

Increasing Threshold Search for Best-Valued Agents

David Sarne

Department of Computer Science
Bar Ilan University
Ramat Gan, 52900, Israel
david.sarne@gmail.com

Simon Shamoun

Department of Computer Science
The Graduate Center, CUNY
New York, NY 10016-4309
srshamoun@yahoo.com

Eli Rata

Department of Computer Science
Bar Ilan University
Ramat Gan, 52900, Israel
rataeli@gmail.com

Abstract

This paper investigates search techniques for multi-agent settings in which the most suitable agent, according to given criteria, needs to be found. In particular, it considers the case where the searching agent incurs a cost for learning the value of an agent and the goal is to minimize the expected overall cost of search by iteratively increasing the extent of search. This kind of search is applicable to various domains, including auctions, first responders, and sensor networks. Using an innovative transformation of the extents-based sequence to a probability-based one, the optimal sequence is proved to consist of either a single search iteration or an infinite sequence of increasing search extents. This leads to a simplified characterization of the the optimal search sequence from which it can be derived. This method is also highly useful for legacy economic-search applications, where all agents are considered suitable candidates and the goal is to optimize the search process as a whole. The effectiveness of the method for both best-valued search and economic search is demonstrated numerically using a synthetic environment.

1. Introduction

Consider a government agency seeking the cheapest contractor to provide a service. It can issue a call for bids and select the lowest bidder. Depending on the costs associated with this process, this may be inefficient from the social welfare perspective (i.e., minimize the overall expense of all parties involved). For example, if there are costs for preparing, submitting, and processing the bids, then requesting all bids may be more costly than necessary. Instead, the agency can publish a threshold on the maximum allowable bid (e.g., a reservation price) and request that contractors submit their bids only if they are below this threshold, potentially reducing the number of bids and their associated costs. However, it must repeat this process with greater reservation prices until at least one bid is received, incurring additional costs for reissuing the call for bids. In many cases, the right sequence of reservation prices results in an overall saving in cost, thereby improving social welfare.

This paper provides a thorough analysis of the problem of deriving the optimal threshold-based search sequence in settings similar to the one above. In the general model, an

agent needs to find the agent with the best value (either the lowest or the highest, depending on the application) in its environment. The agent is acquainted with the distribution of agents' values; however, the process of learning the value of any specific agent incurs a cost (Kephart and Greenwald 2002; Choi and Liu 2000). To guarantee search completeness, all agents must be probed, but not all agents must be compelled to reply. As part of the probing, the agent can publish a threshold, compelling a response only from agents whose values are above or below it, and repeat the process until at least one agent responds. Since the process of probing all agents is also a costly act, the optimal sequence should trade off the expected decrease in the number of replying agents with the increase in the number of search rounds.

There are many multi-agent systems (MAS) that can benefit from such a search technique, here called *increasing threshold search*. For example, a volunteer ambulance corps dispatcher needs to find the closest volunteer to an emergency. She must page the volunteers and request that they call back to learn their locations. Instead of requesting that all volunteers call back, she can request that only volunteers within a certain distance of the emergency call back, and repeat the request with greater distances until at least one volunteer calls. Similarly, in a data-centric sensor network (Intanagonwiwat et al. 2003), the sink may only need the highest sensor reading. Requesting that all sensors send their data significantly depletes the sensors' power supply. Rather, the sink can request only readings above a certain threshold and iteratively decrease the threshold until at least one reply is received.

Search by iteratively increasing the search extent is a common approach to minimizing search costs, such as depth-first iterative-deepening search in artificial intelligence (Korf 1985) and expanding ring search in networking (Chang and Liu 2004). Given a cost structure for the individual search rounds, the overall cost of this type of search depends on the sequence of search extents applied. Most research in optimizing such sequences is in expanding ring search as it applies to networks (Hassan and Jha 2004; Chang and Liu 2004; Cheng and Heinzelman 2005). Although this problem has been extensively analyzed, the results are not applicable to the problem analyzed in this paper. In the expanding ring search problem, the cost of each round

Application	Reservation values	Fixed cost (α)	Variable cost ($\beta(i)$)
Finding the cheapest contractor	Bidding scale values	Cost of issuing a call for bids	Resources to prepare and evaluate each bid
Emergency services	Distance from event	Time to page volunteers	Time to answer each call
Sensor network	dB scale values	Energy to broadcast the request	Energy to transmit response

Table 1: Mapping of sample applications to the increasing threshold search model

increases as the search extent increases. In the increasing threshold search problem, the cost of each search round is fixed (the cost of probing all agents), while the last round incurs an additional variable cost (which increases as the number of agents that comply with the threshold increases).

