of-2 game, all Hider strategies are active if $x = y = z$, but $f(\{1, 2, 3\}) = f(\{1, 2\})$. For a cost function f that is not strictly increasing, there does not seem to be a simple condition which guarantees that all strategies are active.

A pure search strategy corresponds to a path in the Hasse diagram, or equivalently, a permutation of the locations. The strategy space thus has cardinality $n!$. We prove below that we may limit mixed Searcher strategies to random walks on the Hasse diagram. Since the Hasse diagram is a graph with $n2^{n-1}$ edges, this presents a modest reduction on the number of Searcher strategies.

A mixed search strategy is a probability distribution on paths across the Hasse diagram. For each edge AB in the Hasse diagram, let p_{AB} be the probability that the Searcher's path contains AB in this mixed strategy Σ. For a fixed A, it is clear that

$$\sum_{B \text{ covers } A} p_{AB} = \sum_{A \text{ covers } C} p_{CA}$$

since both sums represent the probability that Searcher visits A. In other words, if we take the p_{AB} to represent the flow through the edge AB, then the equation says that in-flow equals out-flow. Thus, a mixed strategy induces a flow on the Hasse diagram. Since the out-flow from the root and the in-flow to the top are equal to 1, it is a flow of unit size.

A flow of unit size corresponds to a directed random walk on the Hasse diagram: if Searcher reaches A, then the probability that she moves to B is proportional to p_{AB}. It is perhaps more convenient to think of a flow as a random walk, since a random walk corresponds to a probability distribution on the paths of the Hasse diagram. So, in turn, a random walk

corresponds to a mixed search strategy. Thus we find that a mixed search strategy induces a random walk on the Hasse diagram, which is a special mixed search strategy. We now show that the random walk produces the same payoffs as the original mixed strategy.

Theorem 1 *For every mixed search strategy there exists a directed random walk on the Hasse diagram that produces the same payoffs.*

Proof It suffices to show that the random walk produces the same payoff against each pure Hider strategy. Suppose the Hider location is $x \in X$. Let $\mathcal{S}$ be the family of all subsets that do not contain x and let $\mathcal{T}$ be the family of subsets that contain x. Then $(\mathcal{S}, \mathcal{T})$ forms a cut in the Hasse diagram. The edges AB that connect $\mathcal{S}$ to $\mathcal{T}$ all add x to A. Let $\mathcal{E}_x$ be the set of all these edges. We claim that the payoff of the mixed search strategy against x is equal to

$$\sum_{AB \in \mathcal{E}_x} p_{AB} f(B).$$

To see why this equation holds, observe that a pure search strategy is a single path. It gives payoff $f(B)$ where B is the first subset of vertices in the path which contains x. In other words, it gives payoff $f(B)$ for the unique edge in $\mathcal{E}_x$ that is in the path. A mixed search strategy is a weighted sum of paths. Each path crosses $\mathcal{E}_x$ at a unique element. The probability that the path crosses AB is equal to p_{AB}, thus we obtain our equation. If we replace the mixed strategy by a random walk, then we do not alter the probabilities p_{AB}, so the random walk produces the same payoff. □

So we have reduced the optimization of the search strategy to a network flow problem. There exists an extensive literature on this topic, but there does not seem to be a ready made solution for our problem. We turn to the analysis of this problem in the next section. We conclude this section with some examples of mixed strategies.

5.1 The double tour

A Searcher pure strategy is a permutation π of the hiding locations. Let π' be the permutation in which Searcher goes through X in the reverse direction, i.e., $\pi'(j) = \pi(n + 1 - j)$. If Searcher's strategy is to select π or π' equiprobably, then we say that she performs a *double tour*. It is analogous to the double tour of a graph, which is an optimal strategy in a wide variety of network search games, see Gal (1979). If the chain $\emptyset = A_0 \subset A_1 \subset \cdots \subset A_n = X$ corresponds to π, then $\emptyset = A_n^c \subset A_{n-1}^c \subset \cdots \subset A_0^c = X$ corresponds to π'. The payoff against a pure hiding strategy $x \in X$ is $\frac{1}{2}\left(f(A_j) + f(A_{j-1}^c)\right)$ for the first index j such that $x \in A_j$. By submodularity $f(A_j) + f(A_{j-1}^c) \geq f(A_j \cup A_{j-1}^c) + f(A_j \cap A_{j-1}^c) = f(X) + f(\{x\})$. So we find that the payoff of a double tour is bounded from below by $\frac{f(X) + \max\{f(\{x\}) : x \in X\}}{2}$. This is close to the lower bound on the game in our following theorem:

