

Enhancing MAS Cooperative Search Through Coalition Partitioning

Efrat Manisterski¹ and David Sarne² and Sarit Kraus¹

¹Computer Science Department, Bar-Ilan University, Ramat Gan, Israel
{maniste,sarit}@cs.biu.ac.il

²Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA
sarned@eecs.harvard.edu

Abstract

In this paper we present new search strategies for agents with diverse preferences searching cooperatively in complex environments with search costs. The uniqueness of our proposed mechanism is in the integration of the coalition's ability to partition itself into sub-coalitions, which continue the search autonomously, into the search strategy (a capability that was neglected in earlier cooperative search models). As we show throughout the paper, this strategy is always favorable in comparison to currently known cooperative and autonomous search techniques: it has the potential to significantly improve the searchers' performance in various environments and in any case guarantees reaching at least as good a performance as that of other known methods. Furthermore, for many common environments we manage to significantly eliminate the consequential added computational complexity associated with the partitioning option, by introducing innovative efficient algorithms for extracting the coalition's optimal search strategy. We illustrate the advantages of the proposed model over currently known cooperative and individual search techniques, using an environment based on authentic settings.

1 Introduction

Cooperative search [Sarne and Kraus, 2005] is a *coalition formation* based process aiming to improve the performance of agents engaged in costly individual searches for opportunities/resources. To illustrate this idea, consider several agents (representing different users) interested in buying a desktop over the internet. Each of the agents can have its own set of preferences for valuing any given offer. Obviously the process of finding and evaluating any potential opportunity for buying the desktop is time consuming (i.e. costly). Therefore each agent sets a strategy for itself in which it decides to what extent to continue its search. As good as an agent's own search may be, often it can gain from forming a coalition and searching cooperatively with others. In this manner a group of agents can share, reuse and re-allocate opportunities (that otherwise might have been discarded) among themselves. A

detailed review of the advantages and potential uses of cooperative search is given in [Sarne and Kraus, 2005].

A fundamental component of the cooperative search model (as in any search model) is the search cost [McMillan and Rothschild, 1994]. Search costs reflect the resources (not necessarily monetary) that need to be invested/consumed for maintaining search-related activities, such as locating opportunities, analyzing and comparing them and negotiating over them with other agents. The existence of such search costs in Multi-Agent systems (MAS) is widely recognized [Bakos, 1997; Choi and Liu, 2000; Kephart and Greenwald, 2002] and used. The importance of these costs increases as a function of the amount and richness of opportunities that can be potentially found in the MAS environment. Thus even in settings where technology can reduce the cost of obtaining a single opportunity (e.g. when buying a specific product over the internet) the cost of evaluating all possible opportunities available becomes an important parameter affecting the agents' search strategy. Alas, when considering a cooperative search one needs to take into consideration additional overhead costs associated with the process of coalition coordination. These costs can be of various forms and are usually modeled as a function of the number of the agents forming the coalition [Sarne and Kraus, 2005].

Former analysis of cooperative search [Sarne and Kraus, 2005] suggests a significant potential improvement for the agents' performance when using the new method in various environments (in comparison to an individual search [McMillan and Rothschild, 1994]). However this improvement cannot be generally guaranteed for all environments. Principally, this is because the strategies proposed for the cooperative search assume the coalition structure is fixed for the entire search process (i.e., force the coalition members to keep the coalition in its initial setting until it is decided that the coalition as a whole should terminate the search). In this paper we show that the coalition can further improve its performance by considering intelligent self-restructuring (i.e., partitioning) strategies along the cooperative search process. Furthermore, the use of our proposed cooperative search strategy guarantees that the search through a coalition outperforms individual search. In fact, as discussed throughout the paper, both the traditional cooperative and individual search techniques, denoted FSCS (Fixed-Structure Cooperative Search) and SAS (Single Agent Search), respectively, can be considered spe-

cific cases of our proposed search technique, termed VSCS (Variable Structure Cooperative Search).

The main advantage of our approach is that it takes into consideration the expected contribution of any subset of agents to the coalition's utility and to its cost of search at any search stage (instead of merely considering the gain in having the coalition maintain the search as a whole). Notice that in many scenarios the expected incremental individual utility from resuming the search (given a set of known opportunities) for the different agents in the coalition highly varies. This is particularly true when different coalition members have non-correlated preferences for the requested product. Here, the coalition may identify opportunities in which it is more beneficial to have some of its members terminate the search (taking advantage of any of the currently known opportunities) while the remaining members resume the search in a reduced coalition structure hence with a smaller coalition overhead. Alternatively, the coalition can partition itself into several sub-coalitions that search more efficiently.

The suggested partitioning process can be seen as coalition reconstruction [Ogston *et al.*, 2003], and is very similar in its nature to the coalition formation process. It is used whenever the coalition can increase its overall utility by acting as several autonomous sub-coalitions (as detailed below). As in other coalition formation models [Ohta *et al.*, 2006; Kahan and Rapoport, 1984], we assume that there is an option for side payments, thus a strategy maximizing the overall coalition utility is always preferred by all members.