In the next two sections, we formally introduce and analyze the optimal increasing threshold search for the class of environments described above. One important result of the analysis is that the optimal search sequence is either a single or an infinite number of reservation values characterized by a common probabilistic property, unlike in expanding ring-like domains, where the optimal strategy is finite due to the discrete nature of the problem (Chang and Liu 2004). This result facilitates the extraction of a generic distribution-independent solution, which can be then mapped to the actual sequence of thresholds for specific distributions of values using a simple transformation. The latter result is of great importance, since dynamic programming and other methodologies used in expanding ring search (Chang and Liu 2004; Cheng and Heinzelman 2005) are inapplicable in our case due to the infinite nature of the strategy.

The remaining sections address other important aspects of the problem. In Section 4, we illustrate the benefit of using the optimal strategy in a synthetic setting with different values of the problem parameters. The results are compared to those of three adaptations of commonly used expanding ring strategies to the problem considered in this paper. In Section 5, we address the usefulness of increasing threshold search in settings where the searcher is not necessarily constrained to finding the best-valued agent, but rather attempts to optimize a function that captures the cost of the overall process. This latter model is the essence of *economic search* (McMillan and Rothschild 1994; Lippman and McCall 1976). We develop the optimal sequential economic sampling strategy for our case and demonstrate, in synthetic settings, how increasing threshold search can lead to a better overall performance from the economic search point of view.

2. Model Formulation

We consider an agent searching in an environment where N other agents, applicable to its search, can be found. Each of the N agents is characterized by its value to the searcher. As in most search-related models, the values are assumed to be randomly drawn from a continuous distribution described by a PDF $f(x)$ and a CDF $F(x)$, defined over the interval $[x_{min}, x_{max}]$ (Chang and Liu 2004). The searcher agent is assumed to be ignorant of the value associated with each of the N agents, but acquainted with the overall utility distribution function, which is assumed to remain constant over time (McMillan and Rothschild 1994). The searcher is interested in finding the agent associated with the “best” value,

which, depending on the application, is either the minimum or the maximum value. For simplicity of exposition, we assume that the best-valued agent is the one associated with the minimum value.

Learning the actual value of an agent incurs some cost. In its most general form, the cost of simultaneously learning the values of i other agents is $\beta(i)$ ($\beta(i)$ is strictly increasing) (Benhabib and Bull 1983; Morgan and Manning 1985). In order to refine the population of agents whose values it plans to learn, the searcher can publish a maximum threshold r for the agents’ value, denoted a *reservation value*, requesting to communicate only with agents that comply with that threshold. If at least one agent complies with r , the search process terminates. Otherwise, the agent sets a new reservation value $r' > r$ and repeats the process. This continues until a non-empty set is found, out of which the agent associated with the minimum value is chosen. A strategy S is therefore a sequence $[r_1, \dots, r_m]$ ($x_{min} < r_i < r_{i+1} \leq x_{max}, \forall 1 \leq i < m$), where r_i denotes the reservation value to be used in the i^{th} search round.

The process of initiating a new search round and communicating the next reservation value to the agents is also associated with a fixed cost α (e.g., the cost of issuing a new call for bids or the cost of broadcasting a message). The overall cost of a search round is thus $\alpha + \beta(j)$, where j is the number of agents that comply with r_i . The expected accumulated cost of finding the best-valued agent when using strategy S is denoted $V(S)$. The searcher’s goal is therefore to derive a strategy S^* that minimizes $V(S)$.

Table 1 maps the problems described in the introduction to the general model introduced in this section.