Theorem 2 *The search value satisfies* $\frac{f(X)}{2} < \mathcal{V} \leq f(X)$

Proof The upper bound is obvious, since $f(X)$ is the maximum payoff. To prove the lower bound, we let $X = \{1, \ldots, n\}$ and let $A_k = \{1, \ldots, k\}$. Define $w_j = f(A_j) - f(A_{j-1})$ and consider the submodular function $w(A) = \sum_{j \in A} w_j$, which is in fact modular (additive). It is well known that $w(A) \leq f(A)$ for all $A \subset X$, see Schrijver (2003, page 771). So the value $\mathcal{W}$ of the game with cost function w satisfies $\mathcal{W} \leq \mathcal{V}$. It suffices to show that $\mathcal{W} \geq \frac{f(X)}{2}$. Consider the mixed strategy in which Hider is at j with probability $p(j) = \frac{w_j}{\sum_j w_j}$. Now if

Searcher goes through the locations 1 to n in some arbitrary order, then the probability of finding Hider at k is equal to $p(k)$, which is proportional to w_k, and the cost is equal to the sum of all w_j over the locations j that have been searched before k. So, regardless of the order of the search, the total cost is $\frac{\sum_{k \geq j} w_k w_j}{\sum_j w_j}$. This can be rewritten as

$$\frac{\frac{1}{2}\left((\sum_j w_j)^2 + \sum_j w_j^2\right)}{\sum_j w_j} = \frac{\frac{1}{2}\left(f(X)^2 + \sum_j w_j^2\right)}{f(X)} > \frac{f(X)}{2}.$$

$\square$

5.2 A remark on the lower bound

We may try to improve the lower bound by optimizing the chain A_i. If $\delta = \min\{f(A_i) - f(A_{i-1})\}$ then $\sum_j w_j^2 \geq \delta \sum_j w_j = \delta f(X)$. The lower bound on $\mathcal{V}$ can thus be improved to $\frac{f(X)+\delta}{2}$. The minimal increment of the chain δ does not exceed $\min\{f(x): x \in X\}$, so this is a small improvement only. However, it is the best possible lower bound: the value of searching 3 for the price of 2 is $\frac{3x}{2}$ if all costs are equal so that $x = y = z$. In this case $f(X) = 2x$ and $\delta = x$.

6 The core of the game

The reader familiar with coalitional games will notice that an optimal Hider strategy $\mathcal{P}$ has properties that are similar to the Shapley value. $\mathcal{P}$ satisfies the following four properties:

(1) $\sum_i \mathcal{P}(i) = 1$
(2) If i and j are such that $f(A \cup \{i\}) = f(A \cup \{j\})$ for every $A \subset X$ not containing i and j then $\mathcal{P}(i) = \mathcal{P}(j)$.
(3) If i is such that $f(A) = f(A \cup i)$ for every $A \subset X$ that does not contain i, then $\mathcal{P}(i) = 0$.
(4) $\mathcal{P}$ is invariant up to a translation of the payoff function, i.e., for any constant c, the strategy remains optimal if f is replaced by $f - c$.

These properties correspond to the four axioms defining the Shapley value Shapley (1953). Our first three axioms are equivalent to Shapley's axioms regarding efficiency, symmetry, and the null player respectively, but the fourth axiom is weaker than the linearity axiom. Note that in property (4) we violate our condition that $f(\emptyset) = 0$, but the empty set is irrelevant since the game continues until the Hider is caught. If $f - c$ remains submodular, which may happen if c is sufficiently small, then we can replace it by the cost function g that is equal to $f - c$ for all non-empty $A \subset X$ and $g(\emptyset) = 0$. This is called a Dilworth truncation of $f - c$, see Lovász (1983). More specifically, the strategy $\mathcal{P}$ remains optimal if the Dilworth truncation of $f - c$ only attains a different value at the empty set.

In analogy to coalitional games we define the core of the game C_f as the polyhedron of Hider strategies that are bounded by the cost function.