The integration of the partitioning option into the coalition's search strategy is computationally difficult. With this option the coalition needs to evaluate the potential benefit from any possible partition of itself as well as additional future partitions of the sub-coalitions created throughout any future sequence of opportunities encountered. This imposes a significant computational challenge. Therefore an important contribution of this paper, alongside modeling and comprehensively analyzing the general case, is the introduction of a computational algorithmic-based means for facilitating the calculation process of such a coalition strategy in common MAS environments. The uniqueness of our algorithm relates to the fact that its computational complexity does not depend on the potential number of coalition partitions

In addition to the formal proofs for the superiority of the new method, we use a simulation to demonstrate its advantages. The simulation is based on authentic settings, fully correlated with a genuine eCommerce specific vertical market. Our computational algorithm runs within less than a second in this setting that contains thousands of opportunities. Thus alongside the theoretical importance of the results, we present a model that has the potential of actually being implemented.

In the following section we address relevant literature in the area of MAS coalition formation and search. Then we formally introduce the VSCS model, present its analysis and introduce appropriate computational means as described above. As expressed at the end of the paper, overcoming the computational complexities associated with extracting the coalition's optimal search strategy in the VSCS model paves the way for further important research.

2 Related Work

Today, search models can be found in various research areas such as eCommerce, real-estate, data mining and social studies. In its most basic form, search theory considers the problem of a searcher seeking to maximize his long term utility by fulfilling his search objective, while operating in a costly environment ([Lippman and McCall, 1976; McMillan and Rothschild, 1994], and references therein). Nevertheless, search theory has been focused mostly on a single searcher, looking for a single opportunity, either as a one sided (taking the environment's reaction to the search strategy used by the agent to be static) or two sided (as a matching model, analyzed from the equilibrium perspective) model. The problem of a group that searches cooperatively has never been addressed in search theory literature. This, is in spite of the fact that cooperative search has been proven [Sarne and Kraus, 2005] to be inherently different from a single agent's search in relation to its complexity, strategy structure and solution methodology. These differences definitely hold when considering the partitioning option, since partitioning is not a feasible alternative in a single agent search.

Research most related to cooperative search can be found in coalition formation literature. Coalition formation is a favorable process in MAS, particularly desirable in environments where group coordination results in more efficient task performance in comparison to individual performance [Lermann and Shehory, 2000]. A review of the extensive literature in this area can be found in [Kahan and Rapoport, 1984]. Three basic stages are common to all coalition formation models [Sandholm *et al.*, 1999; Tsvetovat *et al.*, 2000]: coalition structure generation, executing the coalition task and dividing the generated value among the coalition members. Nevertheless, in most of the proposed models studying coalition formation the focus is on coalition generation and payoff division. The tendency to avoid research of the task execution stage relates to the assumption often made that agents have complete knowledge concerning the opportunities found in their environment.

The formal introduction of cooperative search in MAS and the initial research of this topic can be found in [Sarne and Kraus, 2005; Manisterski *et al.*, 2006]. However, as discussed above, these papers do not allow coalition restructuring, thus the models described therein are limited in the level of performance improvement that they are able to achieve using cooperative search.

3 Model and Analysis

Similar to many other coalition formation models [Tsvetovat *et al.*, 2000; Yamamoto and Sycara, 2001], we adopt the legacy buyers-coalition application for the electronic marketplace and in particular the B2C (Business-to-Consumer) market, where sellers can supply almost any volume of demand. Notwithstanding, we do wish to emphasize that the cooperative search strategy we present in this paper is general and can be applied to any MAS where cooperative search is feasible.

We base our model description and formulation on the definitions given in [Sarne and Kraus, 2005] and extend them to support partitioning as part of the set of actions in the strat-

egy definition. We consider an electronic marketplace where numerous buyer and seller agents can be found, each interested in buying or offering to sell a well defined product. A product can be offered by many different seller agents under various terms and policies (including prices). We assume that while buyer agents are ignorant of individual seller agents' offers, they are acquainted with the overall distribution of opportunities (whereas an opportunity is defined as the option to buy the product under specific terms and policies) in the marketplace.

In the absence of central matching mechanisms or mediators, each agent needs to search for appropriate opportunities to buy its requested product. This process of a single agent search (SAS) is described in [McMillan and Rothschild, 1994]. Throughout its search the buyer agent locates seller agents sequentially (i.e. one at a time) and learns about their offers by interacting with them. Upon learning the new opportunity details, the buyer agent evaluates it using its own utility function. We consider the agents to be heterogeneous; each having its own utility function defined over product attributes, terms and policies as well as reputation and trust factors. Based on the evaluation the agent makes a decision whether to exploit any of the opportunities encountered until this point (i.e. buy from any of the sellers) or resume its search in a similar manner.