3. Analysis

Consider a searcher agent using a strategy $S = [r_1, \dots, r_m = x_{max}]$. (In order to guarantee search completeness when using a finite sequence, the following should hold: $r_m = x_{max}$.) If the agent has to start the i^{th} search round, then there is necessarily no agent found below r_{i-1} . The *a priori* probability of such a scenario is $(1 - F(r_{i-1}))^N$. Furthermore, upon reaching the i^{th} round, the searcher agent can update its beliefs concerning the PDF of the values of the N agents, as it knows that these are necessarily in the interval $(r_{i-1}, r_{max}]$. The PDF of the agents’ values in round i , denoted $f_i(x)$ ($1 \leq i \leq m$), can thus be calculated as ($x_{min} \leq x \leq x_{max}$):

$$f_i(x) = \begin{cases} \frac{f(x)}{1-F(r_{i-1})} & x > r_{i-1} \wedge i > 1 \\ 0 & x \leq r_{i-1} \wedge i > 1 \\ f(x) & i = 1 \end{cases} \quad (1)$$

Similarly, the CDF of any of the agents’ values in round i , denoted $F_i(x)$ ($1 \leq i \leq m$), can be calculated as ($x_{min} \leq x \leq x_{max}$):

$$F_i(x) = \begin{cases} \frac{F(x) - F(r_{i-1})}{1 - F(r_{i-1})} & x > r_{i-1} \wedge i > 1 \\ 0 & x < r_{i-1} \wedge i > 1 \\ F(x) & i = 1 \end{cases} \quad (2)$$

The expected cost of the i^{th} round is thus:

$$\alpha + \sum_{j=1}^N \beta(j) \binom{N}{j} F_i(r_i)^j (1 - F_i(r_i))^{N-j} \quad (3)$$

as it takes into account the cost of initiating the new search round and the learning costs for each possible number of agent values j ($1 \leq j \leq N$). The expected cost of using strategy S is the sum of the expected cost of each of the m search rounds weighted by the probability of reaching that round:

$$V(S) = \sum_{i=1}^m \left(\alpha + \sum_{j=1}^N \beta(j) \binom{N}{j} F_i(r_i)^j (1 - F_i(r_i))^{N-j} \right) (1 - F(r_{i-1}))^N \quad (4)$$

where $r_0 \equiv x_{\min}$. The probability of starting the i^{th} search round can alternatively be formulated as the product of the probability that no agent was found in each of the $i - 1$ previous rounds, expressed as $\prod_{j=1}^{i-1} (1 - F_j(r_j))^N$. Therefore, (4) transforms into:

$$V(S) = \sum_{i=1}^m \left[\alpha + \sum_{j=1}^N \beta(j) \binom{N}{j} F_i(r_i)^j (1 - F_i(r_i))^{N-j} \right] \prod_{j=1}^{i-1} (1 - F_j(r_j))^N \quad (5)$$

For the specific case in which the reservation values are chosen from a finite set $\{x_1, x_2, \dots, x_m\}$, the optimal strategy can be derived with the following dynamic programming formulation:

$$\begin{aligned} C(x_m) &= 0 \\ C(x_i) &= \min_{i+1 \leq l \leq m} \left\{ \alpha + C(x_l) \left(\frac{1 - F(x_l)}{1 - F(x_i)} \right)^N \right. \\ &\quad \left. + \sum_{j=1}^N \beta(j) \binom{N}{j} \left(\frac{F(x_l) - F(x_i)}{1 - F(x_i)} \right)^j \left(\frac{1 - F(x_l)}{1 - F(x_i)} \right)^{N-j} \right\} \end{aligned} \quad (6)$$

where $C(x_i)$ is the cost of continuing the search if a search up to value x_i failed to yield any bids.

For the general case in which the interval $[x_{\min}, x_{\max}]$ is continuous and the process is not constrained by a finite number of rounds, the optimal search strategy must be derived with different methodology since, as we prove in Theorem 1, the optimal search sequence is either a single search round in which the value of all agents is learned or an infinite sequence of reservation values.