Definition 2 Without loss of generality we may suppose that $f(X) = 1$. The core of the game is a polyhedron of probability vectors:

$$C_f = \left\{ (p_1, \ldots, p_n) \colon \sum_{i \in A} p_i \leq f(A) \text{ for all } A \subset X \right\}$$

In a coalitional game, the core may be empty, but the core of a search game is always non-empty, a fact that we have already used in the proof of Theorem 2. We conjecture that the core always contains the optimal Hider strategies.

Conjecture 1 *Optimal Hider strategies are contained in the core.*

If the conjecture holds, then computing the optimal Hider strategy becomes easier. In particular, the theorem below shows that it simplifies the analysis of cost functions that are direct sums Lovász (1983), i.e., functions f for which there exists a non-empty $A \subset X$ such that $f(V) = f(V \cap A) + f(V \cap A^c)$.

Theorem 3 *Suppose the conjecture holds. Suppose that $f(V) = f(V \cap A) + f(V \cap A^c)$ for $A \subset X$ and $f(X) = 1$. Let $\mathcal{V}(A)$ be the value of the search game restricted to A, and similarly let $\mathcal{V}(A^c)$ be the value of the game on A^c. Then*

$$\mathcal{V} = f(A)\mathcal{V}(A) + f(A^c)\mathcal{V}(A^c) + f(A)f(A^c).$$

Before we turn to the proof of this theorem, let us remark why it is useful. Submodular functions can be reduced by repeated Dilworth truncation and decomposition into direct sums, until they become irreducible Cunningham (1983). Optimal strategies are invariant under truncation, while direct sums are easy to handle by this theorem (if our conjecture is true). To find an efficient algorithm to compute the search value of a set, we thus need to focus on games with an irreducible cost function, such as our sample game of searching n for the price of k.

Proof Searcher plays as follows. Either she first exhaustively searches A before searching A^c, or she does the exact opposite and exhaustively searches A^c first. Of course, Searcher adopts the mixed optimal strategies on A and A^c. If Hider hides in A, then the payoff is $\mathcal{V}(A)$ if Searcher selects A first, and $f(A^c) + \mathcal{V}(A)$ if she selects A^c first. If Hider hides in A^c the payoff is similar, switching A and A^c. We represent this using the following matrix game

	Hide in A	Hide in A^c
Search A first	$\mathcal{V}(A)$	$\mathcal{V}(A^c) + f(A)$
Search A^c first	$\mathcal{V}(A) + f(A^c)$	$\mathcal{V}(A^c)$

From this matrix one computes that Searcher exhaustively searches A first with probability $\mathcal{V}(A) + f(A^c) - \mathcal{V}(A^c)$, or exhaustively searches the complement with probability $f(A) - \mathcal{V}(A) + \mathcal{V}(A^c)$. The value of the game, in which Searcher first performs an exhaustive search of A or A^c, is equal to $f(A)\mathcal{V}(A) + f(A^c)\mathcal{V}(A^c) + f(A)f(A^c)$.

We need to show that Searcher cannot improve this strategy. Consider the Hider strategy $\mathcal{P}$ in which he hides in A with probability $f(A)$ and in A^c with probability $f(A^c)$. Of course, his strategy of hiding in these sets is optimal. Let $\mathcal{P}_A$ be an optimal Hider strategy for the game restricted to A. By our conjecture, the optimal Hider strategy $\mathcal{P}_A$ on A satisfies $\mathcal{P}_A(V) \leq f(V)/f(A)$. This implies that the Hider strategy $\mathcal{P}$ is in the core of the game on X. If Searcher first performs an exhaustive search in A or A^c against $\mathcal{P}$, then the expected search cost is $f(A)\mathcal{V}(A) + f(A^c)\mathcal{V}(A^c) + f(A^c)f(A)$. We need to show that Searcher can do no better.

Suppose that a best response to $\mathcal{P}$ is to first search $A_1 \subset A$, then $B_1 \subset A^c$, and then $A_2 \subset A$ etc. More specifically, let $A = A_1 \cup \cdots \cup A_k$ and $A^c = B_1 \cup \cdots \cup B_k$ for disjoint subsets A_i, B_j. Searcher alternately searches A and A^c, switching from A_i to B_i and back to A_{i+1}. Without loss of generality, we may assume that k is as small as possible, and that it is optimal to search A first. If Searcher decides to switch and search B_1 before A_1, then the

cost of searching B_1 decreases by $f(A_1)$ but the cost of searching A_1 increases by $f(B_1)$. Since the search is a best response, we find that

$$\mathcal{P}(A_1)f(B_1) \geq \mathcal{P}(B_1)f(A_1).$$

In the same vein, if Searcher decides to switch and search A_2 before B_1, then the cost of searching A_2 decreases by $f(B_1)$ but the cost of searching B_1 increases by $f(A_1 \cup A_2) - f(A_1)$. Since the search is optimal and since we assume that k is minimal, we find that

$$\mathcal{P}(B_1)\left(f(A_1 \cup A_2) - f(A_1)\right) > \mathcal{P}(A_2)f(B_1).$$