The search activity is assumed to be costly [Choi and Liu, 2000; Kephart and Greenwald, 2002] - each search stage process induces a search cost. We assume utilities and costs are commensurable and additive. Recognizing the benefits of a cooperative search, buyer agents, interested in similar products or interchangeable products, may form coalitions. Any coalition which has been formed conducts its search in a similar manner (i.e. sequentially, encountering a single seller at each search stage, thus accumulating new opportunities along time). There are various methods by which the coalition members can coordinate their cooperative search (e.g. assigning a representative agent that will search on behalf of the coalition or simply taking turns searching), each deriving a different search cost overhead structure. As in [Sarne and Kraus, 2005; Manisterski *et al.*, 2006] the search cost associated with each additional search stage increases as a function of the number of agents forming the coalition.

Once formed, the coalition needs to generate its search strategy. This strategy determines the action to be taken when reaching any potential search stage while in its current structure. The possible actions for each stage include: (a) terminating the search; (b) resuming the search while keeping the current structure; and (c) partitioning the coalition into sub-coalitions (where each member in the partitioned coalition is assigned to one of the sub-coalitions) that set their optimal search strategy from this point onward, independently (i.e. search autonomously or terminate their search). The optimal strategy is the one that maximizes its expected utility (utility obtained from opportunities exploited by the coalition members less the search costs). As in other coalition formation models [Kahan and Rapoport, 1984; Ohta *et al.*, 2006], we assume that there is an option for side payments and set the coalition's goal to maximize the overall coalition utility. Given this goal, the coalition's strategy

is not influenced by the payoff division protocol, nor by stability considerations, but rather influences these two factors [Sarne and Kraus, 2005]. Since the pre-determined payoff division mechanism used for creating the coalition defines the portion of the agent's utility from the overall coalition utility, any increase in this latter value will increase any of the agents' shares. Thus the overall utility maximization strategy is the preferred strategy by all agents at every stage of the search (i.e. no conflict of interests).

3.1 The Search Strategy

In order to formally present our search strategy we use O to denote the space of potential opportunity types the coalition may encounter. The opportunity types' distribution in the marketplace is denoted by the probability function $p(\vec{\sigma})$. We consider a coalition $A_g = \{a_1, a_2, \dots, a_{|A_g|}\}$ of a general size, where a_j is the j -th buyer agent in the coalition. Each buyer agent, a_j , evaluates opportunities using a utility function $U_j : O \rightarrow R$, where $U_j(\vec{\sigma})$ is the agent's utility from opportunity type $\vec{\sigma}$.¹ The search cost associated with having a coalition of size n (i.e. having n agents in the coalition) for each search round is denoted by the function $c(n)$.

We can reduce the number of world states in which the coalition can be, by adopting a representation of states through sets of effective known opportunities. Given a set of known opportunities θ and a coalition A_g , it is sufficient to maintain a subset s of θ to represent the current state of this coalition. Subset s stores the opportunities from θ that maximize the utility of each of the agents in A_g . Formally, we can calculate the state s of a coalition A_g acquainted with a set θ of known opportunities by using the function² $s = state(\theta, A_g) = \{\vec{\sigma}_{a_j}^s | a_j \in A_g, \vec{\sigma}_{a_j}^s \in \theta, U_j(\vec{\sigma}_{a_j}^s) \geq U_j(\vec{\sigma}), \forall \vec{\sigma} \in \theta\}$. We use S_{A_g} to denote the set of all possible states of a coalition A_g . Reaching a state s , the expected utility of a coalition A_g from this point onwards when using its optimal strategy is denoted $V^*(A_g, s)$.

We begin our analysis by developing the appropriate expected utility achieved by the different possible actions the coalition may take, leading to the appropriate equations from which the optimal strategy can be extracted.

The first possible action for the coalition is terminating the search. If the search terminates at state s then the coalition's utility $V_t(A_g, s)$ is the aggregated coalition member's utilities when each coalition member, a_j , is assigned the opportunity $\vec{\sigma}_{a_j}^s$ which maximizes its utility function, U_j , from the set of currently known opportunities in s :

$$V_t(A_g, s) = \sum_{a_j \in A_g} U_j(\vec{\sigma}_{a_j}^s) \quad (1)$$

Next we consider the scenarios in which coalition A_g resumes its search at state s while keeping its current structure. Here we can make use of the analysis methodology given in [Sarne and Kraus, 2005], dividing the opportunities

¹ $\vec{\sigma}$ is noted as a vector since it assigns a value for each product's attributes e.g. an opportunity to buy a calculator can be $\vec{\sigma} = (scientific, 20\$, smallDisplay, pocket)$

²If more than one maximizing opportunity exists, $state()$ will return a single opportunity according to a predefined order.

space into two sub-spaces, containing improving and non-improving opportunities for the coalition's utility, respectively. Hence, the expected utility the group of agents A_g can obtain if all agents in A_g resume the search, $V_r(A_g, s)$, is attained by:

$$V_r(A_g, s) = \sum_{\vec{o} \in O_{improve}^s} p(\vec{o}) V^*(A_g, s') + \sum_{\vec{o} \in O_{stay}^s} p(\vec{o}) V_r(A_g, s) - c(|A_g|) \quad (2)$$

where s' is the new state of coalition A_g after encountering opportunity \vec{o} , $s' = state(s \cup \{\vec{o}\}, A_g)$. $O_{improve}^s$ denotes the set of opportunities that changes the coalition's current state and O_{stay}^s denotes the complementary set (opportunities that do not change the coalition's current state).