Theorem 1. *The optimal sequence of reservation values is either $[r_1 = x_{\max}]$ or the infinite sequence $[r_1, r_2, \dots]$, $x_{\min} < r_i < x_{\max}, \forall i > 0$, where $F_i(r_i) = F_j(r_j) = P$, for some P and $\forall i, j > 0$.*

Proof. Assume the finite sequence $S_1 = [r_1, \dots, r_m]$ is the optimal strategy. We use $S_2 = [r_2, \dots, r_m]$ to denote the optimal strategy to be used if no agent is found in the first search round and denote its expected cost from that point on by $V(S_2)$. Using S_2 , we construct strategy $S'_1 = [r'_1, \dots, r'_m]$ to be applied from the first round, where $F_1(r'_i) = F_2(r_i) \forall 1 < i \leq m$. The new strategy S'_1 has

an expected cost $V(S'_1)$, which equals $V(S_2)$ according to (5). Since S_1 is the optimal strategy, $V(S_1) \leq V(S'_1)$. Now consider a new strategy $S'_2 = [r''_1, \dots, r''_m]$ to be applied from the second round on, where $F_2(r''_i) = F_1(r_i) \forall 1 \leq i \leq m$. We denote the expected cost of S'_2 from that point on by $V(S'_2)$. According to (5), we obtain $V(S'_2) = V(S_1)$. Since S_2 is the optimal strategy from the second round, then $V(S_2) \leq V(S'_2)$, resulting in $V(S_1) \leq V(S'_1) = V(S_2) \leq V(S'_2) = V(S_1)$, which can hold only if $V(S_2) = V(S_1)$. The same logic can be applied to any search round $j \leq m$, resulting in $V(S_j) = V(S_1)$. In particular, $V(S_1) = V(S_m)$. However, the cost on-wards when reaching the last (m^{th}) round, $V(S_m)$, equals $\alpha + \beta(N)$ since $r_m = x_{\max}$; thus, we necessarily find all agents. Therefore, we obtain $V(S_1) = \alpha + \beta(N)$, which is equivalent to $S_1 = [r_1 = x_{\max}]$. Any optimal strategy that has an expected cost less than $\alpha + \beta(N)$ must consist of an infinite sequence and satisfy $F_i(r_i) = F_{i+1}(r_{i+1}) = P$, $\forall i > 0$ and some P . \square

The immediate implication of Theorem 1 is that the optimal search strategy can be expressed as a single value $0 < P \leq 1$, denoted the *reservation probability*. This is the key to deriving the optimal sequence of reservation values. As outlined below, the searcher only needs to derive P and then set the reservation value r_i in each round i such that $F_i(r_i) = P$. Additionally, the derivation of P is distribution independent, such that the same value can be used for problem instances that only differ by their CDF.

First we show how to derive P . Since the optimal sequence is infinite and the expected cost from each round on-wards is stationary, the expected cost of using P is:

$$V(P) = \alpha + \sum_{j=1}^N \left(\beta(j) \binom{N}{j} P^j (1 - P)^{N-j} \right) + (1 - P)^N V(P) \quad (7)$$

Consequently:

$$V(P) = \frac{\alpha + \sum_{j=1}^N (\beta(j) \binom{N}{j}) P^j (1 - P)^{N-j}}{1 - (1 - P)^N} \quad (8)$$

The value $P = P^*$ that minimizes $V(P)$ according to (8) is the optimal reservation probability. This can be solved using numerical approximation. Then, based on (2), the corresponding reservation value to be used in each round can be calculated by solving for r_i in the equation $P = \frac{F(r_i) - F(r_{i-1})}{1 - F(r_{i-1})}$, i.e.,

$$r_i = F^{-1}(P(1 - F(r_{i-1})) + F(r_{i-1})) \quad (9)$$

We consider the special case where the cost of learning the values of j agents is linear in j , i.e., $\beta(j) = cj$. This setting is highly applicable as, in many cases, agents are evaluated individually and independent of one another. Substituting $\beta(j) = cj$, the expression $\sum_{j=1}^N (j \binom{N}{j}) P^j (1 - P)^{N-j}$ in the numerator of (8) is the mean of a binomially distributed random variable, which equals NP . Therefore,

$$V(P) = \frac{\alpha + cNP}{1 - (1 - P)^N} \quad (10)$$

This result enables the proof of Proposition 1, which highlights the nature of the trade-off by which P^* is set.

