Iterative Expanding Search in Multi-Agent Systems

(Extended Abstract)

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 needs to be found and the goal is to minimize the expected cost of search. Given the ability to vary the extent of search, a search strategy is a sequence of search iterations of varying extent.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

General Terms

Economics, Algorithms

Keywords

Economically-motivated agents

1. INTRODUCTION

In many multi-agent systems (MAS), we find problems where an agent needs to find the agent with the lowest (or highest, depending on the application) value, while the process of learning the value of an agent incurs a cost. The agent can minimize costs by publishing a maximum threshold for the agents' value, denoted a *reservation value*, requesting to communicate only with agents that comply with that threshold. The agents repeats the search process with higher reservation values until at least one agent is found. This search technique, which we call "iterative expanding search", is applicable to many problems. For example, consider a police dispatcher that needs to find an available officer within the closest vicinity of an evolving event. The dispatcher can broadcast a request that only officers within a predefined distance to the event reply with their location. Similarly, consider a sensor network in which the sink only needs the highest sensor reading. In this case, it can reduce overall transmission overhead by broadcasting a query for sensor data above a certain threshold. We show how to derive the optimal strategy for such techniques.

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 $[x_{min}, x_{max}]$. 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. The searcher is interested in finding the agent associated with the minimum value.

In its most general form, the cost of simultaneously learning the values of i other agents is $\beta(i)$ ($\frac{d(\beta(i))}{di} > 0$). In order to refine the population of agents whose values it plans to learn, the searcher can publish a reservation value r requesting to communicate only with agents that comply with r. 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 $[r_1, \ldots, r_m]$ ($x_{min} < r_i < r_{i+1} \leq r_{max}, \forall 1 \leq i < m$), 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 broadcasting a query). The overall cost of a search round is thus $\alpha + \beta(i)$, where i 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 $V(S)$.

3. ANALYSIS

Consider a searcher agent using a strategy $S = [r_1, \ldots, r_m = x_{max}]$[1]. If the agent needs 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 $(1 - F(r_{i-1}))^N$. Alternatively, it can be expressed as $\prod_{j=1}^{i-1}(1 - F_j(r_j))^N$, which is the product of the probability that no agent was found in each of the $i-1$ previous rounds. Furthermore, upon reaching the i^{th} round, the searcher agent can update its beliefs concerning the CDF of the values of the N agents, as it knows that these are necessarily in the interval $(r_{i-1}, r_{max}]$. The CDF of any of the agents' values in round i, denoted $F_i(x)$

[1]In order to guarantee search completeness when using a finite sequence, the following should hold: $r_m = x_{max}$.

Cite as: Iterative Expanding Search in Multi-Agent Systems (Extended Abstract), David Sarne, Simon Shamoun and Eli Rata, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 1451-1452

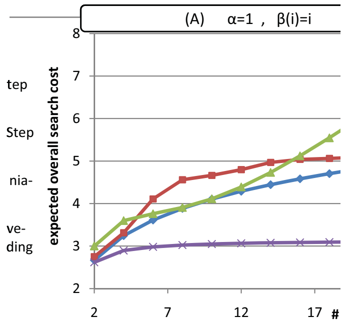

Figure 1: Comparative illustration of the proposed method to expanding ring search

$(0 < i \leq m)$, can be calculated as $(x_{min} \leq x \leq x_{max})$:

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

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 ($r_0 \equiv x_{min}$):

$$V(S) = \sum_{i=1}^{m} \left[\alpha + \sum_{j=1}^{N} \beta(j) \binom{N}{j} F_i(r_i)^j (1 - F_i(r_i))^{N-j} \right] \prod_{j=1}^{i-1} (1 - F_j(r_j))^N \quad (2)$$

For the specific case in which the reservation values are chosen from a finite set $\{x_1, x_2, ..., x_m\}$, the optimal strategy can be derived with dynamic programming. 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 a different methodology, since the optimal search sequence is either a single search round in which the values of all agents are 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* $[r_1, r_2, \ldots]$, $x_{min} < r_i < x_{max}, \forall i > 0$, *where* $F_i(r_i) = F_j(r_j) = P$, *for some* P *and* $\forall i, j > 0$.

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*. The searcher can derive the optimal sequence of reservation values using the following method: First, derive P. Since the optimal sequence is infinite and the expected cost from each round onwards is stationary, the expected cost of using P is:

$$V(P) = \frac{\alpha + \sum_{j=1}^{N} (\beta(j) \binom{N}{j} P^j (1-P)^{N-j})}{1 - (1-P)^N} \quad (3)$$

The value $P = P^*$ that minimizes $V(P)$ according to (3) is the optimal reservation probability. Note that this derivation is distribution independent. Then, based on (1), 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.,

$$r_i = F^{-1}(P(1 - F(r_{i-1})) + F(r_{i-1})) \quad (4)$$

4. THE EXPANDING RING ALTERNATIVE

A technique similar to iterative expanding search, called *expanding ring search*, is widely applied in communication networks to minimize the control overhead associated with broadcast route discovery in on-demand protocols [1, 2]. The search is initially limited to a small radius around the source and expands in rounds until the destination is found. For comparative illustration, we compare our solution to three well studied strategies for expanding ring search. 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. After describing the strategies, we present our results.

Two-Step Rule: A two-step strategy has the form $S = [r_1, r_2 = x_{max}]$. The optimum can be derived from (2).

Fixed-Step Rule: A common design of a multi-round expanding ring search strategy is to use a fixed increment between search extents [2]. For our purposes, we algorithmically derive the optimal m-round strategy $S = [r_1, \ldots, r_m]$, in which $r_i = x_{min} + \frac{ix_{max} - x_{min}}{m}, \forall_{1 \leq i \leq m}$ (i.e., the one which minimizes (2)).

California Split Rule: According to the California Split rule, the search extent is doubled each round. A better solution chooses values randomly from the interval $((\sqrt{2}+1)^{i-1}, (\sqrt{2}+1)^i]$ in each round i [1]. 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 value of r 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 iterative expanding 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). As expected, according to Theorem 1, the performance of iterative expanding search generally dominates the methods inspired by expanding ring search. In some settings, an expanding ring-based method can result in performance close to the one achieved with iterative expanding 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.

5. REFERENCES

[1] N. B. Chang and M. Liu. Revisiting the ttl-based controlled flooding search: optimality and randomization. In *MOBICOM*, pages 85–99, 2004.

[2] J. Hassan and S. Jha. On the optimization trade-offs of expanding ring search. In *IWDC*, pages 489–494, 2004.