Applying some basic mathematic manipulations on the above equation, we obtain:

$$V_r(A_g, s) = \frac{\sum_{\vec{o} \in O_{improve}^s} p(\vec{o}) V^*(A_g, state(s \cup \{\vec{o}\}, A_g)) - c(|A_g|)}{1 - \sum_{\vec{o} \in O_{stay}^s} p(\vec{o})} \quad (3)$$

Since $1 - \sum_{\vec{o} \in O_{stay}^s} p(\vec{o}) = \sum_{\vec{o} \in O_{improve}^s} p(\vec{o})$ we obtain:

$$V_r(A_g, s) = \frac{\sum_{\vec{o} \in O_{improve}^s} p(\vec{o}) V^*(A_g, state(s \cup \{\vec{o}\}, A_g)) - c(|A_g|)}{\sum_{\vec{o} \in O_{improve}^s} p(\vec{o})} \quad (4)$$

The third possible scenario is where coalition $A_g = \{a_1, \dots, a_n\}$ partitions into a set $P = (A_1, \dots, A_k)$ of disjoint non-empty sub-coalitions ($A_i \cap A_j = \emptyset \forall i, j \leq k$, $\bigcup_{i=1}^k A_i = A_g$) that set their search strategies independently. M_{A_g} denotes the set of all possible partitions of coalition A_g . The selected partition will be the one yielding the maximum expected utility, assuming all the sub-coalitions created use their optimal strategies. The expected utility of the partitioned coalition A_g in this case, denoted V_p , is given by:

$$V_p(A_g, s) = \max_{P \in M_{A_g}} \left\{ \sum_{A_i \in P} \max\{V_i(A_i, s'), V_r(A_i, s')\} \right\} \quad (5)$$

where $s' = state(s, A_i)$

As stated in the previous section, the optimal strategy is mapping $(A_g, s) \rightarrow \{resume, terminate, P\}$, maximizing the expected utility $V^*(A_g, s)$ which can now be formulated as: $V^*(A_g, s) = \max\{V_r(A_g, s), V_t(A_g, s), V_p(A_g, s)\}$. This can also be expressed in a more efficient manner as:

$$V^*(A_g, s) = \begin{cases} \max\{V_t(A_g, s), V_r(A_g, s)\} & \text{if } |A_g| = 1 \\ \max\{V_r(A_g, s), \max_{A_i, A_j} \{V_{A_i}^* + V_{A_j}^*\}\} & \text{otherwise} \end{cases}$$

where $V_{A_i}^* = V^*(A_i, state(s, A_i))$, $V_{A_j}^* = V^*(A_j, state(s, A_j))$, $A_i \cup A_j = A_g$, $A_i \cap A_j = \emptyset$, $A_j \neq \emptyset$, $A_i \neq \emptyset$. Notice that in the above equation we simplified the calculation to include only size-two partitions (i.e. partition into two sub-coalitions). This is because every partition $P \in M_{A_g}$ has already been taken into account recursively in one of the size-two partitions as part of the definition of $V^*(A_i, s)$ and $V^*(A_j, s)$. Furthermore, the

latter definition used for representing a partition also covers the option in which all agents terminate the search.

Notice that in the case where partitioning the coalition into sub-coalitions is a costly process, the integration of such split costs is quite straightforward. For example, if we have a cost $C_{split}(k)$ which is a function of the number of sub-coalitions into which A_g partitions, then all we need to do is subtract the value $C_{split}(|P|)$ from the right hand side of equation 5 and the rest of our analysis remains unchanged.

The proposed VSCS search strategy will always be the preferred strategy as proposition 1 states.

Proposition 1. *The VSCS search model is a generalization of both the FSCS and the SAS and weakly dominates them (i.e. guarantees a better or equal overall performance).*

The proof of the proposition is quite straightforward. FSCS and SAS are both specific cases of the VSCS where the coalition always chooses to resume the search in its original structure or partitions into a set of coalitions of size one (i.e. single agents), respectively. Therefore, if any of these two search mechanisms produce the maximum utility for a given environment, the coalition will adopt this structure. Notice that between the two methods FSCS and SAS, neither generally dominates the other (but rather the selection of the optimal one is environment-dependent). The only advantage of these two methods in comparison to the VSCS is in terms of the computational complexity of the optimal strategy. Nevertheless, in many common environments even the calculation complexity can be overcome and reduced to the one similarly obtained for the FSCS model, as we demonstrate in the next section.