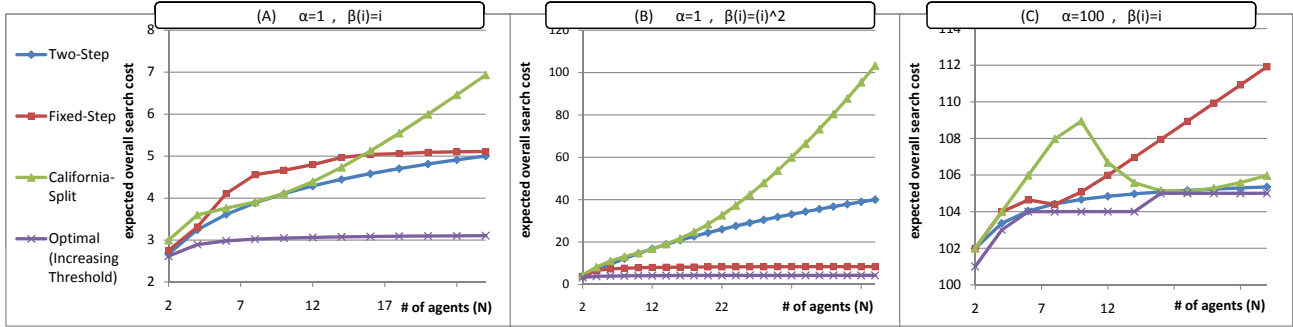


Figure 1: Performance of the proposed method compared to expanding ring search

Proposition 1. When $\beta(j)$ is linear in j , the reservation probability that minimizes $V(P)$, $P = P^*$, satisfies $c = (1 - P^*)^{N-1}V(P^*)$.

Proof. Differentiating (10) with respect to P and setting it to zero obtains:

$$\frac{cN(1 - (1 - P)^N) - N(1 - P)^{N-1}(\alpha + cNP)}{(1 - (1 - P)^N)^2} = 0 \quad (11)$$

Notice that $V(P)(1 - (1 - P)^N) = \alpha + cNP$ according to (10). Substituting the latter expression in (11), we observe that the value P^* which satisfies the equation is given by $cN(1 - (1 - P)^N) - N(1 - P)^{N-1}V(P)(1 - (1 - P)^N) = 0$, which turns into $c = (1 - P)^{N-1}V(P)$. \square

The explanation of Proposition 1 requires understanding the trade-off associated with any increase in P . By increasing P , we increase the chance of finding each of the agents. Each agent found due to the increased chance will incur a cost c . On the benefit side, if the agent found due to the increase is the only agent found during that round, then the increase has actually saved us the expected cost associated with continuing the search $V(P)$. The probability that the latter case holds is $(1 - P)^{N-1}$ (i.e., when all other agents are characterized with a value above $F_i(P)$). Otherwise, the search just ends. Since the incurred cost c is fixed, and the expected benefit $(1 - P)^{N-1}V(P)$ decreases as P increases, the optimal P value is the one for which $(1 - P)^{N-1}V(P) = c$, i.e., when the additional benefit due to the potential saving is offset by the cost incurred by finding that agent.

4. Comparative Illustration

In this section, we study the benefit of using the optimal strategy under various settings. We show how the expected cost changes with N for different combinations of α and β . Note that $F(x)$ does not affect the expected cost, as evident in (8). We also compare the expected cost of the optimal strategy to that of other strategies to understand the importance of choosing the right strategy. Since this problem has not been well addressed in the literature, we adapt to our problem three strategies for expanding ring search that are well studied in networking literature (Cheng and Heinzelman 2005; Hassan and Jha 2004;

Chang and Liu 2004). One reason for choosing expanding ring-based strategies is that, when confronted with a new problem, one might naturally turn to a related problem for solutions. In the following paragraphs, we describe the three strategies and then compare their performance in our context.

Two-Step Rule

The optimal two-step expanding ring strategy $S = [r_1, r_2 = x_{max}]$ in networking was analyzed under specific assumptions about the network structure (Cheng and Heinzelman 2005). The optimal strategy can be derived with (4).