Combining the two inequalities gives

$$\frac{\mathcal{P}(A_1)}{f(A_1)} > \frac{\mathcal{P}(A_2)}{f(A_1 \cup A_2) - f(A_1)}.$$

Extending the argument along the chain of A_j gives

$$\frac{\mathcal{P}(A_1)}{f(A_1)} > \cdots > \frac{\mathcal{P}(A_j)}{f(A_1 \cup \cdots \cup A_j) - f(A_1 \cup \cdots \cup A_{j-1})}$$

for $j = 2, \ldots k$. So we have a decreasing sequence of ratios

$$\frac{\mathcal{P}(A_1)}{f(A_1)}, \frac{\mathcal{P}(A_2)}{f(A_1 \cup A_2) - f(A_1)}, \ldots, \frac{\mathcal{P}(A_k)}{f(A_1 \cup \cdots \cup A_k) - f(A_1 \cup \cdots \cup A_{k-1})}.$$

The denominators of these numbers add up to $f(A)$ and the numerators add up to $\mathcal{P}(A)$. However, in this Hider strategy $\mathcal{P}(A)$ is equal to $f(A)$, so the denominators and the numerators add up to the same number. Since $\frac{a}{b} > \frac{c}{d}$ implies $\frac{a}{b} > \frac{a+c}{b+d}$, we conclude that $\frac{\mathcal{P}(A_1)}{f(A_1)} > 1$, which contradicts the hypothesis that an optimal strategy is in the core (of the game restricted to A). Therefore, k is equal to one. Searcher can do no better than an exhaustive search of A followed by A^c, or vice versa. □

7 Concluding remarks

We have developed a new framework for search games with a discrete number of locations, considering the case of an immobile Hider only. What is called for now is an effective algorithm to solve the game in the case where the cost function is irreducible.

References

Alpern, S., & Asic, M. (1985). The search value of a network. *Networks*, *15*(2), 229–238. doi:10.1002/net. 3230150208.

Alpern, S., & Gal, S. (2003). *Search games and rendezvous*. Berlin: Springer.

Alpern, S., Fokkink, R., Gasieniec, L., & Subrahmanian, V. S. (2013). *Search theory, a game theoretic approach*. Berlin: Springer.

Alpern, S., & Lidbetter, T. (2013). Mining coal or finding terrorists: The expanding search paradigm. *Operations Research*, *61*, 265–279.

Cunningham, W. (1983). Decomposition of submodular functions. *Combinatorica*, *3*, 53–68.

Fokkink, R., Ramsey, D., & Kikuta, K. (2015). A search game with a non-additive cost. *Scientiae Mathematicae Japonicae*, *28*, 17–2015.

Fujishige, S. (2005). *Submodular functions and optimization, Annals of discrete mathematics 58*. Amsterdam: Elsevier.

Gal, S. (1979). Search games with mobile and immobile hider. *SIAM Journal on Control and Optimization*, *17*, 99–122.

Lidbetter, T. (2013). Search for multiple hidden objects. *SIAM Journal on Control and Optimization*, *51*(4), 3056–3074.

Lovász, L. (1983) Submodular functions and convexity, In *Mathematical Programming - the state of the art* (pp. 235–257). Berlin: Springer.

Shapley, L. S. (1953). A value for n-person games, contributions to the theory of games. In H. W. Kuhn & A. W. Tucker (Eds.), *Annals of mathematics studies 28* (II ed., pp. 307–317). Princeton: Princeton University Press.

Shapley, L. S. (1971). Cores of convex games. *International Journal of Game Theory*, *1*, 111–130.

Schrijver, A. (2003) *Combinatorial optimization, polyhedra and efficiency*, Algorithms and Combinatorics 24. Berlin: Springer.