3.2 Reducing Calculation Complexity

In many environments the marginal cooperative search cost (cost of adding an additional agent to a coalition) has a fixed or non-increasing structure (formally, described as: $c(n+2) - c(n+1) \leq c(n+1) - c(n) \forall n \geq 0$). This is typical due to the fact that most of the coalition overhead is associated with communication. A characteristic example of this is where one of the coalition members in each stage of the search conducts the search on behalf of the coalition. The agent conducting the search needs to send the results to the other agents. The other agents do not have to communicate among themselves, therefore the search cost is at most linear and depends on the number of coalition members. In such environments, as we prove and demonstrate in the following paragraphs, many of the calculation complexities induced by allowing coalition partitioning can be overcome. We begin by introducing the following lemma 1 which lays the foundations for our algorithmic-based solution.

Lemma 1. *Given a state s and a coalition A_g there is an optimal strategy for coalition A_g in which at most one sub-coalition resumes the search.*

The above lemma A_g suggests that any strategy in which the coalition partitions into two or more sub-coalitions that resume the search in parallel is weakly dominated by a strategy in which k ($k \leq n$) coalition members terminate their search at the current stage while the rest $n-k$ coalition members continue as a unified coalition in a cooperative search.

Though the proof below is quite detailed, it is intuitive. If the optimal strategy is to have two sub-coalitions searching in parallel then merging them into one coalition for just one additional search stage and then returning to their initial coalition structures will necessarily yield a better performance (since the expected utility will remain the same whereas the aggregated cost of such a move can only decrease). The immediate implication is that in each state an agent needs to decide whether to resume or terminate the search. Since all agents who resume the search conduct the search together, there is no need to decide which coalition to join.

Proof. Assume that according to the best strategy, coalition A_g splits to at least two sub-coalitions that resume the search separately. Consider two scenarios: (1) Each coalition separately draws the opportunity it encounters from distribution P . (2) There is a list of infinite opportunities $\vec{o}^1, \vec{o}^2, \dots$ that were taken randomly from distribution P . All coalitions in their i^{th} search encounter the same opportunity, \vec{o}^i . The probability that a coalition will encounter a given opportunity is exactly the same in both scenarios. Moreover each coalition's utility depends merely on the other opportunities it encounters and doesn't depend on the other coalition's utility. Therefore in both scenarios each coalition has the same expected utility.

Consider that after a coalition A_g splits to sub coalitions, each sub-coalition resumes the search according to the second scenario. If in the next stage of the search all sub-coalitions of A_g conduct the search together instead of conducting the search separately, they can only benefit. This is because their search cost can only decrease. However after conducting the search each sub coalition still can make the same decisions as in the second scenario. Therefore they can only benefit by conducting the search together. In similar way we can eliminate other cases where a coalition splits into sub coalitions that resume the search separately. \square

Our analysis suggests a simple mechanism for determining the agents that will cooperatively continue the search at each state for any given coalition. For this purpose we introduce several supporting definitions and notations. Given a coalition A_g and a state s we use:

- $\hat{V}(s, a_j)$ - the additional expected utility (without including the search cost) that agent $a_j \in A_g$ obtains from terminating its search after conducting one additional search stage rather than terminating the search in the current state s it is in. Thus:

$$\hat{V}(s, a_j) = \sum_{U_j(\vec{o}) \geq U_j(\vec{o}_{a_j}^s)} p(\vec{o})(U_j(\vec{o}) - U_j(\vec{o}_{a_j}^s)) \quad (6)$$

- $A_{order} = \{a'_1, \dots, a'_{|A_g|}\}$ - the list of agents in A_g sorted in a descending order according to their \hat{V} values.
- $A_r = \{a'_1, \dots, a'_k\}$, $A_t = \{a'_{k+1}, \dots, a'_n\}$ - a partition of the sorted list A_{order} , where k is the first index in A_{order} satisfying both conditions: (C1) $\sum_{j=1}^k \hat{V}(s, a'_j) \geq c(k)$; and (C2) $\nexists i, i > k$ that satisfies

$$\sum_{j=k+1}^i V_{one}(s, a'_j) \geq c(i) - c(k) \quad (7)$$

If such a k does not exist then $A_r = \emptyset$ and $A_t = A_g$.

In the above definitions the condition (C1) is used to ensure that the incremental expected utility encapsulated in one additional search stage is smaller than its cost for the sub-coalition A_r . The second condition (C2) ensures that the additional utility obtained from moving any subset of A_t to A_r results in a negative expected net utility. At this point, we have all the necessary tools to establish theorem 1.

Theorem 1. *The optimal strategy of coalition A_g when in state s is to have the agents in A_r resume the search cooperatively and have the rest of the agents in A_g (i.e. the agents in A_t) terminate the search.*

The general sketch of the proof for theorem 1 begins by proving that it is sufficient to consider \hat{V} for the proposed process rather than using the actual values of the additional expected utility for each agent (i.e. the optimal strategy of each agent given the option to resume its search in future states). The proof for this is derived by showing that each agent's marginal expected utility obtained from resuming the search decreases throughout the search whereas the marginal cost of adding the agent to the coalition resuming search can only increase throughout the search (given the search cost structure and the fact that the coalition size throughout the search can only decrease). Therefore if it is not beneficial for the agent to resume its search in the current state given the \hat{V} criteria then certainly this is also the case when using the optimal future strategies. Next we prove that the optimal strategy for all agents in A_r is to resume their search cooperatively as one coalition. This is achieved by showing that under condition (C2), a scenario where one of the agent's expected additional utility from resuming the search is smaller than its own induced cost will not exist.