Fixed-Step Rule

A common design of a multi-round expanding ring search strategy is to use a fixed increment between search extents. In the networking literature, optimal fixed-step strategies for expanding ring search are usually derived empirically (Hassan and Jha 2004). For our purposes, we algorithmically derive the optimal m -round strategy $S = [r_1, \dots, r_m]$, in which $r_i = x_{min} + \frac{i x_{max} - x_{min}}{m}$, $\forall 1 \leq i \leq m$ (i.e., the one which minimizes (4)).

California Split Rule

According to the California Split rule (Baryshnikov et al. 2004), the search extent is doubled each round. This strategy has the best worst-case cost in an expanding ring search. A better solution chooses values randomly from the interval $(\sqrt{2} + 1)^{i-1}, (\sqrt{2} + 1)^i$ in each round i (Chang and Liu 2004). We adapt this method to our problem such that $r_i = x_{min} + r(\sqrt{2} + 1)^{i-1}$ and $r_m = x_{max}$, where r is an arbitrary value. Here we use the r value that minimizes the overall expected cost.

Results

The expected costs of the strategies were calculated under various synthetic settings. Figure 1 depicts the performance (measured as the expected cost of search) of the three expanding ring-based methods and the increasing threshold search as a function of the number of agents N in the environment. The performance is evaluated in three settings that differ in their search costs. The distribution of values used

in all three settings is Gaussian, with $\mu = 50$ and $\sigma = 12.5$, normalized over the interval (0,100). The costs of search used in the different settings are: (A) $\alpha = 1$, $\beta(i) = i$; (B) $\alpha = 1$, $\beta(i) = (i)^2$; and (C) $\alpha = 100$, $\beta(i) = i$. As expected, according to Theorem 1, the performance of the increasing threshold search generally dominates the methods inspired by expanding ring search. Also, as expected, none of the expanding ring-based methods generally dominates any of the other methods. Similarly, the differences in performance between the increasing threshold search technique and the expanding ring search techniques are setting dependent. In some settings, an expanding ring-based method can result in performance close to the one achieved with increasing threshold search (e.g., two-step technique in setting (C)) while it can perform significantly worse in others (e.g., in settings (B) and (A)). Finally, we observe that the number of agents in the environment is a significant factor affecting performance of all methods. Nevertheless, the effect of this factor on the performance of the increasing threshold search is significantly smaller than on the performance of the expanding ring search methods.

5. Implications for Economic Search

Economic search is a widely studied model in the search literature. Unlike in the “best-valued” agent search, in economic search, the searcher is not constrained to finding the best agent. Instead, it attempts to minimize the expected overall cost of the process, defined as the weighted sum of the costs incurred due to the search itself and the lowest value found (assuming costs are additive and can be mapped to the agents’ value scale) ((McMillan and Rothschild 1994; Lippman and McCall 1976), and references therein). For example, consider a buyer agent that is interested in purchasing a product, and that the process of communicating with seller agents to learn their posted prices incurs a cost. Here, the buyer agent can purchase the product from any of the seller agents. Thus, the optimal search strategy derives from the trade-off between the marginal saving of each additional price obtained and the cost of obtaining it.

While economic-search models are inapplicable to our problem, as they do not guarantee finding the best-valued agent, the analysis supplied in this paper is an important contribution to economic search theory. This is because, as we illustrate in the following paragraphs, increasing threshold search, whenever applicable, can result in an overall reduced cost even in comparison to the optimal economic-search strategy. This result is interesting, since increasing threshold search is by definition constrained by the need to find the best-valued agent.

The general economic search model is a multi-stage search in which the number of agents sampled may vary in each stage (Benhabib and Bull 1983; Morgan and Manning 1985). It puts a constraint on the decision horizon, i.e., the number of search rounds allowed, but not on the number of agents that may be evaluated. Therefore, in order to demonstrate the usefulness of increasing threshold search within the economic search context, we first develop the optimal economic search strategy for the case where there is a limited number of agents available to be explored.