Before presenting an algorithm that is based on theorem 1 for computing the coalition's optimal strategy we illustrate this theorem by using the following example:

Example 1. *Suppose there are 4 agents $\{a_1, a_2, a_3, a_4\}$ conducting the search in a market associated with 4 types of opportunities $\{\vec{o}_1, \vec{o}_2, \vec{o}_3, \vec{o}_4\}$. Agents' utilities and opportunities' distribution are given in Table 1. The search cost associated with n agents is $c(n) = 0.4 \ln(n+1)$.*

We compute the coalition's optimal strategy, where the coalition encounters opportunity \vec{o}_1 and its current state is $s = \{\vec{o}_1\}$. First we compute $\forall a_j \in A_j V_{one}(\{\vec{o}_1\}, a_j)$.

$$V_{one}(\{\vec{o}_1\}, a_1) = p(\vec{o}_3)(U_1(\vec{o}_3) - U_1(\vec{o}_1)) + p(\vec{o}_4)(U_1(\vec{o}_4) - U_1(\vec{o}_1)) = 0.1(9 - 5) + 0.6(10 - 5) = 3.4$$

$$V_{one}(\{\vec{o}_1\}, a_2) = p(\vec{o}_2)(U_2(\vec{o}_2) - U_2(\vec{o}_1)) + p(\vec{o}_3)(U_2(\vec{o}_3) - U_2(\vec{o}_1)) = 0.1(4.4 - 4) + 0.1(5 - 4) = 0.14$$

$$V_{one}(\{\vec{o}_1\}, a_3) = p(\vec{o}_3)(U_3(\vec{o}_3) - U_3(\vec{o}_1)) = 0.1(8.5 - 7) = 0.15$$

$$V_{one}(\{\vec{o}_1\}, a_4) = p(\vec{o}_2)(U_4(\vec{o}_2) - U_4(\vec{o}_1)) = 0.1(8.5 - 8) = 0.05$$

The sorted agents list is $A_{order} = (a_1, a_3, a_2, a_4)$.

We start by checking whether the two conditions C1 and C2, are satisfied for $k = 1$. Condition C1 is satisfied as $V_{one}(\{\vec{o}_1\}, a_1) = 3.4 > c(1) = 0.2772588$. To check whether C2 is satisfied we should check that $\forall i > 1$ equation 7 isn't satisfied. For $i = 2$ equation 7 is not satisfied as $V_{one}(\{\vec{o}_1\}, a_2) = 0.15 < c(2) - c(1) = 0.161286$.

Opportunities	probability	a_1	a_2	a_3	a_4
\vec{o}_1	0.2	5	4	7	8
\vec{o}_2	0.1	3	4.4	8.5	7.5
\vec{o}_3	0.1	9	5	3	8.5
\vec{o}_4	0.6	10	4	6	5

Table 1: Agents' utilities for the four opportunities in Example 1

However $i = 3$ satisfies equation 7 as $V_{one}(\{\vec{o}_1\}, a_2) + V_{one}(\{\vec{o}_1\}, a_3) = 0.14 + 0.15 = 0.29 > c(3) - c(1) = 0.2772588$. Consequently $k = 1$ doesn't satisfy condition C2.

Next we check whether $k = 2$ satisfies the two conditions. Condition C1 is satisfied as $V_{one}(\{\vec{o}_1\}, a_1) + V_{one}(\{\vec{o}_1\}, a_2) = 3.4 + 0.15 = 3.55 > c(2) = 0.439944$. To check whether C2 is satisfied we should check that for all $i > 2$ equation 7 isn't satisfied. However $i = 3$ satisfies equation 7 as $V_{one}(\{\vec{o}_1\}, a_3) = 0.14 > c(3) - c(2) = 0.115072$. Consequently $k = 2$ doesn't satisfy condition C2.

The next index $k = 3$ satisfies C1 as $V_{one}(\{\vec{o}_1\}, a_1) + V_{one}(\{\vec{o}_1\}, a_2) + V_{one}(\{\vec{o}_1\}, a_3) = 3.4 + 0.15 + 0.14 = 3.69 > c(2) = 0.439944$. To check whether C2 is satisfied we should check that for $i = 4$ equation 7 isn't satisfied. Indeed $i = 4$ doesn't satisfy equation 7 as $V_{one}(\{\vec{o}_1\}, a_4) = 0.05 < c(4) - c(3) = 0.089$. Therefore $k = 3$ is the first index that satisfies both conditions C1 and C2. From theorem 1 we conclude that $\{a_1, a_2, a_3\}$ should resume the search and $\{a_4\}$ should terminate the search.