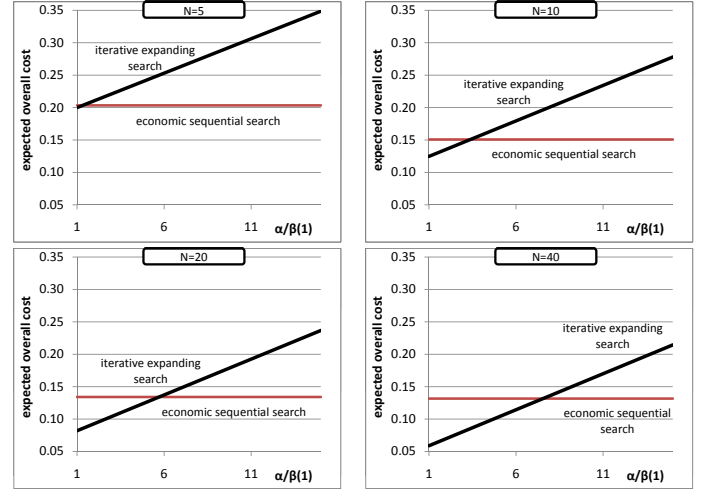


Figure 2: Expected overall cost of economic and increasing threshold search

In the absence of any other information related to the other agents, the optimal economic search strategy is to randomly sample a varying number of other agents in each search round and to decide whether to resume the search based on the best value found so far (Benhabib and Bull 1983; Morgan and Manning 1985). The cost α can be discarded in this case, as the searcher does not need to communicate (publish) its reservation value to the other agents in the environment. The economic search strategy, $S : (x, k) \rightarrow n$, ($0 < n \leq k$), is the mapping from the pair (x, k) to the number of other agents that should be sampled in the next search round, where x is the lowest value among the values of the agents sampled so far and k is the number of remaining non-sampled agents ($k \leq N$). Notice that $S(x, k) = 0$ if the agent decides to terminate its search, i.e., when satisfied with x .

Given the pair (x, k) , we denote the optimal search strategy by $S^*(x, k)$ and the expected overall cost onwards when using $S^*(x, k)$ by $V^*(x, k)$. $V^*(x, k)$ can be calculated using the following recursive equation:

$$V^*(x, k) = \beta(S^*(x, k)) + \int_{y=x_{min}}^x V^*(y, k - S^*(x, k)) f^{S^*(x, k)}(y) dy + (1 - F^{S^*(x, k)}(x))V(x, k - S^*(x, k)) \quad (12)$$

where $f^j(y)$ and $F^j(y)$ are the PDF and CDF of the expected minimum of a sample of size j , respectively.

Notice that in the case where the cost of learning the value of j agents is linear in j , the searcher does not benefit from learning the value of more than one agent in any round. The optimal economic search strategy is thus to learn the value of a single agent at a time and terminate the search once the best value found so far is below a pre-defined reservation value. The optimal reservation value can be derived using the solution to Pandora’s problem (Weitzman 1979). According to this strategy, a reservation value r is calculated for each agent based on the distribution of its value and the cost of learning its value, c . Since all agents share the same distribution of values, they are all assigned the same reservation value, which can be extracted from:

$$r = c + \int_{y=x_{min}}^r y f(y) dy + r \int_{y=r}^{x_{max}} f(y) dy \quad (13)$$

Substituting $\int_{y=r}^{x_{max}} f(y) = 1 - F(r)$ and using integration by parts, we obtain:

$$c = \int_{y=x_{min}}^r F(y)dy \quad (14)$$

For comparative illustration purposes, we use this latter economic search model in a synthetic environment where $\beta(i) = 0.01 * i$ and agent values are uniformly distributed over the interval (0,1). Substituting $F(x) = x$ in (14), we determine that the optimal reservation value to be used in this case is $r = \sqrt{2c}$. For $c = 0.01$, the optimal reservation value is $r = 0.14$. Figure 2 depicts the performance (measured as overall cost) of the optimal economic search and increasing threshold search in our synthetic environment for different values of the ratio $\alpha/\beta(1)$ and the number of agents, N . As observed from Figure 2, whenever the ratio between the fixed cost of each search iteration and the cost of evaluating an agent is sufficiently low, increasing threshold search results in a lower expected overall cost than economic search. This can be intuitively explained as the less it costs to set the limit for distinguishing the agents to be evaluated (in comparison to the cost of evaluating each agent), finer grained increments are used, and, consequently, the better is the result when using our method. Additionally, we observe that the greater the number of agents, the larger the minimal ratio that needs to hold in order for our method to dominate the economic search sampling technique. Here, the intuitive explanation is that, as the number of agents in the system increases, the expected number of search iterations until the best-valued agent is found decreases when using increasing threshold search.