Based on theorem 1 we present algorithm 1 for computing the best strategy for a coalition A_g when reaching a state s . The significance of the algorithm is that it enables us to ex-

Algorithm 1 Computes the best strategy for coalition A_g when reaching a state s

Input: $U = \{U_1, \dots, U_{|A_g|}\}$ - coalition members' utility

Output: (A_r, A_t) - the sub-coalition that needs to resume the search and the complimentary set of agents that needs to terminate the search at the current stage, respectively.

- 1: Generate the sorted descending set A_{order} by computing $\hat{V}(s, a_j) \forall a_j \in A_g$ using equation 6
 - 2: $size \leftarrow 0; \hat{V}_{addCoalition} \leftarrow 0$
 - 3: **for** $index = 1$ to $|A_{order}|$ **do**
 - 4: $\hat{V}_{addCoalition} \leftarrow \hat{V}_{addCoalition} + \hat{V}(s, a'_{index})$
 - 5: **if** $(\hat{V}_{addCoalition} \geq c(index) - c(size))$ **then**
 - 6: $size \leftarrow index; \hat{V}_{addCoalition} \leftarrow 0$
 - 7: **end if**
 - 8: **end for**
 - 9: **return** $(\{a'_1, \dots, a'_{size}\}, \{a'_{size+1}, \dots, a'_{|A_g|}\})$
-

tract the optimal strategy for a coalition without considering all possible states. Its complexity doesn't depend on the number of states and it is polynomial by the market opportunities and agents, which extensively reduce the search space.

4 Illustrative Comparison

In this section we demonstrate the difference in the optimal strategy structure and in the overall performance between the

proposed VSCS model and the FSCS and SAS models. Notice that given proposition 1 the superiority of the VSCS over FSCS and SAS is unquestionable and therefore our goal is mainly to demonstrate different aspects of the optimal search strategy to be used in each of the different search methods.

We base the illustration on authentic environment 1, which was built based on opportunities collected over the internet and utility functions that were defined by human searchers whom we interviewed.

Environment 1. The searching coalition consists of seven agents interested in buying a calculator. Each agent is associated with a different utility function, based on typical attributes of calculators (price, handled/non-handled, display type, scientific functions, warranty, calculator's company, 2\1 line display, etc.). The utility functions of the different agents are constructed according to real preferences of 7 people (which evaluates the different attributes using monetary units) whom we interviewed. For example, a person searching for a handled scientific calculator with a fraction display that worth to him 30\$ (a calculator that doesn't satisfy all these conditions is not useful for this person and worth nothing to him) or a person searching for an handled scientific calculator that worth to him 15\$ and a last digit erase option worth to him 2\$. The opportunities to buy the calculator in this environment are drawn from a distribution that is based on one that can be found on the internet in US-based ecommerce web-sites. The cooperative search is executed by having one of the coalition members conduct the search at each stage and inform the results to the other coalition members. Therefore the search cost of the coalition is equal to the sum of the interaction cost (the cost of locating a seller and communicating with her to learn her offer) and the cost of communicating the search results to the other $n-1$ agents³, i.e. $c(n) = c_{interaction} + c_{communication} * (n - 1)$. The interaction cost we used is 0.01\$.

Figure 1 depicts the average overall utility over 10000 searches using the three methods as a function of the communication coalition search cost (notice that the agents' performance is not affected by this value in SAS). As expected the FSCS performs better than SAS for some $c_{communication}$ values (and SAS performs better for others) while VSCS dominates both methods for any $c_{communication}$ value.

5 Discussion and Conclusions

The ability to maintain an adaptive coalition that can restructure itself is inherent in multi-agent domains [Ogston *et al.*, 2003]. Therefore, when considering the expected performance of a coalition searching cooperatively, one must take into consideration the appropriate beneficial partitioning decisions that may be taken along its search. Having the partitioning option, the coalition should adopt a new strategy, different in its structure in comparison to the optimal strategy used in the fixed structure cooperative search (FSCS) and inherently different from the strategy used when each agent

³If a representative agent conducts the search on behalf of the coalition (substituting $n - 1$ in the search cost with n), we attain similar results.

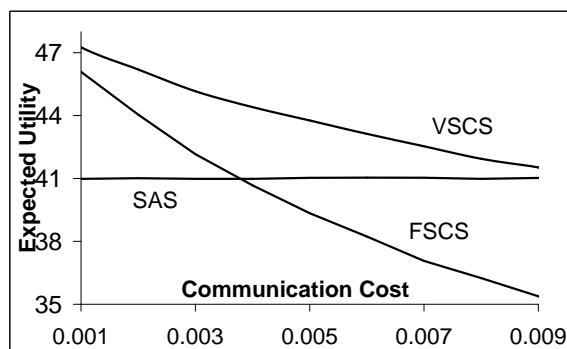


Figure 1: Average utility per buyer agent for the different models searches individually (SAS). Though we use the electronic marketplace as a framework in this paper, the suggested analysis is general and can be applied to various domains where agents can benefit from cooperative search (e.g. searching a large database of potential candidates to fill several positions).