6. Discussion and Conclusions

As illustrated throughout the paper, increasing threshold search is applicable to a wide variety of MAS settings, yet it has not received adequate attention in the literature. Most research of techniques utilizing increasing search extents has focused on other models. As discussed in Section 4, the strategies most studied in those domains are very different from the optimal strategy to our problem, both in structure and quality of the solution obtained. By correlating the reservation values to the respective probabilities in our proof of Theorem 1, we reveal that the essence of the optimal solution is captured by a single reservation probability. Upon calculating the optimal reservation probability using the equations supplied, the searcher can easily derive the optimal sequence of search extents.

One important result is that the optimal reservation probability and corresponding expected cost of search are distribution independent. This substantially simplifies calculations. Furthermore, the solution to one problem instance can be used to derive the optimal search sequence for any other instance that only differs in its distribution of values by merely applying a simple transformation.

As evident in Section 5, increasing threshold search can also be useful in economic search settings where the searcher is not constrained to finding the best-valued agent. This has many implications in the evolving research area known also as *search theory* (McMillan and Rothschild

1994). Among the various possible extensions to existing economic search models are threshold-based enhancements to sequential search models with finite decision horizon and an integrated preliminary increasing threshold search for refining the population on which sequential search takes place.

Other possible extensions of this work include competition and cooperation models for two or more searchers when operating in settings where search is costly and one or more of them is capable of using increasing threshold search.

Acknowledgments

This work was supported in part by the Israeli Ministry of Industry and Trade under project RESCUE and ISF/BSF grants 1401/09 and 2008-404.

References

- Baryshnikov, Y. M.; Jr., E. G. C.; Jelenkovic, P. R.; Momicilovic, P.; and Rubenstein, D. 2004. Flood search under the california split rule. *Oper. Res. Lett.* 32(3):199–206.
- Benhabib, J., and Bull, C. 1983. Job search: The choice of intensity. *Journal of Political Economy* 91:747–764.
- Chang, N. B., and Liu, M. 2004. Revisiting the ttl-based controlled flooding search: optimality and randomization. In Haas, Z. J.; Das, S. R.; and Jain, R., eds., *MOBICOM*, 85–99. ACM.
- Cheng, Z., and Heinzelman, W. B. 2005. Flooding strategy for target discovery in wireless networks. *Wireless Networks* 11(5):607–618.
- Choi, S., and Liu, J. 2000. Optimal time-constrained trading strategies for autonomous agents. In *Proc. of MAMA'2000*.
- Hassan, J., and Jha, S. 2004. On the optimization trade-offs of expanding ring search. In Das, N.; Sen, A.; Das, S. K.; and Sinha, B. P., eds., *IWDC*, volume 3326 of *Lecture Notes in Computer Science*, 489–494. Springer.
- Intanagonwiwat, C.; Govindan, R.; Estrin, D.; Heidemann, J. S.; and Silva, F. 2003. Directed diffusion for wireless sensor networking. *IEEE/ACM Trans. Netw.* 11(1):2–16.
- Kephart, J., and Greenwald, A. 2002. Shopbot economics. *JAAMAS* 5(3):255–287.
- Korf, R. E. 1985. Depth-first iterative-deepening: An optimal admissible tree search. *Artif. Intell.* 27(1):97–109.
- Lippman, S., and McCall, J. 1976. The economics of job search: A survey. *Economic Inquiry* 14:155–189.
- McMillan, J., and Rothschild, M. 1994. Search. In Aumann, R., and Hart, S., eds., *Handbook of Game Theory with Economic Applications*. 905–927.
- Morgan, P., and Manning, R. 1985. Optimal search. *Econometrica* 53(4):923–944.
- Weitzman, M. L. 1979. Optimal search for the best alternative. *Econometrica* 47(3):641–54.