We re-emphasize the fact that the proposed VSCS model is a generalization of FSCS and SAS and its use is always favorable for the coalition. This significantly enhances the importance and usefulness of cooperative search, since the dominance of the traditional method (FSCS) over individual search (SAS) can not be guaranteed.

Naturally the attempt to integrate "search theory" techniques into day-to-day MAS applications raises the applicability question. Justification and legitimacy considerations for the integrated search-MAS models in general and cooperative search model in particular were widely discussed in the literature reviewed throughout the paper, emphasizing both the synthesis of the two domains [Sarne and Kraus, 2005; Manisterski *et al.*, 2006] and applications [Kephart and Greenwald, 2002; Choi and Liu, 2000]. The current paper does not focus on re-arguing applicability, but rather on the improvement of the well established cooperative search model. The mechanisms described in this paper are an important step towards an improved cooperative search, however there are many other aspects of the VSCS that should be addressed, such as: coalition stability, payoff division mechanisms (in particular when partitioning the coalition along its search) and truth telling [Sandholm *et al.*, 1999; Tsvetovat *et al.*, 2000]. Though these were not included in the current paper, we wish to emphasize that not only is the coalition's optimal strategy not influenced by these factors but rather it influences them [Sarne and Kraus, 2005]. The analysis of these important issues is based on the ability to properly derive the coalition's utility given its initial specific self structure (i.e. the number of agents it represents and their reported, not necessarily true, utility functions) and the environment in which it operates. By supplying this functionality, we enable extensive important research in the future and support the integration of various relevant ideas from rich literature in the area of game theory and MAS research [Sandholm *et al.*, 1999; Li *et al.*, 2003] into the proposed model. An additional extension of the model would be to integrate the concept of concurrency into the VSCS model, taking advantage of the results obtained in [Manisterski *et al.*, 2006] for the FSCS model.

6 Acknowledgment

This work was supported in part by ISF #8008. Kraus is also affiliated with UMIACS.

References

- [Bakos, 1997] Y. Bakos. Reducing buyer search costs: Implications for electronic marketplaces. *Mgmt. Science*, 42:1676–92, 1997.
- [Choi and Liu, 2000] S. Choi and J. Liu. Optimal time-constrained trading strategies for autonomous agents. In *Proc. of MAMA'00*, 2000.
- [Kahan and Rapoport, 1984] J.P. Kahan and A. Rapoport. *Theories of coalition formation*. 1984.
- [Kephart and Greenwald, 2002] J. Kephart and A. Greenwald. Shopbot economics. *JAAMAS*, 5(3):255–287, 2002.
- [Lermann and Shehory, 2000] K. Lermann and O. Shehory. Coalition formation for large scale electronic markets. In *ICMAS'2000*, pages 216–222, 2000.
- [Li *et al.*, 2003] C. Li, S. Chawla, U. Rajan, and K. Sycara. Mechanisms for coalition formation and cost sharing in an electronic marketplace. In *ECO3*, pages 68 – 77, 2003.
- [Lippman and McCall, 1976] S. Lippman and J. McCall. The economics of job search: A survey. *Economic Inquiry*, 14:155–189, 1976.
- [Manisterski *et al.*, 2006] E. Manisterski, D. Sarne, and S. Kraus. Integrating parallel interactions into cooperative search. In *AAMAS06*, pages 834–841. 2006.
- [McMillan and Rothschild, 1994] J. McMillan and M. Rothschild. Search. In Robert J. Aumann and Amsterdam Sergiu Hart, editors, *Handbook of Game Theory with Economic Applications*, pages 905–927, 1994.
- [Ogston *et al.*, 2003] E. Ogston, B. Overeinder, M. Van Steen, and F. Brazier. A method for decentralized clustering in large multi-agent systems. In *AAMAS03*, pages 798–796, 2003.
- [Sandholm *et al.*, 1999] T. Sandholm, K. Larson, M. R. Andersson, O. Shehory, and F. Tohme. Coalition structure generation with worst case guarantees. *AIJ*, 111:209–238, 1999.
- [Sarne and Kraus, 2005] D. Sarne and S. Kraus. Cooperative exploration in the electronic marketplace. In *Proc. of AAAI 2005*, pages 158–163, 2005.
- [Tsvetovat *et al.*, 2000] N. Tsvetovat, K. Sycara, Y. Chen, and J. Ying. Customer coalitions in electronic markets. In *AMEC00*, pages 121–138, 2000.
- [Yamamoto and Sycara, 2001] J. Yamamoto and K. Sycara. A stable and efficient buyer coalition formation scheme for e-marketplaces. In *Agents01*, pages 576–583, 2001.
- [Ohta *et al.*, 2006] N. Ohta, A. Iwasaki, M. Yokoo, K. Maruono, V. Conitzer, and T. Sandholm. A compact representation scheme for coalitional games in open anonymous environments. In *AAAI06*, 2006.